6-5-2014

Quantum Resources for Purification and Cooling: Fundamental Limits and Opportunities

Francesco Ticozzi
Dartmouth College

Lorenza Viola
Dartmouth College

SCIENTIFIC REPORTS

OPEN

SUBJECT AREAS:
QUANTUM MECHANICS
THERMODYNAMICS

Received
2 April 2014

Accepted
16 May 2014

Published
5 June 2014

Correspondence and requests for materials should be addressed to F.T. (ticozzi@dei.unipd.it) or L.V. (lorenza.viola@dartmouth.edu)

Quantum resources for purification and cooling: fundamental limits and opportunities

Francesco Ticozzi[1,2] & Lorenza Viola[2]

[1]Dipartimento di Ingegneria dell'Informazione, Università di Padova, via Gradenigo 6/B, 35131 Padova, Italy, [2]Department of Physics and Astronomy, Dartmouth College, 6127 Wilder Laboratory, Hanover, NH 03755, USA.

Preparing a quantum system in a pure state is ultimately limited by the nature of the system's evolution in the presence of its environment and by the initial state of the environment itself. We show that, when the system and environment are initially uncorrelated and arbitrary joint unitary dynamics is allowed, the system may be purified up to a certain (possibly arbitrarily small) threshold if and only if its environment, either natural or engineered, contains a "virtual subsystem" which has the same dimension and is in a state with the desired purity. Beside providing a unified understanding of quantum purification dynamics in terms of a "generalized swap process," our results shed light on the significance of a no-go theorem for exact ground-state cooling, as well as on the quantum resources needed for achieving an intended purification task.

Cooling of quantum systems toward their ground state plays a central role across low-temperature physics and quantum science, by providing the key to unlock novel phases of matter and quantum behavior – as exemplified in settings ranging from laser cooling of atoms and molecules[1–3] to dynamical nuclear polarization in solid- and liquid-state nuclear magnetic resonance[4,5], and cooling of mechanical resonators[6–10]. From a quantum control standpoint, the task of cooling (or "refrigeration," in the language of quantum thermodynamics[11]) may be viewed as an instance of dissipative pure-state preparation, which is in turn closely related to the more general task of *purification* – namely, the ability to steer the system from an arbitrary initial state to a final state with higher purity. Within quantum information processing (QIP), access to pure states is presumed in all quantum computation models that can provably achieve an exponential speed-up over classical ones[12,13], and cold ancilla qubits are critical to the success of fault-tolerant quantum error correction[14]. As a result, schemes for cooling and purification are being actively investigated[15–17], and underlying assumptions and implications formalized with added rigor[18–21].

While in practice a variety of system-dependent imperfections and technological constraints will inevitably hinder the achievable performance, a fundamental question is to determine what ultimate limitations may nevertheless exist on the sole basis of some generic, *system-independent* assumptions on the underlying dynamics. Specifically, assume that *arbitrary* unitary evolution is allowed on the target system S together with its environment E, starting from arbitrary *factorized* initial conditions. To what extent does the initial, typically highly-mixed state of E, limit the degree of purity attainable on S *in principle*? Conversely, if the environment E and its initial state can be controllably engineered, what are the *minimal* resources for purification (cooling) of S to be guaranteed to a prescribed accuracy?

Our main contribution in this work is the identification of *necessary and sufficient conditions* for exact as well as approximate purification and ground-state cooling, given the above ideal scenario. Our starting point is a trivial example: if both S and E are two-dimensional systems (qubits), purification of S is clearly possible in principle only if the initial state of E has a lower entropy, in which case the optimal purification dynamics simply amounts to swapping the two initial states. In a general open-system setting, our strategy is to make precise the intuition that purity can still only be exchanged but not created between subsystems, albeit the latter need no longer coincide with the natural ones. The relevant notion is provided by the concept of a "virtual" subsystem as a *factor of a subspace* of a larger state space, as introduced by Knill *et al.*[22] in the context of quantum error correction and extensively used in QIP[23–28].

Our results complement existing work and advance current understanding in several ways. While a no-go theorem for ground-state cooling under initial system-thermal bath factorization was recently established in

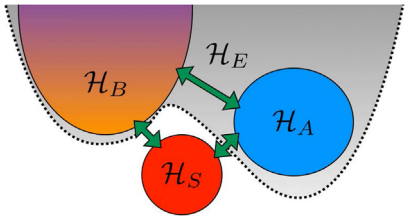

Figure 1 | The system of interest, S, may be generally coupled to a quantum bath, B, and an engineered auxiliary system, A. We collectively refer to the pair (B, A) as the *environment*. The initial state on $\mathcal{H}_{SE} \equiv \mathcal{H}_S \otimes \mathcal{H}_E = \mathcal{H}_S \otimes (\mathcal{H}_B \otimes \mathcal{H}_A)$ is assumed to be *fully factorized* with respect to this partition, i.e., $\rho_{SE} = \rho_S \otimes \rho_E = \rho_S \otimes (\rho_B \otimes \rho_A)$. The joint dynamics is generated by a total Hamiltonian of the form $H \equiv H_0 + H_c(t) = (H_S \otimes \mathbb{I}_E + \mathbb{I}_S \otimes H_E + H_{SE}) + H_c(t)$, where the control Hamiltonian $H_c(t) \equiv \sum_\ell u_\ell(t) H_{c,\ell}$ acts trivially on B. If $\dim(\mathcal{H}_{SE}) < \infty$, complete propagator controllability is ensured in the generic case where the Lie algebra of skew-symmetric operators generated by the control Hamiltonians $\{iH_{c,\ell}\}$, together with the natural "drift" Hamiltonian iH_0, is the whole $\mathfrak{su}(d_S \times d_E)$[35]. If so, there exist some time $T > 0$ and control functions $u_\ell(t)$, $t \in [0, T]$, that allow to reach any element in $\mathfrak{U}(\mathcal{H}_{SE})$ to arbitrary precision. For our discussion, it is not essential to specify how the control actions are enacted. For instance, if $\mathcal{H}_A \simeq \mathbb{C}$, our setting includes open-loop control of S via a semiclassical controller[36,37]. In this case, B is controlled via its interaction with S, yet *indirect controllability* of B given an arbitrary initial state of S still suffices for complete joint controllability, as we assume[38]. If $\dim(\mathcal{H}_A) > 1$, dynamics in the presence of a coherent "quantum controller" and/or an an engineered quantum reservoir[16,37] may be accounted for. In this case, the uncontrollable component B may couple to both S and A in general.

Ref. 20, our analysis further clarifies that such a no-go strictly applies only to *exact* cooling. Notwithstanding initial factorization, no fundamental limit exists to *arbitrarily accurate* purification and ground-state cooling in principle, so long as the environment is effectively infinite-dimensional, and capable of supporting a sufficiently pure virtual subsystem. From a quantum-simulation standpoint, this reinforce the conclusion that a simulated ancillary environment consisting of a single qubit suffices for enacting arbitrary open-system dynamics, so long as it can be measured and *reset* to a sufficiently pure state[29], as recently demonstrated in trapped-ion experiments[30,31]. Conceptually, our analysis points to a *generalized swap process* as the unifying physical mechanism through which *any* purification or groundstate cooling dynamics may ensue from joint unitary evolution, as opposed to known special instances limited to small dimension and/or a fixed (thermal) initial environment state[20,21]. From an open-system quantum-control perspective, our general picture may be exploited to design procedures for purification and ground-state cooling via environment (or "reservoir") engineering, as potentially relevant to a growing number of quantum technologies, see e.g. Refs. 29,32,33 and references therein. Interestingly, within quantum foundations, our results have also implications for dynamical reduction models[34]: in order for the "wave-function collapse" predicted by the standard von Neumann postulates to be consistently reproduced by underlying open-system dynamics, the environment interacting with the system must, again, harbor a sufficiently pure virtual subsystem.

