

Energy-Consumption Advantage of Quantum Computation

Florian Meier^{1,*} and Hayata Yamasaki^{2,†}

¹ Vienna Center for Quantum Science and Technology,
Atominstitut, Technische Universität Wien, 1020 Vienna, Austria

² Department of Physics, Graduate School of Science,
The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan

Energy consumption in solving computational problems has been gaining growing attention as one of the key performance measures for computers. Quantum computation is known to offer advantages over classical computation in terms of various computational resources; however, proving its energy-consumption advantage has been challenging due to the lack of a theoretical foundation linking the physical concept of energy with the computer-scientific notion of complexity for quantum computation. To bridge this gap, we introduce a general framework for studying the energy consumption of quantum and classical computation, based on a computational model conventionally used for studying query complexity in computational complexity theory. Within this framework, we derive an upper bound for the achievable energy consumption of quantum computation, accounting for imperfections in implementation appearing in practice. As part of this analysis, we construct a protocol for Landauer erasure with finite precision in a finite number of steps, which constitutes a contribution of independent interest. Additionally, we develop techniques for proving a nonzero lower bound of energy consumption of classical computation, based on the energy-conservation law and Landauer's principle. Using these general bounds, we rigorously prove that quantum computation achieves an exponential energy-consumption advantage over classical computation for solving a paradigmatic computational problem—Simon's problem. Furthermore, we propose explicit criteria for experimentally demonstrating this energy-consumption advantage of quantum computation, analogous to the experimental demonstrations of quantum computational supremacy. These results establish a foundational framework and techniques to explore the energy consumption of computation, opening an alternative way to study the advantages of quantum computation.

I. INTRODUCTION

With growing interest in the sustainability of our society, energy consumption is nowadays considered an important part of performance measures for benchmarking computers. It is expected that quantum computation will be no exception; its energy efficiency will ultimately be one of the deciding factors as to whether quantum computers will be used on a large scale [1–3]. Originally, quantum computers emerged as a promising platform to solve certain computational problems that would otherwise be unfeasible to solve on classical computers. The advantage of quantum computation is generally examined in terms of computational complexity, which quantifies computational resources required for solving the problems, such as time complexity, communication complexity, and query complexity [4]. Energy is, however, a different computational resource from the above ones. A priori, an advantage in some computational resource does not necessarily imply that in another; for example, quantum computation is believed to achieve an exponential advantage in time complexity over classical computation but does not provide such an advantage in the required amount of memory space [5]. Whether quantum computation can offer a significant energy-consumption advantage over classical computation is a fundamental question

that has been asked since the earliest days of quantum computation [6, 7], but has not been explored rigorously as of now due to the lack of theoretical foundation to relate the known advantages of quantum computation and the advantage in its energy consumption. For example, some existing works [3, 8] numerically evaluate achievable upper bounds of energy consumption of quantum computation, but it is not theoretically proven in these works how the energy consumption of quantum computation scales on large scales in general; even more problematically, it has been challenging to provide a fundamental lower bound of the energy consumption of any classical computation, which the existing works do not rigorously analyze. It has been, therefore, a fundamental open problem to formulate the theoretical framework to compare an upper bound of energy consumption of quantum computation and a lower bound of energy consumption of classical computation fairly in a technology-independent way, so as to prove the energy-consumption advantage of quantum computation over classical computation in a commonly studied setting of computation [2].

Main achievements of our work. In this work, we carry out a rigorous, in-depth study of the energy consumption of quantum and classical computation. Through our work, we introduce a general framework to analyze energy consumption in both quantum and classical computation, which incorporates computational and thermodynamic aspects, as illustrated in Fig. 1.

In Fig. 1, the computation is modeled as a thermodynamic cycle. This model is composed of five parts: the computer, agent, environment, input, and output.

