



Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?

Philip Taranto, Faraj Bakhshinezhad, Andreas Bluhm, Ralph Silva, Nicolai Friis, Maximilian P. E. Lock, Giuseppe Vitagliano, Felix C. Binder, Tiago Debarba, Emanuel Schwarzhans, et al.

► To cite this version:

Philip Taranto, Faraj Bakhshinezhad, Andreas Bluhm, Ralph Silva, Nicolai Friis, et al.. Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?. PRX Quantum, 2023, 4, pp.010332. 10.1103/PRXQuantum.4.010332 . hal-04082955v2

HAL Id: hal-04082955













<https://hal.univ-grenoble-alpes.fr/hal-04082955v2>

Submitted on 26 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?

Philip Taranto ^{1,2,3,*} Faraj Bakhshinezhad ^{2,4,3,†} Andreas Bluhm ^{5,6,‡} Ralph Silva ^{7,‡}
 Nicolai Friis ^{2,3} Maximilian P. E. Lock ^{2,3} Giuseppe Vitagliano ^{2,3} Felix C. Binder ^{8,3,2}
 Tiago Debarba ⁹ Emanuel Schwarzhans ^{2,3} Fabien Clivaz ^{10,3} and Marcus Huber ^{2,3,§}

¹Department of Physics, Graduate School of Science, The University of Tokyo, 7-3-1 Hongo, Bunkyo City, Tokyo 113-0033, Japan

²Atominstytut, Technische Universität Wien, 1020 Vienna, Austria

³Institute for Quantum Optics and Quantum Information - IQOQI Vienna,
Austrian Academy of Sciences, Boltzmanngasse 3, 1090 Vienna, Austria

⁴Department of Physics and Nanolund, Lund University, Box 118, 221 00 Lund, Sweden

⁵Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, 38000 Grenoble, France

⁶QMATH, Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100 Copenhagen, Denmark

⁷Institute for Theoretical Physics, ETH Zürich, Wolfgang-Pauli-Str. 27, Zürich, Switzerland

⁸School of Physics, Trinity College Dublin, Dublin 2, Ireland

⁹Departamento Acadêmico de Ciências da Natureza, Universidade Tecnológica Federal do Paraná (UTFPR),
Campus Cornélio Procopio, Avenida Alberto Carazzai 1640, Cornélio Procopio, Paraná 86300-000, Brazil

¹⁰Institut für Theoretische Physik und IQST, Universität Ulm, Albert-Einstein-Allee 11, D-89069 Ulm, Germany

(Dated: March 29, 2023)

Thermodynamics connects our knowledge of the world to our capability to manipulate and thus to control it. This crucial role of control is exemplified by the third law of thermodynamics, Nernst’s unattainability principle, which states that infinite resources are required to cool a system to absolute zero temperature. But what are these resources and how should they be utilised? And how does this relate to Landauer’s principle that famously connects information and thermodynamics? We answer these questions by providing a framework for identifying the resources that enable the creation of pure quantum states. We show that perfect cooling is possible with Landauer energy cost given infinite time or control complexity. However, such optimal protocols require complex unitaries generated by an external work source. Restricting to unitaries that can be run solely via a heat engine, we derive a novel Carnot-Landauer limit, along with protocols for its saturation. This generalises Landauer’s principle to a fully thermodynamic setting, leading to a unification with the third law and emphasises the importance of control in quantum thermodynamics.

I. INTRODUCTION

What is the cost of creating a pure state? Pure states appear as ubiquitous idealisations in quantum information processing and preparing them with high fidelity is essential for quantum technologies such as reliable quantum communication [1, 2], high-precision quantum parameter estimation [3–5], and fault-tolerant quantum computation [6, 7]. Fundamentally, pure states are prerequisites for ideal measurements [8] and precise timekeeping [9, 10]. To answer the above question, one could turn to Landauer’s principle, stating that erasing a bit of information has an *energy* cost of at least $k_B T \log(2)$ [11]. Alternatively, one could consult Nernst’s unattainability principle (the third law of thermodynamics) [12], stating that cooling a physical system to its ground state requires diverging resources. At the outset, it seems that these statements are at odds with one another. However, Landauer’s protocol requires infinite time, thus identifying *time* as a resource according to the third law [13–17]. *Does this mean either infinite energy or time are needed to prepare a pure state?*

The perhaps surprising answer we give here is: *no*. We show that finite energy and time suffice to perfectly cool any

quantum system and we identify the previously hidden resource—*control complexity*—that must diverge (in the spirit of Nernst’s principle) to do so. Intuitively, the control complexity of a protocol refers to the structure of machine energy gaps that the cooling unitary must couple the system to; we demonstrate that this energy-level spectrum must approximate a continuum in order to cool with minimal time and energy costs. In short, the ultimate limit on the energetic cost of cooling is still provided by the Landauer limit, but in order to achieve it, either time or control complexity must diverge.

At the same time, heat fluctuations and short coherence times in quantum technologies [18] demand that both energy and time are not only finite, but minimal. Therefore, in addition to proving the necessity of diverging control complexity for perfect cooling with minimal time and energy, we develop explicit protocols that saturate the ultimate limits. We demonstrate that mitigating overall heat dissipation comes at the practical cost of controlling fine-tuned interactions that require a *coherent* external work source, i.e., a quantum battery [19–23]. From a thermodynamic perspective, this may seem somewhat unsatisfactory: nonequilibrium resources imply that the total system is not closed, and the optimal protocol (saturating the Landauer bound) is reminiscent of a Maxwellian demon with perfect control.

Accordingly, we also consider an *incoherent* control setting restricted to global energy-conserving unitaries with a heat bath as thermodynamic energy source. This setting corresponds to minimal overall control, where interactions

* philipguy.taranto@phys.s.u-tokyo.ac.jp; P. T. and F. B. contributed equally.

† P. T. and F. B. contributed equally.

‡ A. B. and R. S. contributed equally.

§ marcus.huber@tuwien.ac.at

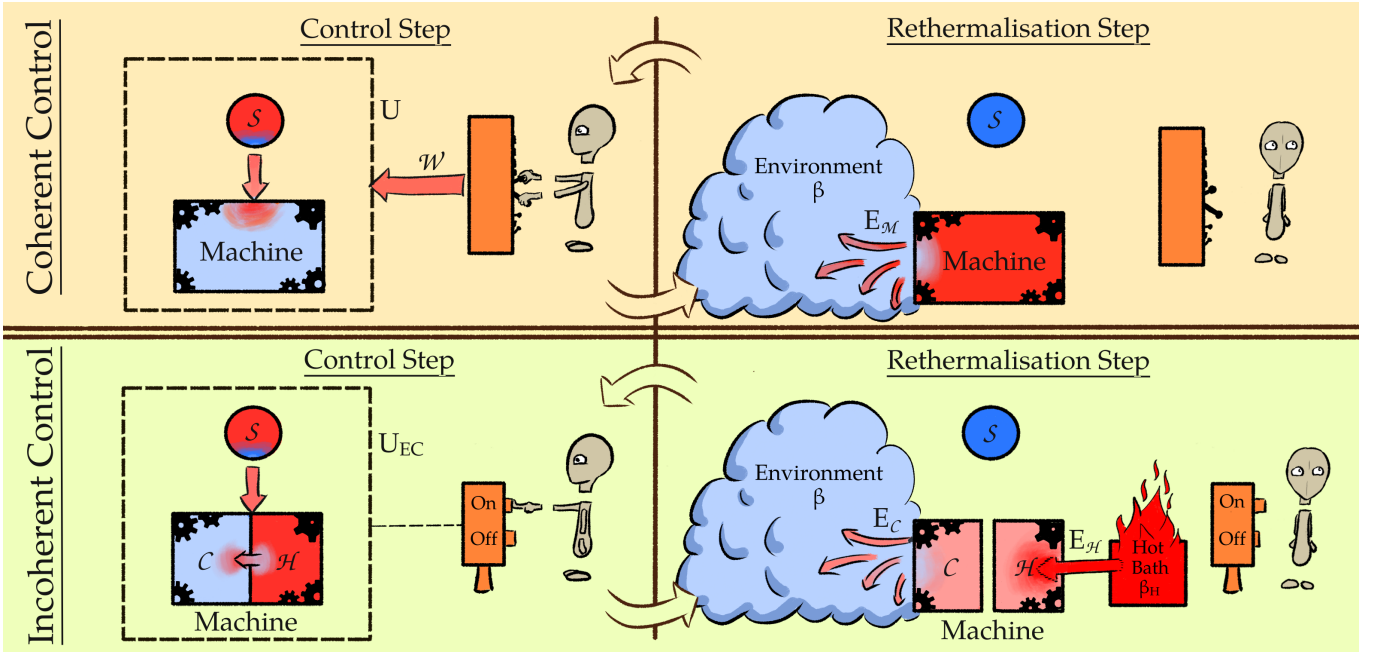


FIG. 1. *Framework.* We consider the task of cooling a quantum system in two extremal control scenarios, with each step of both paradigms comprising two primitives. The top panel depicts the coherent-control scenario: in the control step (left), an agent can use a work source \mathcal{W} to implement any global unitary on the system S and machine M , which both begin thermal at inverse temperature β ; in cooling the target, energy and entropy is transferred to the machine. The machine then rethermalises with its environment (right), thereby dissipating the energy it gained in the control step. The bottom panel depicts the incoherent-control scenario: the machine is bipartitioned into a cold part at inverse temperature β and a hot part at inverse temperature $\beta_H < \beta$. In the control step, the agent switches on an interaction between the three systems, represented by a global energy-conserving unitary U_{EC} . In the rethermalisation step, the interaction is turned off and both subsystems of the machine rethermalise to their respective initial temperatures; the hot part draws energy from the heat bath while the cold part dissipates heat to its environment. In both paradigms, we quantify the control complexity as the effective dimension accessed by the unitary operation in a given control step (i.e., the dimension of the system-machine Hilbert space upon which the unitary acts nontrivially).

need only be switched on and off to generate transformations, i.e., a heat engine alone drives the dynamics [24–28]. The incoherent-control setting is therefore fully thermodynamically consistent inasmuch as both the machine state is assumed to be thermal (and to rethermalize between control steps) and the permitted control operations are those implementable solely via a heat engine. In this paradigm, we show that the Landauer bound is not attainable, subsequently derive a novel limit—which we dub the *Carnot-Landauer bound*—and construct protocols that saturate it, thereby establishing its significance. The Carnot-Landauer bound follows from an equality phrased in terms of entropic and energetic quantities that must hold for any state transformation in the incoherent control paradigm; in this sense, the Carnot-Landauer equality adapts the equality version of Landauer’s principle developed in Ref. [29] to a fully (quantum) thermodynamic setting.

Our work thus both generalises Landauer’s erasure principle and, at the same time, unifies it with the laws of thermodynamics. By accounting for control complexity, we emphasise a crucial resource that is oftentimes overlooked but, as we show, must be taken into account for any operationally meaningful theory of thermodynamics. Here, we focus on the asymptotic setting that allows us to connect this resource with Nernst’s unattainability principle. Beyond the asymp-

totic case, the gained insights also open the door to a better understanding of the intricate relationship between energy, time, and control complexity when all resources are finite, which will be crucial for practical applications; we additionally provide a preliminary analysis to this end. Lastly, our protocols saturating the Carnot-Landauer bound pave the way for thermodynamically driven (i.e., minimal-control) quantum technologies, which, by mitigating the cost of control at the very outset, could lead to tangible advantages.

Overview & Summary of Results

Loosely speaking, there are two types of thermodynamic laws: those, like the second law, that bound (changes of) characteristic quantities during thermodynamic processes, and those, like the third law, which state the impossibility of certain tasks. Landauer’s principle is of the former kind (indeed, it can be rephrased as a version of the second law), associating a minimal heat dissipation to any logically irreversible process, thereby placing a fundamental limit on the energy cost of computation. The paradigmatic logically irreversible process is that of erasing information, i.e., resetting an arbitrary state to a blank register. From a physics perspective, said task

can be rephrased as *perfectly cooling* a system to the ground state, or more generally, taking an initially full-rank state to a rank-deficient one.¹ Note that although there is, in general, a distinction between physical cooling and information erasure, in this paper we focus on erasing quantum information encoded in fundamental degrees of freedom rather than in logical macrostate sectors, and accordingly use the terms somewhat interchangeably. This is justified because in either case, the ultimate limitation (be it cooling to absolute zero or perfectly erasing information) requires a rank-decreasing process, which is what we formally analyse.

Nernst's unattainability principle is of the latter kind of thermodynamic law, stating that perfectly cooling a system requires diverging resources. The resources typically considered are energy and time, whose asymptotic trade-off relation is relatively well established: on the one hand, perfect cooling can be achieved in finite time at the expense of an energy cost that diverges as the ground state is approached; on the other hand, the energy cost can be minimised by implementing a quasistatic process that saturates the Landauer limit but takes infinitely long.²

These two types of thermodynamic laws are intimately related, but details of their interplay have remained elusive: under which conditions can the Landauer bound be saturated and what are the minimal resources required to do so? Which protocols asymptotically create pure states with given (diverging) resources? What type of control do such protocols require and how difficult are they to implement in practice? We address these questions by considering the task of cooling a quantum system in two extremal control paradigms (see Fig. 1): One driven by a *coherent* work source and the other by an *incoherent* heat engine.

After laying out the framework, we proceed to analyse the relationship between the aforementioned three resources for cooling. A core idea of this paper originates from the observation that it is possible to perfectly cool a physical system with both finite energy and time. Although said observation is simple in nature inasmuch as it can be obtained by a shift in perspective of Landauer's original protocol, its consequences run deep: indeed, the apparent tension between Landauer cooling and Nernst's unattainability principle that arises when only energy and time are considered as resources is resolved via the inclusion of control complexity as a consideration. Subsequently, we define a meaningful notion of control complexity in terms of the energy-level structure of the machine that the

system must be coupled to throughout the cooling protocol and demonstrate its thermodynamic consistency by showing that it indeed must diverge to cool the system to the ground state at minimal energy cost, thereby reconciling the viewpoints of Landauer and Nernst.

Having established the trinity of relevant resources, we present three main results:

1. Perfect cooling is possible with coherent control provided either energy, time, or control complexity diverge. In particular, it is possible in finite time and at Landauer energy cost with diverging control complexity.
2. Perfect cooling is possible with incoherent control, i.e., with a heat engine, provided either time or control complexity diverge. On the other hand, it is impossible with both finite time and control complexity, regardless of the amount of energy drawn from the heat bath.
3. No process driven by a finite-temperature heat engine can (perfectly) cool a quantum system at the Landauer limit. Nonetheless, the Carnot-Landauer limit, which we introduce here (as a consequence of a stronger equality), can be saturated for any heat bath, given either diverging time or control complexity.

In the following, we discuss each of these results in turn in more detail and provide a systematic study concerning the asymptotic interplay of energy, time, and control complexity as thermodynamic resources in two extremal control paradigms, as well as develop insight into the finite-resource regime for some special cases. We begin by outlining the framework.

II. FRAMEWORK: COOLING A PHYSICAL SYSTEM

Consider a target system \mathcal{S} in an initial state $\varrho_{\mathcal{S}}$ described by a unit-trace, positive semidefinite operator with associated Hamiltonian $H_{\mathcal{S}}$. An auxiliary machine \mathcal{M} , initially uncorrelated with \mathcal{S} and in equilibrium with a reservoir at inverse temperature $\beta := \frac{1}{k_B T}$, is used to cool the target system. The initial state of \mathcal{M} is thus of Gibbs form,

$$\varrho_{\mathcal{M}} = \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) := \frac{e^{-\beta H_{\mathcal{M}}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})}, \quad (1)$$

where $H_{\mathcal{M}}$ is the machine Hamiltonian and $\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}}) := \text{tr}[e^{-\beta H_{\mathcal{M}}}]$ its partition function. Throughout this paper we consider only Hamiltonians with discrete spectra, i.e., with an associated separable Hilbert space that has a countable energy eigenbasis. Moreover, for the most part we consider finite-dimensional systems (or sequences thereof) and deal with infinite-dimensional systems separately.

As shown in Fig. 1, a single step of a cooling process comprises two subprocedures: first, a joint unitary is implemented during the *control* step; second, the machine *rethermalises* to the ambient temperature. A cooling *protocol* is determined

¹ Low-temperature thermal states correspond to those with low information content, as they have low entropy or small effective support; viewing cooling more broadly (i.e., not restricting to thermal states and allowing for arbitrary Hamiltonians), we see that cooling indeed encompasses information erasure: States with smaller effective support are “colder” than those with greater support according to any meaningful notion of “cool” (see Ref. [30]).

² Note, however, that although the asymptotic trade-off relationship is known, the connection between energy and time in the finite-resource setting remains unresolved: For instance, if one uses twice the amount of energy, it is not clear how much faster a given protocol can be implemented; we provide some preliminary insight to such questions in Sec. VI.

by the initial conditions and any concatenation of such primitives³. We consider two extremal control paradigms corresponding to two classes of allowed global transformations. The *coherent control* paradigm permits arbitrary unitaries on \mathcal{SM} ; in general, these change the total energy but leave the global entropy invariant and thus require an external work source \mathcal{W} . At the other extreme is the *incoherent control* paradigm, where the energy source is a heat bath. Here, the machine \mathcal{M} is bipartitioned: one part, \mathcal{C} , is connected to a *cold* bath at inverse temperature β , which serves as a sink for all energy and entropy flows; the other, \mathcal{H} , is connected to a *hot* bath at inverse temperature $\beta_H \leq \beta$, which provides energy. The composite system \mathcal{SCH} is closed and thus global unitary transformations are restricted to be energy conserving. The temperature gradient causes a natural heat flow away from the hot bath, which carries maximal entropic change with it. Cooling protocols in this setting can be run with minimal external control, i.e., they require only the switching on and off of interactions.

III. COHERENT CONTROL

We begin by considering cooling with coherently controlled resources (see Fig. 1, top panel). We first analyse energy, time, and control complexity as resources that can be traded off against one another in order to optimise cooling performance, before focusing more specifically on the nature and role of control complexity.

A. Energy, Time, and Control Complexity as Resources

In the coherent-control setting, a transformation $\varrho_S \rightarrow \varrho'_S$ is enacted via a unitary U on \mathcal{SM} involving a thermal machine $\varrho_M = \tau_M(\beta, H_M)$, i.e.,

$$\varrho'_S := \text{tr}_M [U(\varrho_S \otimes \varrho_M)U^\dagger]. \quad (2)$$

For such a transformation, there are two energy costs contributing to the total energy change, which must be drawn from a work source \mathcal{W} . The first is the energy change of the target $\Delta E_S := \text{tr}[H_S(\varrho'_S - \varrho_S)]$; the second is that of the machine $\Delta E_M := \text{tr}[H_M(\varrho'_M - \varrho_M)]$, where $\varrho'_M := \text{tr}_S [U(\varrho_S \otimes \varrho_M)U^\dagger]$. The latter is associated with the heat dissipated into the environment and is given by [29]

$$\beta \Delta E_M = \tilde{\Delta} S_S + I(\mathcal{S} : \mathcal{M})_{\varrho'_{SM}} + D(\varrho'_M \| \varrho_M), \quad (3)$$

where $S(\varrho) := -\text{tr}[\varrho \log(\varrho)]$ is the von Neumann entropy, $\tilde{\Delta} S_A := S(\varrho_A) - S(\varrho'_A)$ ⁴, $I(\mathcal{A} : \mathcal{B})_{\varrho_{AB}} := S(\varrho_A) +$

	Energy	Time	Complexity
Qudit	$\rightarrow \infty$	1	$\frac{1}{2}d(d-1)$
	Landauer	$\rightarrow \infty$	$\frac{1}{2}d(d-1)$
	Landauer	1	$\rightarrow \infty$
H.O.	$\rightarrow \infty$	1	$\rightarrow \infty$ (Gaussian)
	Landauer	$\rightarrow \infty$	$\rightarrow \infty$ (Gaussian)
	Finite ($>$ Landauer)	$\rightarrow \infty$	1 (Non-Gaussian)
	Landauer	1	$\rightarrow \infty$ (Gaussian)

TABLE I. *Coherent-control cooling protocols for finite-dimensional (qudit) and harmonic oscillator systems.* Landauer energy cost refers to saturation of Eq. (4) and complexity refers to the proxy measure effective dimension (see Def. 1); time is measured as the number of unitary operations with a fixed complexity. In the qudit case, the system and machine dimensions are equal: $d_S = d_M =: d$.

$S(\varrho_B) - S(\varrho_{AB})$ (with marginals $\varrho_{A/B} := \text{tr}_{B/A}[\varrho_{AB}]$) is the mutual information between \mathcal{A} and \mathcal{B} , and $D(\varrho \| \sigma) := \text{tr}[\varrho \log(\varrho)] - \text{tr}[\varrho \log(\sigma)]$ is the relative entropy of ϱ with respect to σ , with $D(\varrho \| \sigma) := \infty$ if $\text{supp}[\varrho] \not\subseteq \text{supp}[\sigma]$. We derive Eq. (3) and its generalisation to the incoherent-control setting in Appendix A. The mutual information is non-negative and vanishes iff $\varrho_{AB} = \varrho_A \otimes \varrho_B$; similarly, the relative entropy is non-negative and vanishes iff $\varrho = \sigma$. Dropping these terms leads to the Landauer bound [11]

$$\beta \Delta E_M \geq \tilde{\Delta} S_S. \quad (4)$$

The Landauer limit holds *independently* of the protocol implemented, i.e., it assumes only that *some* unitary was applied to the target and thermal machine. For large machines, the dissipated heat is typically much greater than the energy change of the target; nonetheless, the contributions can be comparable at the microscopic scale. We assume that the target begins in equilibrium with the reservoir at inverse temperature β , i.e., in the initial thermal state $\varrho_S = \tau_S(\beta, H_S)$, with no loss of generality since such a relaxation can be achieved for free (by swapping the target with a suitable part of the environment; however, see Ref. [31] for a discussion of initial state dependency of the bound). We track all energetic and entropic quantities and refer to the asymptotic saturation of Eq. (4) with ϱ'_S pure as *perfect cooling at the Landauer limit*.

Although Landauer's limit sets the minimum heat that must be dissipated—and thereby the minimum energy cost—for cooling any physical system, the third law makes no specification that energy must be the resource minimised (or that time must diverge). One might instead consider using a source of unbounded energy to perfectly cool a system as quickly as possible. Additionally, control complexity plays an important role as a resource, inasmuch as its divergence permits perfect cooling at the Landauer limit in finite time (see below). As summarised in Table I, we now present coherently controlled protocols that perfectly cool an arbitrary finite-dimensional target system using thermal machines when any one of the three considered resources—energy, time or control complexity—diverges; moreover, the resources that are kept finite saturate protocol-independent ultimate bounds. The following

³ One could refer to both \mathcal{M} and the transformations applied as the *machine* and call the system \mathcal{M} itself the working *medium* inasmuch as the latter passively facilitates the process, in line with conventional parlance; however, we use the terminology established in the pertinent literature.

⁴ Note the differing sign conventions (denoted by the tilde) that we use for changes in energies, $\Delta E_X := E'_X - E_X$, and in entropies, $\tilde{\Delta} S_X := S_X - S'_X$, such that energy *increases* and entropy *decreases* are positive.

thus provides a comprehensive analysis of cooling with respect to the trinity of resources that can be traded off amongst each other.

B. Perfect Cooling at the Ultimate Limits with Infinite Resources

1. *Diverging Energy.*—We first consider the situation in which time and control complexity are fixed to be finite, while the energy cost is allowed to diverge. Here, we present the following:

Theorem 1. *With diverging energy, any finite-dimensional quantum system can be perfectly cooled using a single interaction of finite complexity.*

The cooling protocol using diverging energy is the simplest. Here, one exchanges all populations of the target system with those of a thermal machine with suitably large energy gaps to sufficiently concentrate the initial machine population in the ground state subspace of the target system. This exchange requires a single system-machine unitary and is of finite complexity (in a sense discussed below). Nonetheless, the energy drawn from the work source in this protocol diverges. Moreover, in addition to being sufficient for perfect cooling with both finite time and control complexity, any protocol that cools perfectly with both finite time and control complexity requires diverging energy. See Appendix B for details.

We now move to consider the situations in which the energy cost is minimised at the expense of either diverging time or control complexity. Equation (3) provides insight for understanding the conditions required for saturating the Landauer bound. Although for finite-dimensional machines only trivial processes of the form $U_{S\mathcal{M}} = U_S \otimes \mathbb{1}_{\mathcal{M}}$ saturate the Landauer limit [29], we show how it can be asymptotically saturated with nontrivial processes by considering diverging machine and interaction properties, as we elaborate on shortly. Any such process must asymptotically exhibit no correlations such that $I(S : \mathcal{M})_{\varrho'_{S\mathcal{M}}} \rightarrow 0$ and effectively not disturb the machine, i.e., yield $\varrho'_{\mathcal{M}} \rightarrow \varrho_{\mathcal{M}}$ such that $D(\varrho'_{\mathcal{M}} \| \varrho_{\mathcal{M}}) \rightarrow 0$. Indeed, any correlations created between initially thermal systems would come at the expense of an additional energetic cost [32–34] whose minimisation is a problem that has so far only been partially resolved [35]. However, it has been shown that for any (strictly) rank nondecreasing process, there exists a thermal machine and joint unitary such that for any $\epsilon > 0$, the heat dissipated satisfies $\beta \Delta E_{\mathcal{M}} \leq \tilde{\Delta} S_S + \epsilon$ [29], thereby saturating the Landauer limit. Here, we present protocols that asymptotically achieve both this and perfect cooling (in particular, effectively decrease the rank), and provide necessary conditions on the underlying resources required to do so.

2. *Diverging Time.*—We now present a protocol that uses a diverging number of operations of finite complexity to asymptotically attain perfect cooling at the Landauer limit [20, 29, 36].

Theorem 2. *With diverging time, any finite-dimensional quantum system can be perfectly cooled at the Landauer limit via interactions of finite complexity.*

Sketch of proof.—We first show that any system can be cooled from $\varrho_S = \tau_S(\beta, H_S)$ to $\tau_S(\beta^*, H_S)$, with $\beta^* \geq \beta$, using only $\beta^{-1} \tilde{\Delta} S_S$ units of energy. Our proof is constructive in the sense that we provide a protocol that achieves the Landauer energy cost as the number of operations diverges. The individual interactions in this protocol are of finite control complexity as they simply swap the target system with one of a sequence of thermal machines with increasing energy gaps. In this way, the final state $\tau_S(\beta^*, H_S)$ can be made to be arbitrarily close to $|0\rangle\langle 0|_S$ for any initial temperature. \square

The proof is presented in Appendix C, along with a more detailed dimension-dependent energy cost function for the special case of equally spaced Hamiltonians.

Through the protocol described above, we see that given a diverging amount of time, the target system can be sequentially coupled with a machine of finite complexity that rethermalizes between control steps in such a way that the final target system state is arbitrarily close to the ground state for any initial temperature. This trade-off between energy and time is well known, and we discuss it only briefly in order to help build intuition and highlight the versatility of our framework. Alternatively, one can compress all the operations applied in the diverging-time protocol into one global unitary that achieves the same final states, thereby achieving perfect cooling at the Landauer limit in a single unit of time but with an infinitely complex interaction. That is, the diverging temporal resource of repeated interactions with a single, finite-size machine is replaced by a single interaction with a larger machine of diverging control complexity.

3. *Diverging Control Complexity.*—By reconsidering the diverging-time protocol above, a trade-off can be made between time and control complexity. As illustrated in Fig. 2, one can consider all of the operations $\{U_k = e^{-iH_k t_k}\}_{k=1, \dots, N}$ required in said protocol to make up one single joint interaction $U_{\text{tot}} := \lim_{N \rightarrow \infty} \prod_{k=1}^N U_k = e^{-iH_{\text{tot}} t_{\text{tot}}}$ acting on a larger machine, thus setting the time required to be unity (in terms of the number of control operations before the machine rethermalises). In other words, for any finite number N of unitary transformations U_k , there exists a total Hamiltonian $H_{\text{tot}}^{(N)}$ and a finite time t_N that generates the overall transformation $U_{\text{tot}}^{(N)} := \prod_{k=1}^N U_k$; since t_N is finite, we can set it equal to one without loss of generality by rescaling the Hamiltonian as $\tilde{H}_{\text{tot}}^{(N)} = t_N H_{\text{tot}}^{(N)}$. Here, we refer to the limit $N \rightarrow \infty$ as diverging control complexity. Compressing a diverging number of finite-complexity operations thus yields a protocol of diverging control complexity. The fact that there exists such an operation that minimises both the time and energy requirements follows from our constructive proof of Theorem 2. We therefore have the following:

Corollary 1. *With diverging control complexity, any finite-dimensional quantum system can be perfectly cooled at the Landauer limit in finite time.*

However, this particular way of constructing complex control protocols is not necessarily unique. It is thus natural to wonder if diverging control complexity is a generic feature necessary to achieve perfect cooling at the Landauer limit in unit time and indeed, how to quantify control complexity

that is operationally meaningful between the extreme cases of being either very small or divergent, as we now turn to discuss. Indeed, the inclusion of an explicit quantifier of control complexity regarding thermodynamic tasks—which, although crucial for practical purposes, is oftentimes overlooked—is one of the main novelties of our present work.

IV. CONTROL COMPLEXITY IN QUANTUM THERMODYNAMICS

Although the protocol described above has diverging control complexity by construction, one need not construct complex protocols in this way, and so the natural concern becomes understanding the generic features that enable perfect cooling at the Landauer limit in unit time. To address this issue, we first provide protocol-independent structural conditions that must be fulfilled by the machine to enable (1) *perfect cooling* and (2) *cooling at Landauer cost*; combined, these independent conditions provide a necessary requirement, namely that the machine must have an unbounded spectrum (from above) and be infinite-dimensional (respectively) for the *possibility* of (3) *perfect cooling at the Landauer limit*. Such properties of the machine Hamiltonian define the *structural complexity*, which sets the potential for how cool the target system can be made and at what energy cost. As the name suggests, this is entailed by the structure of the machine, e.g., the number of energy gaps and their arrangement, and as such provides a static notion of complexity. However, given a machine with particular structural complexity, one may not be able to utilise said potential due to constraints on the dynamics that can be implemented. For instance, one may be restricted to only two-body interactions, or operations involving only a few energy levels at a time. Assuming a sufficient structural complexity at hand, such constraints limit one from optimally manipulating the systems. Thus, the extent to which a machine's potential is utilised depends on properties of the dynamics of a given protocol, i.e., the *control complexity*. We provide a detailed study of structural and control complexity in Appendix D, and here summarise the key methods.

A. Structural & Dynamical Notions of Complexity

We split the consideration of complexity into two parts: first, the protocol-independent *structural* conditions that must be fulfilled by the machine and, second, the dynamic *control complexity* properties of the interaction that implements a given protocol (see Fig. 2).

1. Structural Complexity

Regarding the former, first note that one can lower bound the smallest eigenvalue λ_{\min} of the final state ϱ'_S (and hence how cold the system can become) after *any* unitary interaction

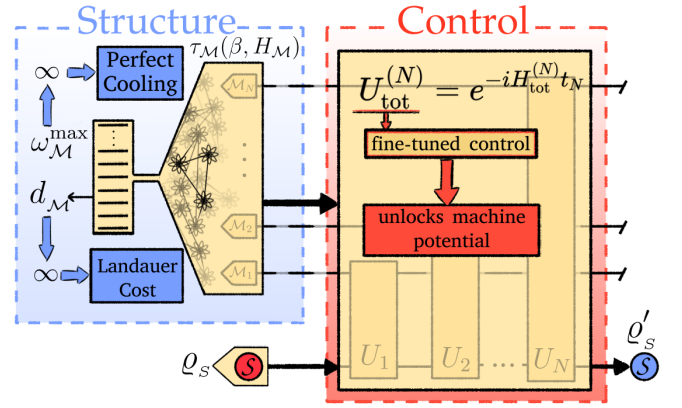


FIG. 2. *Complexity*. We consider structural (left) and control complexity (right). Structural complexity concerns properties of the machine Hamiltonian. For perfect cooling it is necessary that the largest energy gap diverges [see Eq. (5)]. Moreover, an infinite-dimensional machine with particular energy-level structure is required for saturation of the Landauer bound. Control complexity refers to properties of the unitary that represents a protocol. The yellow box in the foreground represents a unitary U involving the entire machine, whereas the smaller yellow columns in the background represent a potential decomposition (e.g., of the diverging-time protocol) into unitaries U_i involving certain subspaces of the overall machine. Not only must the target system interact with all levels of an infinite-dimensional machine for Landauer-cost cooling, it must do so in a fine-tuned way.

with a thermal machine by [29]

$$\lambda_{\min}(\varrho'_S) \geq e^{-\beta \omega_{\mathcal{M}}^{\max}} \lambda_{\min}(\varrho_S), \quad (5)$$

where $\omega_{\mathcal{M}}^{\max} := \max_{i,j} |\omega_j - \omega_i|$ denotes the largest energy gap of the machine Hamiltonian $H_{\mathcal{M}}$ with eigenvalues ω_i . It follows that perfect cooling is only possible under two conditions: either the machine begins in a pure state ($\beta \rightarrow \infty$), or $H_{\mathcal{M}}$ is unbounded, i.e., $\omega_{\mathcal{M}}^{\max} \rightarrow \infty$. Requiring $\beta < \infty$, a diverging energy gap in the machine Hamiltonian is thus a necessary structural condition for perfect cooling. Independently, another condition required to saturate the Landauer limit can be derived for any amount of cooling: in Ref. [29], it was shown that for any finite-dimensional machine, there are correction terms to the Landauer bound which imply that it cannot be saturated; these terms only vanish in the limit where the machine dimension diverges.

We thus have two independent necessary conditions on the structure of the machine that must be asymptotically fulfilled to achieve relevant goals for cooling: the former is required for perfect cooling; the latter for cooling at the Landauer limit. Together, these conditions imply the following:

Corollary 2. *To perfectly cool a target system with energy cost at the Landauer limit using a thermal machine $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$, the machine must be infinite dimensional and $\omega_{\mathcal{M}}^{\max}$, the maximal energy gap of $H_{\mathcal{M}}$, must diverge.*

The unbounded structural properties of the machine support the *possibility* for perfect cooling at the Landauer limit;

we now move to focus on the control properties of the interaction that *realise* said potential (see Fig. 2). This leads to the distinct notion of control complexity, which differentiates between protocols that access the machine in a more or less complex manner. The structural complexity properties are protocol independent and related to the energy spectrum and dimensionality of the machine, whereas the control complexity concerns dynamical properties of the unitary that represents a particular protocol.

2. Control Complexity

Although it is intuitive that a unitary coupling the system to many degrees of freedom of the machine should be considered complex, it is *a priori* unclear how to quantify control complexity in a manner that both

1. corresponds to our intuitive understanding of the word “complex”, meaning “difficult to implement”; and
2. is consistent with Nernst’s third law in the sense that its divergence is necessary to reach a pure state (when all other considered resources are restricted to be finite).

Many notions of complexity put forth throughout the literature to capture the first point above do not necessarily satisfy the second, as we discuss later. Here, we take the opposite approach and seek a *minimal* notion of complexity that is first and foremost consistent with the third law of thermodynamics, which we hope to develop further to incorporate the idea of quantifying how difficult a protocol is to implement.

In the following sections, we begin by demonstrating that any cooling protocol that achieves perfect cooling with minimal time and energy resources requires coupling the target system to an infinite-dimensional machine, thereby capturing a notion of control complexity that satisfies the second point above. However, by subsequently analysing the sufficient conditions for such optimal cooling, we see that such a condition is in general insufficient to achieve said goal; furthermore, coupling to an infinite-dimensional machine is not necessarily difficult to implement in practice in certain experimental platforms. The insights gained here finally motivate our more refined notion of control complexity, namely that the system must be coupled to a spectrum of machine energy gaps that approximate a continuum. This condition is indeed difficult to achieve in all experimental settings and therefore provides a reasonable definition of control complexity inasmuch as it satisfies both desiderata outlined above.

B. Effective Dimension as a Notion of Control Complexity

As a first step in this direction, a good proxy measure of control complexity is the effective dimension of a unitary operation, i.e., the dimension of the subspace of the global Hilbert space upon which the unitary acts nontrivially.

Definition 1. The *effective dimension* is the minimum dimension of a subspace \mathcal{A} of the joint Hilbert space $\mathcal{H}_{S\mathcal{M}}$ in terms of which the unitary can be decomposed as $U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}$:

$$d^{\text{eff}} := \min \dim(\mathcal{A}) : U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}. \quad (6)$$

Intuitively, given any (sufficiently large) machine dimension, the effective dimension captures how much of the machine takes part in the controlled interaction. While any dynamics that requires a high amount of control must accordingly have large effective dimension, the converse does not necessarily follow: there exist dynamics with corresponding large (even infinite) effective dimensions (e.g., Gaussian operations on two harmonic oscillators, such as those enacted by a beam splitter) that are easily implementable and do not require high levels of control, as we discuss further below. Nevertheless, using the definition above, we show that any protocol achieving perfect cooling at the Landauer limit necessarily involves interactions between the target and infinitely many energy levels of the machine. In other words, no interaction restricted to a finite-dimensional subspace suffices.

We begin by demonstrating that the effective dimension (nontrivially) accessed by a unitary (see Def. 1) must diverge to achieve perfect cooling at the Landauer limit, thereby providing a good proxy for control complexity in the sense that it aligns with Nernst’s third law and provides a necessary condition. Intuitively, the effective dimension of a unitary operation is the dimension of the subspace of the global Hilbert space upon which the unitary acts nontrivially, in other words the part of the joint space that is actually accessed by the control protocol. This quantity can be computed by considering a given cooling protocol and finite unit of time T (which we can set equal to unity without loss of generality) with respect to which the target and total machine transform unitarily by decomposing the Hamiltonian in $U_{S\mathcal{M}} = e^{-iH_{S\mathcal{M}}T}$ in terms of local and interaction terms, i.e., $H_{S\mathcal{M}} = H_S \otimes \mathbb{1}_{\mathcal{M}} + \mathbb{1}_S \otimes H_{\mathcal{M}} + H_{\text{int}}$. The effective dimension then corresponds to $\text{rank}(H_{\text{int}})$. With this definition at hand, we have the following:

Theorem 3. *The unitary representing a cooling protocol that saturates the Landauer limit must act nontrivially on an infinite-dimensional subspace of $\text{supp}(H_{\mathcal{M}})$. This implies $d^{\text{eff}} \rightarrow \infty$.*

Intuitively, we show that if a protocol accesses only a finite-dimensional subspace of the machine, then the machine is effectively finite-dimensional inasmuch as a suitable replacement can be made while keeping all quantities relevant for cooling invariant. Invoking the main result of Ref. [29] then implies that there are finite-dimensional correction terms such that the Landauer limit cannot be saturated.

The effective dimension therefore provides a minimal quantifier for control complexity: it is the quantity that must diverge in order to (perfectly) cool at minimal energy cost—thus, it satisfies the above point 2. Moreover, it requires no assumption on the underlying structure of the machine, with the results holding for either collections of finite-dimensional systems or harmonic oscillators without any restrictions on the types of individual operations allowed. This highlights a

certain level of generality regarding the definition put forth, inasmuch as it is not tied to any presupposed structure of the systems at hand or the ability of the agent to control them. Additionally, as we discuss below, in many situations of interest, such as a machine comprising a collection of qubits and/or natural gate-set limitations, said definition also corresponds to protocols that are difficult to implement in practice, therefore also satisfying the above point 1. However, such additional restrictions are by no means generic. Moreover, it is *a priori* unclear if having a diverging effective dimension is enough to permit perfect cooling with minimal time and energy cost. We now move on to discuss the connection to practical difficulty in general before analysing sufficient conditions regarding control complexity.

1. Correspondence to Practical Difficulty

Importantly, if one supposes that the system and machines are finite dimensional, then diverging effective dimension implies diverging circuit complexity, where the latter is defined in terms of the minimum number of gates (from a predetermined set of possibilities) required to implement the overall circuit representing a particular protocol. For instance, considering a qubit system and machines, and the ability to perform arbitrary two-qubit gates, the effective dimension is simply the logarithm of the number of distinct machine qubits that the system interacts with throughout the protocol. For any cooling protocol that achieves Landauer energy cost, it is clear that every one of a diverging number of qubit machines must take part in the overall transformation. Moreover, the particular interactions applied can be taken to be *SWAP* gates, which require the ability for the agent to be able to perform a *CNOT* gate, which in turn permits universal quantum computation with two-qubit interactions. Thus, given a universal two-qubit gate set, the circuit required to perform perfect cooling at minimal energy cost has a complexity that scales with the number of machine qubits. For higher-dimensional architectures or further restrictions on the gate set, any meaningful notion of control complexity will increase accordingly. This means that the task of cooling a finite-dimensional system with finite-dimensional machines at the Landauer limit is—even with a perfect quantum computer—an impossibly difficult task.