Results

Setting. The general setting we consider is depicted in Fig. 1. The target quantum system S, with associated Hilbert space $\mathcal{H}_S$ of dimension d_S, is coupled to a quantum environment E, with associated Hilbert space $\mathcal{H}_E$ of dimension d_E, which may generally include both a component that is not directly controllable (a physical "bath", B) and a fully controllable auxiliary system (or "ancilla", A). We take $d_E \geq d_S$, so that we may decompose $d_E \equiv d_S\, d_F + d_R$, with d_F being is the integer part of d_E/d_S, and $d_R < d_S$ the rest. While we further assume that $d_S < \infty$ in what follows, we may formally extend our results to infinite-dimensional target systems of interest (notably, quantum oscillators) by imposing a finite-energy constraint.

A key assumption is that *no* correlations are initially present between the constituents of the joint system, i.e., the initial state is factorized, $\rho_{SE} = \rho_S \otimes \rho_E$, with ρ_E a trace-class operator in case $d_E = \infty$. Other than that, and unlike in Refs. 20,21, *no* restriction is placed on either ρ_S or ρ_E which, in particular, need not be thermal. We shall denote by $\{\lambda_j(\rho_E)\}$ the eigenvalues of ρ_E, considered with their multiplicity and in non-increasing order.

While the inclusion of both a bath and an ancillary system allows for different physical scenarios to be discussed within the same framework (see caption), the central mathematical assumption is that suitable Hamiltonian control is available on $S + E$ together, so that *any unitary operator in* $\mathfrak{U}(\mathcal{H}_{SE})$ *can be obtained at some time* T. In control-theoretic terms, this is equivalent to assuming complete joint propagator controllability[35,39]. Hence, at any given time T, the joint evolution of $S + E$ is described by some $U_{SE}(T)$ that we are free to choose. The conditions for this to be possible have been extensively investigated within the geometric control framework. At least if $\mathcal{H}_{SE}$ has finite dimension, complete controllability is generic[35], albeit efficient constructive methods for control design are still object of ongoing research, along with controllability conditions for infinite-dimensional quantum systems[40,41].

Starting from factorized initial conditions, the reduced state of the system after the unitary (controlled) evolution takes place is given by

$$\rho'_S \equiv \mathrm{Tr}_E(\rho'_{SE}) = \mathrm{Tr}_E\left(U_{SE}\rho_S \otimes \rho_E U_{SE}^\dagger\right). \tag{1}$$

Exact purification of S is attained if ρ'_S is pure irrespective of the initial state ρ_S, that is, $\rho'_S \equiv |\psi\rangle\langle\psi|$ for some $|\psi\rangle \in \mathcal{H}_S$, so that $\mathrm{Tr}\left(\rho'^2_S\right) = 1$. However, this requirement is too strong in practical situations of interest. We say that *(ε-)approximate purification* of S can be attained if the state of S may be brought to within distance ε from a pure state irrespectively of the initial ρ_S, that is, there exists $|\psi\rangle \in \mathcal{H}_S$ such that

$$d(\rho'_S, |\psi\rangle\langle\psi|) \leq \varepsilon, \quad \forall \rho_S. \tag{2}$$

Here, $d(X,Y) \equiv \frac{1}{2}\mathrm{Tr}(|X-Y|) = \frac{1}{2}\|X-Y\|_1$ is the quantum total-variation distance, which is a natural measure of distinguishability between quantum states[12,26,27]. (Note that the robustness requirement with respect to the system initialization makes our purification notion *stronger* than used in both Refs. 20,21. For given initial ρ_S and ρ_E with *known* spectrum, an upper bound on the purity of the final state ρ'_S may be additionally established[21].) Exact purification is recovered by requesting $\varepsilon = 0$. In the following, we shall consider $0 \leq \varepsilon \ll 1$.

The fact that the joint dynamics $\rho_{SE} \mapsto \rho'_{SE}$ is unitary is equivalent to the preservation of the spectrum of the joint density operator at any time. However, one still intuitively expects purification of a "portion" of the overall system to be possible in an appropriate sense, the limitations on what can be achieved stemming from the initial state of E. Let us first consider a trivial example.

Example 1.– Suppose that both the target system and the environment are a qubit. The factorized initial state can then be parametrized by the maximum eigenvalue of its two components ρ_S, ρ_E, say, $1/2 \leq p_S, p_E \leq 1$ respectively, with the value 1/2 corresponding to a fully mixed state. That is, $\rho_{SE} = \mathrm{diag}(p_S, 1-p_S) \otimes \mathrm{diag}(p_E, 1-p_E)$. Since for qubits the von Neumann's entropy $\mathcal{S}(\rho)$ is completely determined by, and is a decreasing function of, the maximum eigenvalue of the state, we can pursue a direct information-theoretic analysis. Achieving maximal purification is thus equivalent to achieving the (reduced) state ρ'_S in Eq. (1) with minimum entropy with respect

to the choice of U_{SE}. Using the standard definitions of joint and conditional entropy[12], we may write

$$\mathcal{S}(\rho'_S) = \mathcal{S}(S) = \mathcal{S}(S,E) - \mathcal{S}(E|S),$$

where $\mathcal{S}(S,E) = \mathcal{S}(\rho'_{SE}) = \mathcal{S}(\rho_S \otimes \rho_E)$, and the conditional entropy is maximal when the state is factorized. Hence, the maximal purification is attained by either swapping the states (when $p_E > p_S$), or leaving them as they are (when $p_E < p_S$). In other words, some purification is possible *only if* the entropy of the auxiliary qubit is lesser than the one of the system qubit, and exact purification can only be achieved if the former is in a pure state to begin with.

Despite its simplicity, this example suggests a general strategy to tackle the purification problem: given a target system to be purified, if in its environment we may identify a "subsystem" of the same dimension, that is initially in a more pure state, all we need to do is to swap these two states. Formalizing this idea leads to the rigorous conditions we are seeking.

Main result: necessary and sufficient conditions for purification. In common physical situations, subsystems may be naturally identified with (distinguishable) quantum particles and/or degrees of freedom, and their state space directly associated to different factors of the overall tensor-product Hilbert space. This view is not, however, sufficiently general to capture all relevant settings that arise both physically and in the context of QIP applications. Within quantum error correction theory, for instance, "noise-protected" quantum-information-carrying logical degrees of freedom are associated with *virtual subsystems* that typically do not correspond with the original qubit subsystems[22,25,42]. This more general subsystem notion will also be key to our analysis. Mathematically, a *virtual quantum subsystem* $\tilde{S}$ of a larger system E (the environment in our case) is associated with a tensor factor $\mathcal{H}_{\tilde{S}}$ of a *subspace* of $\mathcal{H}_E$[22–25], that is, we may write

$$\mathcal{H}_E = (\mathcal{H}_{\tilde{S}} \otimes \mathcal{H}_F) \oplus \mathcal{H}_R, \tag{3}$$

for some factor $\mathcal{H}_F$ and a (generally non-trivial) remainder space $\mathcal{H}_R$. System E is said to be initialized in subsystem $\tilde{S}$ if its state may be decomposed as $\rho_E = \rho_{\tilde{S}} \otimes \rho_F \oplus 0_R$, where 0_R is the zero operator on $\mathcal{H}_R$ and ρ_F a state on $\mathcal{H}_F$; in particular, *E is initialized in a subsystem pure state* if $\rho_{\tilde{S}} = |\tilde{\varphi}\rangle\langle\tilde{\varphi}|$, for $|\tilde{\varphi}\rangle \in \mathcal{H}_{\tilde{S}}$[26,43,44]. While virtual subsystems are most compactly described in terms of an operator-algebraic characterization[22,23,27,28], a basis with the correct tensor/direct product structure may also be straightforwardly constructed (see Methods). We are now ready to state our central result:

Theorem. *Assume complete unitary controllability and factorized initial conditions* $\rho_S \otimes \rho_E$ *and on* $\mathcal{H}_S \otimes \mathcal{H}_E$. *Then the following conditions hold:*

(i) *For every* $\varepsilon > 0$, ε-approximate purification *of S may be achieved if there exists a decomposition of* $\mathcal{H}_E$ *as in Eq. (3), with* $\mathcal{H}_{\tilde{S}} \simeq \mathcal{H}_S$, *and a pure-state initialization of E in* $\tilde{S}$, $\tilde{\rho}_E = |\tilde{\varphi}\rangle\langle\tilde{\varphi}| \otimes \rho_F \oplus 0_R$, *such that*

$$d(\rho_E, \tilde{\rho}_E) \leq \varepsilon. \tag{4}$$

(ii) Exact purification *of S (ε = 0) may be achieved if and only if the initial state of the environment has exactly the form* $\rho_E = |\tilde{\varphi}\rangle\langle\tilde{\varphi}| \otimes \rho_F \oplus 0_R$, *for some* $|\tilde{\varphi}\rangle \in \mathcal{H}_{\tilde{S}}$.