* florianmeier256@gmail.com

† hayata.yamasaki@gmail.com

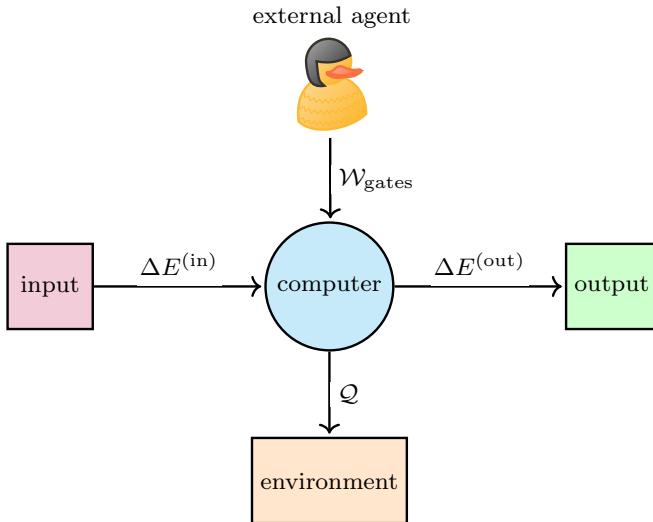


FIG. 1. Our thermodynamic model of quantum and classical computation. To solve a computational problem, an external agent manipulates the computer at a work cost W_{gates} . In the computation, the input for the problem is provided from an input system to the computer, with energy $\Delta E^{(\text{in})}$ flowing into the computer. At the end of the computation, the computer outputs the result of the computation into the output system, with energy $\Delta E^{(\text{out})}$ flowing out; then, the internal state of the computer must be reinitialized into a fixed state $\approx |00\cdots 0\rangle$, so that the computer can be used for solving the next computational task. Throughout manipulating and initializing the computer, energy in the computer may flow into the environment as a heat Q . As explained in the main text, the energy consumption $W_{\text{gates}} + \Delta E^{(\text{in})} - \Delta E^{(\text{out})}$ of the computation is classified into three contributions: the energetic cost, the control cost, and the initialization cost. Our analysis provides a comprehensive account of each contribution to bound the overall energy consumption of quantum and classical computation.

The computer, either quantum or classical, is the physical platform with (qu)bits, where an algorithm is carried out. To execute the algorithm, an external agent invests some work W_{gates} into the computer to perform a sequence of gates on the (qu)bits in the computer. Part of this work is an *energetic cost* used for changing the energy of the internal state of the computer, while the rest is the *control cost* dissipated into the environment due to energy loss arising from, e.g., friction and electric resistance. Moreover, as in a conventional setting of studying query complexity to separate the complexities in quantum and classical computation [4, 9, 10], we structure the algorithm to process the input coming from an oracle — a black-box circuit outside the computer that is unknown to the agent, so that the agent should perform the computation to learn a property of this black box through the queries to the oracle. For this input, we use an additional input system representing the oracle and label the energy of the states transferred from the input system to the computer as $\Delta E^{(\text{in})}$. The result of the entire computation is to be output and stored in

another separate system, the output system. In our case, we solve a decision problem, so the output system should be a single bit. The energy of the state output from the computer incurs an energetic cost $\Delta E^{(\text{out})}$. At the end of the computation, the final state of the computer is reinitialized into its initial state, so that, afterward, the agent should be able to carry out the next computation by the same computer. Aside from the control cost, irreversible operations such as the erasure of the final state of the computer for the initialization also dissipate heat into the environment, which incurs an *initialization cost*. All-in-all, we summarize the heat dissipated into the environment as Q . The energy consumption of the computation is all the energy used for the above process except for that dissipating into the environment, given by

$$W_{\text{gates}} + \Delta E^{(\text{in})} - \Delta E^{(\text{out})}. \quad (1)$$

Our methodology pioneers general techniques to clarify the upper and lower bounds of the energy consumption for the computation in the query-complexity setting. The upper bound of achievable energy consumption of quantum algorithms is obtained by meticulously considering the work cost of implementing gates to finite precision, erasing the computer's memory state for the initialization based on Landauer's principle [11], and reducing the effect of noise by quantum error correction. Under the scalability assumption that each elementary operation has, at most, a polynomial amount of energy consumption, we derive an upper bound of energy consumption of quantum computation by bounding the required number of elementary operations for solving computational problems, including quantum error correction. Also notably, we develop a novel technique for deriving a fundamental lower bound of the energy required for classical algorithms to solve computational problems. In contrast with the upper bound, it is, in general, unknown whether one can find a nonzero lower bound of the energy consumption of each elementary operation; after all, in the idealized limit of zero friction or zero electrical resistance, each elementary operation may consume almost zero energy. To overcome this challenge, we will develop the technique to argue that a lower bound of the query complexity in solving the computational problems can be turned into a lower bound of the Landauer cost, i.e., the energy consumption to erase the information obtained by the queries.