However, although our proposed definition of effective dimension as a notion of control complexity is flexible inasmuch as it applies to arbitrary system-machine structures, the price of such generality comes with the drawback that it tends to overestimate the difficulty of implementing a particular protocol in practice. In other words, without imposing any additional assumptions regarding the situation at hand, the effective dimension does not necessarily satisfy the above point 1. For example, whilst the effective dimension and the circuit complexity coincide for qubits, in higher-dimensional settings, the former overestimates the latter because not all system-machine subspaces are necessarily required to implement a particular protocol (i.e., although using all such subspaces provides one way to achieve it, this is not unique). Thus, the extent to which the circuit complexity is overes-

timated depends on the allowed gate set that is considered “simple” in general. At the extreme end, i.e., for harmonic oscillator systems and machines, this can be seen from the fact that a single beam-splitter operation (which is a two-mode Gaussian operation, corresponding to a simple circuit complexity in the usual sense considered for infinite-dimensional quantum circuit architectures) already has infinite effective dimension, but is far from sufficient to achieve perfect cooling at Landauer cost.

As a representative for infinite-dimensional systems, we treat harmonic oscillator target systems separately in Appendix E. In the infinite-dimensional setting, the difficulty of implementing an operation is often related to the polynomial degree of its generators. Here, we see some friction with respect to Eq. (6): as mentioned above, a generic Gaussian unitary operation (i.e., one generated by a Hamiltonian at most quadratic in the mode operators) between a harmonic oscillator target and machine already implies infinite effective dimensionality. In light of this, we first construct a protocol that achieves perfect cooling at the Landauer limit with diverging time using only sequences of Gaussian operations [i.e., those typically considered to be practically easily implementable (cf. Refs. [22, 37]), but nonetheless with infinite effective dimensionality according to Def. 1]. This result highlights that the polynomial degree of the generators of a particular protocol would—somewhat counterintuitively, since operations corresponding to high polynomial degree are difficult to achieve in practice—*not* provide a suitable measure of control complexity inasmuch as its divergence is not necessary for Landauer-cost cooling. In contrast, we then present a protocol that demonstrates that perfect cooling is possible given diverging time and operations acting on only a finite effective dimensionality (i.e., using non-Gaussian operations), with a finite energy cost that is greater than the Landauer limit; whether or not a similar protocol that saturates the Landauer limit exists in this setting remains an open question.

2. Sufficiency for Optimal Cooling

Thus, in general, accessing an infinite-dimensional machine subspace is not sufficient for reaching the Landauer limit. Indeed, in all of the protocols that we present, the degrees of freedom of the machine must be individually addressed in a fine-tuned manner to permute populations optimally, which intuitively corresponds to complicated multipartite gates and demonstrates that an operationally meaningful notion of control complexity must take into account factors beyond the effective dimensionality accessed by an operation. In particular, the interactions couple the target system to a diverging number of subspaces of the machine corresponding to distinct energy gaps. Moreover, there are a diverging number of energy levels of the machine both above and below the first excited level of the target. These observations highlight that fine-tuned control plays an important role. Indeed, both the final temperature of the target as well as the energy cost required to achieve this depends upon how the global eigenvalues are permuted via the cooling process. First, how cool

the target becomes depends on the sum of the eigenvalues that are placed into the subspace spanned by the ground state. Second, for any fixed amount of cooling, the energy cost depends on the constrained distribution of eigenvalues within the machine. Thus, in general, the optimal permutation of eigenvalues depends upon properties of both the target and machine. To highlight this, in Appendix D, we consider the task of cooling a maximally mixed target system with the additional constraint that the operation implemented lowers the temperature as much as possible. This allows us to derive a closed-form expression for the distribution of machine eigenvalues alone that must be asymptotically satisfied as the machine dimension diverges. Drawing from these insights, in the coming section we propose a stronger notion of control complexity (in the sense that it bounds the effective dimension from below and that it corresponds to practical difficulty in virtual every setting imaginable) in terms of the energy-gap structure of the machine and demonstrate that this measure too must diverge to cool perfectly with minimal time and energy costs. This concept is even more important in the case where all resources are finite, as particular structures of machines and the types of interactions permitted play a crucial role in both how much time or energy is spent cooling a system and how cold the system can ultimately become (see, e.g., Refs. [38–40]).

C. Energy-Gap Variety as a Notion of Control Complexity

This analysis motivates searching for a more detailed notion of control complexity that takes the energy-level structure of the machine into account, which should hold across all platforms and dimension scales. The discussion above illustrates some key challenges in defining a measure of control complexity that satisfies natural desiderata: such a measure should correspond to the difficulty of implementing operations in practice and simultaneously cover all possible physical platforms, including finite-dimensional systems such as, e.g., specific optical transitions of electrons in the shell of trapped ions, and infinite-dimensional systems such as the state-space-specific modes of the electromagnetic field. The effective dimension that we introduce above as a proxy manages to cover all such systems and provides a rigorous mathematical criterion that every physical protocol will necessarily have to fulfil in order to cool at minimal energy cost. As we have seen, however, infinite effective dimension is insufficient for cooling at the Landauer limit and it may not be all that difficult to achieve in continuous-variable setups. This begs the question of how this minimal definition of control complexity can be extended in order to more faithfully represent what permits saturation of the ultimate limitations and is difficult to achieve in practice.

Looking at all of our cooling protocols, a common property that seems to be important in minimising the energy cost of cooling is that the system is coupled to a set of machine energy gaps that are distributed in such a way that they (approximately) densely cover the interval $[\omega_1, \omega^*]$, where ω_1 is the first energy gap of the target system and ω^* is the maximal energy gap, which sets the final achievable temperature of the

system (for perfect cooling to the ground state, note that one requires $\omega^* \rightarrow \infty$). Let us denote the number of distinct energy gaps in a (fixed) interval as the *energy-gap variety*. More formally, we have the following:

Definition 2. Consider an interval $[\omega_a, \omega_b) \subseteq \mathbb{R}$. We define the energy-gap variety in terms of the set of machine energy gaps that lie in said interval, i.e., first construct the set

$$\mathcal{E}_{[\omega_a, \omega_b)} := \{\omega_\gamma := \omega_i - \omega_j \mid \omega_i - \omega_j \in [\omega_a, \omega_b)\}_{\gamma}. \quad (7)$$

The number of distinct elements in such a set is the *energy-gap variety*.

On the one hand, it is clear that coupling a system to a large number of distinct and/or closely spaced energy gaps requires fine-tuned control that is difficult in any experimental setting. On the other, the energy-gap variety lower bounds the effective dimension, and thus it is not clear that it needs to diverge in order to cool at Landauer energy cost. In Appendix D, we demonstrate that the energy-gap variety must indeed diverge and, additionally, that the set of energy gaps must densely cover a relevant interval (whose endpoints set the amount of cooling possible) in order to perfectly cool at the Landauer limit by proving the following:

Theorem 4. *In order to cool $\varrho_S \mapsto |0\rangle\langle 0|$ with a thermal machine $\tau_M(\beta, H_M)$ at Landauer energy cost with a single control operation, the global unitary U must couple the system to a diverging number of distinct energy gaps that densely cover the interval $[\omega_0, \infty)$, where ω_0 is the smallest energy gap of the target system.*

Taken in combination with its sufficiency to achieve said task, this result posits the energy-gap variety as a better quantifier of control complexity than the effective dimension, constituting the best thermodynamically meaningful notion of control complexity that we have put forth so far.

The above theorem establishes the relevance of the energy-gap variety regarding the ultimate limitations of perfect cooling. In reality, of course, experimental imperfections abound, and so naturally the question arises: *how robust is the energy-gap variety and to what extent can it incorporate errors?* Regarding the former: note that the above theorem posits the impossibility of cooling at Landauer energy cost unless one has control over an (infinitely) fine-grained energy-gap structure. Any perturbation away from said structure will result in some additional energy requirement for cooling; however, intuitively, small perturbations will correspond to small increases in energy costs. Properly accounting for such impacts, e.g., by bounding the additional energy cost in terms of a difference from the optimal energy-gap structure, is an important next step to understand the practical limitations of cooling. Regarding the latter point, in reality one never has perfect control over microscopic degrees of freedom. For instance, an immediate experimental imperfection that should be accounted for is the fact that two energy gaps which are very close together will be practically indistinguishable. Although a full-fledged error analysis here would constitute a major follow-up work, note that such cases can be formally

dealt with within our framework by suitably modifying the definition, i.e., by discretising energy bands to suitably capture the indistinguishability of energy gaps and/or error margins.

Aside from introducing and highlighting the important role of control complexity, we now take a step back to consider the notion of overall control at a higher level. It is clear that the protocols that saturate the Landauer limit for the energy cost of cooling require highly controlled microstate interactions between the system and machine; in turn, such transformations necessitate that the agent has access to a versatile *work source*, i.e., either a quantum battery [19–23] or a classical work source with a precise clock [9, 10]. Such control is reminiscent of Maxwell’s demon, who can indeed address all microscopic configurations at hand. This level of control is, however, in some sense at odds with the true spirit of thermodynamics. Indeed, the very reason that the machine is taken to begin as a thermal (Gibbs) state in thermodynamics is precisely because it provides the microscopic description that is *both* consistent with macroscopic observations (in particular, average energy) and makes minimal assumptions regarding the information that the agent has about the initial state; thermodynamics as a whole is largely concerned with what can be done with minimal information requirements. Beginning with this, and then going on to permit dynamical interactions that address the full complex microstructure is somewhat contradictory, at least in essence; indeed, it has been argued that “Maxwell’s demon cannot operate” [41] as an autonomous thermal being. Thus, a more thermodynamically sound setting would also restrict the transformations themselves to be ones that can be driven with minimal overall control. We now move to analyse the task of cooling within such a context.

V. INCOHERENT CONTROL (HEAT ENGINE)

The results presented so far pertain to cooling with the only restriction being that the machines are initially thermal. In particular, there are no restrictions on the allowed unitaries. In general, the operations required for cooling are not energy conserving and require an external work source. With respect to standard considerations of thermodynamics, this may seem somewhat unsatisfactory, as the joint system is, in the coherent setting, open to the universe. When quantifying thermodynamic resources, one typically restricts the permitted transformations to be energy conserving, thereby closing the joint system and yielding a self-contained theory.

We therefore analyse protocols using energy-conserving unitaries. With this restriction, it is in general not possible to cool a target system with machines that are initially thermal at a single temperature, as was considered in the coherent-control paradigm [42]. Instead, cooling can be achieved by partitioning the machine into one cold subsystem \mathcal{C} that begins in equilibrium at inverse temperature β and another hot subsystem \mathcal{H} coupled to a heat bath at inverse temperature $\beta_H < \beta$ [38, 42] (see Fig. 1, bottom panel). In other words, one uses a hot and a cold bath to construct a heat engine that cools the target. As we demonstrate, perfect cooling can be

achieved in this setting as pertinent resources diverge. However, the structure of the hot bath plays a crucial role regarding the resource requirements. In particular, we present a no-go theorem that states that perfect cooling with a heat engine using a single unitary of finite control complexity is impossible, even given diverging energy drawn from the hot bath. This result is in stark contrast to its counterpart in the coherent-control setting, where diverging energy is sufficient for perfect cooling and serves to highlight the fact that the incoherent-control setting is a fundamentally distinct paradigm that must be considered independently. Here, we focus on finite-dimensional systems and leave the analysis of infinite-dimensional ones to future work.

A. Ultimate Limits in the Incoherent Control Paradigm

In the incoherent-control setting, an adaptation of the (equality-form) Landauer bound on the minimum heat dissipated (or, as we phrase it here, the minimum amount of energy drawn from the hot bath) can be derived, which we dub the *Carnot-Landauer limit*:

Theorem 5. *Let $F_\beta(\varrho_x) := \text{tr}[H_x \varrho_x] - \beta^{-1}S(\varrho_x)$ be the free energy of a state ϱ_x with respect to a heat bath at inverse temperature β , $\Delta F_s^{(\beta)} := F_\beta(\varrho'_s) - F_\beta(\varrho_s)$, and let $\eta := 1 - \beta_H/\beta \in (0, 1)$ be the Carnot efficiency with respect to the hot and cold baths. In the incoherent-control setting, the quantity*

$$\Delta F_s^{(\beta)} + \eta \Delta E_{\mathcal{H}} = -\frac{1}{\beta}[\Delta S_s + \Delta S_c + \Delta S_{\mathcal{H}} + D(\varrho'_c || \varrho_c) + D(\varrho'_H || \varrho_H)] \quad (8)$$

satisfies the inequality

$$\Delta F_s^{(\beta)} + \eta \Delta E_{\mathcal{H}} \leq 0. \quad (9)$$

Equation (9) holds due to the non-negativity of the sum of local entropy changes and the relative-entropy terms. The derivation is provided in Appendix A, where we also show that the usual Landauer bound is recovered in the limit of an infinite-temperature heat bath.

The incoherent-control setting is fundamentally distinct from the coherent-control setting in terms of what can (or cannot) be achieved with given resources. For instance, consider the case where one wishes to achieve perfect cooling in unit time and with finite control complexity with diverging energy cost. In the coherent-control setting, this task is possible in principle (see Theorem 1). On the other hand, in the incoherent-control setting, we have the following no-go theorem (see Appendix F for a proof):

Theorem 6. *In the incoherent control scenario, it is not possible to perfectly cool any quantum system of finite dimension in unit time and with finite control complexity, even given diverging energy drawn from the hot bath, for any non-negative inverse temperature heat bath $\beta_H \in [0, \beta < \infty)$.*

This result follows from the fact that in the incoherent-control setting, the target system can only interact with subspaces of the joint hot-and-cold machine with respect to which it is energy degenerate. For any operation of fixed control complexity, there is always a finite amount of population remaining outside of the accessible subspace, implying that perfect cooling cannot be achieved, independent of the amount of energy drawn from the hot bath.

B. Saturating the Carnot-Landauer Limit

The above result emphasises the difference between coherent and incoherent controlling, which means that it is *a priori* unclear if the Carnot-Landauer bound is attainable and, if so, how to attain it. Indeed, the restriction to energy-conserving unitaries generally makes it difficult to tell if the ultimate bounds can be saturated in the incoherent-control setting, and which resources would be required to do so. We present a detailed study of cooling in the incoherent-control setting in Appendix F, where we prove the following results. We begin by demonstrating incoherent cooling protocols that saturate the Landauer bound in the regime where the heat-bath temperature goes to infinity. We do so by fine tuning the machine structure such that the desired cooling transitions between the target system and the cold and hot parts of the machine are rendered energy conserving. In particular, we prove the following:

Theorem 7. *In the incoherent control scenario, for an infinite-temperature hot bath $\beta_H = 0$, any finite-dimensional system can be perfectly cooled at the Landauer limit with diverging time via interactions of finite control complexity. Similarly, the goal can be achieved in unit time with diverging control complexity.*

Following our analysis of infinite-temperature heat baths, we study the more general case of finite-temperature heat baths. In Appendix G, we detail cooling protocols that saturate the Carnot-Landauer limit for any finite-temperature heat bath. More precisely, we prove:

Theorem 8. *In the incoherent control scenario, for any finite-temperature hot bath $0 < \beta_H < \beta$, any finite-dimensional quantum system can be perfectly cooled at the Carnot-Landauer limit given diverging time via finite control complexity interactions. Similarly, the goal can be achieved in unit time with diverging control complexity.*

As in the coherent-control setting, these protocols use either diverging time or control complexity to asymptotically saturate the Carnot-Landauer bound. The results presented in this section therefore provide a comprehensive understanding of the resources required to perfectly cool at minimum energy cost in a setting that aligns with the resource theories of thermodynamics.

VI. IMPERFECT COOLING WITH FINITE RESOURCES

The above results set the ultimate limitations for cooling inasmuch as the protocols saturate optimal bounds by using diverging resources. In reality, however, any practical implementation is limited to having only finite resources at its disposal. According to the third law, a perfectly pure state cannot be achieved in this scenario. Nonetheless, one can prepare a state of finite temperature by investing said resources appropriately. In this finite-resource setting, the interplay between energy, time, and control complexity is rather complicated. First, the cooling performance is stringent upon the chosen figure of merit for the notion of cool—the ground-state population, purity, average energy, or temperature of the nearest thermal state are all reasonable candidates, but they differ in general [42]. Second, the total amount of resources available bounds the reachable temperature in any given protocol. Third, the details of the protocol itself influence the energy cost of achieving a desired temperature. In other words, determining the optimal distribution of resources is an extremely difficult task in general and remains an open question.

We therefore focus here on the paradigmatic special case of cooling a qubit target system by increasing its ground-state population in order to highlight some salient points regarding cooling to finite temperatures. First, we compare the finite performance of two distinct coherent control protocols that both asymptotically saturate the Landauer limit; nonetheless, at any finite time, their performance varies. The first protocol simply swaps the target qubit with one of a sequence of machine qubits whose energy gaps are distributed linearly; the second involves interacting the target with a high-dimensional machine with a particular degeneracy structure. Although the latter cannot be decomposed easily into a qubit circuit (thereby making it more difficult to implement in practice), one can compare the two protocols fairly by fixing the total (and effective) dimension to be equal, i.e., comparing the performance of the linear sequential qubit machine protocol after $N + 1$ qubits have been accessed with that of the latter protocol with machine dimension 2^{N+1} . In doing so, we see that the simpler former protocol outperforms the more difficult latter one in terms of the energy cost at finite times, emphasising the fact that difficulty in practice does not necessarily correspond to complexity as a thermodynamic resource. Additionally, we analyse the cooling rates at which energy and time can be traded off amongst each other in the linear qubit sequence protocol by deriving an analytic expression. Lastly, we compare the performance of a coherent and an incoherent control protocol that use a similar machine structure to achieve a desired final temperature. We see that the price one must pay for running the protocol via a heat engine is that either more steps or more complex operations are required to match the performance of the coherent control setting. This example serves to elucidate the connection between the two extremal control scenarios relevant for thermodynamics.

Although throughout most of the paper we focus on the asymptotic achievability of optimal cooling strategies, the protocols that we construct provide insight into how said asymptotic limits are approached. This facilitates a better un-

derstanding of the more practically relevant questions that are constrained when all resources are restricted to be finite: i) *how cold can the target system be made?* and ii) *at what energy cost?* In line with Nernst's third law, the answer to the former question cannot be perfectly cold (i.e., zero temperature). The answer depends upon how said resources are configured and utilized. For instance, given a single unitary interaction of finite complexity in the coherent-control setting, the ground-state population of the output state can be upper bounded in terms of the largest energy gap of the machine, ω_{\max} [see Eq. (5)]. On the other hand, supposing that one can reuse a single machine system multiple times, then as the number of operation steps increases, the ground-state population of the output state approaches $(1 + e^{-\beta\omega_{\max}})^{-1}$ from below [42]. There is clearly a trade-off relation here between time and complexity, and a systematic analysis of the rate at which these quantities can be traded off against one another warrants further investigation. Similarly, the energy cost to reach a desired final temperature also depends upon the distribution of resources, as we now examine.

Given access to a machine of a certain size (as measured by its dimension), one could ask: *what is the optimal configuration of machine energy spectrum and global unitary to cool a system as efficiently as possible?* Here, we compare two contrasting constructions for the cooling unitary in the coherent-control setting for a qubit target system (with energy gap ω_s)—both of which asymptotically achieve Landauer cost cooling, but whose finite behaviour differs. The first protocol considers a machine of N qubits whose energy gaps increase linearly from the first excited state energy level of the system $\omega_1 = \omega_s$ to some maximum energy level $\omega_N = \omega_{\max}$, which dictates the final achievable temperature. In this protocol, the target system is swapped sequentially with each of the N qubits in order of increasing energy gaps; we hence refer to it as the *linear qubit machine sequence*. The second protocol we consider is presented in full in Appendix D4 and inspired by one presented in Ref. [29] (see Appendix D therein); we hence refer to it as the Reeb & Wolf (RW) protocol. Here, the global unitary acts on the system and a high-dimensional machine with an equally spaced Hamiltonian whose degeneracy doubles with each increasing energy level, i.e., it has a singular ground state, a twofold degenerate first excited state, a fourfold degenerate second excited state, and so on; the final energy level has an extra state so that the total dimension is 2^{N+1} (where N is the number of energy levels). In particular, the unitary performs the permutation that places the maximal amount of population in the ground state of the target system. Due to the structure of both protocols, one can make a fair comparison between them, contrasting the single unitary on a 2^N -dimensional machine in the RW protocol versus the composition of N two-qubit SWAP unitaries in the linear machine sequence, i.e., such that both protocols access a machine of the same size overall.

As shown in Fig. 3, although both protocols asymptotically tend to the Landauer limit, their finite behaviour differs. Indeed, the work cost of the linear qubit machine sequence protocol outperforms that of the RW protocol. This is somewhat surprising, as the latter is a complex high-dimensional unitary




FIG. 3. *Imperfect Cooling.* We compare the cooling performance of a degenerate qubit target system using either N machine qubits of linearly increasing energy accessed sequentially or a single unitary on a 2^N -dimensional machine, the latter being a finite adaptation of a protocol presented in Ref. [29]. We set $\beta = 1$, choose units such that $\hbar = k_B = 1$, and fix $1 - \epsilon$ to be the desired final ground-state population of the target. We plot the inverse of the excess work cost above the Landauer limit, $W - \beta\Delta S_s$ (in units of the smallest machine energy gap, $\omega_{\mathcal{M}}^{\min}$), confirming that the surplus work cost in both cases scales with N^{-1} . Interestingly, we see that the protocol in which the target is sequentially swapped with machine qubits outperforms that which uses a high-dimensional unitary (at equal overall control complexity) in terms of energy cost required to reach a desired temperature.

whereas the former a composition of qubit swaps; although both protocols have the same effective dimension in this comparison overall, this highlights that difficulty in the lab setting need not correspond to resourcefulness in a thermodynamic sense. Indeed, developing optimal finite cooling strategies for arbitrary systems and machines is difficult in general and remains an important open question. Nonetheless, in Appendix H, we derive the rate of resource divergence of the sequential qubit protocol to further clarify the trade-off between time and energy for this protocol.

Finally, we contrast the two extremal thermodynamic paradigms considered by comparing the energy cost of a coherently controlled cooling protocol to an incoherently controlled one that achieves the same final ground-state population. Intuitively, the latter setting requires more resources to achieve the same performance as the former due to the fact that only energy-resonant subspaces can be accessed by the unitary, and hence only a subspace of the full machine is usable. This implies that a greater number of operations (of fixed control complexity) are required to achieve similar results as the coherent setting, as demonstrated in Appendix H explicitly. Indeed, determining the optimal cooling protocols for a range of realistic assumptions remains a major open avenue.

VII. DISCUSSION

Relation to Previous Works

A vast amount of the literature concerning quantum thermodynamics considers resource theories (see Refs. [43, 44] and references therein), whose central question is: *what transformations are possible given particular resources, and how can one quantify the value of a resource?* While this perspective sheds light on what is possible in principle, it does not per se concern itself with the potential implementation of said transformations. Yet, the unitary operations considered in a resource theory will themselves require certain resources to implement in practice. Focusing only on a resource-theoretic perspective would thus overlook the question: *how does one optimally use said resources?* Our results focus on this latter question and highlight the role of control complexity in optimising resource use.

Concurrently, by considering arbitrary unitary operations (akin to our coherent-control paradigm without limitations on machine size) Refs. [20, 36] and [29], studied the potential saturation of the second law of thermodynamics and Landauer’s limit, respectively. References [20] and [36] develop a similar protocol to our diverging time protocol in the context of work extraction and demonstrate its optimality for saturating the second law. However, these works do not discuss the practical viewpoint that the goal can be achieved in a smaller number of operations by allowing the latter to be more complex, as we emphasise. On the other hand, Ref. [29] considers the resources required for saturation of the Landauer limit and show an important result regarding structural complexity, namely that the machine must be infinite dimensional to cool at the Landauer limit. Our analysis regarding complexity begins here and continues to elucidate the key complexity properties that enhance the efficiency of a cooling protocol. In particular, we show that an infinite-dimensional machine is not sufficient unless the controlled unitary indeed accesses the entire machine. This first leads to the notion of “effective dimension”, which provides a good proxy for control complexity that is consistent with Nernst’s third law for all types of quantum machines—from finite-dimensional systems to harmonic oscillators. Moreover, we highlight that the optimal interactions must be fine tuned, i.e., they must couple the system to particular energy gaps of the machine in a specific configuration, paving the way for a more nuanced definition of control complexity that takes into account the complicated and precise level of control required, as we present in terms of the “energy-gap variety”. Lastly, we emphasise that the latter discussion concerns the coherent-control scenario, which is only one of the extremal control paradigms that we consider. In addition, we consider the task of cooling in a more thermodynamically consistent setting, namely the incoherent-control paradigm. There we derive the Carnot-Landauer equality and consequent inequality, which are adaptations of the Landauer equality [29] and inequality [11], respectively, where the protocol can only be run via a heat engine.

On the more practical side, note that our work here concerns erasing quantum information encoded in fundamental

rather than logical degrees of freedom. Our reasoning here is twofold: firstly, the ultimate limitations that we aim to understand are the same whether one wishes to cool a physical system or erase information; in other words, although it may be possible to save some finite trade-off costs for imperfect erasure in the coarse-grained setting, the resources required to perform a rank-reducing process asymptotically diverge in both cases. Secondly, it is much more difficult to create coherent superpositions in the case where information is redundantly encoded in macrostates, as this would require all microstates to be in phase (indeed, this is a major reason why quantum computers aim to encode information in fundamental degrees of freedom). For erasing quantum information using bulk (classical) cooling (i.e., coupling to a suitably engineered cold bath), the relevant condition is nondegeneracy of the ground state; additionally, many original Landauer thought experiments consider degenerate Hamiltonians for the computational states. In contrast, our protocols are based upon directly controlled cooling, which works independently of the target system Hamiltonian and as such bridges the gap between various perspectives. Moving forward, it will be interesting to explore how information can be erased cheaper if it is encoded in a coarse-grained fashion, in order to better square our fundamental results presented here with experimental demonstrations. Doing so would require finite versions of all of the systems and resources that we analyse here, which we leave for future exploration.

Conclusions & Outlook

The results of this work have wide-ranging implications. We have both generalised and unified Landauer’s bound with respect to the laws of thermodynamics. In particular, we have posed the ultimate limitations for cooling quantum systems or erasing quantum information in terms of resource costs and presented protocols that asymptotically saturate these limits. Indeed, while it is well known that heat and time requirements must be minimised to combat the detrimental effects of fluctuation-induced errors and short decoherence times on quantum technologies [18], we have shown that this comes at a practical cost of greater control. In particular, we have demonstrated the necessity of implementing fine-tuned interactions involving a diverging number of energy levels to minimise energy and time costs, which serves to deliver a cautionary message: control complexity must be accounted for to build operationally meaningful resource theories of quantum thermodynamics. This result posits the energy-gap variety accessed by a unitary protocol as a meaningful quantifier of control complexity that is both fully consistent with the third law of thermodynamics and chimes well with what is difficult to achieve in practice. Our analysis of the incoherent-control setting further provides pragmatic ultimate limitations for the scenario where minimal control is required, in the sense that all transformations are driven by thermodynamic energy and entropy flows between two heat baths, which could be viewed as a thermodynamically driven quantum computer [45]. Nevertheless, the intricate relationship between various resources

here will need to be further explored.

Looking forward, we believe it will be crucial to go beyond asymptotic limits. While Landauer erasure and the third law of thermodynamics conventionally deal with the creation of pure states, practical results would need to consider cooling to a finite temperature (i.e., creating approximately pure states) with a finite amount of invested resources [38–40]. In this context, the trade-off between time and control complexity will gain more practical relevance, as realistic quantum technologies have limited coherence times and interaction Hamiltonians are limited to few-body terms. Here, operational measures of control complexity that fit the envisioned experimental setup present an important challenge that must be overcome to apply our results across various platforms.

Our results strengthen the view that, in contrast to classical thermodynamics, the role of control is one of the most crucial issues to address before a true understanding of the limitations and potential of quantum machines is revealed. On the one hand, in classical systems, control is only ever achieved over few bulk degrees of freedom, whereas addressing and designing particular microstate control is within reach of current quantum technological platforms, offering additional routes towards operations enhanced by fine-tuned control. On the other hand, the cost of such control itself can quickly exceed the energy scale of the system, potentially rendering any perceived advantages a mirage. This is exacerbated by the fact that it is not possible to observe (measure) a quantum machine without incurring significant additional thermodynamic costs [8, 46] and non-negligible backaction on the operation of the machine itself [47]. A fully developed theory of quantum thermodynamics would need to take these into account and we hope that our study sheds light on the role of control complexity in this endeavour.

ACKNOWLEDGMENTS

The authors thank Elizabeth Agudelo and Paul Erker for very insightful discussions at the early stages of the project.

P.T. acknowledges support from the Austrian Science Fund (FWF) project: Y879-N27 (START), the European Research Council (Consolidator grant ‘Cocoquest’ 101043705), and the Japan Society for the Promotion of Science (JSPS) by KAKENHI Grant No. 21H03394. F.B. is supported by FQXi Grant No. FQXi-IAF19-07 from the Foundational Questions Institute Fund, a donor advised fund of Silicon Valley Community Foundation. A.B. acknowledges support from the VILLUM FONDEN via the QMATH Centre of Excellence (Grant no. 10059) and from the QuantERA ERA-NET Cofund in Quantum Technologies implemented within the European Union’s Horizon 2020 Programme (QuantAlgo project) via the Innovation Fund Denmark. R.S. acknowledges funding from the Swiss National Science Foundation via an Ambizione grant PZ00P2_185986. N.F. is supported by the Austrian Science Fund (FWF) projects: P 36478-N and P 31339-N27. M.P.E.L. acknowledges financial support by the ESQ (Erwin Schrödinger Center for Quantum Science & Technology) Discovery programme, hosted by the Austrian Academy of Sciences (ÖAW) and TU Wien. G.V. is supported by the Austrian Science Fund (FWF) projects ZK 3 (Zukunftskolleg) and M 2462-N27 (Lise-Meitner). F.C.B. acknowledges support from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 801110 and the Austrian Federal Ministry of Education, Science and Research (BMBWF). This project reflects only the authors’ view, the EU Agency is not responsible for any use that may be made of the information it contains. T.D. acknowledges support from the Brazilian agency CNPq INCT-IQ through the project (465469/2014-0). E.S. is supported by the Austrian Science Fund (FWF) project: Y879-N27 (START). F.C. is supported by the ERC Synergy grant HyperQ (Grant No. 856432). M.H. is supported by the European Research Council (Consolidator grant ‘Cocoquest’ 101043705), the Austrian Science Fund (FWF) project: Y879-N27 (START), and acknowledges financial support by the ESQ (Erwin Schrödinger Center for Quantum Science & Technology) Discovery programme, hosted by the Austrian Academy of Sciences (ÖAW).

-
- [1] Nicolas Gisin, Grégoire Ribordy, Wolfgang Tittel, and Hugo Zbinden, *Quantum cryptography*, *Rev. Mod. Phys.* **74**, 145 (2002), [arXiv:quant-ph/0101098](#).
 - [2] Stefano Pirandola, Ulrik L. Andersen, Leonardo Banchi, Mario Berta, Darius Bunandar, Roger Colbeck, Dirk Englund, Tobias Gehring, Cosmo Lupo, Carlo Ottaviani, Jason L. Pereira, Mohsen Razavi, Jesni Shamsul Shaari, Marco Tomamichel, Vladyslav C. Usenko, Giuseppe Vallone, Paolo Villoresi, and Petros Wallden, *Advances in quantum cryptography*, *Adv. Opt. Photon.* **12**, 1012 (2020), [arXiv:1906.01645](#).
 - [3] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone, *Advances in quantum metrology*, *Nat. Photonics* **5**, 222 (2011), [arXiv:1102.2318](#).
 - [4] Géza Tóth and Iagoba Apellaniz, *Quantum metrology from a quantum information science perspective*, *J. Phys. A: Math. Theor.* **47**, 424006 (2014), [arXiv:1405.4878](#).
 - [5] Rafał Demkowicz-Dobrzański, Marcin Jarzyna, and Janek Kłodyński, *Quantum limits in optical interferometry*, *Prog. Optics* **60**, 345 (2015), [arXiv:1405.7703](#).
 - [6] John Preskill, *Fault-tolerant quantum computation*, in *Introduction to Quantum Computation*, edited by H.-K. Lo, Sandu Popescu, and T. P. Spiller (World-Scientific, Singapore, 1997) Chap. 8, pp. 213–269, [arXiv:quant-ph/9712048](#).
 - [7] John Preskill, *Quantum Computing in the NISQ era and beyond*, *Quantum* **2**, 79 (2018), [arXiv:1801.00862](#).
 - [8] Yelena Guryanova, Nicolai Friis, and Marcus Huber, *Ideal Projective Measurements Have Infinite Resource Costs*, *Quantum* **4**, 222 (2020), [arXiv:1805.11899](#).
 - [9] Paul Erker, Mark T. Mitchison, Ralph Silva, Mischa P. Woods, Nicolas Brunner, and Marcus Huber, *Autonomous Quantum Clocks: Does Thermodynamics Limit Our Ability to Measure Time?* *Phys. Rev. X* **7**, 031022 (2017), [arXiv:1609.06704](#).
 - [10] Emanuel Schwarzhans, Maximilian P. E. Lock, Paul Erker, Nicolai Friis, and Marcus Huber, *Autonomous Tempo-*

- ral Probability Concentration: Clockworks and the Second Law of Thermodynamics*, *Phys. Rev. X* **11**, 011046 (2021), [arXiv:2007.01307](#).
- [11] Rolf Landauer, *Irreversibility and Heat Generation in the Computing Process*, *IBM J. Res. Dev.* **5**, 183 (1961).
 - [12] Walther Nernst, *Über die Beziehung zwischen Wärmeentwicklung und maximaler Arbeit bei kondensierten Systemen*, in *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften* (Berlin, 1906) pp. 933–940.
 - [13] Francesco Ticozzi and Lorenza Viola, *Quantum resources for purification and cooling: fundamental limits and opportunities*, *Sci. Rep.* **4**, 5192 (2014), [arXiv:1403.8143](#).
 - [14] Lluís Masanes and Jonathan Oppenheim, *A general derivation and quantification of the third law of thermodynamics*, *Nat. Commun.* **8**, 14538 (2017), [arXiv:1412.3828](#).
 - [15] Henrik Wilming and Rodrigo Gallego, *Third Law of Thermodynamics as a Single Inequality*, *Phys. Rev. X* **7**, 041033 (2017), [arXiv:1701.07478](#).
 - [16] Nahuel Freitas, Rodrigo Gallego, Lluís Masanes, and Juan Pablo Paz, *Cooling to Absolute Zero: The Unattainability Principle*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A. Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 25, pp. 597–622, [arXiv:1911.06377](#).
 - [17] Jakob Scharlau and Markus P. Müller, *Quantum Horn’s lemma, finite heat baths, and the third law of thermodynamics*, *Quantum* **2**, 54 (2018), [arXiv:1605.06092](#).
 - [18] Antonio Acín, Immanuel Bloch, Harry Buhrman, Tommaso Calarco, Christopher Eichler, Jens Eisert, Daniel Esteve, Nicolas Gisin, Steffen J. Glaser, Fedor Jelezko, Stefan Kuhr, Maciej Lewenstein, Max F. Riedel, Piet O. Schmidt, Rob Thew, Andreas Wallraff, Ian Walmsley, and Frank K. Wilhelm, *The quantum technologies roadmap: a European community view*, *New J. Phys.* **20**, 080201 (2018), [arXiv:1712.03773](#).
 - [19] J. Åberg, *Truly work-like work extraction via a single-shot analysis*, *Nat. Commun.* **4**, 1925 (2013), [arXiv:1110.6121](#).
 - [20] Paul Skrzypczyk, Anthony J. Short, and Sandu Popescu, *Work extraction and thermodynamics for individual quantum systems*, *Nat. Commun.* **5**, 4185 (2014), [arXiv:1307.1558](#).
 - [21] Matteo Lostaglio, David Jennings, and Terry Rudolph, *Description of quantum coherence in thermodynamic processes requires constraints beyond free energy*, *Nat. Commun.* **6**, 6383 (2015), [arXiv:1405.2188](#).
 - [22] Nicolai Friis and Marcus Huber, *Precision and Work Fluctuations in Gaussian Battery Charging*, *Quantum* **2**, 61 (2018), [arXiv:1708.00749](#).
 - [23] Francesco Campaioli, Felix A. Pollock, and Sai Vinjanampathy, *Quantum Batteries*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A. Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 8, pp. 207–225, [arXiv:1805.05507](#).
 - [24] Henry E. D. Scovil and Erich O. Schulz-DuBois, *Three-Level Masers as Heat Engines*, *Phys. Rev. Lett.* **2**, 262 (1959).
 - [25] Ronnie Kosloff and Amikam Levy, *Quantum Heat Engines and Refrigerators: Continuous Devices*, *Annu. Rev. Phys. Chem.* **65**, 365 (2014), [arXiv:1310.0683](#).
 - [26] Raam Uzdin, Amikam Levy, and Ronnie Kosloff, *Equivalence of Quantum Heat Machines, and Quantum-Thermodynamic Signatures*, *Phys. Rev. X* **5**, 031044 (2015), [arXiv:1502.06592](#).
 - [27] Mark T. Mitchison, *Quantum thermal absorption machines: refrigerators, engines and clocks*, *Contemp. Phys.* **60**, 164 (2019), [arXiv:1902.02672](#).
 - [28] Mischa P. Woods, Nelly Huei Ying Ng, and Stephanie Wehner, *The maximum efficiency of nano heat engines depends on more than temperature*, *Quantum* **3**, 177 (2019), [arXiv:1506.02322](#).
 - [29] David Reeb and Michael M. Wolf, *An improved Landauer principle with finite-size corrections*, *New J. Phys.* **16**, 103011 (2014), [arXiv:1306.4352](#).
 - [30] Fabien Clivaz, *Optimal Manipulation Of Correlations And Temperature In Quantum Thermodynamics*, *Ph.D. thesis*, University of Geneva (2020), [arXiv:2012.04321](#).
 - [31] Paul M. Riechers and Mile Gu, *Impossibility of achieving Landauer’s bound for almost every quantum state*, *Phys. Rev. A* **104**, 012214 (2021), [arXiv:2103.02337](#).
 - [32] Marcus Huber, Martí Perarnau-Llobet, Karen V. Hovhannisyán, Paul Skrzypczyk, Claude Klöckl, Nicolas Brunner, and Antonio Acín, *Thermodynamic cost of creating correlations*, *New J. Phys.* **17**, 065008 (2015), [arXiv:1404.2169](#).
 - [33] David E. Bruschi, Martí Perarnau-Llobet, Nicolai Friis, Karen V. Hovhannisyán, and Marcus Huber, *The thermodynamics of creating correlations: Limitations and optimal protocols*, *Phys. Rev. E* **91**, 032118 (2015), [arXiv:1409.4647](#).
 - [34] Giuseppe Vitagliano, Claude Klöckl, Marcus Huber, and Nicolai Friis, *Trade-off between work and correlations in quantum thermodynamics*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A. Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 30, pp. 731–750, [arXiv:1803.06884](#).
 - [35] Faraj Bakhshinezhad, Fabien Clivaz, Giuseppe Vitagliano, Paul Erker, Ali T. Rezakhani, Marcus Huber, and Nicolai Friis, *Thermodynamically optimal creation of correlations*, *J. Phys. A: Math. Theor.* **52**, 465303 (2019), [arXiv:1904.07942](#).
 - [36] Janet Anders and Vittorio Giovannetti, *Thermodynamics of discrete quantum processes*, *New J. Phys.* **15**, 033022 (2013), [arXiv:1211.0183](#).
 - [37] E. G. Brown, N. Friis, and Marcus Huber, *Passivity and practical work extraction using Gaussian operations*, *New J. Phys.* **18**, 113028 (2016), [arXiv:1608.04977](#).
 - [38] Fabien Clivaz, Ralph Silva, Géraldine Haack, Jonatan Bohr Brask, Nicolas Brunner, and Marcus Huber, *Unifying paradigms of quantum refrigeration: Fundamental limits of cooling and associated work costs*, *Phys. Rev. E* **100**, 042130 (2019), [arXiv:1710.11624](#).
 - [39] Philip Taranto, Faraj Bakhshinezhad, Philipp Schüttelkopf, Fabien Clivaz, and Marcus Huber, *Exponential Improvement for Quantum Cooling through Finite-Memory Effects*, *Phys. Rev. Appl.* **14**, 054005 (2020), [arXiv:2004.00323](#).
 - [40] Yi-Zheng Zhen, Dario Egloff, Kavan Modi, and Oscar Dahlsten, *Universal Bound on Energy Cost of Bit Reset in Finite Time*, *Phys. Rev. Lett.* **127**, 190602 (2021), [arXiv:2106.00580](#).
 - [41] L. Brillouin, *Maxwell’s Demon Cannot Operate: Information and Entropy. I*, *J. Appl. Phys.* **22**, 334 (1951).
 - [42] Fabien Clivaz, Ralph Silva, Géraldine Haack, Jonatan Bohr Brask, Nicolas Brunner, and Marcus Huber, *Unifying Paradigms of Quantum Refrigeration: A Universal and Attainable Bound on Cooling*, *Phys. Rev. Lett.* **123**, 170605 (2019), [arXiv:1903.04970](#).
 - [43] Nelly H. Y. Ng and Mischa P. Woods, *Resource Theory of Quantum Thermodynamics: Thermal Operations and Second Laws*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A. Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 26, pp. 625–650, [arXiv:1805.09564](#).
 - [44] Matteo Lostaglio, *An introductory review of the resource theory approach to thermodynamics*, *Rep. Prog. Phys.* **82**, 114001

- (2019), [arXiv:1807.11549](#).
- [45] Charles H. Bennett, *The thermodynamics of computation – a review*, *Int. J. Theor. Phys.* **21**, 905 (1982).
- [46] Tiago Debarba, Gonzalo Manzano, Yelena Guryanova, Marcus Huber, and Nicolai Friis, *Work estimation and work fluctuations in the presence of non-ideal measurements*, *New J. Phys.* **21**, 113002 (2019), [arXiv:1902.08568](#).
- [47] Gonzalo Manzano, Gian-Luca Giorgi, Rosario Fazio, and Roberta Zambrini, *Boosting the performance of small autonomous refrigerators via common environmental effects*, *New J. Phys.* **21**, 123026 (2019), [arXiv:1908.10259](#).
- [48] Wim van Dam and Patrick Hayden, *Rényi-entropic bounds on quantum communication*, [arXiv:quant-ph/0204093](#) (2002).
- [49] Masanori Ohya and Dénes Petz, *Quantum Entropy and Its Use*, 2nd ed. (Springer Berlin, Heidelberg, 1993) Chap. 2, pp. 37–46.
- [50] Masanori Ohya and Noboru Watanabe, *Quantum Entropy and Its Applications to Quantum Communication and Statistical Physics*, *Entropy* **12**, 1194 (2010).
- [51] Armen E. Allahverdyan, Karen V. Hovhannisyan, Dominik Janzing, and Guenter Mahler, *Thermodynamic limits of dynamic cooling*, *Phys. Rev. E* **84**, 041109 (2011), [arXiv:1107.1044](#).
- [52] James Ladyman, James Lambert, and Karoline Wiesner, *What is a complex system?* *Eur. J. Philos.* **3**, 33 (2013).
- [53] Yuriy Holovatch, Ralph Kenna, and Stefan Thurner, *Complex systems: physics beyond physics*, *Eur. J. Phys.* **38**, 023002 (2017), [arXiv:1610.01002](#).
- [54] Albert W. Marshall, Ingram Olkin, and Barry C. Arnold, *Inequalities: Theory of Majorization and its Applications*, 2nd ed., Vol. 143 (Springer New York, NY, 2011) Chap. 9, pp. 297–365.
- [55] Walther Thirring, *Quantum Mathematical Physics: Atoms, Molecules and Large Systems* (Springer Berlin, Heidelberg, 2002) Chap. 7, pp. 423–499.
- [56] George Gamow, *One Two Three... Infinity: Facts and Speculations of Science* (Viking Press, New York, NY, 1947).
- [57] Gerardo Adesso, Sammy Ragy, and Antony R. Lee, *Continuous Variable Quantum Information: Gaussian States and Beyond*, *Open Syst. Inf. Dyn.* **21**, 1440001 (2014), [arXiv:1401.4679](#).
- [58] Ralph Silva, Gonzalo Manzano, Paul Skrzypczyk, and Nicolas Brunner, *Performance of autonomous quantum thermal machines: Hilbert space dimension as a thermodynamical resource*, *Phys. Rev. E* **94**, 032120 (2016), [arXiv:1604.04098](#).