(iii) *ε-approximate purification is always possible provided that* $\varepsilon \geq \tilde{\varepsilon}(\rho_E)$, *where*

$$\tilde{\varepsilon}(\rho_E) \equiv \tilde{\varepsilon} = 1 - \sum_{j=1}^{d_F} \lambda_j(\rho_E) \geq 0. \tag{5}$$

$\tilde{\varepsilon}$*-purification is optimal whenever* $d_R = 0$. *In particular,* arbitrarily accurate purification ($\tilde{\varepsilon} = 0$, $\varepsilon > 0$) *is always possible for* $d_E = \infty$.

Part (i) of the above theorem can be easily proven by considering a unitary operator W_{SE} that at some time T swaps the state of S with the one of its isomorphic copy $\tilde{S}$, which is initially in a pure state $|\tilde{\varphi}\rangle\langle\tilde{\varphi}|$. With the precise definition of W_{SE} being given in the Methods section, the basic observation is to note that if ρ_E satisfies Eq. (4), then it can be written as

$$\rho_E \equiv \tilde{\rho}_E + \Delta\rho_E, \quad \tfrac{1}{2}\mathrm{Tr}(|\Delta\rho_E|) \leq \varepsilon. \tag{6}$$

By implementing the swap dynamics, it thus follows that

$$\begin{aligned} \rho'_S &= \mathrm{Tr}_E\left(W_{SE}\rho_S \otimes \rho_E W_{SE}^\dagger\right) \\ &= \mathrm{Tr}_E\left[|\tilde{\varphi}\rangle\langle\tilde{\varphi}| \otimes (\rho_{\tilde{S}} \otimes \rho_F \oplus 0_R) + W_{SE}\rho_S \otimes \Delta\rho_E W_{SE}^\dagger\right] \\ &\equiv |\tilde{\varphi}\rangle\langle\tilde{\varphi}| + \tilde{\mathcal{E}}(\rho_S \otimes \Delta\rho_E), \end{aligned} \tag{7}$$

where $\tilde{\mathcal{E}}$ is a trace-preserving completely-positive map and hence a trace-norm contraction[12]. Since, together with Eq. (6), this implies that

$$d(\rho'_S, |\tilde{\varphi}\rangle\langle\tilde{\varphi}|) \leq \varepsilon, \quad \forall \rho_S,$$

the target system S is ε-purified, as desired. It is immediate to see that the same argument also applies if $\varepsilon = 0$. In other words, the condition of Eq. (4) is always *sufficient* for ε-purification with $\varepsilon \geq 0$, independently of the dimension and the initial state of E.

Establishing that Eq. (4) remains *necessary* is relatively straightforward for exact purification [as in part (ii)], but more subtle in the approximate case [part (iii)]. While full proofs are given in the Methods section, the gist of the argument showing why ε-purification is indeed always possible for $\varepsilon \geq \tilde{\varepsilon}$ may be summarized as follows. Assume that for an initial state $\rho_{SE} = \rho_S \otimes \rho_E$, the desired purification can be attained at some final time T. Then there exists an orthogonal projector, say, $\Pi_T = |\psi\rangle\langle\psi|_S \otimes I_E$, such that $\mathrm{Tr}(\Pi_T \rho'_{SE}) \geq 1 - \varepsilon$, for all ρ_S. If we define a new projector $\Pi_0 \equiv U_{SE}^\dagger \Pi_T U_{SE}$, this condition clearly also implies that $\mathrm{Tr}(\Pi_0 \rho_{SE}) \geq 1 - \varepsilon$. This inequality shows that a pure subsystem of dimension d_S may be identified to within distance ε from the initial *joint* state as well. The tricky part is to establish that this in turn implies the existence of an $\tilde{\varepsilon}$-pure subsystem in the environment *alone*.

In order to do so, we may consider the worst-case scenario, that is, a fully mixed (infinite temperature) initial state on S, with $\rho_{SE} \equiv \tilde{\rho}_{SE} = (1/d_S) I_S \otimes \rho_E$. The idea is to construct a projector of the form $\tilde{\Pi}_0 \equiv I_{\tilde{S}} \otimes \Pi_1$, where Π_1 is a projector on d_F eigenvectors of ρ_E with highest eigenvalues, which projects on a subspace, say $\mathcal{H}_1 \subseteq \mathcal{H}_E$, of the same dimension of Π_T. This is the best possible strategy whenever $d_R = 0$, and we may show that:

$$\mathrm{Tr}(\Pi_1 \rho_E) = \mathrm{Tr}(\tilde{\Pi}_0 \tilde{\rho}_{SE}) = 1 - \tilde{\varepsilon}, \quad \tilde{\varepsilon} \leq \varepsilon. \tag{8}$$

Accordingly, the subspace $\mathcal{H}_1$, onto which Π_1 projects, collects $(1 - \tilde{\varepsilon})$ of the total probability. The existence of such a subspace may be shown to be equivalent to the existence of *a virtual subsystem* $\tilde{S}$, *such that E is* $\tilde{\varepsilon}$*-close to pure-state initialization in* $\tilde{S}$, as desired.

Our theorem points to an interesting dichotomy between finite- vs. infinite-dimensional environments. If $d_E < \infty$, ε-purification of S may or may not be achievable, depending on whether the conditions on the spectrum of ρ_E imposed by Eq. (8) are fulfilled, for arbitrary ρ_S. If $d_E = \infty$, however, then $\tilde{\varepsilon} = 0$ and for *any* trace-class state of E and *any* fixed $\varepsilon > 0$, a sufficiently pure subsystem always exists. We illustrate how to explicitly construct such a ε-pure subsystem in the case where the target system is a qubit, as the generalization to a higher-dimensional system (qudit) is straightforward.