Based on these bounds, we rigorously prove that a quantum algorithm can attain an exponential energy-consumption advantage over any classical algorithm for an exemplary computational problem—Simon's problem [12, 13]. Furthermore, for this problem, we clarify how to implement the quantum algorithm using a cryptographic primitive to feasibly implement the input system in Fig. 1 and provide explicit criteria for demonstrating the energy-consumption advantage of quantum computation in experiments. In this way, we expand the traditional focus from conventional complexity-theoretic notions in theoretical computer science to en-

ergy consumption—a critical aspect of computational resources that was often overlooked in the existing theoretical studies. These results lay a solid foundation for the exploration of the advantage of quantum computation over classical computation in terms of energy consumption.

Key physical insights. In our analysis of energy consumption, we develop a comprehensive account of the above three types of costs that contribute to energy consumption: the energetic cost, the control cost, and the initialization cost.

The energetic cost to change the energy of the state of the computer would be hard to evaluate at each step of the algorithm, but to avoid this hardness, our analysis evaluates the sum of the energetic cost of the overall computation rather than each step. The hardness arises from the fact that an energetic cost per gate can be both positive and negative depending on the state of the computer; after all, when performing a gate in the algorithm, the same gate may change a low-energy state to a high-energy state and also a high-energy state to a low-energy state. Thus, if one wants to know the required energetic cost at each step of the algorithm, the internal state of the computer at the step should also be kept track of, which would be as hard as simulating each step of the algorithm. By contrast, in our analysis, instead of focusing on each step, we regard the overall computation as a thermodynamic cycle as described in Fig. 1. As a result, the net energetic cost of all the steps sums up to zero due to the energy-conservation law, in both quantum and classical computation, leading to the total energy consumption of the computation coming solely from heat dissipation due to inefficient control and irreversible erasure.

The control cost, arising from the energy loss in implementing each gate, also needs careful analysis. On one hand, the quantum computer is often said to consume a large amount of energy due to the complicated experimental setup under the current technology and, in this regard, might be considered to be more energy-consuming than the classical computer. But this argument overlooks the fact that the energy loss in any scalable implementation of each quantum gate is reasonably upper bounded by, at most, a polynomial factor per physical gate (and usually upper bounded by a constant factor) as the size of problems to be solved by the computer increases; that is, the apparent difference in the energy efficiencies between quantum and classical computers is that of the cost per implementing each physical gate. With this observation, we derive an achievable upper bound of the control cost of quantum computation from the number of gates used for performing the overall algorithm. On the other hand, as opposed to the above saying, it is also often said that the quantum computer could save energy compared to the classical computer due to the large quantum speedup in solving some problems. This logic also overlooks another fact that the control cost per implementing a classical gate can be made extremely small compared to that of a quantum gate. Due to this fact, the quantum advan-

tage in time complexity does not straightforwardly lead to that of energy consumption over classical computation, especially in the case where the control cost of each classical gate is negligibly small; for example, imagine a limit of idealization where the friction and the electric resistance are negligible, and then, no energy loss occurs regardless of the number of gates to be applied. For example, it is indeed proposed to use superconducting materials for classical computation, which may allow for nearly dissipation-free classical computation in the future (see, e.g., Ref. [14]). Thus, unlike the upper bound of the control cost, a lower bound of the control cost can be, in principle, arbitrarily close to zero even if the runtime of the computation is long.

Initialization cost, i.e., the required energy for resetting the state of the computer to a fixed initial state, is the key to our proof of the energy-consumption advantage of quantum computation over classical computation; this cost can provide a fundamentally *nonzero* lower bound of the energy consumption of classical computation, unlike the above energetic and control costs. As described in Fig. 1, after solving a computational problem, the (quantum or classical) computer needs to reset all the (qu)bits to conduct the next task, by erasing the final state of the (qu)bits into a fixed initial state, i.e., $|00\cdots 0\rangle$. However, in the thermodynamic analysis of the energy consumption, we need to be careful about the cost of this initialization due to Nernst's unattainability principle [15]; it is known that the ideal initialization of a qubit exactly into a pure state would require divergent resource costs, e.g., infinite energy consumption or infinitely long runtime [16]. To avoid an unbounded energy consumption in the initialization, we newly construct a finite-step Landauer-erasure protocol for erasing the states of the qubits to a pure state approximately within a desired finite precision, which allows us to achieve a finite energy consumption for all the erasure steps required for the initialization after conducting the computation in Fig. 1.