Supplemental Material

CONTENTS

A. Equality Forms of the (Carnot-)Landauer Limit	17
A1. Coherent-Control Paradigm: The Landauer Limit	17
A2. Incoherent-Control Paradigm: The Carnot-Landauer Limit	18
B. Diverging Energy	19
B1. Sufficiency: Diverging Energy Cooling Protocol	19
B2. Necessity of Diverging Energy for Protocols with Finite Time and Control Complexity	19
C. Diverging Time Cooling Protocol for Finite-Dimensional Systems	20
C1. Proof of Theorem 2	20
C2. Special Case: Equally Spaced Hamiltonians	22
D. Conditions for Structural and Control Complexity	25
D1. Necessary Complexity Conditions	25
D1a. Necessary Structural Conditions	25
D1b. Necessary Control Complexity Conditions	26
D2. Proof of Theorem 3, Corollary 2, and Theorem 11	27
D3. Sufficient Complexity Conditions	28
D4. Fine-Tuned Control Conditions	29
D5. Energy-Gap Variety as a Notion of Control Complexity	36
E. Diverging Time and Diverging Control Complexity Cooling Protocols for Harmonic Oscillators	37
E1. Preliminaries	38
E2. Diverging-Time Cooling Protocol for Harmonic Oscillators	39
E2a. Diverging-Time Protocol using Gaussian Operations (with Diverging Control Complexity)	39
E2b. Diverging-Time Protocol using Non-Gaussian Operations (with Finite Control Complexity)	41
E3. Diverging Control Complexity Cooling Protocol for Harmonic Oscillators	43
F. Cooling Protocols in the Incoherent-Control Paradigm	45
F1. Diverging Energy: Proof of Theorem 6	45
F2. Diverging Time and Diverging Control Complexity	48

F3. Saturating the Landauer Limit with an Infinite-Temperature Heat Bath	49
F4. Analysis of Finite-Temperature Heat Baths	50
G. Perfect Cooling at the Carnot-Landauer Limit in the Incoherent-Control Paradigm	50
G1. Qubit Case	51
G2. Qudit Case	55
H. Comparison of Cooling Paradigms and Resources for Imperfect Cooling	59
H1. Rates of Resource Divergence for Linear Qubit Machine Sequence	59
H2. Comparison of Coherent and Incoherent Control	60

Appendix A: Equality Forms of the (Carnot-)Landauer Limit

In this section, we present lower bounds on the energy change of the machine (or heat dissipated into its environment) in terms of the entropy change of the target system, both in the coherent and incoherent-control settings outlined in the main text. In the coherent setting, this amounts to the well-known Landauer principle [11], whereas the incoherent setting requires an extension of this derivation. These lower bounds are important, because they put limits on the optimal energetic performance of the machines for cooling. Note, finally, that the initial state of the machine is diagonal in its energy eigenbasis and must remain so for any process saturating the (Carnot-)Landauer limit; moreover, the target begins similarly and ends up in the pure state $|0\rangle\langle 0|$ when perfect cooling is achieved. As a result, all quantities relevant to perfect cooling at the (Carnot-)Landauer limit can be computed in terms of their “classical” counterparts, i.e., $\varrho_{\mathcal{X}} \rightarrow p_{\mathcal{X}} := (p_0, \dots, p_d)$ with $p_n = e^{-\beta E_n}$, $\text{tr}[H\varrho_{\mathcal{X}}] \rightarrow \langle E \rangle_{p_{\mathcal{X}}} := \sum_n p_n E_n$, $S(\varrho_{\mathcal{X}}) \rightarrow S(p_{\mathcal{X}}) := -\sum_n p_n \log(p_n)$, $\mathcal{Z}(\beta, H_{\mathcal{X}}) = \sum_n e^{-\beta E_n}$, and so on. Nonetheless, all of the results presented hold for the more general “quantum” properties.

A1. Coherent-Control Paradigm: The Landauer Limit

The coherent setting was already studied in detail in Ref. [29], where the authors derived an equality version of Landauer’s principle. We restate the results here for convenience, since we will also use them in the incoherent paradigm. Recall that the setting we consider consists of two parts, the target system \mathcal{S} and the machine \mathcal{M} . In the beginning, the joint state is $\varrho_{\mathcal{SM}} = \varrho_{\mathcal{S}} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ for some arbitrary (but fixed) Hamiltonian $H_{\mathcal{M}}$ and $\beta \in \mathbb{R}$. Note that any full-rank state ϱ can be associated to some chosen temperature β , which sets the energy scale, and a Hamiltonian $H = -\frac{1}{\beta} \log(\varrho)$; as we consider arbitrary Hamiltonians, we only write the state dependence on these parameters when necessary. If the state is not full rank, the rank can be used to redefine the dimension. We assume that both systems are finite dimensional. Let U be a global unitary on \mathcal{SM} . We write $\varrho'_{\mathcal{SM}} := U[\varrho_{\mathcal{S}} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})]U^\dagger$ and denote by ϱ'_s and ϱ'_m the respective reduced states. The quantity $I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} = S(\varrho'_s) + S(\varrho'_m) - S(\varrho'_{\mathcal{SM}})$ is the final mutual information between \mathcal{S} and \mathcal{M} and $D(\varrho'_m || \varrho_m) = \text{tr}[\varrho'_m \log(\varrho'_m)] - \text{tr}[\varrho'_m \log(\varrho_m)]$ is the relative entropy of the final machine state with respect to its initial state.

Lemma 1 ([29, Lemma 2]). *Let the setting be as above. Then*

$$[S(\varrho'_s) - S(\varrho_s)] + [S(\varrho'_m) - S(\varrho_m)] = I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} \geq 0. \quad (\text{A1})$$

Proof. We note that

$$[S(\varrho'_s) - S(\varrho_s)] + [S(\varrho'_m) - S(\varrho_m)] = S(\varrho'_s) + S(\varrho'_m) - S(\varrho'_{\mathcal{SM}}), \quad (\text{A2})$$

since the von Neumann entropy is additive for product states and invariant under unitary evolution. The assertion follows from the definition of the mutual information and the fact that it is non-negative. \square

Theorem 9 (Equality form of Landauer’s principle, [29, Theorem 3]). *Let the setting be as above. Then*

$$\beta \text{tr}[H_{\mathcal{M}}(\varrho'_m - \varrho_m)] - [S(\varrho_s) - S(\varrho'_s)] = I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} + D(\varrho'_m || \varrho_m) \geq 0. \quad (\text{A3})$$

Proof. From Lemma 1, it follows that

$$[S(\varrho_s) - S(\varrho'_s)] + I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} = S(\varrho'_m) - S(\varrho_m). \quad (\text{A4})$$

Using the fact that $\varrho_m = \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$, we infer that $D(\varrho'_m || \varrho_m) = -S(\varrho'_m) + \beta \text{tr}[H_{\mathcal{M}}\varrho'_m] + \log[\text{tr}(e^{-\beta H_{\mathcal{M}}})]$ and $S(\varrho_m) = \beta \text{tr}[H_{\mathcal{M}}\varrho_m] + \log[\text{tr}(e^{-\beta H_{\mathcal{M}}})]$. Re-expressing the first of these for $S(\varrho'_m)$ and inserting both into Eq. (A4) yields the claimed equality. The inequality results from non-negativity of relative entropy and mutual information. This completes the proof. \square

A2. Incoherent-Control Paradigm: The Carnot-Landauer Limit

Landauer's principle provides a relationship between how much heat must necessarily be dissipated into the thermal background environment upon manipulating the entropy of a given quantum system. Until now, we have assumed that the system of interest can interact arbitrarily with its environment (i.e., the machine); in other words, we have considered general joint unitary interactions between system and machine, without restriction. In doing so, we have tacitly assumed the ability to draw energy from some external resource (i.e., a work source) in order to implement said unitaries, which are in general not energy preserving. The particularities of such a resource are left as an abstraction. However, from a thermodynamicists' perspective, this setting may seem somewhat unsatisfactory, as the joint target-machine system is not energetically closed. In order to provide a more self-contained picture of the cooling procedure, one can explicitly include the energy resource, modelled as a quantum system itself, into the setting.

To this end, note first that said resource must be out of thermal equilibrium with respect to the target and machine in order to perform any meaningful thermodynamic transformation. Furthermore, it is sensible to assume that the energy resource system is in thermal equilibrium with its own environment to begin with. The joint target-machine-resource system is then considered to be energetically closed; as such, global unitaries in this setting are restricted to be energy conserving. In order to act as a resource for cooling the target in this picture, the energy source here must begin in equilibrium with a heat bath that is hotter than the initial temperature of the machine (assuming that the machine and resource both begin in thermal states), such that a natural heat flow is induced that leads the environment of the machine to act as a final heat sink. This setting is what we call the incoherent-control scenario. In this context, Landauer's principle translates to studying the relationship between the heat that is necessarily dissipated into the machine's environment upon manipulating the entropy of the target system. Finally, note that the relationship between the coherent and the incoherent-control paradigms is interesting in itself: while on the one hand the incoherent setting includes an additional system and therefore increases the dimensionality of the overall joint system, on the other hand by restricting the transformations on this larger space to be energy conserving, one limits the orbit of attainable states.

Now let us consider the incoherent-control setting. Here, we have the target system \mathcal{S} and the machine comprises of one part \mathcal{C} coupled to the cold bath and another \mathcal{H} coupled to the hot bath. We assume that all systems are finite-dimensional. Every subsystem \mathcal{A} is associated to a Hamiltonian $H_{\mathcal{A}}$ and \mathcal{C} , \mathcal{H} are initially in a thermal state; the cold bath has inverse temperature β and the hot bath has inverse temperature $\beta_H < \beta$. We assume β, β_H . Thus, the initial joint state is $\varrho_{\mathcal{SCH}} = \varrho_{\mathcal{S}} \otimes \tau_{\mathcal{C}}(\beta, H_{\mathcal{C}}) \otimes \tau_{\mathcal{H}}(\beta_H, H_{\mathcal{H}})$. The global evolution on \mathcal{SCH} is implemented via a unitary U , leading to $\varrho'_{\mathcal{SCH}} = U(\varrho_{\mathcal{SCH}})U^\dagger$. We further assume that the unitary evolution on the joint system is energy conserving, i.e., $[U, H_{\mathcal{S}} + H_{\mathcal{C}} + H_{\mathcal{H}}] = 0$. We write $\Delta S_{\mathcal{A}} := S(\varrho'_{\mathcal{A}}) - S(\varrho_{\mathcal{A}})$ for the entropy change on subsystem \mathcal{A} and $\Delta E_{\mathcal{A}} := \text{tr}[H_{\mathcal{A}}(\varrho'_{\mathcal{A}} - \varrho_{\mathcal{A}})]$ for the average energy change. Moreover, the free energy of a state $\varrho_{\mathcal{A}}$ with respect to the inverse temperature β is $F_{\beta}(\varrho_{\mathcal{A}}) = \text{tr}[H_{\mathcal{A}}\varrho_{\mathcal{A}}] - \beta^{-1}S(\varrho_{\mathcal{A}})$.

In the incoherent setting, it makes sense to look at the energy decrease in the hot bath \mathcal{H} , since the hot bath can be seen as the energetic resource one must to expend in order to cool the system \mathcal{S} (alternatively, as we present after the following theorem, one can consider the energy dissipated into the cold bath \mathcal{C} , which serves as the heat sink).

Theorem 10. *In the above setting, it holds that*

$$\Delta F_{\mathcal{S}}^{(\beta)} + \eta \Delta E_{\mathcal{H}} = -\frac{1}{\beta} [\Delta S_{\mathcal{S}} + \Delta S_{\mathcal{C}} + \Delta S_{\mathcal{H}} + D(\varrho'_{\mathcal{C}} || \varrho_{\mathcal{C}}) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}})] \leq 0, \quad (\text{A5})$$

where $(0, 1) \ni \eta := 1 - \beta_H/\beta$ is the Carnot efficiency and $\Delta F_{\mathcal{S}}^{(\beta)} = F_{\beta}(\varrho'_{\mathcal{S}}) - F_{\beta}(\varrho_{\mathcal{S}})$.

Proof. Let us consider

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} := S(\varrho'_{\mathcal{S}}) + S(\varrho'_{\mathcal{C}}) + S(\varrho'_{\mathcal{H}}) - S(\varrho'_{\mathcal{SCH}}) \geq 0. \quad (\text{A6})$$

Note that the quantity $I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}}$, which quantifies the tripartite mutual information of the state $\varrho'_{\mathcal{SCH}}$, is non-negative via subadditivity $S(\varrho_{\mathcal{A}}) + S(\varrho_{\mathcal{B}}) \geq S(\varrho_{\mathcal{AB}})$ for any state $\varrho_{\mathcal{AB}}$. Furthermore, since the von Neumann entropy is invariant under unitary transformations and additive for tensor product states, we have

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} = \Delta S_{\mathcal{S}} + \Delta S_{\mathcal{C}} + \Delta S_{\mathcal{H}}. \quad (\text{A7})$$

We also have that

$$\Delta S_{\mathcal{C}} = \beta \Delta E_{\mathcal{C}} - D(\varrho'_{\mathcal{C}} || \varrho_{\mathcal{C}}) \quad (\text{A8})$$

and

$$\Delta S_{\mathcal{H}} = \beta_H \Delta E_{\mathcal{H}} - D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A9})$$

Thus,

$$I(S : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} = \Delta S_s + \beta \Delta E_c - D(\varrho'_c || \varrho_c) + \beta_H \Delta E_{\mathcal{H}} - D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A10})$$

Since the unitary is energy conserving, we infer that $\Delta E_s + \Delta E_c + \Delta E_{\mathcal{H}} = 0$. Hence, we have

$$\Delta S_s - \beta \Delta E_s + (\beta_H - \beta) \Delta E_{\mathcal{H}} = I(S : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} + D(\varrho'_c || \varrho_c) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A11})$$

Using the free energy, we can rewrite this as

$$-\beta[F_{\beta}(\varrho'_s) - F_{\beta}(\varrho_s)] - (\beta - \beta_H) \Delta E_{\mathcal{H}} = I(S : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} + D(\varrho'_c || \varrho_c) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A12})$$

Dividing by $-\beta$, we obtain the assertion, since, in particular, $I(S : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} + D(\varrho'_c || \varrho_c) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}) \geq 0$ by the non-negativity of each term. \square

In particular, we have shown that the energy extracted from the hot bath is lower-bounded by the increase in free energy, weighted by the inverse Carnot efficiency:

$$\text{tr}[H_{\mathcal{H}}(\varrho_{\mathcal{H}} - \varrho'_{\mathcal{H}})] \geq \frac{1}{\eta}[F_{\beta}(\varrho'_s) - F_{\beta}(\varrho_s)]. \quad (\text{A13})$$

Note that if $\varrho_s = \tau_s(\beta, H_s)$, the r.h.s. is non-negative for any nontrivial thermodynamic process, i.e., any for which the target system is heated or—of particular relevance for us—cooled. This follows by the Gibbs variational principle, which states that the free energy of ϱ is minimal iff ϱ is the corresponding Gibbs state.

Finally, in order to make a more concrete connection to the spirit of Landauer's original derivation, note that one can consider bounding the heat dissipated into the cold bath, rather than that drawn from the hot bath. Substituting $\Delta E_{\mathcal{H}} = -(\Delta E_s + \Delta E_c)$ into Eq. (A10) leads to

$$-\tilde{\Delta} S_s - \beta_H \Delta E_s + (\beta - \beta_H) \Delta E_c \geq 0, \quad (\text{A14})$$

which recovers the standard Landauer bound for the dissipated heat in the limit of an infinitely hot heat bath, i.e., $\beta_H \rightarrow 0$.

Appendix B: Diverging Energy

B1. Sufficiency: Diverging Energy Cooling Protocol

This cooling protocol is arguably the simplest of those presented. The thermal populations of any target system can be exchanged with a machine system of the same dimension, in the thermal state of $H_{\mathcal{M}} = \omega_{\mathcal{M}} \sum_{n=0}^{d-1} n |n\rangle\langle n|$. As $\omega_{\mathcal{M}} \rightarrow \infty$, the machine state $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ approaches $|0\rangle\langle 0|_{\mathcal{M}}$ independently of β (as long as $\beta \neq 0$). Such a population-exchange operation is a single interaction (i.e., the protocol occurs in unit time), which is of finite complexity (in a sense that we discuss below). However, the energy drawn from the resource \mathcal{W} upon performing said SWAP operation is at least $E = (p_s^{(1)} - p_{\mathcal{M}}^{(1)})(\omega_{\mathcal{M}} - \omega_s^{(1)})$, where $p_s^{(1)}$ is the initial population of the first excited level of system \mathcal{X} and $\omega_s^{(1)}$ is the first energy eigenvalue of the target system. Denoting by $\omega_s^{(k)}$ the energy eigenvalue of the k^{th} excited level of the target system, we have above assumed that $\omega_s^{(0)} = 0$ (which we do for all Hamiltonians without loss of generality) and $\omega_{\mathcal{M}} > \omega_s^{(d-1)}$. As such, perfect cooling will incur diverging energy cost.

B2. Necessity of Diverging Energy for Protocols with Finite Time and Control Complexity

Consider the following Hamiltonians for the target system and machine with finite but otherwise arbitrary energy levels, $H_s = \sum_{n=0}^{d_s-1} \omega_s^{(n)} |n\rangle\langle n|_s$ and $H_{\mathcal{M}} = \sum_{n=0}^{d_{\mathcal{M}}-1} \omega_{\mathcal{M}}^{(n)} |n\rangle\langle n|_{\mathcal{M}}$, respectively. For any finite inverse temperature β , the initial thermal states $\tau_s(\beta, H_s)$ and $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ are of full rank. Suppose now that one can implement a single unitary transformation (i.e., a unit time protocol) of finite control complexity on the joint target and machine, yielding the joint output state $\varrho'_{s\mathcal{M}} = \text{tr}_{\mathcal{M}}[U(\tau_s(\beta, H_s) \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}))U^{\dagger}]$, and wishes to attain perfect cooling of the target in doing so. By invariance of the rank under unitary transformations and the fact that the system and machine begin uncorrelated, we have

$$\text{rank}[\tau_s(\beta, H_s)] \text{rank}[\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})] = \text{rank}[\varrho'_{s\mathcal{M}}] \leq \text{rank}[\varrho'_s] \text{rank}[\varrho'_{\mathcal{M}}], \quad (\text{B1})$$

where the inequality follows from the subadditivity of the Rényi-zero entropy [48], which is the logarithm of the rank. To achieve perfect cooling of the target, one must (at least asymptotically) attain $\text{rank}[\varrho'_S] < \text{rank}[\tau_S(\beta, H_S)]$, which implies that $\text{rank}[\varrho'_M] > \text{rank}[\tau_M(\beta, H_M)]$. However, if this condition is achieved, then $D[\varrho'_M \| \tau_M(\beta, H_M)]$ diverges, implying a diverging energy cost by Eq. (3). The above argument already appears in Ref. [29].

The other situation that one must consider is the case where one attains a ϱ'_S such that $\text{rank}[\varrho'_S] = \text{rank}[\tau_S(\beta, H_S)]$ but nonetheless ϱ'_S is arbitrarily close to a pure state, as is the case, for instance, in the protocols that we present. Consider a sequence of machines $\varrho_M^{(i)}$ and unitaries $U^{(i)}$ such that $\varrho_M^{(i)} \rightarrow \varrho_M$ and $U^{(i)} \rightarrow U$. Note that since we fixed the dimensions of S and M , any sequence of machines has a converging subsequence by the Bolzano-Weierstrass theorem and the fact that the set of quantum states is compact. Here, ϱ_M and U achieve perfect cooling. If we fix ϱ_S , we obtain a corresponding sequence $(\varrho'_M)^{(i)}$ such that $(\varrho'_M)^{(i)} \rightarrow \varrho'_M$. Crucially, here, since we restrict the unitary transformation to be of finite control complexity, the states ϱ_M and ϱ'_M are effectively finite dimensional, in the sense that whatever their true dimension, they can be replaced by finite-dimensional versions without changing any of the relevant quantities (see Appendix D). Since the relative entropy $(\varrho, \sigma) \mapsto D(\varrho \| \sigma)$ is lower semicontinuous [49, 50] and since $D(\varrho'_M \| \varrho_M) \rightarrow \infty$ by the arguments above, we infer that $D[(\varrho'_M)^{(i)} \| \varrho_M^{(i)}] \rightarrow \infty$ as $i \rightarrow \infty$. This argument holds independently of $\text{rank}[\varrho'_S]$; in particular, for the special case $\text{rank}[\varrho'_S] = \text{rank}[\tau_S(\beta, H_S)]$ that we are considering here. Thus, to approach perfect cooling in finite time and with finite control complexity, one would need a diverging energy cost. Thus, we see that within the resource trinity of energy, time, and control complexity, if the latter two are finite, then energy must diverge to asymptotically achieve a pure state. Whether or not there exist other (unaccounted for) resources that allow one to achieve this with all three of the aforementioned resources being finite remains an open question.

Importantly, the above argument no longer holds if the time or control complexity is allowed to diverge. In such cases, both ϱ_M and ϱ'_M can be infinite dimensional, and because of this the rank argument no longer applies and the relative entropy does not necessarily diverge in the limit of perfect cooling. In contrast, as we show, it is even possible to saturate the Landauer bound.

Appendix C: Diverging Time Cooling Protocol for Finite-Dimensional Systems

C1. Proof of Theorem 2

Proof. Consider a target system S of dimension d with associated Hamiltonian

$$H_S = \sum_{k=0}^{d-1} \omega_k |k\rangle\langle k|_S, \quad (\text{C1})$$

where we also set $\omega_0 = 0$ without loss of generality. Consider also the machine M to be composed of N subsystems, $\{\mathcal{M}_n\}_{n=1,\dots,N}$, each of the same dimension d as the target, whose local Hamiltonians are

$$H_M^{(n)} = (1 + n\epsilon)H_S, \quad (\text{C2})$$

where $\epsilon = (\beta_{\max} - \beta)/(N\beta)$. We first cool the system initially at nonzero β to some fixed, finite β_{\max} , which we eventually take $\beta_{\max} \rightarrow \infty$ in order to asymptotically achieve perfect cooling. We treat the case $\beta = 0$ as a limiting case of $\beta \rightarrow 0$: here, as $\beta \rightarrow 0$, we let $N \rightarrow \infty$ such that $N\beta \rightarrow \infty$, e.g., we specify a suitable function $N(\beta)$ such that $N(\beta) \rightarrow \infty$ “faster” than $\beta \rightarrow 0$.

We now show that, given the ability to perform a diverging number of operations on such a configuration, one can reach the target state $\tau_S(\beta_{\max}, H_S)$. In particular, we show that the protocol presented uses the minimal amount of energy to do so, and explicitly calculate this to be $\beta^{-1}\tilde{\Delta}S$ units of energy, where $\tilde{\Delta}S := S[\tau_S(\beta, H_S)] - S[\tau_S(\beta_{\max}, H_S)]$. In other words, as the number of operations in the protocol diverges, we approach perfect cooling at the Landauer limit, thereby saturating the ultimate bound.

The diverging time cooling protocol is as follows. At each step, the target system interacts with a single machine labelled by n via the SWAP operator

$$\mathbb{S}_{S\mathcal{M}_n}^d := \sum_{i,j=0}^{d-1} |i, j\rangle\langle j, i|_{S\mathcal{M}_n}. \quad (\text{C3})$$

As the target and machine subsystems considered here are of the same dimension, we drop the subscript on the states associated to each subsystem, for ease of notation. Such a transformation is, in general, not energy conserving, but one can calculate the energy change for both the target system and the machine due to the n^{th} interaction as

$$\Delta E_S^{(n)} = \text{tr} \left[H_S \tau(\beta, H_M^{(n)}) \right] - \text{tr} \left[H_S \tau(\beta, H_M^{(n-1)}) \right], \quad (\text{C4})$$

and so the total energy change of the system over the entire N -step protocol is given by

$$\Delta E_S = \sum_{n=1}^N \Delta E_S^{(n)} = \text{tr} \left[H_S \tau(\beta, H_{\mathcal{M}}^{(N)}) \right] - \text{tr} \left[H_S \tau(\beta, H_{\mathcal{M}}^{(0)}) \right]. \quad (\text{C5})$$

The energy change of the machine subsystem that is swapped with the target system at each step is given by

$$\Delta E_{\mathcal{M}}^{(n)} = \text{tr} \left[H_{\mathcal{M}}^{(n)} \tau(\beta, H_{\mathcal{M}}^{(n-1)}) \right] - \text{tr} \left[H_{\mathcal{M}}^{(n)} \tau(\beta, H_{\mathcal{M}}^{(n)}) \right] = \sum_{k=0}^{d-1} (1 + n\epsilon) \omega_k \left[p_k(\beta, H_{\mathcal{M}}^{(n-1)}) - p_k(\beta, H_{\mathcal{M}}^{(n)}) \right], \quad (\text{C6})$$

where $p_k(\beta, H_{\mathcal{M}}^{(n)}) = e^{-\beta(1+n\epsilon)\omega_k} / \mathcal{Z}_{\mathcal{M}_n}(\beta, H_{\mathcal{M}}^{(n)})$ is the population in the k^{th} energy level of the thermal state of the n^{th} machine subsystem at inverse temperature β , with $\mathcal{Z}_{\mathcal{M}_n}(\beta, H_{\mathcal{M}}^{(n)}) = \text{tr} \left[e^{-\beta H_{\mathcal{M}}^{(n)}} \right]$ being the partition function.

By summing the contributions of the energy changes in each step, one can obtain the total energy change for the overall machine throughout the entire process:

$$\Delta E_{\mathcal{M}}^{(N)} = \sum_{n=1}^N \Delta E_{\mathcal{M}}^{(n)} = \sum_{n=1}^N \sum_{k=0}^{d-1} (1 + n\epsilon) \omega_k \left[p_k(\beta, H_{\mathcal{M}}^{(n-1)}) - p_k(\beta, H_{\mathcal{M}}^{(n)}) \right], \quad (\text{C7})$$

In general, it is complicated to calculate the energy cost for the protocol up until a finite time step N , since this depends on the full energy structure of the target system and machine subsystems involved (we return to resolve this problem for the special case of equally spaced system and machine Hamiltonians in the coming section). Here, we focus on a special case in which $N \rightarrow \infty$, i.e., there is a diverging number of machine subsystems that the target system interacts with throughout the protocol. This limit physically corresponds to that of requiring a diverging amount of time (in terms of the number of steps). Furthermore, we take the limit $\epsilon \rightarrow 0$ for any fixed β, β_{\max} . Considering the differentials

$$\Delta p_k^{(n)} := p_k(\beta, H_{\mathcal{M}}^{(n)}) - p_k(\beta, H_{\mathcal{M}}^{(n-1)}), \quad (\text{C8})$$

and

$$\Delta x_n := x_n - x_{n-1} \quad \text{with} \quad x_n := 1 + n\epsilon. \quad (\text{C9})$$

In order for x_n to become infinitesimal, and noting the explicit form of the machine subsystem Hamiltonians $H_{\mathcal{M}}^{(n)} = (1+n\epsilon)H_S$, we can make the replacement

$$-\frac{\Delta p_k^{(n)}}{\Delta x_n} \Delta x_n \rightarrow -\frac{\partial p_k(\beta, xH_S)}{\partial x} dx \quad (\text{C10})$$

where $x := 1 + n\epsilon$ has become a continuous parameter. This way we can express the limit $N \rightarrow \infty$ of Eq. (C7) as a Riemann integral in the following form:

$$\lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} = - \int_1^{x_{\max}} \sum_{k=0}^{d-1} x \omega_k \frac{\partial p_k(\beta, xH_S)}{\partial x} dx, \quad (\text{C11})$$

where $x_{\max} := \beta_{\max}/\beta$. Both the summation and the integral converge, so one can swap the order of their evaluation. Integrating by parts then gives

$$\begin{aligned} \lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} &= \sum_{k=0}^{d-1} \left[-x \omega_k p_k(\beta, xH_S) \Big|_1^{x_{\max}} + \int_1^{x_{\max}} \omega_k p_k(\beta, xH_S) dx \right] \\ &= \sum_{k=0}^{d-1} \left[-x \omega_k p_k(\beta, xH_S) \Big|_1^{x_{\max}} \right] - \int_1^{x_{\max}} \frac{1}{\beta} \frac{\partial}{\partial x} [\log \mathcal{Z}(\beta, xH_S)] dx \\ &= E[\tau(\beta, H_S)] - E[\tau(\beta, x_{\max} H_S)] - \frac{1}{\beta} \log \mathcal{Z}(\beta, x_{\max} H_S) + \frac{1}{\beta} \log \mathcal{Z}(\beta, H_S), \end{aligned} \quad (\text{C12})$$

where in the second line we again swap the order of the integral and the sum to write $\sum_{k=0}^{d-1} \omega_k p_k(\beta, xH_S) = -\frac{1}{\beta} \frac{\partial}{\partial x} [\log \mathcal{Z}(\beta, xH_S)]$ and in the last line we invoke $E[\tau(\beta, xH)] = \text{tr} [xH \tau(\beta, xH)]$. Finally, writing the partition function in terms of the average energy and entropy, i.e., $\log[\mathcal{Z}(\beta, xH)] = -\beta E[\tau(\beta, xH)] + S[\tau(\beta, xH)]$, the total energy change of the machine is given by

$$\lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} = E[\tau(\beta, H_S)] - E[\tau(\beta, x_{\max} H_S)] + E[\tau(\beta, x_{\max} H_S)] - \frac{1}{\beta} S[\tau(\beta, x_{\max} H)] - E[\tau(\beta, H_S)] + \frac{1}{\beta} S[\tau(\beta, H_S)]$$

$$= \frac{1}{\beta} \{S[\tau(\beta, H_S)] - S[\tau(\beta_{\max}, H_S)]\} = \frac{1}{\beta} \tilde{\Delta} S_S, \quad (\text{C13})$$

where we make use of the property $\tau_S(\beta, x_{\max} H_S) = \tau_S(\beta_{\max}, H_S)$ and the entropy decrease of the target system corresponds to that associated with the transformation $\tau(\beta, H_S) \rightarrow \tau(\beta_{\max}, H_S)$. Thus, as the number of timesteps diverges, this cooling process saturates the Landauer limit for the heat dissipated by the machine. In order to achieve perfect cooling at the Landauer limit, i.e., the final target state to approach $|0\rangle\langle 0|$ and thus prove Theorem 2, we can now take the limit $\beta_{\max} \rightarrow \infty$. \square

The above proof holds for systems and machines of arbitrary (but equal) dimension, either finite or infinite, with arbitrary Hamiltonians. We now present some more detailed analysis regarding the special case where the Hamiltonians of the target system and all machine subsystems are equally spaced; this provides an opportunity both to derive a more detailed formula for the energy costs involved and to build intuition regarding some of the important differences between the finite- and infinite-dimensional settings.

C2. Special Case: Equally Spaced Hamiltonians

Consider a finite d -dimensional target system beginning at inverse temperature β with an equally spaced Hamiltonian $H_S(\omega_S) = \omega_S \sum_{n=0}^{d-1} n|n\rangle\langle n|_S$. In this case, we can derive a more precise dimension-dependant function for the energy cost dissipated by the machines throughout the optimal cooling protocol presented above.

Consider an initial target system $\tau_S(\beta, H_S)$ and a diverging number N of machines $\{\mathcal{M}_\alpha\}_{\alpha=0,\dots,N}$ of the same dimension d as the target, which all begin in a thermal state at inverse temperature β with respect to an equally spaced Hamiltonian whose gaps between neighbouring energy levels $\omega_{\mathcal{M}_\alpha}$ are ordered non-decreasingly. Each machine is used once and then discarded; the particular interaction is the aforementioned SWAP between the target system and the n^{th} qudit machine, i.e., that represented by the unitary $\mathbb{S}_{S\mathcal{M}_\alpha}^d := \sum_{i,j=0}^{d-1} |i, j\rangle\langle j, i|_{S\mathcal{M}_\alpha}$. After applying such an operation, the state of the target system is given by

$$\tau_S(\beta, \omega_\alpha) := \frac{e^{-\beta H_S(\omega_\alpha)}}{\mathcal{Z}_S(\beta, \omega_\alpha)}, \quad (\text{C14})$$

where $H_S(\omega_\alpha) := \omega_\alpha \sum_{n=0}^{d-1} n|n\rangle\langle n|_S$ and $\mathcal{Z}_S(\beta, \omega_\alpha) := \text{tr}[e^{-\beta H_S(\omega_\alpha)}]$.

We now calculate the energy cost explicitly for the diverging time cooling protocol, which saturates the Landauer bound in the asymptotic limit. In order to minimise the energy cost of cooling, the target system must be cooled by the qudit system in the machines with the smallest gap between neighbouring energy levels (that permits cooling) as much as possible at each stage. In order to optimally use the given machine structure at hand, we thus order the set of energy gaps ω_α in non-decreasing order. In addition, the protocol to reach the Landauer erasure bound, i.e., minimal energy cost, dictates that one must infinitesimally increase ω_α of the machines in order to dissipate as little heat as possible throughout the interactions. Since we are here considering a diverging time limit, we have access to a diverging number of qudit machine with distinct energy gap ω_α at our disposal; the task is then to use these in an energy-optimal manner.

It is straightforward to see that to minimise the total energy cost, one must apply the sequence of unitaries $\mathbb{S}_{S\mathcal{M}_\alpha}^d$ such that $\mathbb{S}_{S\mathcal{M}_0}^d$ is first applied to reach the optimally cool $\tau_S(\beta, \omega_0)$, then $\mathbb{S}_{S\mathcal{M}_1}^d$ to reach $\tau_S(\beta, \omega_1)$, and so on. The heat dissipated by the reset machines in each stage of such a cooling protocol (i.e., for each value of α) can thus be calculated as

$$\begin{aligned} \Delta E_{\mathcal{M}_\alpha}(\omega_\alpha) &= -\{\text{tr}[H_{\mathcal{M}_\alpha}(\omega_\alpha)\tau_{\mathcal{M}_\alpha}(\beta, \omega_\alpha)] + \text{tr}[H_{\mathcal{M}_\alpha}(\omega_\alpha)\tau_{\mathcal{M}_\alpha}(\beta, \omega_{\alpha-1})]\} \\ &= -\text{tr}[H_S(\omega_\alpha)[\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})]]. \end{aligned} \quad (\text{C15})$$

In the second line, we have made use of the fact that the Hamiltonians of both the target system and each of machine are d -dimensional and equally spaced. So far, we have obtained the energy dissipated by the reset machines. To investigate the total energy cost of cooling in such a process, we also must consider the contribution of energy transferred to the target system S , which is characterised via its local Hamiltonian H_S and calculated via

$$\Delta E_S(\omega_\alpha) = \text{tr}[H_S(\omega_S)\tau_S(\beta, \omega_\alpha)] - \text{tr}[H_S(\omega_S)\tau_S(\beta, \omega_{\alpha-1})], \quad (\text{C16})$$

in which we set $\omega_0 = \omega_S$. Using Eqs. (C15, C16), the total energy cost for each stage of cooling is given by

$$\Delta E_{S\mathcal{M}}(\omega_\alpha) = \Delta E_S(\omega_\alpha) + \Delta E_{\mathcal{M}}(\omega_\alpha) = \text{tr}\{[H_S(\omega_S) - H_S(\omega_\alpha)][\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})]\}, \quad (\text{C17})$$

which leads to the overall energy cost after N stages, where N is the number of non-zero distinct energy gaps of the reset machines, as

$$\Delta E_{S\mathcal{M}}^{(N)} = \sum_{\alpha=1}^N \Delta E_{S\mathcal{M}}(\omega_\alpha) = \sum_{\alpha=1}^N \text{tr}\{[H_S(\omega_S) - H_S(\omega_\alpha)][\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})]\}. \quad (\text{C18})$$

Now, we can obtain the total energy cost for each stage of the protocol (i.e., each value of α considered) in terms of the transformation of the target system alone. Note that in this protocol, each stage corresponding to each of the N distinct energy gaps $\{\omega_\alpha\}$ in itself requires only one operation to perfectly reach $\tau_s(\beta, \omega_\alpha)$. The end result of this protocol is that the target system is cooled from the initial thermal state $\tau_s(\beta, \omega_s)$, where ω_s is the energy gap between each pair of adjacent energy levels in the system, to $\tau_s(\beta, \omega_{\max})$ in the energy-optimal manner.