Let ρ_E be a trace-class environment state, and consider its spectral representation, say, $\rho_E = \sum_{j=1}^{\infty} p_j |j\rangle\langle j|$, $\sum_{j=1}^{\infty} p_j = 1$. The identification of the desired ε-pure subsystem may be accomplished by

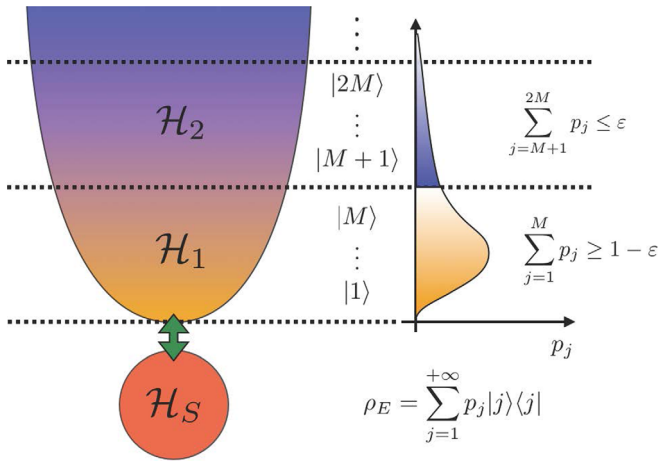

Figure 2 | The target system (with d_S-dimensional state space $\mathcal{H}_S$) is coupled to an infinite-dimensional quantum bath (with state space $\mathcal{H}_B$), initially in an arbitrary state ρ_B. To construct a subsystem of B which is arbitrarily (yet not perfectly) pure, we identify a finite-dimensional subspace $\mathcal{H}_1$ that collects the first M eigenvectors of ρ_B accounting for $(1-\varepsilon)$ of the total probability. To complete this virtual subsystem, we only need to identify (d_S-1) orthogonal subspaces $\mathcal{H}_i$, each of dimension M. Purification is then attained by swapping the virtual subsystem's state with the one of the target system.

identifying two orthogonal subspaces $\mathcal{H}_1, \mathcal{H}_2 \subsetneq \mathcal{H}_E$ each of dimension M, one of which accounts for (at least) $(1-\varepsilon)$ probability. Since ρ_E is trace class, hence its spectrum is absolutely summable, for any $\varepsilon > 0$ there exists an M large enough such that $\sum_{j>M} p_j < \varepsilon$. Define $\mathcal{H}_1 \equiv \mathrm{span}\{|j\rangle\}_{j=1,\ldots,M}$, and $\mathcal{H}_2$ any M-dimensional subspace orthogonal to $\mathcal{H}_1$. From these two subspaces, we can easily construct a subsystem decomposition as in Eq. (3), with $(\mathcal{H}_{\tilde{S}}) = 2$, $\dim(\mathcal{H}_F) = M$, such that the final reduced state ρ'_S is ε-close to a pure state. The strategy is pictorially illustrated in Fig. 2. The general qudit case can be obtained along the same lines, by considering d_S copies of the M-dimensional subspace, where again only one accounts for (at least) $(1-\varepsilon)$ of the total probability. Similar considerations also apply to typical physical scenarios where $d_F \approx d_E/d_S \approx d_E \gg d_S$, in which case nearly arbitrary accuracy $\tilde{\varepsilon} \approx 0$ may still be achieved in principle.

Our results show how there is *no* fundamental limit to arbitrarily accurate purification when coupling the target system to an effectively infinite-dimensional environment. Exact purification, on the other hand, would require a sufficiently large number of eigenvalues of ρ_E to be *precisely zero*. Since this is *not* a generic condition, in particular it cannot be obeyed if ρ_E is thermal, the no-go theorem of Ref. 20 is recovered. With this general conceptual framework in hand, we next proceed to examine in more detail relevant applications, beginning from the special important case of ground-state cooling.

Ground-state cooling given initial system-bath factorization. Consider a setting where, as in Fig. 2, the environment consists of a physical bath ($E \equiv B$), and let H_S denote the (free) Hamiltonian of the target system S, so that the corresponding initial energy is $\mathrm{Tr}(H_S\rho_S)$.

Assume first that the minimum eigenvalue $E_{\min}$ of H_S is not degenerate, in which case exact cooling of S to its ground state entails preparing it in the unique pure state $|\psi_{\mathrm{gs}}\rangle$ corresponding to eigenvalue $E_{\min}$. It is then a straightforward corollary of our theorem that *exact ground-state cooling can be obtained* only if *the environment contains a virtual subsystem of the same dimension of the target, which is initialized in a pure state*. Under the complete joint unitary controllability assumption, however, the ability to prepare a given pure state also imply the ability to prepare any pure state in $\mathcal{H}_S$. Hence, *the existence of a pure virtual subsytem of the environment is also* necessary *for exact cooling*, fully consistent with the conclusions reached in Ref. 20.

On the other hand, suppose that only ε-approximate purification may be achieved in the sense of Eq. (2), so that the state of S can only be cooled down to within distance $\varepsilon > 0$ from the unique ground state $|\psi_{\mathrm{gs}}\rangle$ of H_S. Then the final energy of the system may be estimated as

$$\mathrm{Tr}\left[H_S\rho'_S\right] = \mathrm{Tr}\left[H_S\Big((1-\varepsilon)|\psi_{\mathrm{gs}}\rangle\langle\psi_{\mathrm{gs}}| + \varepsilon\tau_{\mathrm{ex}}\Big)\right]$$
$$\leq (1-\varepsilon)E_{\min} + \varepsilon E_{\max},$$

where τ_{ex} and $E_{\max}$ denote some state in the orthogonal complement to the ground manifold and the maximal eigenvalue of H_S, respectively. Accordingly, approximate ground-state cooling may be attained with an "excess" energy that is upper-bounded by $\varepsilon E_{\max}$. We already observed that when E is infinite-dimensional, ε can in principle be chosen arbitrarily small, albeit not zero. Thus, as soon as one allows for approximate yet arbitrarily good cooling, *the no-go theorem can be effectively evaded*[20].

If $E_{\min}$ has degeneracy $d_{\mathrm{gs}} > 1$, being able to prepare a pure state still suffices for exact ground-state cooling, but is no longer needed. Sufficient and necessary conditions for approximate cooling in a degenerate subspace may be derived using the same reasoning used in establishing necessity of our condition for $\varepsilon > \tilde{\varepsilon}$ – by finding a virtual d_S-dimensional subsystem $\tilde{S}$ such that E is $\tilde{\varepsilon}$-close to initialization in a subspace of dimension d_{gs} in $\mathcal{H}_{\tilde{S}}$.

Arbitrary purification and ground-state cooling with an engineered qubit reservoir. From an open-system quantum-control perspective, our theorem may be used to explicitly characterize what quantum resources may suffice to arbitrarily purify/cool the target system, by coupling it to a suitably engineered environment ($E \equiv A$). Let us focus on the simplest yet paradigmatic case in which S is a single qubit and A consists itself of N qubits, so that $\mathcal{H}_A \simeq \left(\mathbb{C}^2\right)^{\otimes N}$.

Building on the previous discussion, identifying the desired virtual qubit-subsystem entails to split $\mathcal{H}_A$ into two isomorphic, orthogonal subspaces. For the resulting "virtual state" to be approximately pure, we further require the probability for A to be found in one of such subspaces to be much higher than the one for its complementary. A natural approach is to invoke a "typical subspace" argument. Let each auxiliary qubit be prepared in the same state, say, $\rho \equiv \mathrm{diag}(q, 1-q)$, $1/2 \leq q \leq 1$, with respect to a standard basis $\{|0\rangle, |1\rangle\}$, so that the joint initial state $\rho_{SA} \equiv \rho_S \otimes \rho^{\otimes N}$. As N grows, the state of A will populate with increasing probability the ϵ-*typical subspace*. Recall that a sequence $x(N)$ of N zeroes and ones, in which each entry is chosen independently at random with probability $\mathbb{P}(0) = q$, $\mathbb{P}(1) = 1-q$, is ϵ-typical if[12]

$$2^{-N(\mathcal{S}(x)+\epsilon)} \leq \mathbb{P}(x(N)) \leq 2^{-N(\mathcal{S}(x)-\epsilon)},$$

or, equivalently, its total Shannon entropy is ϵ-close to N times the binary entropy of the single symbol. Let $\mathcal{T}(N,\epsilon)$ be the set of ϵ-typical sequences. In the quantum case, such a set naturally generalizes to the ϵ-typical subspace: in our qubit setting, the latter is spanned by those computational basis states that include (approximately) qN zeroes and $(1-q)N$ ones:

$$\mathcal{H}_{\mathcal{T}(N,\epsilon)} \equiv \mathrm{span}\{|x(N)\rangle | x(N) \in \mathcal{T}(N,\epsilon)\}.$$