Apart from this explicit protocol construction for showing an achievable upper bound of energy consumption, we discover a way to prove a nonzero lower bound of the initialization cost by applying two fundamental concepts in physics: the energy-conservation law in thermodynamics, and Landauer's principle [11] that relates the irreversible process to energy consumption. Landauer's principle shows that we need more energy consumption to initialize more information stored as states of (qu)bits, which is measured by the entropy of the states. In particular, to solve Simon's problem [12, 13], we see that any classical algorithm needs to make an exponentially large number of queries to obtain an exponentially long bit string obtained from the input system. In this case, our analysis shows that the initialization cost required for erasing this input information, in terms of entropy, also becomes exponentially large, which is one of our key techniques for proving the exponential energy-consumption advantage of quantum computation.

Consequently, with our formulation of the computation

as the thermodynamic cycle in Fig. 1, we provide an unprecedented detailed analysis of all three types of costs, establishing the novel techniques for showing upper and lower bounds of the energy consumption of quantum and classical computation. Our analysis shows that the upper bound on the energy consumption of quantum computation is governed by the gate and query complexities and the heat dissipation achievable through finite-step Landauer erasure. Therefore, to ensure an upper bound of energy consumption within this framework, it suffices to feasibly bound each of these factors. In contrast, the provably nonzero lower bound on the energy consumption of classical computation arises from the minimum cost of Landauer erasure, quantified by entropy. Although we analyze this lower bound using techniques originally developed for studying query complexity, the lower bound is ultimately determined by the Landauer cost—the entropic quantity that depends both on the query complexity and the probability distribution of the oracle. These techniques make it possible to rigorously analyze the energy-consumption advantage of quantum computation with a solid theoretical foundation.

Impact. In today’s world, where sustainability is a key concern, energy efficiency has become an essential performance metric for modern computers, along with computational speed and memory consumption. Our theoretical framework provides a new pathway for examining the relevance of quantum computation in terms of energy efficiency. Our study provides a fundamental framework and techniques for exploring a novel quantum advantage in computation in terms of energy consumption. The framework is designed in such a way that a quantum advantage in query complexity can be employed to prove the energy-consumption advantage of quantum computation over classical computation. These results also clarify the physical meaning of the quantum advantage in the query-complexity setting in terms of energy consumption.

The query-complexity setting is a widely studied setting of computation by itself, but also, from a broader perspective, quantum computers may have promising applications in learning properties of physical dynamics described by an unknown quantum-mechanical map [17–22]. In these learning problems, the physical dynamics can be considered a black-box oracle whose properties are to be learned, as in the query-complexity setting. In other words, the query-complexity setting can also be regarded as a variant of these learning problems of unknown physical dynamics represented by the oracles. Furthermore, we also clarify how to implement and demonstrate the query-complexity setting in an experimental setup. The framework and techniques developed in this work serve as a theoretical foundation to realize the energy-consumption advantage of quantum computation in query-complexity settings with these physically well-motivated applications.

Organization of this article. The rest of this article is organized as follows.

- In Sec. II, we formulate a framework within which the energy consumption of quantum and classical computation can be rigorously studied. We will describe its core idea based on Fig. 1, and further detail of the framework will be illustrated in Fig. 3.
- In Sec. III, we show general upper and lower bounds of quantum and classical computation within the framework of Sec. II. In Sec. IIIA, we derive a general upper bound on the energy consumption that is achievable by the quantum computation in our framework. For this analysis, extending upon the existing asymptotic results [16, 23–26] on Landauer erasure [11], we show the achievability of finite-fidelity and finite-step Landauer erasure (Theorem 2). Apart from the finite Landauer erasure, we derive the upper bound of achievable energy consumption using complexity-theoretic considerations and also taking into account overheads from quantum error correction, so as to cover all the energy consumption in the framework (Theorem 4). Moreover, in Sec. IIIB, we develop a technique for obtaining an implementation-independent lower bound on the energy consumption of the classical computation in our framework. The lower bound is derived using energy conservation and the Landauer-erasure bound (Theorem 5).
- In Sec. IV, we show an explicit example where the energy consumption of quantum and classical computation for solving a computational problem is exponentially separated. To prove this exponential energy-consumption advantage of quantum computation rigorously, we apply the above upper and lower bounds to Simon’s problem [12, 13] (Corollaries 7 and 9).
- In Sec. V, we propose a setting to demonstrate the energy-consumption advantage of quantum computation in quantum experiments.
- In Sec. VI, we provide a discussion and outlook based on our findings.