Starting from Eq. (C18), we have

$$\begin{aligned}\Delta E_{s\mathcal{M}}^{(N)} &= \sum_{\alpha=1}^N \text{tr} \{ [H_s(\omega_s) - H_s(\omega_\alpha)] [\tau_s(\beta, \omega_\alpha) - \tau_s(\beta, \omega_{\alpha-1})] \} \\ &= \sum_{\alpha=1}^N (\omega_s - \omega_\alpha) \left[\left(\frac{e^{-\beta\omega_\alpha}}{1 - e^{-\beta\omega_\alpha}} - \frac{e^{-\beta\omega_{\alpha-1}}}{1 - e^{-\beta\omega_{\alpha-1}}} \right) - \left(\frac{d e^{-\beta d\omega_\alpha}}{1 - e^{-\beta d\omega_\alpha}} - \frac{d e^{-\beta d\omega_{\alpha-1}}}{1 - e^{-\beta d\omega_{\alpha-1}}} \right) \right] \\ &= \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N (\omega_s - \omega_\alpha) \sum_{k=0}^K [(e^{-\beta(k+1)\omega_\alpha} - e^{-\beta(k+1)\omega_{\alpha-1}}) - d(e^{-\beta(k+1)d\omega_\alpha} - e^{-\beta(k+1)d\omega_{\alpha-1}})] \\ &= \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N (\omega_s - \omega_\alpha) \sum_{k=0}^K [e^{-\beta(k+1)\omega_\alpha} (1 - e^{-\beta(k+1)(\omega_{\alpha-1} - \omega_\alpha)}) - d e^{-\beta d(k+1)\omega_\alpha} (1 - e^{-\beta d(k+1)(\omega_{\alpha-1} - \omega_\alpha)})].\end{aligned}\quad (\text{C19})$$

Here, since both $H_{\mathcal{M}_\alpha}$ and H_s are equally spaced Hamiltonians, the average energy can be written as

$$E(\omega_x, \omega_y) = \text{tr} [H_s(\omega_x) \tau_s(\beta, \omega_y)] = \frac{\sum_{n=0}^{d-1} n \omega_x e^{-n\beta\omega_y}}{\sum_{n=0}^{d-1} e^{-n\beta\omega_y}} = \omega_x \left(\frac{e^{-\beta\omega_y}}{1 - e^{-\beta\omega_y}} - \frac{d e^{-\beta d\omega_y}}{1 - e^{-\beta d\omega_y}} \right) \quad (\text{C20})$$

by evaluating the geometric series

$$\mathcal{Z}(\beta, \omega_y) = \sum_{n=0}^{d-1} e^{-\beta n \omega_y} = \frac{1 - e^{-\beta d \omega_y}}{1 - e^{-\beta \omega_y}} \quad (\text{C21})$$

and writing

$$E(\omega_x, \omega_y) = \sum_{n=0}^{d-1} n \omega_x \frac{e^{-\beta n \omega_y}}{\mathcal{Z}(\beta, \omega_y)} = \frac{\omega_x}{\omega_y} \left\{ -\frac{\partial}{\partial \beta} \log [\mathcal{Z}(\beta, \omega_y)] \right\} = -\frac{\omega_x}{\omega_y} \frac{\partial}{\partial \beta} [\log (1 - e^{-\beta d \omega_y}) - \log (1 - e^{-\beta \omega_y})] \quad (\text{C22})$$

as we do in the second line of Eq. (C19) and then using the infinite series expression $(1 - x)^{-1} = \lim_{K \rightarrow \infty} \sum_{k=0}^K x^k$ for any $|x| < 1$ as per the third line.

As we will see in Appendix E2a, the energy cost for cooling an infinite-dimensional system when both target and machines have equally spaced Hamiltonians (i.e., harmonic oscillators) is similar to the form of Eq. (C19). Importantly, the second term in square parenthesis vanishes as $d \rightarrow \infty$, simplifying the expression even further.

We now assume that the energy gaps of the machine are given by $\omega_\alpha = \omega_s + \epsilon \alpha$ and so the total energy cost can be written as follows:

$$\begin{aligned}\Delta E_{s\mathcal{M}}^{(N)} &= - \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N \alpha \epsilon \sum_{k=0}^K e^{-\beta k(\omega_s + \alpha \epsilon)} (1 - e^{\beta k \epsilon}) + \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N \alpha d \epsilon \sum_{k=0}^K e^{-\beta k d(\omega_s + \alpha \epsilon)} (1 - e^{\beta k d \epsilon}) \\ &= \lim_{K \rightarrow \infty} \sum_{k=0}^K [e^{-\beta k \omega_s} (e^{\beta k \epsilon} - 1) \left(\sum_{\alpha=1}^N \alpha \epsilon e^{-\beta k \alpha \epsilon} \right)] - \lim_{K \rightarrow \infty} \sum_{k=0}^K e^{-\beta k d \omega_s} [(e^{\beta k d \epsilon} - 1) \left(\sum_{\alpha=1}^N d \alpha \epsilon e^{-\beta k d \alpha \epsilon} \right)],\end{aligned}\quad (\text{C23})$$

where we can swap the order of summation since both sums converge and the summands are non-positive. This can be seen from the first line above, using the fact that $e^{-\alpha x}(1 - e^x) \in [-1, 0]$ for all $\alpha \geq 1$ and $x \geq 0$. We now calculate the sum over α .

$$\begin{aligned}\sum_{\alpha=1}^N \alpha \epsilon e^{-\beta \alpha \epsilon} &= -\frac{\partial}{\partial \beta} \sum_{\alpha=0}^N e^{-\beta \alpha \epsilon} = -\frac{\partial}{\partial \beta} \left(\frac{1 - e^{-\beta(N+1)\epsilon}}{1 - e^{-\beta \epsilon}} \right) \\ &= - \left(\frac{(N+1)\epsilon e^{-\beta(N+1)\epsilon} - (N+1)\epsilon e^{-\beta(N+2)\epsilon} - \epsilon e^{-\beta \epsilon} + \epsilon e^{-\beta(N+2)\epsilon}}{(1 - e^{-\beta \epsilon})^2} \right)\end{aligned}$$

$$\begin{aligned}
&= \frac{\epsilon e^{-\beta\epsilon}}{(1 - e^{-\beta\epsilon})^2} (1 - (N+1)e^{-\beta N\epsilon} + N e^{-\beta(N+1)\epsilon}) \\
&= \frac{\epsilon e^{-\beta\epsilon}}{(1 - e^{-\beta\epsilon})^2} (1 - e^{-\beta N\epsilon} - N e^{-\beta N\epsilon} (1 - e^{-\beta\epsilon})).
\end{aligned} \tag{C24}$$

Combining Eqs. (C23) and (C24), we arrive at

$$\begin{aligned}
\Delta E_{\mathcal{SM}}^{(N)} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{k} \frac{k \epsilon (1 - e^{-\beta N k \epsilon})}{(1 - e^{-\beta k \epsilon})} - N \epsilon e^{-\beta k (\omega_S + N \epsilon)} \right] \\
&\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{k} \frac{k d \epsilon (1 - e^{-\beta N k d \epsilon})}{(1 - e^{-\beta k d \epsilon})} - N d \epsilon e^{-\beta k d (\omega_S + N \epsilon)} \right].
\end{aligned} \tag{C25}$$

In order to optimise the energy cost, we now assume that the energy gaps of the machines can be chosen to be smoothly increasing in such way that $\epsilon = \Delta\omega/N := (\omega_{\max} - \omega_S)/N$. Substituting this expression for ϵ into the above equation yields

$$\begin{aligned}
\Delta E_{\mathcal{SM}}^{(N)} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{k} \frac{k \Delta\omega (1 - e^{-\beta k \Delta\omega})}{N (1 - e^{-\beta k \frac{\Delta\omega}{N}})} - \Delta\omega e^{-\beta k (\omega_S + \Delta\omega)} \right] \\
&\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{k} \frac{k d \Delta\omega (1 - e^{-\beta k d \Delta\omega})}{N (1 - e^{-\beta k d \frac{\Delta\omega}{N}})} - d \Delta\omega e^{-\beta k d (\omega_S + \Delta\omega)} \right].
\end{aligned} \tag{C26}$$

We now wish to take the limit of $N \gg K \rightarrow \infty$. This assumption means that energy change of the system is approximately equal to its free energy change; in other words, the process occurs quasi-adiabatically. The ability to switch the order of taking the limits of K and N going to ∞ follows from the monotonic convergence of the sum over k . In particular, note that the term inside square parentheses in each summand converges and the first term in each summation (which is the only part that depends on N) is positive and bounded.

Under this assumption, we can use the approximation $\lim_{\beta x \rightarrow 0} \frac{x}{1 - e^{-\beta x}} = \frac{1}{\beta}$; since $0 < e^{-\beta x} < 1$ for any positive x , the sum over k converges to a finite value. In general, this approximation introduces a correction term for the energy change, however under said assumption the error incurred becomes negligible. Then, the total energy change $\Delta E_{\mathcal{SM}}^{\text{tot}}$ for the transformation $\tau_S(\beta, \omega_S) \rightarrow \tau_S(\beta, \omega_{\max})$ throughout the overall process is

$$\begin{aligned}
\Delta E_{\mathcal{SM}}^{\text{tot}} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{\beta k} - \frac{e^{-\beta k \omega_{\max}}}{\beta k} - (\omega_{\max} - \omega_S) e^{-\beta k \omega_{\max}} \right] \\
&\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{\beta k} - \frac{e^{-\beta k d \omega_{\max}}}{\beta k} - d (\omega_{\max} - \omega_S) e^{-\beta k d \omega_{\max}} \right].
\end{aligned} \tag{C27}$$

As a side remark, note that here one can see that in the special case of equally spaced Hamiltonians, one indeed requires a diverging number of machine subsystems to attain perfect cooling at the Landauer limit, as this is the only way to fulfil the condition of Theorem 3. This follows from the fact that the approximation $\frac{x}{1 - e^{-\beta x}} \approx \frac{1}{\beta}$ only holds for small βx and in general one would need to include higher-order terms that lead to an increase in energy cost.

We then have, using the expression for $E(\omega_x, \omega_y)$ derived earlier:

$$\begin{aligned}
\Delta E_{\mathcal{SM}}^{\text{tot}} &= -\frac{1}{\beta} \log(1 - e^{-\beta \omega_S}) + \frac{1}{\beta} \log(1 - e^{-\beta \omega_{\max}}) - \frac{(\omega_{\max} - \omega_S) e^{-\beta \omega_{\max}}}{1 - e^{-\beta \omega_{\max}}} \\
&\quad + \frac{1}{\beta} \log(1 - e^{-\beta d \omega_S}) - \frac{1}{\beta} \log(1 - e^{-\beta d \omega_{\max}}) + \frac{d(\omega_{\max} - \omega_S) e^{-\beta d \omega_{\max}}}{1 - e^{-\beta d \omega_{\max}}} \\
&= \frac{1}{\beta} \log \left(\frac{1 - e^{-\beta d \omega_S}}{1 - e^{-\beta \omega_S}} \right) - \frac{1}{\beta} \log \left(\frac{1 - e^{-\beta d \omega_{\max}}}{1 - e^{-\beta \omega_{\max}}} \right) - (\omega_{\max} - \omega_S) \left(\frac{e^{-\beta \omega_{\max}}}{1 - e^{-\beta \omega_{\max}}} - \frac{d e^{-\beta d \omega_{\max}}}{1 - e^{-\beta d \omega_{\max}}} \right) \\
&= \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_S)] - \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_{\max})] - \text{tr}[H_S(\omega_{\max}) \tau_S(\beta, \omega_{\max})] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_{\max})] \\
&= \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_S)] - \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_{\max})] \\
&\quad - \text{tr}[H_S(\omega_{\max}) \tau_S(\beta, \omega_{\max})] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_S)] - \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_S)] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_{\max})]
\end{aligned}$$

$$= \frac{1}{\beta} \Delta S_s + \Delta E_s, \quad (\text{C28})$$

where we have explicitly written the von Neumann entropy $S(\varrho) = -\text{tr}[\varrho \log(\varrho)]$ of a thermal state at inverse temperature β as $S[\tau_s(\beta, \omega)] = \log[\mathcal{Z}_s(\beta, \omega)] + \beta E[\tau_s(\beta, \omega)]$. Since the energy change of the target system only concerns its local Hamiltonian, we immediately see that the heat dissipated by the resetting of machines in such a cooling process, i.e., $\Delta E_{\mathcal{M}}$, saturates the Landauer bound as it is equal to $\beta^{-1} \Delta S_s$. The process described is thus energy-optimal.

Appendix D: Conditions for Structural and Control Complexity

Here we begin by considering the protocol-independent structural conditions that must be fulfilled by the machine Hamiltonian to enable (1) *perfect cooling* and (2) *cooling at Landauer cost*; combined, these independent conditions provide a necessary requirement, namely that the machine must be infinite-dimensional with a spectrum that is unbounded (from above) for the possibility of (3) *perfect cooling at the Landauer limit*. We then turn to analyse the control complexity, which concerns the properties of the interaction that implements a given protocol. The properties of the machine Hamiltonian define the *structural complexity*, which set the potential for how cool the target system can be made and at what energy cost; the extent to which a machine's potential is utilised in a particular protocol then depends on the properties of the joint unitary, i.e., the *control complexity*. Here, we show that it is necessary that any protocol achieving perfect cooling at the Landauer limit involves interactions between the target and infinitely-many levels of the machine to realise the full cooling potential. We then analyse some sufficient conditions that arise as observations from our diverging control complexity protocols. This then leads us to demonstrate that individual degrees of freedom of the machine must be addressed in a fine-tuned manner to permute populations, highlighting that an operationally meaningful notion of control complexity must take into account factors beyond the effective dimensionality.

D1. Necessary Complexity Conditions

D1a. Necessary Structural Conditions

1. *Perfect Cooling*.—Let us consider the task of perfect cooling, independently from protocol-specific constraints, in the envisaged setting. One can lower bound the smallest eigenvalue λ_{\min} of the final state ϱ'_s (and hence how cold the system can become) after *any* unitary interaction with a thermal machine by [29]

$$\lambda_{\min}(\varrho'_s) \geq e^{-\beta \omega_{\mathcal{M}}^{\max}} \lambda_{\min}(\varrho_s), \quad (\text{D1})$$

where $\omega_{\mathcal{M}}^{\max} := \max_{i,j} |\omega_j - \omega_i|$ denotes the largest energy gap of the machine Hamiltonian $H_{\mathcal{M}}$ with eigenvalues ω_i . Without loss of generality, throughout this paper we set the ground-state energy of any system to be zero, i.e., $\omega_0 = 0$, such that the largest energy gap coincides with the largest energy eigenvalue. As we make no restrictions on the size or structure of the target or machine, the above inequality pertains to cooling protocols that could, for instance, be realised via sequences of unitaries on the target and parts of the machine. It follows that perfect cooling is only possible under two conditions: either the machine begins in a pure state ($\beta \rightarrow \infty$), or $H_{\mathcal{M}}$ is unbounded, i.e., $\omega_{\mathcal{M}}^{\max} \rightarrow \infty$. Requiring $\beta < \infty$, a diverging energy gap in the machine Hamiltonian is thus a necessary structural condition for perfect cooling. Indeed, the largest energy gap of the machine plays a crucial role in limiting how cool the target system can be made (see also, e.g., Refs. [42, 51]). We now detail an independent property that is required for cooling with minimal energetic cost.

2. *Cooling at the Landauer Limit*.—Suppose now that one wishes to cool an initial target state $\tau_s(\beta, H_s)$ to any thermal state $\tau'_s(\beta^*, H_s)$ with $\beta^* > \beta$ (not necessarily close to a pure state), at an energy cost saturating the Landauer limit. In Ref. [29], it was shown that for any finite-dimensional machine, there are correction terms to the Landauer bound, which imply that it cannot be saturated; these terms vanish only in the limit where the machine dimension diverges. Thus, a necessary condition for achieving cooling with energy cost at the Landauer limit is provided by the following:

Theorem 11. *To cool a target system $\tau_s(\beta, H_s)$ to $\tau_s(\beta^*, H_s)$, with $\beta^* > \beta$, using a machine in the initial state $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ with energy cost at the Landauer limit, the machine must be infinite dimensional.*

As we will discuss below, this minimal requirement for the notion of complexity is far from sufficient to achieve cooling at Landauer cost.

3. *Perfect Cooling at the Landauer Limit*.—We have two independent necessary conditions on the structure of the machine that must be asymptotically achieved to enable relevant goals for cooling: the former is required to achieve perfect cooling; the latter for cooling at the Landauer limit. Together, these conditions imply that in order to achieve perfect cooling at the Landauer limit, one must have an infinite-dimensional machine with a spectrum that is unbounded (from above), as stated in Corollary 2.

Henceforth, we assume that these conditions are satisfied by the machine. The question then becomes: *how does one engineer an interaction between the target system and machine to achieve perfect cooling at Landauer cost?*

D1b. Necessary Control Complexity Conditions

The unbounded structural properties of the machine support the *possibility* for perfect cooling at the Landauer limit; however, we now focus on the control properties of the interaction that *realise* said potential (see Fig. 2). This leads to the distinct notion of *control complexity*, which aims to differentiate between protocols that access the machine in a more or less complex manner. The structural complexity properties are protocol independent and related to the energy spectrum and dimensionality of the machine, whereas the control complexity concerns properties of the unitary that represents a particular protocol. For instance, the diverging-time protocol previously outlined comprises a sequence of interactions, each of which is individually not very complex; at the same time, the unconstrained control complexity protocol accesses the total (overall infinite-dimensional) machine “at once”, and thus the number of (nontrivial) terms in the interaction Hamiltonian, or the effective dimensionality of the machine accessed by the unitary, becomes unbounded. Nonetheless, the net energy cost of this protocol with unconstrained control complexity remains in accordance with the Landauer limit, as the initial and final states of both the system and machine are identical to those in the diverging-time protocol.

Effective Dimensionality.—We begin by considering the effective dimensionality accessed (nontrivially) by a unitary, whose divergence is necessary but insufficient for achieving perfect cooling at the Landauer limit, as we show in the next section. This in turn motivates the desire for a more detailed notion of control complexity that takes into account the energy-level structure of the machine.

We define the effective dimension as the dimension of the subspace of the global Hilbert space upon which the unitary acts nontrivially, which can be quantified via the minimum dimension of a subspace \mathcal{A} of the joint Hilbert space $\mathcal{H}_{S\mathcal{M}}$ in terms of which the unitary can be decomposed as $U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}$, i.e.,

$$d^{\text{eff}} := \min \dim(\mathcal{A}) : U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}. \quad (\text{D2})$$

One can relate this quantity to properties of the Hamiltonian that generates the evolution in a finite unit of time T (which we can set equal to unity without loss of generality) by considering the interaction picture. In general, any global unitary $U_{S\mathcal{M}} = e^{-iH_{S\mathcal{M}}T}$ is generated by a Hamiltonian of the form $H_{S\mathcal{M}} = H_S \otimes \mathbb{1}_{\mathcal{M}} + \mathbb{1}_S \otimes H_{\mathcal{M}} + H_{\text{int}}$. However, all protocols considered in this work have vanishing local terms, i.e., $H_S = H_{\mathcal{M}} = 0$. More generally, one can argue that the local terms play no role in how the machine is used to cool the target. As such, one can consider unitaries generated by only the nontrivial term H_{int} to be those representing a particular protocol of interest. That is, we can restrict our attention to $U_{S\mathcal{M}} = e^{-iH_{\text{int}}T}$, where H_{int} is a Hermitian operator on $\mathcal{H}_{S\mathcal{M}}$ of the form $\sum_i A_S^i \otimes B_{\mathcal{M}}^i$ such that none of the $A_S^i, B_{\mathcal{M}}^i$ are proportional to the identity operator. In doing so, it follows that the effective dimension corresponds to $\text{rank}(H_{\text{int}})$. Lastly, note that the above definition in terms of a direct sum decomposition provides an upper bound on any similar quantification of effective dimensionality based on other tensor factorisations of the joint Hilbert space considered and makes no assumption about the underlying structure. On the other hand, knowledge of said structure would permit a more meaningful notion of complexity to be defined. For instance, the effective dimensionality of a unitary acting on a many qubit system is better captured by considering its decomposition into a tensor product factorisation rather than the direct sum. We leave the exploration of such considerations to future work.

The effective dimensionality provides a minimal quantifier for a notion of control complexity, insofar as its divergence is necessary for saturating the Landauer bound, as we prove in the next section. In fact, we prove a slightly stronger statement, namely that the dimension of the machine Hilbert space to which the unitary (nontrivially) couples the target system to must diverge. However, as we discuss below, $d^{\text{eff}} \rightarrow \infty$ is generally insufficient to achieve said goal, and fine-tuned control is required. Nonetheless, the manifestation of such control seems to be system dependent, precluding our ability (so far) to present a universal quantifier of control complexity. Thus, even though further conditions need to be met to achieve perfect cooling at minimal energy cost in unit time (see Theorem 12), whenever we talk of an operation with finite control complexity, we mean those represented by a unitary that acts (nontrivially) only on a finite-dimensional subspace of the target system and machine. In contrast, by diverging control complexity, we mean a unitary that couples the target (nontrivially) to a full basis of the machine’s Hilbert space, whose dimension diverges. With this notion at hand, we have Theorem 3, which is proven below. Intuitively, we show that if a protocol accesses only a finite-dimensional subspace of the machine, then the machine is effectively finite dimensional inasmuch as a suitable replacement can be made while keeping all quantities relevant for cooling invariant. Invoking then the main result of Ref. [29], there are finite-dimensional correction terms that then imply that the Landauer limit cannot be saturated.

Note finally that in Theorem 3 no particular structure of the systems is presupposed and the effective dimensionality relates to various notions of complexity put forth throughout the literature (see, e.g., Refs. [52, 53]). For instance, for a finite-dimensional target system with equally spaced energy levels ω_S , suppose that the machine structure is decomposed as N qubits with energy gaps $\omega_{\mathcal{M}_n} \in \{\omega_S + n\epsilon\}_{n=1,\dots,N}$, with arbitrarily small $\epsilon > 0$ and $N \rightarrow \infty$. Then the overall unitary that approaches perfect cooling at the Landauer limit has circuit complexity equal to the diverging N .

D2. Proof of Theorem 3, Corollary 2, and Theorem 11

Here we prove Theorem 3, which implies Theorem 11 and leads to Corollary 2.

Proof. Let $\mathcal{H}_{\mathcal{X}}$ be a separable Hilbert space associated with the system \mathcal{X} . Consider

$$H_{\mathcal{M}} = \sum_{n=0}^{\infty} \omega_n |n\rangle\langle n| \quad \text{and} \quad \mathcal{H}_{\mathcal{M}'} = \text{span}_{n \leq m} \{|n\rangle\}, \quad (\text{D3})$$

for some finite m . In other words, $\mathcal{H}_{\mathcal{M}'}$ is a finite-dimensional restriction of $\mathcal{H}_{\mathcal{M}}$. We show that any unitary that (nontrivially) interacts the target system with only a subspace spanned by finitely many eigenstates of $H_{\mathcal{M}}$ cannot attain Landauer's bound. Consider a general unitary U . Suppose that U couples only $\mathcal{H}_{\mathcal{S}}$ with $\mathcal{H}_{\mathcal{M}'}$; whenever we talk of an operation with finite effective dimension in this paper, we mean specifically such a U , and by diverging effective dimension we mean a unitary that couples the target to any subspace of $\mathcal{H}_{\mathcal{M}}$ whose dimension diverges. Since

$$\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}} = \mathcal{H}_{\mathcal{S}} \otimes (\mathcal{H}_{\mathcal{M}'} \oplus \mathcal{H}_{\mathcal{M}'}^{\perp}) \simeq (\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}) \oplus (\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}^{\perp}), \quad (\text{D4})$$

we can associate the subspace $\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}$ with the label \mathcal{A} and $\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}^{\perp}$ with \mathcal{B} and write $U = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{B}}$. Then the initial configuration can be expressed as

$$\varrho_{\mathcal{S}} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \begin{bmatrix} \varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'} & 0 \\ 0 & \varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}^{\perp} \end{bmatrix}, \quad (\text{D5})$$

where

$$\varrho_{\mathcal{M}'} := \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n \leq m} e^{-\beta \omega_n} |n\rangle\langle n| \quad \text{and} \quad \varrho_{\mathcal{M}'}^{\perp} := \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n > m} e^{-\beta \omega_n} |n\rangle\langle n| \quad (\text{D6})$$

add up to a (normalised) thermal state. Now consider the state

$$\tilde{\varrho}_{\mathcal{M}} = \begin{bmatrix} \varrho_{\mathcal{M}'} & 0 \\ 0 & \text{tr}[\varrho_{\mathcal{M}'}^{\perp}] \end{bmatrix}. \quad (\text{D7})$$

It is straightforward to check that is indeed a quantum state; moreover, it is the Gibbs state (at inverse temperature β) associated with the Hamiltonian

$$\tilde{H}_{\mathcal{M}} = \sum_{n \leq m} \omega_n |n\rangle\langle n| - \frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right) |m+1\rangle\langle m+1|. \quad (\text{D8})$$

To see this, note that $\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \mathcal{Z}_{\mathcal{M}}(\beta, \tilde{H}_{\mathcal{M}})$ and that

$$\exp \left\{ -\beta \left[-\frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right) \right] \right\} = \sum_{n > m} e^{-\beta \omega_n}. \quad (\text{D9})$$

Thus $\tilde{\varrho}_{\mathcal{M}} = \tau_{\mathcal{M}}(\beta, \tilde{H}_{\mathcal{M}})$. To ease notation in what follows, we write $\tilde{\omega}_{m+1} := -\frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right)$. In the rest of the proof, we show that the unitary U and the Hamiltonian $H_{\mathcal{M}}$ can be replaced by finite-dimensional versions without changing the quantities relevant for Landauer's principle.

Let $\tilde{U} = U_{\mathcal{A}} \oplus (\mathbb{1}_{\mathcal{S}} \otimes |m+1\rangle\langle m+1|)$. We then have

$$\tilde{U} (\varrho_{\mathcal{S}} \otimes \tilde{\varrho}_{\mathcal{M}}) \tilde{U}^{\dagger} = \begin{bmatrix} U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger} & 0 \\ 0 & \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_{\mathcal{S}} \end{bmatrix} \quad (\text{D10})$$

and

$$\text{tr}_{\mathcal{M}} \left[\tilde{U} (\varrho_{\mathcal{S}} \otimes \tilde{\varrho}_{\mathcal{M}}) \tilde{U}^{\dagger} \right] = \text{tr}_{\mathcal{M}'} \left[U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger} \right] + \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_{\mathcal{S}} \quad (\text{D11})$$

Compare this to the expression

$$\begin{aligned}
\text{tr}_{\mathcal{M}} [U(\varrho_S \otimes \varrho_{\mathcal{M}})U^\dagger] &= \text{tr}_{\mathcal{M}'} \begin{bmatrix} U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger & 0 \\ 0 & \varrho_S \otimes \varrho_{\mathcal{M}'}^\perp \end{bmatrix} \\
&= \text{tr}_{\mathcal{M}'} [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \text{tr} [\varrho_{\mathcal{M}'}^\perp] \varrho_S \\
&= \text{tr}_{\mathcal{M}'} [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \frac{e^{-\beta\tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_S,
\end{aligned} \tag{D12}$$

since $\text{tr} [\varrho_{\mathcal{M}'}^\perp] = \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n>m} e^{-\beta\omega_n}$. Thus, the final system state is the same as it would be if we replaced the full initial machine state with $\tilde{\varrho}_{\mathcal{M}}$; in particular, the entropy decrease of the system for any unitary that cools it is also unchanged.

The last thing we need to check is that the energy change of the machine similarly remains invariant. To that end, we have that

$$\begin{aligned}
\tilde{\varrho}'_{\mathcal{M}} &= \text{tr}_S [\tilde{U}(\varrho_S \otimes \tilde{\varrho}_{\mathcal{M}})\tilde{U}^\dagger] = \text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \frac{e^{-\beta\tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} |m+1\rangle\langle m+1| \\
\tilde{\varrho}_{\mathcal{M}} &= \varrho_{\mathcal{M}'} + \frac{e^{-\beta\tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} |m+1\rangle\langle m+1|.
\end{aligned} \tag{D13}$$

Thus, we have

$$\text{tr} [\tilde{H}_{\mathcal{M}}(\tilde{\varrho}'_{\mathcal{M}} - \tilde{\varrho}_{\mathcal{M}})] = \text{tr} \{H_{\mathcal{M}} [\text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] - \varrho_{\mathcal{M}'}]\}, \tag{D14}$$

since $U_{\mathcal{A}}$ only acts on $\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}'}$ and $\tilde{H}_{\mathcal{M}|\mathcal{M}'} = H_{\mathcal{M}|\mathcal{M}'}$. In the same way, we have

$$\begin{aligned}
\text{tr}_S [U(\varrho_S \otimes \varrho_{\mathcal{M}})U^\dagger] &= \text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \varrho_{\mathcal{M}'}^\perp \\
\varrho_{\mathcal{M}} &= \varrho_{\mathcal{M}'} + \varrho_{\mathcal{M}'}^\perp.
\end{aligned} \tag{D15}$$

Thus, the energy difference is also

$$\text{tr} \{H_{\mathcal{M}} [\text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] - \varrho_{\mathcal{M}'}]\}. \tag{D16}$$

Hence, we show that one can replace (a potentially infinite-dimensional) \mathcal{M} by some (finite) $m+1$ -dimensional machine $\tilde{\mathcal{M}}$ if the joint unitary U acts only on m levels of $H_{\mathcal{M}}$. By Theorem 6 of Ref. [29], there are finite-dimensional corrections to the Landauer bound, which then imply that it cannot be reached for finite m . Thus, the effective machine dimension, i.e., that which is actually (nontrivially) accessed throughout the interaction, must diverge in order for cooling to be possible at the Landauer limit. This proves Theorem 3, which implies Theorem 11. \square

D3. Sufficient Complexity Conditions

Having shown the necessary requirements for cooling at Landauer cost, namely a control interaction that acts nontrivially on an infinite-dimensional (sub)space of the machine's Hilbert space, let us now return to emphasise the properties of the machine and cooling protocol that are sufficient to achieve perfect cooling at Landauer cost. For simplicity, we consider the case of a qubit, which exemplifies the discussion of finite-dimensional systems. The case of infinite-dimensional systems is treated independently in the next Appendix.

We first consider the structural properties of the machine. The diverging-time protocol discussed in Appendix C makes use of a diverging number N of machines. Thus, the machine begins in the thermal state $\tau(\beta, H_{\mathcal{M}}^{\text{tot}})$ of a (2^N) -dimensional system (with N eventually diverging), with energy-level structure given by the sum of the Hamiltonians in Eq. (C2), i.e.,

$$H_{\mathcal{M}}^{\text{tot}} = \sum_{n=1}^N H_{\mathcal{M}_n}^{(n)} = \sum_n (1 + n\epsilon) H_S^{(n)}, \tag{D17}$$

that acts on the full Hilbert space (we use the usual convention that it acts as identity on unlabelled subspaces, e.g., $H_{\mathcal{M}}^{(1)} \equiv H_{\mathcal{M}}^{(1)} \otimes \mathbb{1}^{(2)} \otimes \dots \otimes \mathbb{1}^{(N)}$). Let us analyse in detail the properties of this Hamiltonian. The ground state is $|0\rangle^{\otimes N}$, which is set at zero energy. More generally, the energy eigenvalue corresponding to an eigenstate $|i_0, i_1, \dots, i_N\rangle$ is given by ω_1 multiplied by the number of indices i_k that are equal to 1, plus a sum of terms $k\epsilon$ where k is the label of each index equal to 1. Thus, the energy

eigenvalue of the eigenstate $|1, \dots, 1\rangle$ diverges as the number of subsystems diverges. At the same time, letting the factor ϵ go to zero renders all eigenstates with the same (constant) number of indices such that $i_k = 1$ approach the same energy. Thus, in the limit $\epsilon \rightarrow 0$, one obtains subspaces of energy $E_{\mathcal{M}}^{(k)} = k\omega_1$ with degeneracy given by $D_k = \binom{N}{k}$, which also diverges for each constant k and diverging N . Therefore, in addition to satisfying the structural conditions that are necessary for perfect cooling, as stated in Theorem 11, the machine used here features additional properties, which are crucially important for this particular protocol, in particular because they are sufficient for perfect cooling at Landauer cost. As a remark, we also emphasise that for fixed (large) N and (small) ϵ , the machine is finite dimensional and has a nondegenerate Hamiltonian without any energy levels formally at infinity.

Concerning the control complexity properties of the unitary that achieves perfect cooling in unit time, note that it is a cyclic shift operator, which can be written as

$$\begin{aligned} U_{S\mathcal{M}} &= \Pi_{n=1}^N \mathbb{S}_{S\mathcal{M}_n}^2 = \Pi_n \left(\sum_{i,j_n=0}^1 |i, j_1, \dots, j_n, \dots, j_N\rangle \langle j_n, j_1, \dots, i, \dots, j_N|_{S\mathcal{M}} \right) \\ &= \sum_{i,j_1 \dots j_N=0}^1 |i, j_1, \dots, j_N\rangle \langle j_N, i, j_1, \dots, j_{N-1}|_{S\mathcal{M}}. \end{aligned} \quad (\text{D18})$$

As it is evident from its form, this unitary acts nontrivially on all of the (divergingly many) energy levels of the machine. The only basis vectors of the system-plus-machine Hilbert space that are left invariant are $|i=0, j_1=0, \dots, j_N=0\rangle$ and $|i=1, j_1=1, \dots, j_N=1\rangle$.

D4. Fine-Tuned Control Conditions

Theorem 3 captures a notion of control complexity as a resource in a thermodynamically consistent manner, i.e., in line with Nernst's unattainability principle. However, following the discussion around Theorem 11 and that above, the protocols that we present that achieve perfect cooling at Landauer cost make use of machines and interactions with a far more complicated structure than suggested by the necessary condition of diverging effective dimensionality. In particular, we note that the interactions couple the target system to a diverging number of subspaces of the machine corresponding to distinct energy gaps in a fine-tuned manner. Moreover, there are a diverging number of energy levels of the machine both above and below the first excited level of the target. In this section, we begin by outlining the general conditions that perfect cooling at the Landauer limit entails, before presenting a more nuanced notion of control complexity in terms of the variety of distinct energy gaps in the machine in Appendix D5.

This suggests that an operationally meaningful quantifier of control complexity must take into account the energy-level structure of the machine that is accessed throughout any given protocol; additionally that of the target system plays a role. Indeed, both the final temperature of the target as well as the energy cost required to achieve this depends upon how the global eigenvalues are permuted via the cooling process. First, how cool the target becomes depends on the sum of the eigenvalues that are placed into the subspace spanned by the ground state. Second, for any fixed cooling amount, the energy cost depends on the constrained distribution of eigenvalues within the machine. Thus, in general, the optimal permutation of eigenvalues depends upon properties of both the target and machine.

For instance, consider an arbitrary initially thermal target qubit, whose state is given by $\text{diag}(p, 1-p)$ and a thermal machine of dimension $d_{\mathcal{M}}$ with spectrum $\{\lambda_{\mathcal{M}}^i\}_{i=0, \dots, d_{\mathcal{M}}-1}$. Now consider the decomposition of the joint Hilbert space into two orthogonal subspaces, \mathcal{B}_0 and \mathcal{B}_1 , corresponding to the ground and excited eigenspaces of the target. The initial joint state is $p \text{diag}(\lambda_{\mathcal{B}_0}^i) \oplus (1-p) \text{diag}(\lambda_{\mathcal{B}_1}^i)$, where we write $\lambda_{\mathcal{B}_j}^i$ to denote the i^{th} machine eigenvalue in the subspace \mathcal{B}_j . The total population in the subspaces \mathcal{B}_0 and \mathcal{B}_1 are p and $(1-p)$ respectively. To achieve perfect cooling one must permute the eigenvalues such that approximately a net transfer of population $(1-p)$ is moved from \mathcal{B}_1 to \mathcal{B}_0 . To do this, one can take any subset K of $\{\lambda_{\mathcal{B}_1}^i\}$ such that as $d_{\mathcal{M}} \rightarrow \infty$, $\sum_{i \in K} \lambda_{\mathcal{B}_1}^i \rightarrow (1-p)$ and a subset K' (with $|K| = |K'|$) from $\{\lambda_{\mathcal{B}_0}^i\}$ such that $\sum_{i \in K'} \lambda_{\mathcal{B}_0}^i \rightarrow 0$ and exchange them. Although the choice of eigenvalues permuted is nonunique, the requirement must be fulfilled for some sets to perfectly cool the target. For any pair of eigenvalues exchanged between the subspaces, demanding that the exchange costs minimal energy amounts to a fine-tuning condition of the form $\lambda_{\mathcal{M}}^i \rightarrow p\lambda_{\mathcal{B}_0}^i + (1-p)\lambda_{\mathcal{B}_1}^i$ that must be satisfied. In general, the fine-tuned eigenvalue conditions that must be asymptotically attained depend upon target and machine eigenvalues, making it difficult to derive a closed-form expression. However, in the restricted scenario in which the target qubit begins maximally mixed (i.e., at infinite temperature), the machine begins thermal at some $\beta > 0$ and of dimension $d_{\mathcal{M}}$, and that the unitary implemented is such that the target is cooled as much as possible, one can derive precise conditions in terms of the machine structure alone, as we demonstrate below. The case for higher-dimensional target systems is similar.

This discussion highlights the importance of capturing properties beyond the effective dimensionality, e.g., those regarding the distribution of machine (and, more generally, target system) eigenvalues, in order to meaningfully quantify control complexity

in thermodynamics. Our protocols display similar behaviour to that discussed above asymptotically. Moreover, the machines exhibit an energy-level structure such that every possible energy gap is present, i.e., the set of machine energy gaps $\{\omega_{ij} = \omega_i - \omega_j\}$ densely covers the interval $[\omega_S, \infty)$, where ω_S is the energy of the first excited level of the target. In Appendix D5, we demonstrate that indeed this condition is necessary for minimal-energy cost cooling.

Before doing so, we here first derive the fine-tuned control conditions that are asymptotically required for cooling at the Landauer limit. We begin with some general considerations before focusing on a special case for which an analytic expression can be derived. Furthermore, we demand that the unitary implemented is such that the target is cooled as much as possible: this does not preclude the possibility for cooling the target system less (albeit still close to a pure state) at a cost closer to the Landauer bound without satisfying all of the fine-tuning conditions. Nonetheless, in general there are a number of such conditions to be satisfied, and the special case serves as a pertinent example that demonstrates how the particular set of fine-tuning conditions for any considered scenario can be similarly derived.

Consider an arbitrary thermal target system and machine of finite dimensions, with respective spectra $\lambda_S := \{\lambda_S^0, \dots, \lambda_S^{d_S-1}\}$ and $\lambda_M := \{\lambda_M^0, \dots, \lambda_M^{d_M-1}\}$. The states begin uncorrelated, so the global spectrum of the initial joint state is $\lambda_{SM} := \{\lambda_{SM}^0, \dots, \lambda_{SM}^{d_S d_M - 1}\} = \{\lambda_S^0 \lambda_M^0, \lambda_S^0 \lambda_M^1, \dots, \lambda_S^{d_S-1} \lambda_M^{d_M-1}\}$. Consider now a global unitary transformation; such a transformation cannot change the values of the spectrum, but merely permute them. In other words, the spectrum of the final global state after any such unitary is invariant and we have equivalence of the (unordered) sets λ'_{SM} and λ_{SM} .