Let now $\Pi_{\mathcal{T}(N,\epsilon)}$ denote the orthogonal projection onto the ϵ-typical subspace. Then the following asymptotic result holds (see e.g. Theorem 6.3 in Ref. 45):

$$\lim_{N\to\infty} q_{\rm typ}(N) \equiv \lim_{N\to\infty} {\rm Tr}\left(\Pi_{\mathcal{T}(N,\epsilon)}\rho^{\otimes N}\right) = 1. \qquad (9)$$

Furthermore, for any fixed $\epsilon > 0$ and a sufficiently large N, the size of the typical subspace satisfies:

$$\dim\left(\mathcal{H}_{\mathcal{T}(N,\epsilon)}\right) \leq 2^{N(\mathcal{S}(\rho)+\epsilon)}.$$

Hence, if ϵ is sufficiently small, the dimension of the ϵ-typical subspace becomes less or equal than half of the total space dimension as soon as $N\mathcal{S}(\rho) < N-1$, or, $\mathcal{S}(\rho) < 1 - 1/N$. Therefore, provided that the entropy of each of the auxiliary qubits is strictly less than one, namely $\rho \neq \frac{1}{2}I$, the typical subspace's dimension will become less than half of the dimension of $\mathcal{H}_A$ for large enough N. If so, we know how to explicitly construct a unitary transformation W_{SA} that achieves (optimal) $\tilde{\varepsilon}$-purification in principle: it suffices to swap the state of S with the state of a virtual qubit system $\tilde{S}$ that exploits the typical-subspace structure. We further illustrate this strategy by specializing, again, to ground-state cooling.

Example 2.– Assume, similar to Example 1, that the initial state of the target system $\rho_S \equiv {\rm diag}(p_S, 1-p_S)$, with respect to the qubit energy basis, say, $\{|\psi_\ell\rangle\} \equiv \{|\psi_{\rm gs}\rangle, |\psi_{\rm ex}\rangle\}$ and $p_S < q$. The action of the unitary transformation W_{SA} may be explicitly described by introducing a factorized basis $\{|\psi_\ell\rangle \otimes |j(N)\rangle\}$ on $\mathcal{H}_S \otimes \mathcal{H}_A$, where $\{|j(N)\rangle \approx |j_{\rm typ}\rangle, |j_{\rm ntyp}\rangle\}$ in the large-N limit, with $\{|j_{\rm typ}\rangle\}$ and $\{|j_{\rm ntyp}\rangle\}$ denoting orthonormal bases for the typical subspace and its orthogonal complement, respectively. The idea is then to swap $\approx 2^{N\mathcal{S}(\rho)}$ typical basis states which have non-zero probability and are associated to $|\psi_{\rm ex}\rangle$, with $\approx 2^{N\mathcal{S}(\rho)}$ non-typical basis states which are in tensor product to $|\psi_{\rm gs}\rangle$ but are associated to low probability. If we compute the final energy of the system, by using Eq. (9) we obtain ${\rm Tr}[H_S\rho'_S] \approx q_{\rm typ}(N)E_{\rm min} = (1-\varepsilon)E_{\rm min}$, with arbitrarily small $\tilde{\varepsilon}$ (hence ε) as $N \to \infty$, as desired.

Altogether, our results imply that, for a target qubit system, arbitrary accuracy in purification and cooling may be achieved through fully coherent (unitary) interaction with sufficiently many copies of *any auxiliary qubit state which is not the completely mixed one*. It is interesting to notice, however, that repeated interactions with an identically prepared qubit do not suffice in general: the generalized swap operation needs to simultaneously operate on multiple qubits of the engineered environment, pointing to an intrinsic non-Markovian action.

Robust pure-state preparation with finite control iterations. As a final application of our framework, our main theorem may be used to understand and characterize the control resources involved in a stronger form of purification, whereby the goal is to bring the state of S to a *predetermined* target pure state $|\psi\rangle_{\rm target} \in \mathcal{H}_S$, not necessarily related to the system's ground state – so-called "global asymptotic stabilization" in control-theoretic parlance[33,35,39,44]. In particular, the case where $|\psi\rangle_{\rm target}$ is an *entangled* pure state on a multipartite n-qubit target system provides an important quantum-stabilization benchmark. While it is well-known that access to a single *resettable* ancillary qubit A, along with complete unitary control over S and fully coherent "conditional" interactions between A and S, suffices to engineer arbitrary dynamics on S[29,37] (and hence achieve the desired stabilization task) in principle, our result sheds light on the thermodynamical foundation of this result. With reference to the general setting of Fig. 1, suppose for simplicity that no uncontrollable bath is coupled to S ($H_{SB} \equiv 0$), and that B represents the physical degrees of freedom which enact, possibly together with coherent control between S and A, the resetting process on A. Then, in order for stabilization of S to be achievable with arbitrary accuracy ε starting from an *arbitrary* environment state $\rho_E \equiv \rho_B \otimes \rho_A$, an effectively infinite-dimensional environment is necessary. Furthermore, *exact* pure-state stabilization is only achievable provided that A may be perfectly refreshed, which in turn requires B to be perfectly initialized in a pure, two-level virtual subsystem. Remarkably, if these conditions are met, an arbitrary n-qubit pure state $|\psi\rangle_{\rm target}$ may in fact be dissipatively prepared by using a *finite* number, n, of suitably defined control iterations[46].

Experimentally, controlled dissipation mechanisms are becoming available in a growing number of scalable platforms for universal "digital" open-system quantum simulators, including trapped-ion[30,31] and superconducting qubit technologies[47]. In the above-mentioned experiments on $^{40}Ca^{+}$ ions, for instance, the required re-initialization dynamics of the ancilla qubit to a reference pure state was realized through a combination of coherent control on A, in conjunction with optical pumping followed by spontaneous emission. While a number of details are important and require careful consideration in practice, conceptually it is this step that ultimately grants access to virtual subsystems whose states are sufficiently pure, and can thus be swapped with those of the physical degrees of freedom to be purified and/or cooled.

Discussion

We have identified sufficient conditions for purification and ground-state cooling of a quantum system of interest to be achievable in principle, under the two assumptions of *initial system-environment factorization* and *complete unitary controllability*. Such conditions are also necessary in most realistic situations, where the environment is much larger than the target system. While in essence these conditions make rigorous an intuition that is both compelling and natural in retrospect – namely, that purity can only be "swapped" across appropriately defined quantum subsystems – we have shown how these conditions allow to both elucidate fundamental limitations in harnessing open-system dynamics as well as identify new opportunities for control engineering. In particular, our analysis makes it clear that *arbitrarily accurate* purification and/or ground-state cooling is always possible in principle as long as the relevant environment is effectively infinite-dimensional, with a no-go result[20] only emerging in the limiting case of zero error.

From a control-theory standpoint, an interesting direction for further study is to characterize what (more stringent) limitations on quantum purification and cooling may arise upon relaxing the assumption of complete controllability for $S + E$. We envision that the existence of a sufficiently pure virtual subsystem in the environment will still be a necessary and sufficient condition, albeit identification of the relevant subsystem structure will be carried out in this case by exploiting the dynamical-symmetry decomposition associated to the reachable control sub-algebra, in analogy to dynamical error-control strategies and encoded tensoriality in QIP[48,49].

Lastly, it is interesting to comment on our results in relationship to the third law of thermodynamics in its dynamical formulation – the so-called "unattainability principle", namely, the impossibility to cool a system to absolute zero temperature in finite time[11]. Throughout our discussion, we have deliberately made no explicit statement on the time T needed to implement the required generalized swap transformation $W_{SE}(T)$. For a standard thermodynamic setting where the bath is given, and is initially in a generic trace-class state (say, thermal at non-zero temperature), we have showed that arbitrarily small cooling error, $\varepsilon > 0$, may be achieved only if a sufficiently large subspace of the bath can correspondingly account for *less than* ε probability. This, in turn, translates into an increasingly complex (energetically "delocalized") action of the swap transformation $W_{SE}(T)$ to be implemented. Since realistic control Hamiltonians are inevitably *constrained* (e.g., bounded in amplitude and/or speed, as stressed in Refs. 15,16), the limit of perfect accuracy, $\varepsilon \to 0$, can only be approached *asymptotically* in time, $T \to \infty$. While

this supports the validity of the third law under typical conditions, it is our hope that our general subsystem-based approach may prove useful to deepen our understanding of fundamental performance bounds in more complex thermodynamic scenarios, including "quantum-enhanced" refrigeration as recently proposed in Ref. 50.