II. FORMULATION OF COMPUTATIONAL FRAMEWORK AND ENERGY CONSUMPTION

In this section, we present the setting of our analysis. We start with introducing the thermodynamic model of the computation in Sec. IIA. In Sec. IIB, we formulate and elaborate on the framework used for analyzing the energy consumption of the computation (Fig. 3). In Sec. IIC, we introduce all the relevant costs that contribute to the energy consumption of the computation in this framework.

A. Thermodynamic model of computation

Computation is a physical process in a closed physical system that receives a given input and returns the corresponding output of a mathematical function. Various models can be considered as the physical processes for the computation. Conventionally, classical computation uses the Turing machine, or equivalently a classical logic circuit, which works based on classical mechanics; quantum computation can be modeled by a quantum circuit based on the law of quantum mechanics [4, 9]. Under restrictions on some computational resources, the differences between classical and quantum computation may appear. For example, in terms of running time, quantum computation is considered to be advantageous over classical computation [27]. Another type of quantum advantage can be shown by query complexity [10], where we have a black-box function, or an *oracle*, that can be called multiple times during executing the computation. In this query-complexity setting, the required number of queries to the oracle for achieving the computational task is counted as a computational resource of interest. This setting may concern the fundamental structure of computation rather than space and time resources. The advantage of quantum computation over classical computation can also be seen in the required number of queries for solving a computational task.

Progressing beyond the studies of these various computational resources, we here formulate and analyze the fundamental requirements of energy consumption in classical and quantum computation, by introducing the model illustrated in Fig. 1 for our thermodynamic analysis of computation. To clarify the motivation for introducing this model, we start this section by going through the main difficulties in proving a potential advantage in the energy consumption of quantum computation over classical computation. Then, in the next step, we will provide additional details to Fig. 1 and formally define what we refer to as the energy consumption of a computation.

Challenges in studying energy-consumption advantage of quantum computation. To analyze the energy-consumption advantage of quantum computation over classical computation, there are two inequalities that have to be shown: first, an achievable upper bound for the energy consumption of performing a quantum algorithm, and second, a lower bound on the energy consumption for *all* classical algorithms for solving the same problem.

As for the former, previous works mostly investigate energy consumption in implementing a single quantum gate [28–35], but the analysis of the quantum advantage requires an upper bound of energy consumption of the overall quantum computation composed of many quantum operations. Such an analysis has been challenging because, to account for the energy consumption of quantum computation, we need to take into account all the operations included in the computation, e.g., not only the gates but also the cost of initializing qubits and performing

quantum error correction. A challenge here arises since, unlike quantum gates, some other quantum operations, such as measurements and initialization, may consume an infinitely large amount of energy as we increase the accuracy in their implementation [16, 36]. Thus, we need to formulate the framework of quantum computation properly to avoid the operations requiring infinite energy consumption and to establish finite achievability results of energy consumption for all the operations used for the quantum computation within the framework.

The latter is even more challenging since it is, in general, hard to derive a lower bound of energy consumption of overall computation from the lower bound of each individual operation; after all, we may be able to perform multiple operations in the computation collectively to save energy consumption. In the first place, the energetic cost of performing reversible operations, be it on a quantum or classical computer, can be either positive or negative. To see this fact, consider the following illustrative example: what is the lower bound of performing a bit-flip operation? The answer may depend on the implementation and the initial state. Let us choose, for example, a two-level system with states 0 and 1 and energies $E_0 < E_1$. If the bit initially starts in the state 0, then we need to invest a positive amount of energy $E_1 - E_0 > 0$ to obtain 1 by the bit-flip gate. By contrast, at best, one can also gain energy $E_1 - E_0$ from the bit-flip gate, if the bit initially starts in the state 1 and is flipped into 0; in other words, the energetic cost of performing a reversible operation can be negative in general. Thus, for the initial input state 1, the operation could even extract work because energy-efficient gate implementations are possible in principle, especially for macroscopic physical systems.¹ Apart from this energetic cost, the implementation of the bit-flip gate may also require an additional control cost arising from heat dissipation, which is caused by, e.g., friction and electrical resistance. The control cost may be positive in reality, but in the limit of energy-efficient implementation, it is hard to rule out the possibility that the control cost may be negligibly small; in particular, it is unknown whether the infimum of the control cost over any possible implementation of computation can still be lower bounded by a strictly positive constant gapped away from zero.