The transformation that cools the target system as much as possible⁵ is the one that places the d_M largest of the global eigenvalues into the subspace spanned by the ground state of the target, the second d_M largest into that spanned by the first excited state of the target, and so forth, with the smallest d_M global eigenvalues placed in the subspace corresponding to the highest energy eigenstate of the target system (we prove this statement shortly). More precisely, we denote by λ^\downarrow the nonincreasing ordering of the set λ . Since the target and machine begin thermal, the local spectra λ_S and λ_M are already ordered in this way with respect to their energy eigenbases, which we consider to be labelled in nondecreasing order. Cooling the target system as much as possible amounts to achieving the final reduced state of the target

$$\varrho'_S = \sum_{i=0}^{d_S-1} \left(\sum_{j=0}^{d_M-1} \lambda_{SM}^{\downarrow id_M + j} \right) |i\rangle\langle i|_S. \quad (D19)$$

As a side remark, note that since each of the global eigenvalues are a product of the initial local eigenvalues (due to the initial tensor product structure), which are in turn related to the energy-level structure of the target system and machine (as they begin as thermal states), one can already see here that in order to approach perfect cooling, the machine must have some diverging energy gaps, such that the (finite) sum of the global eigenvalues contributing to the ground-state population of the target approaches 1.

Of course, there is an equivalence class of unitaries that can achieve the same amount of cooling; in particular, any permutation of the set of the d_M global eigenvalues *within* each energy eigenspace of the target system achieves the same amount of cooling, since it is the sum of these values that contribute to the total population in each subspace. Importantly, although such unitaries cool the target system to the same extent, their effect on the machine differs, and therefore so too does the energy cost of the protocol. However, demanding that such cooling is achieved at minimal energy cost amounts to a unique constraint on the global post-transformation state, namely that it must render the machine energetically passive, leading to the form:

$$\varrho'_{SM} = \sum_{i=0}^{d_S-1} \sum_{j=0}^{d_M-1} \lambda_{SM}^{\downarrow id_M + j} |ij\rangle\langle ij|_{SM}. \quad (D20)$$

We can derive the above form of the final joint state as follows. Consider the following ordering for the energy eigenbasis of SM chosen to match the above form

$$\{|00\rangle_{SM}, |01\rangle_{SM}, \dots, |0, d_M - 1\rangle_{SM}, |10\rangle_{SM}, \dots, |1, d_M - 1\rangle_{SM}, \dots, |d_S - 1, 0\rangle_{SM}, \dots, |d_S - 1, d_M - 1\rangle_{SM}\}. \quad (D21)$$

This ordering is monotonically nondecreasing primarily with respect to the energy of S , and secondarily w.r.t. M . We take the final state ϱ'_{SM} to be expressed in this basis. To maximise the cooling in a single unitary operation, we maximise the sum of the first $k \cdot d_M$ diagonal elements, for each $k \in \{1, 2, \dots, d_S\}$, as each sum corresponds to the total population in the k^{th} lowest energy eigenstate of S . The initial state ϱ_{SM} is diagonal in this basis, so the vector of initial diagonal elements, which we label $\theta := \text{diag}(\varrho_{SM})$, is also the vector of eigenvalues, λ_{SM} , i.e., $\theta = \lambda_{SM}$. Furthermore, since the unitary operation leaves the set of eigenvalues invariant, we have via the Schur-Horn lemma [54] that the vector of final diagonal elements, which we label

⁵ We take majorisation among passive states to be the measure of cooling; this implies the highest possible ground state population and purity, and

lowest possible entropy and average energy via Schur convexity.

$\theta' := \text{diag}(\varrho'_{S\mathcal{M}})$, is majorised by the vector of initial ones, i.e., $\theta' \prec \theta$. It follows that the partial sums we wish to maximise are upper bounded by the corresponding partial sums of the $k \cdot d_{\mathcal{M}}$ largest diagonal elements of the initial state. We claim that the unitary that cools this maximal cooling amount at minimum energy cost is the one that permutes the diagonal elements to be ordered w.r.t. the basis ordering in Eq. (D21).

More precisely, via the Schur-Horn lemma, one can always write $\theta' = D\theta$, with D a doubly stochastic matrix. The partial sums of the $k \cdot d_{\mathcal{M}}$ first elements are linear functions of the elements of θ . Thus the maximum values are obtained at the extremal points of the convex set of doubly stochastic matrices, which are the permutation matrices, via the Birkhoff-von Neumann theorem [54]. One can see by inspection that the optimal permutation matrices are the ones that place the largest $d_{\mathcal{M}}$ diagonal elements in the first block (i.e., the ground-state eigenspace of S), the next largest $d_{\mathcal{M}}$ elements in the second block (i.e., the first excited-state eigenspace of S), and so on. Within each block, the ordering does not affect the cooling of the target, so there is an equivalence class of permutations that satisfy the maximal cooling criterion. However, adding the optimisation over the energy cost eliminates this freedom. We may consider the reduced set of stochastic matrices that satisfy maximal cooling, generated by the permutations described above. Since the average energy of the final state is again a linear function of the diagonal elements, here too the minimum corresponds to a permutation matrix. Clearly the permutation that minimises the average energy is the one that orders the elements within each block to be decreasing w.r.t. the energies of \mathcal{M} . Thus, the unique⁶ stochastic matrix D that leads to maximal cooling at the least energy cost possible is the one that permutes the energy eigenvalues to be ordered decreasing primarily w.r.t. the system energies, and secondarily w.r.t. the machine energies. The action of the stochastic matrix on diagonal elements of the state is related to the unitary operation on the entire quantum state by $|U_{ij}|^2 = D_{ij}$, so that the unitary operation is also a permutation (up to an energy-dependent phase, which is irrelevant since the initial and final states are diagonal).

We may understand this optimal operation through the notion of passivity, by noting that it cools at minimal energy cost by rendering the machine into the most energetically passive reduced state in the joint unitary orbit with respect to the cooling constraint on the target. Intuitively, one has cooled the target system maximally at the expense of heating the machine as little as possible. The final reduced state of the machine corresponding to this energetically optimal cooling transformation is

$$\varrho'_{\mathcal{M}} = \sum_{j=0}^{d_{\mathcal{M}}-1} \left(\sum_{i=0}^{d_S-1} \lambda_{S\mathcal{M}}^{\downarrow i d_{\mathcal{M}} + j} \right) |j\rangle\langle j|_{\mathcal{M}}. \quad (\text{D22})$$

In general, any unitary that achieves these desired conditions simultaneously depends upon the energy-level structure of both the target system and machine, precluding a closed-form set of conditions that can be expressed only in terms of the machine. However, for the special case of a maximally mixed initial target state (i.e., cooling a thermal state at infinite temperature or erasing quantum information from its most entropic state), one can deduce this ordering precisely and moreover relate it directly to properties of the machine Hamiltonian, as we now demonstrate. In the following, we assume that $d_{\mathcal{M}}$ is even; the case for odd $d_{\mathcal{M}}$ can be derived similarly.

Theorem 12. *Consider the target system to begin in the maximally mixed state and a thermal machine at temperature $\beta > 0$, whose eigenvalues are labelled in nonincreasing order, $\{\lambda_{\mathcal{M}}^{\downarrow i}\}_{i=0, \dots, d_{\mathcal{M}}-1}$. In order to cool the target perfectly, with the restriction that the target must be cooled as much as possible, at an energy cost that saturates the Landauer limit, the machine eigenvalues must satisfy*

$$\sum_{i=0}^{\frac{d_{\mathcal{M}}}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1, \quad \sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0, \quad (\text{D23})$$

and

$$\frac{\frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor} \right)}{\lambda_{\mathcal{M}}^{\downarrow i}} \rightarrow 1 \quad (\text{D24})$$

for all $i \in \{0, \dots, d_{\mathcal{M}} - 1\}$, where $\lfloor \cdot \rfloor$ denotes the floor function and \rightarrow denotes that the condition is satisfied asymptotically, i.e., as $d_{\mathcal{M}} \rightarrow \infty$ ⁷.

⁶ Note that degeneracies in energy eigenvalues would lead to sets of equal diagonal elements, and prevent one from choosing a unique permutation. However, as the state in such degenerate subspaces is proportional to the identity matrix, we may take any unitary that is block diagonal w.r.t. the degeneracies without affecting the state, and hence the final cooling or av-

erage energy change.

⁷ Strictly speaking, in the limit $d_{\mathcal{M}} \rightarrow \infty$ the conditions in Eq. (D24) must only be satisfied for almost all i , i.e., for all but a small subset that contributes negligibly to the relative entropy, as we discuss below.

Proof. We consider a qubit for simplicity, but the generalisation to cooling an arbitrary-dimensional maximally mixed state is straightforward. The initial joint spectrum of the system and machine is

$$\lambda_{sM} = \frac{1}{2} \{\lambda_M^\downarrow, \lambda_M^\downarrow\} = \frac{1}{2} \{\lambda_M^{\downarrow 0}, \lambda_M^{\downarrow 1}, \dots, \lambda_M^{\downarrow d_M-1}, \lambda_M^{\downarrow 0}, \lambda_M^{\downarrow 1}, \dots, \lambda_M^{\downarrow d_M-1}\}. \quad (\text{D25})$$

As each $\lambda_M^{\downarrow i} = \frac{1}{Z_M(\beta, H_M)} e^{-\beta \omega_i}$ for any thermal state with Hamiltonian $H_M = \sum_i \omega_i |i\rangle\langle i|_M$ written with respect to nondecreasing energy eigenvalues, it follows that the globally ordered spectrum is

$$\lambda_{sM}^\downarrow = \frac{1}{2} \{\lambda_M^{\downarrow 0}, \lambda_M^{\downarrow 0}, \lambda_M^{\downarrow 1}, \lambda_M^{\downarrow 1}, \dots, \lambda_M^{\downarrow d_M-1}, \lambda_M^{\downarrow d_M-1}\}. \quad (\text{D26})$$

Expressing the global states with respect to the product of local energy eigenbases, we have that the initial joint state is $\frac{1}{2} \otimes \tau_M(\beta, H_M) = \text{diag}(\lambda_{sM})$ [see Eq. (D25)] and the unitary that cools the target as much as possible at minimum energy cost is the one achieving the globally passive final joint state $\varrho'_{sM} = \text{diag}(\lambda_{sM}^\downarrow)$. This leads to the following reduced states

$$\varrho'_s = \left(\sum_{i=0}^{\frac{d_M}{2}-1} \lambda_M^{\downarrow i} \right) |0\rangle\langle 0|_s + \left(\sum_{i=\frac{d_M}{2}}^{d_M-1} \lambda_M^{\downarrow i} \right) |1\rangle\langle 1|_s, \quad (\text{D27})$$

$$\begin{aligned} \varrho'_M &= \frac{1}{2} \left(\lambda_M^{\downarrow 0} + \lambda_M^{\downarrow \frac{d_M}{2}} \right) |0\rangle\langle 0|_M + \frac{1}{2} \left(\lambda_M^{\downarrow 0} + \lambda_M^{\downarrow \frac{d_M}{2}} \right) |1\rangle\langle 1|_M + \\ &+ \frac{1}{2} \left(\lambda_M^{\downarrow 1} + \lambda_M^{\downarrow \frac{d_M}{2}+1} \right) |2\rangle\langle 2|_M + \frac{1}{2} \left(\lambda_M^{\downarrow 1} + \lambda_M^{\downarrow \frac{d_M}{2}+1} \right) |3\rangle\langle 3|_M + \dots \end{aligned} \quad (\text{D28})$$

Intuitively, the reduced target state has the larger half of the initial machine eigenvalues in the ground state and the smaller half in the excited state; the reduced machine state has the sum of the largest elements from each of these halves in its ground state, the next largest element from each half (which, in this case, is equal to the first) in its first excited state, and so forth.

Let us denote the spectrum of the final state of the machine by $\lambda_M^{\downarrow} := \{\lambda_M^{\downarrow 0}, \lambda_M^{\downarrow 1}, \dots, \lambda_M^{\downarrow d_M-1}\} = \frac{1}{2} \{\lambda_M^{\downarrow 0} + \lambda_M^{\downarrow \frac{d_M}{2}}, \lambda_M^{\downarrow 0} + \lambda_M^{\downarrow \frac{d_M}{2}+1}, \dots, \lambda_M^{\downarrow \frac{d_M}{2}-1} + \lambda_M^{\downarrow d_M-1}\}$. Importantly, by construction, the reduced state of the final machine has its local eigenvalues in nonincreasing order, i.e., it is energetically passive.

We therefore have the final reduced states of the protocol that cools the initially maximally mixed target as much as possible at minimal energy cost, in particular with minimal heat dissipation by the machine, given the structural resources at hand. We can now analyse the properties that are required to saturate the Landauer limit by considering the terms on the r.h.s. of Eq. (3) for any fixed initial inverse temperature of the machine $\beta \geq 0$.

First note that cooling the target system by any amount fixes the change in entropy of the target system, so the first term is irrelevant. The second term concerns the mutual information built up between the target system and machine. In general, this is nonvanishing, although one can achieve any desired amount of cooling without generating such correlations (as per our constructions). Furthermore, in the case where one wants to consider attaining a perfectly cool final state, as we do here, the final reduced state of the target is approximately pure and so $I(S : M)_{\varrho'_{sM}} \rightarrow 0$. In terms of the reduced states above, this means that $\sum_{i=0}^{\frac{d_M}{2}-1} \lambda_M^{\downarrow i} \rightarrow 1$ and $\sum_{i=\frac{d_M}{2}}^{d_M-1} \lambda_M^{\downarrow i} \rightarrow 0$, which can occur only if the largest half of energy eigenvalues of the machine, i.e., ω_i for all $i \geq \frac{d_M}{2}$, diverge (since the summation contains only non-negative summands).

The final term that must be minimised to saturate the Landauer limit is the relative entropy of the final with respect to the initial machine state, $D(\varrho'_M \| \varrho_M)$. Here one can already see that an infinite-dimensional machine is required to saturate the Landauer bound: from Ref. [29], $D(\varrho'_M \| \varrho_M) \geq f(\Delta S_M, d_M)$, where f is a dimension-dependant function of the entropy difference of the machine that exhibits non-negative correction terms that vanish only in the limit $d_M \rightarrow \infty$. The relative entropy vanishes iff $\varrho_M = \varrho'_M$; moreover, by Pinsker's inequality one has $\frac{1}{2} \|\varrho_M - \varrho'_M\|_1^2 \leq D(\varrho_M \| \varrho'_M)$, so one can bound the trace distance between the initial and final state of the machine for any desired value of the relative entropy. Although $\varrho_M = \varrho'_M$ implies a trivial process that cannot cool the (initially thermal) target system, as our protocols that saturate the Landauer limit demonstrate, there are processes that asymptotically display the behaviour $\varrho'_M \rightarrow \varrho_M$ and cool the target system. For the asymptotic machine states to converge, in particular, their eigenvalues must become approximately equal asymptotically. Demanding this on the spectrum in Eq. (D28) leads to a generic term that must be asymptotically satisfied of the form:

$$\frac{\frac{1}{2} \left(\lambda_M^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_M^{\downarrow \frac{d_M}{2} + \lfloor \frac{i}{2} \rfloor} \right)}{\lambda_M^{\downarrow i}} \rightarrow 1 \quad \forall i \in \{0, \dots, d_M - 1\}. \quad (\text{D29})$$

□

In order to achieve perfect cooling at the Landauer limit, one thus must simultaneously satisfy the conditions outlined in Theorem 12. In other words, to minimise the relative-entropy term with the additional constraints $\sum_{i=0}^{\frac{d_{\mathcal{M}}}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1$ and $\sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0$. The first thing to note is that since the eigenvalues $\lambda_{\mathcal{M}}^{\downarrow i}$ contribute to different sums depending on whether i is in the larger half $\{0, \dots, \frac{d_{\mathcal{M}}}{2} - 1\}$ or smaller half $\{\frac{d_{\mathcal{M}}}{2}, \dots, d_{\mathcal{M}}\}$, one cannot have $\lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor} = \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} \forall i$ (i.e., a completely degenerate machine), since then both summations would be over identical values and there is no way for them to converge to distinct values. This precludes the trivial solution that satisfies the constraints of Eq. (D24) alone, namely the maximally mixed machine state, which cannot be used to perform any cooling [as, in particular, it does not satisfy the constraints of Eq. (D23)]. For the conditions to be simultaneously satisfied, we intuitively require that, although they must be distinct, for each i both $\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor}$ and $\lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor}$ become “close” to each other, but with a difference that decays rapidly as $d_{\mathcal{M}} \rightarrow \infty$, such that in the infinite-dimensional limit the larger “half” of the eigenvalues sum to one and the smaller “half” sum to zero. A subtle point to note is that because the relative entropy involves the ratio of final to original eigenvalues it is not enough that the absolute difference $|\lambda_{\mathcal{M}}^{\downarrow i} - \lambda_{\mathcal{M}}^{\downarrow i}|$ goes to zero, as in the infinite $d_{\mathcal{M}}$ limit, it is possible for this to happen for all of the eigenvalues approaching zero without the ratios of final to initial eigenvalues approaching unity (and hence the relative entropy not vanishing). One manner of satisfying such a constraint, as evidenced by the construction we proceed with next, is for the ratios of final to initial eigenvalues go to unity for all but a small number energy levels, with the population in this exceptional subspace going to zero in the infinite $d_{\mathcal{M}}$ limit (along with the ratios not diverging within said subspace).

The natural question that arises here is whether or not it is possible to satisfy these constraints concurrently. (Note that none of the cooling protocols provided throughout this paper use the max-cooling operation, so do not necessarily serve as examples.) To this end, we now construct a family of machine Hamiltonians $H_{\mathcal{M}}$ of increasing dimension that in the limit $d_{\mathcal{M}} \rightarrow \infty$ manages to attain both perfect cooling of a maximally mixed qubit and the Landauer limit for the energy cost using the maximal cooling operation discussed above. The form of the Hamiltonian is instructive regarding the complexity requirements for perfect cooling at the Landauer limit. The construction is inspired by the infinite-dimensional Hamiltonian found in Ref. [29] (Appendix D), therein used to perfectly cool a qubit with energy cost arbitrarily close to the Landauer limit. Their construction already begins with infinitely many machine eigenvalues, as well as infinitely many of them corresponding to diverging energy levels. In the following, we demonstrate that one can arbitrarily closely attain perfect cooling and the Landauer limit with finite-dimensional Hamiltonians, and by taking the limit $d_{\mathcal{M}} \rightarrow \infty$, recover the result of Ref. [29].

The Hamiltonian of the machine is $d_{\mathcal{M}} := 2^{N+1}$ dimensional,

$$H_{\mathcal{M}} = \sum_{n=0}^N \sum_{j=1}^{2^n} \left(n\Delta |n; j\rangle\langle n; j|_{\mathcal{M}} \right) + N\Delta |N; 2^N+1\rangle\langle N; 2^N+1|_{\mathcal{M}} \quad (\text{D30})$$

Here, each energy eigenvalue labelled by n is 2^n -fold degenerate. Thus the ground state is unique, the first excited state is twofold degenerate, the second excited state fourfold degenerate, and so on, with the degeneracy doubling every energy level. In order to make the Hamiltonian of even dimensionality for convenience, we add an extra degenerate state to the final level [which makes this level $(2^N + 1)$ -fold degenerate]. Also note that the Hamiltonian is equally spaced with energy gap Δ . In the following, we use the index n to denote any one of the degenerate states in the n^{th} energy level from $n = 0$ to $n = N$, and the index i to denote individual energy eigenstates from $i = 1$ to $i = 2^{N+1}$ (note that in contrast to the previous section, we are here beginning with $i = 1$ in order to simplify some future notation). With these indices, the eigenvalues are related by

$$\lambda_{\mathcal{M}}^{\downarrow i} = e^{-\beta\Delta} \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} \quad \forall i \in \{2, \dots, d_{\mathcal{M}} - 1\}, \quad (\text{D31})$$

$$\lambda_{\mathcal{M}}^{\downarrow n} = e^{-\beta\Delta} \lambda_{\mathcal{M}}^{\downarrow n-1} \quad \forall n \in \{1, \dots, N\}. \quad (\text{D32})$$

We introduce a parameter ϵ to express the Gibbs ratio as

$$e^{-\beta\Delta} = \frac{1 - \epsilon}{2}, \quad (\text{D33})$$

where $0 < \epsilon < 1$, and we eventually take the limit $\epsilon \rightarrow 0$ appropriately as the dimension diverges. Note that this constrains the Gibbs ratio to be smaller than $\frac{1}{2}$, which in turn ensures that the total population over all of the degenerate eigenstates in the n^{th} level is smaller than that in the $(n - 1)^{\text{th}}$ level (as it has twice the number of eigenstates, but less than half the population in each). If this constraint failed to hold, then in the asymptotic limit, all of the population would lie in energy levels that diverge.

We now consider using this machine to cool a maximally mixed qubit target. The final ground-state population of the qubit under the maximal cooling operation is the sum over the larger half of the eigenvalues of the machine, corresponding to the eigenvalues from $i = 1$ to $i = 2^N$ (equivalently, from $n = 0$ to $n = N - 1$ plus a single eigenvalue from the $n = N$ energy level), and is thus given by

$$p'_0 = \frac{1}{\mathcal{Z}_{\mathcal{M}}} \left(\sum_{n=0}^{N-1} 2^n \left(\frac{1 - \epsilon}{2} \right)^n + \left(\frac{1 - \epsilon}{2} \right)^N \right), \quad (\text{D34})$$

$$\text{where } \mathcal{Z}_{\mathcal{M}} = \sum_{n=0}^N 2^n \left(\frac{1-\epsilon}{2} \right)^n + \left(\frac{1-\epsilon}{2} \right)^N \quad (\text{D35})$$

is the partition function of the machine. The geometric series above evaluates to

$$p'_0 = \left(1 + \frac{\epsilon(1-\epsilon)^N}{1 - (1-\epsilon)^N + \epsilon(1-\epsilon)^N 2^{-N}} \right)^{-1}. \quad (\text{D36})$$

As an ansatz, supposing that ϵ scales inversely with N as $\epsilon := \frac{\theta}{N}$ leads to the simplification $(1-\epsilon)^N \rightarrow e^{-\theta}$ as $d_{\mathcal{M}}$ (and hence N) diverges. The asymptotic behaviour of the ground-state population is thus

$$p'_0 = 1 - \frac{1}{N} \left(\frac{\theta}{e^\theta - 1} \right) + O\left(\frac{1}{N^2}\right), \quad (\text{D37})$$

and so $p'_0 \rightarrow 1$ in the $N \rightarrow \infty$ limit.

We now move to calculate the energy cost. Rather than considering the optimal max-cooling operation described above, we consider a slight modification in order to make the connection to the construction in Ref. [29] clear as well as to simplify notation. Nonetheless, the energy cost of this modified protocol upper bounds that of the max-cooling operation (for the same achieved ground-state population), and so showing that the Landauer limit is attained for the modified protocol implies that it would be too for the max-cooling protocol. The modification is simply to relabel the smallest eigenvalue of the machine $\lambda_{\mathcal{M}}^{2^{N+1}}$ as $\lambda_{\mathcal{M}}^0$, and treat it as the ground-state eigenvalue in the max-cooling operation. For general machine states, such a switch would lead to less cooling (if the same unitary were applied), but in this case it does not because the sum of the first half of the machine eigenvalues, from $i = 0$ to $i = 2^N - 1$, is the same as the original sum from $i = 1$ to $i = 2^N$, due to the relabelling $\lambda_0 = \lambda_{2^N}$, since they are both eigenvalues of states corresponding the maximum excited energy level of the machine spectrum. The spectrum of the final state of the machine is then given by

$$\lambda_{\mathcal{M}}^{\downarrow i} = \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}} \right) \quad \forall i \in \{0, \dots, d_{\mathcal{M}} - 1\}, \quad (\text{D38})$$

which leads to

$$\begin{aligned} \lambda_{\mathcal{M}}^{\downarrow 0} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow 2^N} \right) = \lambda_{\mathcal{M}}^{\downarrow 0}, & \lambda_{\mathcal{M}}^{\downarrow 1} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow 2^N} \right) = \lambda_{\mathcal{M}}^{\downarrow 0}, \\ \lambda_{\mathcal{M}}^{\downarrow i} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}} \right) & \forall i \in \{2, \dots, d_{\mathcal{M}} - 1\} \\ &= \frac{1}{2} \left(\frac{2}{1-\epsilon} \lambda_{\mathcal{M}}^{\downarrow i} + \lambda_{\mathcal{M}}^{n=N} \right) &= \frac{1}{2} \frac{1}{\mathcal{Z}_{\mathcal{M}}} \left[\left(\frac{1-\epsilon}{2} \right)^{n-1} + \left(\frac{1-\epsilon}{2} \right)^N \right], \end{aligned} \quad (\text{D39})$$

where we observe that the index $\lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}$ corresponds to the largest energy level of the machine for all i , and we use Eq. (D31) for the spectrum of initial eigenvalues. Using the index n instead to denote a generic eigenvalue of the n^{th} energy level, we have the simpler expression

$$\lambda_{\mathcal{M}}^{\downarrow n} = \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow n-1} + \lambda_{\mathcal{M}}^{\downarrow N} \right), \quad \forall n \in \{1, 2, \dots, N\}. \quad (\text{D40})$$

The energy cost can now be simply calculated from the difference in the average energy of the machine state,

$$\Delta E_{\mathcal{M}} = \sum_{i=0}^{d_{\mathcal{M}}-1} (\lambda_{\mathcal{M}}^{\downarrow i} - \lambda_{\mathcal{M}}^{\downarrow i}) \omega_i, \quad (\text{D41})$$

where we denote the i^{th} energy eigenvalue by ω_i . $\lambda_{\mathcal{M}}^{\downarrow 0}$ is unchanged, and although $\lambda_{\mathcal{M}}^{\downarrow 1}$ does change, $\omega_1 = 0$ corresponds to the ground state and thus this eigenvalue change does not affect the energy cost. We can thus express the energy cost in terms of the index n instead, starting from $n = 1$ (corresponding to $i = 2$ onward), as

$$\Delta E_{\mathcal{M}} = \sum_{n=1}^N (\lambda_{\mathcal{M}}^{\downarrow n} - \lambda_{\mathcal{M}}^{\downarrow n}) \omega_n = \frac{1}{\beta} \left[1 - \frac{2(1-\epsilon)^N}{1 - (1-\epsilon)^N + (1-2^{-N})(1-\epsilon)^N \epsilon} \right] \log \left(\frac{2}{1-\epsilon} \right). \quad (\text{D42})$$

As we did above, we parameterise $\epsilon = \frac{\theta}{N}$. The asymptotic behaviour of the energy cost is then

$$\beta \Delta E_{\mathcal{M}} = \log(2) + \frac{1}{N} \left(1 - \frac{2 \log(2)}{e^\theta - 1} \right) \theta + O\left(\frac{1}{N^2}\right), \quad (\text{D43})$$

or in terms of the decrease in entropy of the system,

$$\beta \Delta E_{\mathcal{M}} = \tilde{\Delta} S_{\mathcal{M}} + \frac{\log N}{N} \left(\frac{\theta}{e^\theta - 1} \right) + O\left(\frac{1}{N}\right). \quad (\text{D44})$$

Combining (D37) and (D43), we have that in the limit $N \rightarrow \infty$, which is also $d_{\mathcal{M}} \rightarrow \infty$, the ground-state population approaches 1—corresponding to perfect cooling—and the energy cost approaches $\beta^{-1} \log(2)$, which is the Landauer limit for the perfect erasure of a maximally mixed qubit.

To connect this construction to the constraints of Eq. (D24), note that in the limit $N \rightarrow \infty$ (recalling that $\epsilon = \frac{\theta}{N}$),

$$\frac{\lambda'^{\downarrow n}}{\lambda^{\downarrow n}} = \lim_{N \rightarrow \infty} \frac{\frac{1}{2} \left[\left(\frac{1-\epsilon}{2} \right)^{n-1} + \left(\frac{1-\epsilon}{2} \right)^N \right]}{\left(\frac{1-\epsilon}{2} \right)^n} = \lim_{N \rightarrow \infty} \left[\frac{1}{1-\epsilon} + \frac{1}{2^{N-n+1}} \left(\frac{e^{-\theta}}{(1-\epsilon)^n} \right) \right] = 1, \quad (\text{D45})$$

for all $n \geq 1$, leaving only the ground-state eigenvalue (corresponding to $n = 0$ and $i = 1$) not satisfying the condition. However, this term is actually a negative contribution to the relative entropy as this eigenvalue decreases, and in any case can be verified independently to approach zero.

To see this, note that a necessary condition that ensures the contribution of any set of eigenvalues that do not satisfy Eq. (D24) to the relative entropy to be negligible is that the total population of the relevant subspace is vanishingly small. Writing the relative entropy between two states in terms of their eigenvalues, we have $D(\varrho' \parallel \varrho) = \sum_n \lambda'_n \log \left(\frac{\lambda'_n}{\lambda_n} \right)$, which we split up into two sets: S_0 containing all n for which Eq. (D24) is satisfied and S_{\pm} containing the all n for which Eq. (D24) is not satisfied. The contribution of the first term to the relative entropy is asymptotically zero, so we are left with $D(\varrho' \parallel \varrho) = \sum_{n \in S_{\pm}} \lambda'_n \log \left(\frac{\lambda'_n}{\lambda_n} \right)$. For each term in the sum here, one can write $\lambda'_n = \lambda(1 + \Delta_n)$ with the condition $|\Delta_n| \geq \theta > 0$ for some θ , i.e., the ratio of eigenvalues is bounded away from unity (on either side) by at least θ . This leads to the expression

$$D(\varrho' \parallel \varrho) = - \sum_{n \in S_{\pm}} \lambda'_n \log(1 + \Delta_n) = -N_{\pm} \sum_{n \in S_{\pm}} p_n \log(1 + \Delta_n), \quad (\text{D46})$$

where we renormalise the eigenvalues (which here correspond to a subnormalised probability distribution) by writing $\lambda'_n = N_{\pm} p_n$, with $N_{\pm} := \sum_{n \in S_{\pm}} \lambda'_n$ being the total population of the subspace S_{\pm} and $\{p_n\}$ here forming a probability distribution. Note that the ratio of eigenvalues going to unity in the S_0 subspace implies that the total populations of initial and final eigenvalues in this subspace are equal, i.e., $\sum_{n \in S_0} \lambda_n = \sum_{n \in S_0} \lambda'_n$, which in turn implies that the same is true for the S_{\pm} subspace, leading to $\sum_{n \in S_{\pm}} p_n \Delta_n = 0$.

We argue from the concavity of the logarithm function that

$$\frac{1}{2} \log(1 + \theta) + \frac{1}{2} \log(1 - \theta) \geq \sum_{n \in S_{\pm}} p_n \log(1 + \Delta_n). \quad (\text{D47})$$

Visualising the graph of the function $y = \log(1 + x)$, the latter expression above must evaluate to a point that lies within the intersection of the convex hull of $(\Delta_n, \log(1 + \Delta_n))$ and the linear equality $\sum_{n \in S_{\pm}} p_n \Delta_n = 0$, the latter of which is the line $x = 0$. By the concavity of the logarithm, the aforementioned convex hull lies entirely below the line segment connecting $(1 - \theta, \log(1 - \theta))$ to $(1 + \theta, \log(1 + \theta))$, and thus the expression is upper bounded by the intersection of this line segment with $x = 0$, which is precisely the l.h.s. of the inequality above. Thus we have the inequality

$$D(\varrho' \parallel \varrho) \geq -N_{\pm} \left[\frac{1}{2} \log(1 + \theta) + \frac{1}{2} \log(1 - \theta) \right] = -\frac{N_{\pm}}{2} \log(1 - \theta^2) \geq \frac{N_{\pm}}{2} \theta^2, \quad (\text{D48})$$

where we use $\log(1 - \theta^2) \leq -\theta^2$ for all $\theta \in [-1, 1]$. As $\theta > 0$, the only way that this contribution to the relative entropy by the eigenvalues that do not satisfy Eq. (D24) can be asymptotically negligible is if the total population of their associated subspace N_{\pm} goes to zero.

Finally note that, as mentioned in the main text, the above result pertains to the restricted setting where the target system is cooled as much as possible. However, this is not the only way to approach perfect cooling at the Landauer cost: instead of the largest half of global eigenvalues being placed into the ground-state subspace of the target system, any amount of them such that their sum is sufficiently close to one would suffice. Although it is complicated to derive an exact set of conditions that would

need to be satisfied in such cases (since it depends upon exactly which eigenvalues are permuted to which subspaces), the fact that fine-tuned control over particular degrees of freedom is required remains. Lastly, note that even in the restricted setting of cooling the target as much as possible, the situation becomes even more complicated when considering target systems that begin at a finite temperature. Here, the choice of which global eigenvalues should be permuted to which subspaces to cool the system as much as possible at minimal energy cost depends on the microscopic structure of both the system and machine. This means that one can no longer determine the final eigenvalue distributions of the reduced states in terms of the initial machine eigenvalues alone, as we were able to do for the maximally mixed state. In turn, one can no longer derive a condition on properties of the machine itself, independently of the target system. Nonetheless, again, the key message that cooling at minimal energy cost requires fine-tuned control to access precisely distributed populations still holds true. We leave the further exploration of such scenarios, for instance constructing optimal machines for particular initial target systems, to future work.

D5. Energy-Gap Variety as a Notion of Control Complexity

The insights drawn above regarding sufficient conditions for cooling a system at the Landauer limit lead us to propose a more nuanced notion of control complexity than the preliminary effective dimension that satisfies the natural desiderata outlined in the main text. In particular, here we demonstrate that the *energy-gap variety* (see Definition 2) provides a good measure of control complexity, both from a theoretical, thermodynamic standpoint as well as a practical one.

Firstly, it is quite clear that coupling a system to a diverging number of distinct machine energy gaps is a difficult task to achieve in almost any conceivable physical platform, especially when the energy gaps are closely spaced; thus, this definition indeed corresponds to our intuitive understanding of “complex” as an operation that is inherently difficult to perform in practice. Secondly, from all of the optimal cooling protocols that we outline in this paper, we see that, in contrast to the effective dimension, having a diverging energy-gap variety that densely covers an appropriate interval is sufficient for saturating the Landauer limit, thereby making it a better quantifier of control complexity. The remaining point is to show that its divergence is necessary to cool a system to the ground state using a single control operation with energy cost saturating the Landauer bound, so that it is fully consistent also with Nernst’s unattainability principle. We argue that this is indeed the case below by proving Theorem 4.

Proof: First of all, note that how cold the final system state can be made is bounded by the inequality:

$$\lambda_{\min}(\varrho'_S) \geq e^{-\beta \omega_{\mathcal{M}}^{\max}} \lambda_{\min}(\varrho_S), \quad (\text{D49})$$

where λ_{\min} denotes the minimal eigenvalue. For a pure final system state, the l.h.s. of the above equation goes to 0; thus, for any nontrivial initial system state [i.e., such that $\lambda_{\min}(\varrho_S) > 0$] and finite temperature $\beta < \infty$, we must have $\omega_{\mathcal{M}}^{\max} \rightarrow \infty$. This determines the upper limit of the required interval of energy gaps. The lower limit of the required interval comes from the fact that the only subspaces of the machine that are relevant for cooling the target system are those associated to energy gaps that are at least as large as the smallest energy gap of the target, ω_0 [38].

Next, recall the equality form of the Landauer limit, which holds true for *any* global unitary transformation with a thermal machine:

$$\beta \Delta E_{\mathcal{M}} = \tilde{\Delta} S_S + I(\mathcal{S} : \mathcal{M})_{\varrho'_{S\mathcal{M}}} + D(\varrho'_{\mathcal{M}} \| \varrho_{\mathcal{M}}), \quad (\text{D50})$$

Cooling the target system to a pure state necessitates that the final system and machine are uncorrelated and we therefore have $I(\mathcal{S} : \mathcal{M})_{\varrho'_{S\mathcal{M}}} = 0$ for the optimal process. We thus need to focus on minimising the relative-entropy term, which we do in the following steps.

Consider for simplicity the target system to be a qubit initially in the maximally mixed state. A generic cooling machine should be able to cool *any* system state, include the maximally mixed one; therefore the following insights pertaining to this special case apply generically. In this case, the initial joint spectrum of the system and machine is given by Eq. (D25). Cooling the target system to the ground state necessitates taking a subset \mathcal{A} of these global eigenvalues such that $\sum_{i \in \mathcal{A}} \lambda_{\mathcal{M}}^i = 1 - \epsilon$ for arbitrarily small ϵ and placing these into the ground-state subspace of the target, with the remaining (small) ϵ amount of population contributing only to any higher-energy eigenstates [this is essentially a generalisation of the conditions put forth in Eqs. (D23), accounting for an arbitrarily small cooling error].

As discussed previously, there are many possible ways to achieve such a configuration, but there is a *unique* one that minimises the total energy cost of doing so: namely, that in which the reduced final state of the machine is rendered passive. This is because if one compares two protocols achieving the same cooling for the target system, one in which the final machine is passive and any other in which it is not, then the former protocol has the smaller energy cost since a positive amount of energy can be (unitarily) extracted from the latter machine in order to render it passive.

Thus, for any protocol saturating the Landauer limit, the final machine state must be arbitrarily close to a passive state, which implies that it must be approximately diagonal in the local machine energy eigenbasis with the globally ordered spectrum as per Eq. (D26). Moreover, in order to minimise the relative-entropy term and therefore saturate the Landauer limit, the final machine

state must be arbitrarily close to the initial (i.e., thermal) machine state; following the argumentation put forth in the previous Appendix, this leads to the set of conditions outlined in Eq. (D29), which must be satisfied up to arbitrary precision.

Since we have the exact relationship between the initial and final machine eigenvalues, the contribution to the energy cost from the relative-entropy term can be calculated explicitly, i.e., $D(\varrho' \parallel \varrho) = \sum_n \lambda'_n \log \left(\frac{\lambda'_n}{\lambda_n} \right)$. Following the argumentation from Eq. (D46) until Eq. (D48) in the previous appendix, we see that by assuming a finite deviation from *any* of the conditions of Eq. (D29), i.e., writing $\lambda'_n = \lambda_n(1 + \Delta_n)$ with $|\Delta_n| \geq \theta > 0$ for some θ , one can derive a lower bound on the relative entropy:

$$D(\varrho' \parallel \varrho) \geq \frac{N_{\pm}}{2} \theta^2, \quad (\text{D51})$$

where N_{\pm} is the total population of the subspaces corresponding to the terms in the sum such that $\frac{\lambda'_n}{\lambda_n}$ differs from unity by at least θ . In other words, these are the relevant additional contributions to the energy cost; whenever N_{\pm} is nonzero, the Landauer bound *cannot* be approached arbitrarily closely.

The final piece is to relate the machine eigenvalues to its energy-gap spectrum, which can be done straightforwardly due to the initial thermality of the machine, i.e., $\lambda_{\mathcal{M}}^{\downarrow i} = e^{-\beta \omega_i} / \mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})$. We now argue that if there is ever a finite “jump” in the energy-gap structure of the machine, then one cannot achieve both a ground-state population of the target that is arbitrarily close to unity *and* have N_{\pm} be arbitrarily close to zero concurrently. Suppose now that one has a machine with a dense energy-gap structure from ω_0 up until some (finite) ω_a , followed by a finite jump until the energy level $\omega_a + \Omega$ (for some strictly finite $\Omega > 0$), and then again a dense set of energy gaps throughout the interval $[\omega_a + \Omega, \infty)$. Then, one can utilise the energy-gap structure in the “lower band” $[\omega_0, \omega_a)$ in an optimal fashion in order to cool the target system to a minimum temperature (set by ω_a) at Landauer energy cost [38, 42]. However, assuming that the jump in the energy-gap structure begins at some finite ω_a , then there is always a finite amount of population in the machine that is supported on the energy levels corresponding to the “upper band” $[\omega_a + \Omega, \infty)$. To cool the target to arbitrarily close to the ground state, one must therefore access this population and transfer it to the ground-state subspace of the target; this precisely corresponds to the N_{\pm} that contributes to the excess energy cost in a non-negligible manner for finite population exchanges. In particular, we have the bound $N_{\pm} \geq \min\left(\frac{e^{-\beta \omega_a}}{1 + e^{-\beta \omega_a}}, \frac{1}{1 + e^{-\beta \omega_a}}\right)$. Thus, whenever ω_a takes a finite value, N_{\pm} is a strictly positive number. The only way that the relative-entropy term can vanish then is if θ vanishes; however, this can occur only if $\Omega \rightarrow 0$, because for any finite Ω , the ratio $\frac{\lambda'_n}{\lambda_n}$ for at least one value of n differs from 1 by a finite amount as argued above, which finally leads to a nonzero lower bound in Eq. (D51) and implies that the Landauer limit cannot be saturated. In other words, the endpoints of the lower and upper energy gap intervals considered above must coincide (up to arbitrary precision) in order to saturate the Landauer bound. This implies that the energy-gap variety must diverge and moreover, since the above logic holds for arbitrary ω_a , which can be smoothly varied as a parameter, it follows that the diverging number of energy gaps must additionally approximately densely cover the interval in question.