Methods

Subsystem construction and generalized swap operation. Starting from a general d-dimensional Hilbert space $\mathcal{H}$, a "virtual subsystem structure" as used in the main text can be identified by constructing a basis with the correct tensor/direct sum structure. The main steps may be summarized as follows:

- Identify a $(d_1 \times d_2)$-dimensional subspace $\mathcal{H}_{1,2}$, so that $\mathcal{H} \simeq \mathcal{H}_{1,2} \oplus \mathcal{H}_R$, where $\mathcal{H}_R = \mathcal{H}_{1,2}^{\perp} \equiv \mathcal{H} \ominus \mathcal{H}_{1,2}$.
- Inside such a subspace, choose d_1 mutually-orthogonal subspaces $\mathcal{H}_{j,2}$, each of dimension d_2, so that we may decompose $\mathcal{H}_{1,2} \simeq \oplus_{j=1}^{d_1} \mathcal{H}_{j,2}$.
- Pick an orthonormal basis in each of the summands, say, $\{|\phi^j\rangle_k, k = 1, \ldots, d_2\}$, for $j = 1, \ldots, d_1$. We can then establish the following identification:

$$|\phi^j\rangle_k = |\phi_j\rangle_S \otimes |\phi_k\rangle_F, \quad \mathcal{H}_{1,2} \simeq \mathcal{H}_S \otimes \mathcal{H}_F,$$

and obtain the desired subsystem structure, with $\dim(\mathcal{H}_S) = d_1$, $\dim(\mathcal{H}_F) = d_2$, respectively.

Consider now, specifically, a subsystem structure as given in Eq. (3) on the environment Hilbert space, namely, $\mathcal{H}_E = (\mathcal{H}_{\tilde{S}} \otimes \mathcal{H}_F) \oplus \mathcal{H}_R$, and let $\{|\psi_j\rangle_S\}$, $\{|\phi_k\rangle_{\tilde{S}}\}$, $\{|\xi_\ell\rangle_F\}$, $\{|\chi_m\rangle_R\}$ be orthonormal (ordered) bases for $\mathcal{H}_S$, $\mathcal{H}_{\tilde{S}}$, $\mathcal{H}_F$, $\mathcal{H}_R$, respectively. We may define the required *generalized swap* unitary operator W_{SE} through its action on the element of an orthonormal basis. That is, consider the (ordered) basis of $\mathcal{H}_S \otimes \mathcal{H}_E$ given by:

$$\left\{|\psi_j\rangle_S \otimes |\phi_k\rangle_{\tilde{S}} \otimes |\xi_\ell\rangle_F\right\} \bigcup \left\{|\psi_j\rangle_S \otimes |\chi_m\rangle_R\right\},$$

for all j, k, ℓ, m. The action of W_{SE} is then defined by:

$$\begin{cases} W_{SE}\left(|\psi_j\rangle_S \otimes \left(|\phi_k\rangle_{\tilde{S}} \otimes |\xi_\ell\rangle_F\right)\right) = |\psi_k\rangle_S \otimes |\phi_j\rangle_{\tilde{S}} \otimes |\xi_\ell\rangle_F, \\ W_{SE}\left(|\psi_j\rangle_S \oplus |\chi_m\rangle_R\right) = |\psi_j\rangle_S \oplus |\chi_m\rangle_R. \end{cases}$$

Proof the main theorem. Assume that, as in the main text, we write $d_E = d_S\, d_F + d_R$, with $d_R < d_S \leq d_E$, and let $\rho_{SE} \equiv \rho_S \otimes \rho_E$ denote an arbitrary joint initial state on $\mathcal{H}_S \otimes \mathcal{H}_E$.

Proof of part (ii). The fact that the existence of an ε-pure subsystem in the environment *suffices* for ε-purification ($\varepsilon \geq 0$) has already been proved in the text. We show here that for the case of *exact* purification ($\varepsilon = 0$), the existence of a purely-initialized, d_S-dimensional subsystem in $\mathcal{H}_E$ is indeed also *necessary*.

Recall that exact purification is equivalent to the existence of an orthogonal projector, $\Pi_T = |\psi\rangle\langle\psi|_S \otimes I_E$, such that $\mathrm{Tr}(\Pi_T \rho'_{SE}) = 1$, and that upon defining $\Pi_0 \equiv U_{SE}^{\dagger} \Pi_T U_{SE}$, this also implies that

$$\mathrm{Tr}(\Pi_0 \rho_{SE}) = \mathrm{Tr}(\Pi_0 \rho_S \otimes \rho_E) = 1, \quad \forall \rho_S. \tag{10}$$

This in particular means that the support of $\rho_S \otimes \rho_E$ is included in the range of Π_0. Let us consider d_S specially chosen initial states ρ_S, associated to an orthonormal basis $\{|\phi_j\rangle_S\}$ of $\mathcal{H}_S$, that is,

$$(\rho_S \otimes \rho_E)_j = |\phi_j\rangle\langle\phi_j|_S \otimes \rho_E, \quad j = 1, \ldots, d_S.$$

Since their supports are mutually orthogonal, and each of them has dimension $\mathrm{rank}(\rho_E)$, it follows that:

$$\mathrm{rank}(\Pi_0) \geq \sum_{j=1}^{d_S} \mathrm{rank}\left(|\phi_S^j\rangle\langle\phi_S^j| \otimes \rho_E\right) = d_S\, \mathrm{rank}(\rho_E). \tag{11}$$

On the other hand, since $\mathrm{rank}(\Pi_T) = d_E$, we also have $\mathrm{rank}(\Pi_0) = d_E$. Together with Eq. (11), this implies:

$$\mathrm{rank}(\rho_E) \leq \frac{d_E}{d_S},$$

and hence, being an integer, $\mathrm{rank}(\rho_E) \leq d_F$. Call $\mathcal{H}_1 \equiv \mathrm{supp}(\rho_E) \subset \mathcal{H}_E$, and construct a d_S-dimensional virtual subsystem of $\mathcal{H}_E$ as described above. By construction, ρ_E is purely initialized in the first elements of the basis associated to the d_S-dimensional subsystem $\tilde{\mathcal{H}}_S$, leading to the desired conclusion.

Proof of part (iii). Let us define the following two quantities [see also Eq. (5)]:

$$\tilde{\varepsilon}(\rho_E) \equiv \tilde{\varepsilon} = 1 - \sum_{j=1}^{d_F} \lambda_j(\rho_E),$$

$$\varepsilon_R(\rho_E) \equiv \varepsilon_R = \frac{d_R}{d_S} \lambda_{d_F+1}(\rho_E).$$

We next proceed to show that:

1. A lower bound ε_0 exists for purification of S, independently of the initial state ρ_S;
2. Purification up to $\bar{\varepsilon} = \varepsilon_0 + \varepsilon_R$ is always possible by properly identifying a subsystem in $\mathcal{H}_E$ alone and then swapping it with the target.

1. *Determining* ε_0.– We look for necessary conditions on $\varepsilon > 0$, so that ε-purification of S can be attained at time T by some joint unitary transformation U_{SE}. Again, this means that there exists an orthogonal projector, $\Pi_T = |\psi\rangle\langle\psi|_S \otimes I_E$, such that $\mathrm{Tr}(\Pi_T \rho'_{SE}) \geq 1 - \varepsilon$, for all ρ_S. Upon defining $\Pi_0 \equiv U_{SE}^{\dagger} \Pi_T U_{SE}$ as above, this also implies that

$$\mathrm{Tr}(\Pi_0 \rho_{SE}) = \mathrm{Tr}(\Pi_0 \rho_S \otimes \rho_E) \geq 1 - \varepsilon, \quad \forall \rho_S. \tag{12}$$

Thus, a pure subsystem of dimension d_S may be identified to within ε-distance from the *joint* initial state as well. While Eq. (12) must hold for all ρ_S, in order to determine the desired lower bound we consider a worst-case scenario where $\rho_S = (1/d_S) I_S$ and, correspondingly, the initial joint state $\rho_{SE} \equiv \tilde{\rho}_{SE} = (1/d_S) I_S \otimes \rho_E$.

In fact, consider a basis in which $\tilde{\rho}_{SE}$ is diagonal, ordered in such a way that its eigenvalues are non-increasing. The eigenvalues of $\tilde{\rho}_{SE}$ are the eigenvalues of ρ_E, each multiplied by $(1/d_S)$ and repeated d_S times. Given that Π_0 has rank d_E, the maximal purification achievable in this case correspond to Π_0 projecting on the first d_E eigenvalues. It is then easy to show that:

$$\sum_{k=1}^{d_E} \lambda_k(\tilde{\rho}_{SE}) = \sum_{j=1}^{d_F} \lambda_j(\rho_E) + \frac{d_R}{d_S} \lambda_{d_F+1}(\rho_E). \tag{13}$$

Since, to guarantee ε-purification, the d_E-ranked projector Π_0 must satisfy Eq. (12) in particular for $\rho_{SE} = \tilde{\rho}_{SE}$, Eq. (13) implies the following lower bound ε_0:

$$\varepsilon \geq \varepsilon_0 \equiv \bar{\varepsilon} - \varepsilon_R. \tag{14}$$

We remark that so far nothing guarantees that ε_0-purification is attainable for *any* initial state.

2. *Attaining* $\bar{\varepsilon}$-purification.– From Eq. (13), we infer that there exists a subspace $\mathcal{H}_1$ of $\mathcal{H}_E$ *alone*, with dimension d_F, that accounts for $1 - \bar{\varepsilon} = 1 - \varepsilon_0 - \varepsilon_R$ of the trace of ρ_E. We can thus consider the subspace $\mathcal{H}_1 \subseteq \mathcal{H}_E$ that collects *only* the one-dimensional eigenspaces corresponding to the first d_F eigenvectors of ρ_E. The last step is to start from $\mathcal{H}_1$ to construct a virtual subsystem $\tilde{S}$, such that E is $\bar{\varepsilon}$-close to pure-state initialization in $\tilde{S}$.

We can in fact identify additional $(d_S - 1)$ orthogonal subspaces in $\mathcal{H}_E$, say, $\{\mathcal{H}_j\}_{j=2}^{d_S}$, all isomorphic to $\mathcal{H}_1$ and composed of eigenspaces of ρ_E, so that, by following the general subsystem construction described above, have:

$$\mathcal{H}_E = \left(\overset{d_S}{\underset{j=1}{\oplus}} \mathcal{H}_j\right) \oplus \mathcal{H}_R \simeq (\mathcal{H}_{\tilde{S}} \otimes \mathcal{H}_F) \oplus \mathcal{H}_R,$$

where $\mathcal{H}_{\tilde{S}} \simeq \mathcal{H}_S$, $\dim(\mathcal{H}_F) = d_F$, and $\mathcal{H}_R = \mathcal{H}_E \ominus \oplus_{j=1}^{d_S} \mathcal{H}_j$, $\dim(\mathcal{H}_R) = d_R$. Let Π_1 be the orthogonal projector onto $\mathcal{H}_1$, and define $\tilde{\Pi}_0 \equiv I_S \otimes \Pi_1$. By construction, $\tilde{\Pi}_0$ has rank $d_S d_F \leq d_E$. It thus follows that:

$$\mathrm{Tr}\left(\tilde{\Pi}_0 \tilde{\rho}_{SE}\right) = Tr(\Pi_1 \rho_E) = 1 - \bar{\varepsilon}.$$

Now notice that with respect to the subsystem decomposition above, we may write $\Pi_1 = |\tilde{\varphi}\rangle\langle\tilde{\varphi}| \otimes I_F \oplus 0_R$ for some $|\tilde{\varphi}\rangle \in \mathcal{H}_{\tilde{S}}$, and

$$\rho_E \equiv \tilde{\rho}_E + \Delta\rho_E = |\tilde{\varphi}\rangle\langle\tilde{\varphi}| \otimes \tau_F \oplus 0_R + \Delta\rho_E,$$

with $\tau_F = \frac{1}{1-\bar{\varepsilon}} \mathrm{diag}(\lambda_1(\rho_E), \ldots, \lambda_{d_F}(\rho_E))$. Accordingly, with respect to the decomposition $\mathcal{H}_E = \mathcal{H}_1 \oplus \mathcal{H}_1^{\perp}$, we may write $\Delta\rho_E = \Delta\rho_1 \oplus \Delta\rho_1^{\perp}$, with

$$\Delta\rho_1 = \frac{-\bar{\varepsilon}}{1-\bar{\varepsilon}} \mathrm{diag}(\lambda_1(\rho_E), \ldots, \lambda_{d_F}(\rho_E)),$$

$$\Delta\rho_1^{\perp} = \mathrm{diag}(\lambda_{d_F+1}(\rho_E), \ldots, \lambda_{d_E}(\rho_E)).$$

Since these matrices correspond to the positive and negative-semidefinite part of $\Delta\rho_E$, it follows that

$$\frac{1}{2}\mathrm{Tr}(|\Delta\rho_E|)=\frac{1}{2}\left[-\mathrm{Tr}(\Delta\rho_1)+\mathrm{Tr}\left(\Delta\rho_1^{\perp}\right)\right]=\tilde{\varepsilon}.$$

We may thus conclude that ρ_E admits a $\tilde{\varepsilon}$-pure subsystem, and by using the generalized swapping we can guarantee $\tilde{\varepsilon}$-purification of the target, as claimed.

Note that whenever $d_R = 0$, we have $\tilde{\varepsilon} = \varepsilon_0$ and thus our generalized swap operator attains the best possible purification. If, additionally, $d_E = \infty$, this also formally corresponds to $d_F = \infty$ hence $\tilde{\varepsilon} = 0$. We have then explicitly shown in the main text how to achieve purification up to arbitrary finite accuracy $\varepsilon > 0$.