Appendix E: Diverging Time and Diverging Control Complexity Cooling Protocols for Harmonic Oscillators

We now analyse the case of cooling infinite-dimensional quantum systems in detail. More specifically, we consider ensembles of harmonic oscillators. For the sake of completeness, we first briefly present some key concepts that will become relevant throughout this analysis. Following this, in Appendix E2a, we construct a protocol that achieves perfect cooling at the Landauer limit using a diverging number of Gaussian operations. Although such operations are typically considered to be relatively “simple” both when it comes to experimental implementation and theoretical description, according to the effective dimension notion of control complexity that we have shown must necessarily diverge to cool at the Landauer limit [see Eq. (6)], such Gaussian operations have infinite control complexity. Subsequently, in Appendix E2b, we consider the task of perfect cooling with diverging time but restricting the individual operations to be of finite control complexity. In particular, note that such operations are non-Gaussian in general. Here, we present a protocol that approaches perfect cooling of the target system as the number of operations diverges, with finite energy cost—albeit not at the Landauer limit. Whether or not a similar protocol exists that also saturates the Landauer bound remains an open question. Finally, in Appendix E3, we reconsider the protocol from Appendix E2b in terms of a single transformation, i.e., unit time. By explicitly constructing the joint unitary transformation that is applied throughout the entire protocol, we show this to be a multimode Gaussian operation acting on a diverging number of harmonic oscillators. The key message to be taken away from these protocols is that, while the distinction between Gaussian and non-Gaussian operations is a significant one in terms of experimental feasibility, and it certainly plays a role regarding the task of cooling—in particular, the energy cost incurred—these concepts alone cannot be used to characterise a notion of control complexity that must diverge to approach perfect cooling at the Landauer limit. On the other hand, the effective dimension of the machine used does precisely that; however, in a manner that is far from sufficient (for the case of harmonic oscillators), as even a single two-mode swap, which cannot cool perfectly at Landauer cost, would have infinite control complexity. Indeed, a more nuanced characterisation of control complexity in the infinite-dimensional setting, which takes more structure regarding the operations and energy levels into account, remains an open problem to be addressed.

E1. Preliminaries

We consider ensembles of N harmonic oscillators (i.e., infinite-dimensional systems consisting of N bosonic modes), which are associated to a tensor product Hilbert space $\mathcal{H}_{\text{tot}} = \bigotimes_{j=1}^N \mathcal{H}_j$ and (respectively: lowering, raising) mode operators $\{a_k, a_k^\dagger\}$ satisfying the bosonic commutation relations:

$$[a_k, a_{k'}^\dagger] = \delta_{kk'}, \quad [a_k, a_{k'}] = 0, \quad \forall k, k' = 1, 2, \dots, N. \quad (\text{E1})$$

The free Hamiltonian of any such system can be written as $H_{\text{tot}} = \sum_{k=1}^N \omega_k a_k^\dagger a_k$, where ω_k represents the energy gap of the k -th mode (in units where $\hbar = 1$). Position- and momentum-like operators for each mode can be defined as follows (for simplicity, we use the rescaled version below where the ω_k are omitted from the prefactors)

$$q_k := \frac{1}{\sqrt{2}}(a_k + a_k^\dagger), \quad p_k := \frac{1}{i\sqrt{2}}(a_k - a_k^\dagger). \quad (\text{E2})$$

As a consequence of the commutation relations in Eq. (E1), the generalised position and momentum operators satisfy the canonical commutation relations

$$[q_k, p_l] = i\delta_{kl}. \quad (\text{E3})$$

To simplify notation, one may further introduce the vector of quadrature operators $\mathbb{X} := (q_1, p_1, \dots, q_N, p_N)$; then, the commutation relations can be expressed succinctly as

$$[\mathbb{X}_k, \mathbb{X}_l] = i\Omega_{kl}, \quad (\text{E4})$$

where the Ω_{kl} are the components of the symplectic form

$$\Omega = \bigoplus_{j=1}^N \Omega_j, \quad \Omega_j = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}. \quad (\text{E5})$$

The density operator associated to N harmonic oscillators can be written in the so-called *phase-space representation* as

$$\varrho = \frac{1}{(2\pi)^N} \int \chi(\Omega\xi) \mathcal{W}(-\Omega\xi) d^{2N}\xi, \quad (\text{E6})$$

where $\mathcal{W}(\xi) := e^{i\xi^T \mathbb{X}}$ is the Weyl operator and $\chi(\xi) := \text{tr}[\varrho \mathcal{W}(\xi)]$ is called the characteristic function.

Throughout our analysis, we see that a particular class of states and operations, namely those that are known as *Gaussian*, are of particular importance. A Gaussian state is one for which the characteristic function is Gaussian

$$\chi(\xi) = e^{-\frac{1}{4}\xi^T \Gamma \xi + i\bar{\mathbb{X}}^T \xi}. \quad (\text{E7})$$

Here, $\bar{\mathbb{X}} := \langle \mathbb{X} \rangle_\varrho$ is the *displacement vector* or *vector of first moments*, and Γ is a real symmetric matrix that collects the *second statistical moments* of the quadratures, which is known as the *covariance matrix*. Its entries are given by

$$\Gamma_{mn} := \langle \mathbb{X}_m \mathbb{X}_n + \mathbb{X}_n \mathbb{X}_m \rangle_\varrho - 2\langle \mathbb{X}_n \rangle_\varrho \langle \mathbb{X}_m \rangle_\varrho. \quad (\text{E8})$$

We see that any Gaussian state is thus uniquely determined by its first and second moments. As an example of specific interest here, we recall that any thermal state τ of a harmonic oscillator with frequency ω is a Gaussian state and has vanishing first moments, $\bar{\mathbb{X}} = 0$. Here and throughout this article, we are assuming that the infinite-dimensional thermal state is well defined (see, e.g., Ref. [55] for discussion). The covariance matrix of a thermal state is proportional to the 2×2 identity, and given by $\Gamma[\tau(\beta, H)] = \coth\left(\frac{\beta\omega}{2}\right) \mathbb{I}_2$.

Gaussian operations are transformations that map the set of Gaussian states onto itself. Such operations, which include, e.g., beam-splitting and phase-space displacement, are generally considered to be relatively easily implementable in the laboratory. Although nonunitary Gaussian operations exist as well, all of the examples mentioned above are Gaussian unitaries. Such Gaussian unitaries are generated by Hamiltonians that are at most quadratic in the raising and lowering operators. Conversely, any Hamiltonian that can be expressed as a polynomial of at most second order in the mode operators generates a Gaussian unitary. Any unitary Gaussian transformation can be represented by an affine map (M, κ) ,

$$\mathbb{X} \mapsto M\mathbb{X} + \kappa, \quad (\text{E9})$$

where $\kappa \in \mathbb{R}^{2N}$ is a displacement vector in the phase-space representation and M is a symplectic $2N \times 2N$ matrix that leaves the symplectic form Ω invariant, i.e.,

$$M \Omega M^T = \Omega. \quad (\text{E10})$$

Under such a mapping, the first and second moments transform according to

$$\overline{\mathbb{X}} \mapsto M \overline{\mathbb{X}} + \kappa, \quad \Gamma \mapsto M \Gamma M^T. \quad (\text{E11})$$

Lastly, note that the energy of a Gaussian state ϱ_G with respect to its free Hamiltonian $H = \sum_k \omega_k a_k^\dagger a_k$ can be calculated in terms of the first and second moments as follows [22]

$$E(\varrho_G) = \sum_k \omega_k \left(\frac{1}{4} \text{tr} [\Gamma^{(k)} - 2] + \frac{1}{2} \|\overline{\mathbb{X}}^{(k)}\|^2 \right), \quad (\text{E12})$$

where $\|\cdot\|$ denotes the Euclidean norm. Here, $\Gamma^{(k)}$ indicates the (2×2) submatrix of the full covariance matrix Γ corresponding to the reduced state of the k^{th} mode. Similarly $\overline{\mathbb{X}}^{(k)}$ denotes the two-component subvector of first moments for the k^{th} mode of the displacement vector $\overline{\mathbb{X}}$.

E2. Diverging-Time Cooling Protocol for Harmonic Oscillators

E2a. Diverging-Time Protocol using Gaussian Operations (with Diverging Control Complexity)

We now consider a simple protocol for lowering the temperature of a single-mode system within the coherent-control paradigm using a single harmonic oscillator machine. This protocol will form the basic step of a protocol for achieving perfect cooling at the Landauer limit using diverging time, which we subsequently present.

In the situation we consider here, the target system \mathcal{S} to be cooled is a harmonic oscillator with frequency ω_s interacting with a harmonic oscillator machine \mathcal{M} at frequency $\omega_{\mathcal{M}} \geq \omega_s$ via a (non-energy-conserving) unitary acting on the joint system \mathcal{SM} initialised as a tensor product of thermal states $\tau_s(\beta, H_s) \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ at inverse temperature β with respect to their local Hamiltonians H_s and $H_{\mathcal{M}}$, respectively. The joint covariance matrix of the system and machine modes is block diagonal since the initial state is of product form, i.e.,

$$\Gamma[\tau_s(\beta, H_s) \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})] = \Gamma[\tau_s(\beta, H_s)] \oplus \Gamma[\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})], \quad (\text{E13})$$

and the 2×2 blocks of the individual modes are also diagonal, with the explicit expression $\Gamma[\tau_X(\beta, H_X)] = \coth\left(\frac{\beta\omega_X}{2}\right) \mathbb{1}_2$.

In this setting, it has been shown that the minimum reachable temperature of the target system is given by $T_{\min} = \frac{\omega_s}{\omega_{\mathcal{M}}} T$ (for the case $\omega_{\mathcal{M}} \geq \omega_s$) [38]. The non-energy-conserving unitary transformation that achieves this is of the form

$$U = e^{-i\frac{\pi}{2}(a^\dagger b + ab^\dagger)}, \quad (\text{E14})$$

where the operators a (a^\dagger) and b (b^\dagger) denote the annihilation (creation) operators of the target system and machine, respectively. This beam-splitter-like unitary acts as a SWAP with a relative phase factor imparted on the resultant state; nonetheless, this phase is irrelevant at the level of the covariance matrix, which fully characterises the (Gaussian) thermal states considered, and transforms it according to a standard swapping of the systems. After acting with such a SWAP operator, which is a Gaussian operation, the first moment remains vanishing and the covariance matrix transforms as [see Eq. (E11)]

$$\begin{bmatrix} \coth\left(\frac{\beta\omega_s}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta\omega_{\mathcal{M}}}{2}\right) \mathbb{1}_2 \end{bmatrix} \xrightarrow{\text{SWAP}} \begin{bmatrix} \coth\left(\frac{\beta\omega_{\mathcal{M}}}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta\omega_s}{2}\right) \mathbb{1}_2 \end{bmatrix}. \quad (\text{E15})$$

This means that both the output target system and machine are thermal states at different temperatures $T'_s = \frac{\omega_s}{\omega_{\mathcal{M}}} T$ and $T'_{\mathcal{M}} = \frac{\omega_{\mathcal{M}}}{\omega_s} T$. Making use of Eq. (E12), we can calculate the energy change for the system and machine as

$$\Delta E_s = E\left[\tau_s\left(\frac{\omega_{\mathcal{M}}}{\omega_s} \beta, H_s\right)\right] - E[\tau_s(\beta, H_s)] = \frac{\omega_s}{2} \left[\coth\left(\frac{\beta\omega_{\mathcal{M}}}{2}\right) - \coth\left(\frac{\beta\omega_s}{2}\right) \right],$$

$$\Delta E_{\mathcal{M}} = E \left[\tau_{\mathcal{M}} \left(\frac{\omega_s}{\omega_{\mathcal{M}}} \beta, H_{\mathcal{M}} \right) \right] - E [\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})] = \frac{\omega_{\mathcal{M}}}{2} \left[\coth \left(\frac{\beta \omega_s}{2} \right) - \coth \left(\frac{\beta \omega_{\mathcal{M}}}{2} \right) \right]. \quad (\text{E16})$$

The total energy cost associated to this SWAP operation is thus

$$\Delta E_{s,\mathcal{M}} = \Delta E_s + \Delta E_{\mathcal{M}} = \frac{(\omega_{\mathcal{M}} - \omega_s)}{2} \left[\coth \left(\frac{\beta \omega_s}{2} \right) - \coth \left(\frac{\beta \omega_{\mathcal{M}}}{2} \right) \right] = (\omega_{\mathcal{M}} - \omega_s) \frac{e^{-\beta \omega_s} (1 - e^{-\beta(\omega_{\mathcal{M}} - \omega_s)})}{(1 - e^{-\beta \omega_s})(1 - e^{-\beta \omega_{\mathcal{M}}})}. \quad (\text{E17})$$

Note that this form is similar to that for finite-dimensional systems with equally spaced Hamiltonian [cf., Eq. (C19)]; the dimension-dependent term vanishes as $d \rightarrow \infty$, simplifying the expression in the infinite-dimensional case.

With this simple protocol for lowering the temperature of a harmonic oscillator target using a single harmonic oscillator machine at hand, we are now in a position to describe an energy-optimal (in the sense of saturating the Landauer bound) cooling protocol when a diverging number of operations, i.e., diverging time, is permitted. In other words, we now show how to achieve perfect cooling with minimal energy at the expense of requiring diverging time, i.e., infinitely many steps of finite duration. As mentioned above, in this specific protocol, the control complexity as per Eq. (6) is infinite in each of these infinitely many steps. As we argue after having presented the protocol, this is an artefact of the simple structure of the Gaussian operations used. Indeed, we later present a non-Gaussian diverging-time protocol for cooling a single harmonic oscillator to the ground state using finite control complexity in each of the infinitely many steps, and at an overall finite (albeit not minimal, i.e., not at the Landauer limit) energy cost. Before presenting this non-Gaussian protocol, let us now discuss the details of the Gaussian diverging-time protocol for cooling at the Landauer limit.

We consider a harmonic oscillator with the frequency ω_s as the target system and the machine to comprise N harmonic oscillators, where the n^{th} oscillator has frequency $\omega_{\mathcal{M}_n} = \omega_s + n\epsilon$. In addition, we assume that all modes are initially uncorrelated and in thermal states at the same inverse temperature β with respect to their free Hamiltonians, i.e., the target system is $\tau_s(\beta, H_s)$ and the multimode thermal machine is $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \bigotimes_{n=1}^N \tau_{\mathcal{M}_n}(\beta, H_{\mathcal{M}_n})$.

In this case, the cooling process is divided into N time steps. During each step, there is an interaction between the target system and one of the harmonic oscillators in the machine. Here, we assume that at the n^{th} time step, the target system interacts only with the n^{th} harmonic oscillator, which has frequency $\omega_s + n\epsilon$. To obtain the minimum temperature for the target system, we perform the previously outlined cooling process at each step, which is given by swapping the corresponding two modes. Using Eq. (E15), the covariance matrix transformation of the two-mode process at the first time step takes the form

$$\Gamma^{(1)}(\tau_s(\beta) \otimes \tau_{\mathcal{M}_1}(\beta)) = \begin{bmatrix} \coth \left(\frac{\beta \omega_s}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta(\omega_s + \epsilon)}{2} \right) \mathbb{1}_2 \end{bmatrix} \xrightarrow{\text{SWAP}} \Gamma_{\text{opt}}^{(1)} = \begin{bmatrix} \coth \left(\frac{\beta(\omega_s + \epsilon)}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta \omega_s}{2} \right) \mathbb{1}_2 \end{bmatrix}. \quad (\text{E18})$$

By repeating this process on each of the harmonic oscillators in the machine, after the $(n-1)^{\text{th}}$ step, the 2×2 block corresponding to the target system S in the covariance matrix is given by $\coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \mathbb{1}_2$. Therefore, one can show inductively that the covariance matrix transformation associated to the n^{th} interaction is given by

$$\begin{aligned} \Gamma^{(n)}(\tau_s(\beta) \otimes \tau_{\mathcal{M}_n}(\beta)) &= \begin{bmatrix} \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) \mathbb{1}_2 \end{bmatrix} \\ &\xrightarrow{\text{SWAP}} \Gamma_{\text{opt}}^{(n)} = \begin{bmatrix} \coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \mathbb{1}_2 \end{bmatrix}. \end{aligned} \quad (\text{E19})$$

Based on this process, after N steps (i.e., after the system has interacted with all N harmonic oscillators), the minimal achievable temperature of the target system is $T_{\text{min}}^{(N)} = \frac{\omega_s}{\omega_s + N\epsilon} T$. Moreover, by using Eq. (E16), one can calculate the energy changes of the target system and the machine at each time step as

$$\begin{aligned} \Delta E_s^{(n)} &= \frac{\omega_s}{2} \left[\coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \right], \\ \Delta E_{\mathcal{M}_n}^{(n)} &= \frac{(\omega_s + n\epsilon)}{2} \left[\coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) \right]. \end{aligned} \quad (\text{E20})$$

The total energy change for the target system during the overall process (i.e., throughout the N steps) is thus given by

$$\Delta E_s = \sum_{n=1}^N \Delta E_s^{(n)} = \sum_{n=1}^N \frac{\omega_s}{2} \left[\coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \right]$$

$$= \frac{\omega_S}{2} \left[\coth \left(\frac{\beta(\omega_S + N\epsilon)}{2} \right) - \coth \left(\frac{\beta\omega_S}{2} \right) \right] = \omega_S \left[\frac{e^{-\beta(\omega_S + N\epsilon)}}{1 - e^{-\beta(\omega_S + N\epsilon)}} - \frac{e^{-\beta\omega_S}}{1 - e^{-\beta\omega_S}} \right]. \quad (\text{E21})$$

Here, we write $\coth(x) = 1 + (2e^{-2x})/(1 - e^{-2x})$. Similarly, one can obtain the total energy change of the overall machine

$$\begin{aligned} \Delta E_{\mathcal{M}} &= \sum_{n=1}^N \Delta E_{\mathcal{M}_n}^{(n)} = \sum_{n=1}^N \frac{\omega_S + n\epsilon}{2} \left[\coth \left(\frac{\beta(\omega_S + (n-1)\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_S + n\epsilon)}{2} \right) \right] \\ &= \sum_{n=1}^N (\omega_S + n\epsilon) \left[\frac{e^{-\beta(\omega_S + (n-1)\epsilon)}}{1 - e^{-\beta(\omega_S + (n-1)\epsilon)}} - \frac{e^{-\beta(\omega_S + n\epsilon)}}{1 - e^{-\beta(\omega_S + n\epsilon)}} \right]. \end{aligned} \quad (\text{E22})$$

It is straightforward to check that the total energy change, i.e., the sum of Eqs. (E21) and (E22), is equal to the energy cost obtained in Eq. (C19) with $d \rightarrow \infty$. In particular, this can be seen by considering the second line of Eq. (C19), where the second term in round parenthesis vanishes as $d \rightarrow \infty$ for any value of N . Thus, when the number of operations diverges $N \rightarrow \infty$ and $\epsilon = (\omega_{\max} - \omega_S)/N \rightarrow 0$, where $\omega_{\max} := \frac{\beta_{\max}}{\beta} \omega_S$ is the maximum frequency of the machines, the heat dissipated by the machines throughout the process saturates the Landauer bound and is therefore energetically optimal. Moreover, by taking $\omega_{\max} \rightarrow \infty$ one approaches perfect cooling.

At this point, a comment on the notion of control complexity is in order. According to Eq. (6), the effective dimension of the machine in the protocol we consider here diverges in addition to time. Indeed, the notion of control complexity thusly defined diverges for *any* Gaussian operation acting on the machine, in particular, it diverges for any single one of the infinitely many steps of the protocol, as each operation is a two-mode Gaussian operation. At first glance, this appears to be in contrast to the common conception that Gaussian operations are typically easily implementable (cf. Refs. [22, 37]). However, an alternative way of interpreting this protocol is that, exactly because of the simple structure of Gaussian operations, reaching the ground state at finite energy cost requires a diverging number of two-mode Gaussian unitaries, and thus divergingly many modes on which to act (see also Appendix E3). In fact, if non-Gaussian unitaries are employed, then the ground state can be approached at finite energy cost using just a single harmonic oscillator machine, as we now show.

E2b. Diverging-Time Protocol using Non-Gaussian Operations (with Finite Control Complexity)

We now consider a protocol for cooling a single harmonic oscillator at frequency ω_S to the ground state using a diverging amount of time, but requiring only a finite overall energy input as well as finite control complexity in each of the diverging number of steps of the protocol. In this protocol, the machine \mathcal{M} is also represented by a single harmonic oscillator whose frequency matches that of the target oscillator that is to be cooled, $\omega_{\mathcal{M}} = \omega_S =: \omega$. The initial states of both the target system \mathcal{S} and machine \mathcal{M} are assumed to be thermal at the same inverse temperature β , and are hence both described by thermal states of the form

$$\tau(\beta) = \frac{e^{-\beta H}}{\text{tr}[e^{-\beta H}]} = \sum_{n=0}^{\infty} e^{-\beta\omega n} (1 - e^{-\beta\omega}) |n\rangle\langle n| = \sum_{n=0}^{\infty} p_n |n\rangle\langle n|_{\mathcal{S}\mathcal{M}}, \quad (\text{E23})$$

where the Hamiltonian H is given by $H = \sum_{n=0}^{\infty} n\omega |n\rangle\langle n|$ and the $p_n = e^{-\beta\omega n} (1 - e^{-\beta\omega})$ are the eigenvalues of τ . The joint initial state is a product state that we can then write as

$$\tau_{\mathcal{S}}(\beta) \otimes \tau_{\mathcal{M}}(\beta) = \sum_{m,n=0}^{\infty} p_m p_n |m\rangle\langle m|_{\mathcal{S}} \otimes |n\rangle\langle n|_{\mathcal{M}} = \sum_{m,n=0}^{\infty} \tilde{p}_{m+n} |m, n\rangle\langle m, n|, \quad (\text{E24})$$

where we define $\tilde{p}_k := e^{-\beta\omega k} (1 - e^{-\beta\omega})^2$. We then note that the eigenvalues \tilde{p}_k of the joint initial state have degeneracy $k + 1$. For instance, the largest value $\tilde{p}_0 = p_0 p_0$, corresponding to both the system and machine being in the ground state, is the single largest eigenvalue, but there are two eigenstates, $|0, 1\rangle$ and $|1, 0\rangle$, corresponding to the second largest eigenvalue \tilde{p}_1 , three states, $|0, 2\rangle$, $|1, 1\rangle$, and $|2, 0\rangle$ for the third largest eigenvalue \tilde{p}_2 , and so forth. Obviously, not all of these eigenvalues correspond to eigenstates for which the target system is in the ground state.

In order to increase the ground-state population of the target system oscillator, we can now apply a sequence of ‘two-level’ unitaries, i.e., unitaries that act only on a subspace spanned by two particular eigenstates and exchange their respective populations. The two-dimensional subspaces are chosen such that one of the two eigenstates corresponds to the system \mathcal{S} being in the ground state, $|0, k\rangle$, while the other eigenstate corresponds to \mathcal{S} being in an excited state, $|i \neq 0, j\rangle$. In addition, these pairs of levels are selected such that, at the time the unitary operation is to be performed, the population of $|0, k\rangle$ is smaller than that of $|i \neq 0, j\rangle$, such that the two-level exchange increases the ground-state population of \mathcal{S} at each step.

More specifically, at the k^{th} step of this sequence, the joint system \mathcal{SM} is in the state $\varrho_{\mathcal{SM}}^{(k)}$ and one determines the set Ω_k of index pairs $(i \neq 0, j)$ such that $\tilde{p}_k < \langle i, j | \varrho_{\mathcal{SM}}^{(k)} | i, j \rangle$, i.e., the set of eigenstates for which \mathcal{S} is not in the ground state and which have a larger associated population (at the beginning of the k^{th} step) than $|0, k\rangle$. One then determines an index pair (m_k, n_k) for which this population is maximal, i.e., $\langle m_k, n_k | \varrho_{\mathcal{SM}}^{(k)} | m_k, n_k \rangle = \max\{\langle i, j | \varrho_{\mathcal{SM}}^{(k)} | i, j \rangle | (i, j) \in \Omega_k\}$, and performs the unitary

$$U_{\mathcal{SM}}^{(k)} = \mathbb{1}_{\mathcal{SM}} - |0, k\rangle\langle 0, k| - |m_k, n_k\rangle\langle m_k, n_k| + \left(|0, k\rangle\langle m_k, n_k| + |m_k, n_k\rangle\langle 0, k|\right). \quad (\text{E25})$$

If there is no larger population that is not already in the subspace of the ground state of the target system, i.e., when $\Omega_k = \emptyset$, which is only the case for the first step ($k = 1$), then no unitary is performed. After the k^{th} step, the joint state $\varrho_{\mathcal{SM}}^{(k+1)}$ is still diagonal in the energy eigenbasis, and the subspace of the joint Hilbert spaces for which \mathcal{S} is in the ground state is populated with the $k + 1$ largest eigenvalues \tilde{p}_i in nonincreasing order with respect to nondecreasing energy eigenvalues of the subspace's basis vectors $|0, i\rangle$. That is, for all $i \in \{0, 1, 2, \dots, k\}$ and for all $j \in \mathbb{N}$ with $j > i$, we have $\langle 0, i | \varrho_{\mathcal{SM}}^{(k+1)} | 0, i \rangle \geq \langle 0, j | \varrho_{\mathcal{SM}}^{(k+1)} | 0, j \rangle$.

Since the Hilbert spaces of both \mathcal{S} and \mathcal{M} are infinite dimensional, we can continue with such a sequence of two-level exchanges indefinitely, starting with $k = 1$ and continuing step by step as $k \rightarrow \infty$. Here we note that the choice of (m_k, n_k) is generally not unique at the k -th step, but as $k \rightarrow \infty$, the resulting final state is independent of the particular choices of (m_k, n_k) made along the way. In particular, in a fashion that is reminiscent of the famed Hilbert hotel paradox (see, e.g., Ref. [56, p. 17]), this sequence places *all* of the infinitely many eigenvalues \tilde{p}_k of the joint state of \mathcal{SM} (which must hence sum to one) into the subspace where \mathcal{S} is in the ground state. In other words, in the limit of infinitely many steps, the population of the ground-state subspace evaluates to

$$\sum_{k=0}^{\infty} (k+1) \tilde{p}_k = \sum_{k=0}^{\infty} (k+1) e^{-\beta\omega k} (1 - e^{-\beta\omega})^2 = 1, \quad (\text{E26})$$

where we take into account the $(k+1)$ -fold degeneracy of the k^{th} eigenvalue \tilde{p}_k . We thus have $\lim_{k \rightarrow \infty} \text{tr}_{\mathcal{M}} [\varrho_{\mathcal{SM}}^{(k)}] = |0\rangle\langle 0|_{\mathcal{S}}$, the reduced state of the system is asymptotically the pure state $|0\rangle_{\mathcal{S}}$.

As per our requirement on the structural complexity (see Appendix D), the Hilbert space of the machine required to achieve this is infinite-dimensional, and since each step of the protocol is assumed to take a finite amount of time, the overall time for reaching the ground state diverges. At the same time, the control complexity for each individual step is finite, since each U_k acts nontrivially only on a two-dimensional subspace. To see that also the energy cost for this protocol is finite, we first note that the protocol results in a final state of the machine that is diagonal in the energy eigenbasis $|n\rangle_{\mathcal{M}}$, with probability weights \tilde{p}_k decreasing (but not strictly) with increasing energy. Due to the degeneracy of the eigenvalues \tilde{p}_k , each one appears $(k+1)$ times on the diagonal (w.r.t. the energy eigenbasis) of the resulting machine state, populating adjacent energy levels. The label $n(k)$ of the lowest energy level that is populated by a particular value \tilde{p}_k can be calculated as

$$\tilde{n}(k) := \sum_{n=0}^{k-1} (n+1) = \frac{1}{2} k(k+1), \quad (\text{E27})$$

while the largest energy populated by \tilde{p}_k is given by $\tilde{n}(k+1) - 1$. With this, we calculate the energy of the machine after the protocol, which evaluates to

$$\frac{E_{\mathcal{M}}^{\text{final}}}{\omega} = \sum_{k=1}^{\infty} e^{-\beta\omega k} (1 - e^{-\beta\omega})^2 \sum_{n=\tilde{n}(k)}^{\tilde{n}(k+1)-1} n = \sum_{k=1}^{\infty} e^{-\beta\omega k} (1 - e^{-\beta\omega})^2 \frac{1}{2} k(k+1)(k+2) = \frac{3}{4} \text{cosech}^2\left(\frac{\beta\omega}{2}\right). \quad (\text{E28})$$

Since the energy of the initial thermal state is given by

$$\frac{E[\tau(\beta)]}{\omega} = \sum_{n=0}^{\infty} n e^{-\beta\omega n} (1 - e^{-\beta\omega}) = \frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}, \quad (\text{E29})$$

we thus arrive at the energy cost

$$\frac{\Delta E_{\mathcal{M}}}{\omega} = \frac{E_{\mathcal{M}}^{\text{final}} - E[\tau(\beta)]}{\omega} = \frac{e^{-\beta\omega} (2 + e^{-\beta\omega})}{(1 - e^{-\beta\omega})^2}. \quad (\text{E30})$$

We thus see that this energy cost is finite for all finite initial temperatures (although note that the energy cost diverges when $\beta \rightarrow 0$).

However, as we show next, the energy cost for attaining the ground state is not minimal, i.e., the protocol achieves perfect cooling (with finite energy and control complexity, but infinite time) but not at the Landauer limit. To see this, we first observe that the entropy of the final pure state of the system \mathcal{S} vanishes, such that $\tilde{\Delta}S_s = S[\tau(\beta)]$. Evaluating this entropy, one obtains

$$\begin{aligned} S[\tau(\beta)] &= -\text{tr}[\tau \log(\tau)] = -\sum_{n=0}^{\infty} e^{-\beta\omega n} (1 - e^{-\beta\omega}) \log[e^{-\beta\omega n} (1 - e^{-\beta\omega})] = -\sum_{n=0}^{\infty} e^{-\beta\omega n} (1 - e^{-\beta\omega}) [-\beta\omega n + \log(1 - e^{-\beta\omega})] \\ &= \frac{\beta\omega e^{-\beta\omega}}{1 - e^{-\beta\omega}} + \beta\omega + \log\left(\frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}\right) = \frac{\beta\omega}{1 - e^{-\beta\omega}} + \log\left(\frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}\right). \end{aligned} \quad (\text{E31})$$

Using the results from Eqs. (E30) and (E31), we can thus compare the expressions for $\beta\Delta E_{\mathcal{M}}$ and $\tilde{\Delta}S_s$, and we find that $\beta\Delta E_{\mathcal{M}} - \tilde{\Delta}S_s > 0$ for all nonzero initial temperatures. The origin of this difference is easily identified: although the protocol results in an uncorrelated final state because the system is left in a pure state, that is, $I(\mathcal{S} : \mathcal{M})_{\rho'_{\mathcal{SM}}} = 0$, the last term $D(\rho'_{\mathcal{SM}} \| \tau_{\mathcal{M}})$ in Eq. (3) is nonvanishing for nonzero temperatures because the protocol does not result in a thermal state of the machine.

With this, we thus show that perfect cooling is indeed possible using a finite energy cost and a finite control complexity in every one of infinitely many steps (thus using diverging time). As we have seen, the structural requirement of an infinite-dimensional effective machine Hilbert space can be met by realising \mathcal{M} as a single harmonic oscillator. Although the presented protocol does not minimise the energy cost to saturate the Landauer bound, we cannot at this point conclusively say that it is not possible to do so in this setting. However, we suspect that a more complicated energy-level structure of the machine is necessary.

Finally, let us comment again on the notion of control complexity in terms of effective machine dimension as opposed to the notion of complexity that is often (loosely) associated with the distinction between Gaussian and non-Gaussian operations. As we see from the protocols presented here, the concept of control complexity based on the nontrivially accessed Hilbert-space dimension of the machine indeed captures the resource that must diverge in order to reach the ground state, while the intuition of complexity associated with (non)-Gaussian operations, albeit valid as a characterisation of a certain practical difficulty in realising such operations, seems to be irrelevant for determining if the ground state can be reached. In the protocol presented in this section, non-Gaussian operations with finite control complexity are used in each step to reach the ground state. Infinitely many steps (i.e., diverging time) could then be traded for a single (also non-Gaussian) operation of infinite control complexity, performed in unit time. In the previous protocol based on Gaussian operations (Appendix E2a), the control complexity diverges in every single step of the cooling protocol, but only when there are infinitely many such steps (diverging time) or one operation in unit time on infinitely many modes (see below), can we reach the ground state. However, in the latter case, the operation, although acting on a diverging number of harmonic oscillators, remains Gaussian, as we now show explicitly.

E3. Diverging Control Complexity Cooling Protocol for Harmonic Oscillators

Here we give a protocol for perfectly cooling a harmonic oscillator in unit time and with the minimum energy cost, but with diverging control complexity. In accordance with Theorem 3, the machines used to cool the target system will likewise be harmonic oscillators. Let the operators a (a^\dagger) and b_k (b_k^\dagger), respectively, denote the annihilation (creation) operators of the target system and a machine subsystem labelled k . We then consider the unitary transformation in Eq. (E14), namely

$$U_k := e^{i\frac{\pi}{2}(a^\dagger b_k + a b_k^\dagger)}. \quad (\text{E32})$$

One can then apply the diverging-time cooling protocol from Appendix E2a to cool the system to the ground state at the Landauer limit via the total unitary transformation

$$U_{\text{tot}} := \lim_{N \rightarrow \infty} U_{(N)}, \quad \text{with} \quad U_{(N)} := \prod_{k=1}^N U_k. \quad (\text{E33})$$

We now seek the Hamiltonian that generates U_{tot} . First note that $U_{(N)} a U_{(N)}^\dagger = i b_1$ and

$$U_{(N)} b_k U_{(N)}^\dagger = \begin{cases} -b_{k+1}, & \text{for } k < N \\ i a, & \text{for } k = N \\ b_k, & \text{for } k > N \end{cases}, \quad (\text{E34})$$

which can be proven by induction. In contrast with Appendix E2a, here we use the complex representation of the symplectic group to describe the transformation, i.e., the set of matrices S satisfying $SKS^\dagger = K$, where $K := \mathbb{1}_N \oplus$

$(-\mathbb{1}_N)$. Gathering the raising and lowering operators of the target system and the first N machines into the vector $\vec{\xi} := (a \ b_1 \ b_2 \ \dots \ b_N \ a^\dagger \ b_1^\dagger \ b_2^\dagger \ \dots \ b_N^\dagger)^\top$, we can write the transformation above as $U_{(N)}\vec{\xi}U_{(N)}^\dagger = S^\top\vec{\xi}$ [57], where

$$S = \begin{pmatrix} \alpha_{(N)} & 0 \\ 0 & \alpha_{(N)} \end{pmatrix}, \quad \text{with} \quad \alpha_{(N)} := \begin{pmatrix} 0 & 0 & 0 & \dots & 0 & i \\ i & 0 & 0 & \dots & 0 & 0 \\ 0 & -1 & 0 & \dots & 0 & 0 \\ 0 & 0 & -1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -1 & 0 \end{pmatrix}. \quad (\text{E35})$$

Now, defining the matrix of Hamiltonian coefficients $h_{(N)}$ implicitly by $U_{(N)} =: \exp(-i\vec{\xi}^\dagger \cdot h_{(N)} \cdot \vec{\xi})$, we have that $S = \exp(-iKh_{(N)})$ [57], i.e., $h_{(N)} = iK \log(S^\top) = iK \log(S)^\top$, where we take the principal logarithm. To calculate this, we must diagonalise the matrix $\alpha_{(N)}$ in Eq. (E35). The eigenvalues of $\alpha_{(N)}$ are

$$\lambda_k := -e^{-i\pi \frac{2k-1}{N+1}}, \quad \text{with} \quad k \in \{1, 2, \dots, N+1\}, \quad (\text{E36})$$

i.e., the negative of the $(N+1)^{\text{th}}$ roots of -1 , and it is diagonalised by the unitary matrix V constructed from the eigenvectors \vec{v}_k :

$$V := \begin{pmatrix} \vec{v}_1 & \vec{v}_2 & \vec{v}_3 & \dots & \vec{v}_{N+1} \end{pmatrix} \quad \text{with} \quad \vec{v}_k := \frac{-1}{\sqrt{N+1}} \begin{pmatrix} i(-\lambda_k)^{-1} \\ (-\lambda_k)^{-2} \\ (-\lambda_k)^{-3} \\ \vdots \\ (-\lambda_k)^{-(N+1)} \end{pmatrix}. \quad (\text{E37})$$

Specifically, $\alpha_{(N)} = VDV^\dagger$, where $D := \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{N+1})$, and thus

$$h_{(N)}^\top = iK \log \begin{pmatrix} VDV^\dagger & 0 \\ 0 & VDV^\dagger \end{pmatrix} = iK \begin{pmatrix} V & 0 \\ 0 & W \end{pmatrix} \begin{pmatrix} \log(D) & 0 \\ 0 & \log(D) \end{pmatrix} \begin{pmatrix} V^\dagger & 0 \\ 0 & V^\dagger \end{pmatrix} =: \begin{pmatrix} A & 0 \\ 0 & -A \end{pmatrix} \quad (\text{E38})$$

for some matrix A . By direct calculation, one finds that

$$A_{jk} = i^{\delta_{j1}} i^{\delta_{k1}} \frac{\pi}{(N+1)^2} \sum_{p=1}^{N+1} (2p-2-N) e^{-i\pi \frac{2p-1}{N+1}(j-k)}. \quad (\text{E39})$$

Now, considering the identity

$$\sum_{p=1}^{N+1} e^{i\theta p} = \frac{e^{i\theta(N+1)} - 1}{1 - e^{i\theta}} \quad (\text{E40})$$

for $\theta \in \mathbb{R}$, as well as its derivative with respect to θ , one can calculate the sum in Eq. (E39). We then have

$$\lim_{N \rightarrow \infty} A_{jk} = \begin{cases} 0, & \text{for } j = k \\ ii^{\delta_{j1}} i^{\delta_{k1}} \frac{1}{j-k}, & \text{for } j \neq k \end{cases}. \quad (\text{E41})$$

Then, finally, we have that $U_{\text{tot}} = e^{-iH_{\text{tot}}}$, where $H_{\text{tot}} = \lim_{N \rightarrow \infty} (\vec{v}^\dagger \cdot h_{(N)} \cdot \vec{v})$, i.e.,

$$H_{\text{tot}} = - \sum_{j=2}^{\infty} \left(\frac{1}{j-1} b_j^\dagger a + \text{H.c.} \right) + \sum_{j,k=1; j \neq k}^{\infty} \frac{i}{j-k} b_j^\dagger b_k. \quad (\text{E42})$$

Thus, the system is cooled to the ground state at an energy cost saturating the Landauer bound, and in unit time, but via a procedure that implements a multimode Gaussian unitary on a diverging number of modes.

Appendix F: Cooling Protocols in the Incoherent-Control Paradigm

In this section, we investigate the required resources to cool the target system within the incoherent-control paradigm. For simplicity, we consider only the finite-dimensional setting. Here, we have a qudit target system \mathcal{S} interacting resonantly (i.e., in an energy-conserving manner) with a qudit machine \mathcal{M} , which is partitioned into one part, \mathcal{C} , in thermal contact with the ambient environment at inverse temperature β and another part, \mathcal{H} , in contact with a hot bath at inverse temperature $\beta_H < \beta$. The Hamiltonians for each subsystem are $H_{\mathcal{X}} = \sum_{n=0}^{d_{\mathcal{X}}-1} n \omega_{\mathcal{X}} |n\rangle\langle n|_{\mathcal{X}}$; the energy resonance condition enforces that $\omega_{\mathcal{H}} = \omega_{\mathcal{C}} - \omega_{\mathcal{S}}$. For the most part in this section, we focus on equally spaced Hamiltonians for simplicity; we comment specifically whenever we consider otherwise.