1. Ketterle, W. & Pritchard, D. E. Atom cooling by time-dependent potentials. *Phys. Rev. A* **46**, 4051–4054 (1992).
2. Bartana, A., Kosloff, R. & Tannor, D. J. Laser cooling of molecular internal degrees of freedom by a series of shaped pulses. *J. Chem. Phys.* **99**, 196–210 (1993).
3. Cohen-Tanoudji, C. & Guéry-Odelin, D. *Advances in Atomic Physics* (World Scientific, Singapore, 2011).
4. Abragam, A. & Goldman, M. Principles of dynamic nuclear polarisation. *Rep. Progr. Phys.* **41**, 395–467 (1978).
5. Maly, T. *et al.* Dynamic nuclear polarization at high magnetic fields. *J. Chem. Phys.* **128**, 052211/1–20 (2008).
6. Mancini, S., Vitali, D. & Tombesi, P. Optomechanical cooling of a macroscopic oscillator by homodyne feedback. *Phys. Rev. Lett.* **80**, 688–691 (1998).
7. Hopkins, A., Jacobs, K., Habib, S. & Schwab, K. Feedback cooling of a nanomechanical resonator. *Phys. Rev. B* **68**, 235328/1–10 (2003).
8. LaHaye, M. D., Buu, O., Camarota, B. & Schwab, K. C. Approaching the quantum limit of a nanomechanical resonator. *Science* **304**, 74–77 (2004).
9. O'Connell, A. D. *et al.* Quantum ground state and single-phonon control of a mechanical resonator. *Nature* **464**, 697–703 (2010).
10. Safavi-Naeini, A. H. *et al.* Observation of quantum motion of a nanomechanical resonator. *Phys. Rev. Lett.* **108**, 033602/1–5 (2012).
11. Kosloff, R. Quantum thermodynamics: a dynamical viewpoint. *Entropy* **15**, 2100–2128 (2013).
12. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Information* (Cambridge University Press, Cambridge, 2002).
13. Knill, E. & Laflamme, R. On the power of one bit of quantum information. *Phys. Rev. Lett.* **81**, 5672–5675 (1998).
14. Lidar, D. A. & (Eds.), T. A. B. *Quantum Error Correction* (Cambridge University Press, Cambridge, 2013).
15. Wang, X., Vinjanampathy, S., Strauch, F. W. & Jacobs, K. Absolute dynamical limit to cooling weakly coupled quantum systems. *Phys. Rev. Lett.* **110**, 157207/1–5 (2013).
16. Horowitz, J. M. & Jacobs, K. A quantum advantage in the thermodynamic efficiency of feedback control. *arXiv*,1311.2920 (2013).
17. Hamerly, R. & Mabuchi, H. Advantages of coherent feedback for cooling quantum oscillators. *Phys. Rev. Lett.* **109**, 173602/1–5 (2012).
18. Allahverdyan, A. A., Hovhannisyan, K. V., Janzing, D. & Mahler, G. Thermodynamic limits of dynamic cooling. *Phys. Rev. A* **84**, 041109/1–16 (2011).
19. Franco, C. D. & Paternostro, M. A no-go result on the purification of quantum states. *Sci. Rep.* **3**, 1387; DOI:10.1038/srep01387 (2013).
20. Wu, L.-A., Segal, D. & Brumer, P. No-go theorem for ground state cooling given initial system-thermal bath factorization. *Sci. Rep.* **3**, 1824; DOI:10.1038/srep01824 (2012).
21. Reeb, D. & Wolf, M. M. (Im-)proving Landauer's principle. *arXiv*,1306.4352 (2013).
22. Knill, E., Laflamme, R. & Viola, L. Theory of quantum error correction for general noise. *Phys. Rev. Lett.* **84**, 2525–2528 (2000).
23. Zanardi, P. Virtual quantum subsystems. *Phys. Rev. Lett.* **87**, 077901/1–4 (2001).
24. Viola, L., Knill, E. & Laflamme, R. Constructing qubit in physical systems. *J. Phys. A* **34**, 7067–7079 (2001).
25. Knill, E. On protected realization of quantum information. *Phys. Rev. A* **74**, 042301/1–11 (2006).
26. Ticozzi, F. & Viola, L. Quantum information encoding, protection and correction via trace-norm isometries. *Phys. Rev. A* **81**, 032313/1–9 (2010).
27. Blume-Kohout, R., Ng, H. K., Poulin, D. & Viola, L. The structure of preserved information in quantum processes. *Phys. Rev. Lett.* **100**, 030501/1–4 (2008).
28. Blume-Kohout, R., Ng, H. K., Poulin, D. & Viola, L. Information preserving structures: A general framework for quantum zero-error information. *Phys. Rev. A* **82**, 062306/1–25 (2010).
29. Lloyd, S. & Viola, L. Engineering quantum dynamics. *Phys. Rev. A* **65**, 010101/1–4 (2001).
30. Barreiro, J. T. *et al.* An open-system quantum simulator with trapped ions. *Nature* **470**, 486–491 (2011).
31. Schindler, P. *et al.* Quantum simulation of open-system dynamical maps with trapped ions. *Nature Phys.* **9**, 361–367 (2013).
32. Poyatos, J. F., Cirac, J. I. & Zoller, P. Quantum reservoir engineering with laser cooled trapped ions. *Phys. Rev. Lett.* **77**, 4728–4731 (1996).
33. Ticozzi, F. & Viola, L. Steady-state entanglement by engineered quasi-local Markovian dissipation. *Quantum Inf. Comput.* **14**, 0265–0294 (2014).
34. Ghirardi, G. C., Rimini, A. & Weber, T. Unified dynamics for microscopic and macroscopic systems. *Phys. Rev. D* **34**, 470–491 (1986).
35. D'Alessandro, D. *Introduction to Quantum Control and Dynamics.* Applied Mathematics & Nonlinear Science (Chapman & Hall/CRC, 2007).
36. Viola, L., Knill, E. & Lloyd, S. Dynamical decoupling of open quantum systems. *Phys. Rev. Lett.* **82**, 2417–2421 (1999).
37. Lloyd, S. Coherent quantum feedback. *Phys. Rev. A* **62**, 022108/1–12 (2000).
38. D'Alessandro, D. Equivalence between indirect controllability and complete controllability for quantum systems. *Syst. Control Lett.* **62**, 188–193 (2013).
39. Altafini, C. & Ticozzi, F. Modeling and control of quantum systems: an introduction. *IEEE Trans. Aut. Control* **57**, 1898–1917 (2012).
40. Bloch, A. M., Brockett, R. W. & Rangan, C. Finite controllability of infinite-dimensional quantum systems. *IEEE Trans. Aut. Control* **55**, 1797–1805 (2010).
41. Boscain, U., Gauthier, J.-P., Rossi, F. & Sigalotti, M. Approximate controllability, exact controllability, and conical eigenvalue intersections for quantum mechanical systems. *arXiv:1309.1970.*
42. Knill, E. & Laflamme, R. Theory of quantum error-correcting codes. *Phys. Rev. A* **55**, 900–911 (1997).
43. Shabani, A. & Lidar, D. A. Theory of initialization-free decoherence-free subspaces and subsystems. *Phys. Rev. A* **72**, 042303/1–14 (2005).
44. Ticozzi, F. & Viola, L. Quantum Markovian subsystems: invariance, attractivity and control. *IEEE Trans. Aut. Contr.* **53**, 2048–2063 (2008).
45. Petz, D. *Quantum Information Theory and Quantum Statistics* (Springer Verlag, Berlin, 2008).
46. Baggio, G., Ticozzi, F. & Viola, L. State preparation by controlled dissipation in finite time: From classical to quantum controllers. In *Proc. of the 51st IEEE Conf. on Decision and Control*, 1072–1077 (2012).
47. Shankar, S. *et al.* Autonomously stabilized entanglement between two superconducting quantum bits. *Nature* **504**, 419–422 (2013).
48. Viola, L., Knill, E. & Lloyd, S. Dynamical generation of noiseless quantum subsystems. *Phys. Rev. Lett.* **85**, 3520–3523 (2000).
49. Zanardi, P., Lidar, D. A. & Lloyd, S. Quantum tensor product structures are observable induced. *Phys. Rev. Lett.* **92**, 060402/1–4 (2004).
50. Correa, L. A., Palao, J. P., Alonso, D. & Adesso, G. Quantum-enhanced absorption refrigerators. *Sci. Rep.* **4**, 3949; DOI:10.1038/srep03949 (2014).

Acknowledgments

It is a pleasure to thank David Reeb for bringing Ref. 21 to our attention and for pointing out a technical problem with an early version of the manuscript, as well as Peter Johnson and Alireza Seif for a critical reading of the manuscript. Work at Dartmouth was supported in part by the US ARO under contract No. W911NF-11-1-0068 and the Constance and Walter Burke Special Project Fund in *Quantum Information Science*. F.T. acknowledges partial support from the QUINTET and QFUTURE projects of the University of Padua.

Author contributions

F.T. and L.V. jointly contributed to the concept, execution, and interpretation of the work, and preparation of the manuscript.

Additional information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Ticozzi, F. & Viola, L. Quantum resources for purification and cooling: fundamental limits and opportunities. *Sci. Rep.* **4**, 5192; DOI:10.1038/srep05192 (2014).