In order to cool the target system, we aim to compress as much population as possible into the its lowest energy eigenstates via interactions that are restricted to the energy-degenerate subspaces of the joint \mathcal{SCH} system. Thus we are restricted to global energy-conserving unitaries U_{EC} that satisfy

$$[H_{\mathcal{S}} + H_{\mathcal{C}} + H_{\mathcal{H}}, U_{\text{EC}}] = 0. \quad (\text{F1})$$

In Ref. [38], it was shown that for the case where all three subsystems are qubits, the optimal global unitary in this setting (inasmuch as they render the target system in the coldest state possible given the restrictions) is

$$U_{\text{EC}} = |0, 1, 0\rangle\langle 1, 0, 1|_{\mathcal{SCH}} + |1, 0, 1\rangle\langle 0, 1, 0|_{\mathcal{SCH}} + \bar{\mathbb{I}}, \quad (\text{F2})$$

where $\bar{\mathbb{I}}$ denotes the identity matrix on all subspaces that are not energy degenerate. Considering the generalisation to qudit subsystems, it is straightforward to see that, for equally spaced Hamiltonians, the optimal global unitaries must be of the form

$$U_{\text{EC}} = \left[\sum_{m,n,l=0}^{d-2} |m, n+1, l\rangle\langle m+1, n, l+1|_{\mathcal{SCH}} + |m+1, n, l+1\rangle\langle m, n+1, l|_{\mathcal{SCH}} \right] + \bar{\mathbb{I}}. \quad (\text{F3})$$

For the most general case where the Hamiltonians of each subsystem are arbitrary, it is not possible to write down a generic form of the optimal unitary, since the energy-resonant transitions that lead to cooling the target now depend on the microscopic details of the energetic structure. Nonetheless, in Appendix G, we provide a protocol (i.e., not the unitary *per se*, but a sequence of steps) in this setting that attains perfect cooling and saturates the Carnot-Landauer limit.

Intuitively, the above types of unitaries simply reshuffle populations that are accessible through resonant transitions. For the purpose of cooling, one wishes to do this in such a way that the largest population is placed in the lowest energy eigenstate of the target system, the second largest in the second lowest energy eigenstate, and so on (in line with the optimal unitaries in the coherent-control setting); indeed, on the energy-degenerate subspaces accessible, such unitaries act precisely in this way. It is straightforward to show that interactions of this form satisfy Eq. (F1).

For the sake of simplicity, we now focus on the case where all systems are qubits, although the results generalise to the qudit setting. Consider the initial joint state $\varrho_{\mathcal{SCH}} = \sum_{m,n,l=0}^1 p_{mnl} |m, n, l\rangle\langle m, n, l|_{\mathcal{SCH}}$. By applying a unitary U_{EC} of the form given in Eq. (F3), the post-transformation joint state is

$$\varrho'_{\mathcal{SCH}} = U_{\text{EC}} \varrho_{\mathcal{SCH}} U_{\text{EC}}^\dagger = \varrho_{\mathcal{SCH}} + \Delta p |0, 1, 0\rangle\langle 0, 1, 0|_{\mathcal{SCH}} - \Delta p |1, 0, 1\rangle\langle 1, 0, 1|_{\mathcal{SCH}}, \quad (\text{F4})$$

where $\Delta p := p_{101} - p_{010}$ indicates the amount of population that has been transferred from the excited state of the target system to the ground state throughout the interaction. Naturally, in order to cool the target system, $\Delta p \geq 0$, i.e., the initial population p_{101} must be at least as large as p_{010} .

Due to the energy-conserving nature of the global interaction, the energy exchanged between the subsystems throughout a single such interaction, $\Delta E_{\mathcal{X}} = \text{tr} [H_{\mathcal{X}}(\varrho'_{\mathcal{X}} - \varrho_{\mathcal{X}})]$, can be calculated via

$$\Delta E_{\mathcal{S}} = -\omega_{\mathcal{S}} \Delta p, \quad \Delta E_{\mathcal{C}} = \omega_{\mathcal{C}} \Delta p, \quad \Delta E_{\mathcal{H}} = -\omega_{\mathcal{H}} \Delta p. \quad (\text{F5})$$

Thus, for a fixed energy-level structure of all subsystems (i.e., given the local Hamiltonians), one requires only knowledge of the pre- and post-transformation state of any one of the subsystems to calculate the energy change for all of them.

F1. Diverging Energy: Proof of Theorem 6

The first thing to note is that in the incoherent-control paradigm, even when one allows for the energy cost, i.e., the heat drawn from the hot bath, to be diverging, it is not possible to perfectly cool the target system, as presented in Theorem 6. The intuition behind this result is that the target system can interact only with energy-degenerate *subspaces* of the hot and cold machine subsystems. The optimal transformation that one can do here to achieve cooling is to transfer the highest populations

of any such subspace to the lowest energy eigenstate of the target system; however, any such subspace has population strictly less than one for any $0 \leq \beta_H \leq \beta < \infty$ independently of the energy structure. Moreover, the difference from one can be bounded by a finite amount that does not vanish independent of the energy-level structure of any machine of finite dimension. This makes it impossible to attain a subspace population of one even as the energy cost diverges for any fixed and finite control complexity. It follows that the ground-state population of the target system can never reach unity in a single operation of finite control complexity and hence perfect cooling cannot be achieved.

Precisely, we show the following. Let \mathcal{S} be a finite-dimensional system of dimension d_S with associated Hamiltonian with finite but otherwise arbitrary energy gaps $H_S = \sum_{i=0}^{d_S-1} \omega_S^i |i\rangle\langle i|_S$, and let d_C and d_H be integers denoting the dimensions of the cold and hot parts of the machine respectively. Then it is impossible to cool the system \mathcal{S} in the incoherent-control paradigm, i.e., using energy-conserving unitaries involving \mathcal{C} and \mathcal{H} at some initial inverse temperatures β, β_H respectively, arbitrarily close to the ground state. Note that, in particular, this result holds irrespective of the energy-level structure of \mathcal{C} and \mathcal{H} and no matter how much energy is drawn from the hot bath as a resource.

In order to set notation for the following, we assume $\omega_X^i \geq \omega_X^j$ for $i \geq j$ and $\omega_X^0 = 0$, where ω_X^i denotes the i^{th} energy eigenvalue of system \mathcal{X} with $\mathcal{X} \in \{\mathcal{S}, \mathcal{C}, \mathcal{H}\}$. We also assume the initial states on \mathcal{S} and \mathcal{C} to be thermal at inverse temperature β , and \mathcal{H} is assumed to be initially in a thermal state at inverse temperature $\beta_H \leq \beta$. We denote by p_X^i the i^{th} population of system \mathcal{X} in a given state, i.e., $p_X^i = \langle i | \rho_X | i \rangle$, where $|i\rangle$ denotes the i^{th} energy eigenstate of ρ_X . We also write $p_{ijk} := p_S^i p_C^j p_H^k$.

The intuition behind the proof is as follows. The global ground-state level of the joint hot-and-cold machine has some nonzero initial population for any finite-dimensional machine; in particular it can always be lower bounded by $\frac{1}{d_C d_H}$ for any Hamiltonians and initial temperatures, which is strictly greater than zero as long as the dimensions remain finite. Fixing the control complexity of any protocol considered here to be finite in value thus implies a lower bound on the initial ground-state population of the total machine that is larger than zero by a finite amount. Depending on the energy-level structure of the hot and cold parts of the machine, there may be other nonzero initial populations, but in order to cool the target system \mathcal{S} perfectly, at least all of the previously mentioned populations must be transferred into spaces spanned by energy eigenstates of the form $|0jk\rangle_{SC\mathcal{H}}$. This intuition is formalised via Lemma 2, where we show that independent of the energy structure of \mathcal{C} and \mathcal{H} , one must be able to make such transfers of population in order to perfectly cool \mathcal{S} . However, in order to make such transfers in an energy-conserving manner, all energy eigenstates of the form $|i00\rangle_{SC\mathcal{H}}$ must be degenerate with some of the form $|0jk\rangle_{SC\mathcal{H}}$. This degeneracy condition, in turn, also implies that every energy eigenstate of the form $|0jk\rangle_{SC\mathcal{H}}$ has an associated initial population p_{0jk} that is nonvanishing for all machines of finite dimension (i.e., for all protocols with finite control complexity). Thus, upon transferring some population p_{i00} into the subspace spanned by $|0jk\rangle_{SC\mathcal{H}}$, i.e., one of a relevant form for the population to contribute to the final ground-state population of the target, one inevitably transfers some finite amount of population away from the relevant space and into $|i00\rangle_{SC\mathcal{H}}$, which does not contribute to the final ground-state population of the target. We formalise this intuition in the discussion following Lemma 2. In this way, no matter what one does, there is always a finite amount of population, which is lower bounded by some strictly positive number due to the constraint on control complexity, that does not contribute to the final ground-state population of the target, implying that perfect cooling is not possible.

The formal proof occurs in two steps. We first show that some specific degeneracies in the joint $SC\mathcal{H}$ system must be present in order to be able to even potentially cool \mathcal{S} arbitrarily close to the ground state. We then prove that, given the above degeneracies, one cannot cool the system \mathcal{S} beyond a fixed ground-state population that is independent of the energy structure of \mathcal{C} and \mathcal{H} ; in particular, one can draw as much energy from the hot bath as they like and still do no better. We begin with the following lemma.

Lemma 2. *Given \mathcal{S} , d_C , and d_H as above, one can reach a final ground-state population of the system \mathcal{S} arbitrarily close to one in the incoherent-control setting only if each $|i00\rangle_{SC\mathcal{H}}$, where $i \in \{1, \dots, d_S - 1\}$, energy eigenstate is degenerate with at least one $|0jk\rangle_{SC\mathcal{H}}$ energy eigenstate, where $j \in \{0, \dots, d_C - 1\}, k \in \{0, \dots, d_H - 1\}$.*

Proof. Suppose that there exists an $i^* \in \{1, \dots, d_S - 1\}$ such that $|i^*00\rangle_{SC\mathcal{H}}$ is not degenerate with any $|0jk\rangle_{SC\mathcal{H}}$, where $j \in \{0, \dots, d_C - 1\}, k \in \{0, \dots, d_H - 1\}$. We show that, then, one cannot cool \mathcal{S} arbitrarily close to zero.

Let B_i denote the degenerate subspace of the total Hamiltonian $H_S + H_C + H_H$, where H_X denotes the Hamiltonian of system $\mathcal{X} \in \{\mathcal{S}, \mathcal{C}, \mathcal{H}\}$, that contains the eigenvector $|i00\rangle_{SC\mathcal{H}}$. Then, any energy-conserving unitary U_{EC} used to cool the system in the incoherent-control paradigm must act within such B_i subspaces, i.e., $U_{EC} = \bigoplus_i U_{B_i}$ (this is a direct consequence of $[U_{EC}, H_S + H_C + H_H] = 0$, see, e.g., Lemma 5 of Ref. [30]). This means, in particular, that the initial population of $|i^*00\rangle_{SC\mathcal{H}}$ can only be distributed within B_{i^*} , and as no eigenvector of the form $|0jk\rangle_{SC\mathcal{H}}$ is contained in B_{i^*} by assumption, it can never contribute to the final ground-state population of \mathcal{S} , which we denote \tilde{p}_S^0 . So we have

$$\tilde{p}_S^0 \leq 1 - p_{i^*00}. \quad (\text{F6})$$

Now, as for $\mathcal{X} \in \{\mathcal{C}, \mathcal{H}\}$, with any $\{\omega_X^i\}$ such that each $\omega_X^i \geq 0$ with $\omega_X^0 = 0$ and any inverse temperature $\beta \geq 0$, we have for the partition function \mathcal{Z}_S that

$$\mathcal{Z}_X = 1 + e^{-\beta\omega_X^1} + \dots + e^{-\beta\omega_X^{d_X-1}} \leq d_X, \quad (\text{F7})$$

and so we have the following bound on the initial populations associated to each eigenvector $|i00\rangle_{SC\mathcal{H}}$

$$p_{i00} = \frac{e^{-\beta\omega_S^i}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-\beta\omega_S^i}}{\mathcal{Z}_S d_C d_\mathcal{H}} > 0 \quad \forall i \in \{1, \dots, d_S - 1\}. \quad (\text{F8})$$

Combining the above, we have that

$$\hat{p}_S^0 \leq 1 - \frac{e^{-\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}} < 1. \quad (\text{F9})$$

So as desired, we show that one cannot cool beyond $1 - \frac{e^{-\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}}$, a bound strictly smaller than 1 for any finite-dimensional machine (i.e., for any protocol using only finite control complexity) and independent of the energies of \mathcal{C} and \mathcal{H} . \square

We can now proceed to the second step of the proof of Theorem 6.

Proof. To this end, consider any $i^* \in \{1, \dots, d_S - 1\}$. If $|i^*00\rangle_{SC\mathcal{H}}$ is not degenerate with any $|0jk\rangle_{SC\mathcal{H}}$, our assertion is proven by Lemma 2. On the other hand, if there is a $j^* \in \{0, \dots, d_C - 1\}$ and a $k^* \in \{0, \dots, d_\mathcal{H} - 1\}$ for which $|i^*00\rangle_{SC\mathcal{H}}$ and $|0j^*k^*\rangle_{SC\mathcal{H}}$ are degenerate, then B_{i^*} , the degenerate subspace containing $|i^*00\rangle_{SC\mathcal{H}}$, also contains $|0j^*k^*\rangle$. Now B_{i^*} may also contain other eigenvectors of the form $|0jk\rangle_{SC\mathcal{H}}$, i.e., some other $|0j'k'\rangle_{SC\mathcal{H}}$ with $j' \in \{0, \dots, d_C - 1\}$, $k' \in \{0, \dots, d_\mathcal{H} - 1\}$. Crucially, each such eigenvector in B_{i^*} must have an associated minimal amount of initial population as long as the machine is finite dimensional. Indeed, for any such $|0j^*k^*\rangle_{SC\mathcal{H}}$ in B_{i^*} , we have the condition $\omega_C^{j^*} + \omega_\mathcal{H}^{k^*} = \omega_S^{i^*}$ and so $\omega_C^{j^*} \leq \omega_S^{i^*}$, $\omega_\mathcal{H}^{k^*} \leq \omega_S^{i^*}$, implying that $\beta\omega_C^{j^*} \leq \beta\omega_S^{i^*}$ and $\beta\omega_\mathcal{H}^{k^*} \leq \beta\omega_S^{i^*}$. Thus we have the bound

$$p_{0j^*k^*} = \frac{e^{-\beta\omega_C^{j^*}} e^{-\beta\omega_\mathcal{H}^{k^*}}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}} =: q_{i^*}. \quad (\text{F10})$$

Now, take any particular $i^* \in \{1, \dots, d_S - 1\}$ and let π_{i^*} be the dimension of B_{i^*} , μ the number of energy eigenstates of the form $|0jk\rangle_{SC\mathcal{H}}$ that B_{i^*} contains and $\nu = \pi_{i^*} - \mu$ the number of energy eigenstates of the form $|ijk\rangle_{SC\mathcal{H}}$, where $i \neq 0$, that B_{i^*} contains. So

$$B_{i^*} = \text{span}\{|0jk\rangle, |0j_2k_2\rangle, \dots, |0j_\mu k_\mu\rangle, |i^*00\rangle, |i_2\ell_2 m_2\rangle, \dots, |i_\nu \ell_\nu m_\nu\rangle\}. \quad (\text{F11})$$

Let $\mathbf{v} = \{p_{0jk}, p_{0j_2k_2}, \dots, p_{0j_\mu k_\mu}, p_{i^*00}, p_{i_2\ell_2 m_2}, \dots, p_{i_\nu \ell_\nu m_\nu}\}$ be the vector of initial populations associated to the eigenvectors of B_{i^*} , and \mathbf{v}^\uparrow be the vector whose components are those of \mathbf{v} arranged in nondecreasing order. Using Schur's theorem [54], we know that after applying any unitary transformation $U_{B_{i^*}}$ on the relevant energy-degenerate subspace, then the vector of transformed populations, $\tilde{\mathbf{v}}$, is majorised by \mathbf{v} . In particular, labelling the vector elements by \mathbf{v}_α , we have

$$\tilde{p}_{i^*00} + \sum_{\alpha=2}^{\nu} \tilde{p}_{i_\alpha \ell_\alpha m_\alpha} \geq \sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow. \quad (\text{F12})$$

We now claim that $\sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow \geq q_{i^*}$ from Eq. (F10). Indeed, as \mathbf{v} has at most $\nu - 1$ elements that do not belong to the set $A := \{p_{0jk}, p_{0j_2k_2}, \dots, p_{0j_\mu k_\mu}, p_{i^*00}\}$, at least one element of A must contribute to the sum $\sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow$. Let x be that element. As $\mathbf{v}_\alpha^\uparrow \geq 0$ for all $\alpha = 1, \dots, \pi = \mu + \nu$, we have

$$\sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow \geq x. \quad (\text{F13})$$

Now as $p_{0j_\gamma k_\gamma} \geq q_{i^*}$ for all $\gamma = 2, \dots, \mu$, we have

$$x \geq \min(q_{i^*}, p_{i^*00}) = q_{i^*}, \quad (\text{F14})$$

where $p_{i^*00} \geq q_{i^*}$ can be seen from Eq. (F10), as claimed.

As the l.h.s. of Eq. (F12) represents the amount of population in the subspace B_{i^*} that does *not* contribute to the final ground-state population of the target system, we have

$$\hat{p}_S^0 \leq 1 - \left(\tilde{p}_{i^*00} + \sum_{\alpha=2}^{\nu} \tilde{p}_{i_\alpha \ell_\alpha m_\alpha} \right) \leq 1 - q_{i^*} = 1 - \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}}. \quad (\text{F15})$$

So, for any finite-dimensional machine, one cannot cool the system \mathcal{S} beyond $1 - \frac{e^{-\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}}$, a bound strictly smaller than 1 and independent of the energy structure of \mathcal{C} and \mathcal{H} , as desired. \square

As a concrete example, consider the case where all systems are qubits. The initial joint state is

$$\varrho_{sc\mathcal{H}}^{(0)} = \frac{(|0\rangle\langle 0| + e^{-\beta\omega_S}|1\rangle\langle 1|)_S \otimes (|0\rangle\langle 0| + e^{-\beta\omega_C}|1\rangle\langle 1|)_C \otimes (|0\rangle\langle 0| + e^{-\beta_H\omega_H}|1\rangle\langle 1|)_H}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_H(\beta_H, \omega_H)}. \quad (\text{F16})$$

The only energy-conserving unitary interaction that is relevant for cooling is the one that exchanges the populations in the levels spanned by $|010\rangle$ and $|101\rangle$, which have initial populations $\frac{e^{-\beta\omega_C}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_H(\beta_H, \omega_H)}$ and $\frac{e^{-\beta\omega_S}e^{-\beta_H\omega_H}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_H(\beta_H, \omega_H)}$ respectively, which are both strictly less than one. The necessary condition for any cooling to be possible implies that $e^{-\beta\omega_S}e^{-\beta_H\omega_H} \geq e^{-\beta\omega_C}$; now, performing the optimal cooling unitary leads to the final ground-state population of the target system

$$p'_s(0) = \langle 0 | \text{tr}_{c\mathcal{H}} [U\varrho_{sc\mathcal{H}}^{(0)}U^\dagger] | 0 \rangle_s = \frac{1 + e^{-\beta_H\omega_H}(1 + e^{-\beta\omega_S} + e^{-\beta\omega_C})}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_H(\beta_H, \omega_H)} < 1. \quad (\text{F17})$$

Indeed, using $e^{-\beta\omega_S}e^{-\beta_H\omega_H} \geq e^{-\beta\omega_C}$,

$$p'_s(0) \leq \frac{1 + e^{-\beta_H\omega_H}e^{-\beta\omega_S}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)} \leq \frac{1}{\mathcal{Z}_C(\beta, \omega_C)} \leq 1. \quad (\text{F18})$$

The second inequality is strict unless $\beta_H = 0$ or $\omega_H = 0$. In the both cases, for equality in the first inequality, we need $\beta\omega_S = \beta\omega_C$. If $\beta = 0$, then $\mathcal{Z}_C(\beta, \omega_C) = 2$ and the last inequality is strict. If $\omega_S = \omega_C$, no cooling is possible; hence $p'_s(0) = p_s(0) < 1$.

F2. Diverging Time and Diverging Control Complexity

We now move to analyse the case where diverging time is allowed, where we wish to minimise the energy cost and control complexity throughout the protocol over a diverging number of energy-conserving interactions between the target system and the hot and cold subsystems of the machine. We again consider all three systems to be qubits, but the results generalise to arbitrary (finite) dimensions. Here, the machines and ancillas begin as thermal states with initial inverse temperatures β and $\beta_H \leq \beta$ respectively. Just as in the diverging time cooling protocol in the coherent-control setting presented in Appendix C, we consider a diverging number of machines, with slightly increasing energy gaps, in a configuration such that the target system interacts with the n^{th} machine at time step n . Suppose that after n steps of the protocol, the target qubit has been cooled to some inverse temperature $\beta_n > \beta$; equivalently, this can be expressed as a thermal state with corresponding energy gap $\omega_n = \frac{\beta_n}{\beta}\omega_S$. We now wish to interact the target system $\tau_S(\beta_n, \omega_S)$ with a machine \mathcal{M}_{n+1} with slightly increased energy gaps with respect to the most recent one \mathcal{M}_n , i.e., we increase the energy gaps of the cold subsystem \mathcal{C} from ω_n to $\omega_{n+1} = \omega_n + \epsilon_n$; the resonance condition enforces the energy gap of the hot subsystem \mathcal{H} to be similarly increased to $\omega_n + \epsilon_n - \omega_S$. Thus, the next step of the protocol is a unitary acting on the global state

$$\varrho_{sc\mathcal{H}}^{(n)} = \tau_S(\beta_n, \omega_S) \otimes \tau_C(\beta, \omega_n + \epsilon_n) \otimes \tau_H(\beta_H, \omega_n + \epsilon_n - \omega_S). \quad (\text{F19})$$

In order to cool the target system via said unitary, we must have that $p_{101} \geq p_{010}$ for the state in Eq. (F19), which implies that ϵ_n must satisfy the following condition:

$$e^{-\beta\omega_n - \beta_H(\omega_n + \epsilon_n - \omega_S)} \geq e^{-\beta(\omega_n + \epsilon_n)} \Rightarrow \epsilon_n \geq \gamma(\omega_n - \omega_S) \quad \text{where} \quad \gamma := \frac{\beta_H}{\beta - \beta_H}. \quad (\text{F20})$$

This condition is crucial. It means that if the hot subsystem \mathcal{H} is coupled to a heat bath at any finite temperature, i.e., $\beta_H > 0$, ϵ_n depends linearly on the inverse temperature of the target system at the previous step β_n , and can thus not be taken to be arbitrarily small. As we now show, this condition prohibits the ability to perfectly cool the target system at the Landauer limit for the energy cost whenever the heat bath is at finite temperature.

On the other hand, for infinite-temperature heat baths, perfect cooling at the Landauer limit is seemingly achievable; here, $\beta_H \rightarrow 0$ and so $\gamma \rightarrow 0$, leading to the trivial constraint $\epsilon_n \geq 0$ which allows it to be arbitrarily small, as is required. Nonetheless, the explicit construction of any protocol doing so in the incoherent-control setting is *a priori* unclear, as the restriction of energy conservation makes for a fundamentally different setting from the coherent-control paradigm. We now explicitly derive the optimal diverging-time protocol to perfectly cool at the Landauer limit for an infinite-temperature heat bath, thereby proving Theorem 7.

F3. Saturating the Landauer Limit with an Infinite-Temperature Heat Bath

Before calculating the energy cost, we briefly discuss the attainability of the optimally cool target state. We begin with all subsystems as qubits, for the sake of simplicity, but the logic generalises to higher dimensions. In the incoherent paradigm, the target system \mathcal{S} interacts with a virtual qubit of the total machine $\mathcal{M} = \mathcal{CH}$ that consists of the energy eigenstates $|0, 1\rangle_{\mathcal{CH}}$ and $|1, 0\rangle_{\mathcal{CH}}$, with populations $p_{0_{\mathcal{C}}1_{\mathcal{H}}}$ and $p_{1_{\mathcal{C}}0_{\mathcal{H}}}$ respectively. Suppose that at step $n + 1$ the cold subsystem involved in the interaction has energy gap $\omega_n + \epsilon_n$. In Ref. [38], it is shown that by repeating the incoherent cooling process (i.e., implementing the unitary in Eq. (F3)) and taking the limit of infinite cycles, this scenario equivalently corresponds to the general (coherent) setting where arbitrary unitaries are permitted and the target system interacts with a virtual qubit machine with effective energy gap ω_n^{eff} given by

$$e^{-\beta\omega_n^{\text{eff}}} := \frac{p_{1_{\mathcal{C}}0_{\mathcal{H}}}}{p_{0_{\mathcal{C}}1_{\mathcal{H}}}} = e^{-\beta(\omega_n + \epsilon_n)} e^{\beta_H(\omega_n + \epsilon_n - \omega_S)} \Rightarrow \omega_n^{\text{eff}} = \omega_n + \epsilon_n - \frac{\beta_H}{\beta}(\omega_n + \epsilon_n - \omega_S). \quad (\text{F21})$$

It is clear that for finite-temperature heat baths, i.e., $\beta_H > 0$, the effective energy gap ω_n^{eff} is always smaller than the energy gap of the machine at any given step, i.e., $\omega_n^{\text{eff}} \leq \omega_n + \epsilon_n$; on the other hand, equality holds iff the heat bath is at infinite temperature, i.e., $\beta_H \rightarrow 0$. Thus, in the infinite-temperature case, given a target system beginning at some step of the protocol in the state $\varrho_S^*(\beta, \omega_n)$, it is possible to get close to the asymptotic state $\varrho_S^*(\beta, \omega_n + \epsilon_n)$; if the temperature is finite, however, this state is not attainable (even asymptotically). Following the arguments in Appendix C, i.e., considering a diverging number of machines, each of which having the part connected to the cold bath with energy gap $\omega_{\mathcal{C}_n} = \omega_n + \epsilon_n$ and taking the limit of $\epsilon_n \rightarrow 0$, which one can *only* do if the hot-bath temperature is infinite, allows one to cool perfectly in diverging time in the incoherent paradigm at the Landauer limit.

We now calculate the energy cost explicitly for the infinite-temperature heat bath case, precisely demonstrating attainability of the Landauer limit. We use a similar approach to that described in Appendix C: we have a diverging number of cold machines for each energy gap ω_n , with which the target system at the $n - 1^{\text{th}}$ time step interacts; for an infinite-temperature heat bath, i.e., \mathcal{H} is in the maximally mixed state independent of its energy structure, the state of the target system at each step $\varrho_S^*(\beta, \omega_{n-1})$ is achievable. From Eq. (F5), the energy change between all subsystems for a given step of the protocol, i.e., taking $\varrho_S^*(\beta, \omega_{n-1}) \rightarrow \varrho_S^*(\beta, \omega_n)$, can be calculated as

$$\begin{aligned} \Delta E_S^{(n)} &= \text{tr} [H_S(\omega_S)(\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1}))] \\ \Delta E_C^{(n)} &= -\text{tr} [H_C(\omega_n)(\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1}))] \\ \Delta E_{\mathcal{H}}^{(n)} &= \text{tr} [H_{\mathcal{H}}(\omega_n - \omega_S)(\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1}))] \end{aligned} \quad (\text{F22})$$

In general, i.e., for finite-temperature heat baths, we would have $\omega_n = \omega_{n-1} + \epsilon_{n-1}$, with a lower bound on ϵ_{n-1} for cooling to be possible [in accordance with Eq. (F20)]. However, for infinite-temperature heat baths, this lower bound trivialises since the energy structure of the hot-machine subsystem plays no role in its state; thus we can choose the energy gap structure for the machines as $\{\omega_n = \omega_S + n\epsilon\}_{n=1}^N$ with ϵ arbitrarily small. Taking the limit $\epsilon \rightarrow 0$, the diverging time limit $N \rightarrow \infty$, and writing $\omega_N = \omega_{\text{max}}$ for the maximum energy gap of the cold-machine subsystems, the energy exchanged throughout the entire cooling protocol here is given by

$$\begin{aligned} \Delta E_S &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_S^{(n)} = \text{tr} [H_S(\omega_S)(\varrho_S^*(\beta, \omega_{\text{max}}) - \varrho_S^*(\beta, \omega_S))] \\ \Delta E_C &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_C^{(n)} = \frac{1}{\beta} \{S[\varrho_S^*(\beta, \omega_S)] - S[\varrho_S^*(\beta, \omega_{\text{max}})]\} = \frac{1}{\beta} \tilde{\Delta} S_S \\ \Delta E_{\mathcal{H}} &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_{\mathcal{H}}^{(n)} = -\Delta E_S - \Delta E_C. \end{aligned} \quad (\text{F23})$$

Here, the expression for ΔE_C can be derived using the same arguments as presented in Appendix C1. In particular, the heat dissipated by the cold part of the machine, which is naturally connected to the heat sink in the incoherent setting as an infinite-temperature heat bath can be considered a work source since any energy drawn comes with no entropy change, is in accordance with the Landauer limit. It is straightforward to obtain the same result for qudit systems. Lastly, in a similar way to the other protocols we have presented, one could compress all of the diverging number of operations into a single one whose control complexity diverges, thereby trading off between time and control complexity.

F4. Analysis of Finite-Temperature Heat Baths

We now return to the more general consideration of finite-temperature heat baths, i.e., $0 < \beta_H \leq \beta$. In the case where $\beta_H = \beta$, from Eq. (F21), it is straightforward to see that for any machine energy gap ω_n , the effective gap ω_n^{eff} is equal to the gap of the target system, which means that no cooling can be achieved in the incoherent paradigm. Nonetheless, for any \mathcal{H} subsystem coupled to a heat bath of inverse temperature $\beta_H < \beta$, cooling is possible. We first provide more detail regarding why cooling at the Landauer limit is not possible in this setting, before deriving the minimal energy cost in accordance with the Carnot-Landauer limit presented in Theorem 5; in Appendix G, we provide explicit protocols that saturate this bound for any finite-temperature heat bath and arbitrary finite-dimensional systems and machines.

Suppose that at some step n one has the initial joint state of Eq. (F19), where $\epsilon_n = \gamma(\omega_n - \omega_S) + \epsilon$ and $\omega_n = \omega_S + n\epsilon$. Here, γ is as in Eq. (F20). We now wish to cool the target system to $\varrho_S^*(\beta, \omega_n + \epsilon)$. For cooling to be possible in the incoherent setting here, we need the cold-machine subsystem to have an energy gap of at least $\omega_n + \epsilon_n$; moreover, with a finite-temperature heat bath, this energy gap is insufficient to achieve the desired transformation [see Eq. (F20)]. Based on Eq. (F5), we can see that nonetheless, if we calculate the *hypothetical* energy change in this scenario if it were possible, we can derive a lower bound for the actual energy cost incurred. Employing Eq. (F22), we have

$$\begin{aligned} \Delta E_c^{(n+1)} &\geq -\text{tr}\{H_c(\omega_n + \epsilon_n)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= -\text{tr}\{H_c[(\gamma + 1)\omega_n - \gamma\omega_S + \epsilon][\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= -\text{tr}\{H_c[(\gamma + 1)\omega_n - \gamma\omega_S + \epsilon + \gamma\epsilon - \gamma\epsilon][\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= -(\gamma + 1)\text{tr}\{H_c(\omega_n + \epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} + \gamma\text{tr}\{H_c(\omega_S + \epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= (\gamma + 1)\Delta E_c^{*(n+1)} + \gamma\Delta E_S^{*(n+1)} + \gamma\text{tr}\{H_c(\epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\}, \end{aligned} \quad (\text{F24})$$

where we make use of the fact that for equally spaced Hamiltonians, the structure of the Hamiltonians on each subsystem take the same form [i.e., we can write, with slight abuse of notation, $H_c(\omega + \omega_S) = H_c(\omega) + H_S(\omega_S)$]. We use the star in $\Delta E_{\mathcal{A}}^*$ to denote the idealised energy cost [i.e., that corresponding to what would be achievable in the infinite-temperature setting; see Eq. (F22)] and the energy costs without the star to represent those for when the temperature of the heat bath is finite. The additional term $\text{tr}\{H(\gamma\epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\}$ vanishes for $\epsilon \rightarrow 0$.

Summing up these contributions for a diverging number of steps gives the lower bound for the heat dissipated throughout the entire protocol for cooling an initial state $\tau_S(\beta, \omega_S)$ to some final $\tau_S(\beta_{\text{max}}, \omega_S)$ is given by

$$\begin{aligned} \Delta E_c &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_c^{(n+1)} \\ &\geq (\gamma + 1) \frac{1}{\beta} \tilde{\Delta} S_S + \gamma \Delta E_S \\ &= \frac{1}{\beta} \tilde{\Delta} S_S + \gamma \left(\Delta E_S + \frac{1}{\beta} \tilde{\Delta} S_S \right). \end{aligned} \quad (\text{F25})$$

Note that for infinite-temperature heat baths, $\gamma \rightarrow 0$ and the usual Landauer limit is recovered; nonetheless, for finite-temperature heat baths, $\gamma > 0$ and there is an additional energy contribution, implying that the Landauer limit cannot be achieved. Moreover, note that the expression inside the parenthesis in the second term above is always non-negative, as it is the free energy difference of the system during the cooling process. Lastly, it is straightforward to show that this lower bound is equivalent to the Carnot-Landauer limit in Eq. (A14), which was derived in a protocol-independent manner as the ultimate limit in the incoherent-control setting. We now present explicit protocols that saturate this bound.

Appendix G: Perfect Cooling at the Carnot-Landauer Limit in the Incoherent-Control Paradigm

The precise statement that we wish to prove regarding saturation of the Carnot-Landauer limit is the following:

Lemma 3. *For any $\beta^* \geq \beta > \beta_H$ and $\epsilon_{1,2} > 0$, there exists a cooling protocol in the incoherent-control setting comprising a number of unitaries of finite control complexity, which, when the number of operations diverges, cools to some final temperature β' that is arbitrarily close to the ideal temperature value β^* , i.e.,*

$$|\beta' - \beta^*| < \epsilon_1, \quad (\text{G1})$$

with an energy cost, measured as heat drawn from the hot bath, that is arbitrarily close to the ideal Carnot-Landauer limit, i.e.,

$$\left| \Delta E_{\mathcal{H}} - \eta^{-1} \tilde{\Delta} F_S^{(\beta)} \right| < \epsilon_2, \quad (\text{G2})$$

where $\eta = 1 - \beta_H/\beta$ and $\Delta F_S^{(\beta)} = F_\beta(\varrho'_S) - F_\beta(\varrho_S)$ is the free energy difference between the initial $\varrho_S = \tau_S(\beta, H_S)$ and final $\varrho'_S = \tau_S(\beta^*, H_S)$ system states (w.r.t. inverse temperature β).

We begin by presenting the diverging-time protocol that saturates the Carnot-Landauer limit when all three subsystems $\mathcal{S}, \mathcal{C}, \mathcal{H}$ are qubits. The simplicity of this special case allows us to calculate precisely bounds on the number of operations required to reach any chosen error threshold. Building on this intuition, we then present the generalisation to the case where all systems are qudits. The protocols with diverging control complexity follow directly via the same line of reasoning presented in the main text.

G1. Qubit Case

We begin with setting some notation and intuition for the proof, before expanding on mathematical details.

Sketch of Protocol.—The protocol consists of the following. There are N stages, each labelled by $n \in \{1, 2, \dots, N\}$. Each stage proceeds as follows:

- A qubit with energy gap $\omega_S + n\theta$ is taken from the cold part of the machine, and a qubit with energy gap $n\theta$ is taken from the hot part (see below). The initial state of the machine at the beginning of the n^{th} stage is thus $\tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta)$.
- The energy-preserving three qubit unitary cycle in the $\{010, 101\}_{\mathcal{SCH}}$ subspace is performed [see Eq. (F3)], after which the cold and hot qubits are rethermalised to their respective initial temperatures.
- The above steps are repeated m_n times.

The energy increment θ is defined as

$$\theta := \frac{\omega_S}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right), \quad (\text{G3})$$

while the number of repetitions within each stage is given by

$$m_n = \left\lceil \frac{\log(\delta)}{\log(1 - N_V^{(n)})} \right\rceil. \quad (\text{G4})$$

$\lceil \cdot \rceil$ is the ceiling function, and $N_V^{(n)}$ is the sum of the initial thermal populations in the $\{01, 10\}_{\mathcal{CH}}$ subspace of the machine, i.e.,

$$N_V^{(n)} := \langle 01 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 01 \rangle + \langle 10 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 10 \rangle. \quad (\text{G5})$$

The parameter δ is chosen appropriately to complete the proof ($\delta = 1/N^2$ works).

The intuition for the proof is as follows. We first consider how the populations of the target system changes in the *idealised* protocol where $m_n \rightarrow \infty$, so that in each stage, the system reaches the virtual temperature determined by the \mathcal{CH} qubits. We can use this ideal setting to find expressions for the final temperature and energy cost, which serves as a baseline that we wish to attain to within arbitrary precision. We then consider the protocol as constructed above with a finite number of repetitions m_n in each stage, and show that its expressions for temperature and work cost are close (w.r.t. $1/N$) to the original expressions, and by taking N to be sufficiently large but still finite (i.e., in the diverging time limit), we prove that the protocol can be arbitrarily close in temperature and energy cost to the ideal values.

Proof. We label the population in the *excited* state of the target system at the end of stage n as p_n . Thus p_0 is the initial population and p_N is the final population in the excited level of the target system qubit, i.e., that spanned by $|1\rangle\langle 1|_{\mathcal{S}}$. We also label by q_n what the corresponding population p_n would hypothetically be in the limit $m_n \rightarrow \infty$. This value can be calculated by matching the temperature of the target system qubit to the temperature of the $\{01, 10\}_{\mathcal{CH}}$ virtual qubit within the machine (see Appendix G in Ref. [38]). Thus q_n is defined via the Gibbs ratio

$$\frac{q_n}{1 - q_n} = e^{-\beta(\omega_S + n\theta)} e^{+\beta_H n\theta} = e^{-\beta\omega_S} e^{-(\beta - \beta_H)n\theta}. \quad (\text{G6})$$

Note that

1. $\{p_n\}, \{q_n\}$ are both monotonically decreasing sequences, as each stage cools the target qubit further.
2. $p_n > q_n$ for all n , as more repetitions within each stage keep cooling the target qubit further.

To keep track of the energetic resource cost, which we take here to be the total heat drawn from the hot bath, we must sum the energetic contribution from each time the hot qubit is rethermalised to β_H after the application of the three-party cycle unitary. Due to the fact that the only manner in which the population of the hot qubit changes is due to the $\{010, 101\}_{SC\mathcal{H}}$ exchange, it follows that any population change in the hot qubit is identical to the population change in the target system qubit.

Focusing on a single stage, where the machine qubits are fixed in energy gap, the total population change in the hot qubit that must be restored by the hot bath is therefore equal to the population change in the target system throughout that stage. The heat drawn from the hot bath throughout the entire stage is therefore

$$\tilde{\Delta}E_{\mathcal{H}}^{(n)} = \omega_{\mathcal{H}}^{(n)}(p_{n-1} - p_n) = n\theta(p_{n-1} - p_n). \quad (\text{G7})$$

With these expressions derived, we can study the properties of the abstract protocol where the number of repetitions within each stage goes to infinity: $m_n \rightarrow \infty$. First, the final temperature asymptotically achieved here is given by finding the temperature $\tilde{\beta}$ associated with the qubit with excited-state population q_N

$$\frac{q_N}{1 - q_N} = e^{-\tilde{\beta}\omega_S} \Rightarrow e^{-\beta\omega_S} e^{-(\beta - \beta_H)N\theta} = e^{-\tilde{\beta}\omega_S} \Rightarrow \tilde{\beta} = \beta^*, \quad (\text{G8})$$

where we make use of the definition of θ in Eq. (G3). We can thus identify $q_N = q^*$, since it is the population associated with the ideal final temperature β^* .

We also have the following expression for the total energetic cost of the ideal protocol after N stages

$$\tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N n\theta(q_{n-1} - q_n), \quad (\text{G9})$$

which can alternatively be expressed as

$$\tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N [(n-1)\theta(q_{n-1} - q_n)] + \theta(q_0 - q_N) \quad (\text{G10})$$

The sums appearing in the two alternative expressions are the left and right Riemann sums of the integral of the variable $y = n\theta$ integrated with respect to the variable q , i.e.,

$$I := - \int_{q_0}^{q^*} y \, dq, \quad \text{where} \quad \frac{q(y)}{1 - q(y)} = e^{-\beta\omega_S} e^{-(\beta - \beta_H)y}, \quad (\text{G11})$$

from Eq. (G6). For $y > 0$, $q(y)$ is monotonically decreasing and so the converse is also true, i.e., y is monotonically decreasing w.r.t. $q(y)$. This implies that the integral is bounded by the left and right Riemann sums, so we have

$$\sum_{n=1}^N (n-1)\theta(q_{n-1} - q_n) \leq \int_{q_0}^{q^*} y \, dq \leq \sum_{n=1}^N n\theta(q_{n-1} - q_n), \quad (\text{G12})$$

from which we can deduce that the value of $\Delta E_{\mathcal{H}}^*$ is itself is bounded both ways from Eqs. (G9) and (G10):

$$\int_{q_0}^{q^*} y \, dq \leq \tilde{\Delta}E_{\mathcal{H}}^* \leq \int_{q_0}^{q^*} y \, dq + \theta(q_0 - q^*). \quad (\text{G13})$$

The integral itself can be expressed in terms of the free energy of the qubit target system with respect to the environment inverse temperature β . Expressing the free energy as a function of the excited-state population q and differentiating w.r.t. q gives

$$F(q) = \langle E \rangle(q) - \frac{S(q)}{\beta} = q\omega_S + \frac{1}{\beta} [q \log(q) + (1-q) \log(1-q)]. \quad (\text{G14})$$

$$\frac{\partial F}{\partial q} = \omega_S + \frac{1}{\beta} \log\left(\frac{q}{1-q}\right) = \left(\omega_S + \frac{1}{\beta} (-\beta\omega_S - (\beta - \beta_H)y)\right) = -\frac{\beta - \beta_H}{\beta} y. \quad (\text{G15})$$

Using the above expression, the definite integral in Eq. (G11) amounts to

$$I = \frac{1}{\eta} [F(q^*) - F(q_0)] =: \frac{1}{\eta} (F^* - F_0), \quad (\text{G16})$$

where we identify the Carnot efficiency $\eta = 1 - \beta_H/\beta$ and for ease of notation written $F^* := F(q^*)$ and $F_0 := F(q_0)$. Thus we can bound $\tilde{\Delta}E_{\mathcal{H}}^*$ on both sides

$$\frac{1}{\eta} (F^* - F_0) \leq \tilde{\Delta}E_{\mathcal{H}}^* \leq \frac{1}{\eta} (F^* - F_0) + \theta(q_0 - q^*) \leq \frac{1}{\eta} (F^* - F_0) + \frac{\omega_S}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right), \quad (\text{G17})$$

where the inequality in the second line follows from the fact that $\{q_n\}$ forms a decreasing sequence.

We now proceed to consider the cooling protocol with a finite number of repetitions m_n within each stage. We first bound the difference between p_n and q_n . Using the properties of the exchange unitary under repetitions [38, 58] (in particular, see Appendix G in Ref. [38]), we have that in each stage

$$\frac{p_n - q_n}{p_{n-1} - q_n} = \left(1 - N_V^{(n)}\right)^{m_n}. \quad (\text{G18})$$

Thus, the population difference to the asymptotically achievable population given by the virtual temperature shrinks as a power law w.r.t. the number of repetitions. Since $0 < N_V^{(n)} < 1$ (all strict inequalities), three points follow: first, the population q_n can never be attained with a finite number of steps within the stage n ; second, that every repetition cools the system further by some finite amount; third, that one can get arbitrarily close to q_n by taking m_n sufficiently large. In fact, by our definition of m_n , we have that

$$\frac{p_n - q_n}{p_{n-1} - q_n} \leq \delta. \quad (\text{G19})$$

From this, we can prove that

$$p_n - q_n \leq \delta^n q_0 - \delta q_n + (1 - \delta) \delta \sum_{j=1}^{n-1} \delta^{n-j-1} q_j. \quad (\text{G20})$$

The proof is by induction. For $n = 0$, $p_0 = q_0$ (initial state), and for $n = 1$, using Eq. (G19)

$$\begin{aligned} p_1 - q_1 &\leq \delta(p_0 - q_1) \\ &= \delta(q_0 - q_1). \end{aligned} \quad (\text{G21})$$

Suppose that the above statement holds true for p_k . Then from Eq. (G19)

$$\begin{aligned} p_{k+1} - q_{k+1} &\leq \delta(p_k - q_{k+1}) \\ &= \delta(p_k - q_k + q_k - q_{k+1}) \\ &\vdots \\ &\leq \delta^{k+1} q_0 - \delta q_{k+1} (1 - \delta) \delta + \sum_{j=1}^{(k+1)-1} \delta^{(k+1)-j-1} q_j. \end{aligned} \quad (\text{G22})$$

With this result, we can now bound the difference between the energy cost of this finite-repetition protocol and that of the idealised one. We now proceed to prove that

$$\tilde{\Delta}E_{\mathcal{H}} - \tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N n\theta(p_{n-1} - p_n) - \sum_{n=1}^N n\theta(q_{n-1} - q_n) \leq \theta \left(q_0 \sum_{j=1}^{N-1} \delta^{N-j} - \sum_{j=1}^{N-1} \delta^{N-j} q_j \right). \quad (\text{G23})$$

We again use proof by induction. First note that we can rewrite

$$\sum_{n=1}^N n\theta(f_{n-1} - f_n) = \theta \left(\sum_{n=1}^N f_{n-1} \right) - N\theta f_N, \quad (\text{G24})$$

for $f_n \in \{p_n, q_n\}$. Therefore, we can rewrite the difference

$$\tilde{\Delta}E_{\mathcal{H}} - \tilde{\Delta}E_{\mathcal{H}}^* = \theta \sum_{n=1}^N (p_{n-1} - q_{n-1}) - N\theta(p_N - q_N) \leq \theta \left(\sum_{n=1}^N (p_{n-1} - q_{n-1}) \right), \quad (\text{G25})$$

since the last subtracted term is always strictly positive. Consider now the partial sum

$$\mathcal{E}_k = \sum_{n=1}^k (p_{n-1} - q_{n-1}). \quad (\text{G26})$$

For $k = 1$, $\mathcal{E}_1 = 0$, since $p_0 = q_0$. For $k = 2$, we have

$$\mathcal{E}_1 = (p_1 - q_1) \leq \delta(q_0 - q_1) = \left(q_0 \sum_{j=1}^1 \delta^{2-j} - \sum_{j=1}^1 \delta^{2-j} q_j \right), \quad (\text{G27})$$

which matches the hypothesis of Eq. (G23). Assuming that the same holds true for \mathcal{E}_k , then for \mathcal{E}_{k+1} , we have

$$\begin{aligned} \mathcal{E}_{k+1} &= \mathcal{E}_k + (p_k - q_k) \\ &\leq \left(q_0 \sum_{j=1}^{k-1} \delta^{k-j} - \sum_{j=1}^{k-1} \delta^{k-j} q_j \right) + \left(\delta^k q_0 + (1 - \delta) \delta \sum_{j=1}^{k-1} \delta^{k-j-1} q_j - \delta q_k \right) \\ &\vdots \\ &= q_0 \sum_{j=1}^k \delta^{k+1-j} - \sum_{j=1}^k \delta^{k+1-j} q_j. \end{aligned} \quad (\text{G28})$$

Then, by dropping the second sum, which is a strictly positive quantity, the difference in Eq. (G23) can be further simplified to

$$\tilde{\Delta} E_{\mathcal{H}} - \tilde{\Delta} E_{\mathcal{H}}^* \leq \theta q_0 \sum_{j=1}^{N-1} \delta^{N-j} = \theta q_0 \delta \sum_{k=0}^{N-2} \delta^k < \theta q_0 \delta (N-1) < \theta q_0 \delta N < \omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \delta, \quad (\text{G29})$$

where we use that $\delta < 1$. Finally, to upper bound the number of operations required in the protocol, we bound the number of repetitions within each stage by bounding the total population of the virtual qubit spanned by the levels $\{01, 10\}_{\mathcal{CH}}$ as follows:

$$\begin{aligned} N_V^{(n)} &= \langle 01 | \tau_c(\beta, \omega_S + n\theta) \otimes \tau_{\mathcal{H}}(\beta_H, n\theta) | 01 \rangle + \langle 10 | \tau_c(\beta, \omega_S + n\theta) \otimes \tau_{\mathcal{H}}(\beta_H, n\theta) | 10 \rangle \\ &= \frac{e^{-\beta_H n\theta} + e^{-\beta(\omega_S + n\theta)}}{(1 + e^{-\beta_H n\theta})(1 + e^{-\beta(\omega_S + n\theta)})} \\ &> \frac{e^{-\beta(\omega_S + n\theta)}}{4}. \\ \Rightarrow \log \left[1 - N_V^{(n)} \right] &< \log \left[1 - \frac{e^{-\beta(\omega_S + n\theta)}}{4} \right] \\ &< -\frac{e^{-\beta(\omega_S + n\theta)}}{4} \quad \text{if } x \in (0, 1) \Rightarrow \log(1 - x) < -x. \\ \Rightarrow -\frac{1}{\log \left[1 - N_V^{(n)} \right]} &< 4e^{+\beta(\omega_S + n\theta)} \end{aligned} \quad (\text{G30})$$

Thus we can bound the number of repetitions in each stage from Eq. (G4). Noting that $\log(\delta) < 0$, we have

$$m_n < 4 \log(1/\delta) e^{+\beta(\omega_S + n\theta)} + 1. \quad (\text{G31})$$

For a crude bound, we can replace n by its maximum value N , and sum over all the stages to find an upper bound on the total number of three-qubit exchange unitaries implemented throughout the entire protocol, which gives

$$M = \sum_{n=1}^N m_n < N \left[4 \log(1/\delta) e^{+\beta(\omega_S + N\theta)} + 1 \right] = N \left[4 \log(1/\delta) e^{\omega_S(\beta^* - \beta_H)/\eta} + 1 \right]. \quad (\text{G32})$$

Also, note that $\lim_{\delta \rightarrow 0} p_N = q_N = q^*$. More precisely, using Eq. (G20), we have

$$p_N - q^* < \delta \left(\delta^{N-1} q_0 + (1 - \delta) \sum_{j=1}^{N-1} \delta^{N-j-1} q_j - q_N \right)$$

$$< \delta (1 + (1 - \delta)(N - 1)) < \delta N. \quad (\text{G33})$$

In summary, we have the following bounds on the protocol in which each stage consists of a finite number of steps

$$\begin{aligned} p_N - q^* &< \delta N \\ \tilde{\Delta} E_{\mathcal{H}} &< \frac{1}{\eta} (F^* - F_0) + \omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \left(\frac{1}{N} + \delta \right), \end{aligned} \quad (\text{G34})$$

where we combine Eqs. (G17) and (G29) for the second expression. For simplicity, we choose $\delta = 1/N^2$, so that

$$\begin{aligned} p_N - q^* &< \frac{1}{N} \\ \tilde{\Delta} E_{\mathcal{H}} &< \frac{1}{\eta} (F^* - F_0) + \omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \left(\frac{2}{N} \right). \end{aligned} \quad (\text{G35})$$

Thus, given any final temperature (encoded by the population q^*), and allowed errors ϵ_1 and ϵ_2 for the final population and energy cost respectively, one can always choose N large enough so that both quantities are within the error threshold. Specifically, choosing N as

$$N = \left\lceil \max \left\{ \epsilon_1^{-1}, 2\omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \epsilon_2^{-1} \right\} \right\rceil, \quad (\text{G36})$$

we automatically have that $p_N - q^* < \epsilon_1$ and $\Delta E_{\mathcal{H}} < (F^* - F_0)/\eta + \epsilon_2$. The total number of unitary operations (each of which is followed by rethermalisation of the machine) is then bounded by Eq. (G32)

$$M < N \left(8 \log[N] e^{\omega_S(\beta^* - \beta_H)/\eta} + 1 \right). \quad (\text{G37})$$

We can see from Theorem 10 that the protocol is asymptotically optimal with respect to the energy extracted from the hot bath. \square

G2. Qudit Case

The extension of the proof above to the case of qudits is nontrivial. This is because, while for qubits there is only one energy-resonant subspace that leads to cooling and hence a unique protocol [see Eq. (F3)] that asymptotically attains perfect cooling at the Carnot-Landauer bound, this is no longer the case for higher-dimensional systems; here, there can be a number of energy-resonant subspaces that cool the target and the question of optimality hinges crucially on the complex energy-level structure of all systems involved. Hence, it is not possible to provide a unique unitary that generates the optimal protocol independently of the subsystem Hamiltonians. Nonetheless, we slightly modify the protocol for the qubit case above to be implemented on a number of particular three-qubit subspaces of the three-qudit global state such that, at the end of each stage, the state of the target system is arbitrarily close to the (known) state, which would be achieved in an abstract protocol in the diverging-time limit. This asymptotically attainable state is precisely that which would be achieved in the coherent-control paradigm with a machine the same dimension as the joint hot-cold qudits. Thus, we first begin by presenting the necessary steps for the proof in the coherent-control setting, which we then adapt as appropriate for the incoherent setting control. Finally, summing the overall energy cost of said protocol over all stages saturates the Carnot-Landauer bound, as required.

Proof. An idealised sequence of temperatures and system states. We construct the incoherent protocol in the following manner. We seek to take the system through a sequence of thermal states starting at inverse temperature β and ending at inverse temperature β^* with N equally spaced intermediary steps, i.e.,

$$\beta_n = \beta + n\theta (\beta - \beta_{\mathcal{H}}), \quad (\text{G38})$$

$$\theta = \frac{1}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_{\mathcal{H}}} \right), \quad (\text{G39})$$

so that $\beta_N = \beta^*$ by construction. This corresponds to taking the system through the following sequence of thermal states

$$\varrho_S^{(n)} = \frac{e^{-\beta_n H_S}}{\mathcal{Z}_S(H_S, \beta_n)}. \quad (\text{G40})$$

Note that, in contrast to the coherent protocol where such a sequence can be traversed by simply swapping the target system with a sequence of appropriate machines, in the incoherent setting such a protocol is generally not possible as such swaps are not energy conserving. Nonetheless, we develop a modified protocol that is energy conserving and mimics this idealised one.

Corresponding to each step in the sequence, we define the following quantity, which we eventually show to be related to the heat drawn from the hot bath:

$$G^{(n)} = -n\theta\Delta E_S^{(n)} = -n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n)} - \varrho_S^{(n-1)} \right) \right]. \quad (\text{G41})$$

We proceed to show that the total $\sum_n G^{(n)}$ that we label *the idealised heat cost* $\tilde{\Delta}E_{\mathcal{H}}^*$ is close to the free energy difference over the entire sequence. We have

$$\begin{aligned} \tilde{\Delta}E_{\mathcal{H}}^* &= \sum_{n=1}^N G^{(n)} \\ &= \sum_{n=1}^N n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \end{aligned} \quad (\text{G42})$$

$$= \left\{ \sum_{n=1}^N (n-1)\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \right\} + \theta \operatorname{tr} \left[H_S \left(\varrho_S^{(0)} - \varrho_S^{(N)} \right) \right]. \quad (\text{G43})$$

The sums on the second and third lines above, Eqs. (G42) and (G43) respectively, are the *right* and *left* Riemann sums corresponding to the following integral:

$$\begin{aligned} I &= \int_{q_i}^{q_f} q (-dx) = \int_{q_f}^{q_i} q dx, \\ \text{where } n\theta &\rightarrow q, \\ x &= \operatorname{tr} [H_S \varrho_S(q)], \\ \varrho_S(q) &= \frac{e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}}{\operatorname{tr} [e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}]}. \end{aligned} \quad (\text{G44})$$

We observe that x is the average energy of the thermal state of temperature $\beta + q(\beta - \beta_{\mathcal{H}})$, and thus x and q are strictly monotonically decreasing w.r.t. each other (which explains why the left and right sums are switched). It follows that the Riemann sums bound the integral

$$\sum_{n=1}^N (n-1)\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \leq \int_{q_f}^{q_i} q dx \leq \sum_{n=1}^N n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right]. \quad (\text{G45})$$

We can thus bound the idealised heat cost in both directions via

$$I \leq \tilde{\Delta}E_{\mathcal{H}}^* \leq I + \theta \operatorname{tr} \left[H_S \left(\varrho_S^{(0)} - \varrho_S^{(N)} \right) \right]. \quad (\text{G46})$$

The integral in Eq. (G44) can be shown to be equal to the change in free energy of the target system (w.r.t. inverse temperature β)

$$\begin{aligned} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} [H_S \varrho_S(q)] + \frac{1}{\beta} \operatorname{tr} [\varrho_S(q) \log \varrho_S(q)], \\ \frac{d}{dq} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} \left[\left(H_S + \frac{\mathbb{1}_S + \log \varrho_S(q)}{\beta} \right) \frac{d\varrho_S(q)}{dq} \right]. \end{aligned} \quad (\text{G47})$$

Note that $\varrho_S(q)$ and $d\varrho_S(q)$ are both always diagonal in H_S and full rank for all $q \in \mathbb{R}$, so we have no problems with $\log \varrho_S(q)$, and all of the operators in the expression are well defined and commute. Proceeding, we repeatedly use $\operatorname{tr} [d\varrho_S(q)] = d \operatorname{tr} [\varrho_S(q)] = 0$ and label the partition function $\mathcal{Z}(q) := \operatorname{tr} [e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}]$ to obtain

$$\begin{aligned} \frac{d}{dq} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} \left[\left(H_S + \frac{\log \varrho_S(q)}{\beta} \right) \frac{d\varrho_S(q)}{dq} \right] \\ &= \operatorname{tr} \left[\left(H_S - \frac{\beta + q(\beta - \beta_{\mathcal{H}})}{\beta} H_S - \mathbb{1}_S \frac{\log \mathcal{Z}(q)}{\beta} \right) \frac{d\varrho_S(q)}{dq} \right] \end{aligned}$$

$$= -q \left(1 - \frac{\beta_{\mathcal{H}}}{\beta}\right) \frac{d}{dq} \text{tr} [H_s \varrho_s(q)] = -q \eta \frac{dx}{dq}, \quad (\text{G48})$$

where we identify the Carnot efficiency η for an engine operating between β and $\beta_{\mathcal{H}}$. The integral thus simplifies to

$$I = \eta^{-1} (F_{\beta}[\varrho_s(q_f)] - F_{\beta}[\varrho_s(q_i)]) =: \eta^{-1} \Delta F_s^{(\beta)}. \quad (\text{G49})$$

The idealised heat cost is thus bounded by

$$\eta^{-1} \Delta F_s^{(\beta)} \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq \eta^{-1} \Delta F_s^{(\beta)} + \theta \text{tr} \left[\mathcal{H}_s \left(\varrho_s^{(0)} - \varrho_s^{(N)} \right) \right]. \quad (\text{G50})$$

The left inequality is Landauer's bound applied to cooling a target system with Hamiltonian H_s (see Theorem 5), and the error term on the right can be bounded quite easily; for instance, for $\beta > 0$, we have

$$\begin{aligned} \text{tr} \left[\mathcal{H}_s \left(\varrho_s^{(0)} - \varrho_s^{(N)} \right) \right] &= \text{tr} \left[\left(\mathcal{H}_s - E_s^{\min} \mathbb{1}_s \right) \left(\varrho_s^{(0)} - \varrho_s^{(N)} \right) \right] \\ &\leq \text{tr} \left[\left(\mathcal{H}_s - E_s^{\min} \mathbb{1}_s \right) \varrho_s^{(0)} \right] && \text{since } \mathcal{H}_s - E_s^{\min} \mathbb{1}_s \text{ is a positive operator,} \\ &\leq \text{tr} \left[\left(\mathcal{H}_s - E_s^{\min} \mathbb{1}_s \right) \frac{\mathbb{1}_s}{d_s} \right] \leq \frac{\omega_s^{\max}}{d_s}, \end{aligned} \quad (\text{G51})$$

where $\omega_s^{\max} := E_s^{\max} - E_s^{\min}$ is the largest energy gap in the target system Hamiltonian and d_s is the system dimension. We use the fact that since $\varrho_s^{(0)}$ is a thermal state of positive temperature, its average energy is less than that of the infinite temperature thermal state, $\mathbb{1}_s/d_s$. Since $\theta \propto 1/N$, it follows that one can always find an N large enough such that the error is smaller than a given value, thereby saturating the Landauer bound.

A sequence of machine Hamiltonians to mimic the idealised sequence. Next we construct a protocol that mimics the above sequence and obeys the global energy conservation condition imposed in the incoherent-control setting. The protocol is split into N stages (like above). In each stage, the Hamiltonian of the machine is fixed. The machine here comprises to two parts: the “cold” part and the “hot” part. The cold part is chosen to begin in a thermal state at temperature β of the Hamiltonian

$$H_c = (1 + n\theta) H_s \quad (\text{G52})$$

At this point we note that this sequence of cold-machine states is exactly the same as in the coherent protocol, which would proceed by simply swapping the full state of target system and machine in each stage. However, that is not possible here since this is not an energy-preserving operation. To allow for energy-preserving operations, the hot part of the machine consists of $d_s(d_s - 1)/2$ qubits, each corresponding to a pair of levels (i, j) of the target system (henceforth we take $i < j$ to avoid double counting), whose energy gap is equal to the difference in energies of the target and cold qubit subspaces (hence rendering the desired exchange energy resonant)

$$H_{\mathcal{H}}^{(ij)} = [\omega_i + (1 + n\theta)\omega_j - (\omega_j + (1 + n\theta)\omega_i)] |1\rangle\langle 1|_{\mathcal{H}}^{(ij)} = n\theta (\omega_j - \omega_i) |1\rangle\langle 1|_{\mathcal{H}}^{(ij)}, \quad (\text{G53})$$

where we label the energy eigenvalues of H_s by $\{\omega_i\}$. Each of these hot qubits begins at inverse temperature $\beta_{\mathcal{H}}$. After every unitary operation, the cold and hot parts of the machine are rethermalised to their respective initial temperatures.

To understand the choice of machine Hamiltonians, consider the following two energy eigenstates of the machine: $|i\rangle_c \otimes |1\rangle_{\mathcal{H}}^{(ij)}$ and $|j\rangle_c \otimes |0\rangle_{\mathcal{H}}^{(ij)}$. The energy difference is

$$\Delta^{(ij)} = \omega_j(1 + n\theta) - \omega_i(1 + n\theta) - n\theta(\omega_j - \omega_i) = \omega_j - \omega_i, \quad (\text{G54})$$

matching the energy difference between the corresponding pair of energy eigenstates of the target system. Furthermore, calculating the ratio of populations of the two levels we find

$$g^{(ij)} = \frac{e^{-\beta\omega_j(1+n\theta)}}{e^{-\beta\omega_i(1+n\theta)} e^{-\beta_{\mathcal{H}}n\theta(\omega_j-\omega_i)}} = e^{-(\omega_j-\omega_i)(\beta+n\theta(\beta-\beta_{\mathcal{H}}))}. \quad (\text{G55})$$

This corresponds to the Gibbs ratio of a qubit at the temperature $\beta + n\theta(\beta - \beta_{\mathcal{H}})$, which is the temperature that defines stage n [see Eq. (G38)]. In summary, we construct a machine featuring $d_s(d_s - 1)/2$ qubit subspaces (or virtual qubits), each of the same energy gap as one pair of energy eigenstates of the system, and all of which have a Gibbs ratio (or virtual temperature) corresponding to the n^{th} temperature of our desired sequence.

A single step of the protocol: The max exchange. Within each stage of the protocol, a single step consists of a unitary operation on $SC\mathcal{H}$, followed by the rethermalisation of the machine parts to their respective initial temperatures. We construct the unitary operation as follows: for every pair (i, j) of system energy levels, one can calculate the absolute value of the difference in populations of the following two degenerate eigenstates $|i\rangle_s |j\rangle_c |0\rangle_{\mathcal{H}}^{(ij)}$ and $|j\rangle_s |i\rangle_c |1\rangle_{\mathcal{H}}^{(ij)}$. This value corresponds to the amount of population that would move under an exchange $|i\rangle_s |j\rangle_c |0\rangle_{\mathcal{H}}^{(ij)} \leftrightarrow |j\rangle_s |i\rangle_c |1\rangle_{\mathcal{H}}^{(ij)}$. We then choose the pair with the largest absolute value of this difference and perform that exchange, with an identity operation applied to all other subspaces. We call this unitary operation the *max exchange*. We proceed to prove two statements about the max-exchange operation. First, that the heat extracted from the hot bath is proportional to the change in average energy of the system; and second, that system state under repetition of said operation converges to the thermal state of the temperature that defines the stage n .

Consider the change in average energy of the target system under the exchange unitary. The only two populations that change are those of the $|i\rangle_s$ and $|j\rangle_s$. We label the increase in the population of $|i\rangle_s$ as δp . Then, we have

$$\Delta E_s = \text{tr} [H_s (\varrho'_s - \varrho_s)] = -\delta p (\omega_j - \omega_i). \quad (\text{G56})$$

On the other hand, the populations of the corresponding hot qubit (i.e., tracing out the target system and cold machine) change by the same amount, i.e., there is a move of δp from $|1\rangle_{\mathcal{H}}^{(ij)}$ to $|0\rangle_{\mathcal{H}}^{(ij)}$. In order to rethermalise the hot qubit, the heat drawn from the hot bath is thus

$$\tilde{\Delta} E_{\mathcal{H}} = \delta p n \theta (\omega_j - \omega_i) = -n \theta \Delta E_s. \quad (\text{G57})$$

This is an expression conveniently independent of the pair (i, j) that applies after an arbitrary number of repetitions of the max-exchange operation (which will use different pairs in general).

Convergence of the max-exchange protocol to the virtual temperature. To show that the max-exchange protocol indeed converges to the desired system state in each stage of the protocol, we first prove a rather general statement: given a state ϱ diagonal in the energy eigenbasis, if we exchange any qubit subspace within this system with a virtual qubit of a particular virtual temperature, then the relative entropy of the target system w.r.t. the thermal state of that (virtual) temperature decreases.

To this end, consider the relative entropy of a state ϱ that is diagonal in the energy eigenbasis to a thermal state τ . Labelling the populations of ϱ as p_i and those of τ as q_i , this can be expressed as

$$D(\varrho||\tau) = \sum_k p_k \log \left(\frac{p_k}{q_k} \right). \quad (\text{G58})$$

We now focus on a single-qubit subspace labelled by $\{i, j\}$, which leads to

$$\begin{aligned} D(\varrho||\tau) &= p_i \log \left(\frac{p_i}{q_i} \right) + p_j \log \left(\frac{p_j}{q_j} \right) + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right) \\ &= (p_i + p_j) \left[\frac{p_i}{p_i + p_j} \log \left(\frac{\frac{p_i}{q_i}}{\frac{p_i + p_j}{q_i + q_j}} \right) + \frac{p_j}{p_i + p_j} \log \left(\frac{\frac{p_j}{q_j}}{\frac{p_i + p_j}{q_i + q_j}} \right) \right] + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right) \\ &= N \left(\bar{p}_i \log \frac{\bar{p}_i}{\bar{q}_i} + \bar{p}_j \log \frac{\bar{p}_j}{\bar{q}_j} + \log \frac{N}{N_v} \right) + \sum_{k \notin \{i, j\}} p_k \log \frac{p_k}{q_k}. \end{aligned} \quad (\text{G59})$$

In the last line we renormalise the populations within the qubit subspace and labelled the total populations of the system and thermal state qubit subspaces of interest by N and N_v , respectively. Labelling the normalised states within these subspaces as ϱ_v and τ_v respectively, we have

$$D(\varrho||\tau) = N \left[D(\varrho_v||\tau_v) + \log \left(\frac{N}{N_v} \right) \right] + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right). \quad (\text{G60})$$

Suppose now that this qubit subspace of the target system is exchanged with a qubit subspace of any machine that has the same temperature as the thermal state above. The only object that changes in the the above expression is ϱ_v , since the norm N remains the same. In addition, ϱ_v always gets closer to τ_v under such an exchange [38, 58], implying that the relative entropy always strictly decreases under such an operation.

Returning to the max-exchange protocol, note that by construction, every virtual qubit in the machine that is exchanged with the qubit subspace $\{i, j\}$ of the target system in a given stage n has the same virtual temperature, $\beta_n = \beta + n\theta(\beta - \beta_{\mathcal{H}})$. Thus the relative entropy of the system to the thermal state at this temperature always decreases under this operation, unless the operation

does not shift any population, which happens only at the unique fixed point where every qubit subspace of the system is already at the virtual temperature β_n . By monotone convergence, the relative entropy must converge, and moreover converge to the value that it has at the fixed point of the operation, which is the thermal state at inverse temperature β_n . Note that rather than choosing the qubit subspace with maximum population difference to exchange we could also have picked at random from among the pairs $\{i, j\}$ and convergence would still hold; the max-exchange protocol simply ensures the fastest rate of convergence among these choices.

Choosing a large enough number of repetitions in each stage so that the overall heat cost is close to the idealised heat cost. Given that the max-exchange protocol in stage n converges to the thermal state that we label $\varrho_s^{(n)}$, given any error δ_E , we choose a number of repetitions m_n that is large enough so that the difference between the average energy of the actual final state of this stage, which we label $\tilde{\varrho}_s^{(n)}$, and that of the ideal state $\varrho_s^{(n)}$ is less than δ_E . In this case, the total heat cost over all stages is close to the idealised heat cost

$$\begin{aligned} \left| \tilde{\Delta}E_{\mathcal{H}} - \tilde{\Delta}E_{\mathcal{H}}^* \right| &= \left| \sum_{n=1}^N \left\{ -n\theta \operatorname{tr} \left[H_s \left(\tilde{\varrho}_s^{(n)} - \tilde{\varrho}_s^{(n-1)} \right) \right] \right\} - \sum_{n=1}^N \left\{ -n\theta \operatorname{tr} \left[H_s \left(\varrho_s^{(n)} - \varrho_s^{(n-1)} \right) \right] \right\} \right| \\ &= \left| \sum_{n=0}^{N-1} \theta \operatorname{tr} \left[H_s \left(\tilde{\varrho}_s^{(n)} - \varrho_s^{(n)} \right) \right] - N\theta \left(\tilde{\varrho}_s^{(N)} - \varrho_s^{(N)} \right) \right| \\ &\leq 2N\theta\delta_E = 2 \left(\frac{\beta^* - \beta}{\beta - \beta_{\mathcal{H}}} \right) \delta_E. \end{aligned} \quad (\text{G61})$$

The number of repetitions in each stage m_n required depends only upon the initial choice of β^* and N .

Completing the proof. Finally, suppose that one is given any target temperature β^* and two arbitrarily small errors, ϵ_β for the cooling and ϵ_E for the heat cost, and asked to cool incoherently in such a way that achieves

$$|\beta' - \beta^*| \leq \epsilon_\beta, \quad (\text{G62})$$

$$\left| \tilde{\Delta}E_{\mathcal{H}} - \eta^{-1} \Delta F_s^{(\beta)} \right| \leq \epsilon_E. \quad (\text{G63})$$

We proceed by first choosing a number of stages N so that the idealised heat cost $\tilde{\Delta}E_{\mathcal{H}}^*$ is within $\frac{\epsilon_E}{2}$ to the Carnot-Landauer bound above. The idealised sequence of temperatures satisfies $\beta_N = \beta^*$ by construction. Once N is fixed, for each stage from $n = 1$ to $N - 1$ we choose a number of repetitions for each stage m_n such that the actual heat cost is within $\frac{\epsilon_E}{2}$ of the idealised heat cost, as discussed above. This ensures that the total heat cost is within ϵ_E of the bound. Finally, we check that the number of repetitions of the last stage m_N is large enough for us to be within ϵ_β of β^* . If not, we increase the number of repetitions (this can only decrease the error in the heat cost anyway) until we are close enough, as required. \square

Appendix H: Comparison of Cooling Paradigms and Resources for Imperfect Cooling

Although we have looked at a number of cooling protocols throughout to demonstrate the ability for perfect cooling in the asymptotic limit, here we focus on imperfect cooling behaviour, i.e., when all resources are restricted to be finite and thus a perfectly pure state cannot be attained. We have three main goals in doing so.

1. To illustrate the finite trade-offs between the trinity of resources (energy, time, control complexity).
2. To compare the behaviour of different constructions of the cooling unitary for machines of the same size (i.e., analysing the energy-time trade-off for fixed control complexity).
3. To demonstrate the increase in resources required for cooling in the thermodynamically self-contained paradigm of energy-preserving unitaries (i.e., incoherent control), as compared to coherently driven unitaries.

H1. Rates of Resource Divergence for Linear Qubit Machine Sequence

Consider cooling a qubit target system with energy gap ω_s by swapping it sequentially with a sequence of N machine qubits of linearly increasing energy gaps. In Appendix G1, we derived the deviation from the idealised heat dissipation in the incoherent

control setting for a sequence of N machines [see Eq. (G17)], which we repeat below:

$$\frac{1}{\eta} (F^* - F_0) \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq \frac{1}{\eta} (F^* - F_0) + \frac{\omega_s}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right). \quad (\text{H1})$$

We can immediately adapt this result to the paradigm of coherent control by taking $\beta_H = 0$ and replacing the heat by work, which yields

$$\Delta F_s \leq W \leq \Delta F_s + \frac{\omega_s}{N} \left(\frac{\beta^*}{\beta} - 1 \right). \quad (\text{H2})$$

Since the above inequalities are derived from the left and right Riemann sums of an integral, as N becomes large, one can expect that W lies roughly halfway between both extremes; we can thus cast the scaling in the approximate form

$$\left[\frac{W - \Delta F}{\omega_s} \right] N \sim \frac{1}{2} \left(\frac{\beta^*}{\beta} - 1 \right). \quad (\text{H3})$$

Thus, we see that the relevant quantifier of the energy resource here is the extra work cost above the Landauer limit relative to the system energy. Additionally, the quantifier of how much said resource is required (per machine qubit) is $\beta^*/\beta - 1$, which, for cold enough final temperatures, is approximately the ratio β^*/β .

Returning to the incoherent control paradigm, analysing the scaling behaviour between energy and time is more complicated. On the one hand, the expression above is only slightly modified, with the work being replaced by the heat dissipated multiplied by the Carnot factor:

$$\left[\frac{\eta \Delta E_{\mathcal{H}} - \Delta F}{\omega_s} \right] N \sim \frac{1}{2} \left(\frac{\beta^*}{\beta} - 1 \right), \quad (\text{H4})$$

which is consistent with the work-to-heat efficiency of a Carnot engine. However, in the case of incoherent control, since the population swap only takes place within a subspace of the two-qubit machine, the total population is not completely exchanged in a single operation (in contrast to that in the coherent control setting). Thus the number of operations here required to transfer a desired amount of population to the ground state of the target is greater than the number of machine qubits N . To make a fair comparison, one could either compare the same number of machine qubits but swap repeatedly (with rethermalisation of the machine in between operations)—thereby fixing the control complexity at the expense of longer time—or one could increase the number of machine qubits and count time by the number of two-level swaps—thereby fixing time to be equal at the expense of increased control complexity overall. We investigate both methods in the coming section.

H2. Comparison of Coherent and Incoherent Control

Intuitively, the incoherent control paradigm requires the utilisation of a greater amount of resources (albeit less overall control in general) than the coherent control counterpart because of two distinct disadvantages. First, the temperature of the baths plays a substantial role in cooling performance. Consider the example of a SWAP gate applied between a system and machine qubit: in the coherent control case, this operation transforms the target system to the state of the thermal machine qubit, characterised by the Gibbs ratio of ground-state to excited-state population. In the incoherent control case, one requires the addition of a thermal qubit from the hot bath to render said operation energy preserving; as a result, the Gibbs ratio of the virtual qubit that the target system swaps with is, in general, worse than that of the coherent control setting, and only becomes equal in the limit of an infinite temperature hot bath. This is the first disadvantage. The second disadvantage is that in the incoherent control setting, the target system swaps with only a subspace of the machine rather than the entire one, i.e., it is swapped with a virtual qubit. Thus, the exchange of population is only partial as compared to the coherent control case: in the limiting case of an infinite temperature hot bath, said factor goes to $\frac{1}{2}$ for all relevant two-level subspaces. This implies that a greater number of operations, and thus time, is required in the incoherent control paradigm in order to achieve a similar result as its coherent control counterpart.

We illustrate this behaviour via the following example. The system is a degenerate qubit (beginning in the maximally mixed state), and we fix the final target ground-state population ($p = 0.99$, corresponding to $\epsilon = 1 - p = 0.01$). Even in this simple case, the optimal finite-resource protocols with coherent and incoherent control are not known; we therefore compare protocols from each setting that make use of machines of a similar structure, namely swapping with machine qubits (virtual ones, in the incoherent control setting) of linearly increasing energy gaps.

More specifically, the coherent control cooling protocol employed is that of a sequence of swaps with machine qubits of linearly increasing energy gaps, and for the fixed target population, we can calculate the surplus work cost over the Landauer limit as a function of the number N of operations (which corresponds in this case to the number of machine qubits). In the

incoherent control case, we take the hot bath to be at infinite temperature, allowing for the potential saturation of the Landauer limit as in the coherent case. In this way we isolate the disadvantage that arises due to working in degenerate subspaces in our analysis. Here too we take a linear sequence of energy gaps for the cold (and hot) baths, with a single operation step corresponding to a three-level energy-conserving exchange involving the qubit taken from each of the hot and cold parts of the machine, i.e., $|1\rangle_s|0\rangle_c|0\rangle_h \leftrightarrow |0\rangle_s|1\rangle_c|1\rangle_h$. As mentioned previously, for an incoherent control protocol of fixed overall machine size, there are essentially two extremal methods of implementation. The first is to identify N two-level subspaces of the total machine with distinct energy gaps and perform the sequence of virtual swaps between them and the target; in the language of Appendix G, we therefore have N different stages with a single step within each stage (no repetitions) before moving on to the next stage. The second is to take N/m two-level subspaces and swap the target with each virtual qubit m times before moving on to the next; in other words, we here have N/m different stages with m steps (repetitions) within each stage. For the same fixed ground-state population, we plot the surplus work cost (energy drawn from the hot bath in the case of incoherent control) against the total machine size and number of two-level unitary swaps, as characterised by N , for both of these incoherent control adaptations, comparing them to the coherent control paradigm in Fig. 4.

In both control paradigms, we see that the deviation of the energy cost above the Landauer limit scales inversely with the number of operations [as expected from Eqs. (H3) and (H4)], but the proportionality constant is worse in the case of incoherent control. Moreover, the incoherent control paradigm with no repetitions within stages outperforms that with multiple repetitions, as intuitively expected since the former protocol corresponds to one for which the spacing between distinct energy gaps that are utilised is smaller, allowing us to stay closer to the reversible limit in each step. In our example, the no repetition incoherent control protocol is around 3 times worse than the coherent control protocol and the incoherent control protocol with $m = 5$ repetitions is around 5.3 times worse, implying that one would require that many times the number of operations (i.e., that much more time) to achieve the same performance with incoherent control paradigm as with coherent control.




FIG. 4. *Imperfect Cooling with Coherent and Incoherent Control.* We compare the performance of coherent and incoherent control protocols for cooling a degenerate qubit target by swapping it with machine qubits with linearly increasing energy. The final ground-state population is fixed to be 0.99. The inverse of the surplus work cost $W - \beta \Delta S_s$ (with $\beta = 1$) is plotted (in units of the smallest machine energy gap, $\omega_{\mathcal{M}}^{\min}$) against the total number of unitary operations, with the temperature of the hot bath in the incoherent control protocols set to $\beta_H = \infty$ in order to make meaningful comparison to the coherent control case. We see that the coherent control protocol (blue) outperforms the two incoherent ones (purple, red) at any given time. As discussed in the text, there are two choices for how to implement an incoherent control protocol of this type with fixed control complexity: The red line corresponds to a protocol in which a machine (subspace) with the same energy gap is reused 5 times before moving on to the next; on the other hand, the purple line depicts the case where there are no repetitions within each stage defined by a distinct energy gap in the machine. By inspection, the single-use incoherent protocol (purple) requires approximately 3 times more unitaries to achieve the same efficiency as the coherent one (blue), whereas the five-repetition incoherent protocol (red) requires approximately 5.3 times as many unitaries as the coherent one.