As far as the energetic and control costs are concerned, it is challenging to rule out the possibility that an energy-efficient implementation would achieve zero work cost (or

¹ Two concrete examples would be the Japanese *Soroban* and the *Abacus*. Suppose that these devices are standing upright in the earth's gravitational field, and the external agent keeps the locations of the marbles of these devices fixed by some means. The locations of the marbles are used for representing states of bits, and bit-flip gates can be applied by physically moving the marbles. In this case, the information-bearing degrees of freedom carry potential energy. Moving up a marble requires an investment of work. On the other hand, moving down the marble yields a gain of energy, i.e., the extraction of the work.

arbitrarily small work cost) per gate even in the worst case over all possible input states. For example, in the limit of realizing 0 and 1 with degenerate energy levels $E_0 = E_1$, the bit-flip gate does not come at a fundamentally positive energetic cost. Another example may exist in the limit of ideal implementation where friction, electrical resistance, and other sources of energy loss are negligible, and if this is the case, the control cost also becomes negligibly small. As long as one considers these costs of an individual gate, it is hard to argue that each gate requires a positive work cost. Thus, the analysis requires a novel technique for deriving a nonzero lower bound of energy consumption of the overall computation, which needs to be applied independently of the detail of the implementation; in our case, we derive such a nonzero lower bound based on taking into account the initialization cost, energy consumption for erasing the state of the computer at the end of the computation into the fixed initial state 0.

Model. In particular, to establish the framework for studying energy consumption, we view quantum and classical computation as thermodynamic processes on their respective physical implementation platforms, as shown in Fig. 1. We here describe our model in detail. The agent uses a computer to carry out the main part of the computation to solve a decision problem. Performing the operations that make up an algorithm comes at a work cost for this external agent, which we summarize as $\mathcal{W}_{\text{gates}}$. This cost includes, on the one hand, energetic costs arising from the change of energy of the internal states of the computer and, on the other hand, control costs caused by energetic losses in implementing these operations due to heat dissipation.

As in a conventional setting of studying query complexity to separate the complexity of quantum and classical computation [4, 9, 10], we use an oracle as an input model, which provides the input to the computer via oracle queries. In our model, the oracle is a black box outside the computer whose internals are unknown to the agent performing the computation, as we will define more precisely in Sec. II B. There is some transfer of energy between the oracle and the computer, namely $\Delta E^{(\text{in})}$, due to the states the oracle generates and inputs into the computer.

On the other hand, working on decision problems means that the output is a two-level system where the decision is stored as either 0 or 1. Generating this single-bit output comes at an energy exchange of $\Delta E^{(\text{out})} = O(1)$ and is practically negligible compared to the growing sizes of the computer, the input, and the thermal environment surrounding around the computer, which may usually scale polynomially (or can scale even exponentially in our framework) as the problem size increases.

Throughout the computation, the energy corresponding to the control cost may flow into the environment, which in parts contributes to the heat \mathcal{Q} dissipated to the environment. The other contribution to \mathcal{Q} arises from reinitializing the internal state of the computer at the end

of the computation into the fixed initial state as it was before performing the computation. The resetting of the memory after the computation is necessary because we demand that the computation is cyclic on the computer so that this computer can be used for solving another task after conducting the current computation. In contrast to the control cost that may approach zero in the limit of energy-efficient implementation, the energy consumption in reinitializing the computer may be nonzero due to Landauer's principle [11] if part of the computation is irreversible, as is the case for the oracle queries in classical computation.

With this model, we define the *energy consumption* \mathcal{W} of the computer as the sum of all the contributions of the external agent, the input, and the output in Fig. 1, as presented in the following definition. The energetic contributions to \mathcal{W} will be detailed further for our framework (Fig. 3) introduced in Sec. II C.

Definition 1 (Energy consumption). *We define the energy consumption of computation modeled in Fig. 1 as*

$$\mathcal{W} := \mathcal{W}_{\text{gates}} + \Delta E^{(\text{in})} - \Delta E^{(\text{out})}. \quad (2)$$

The demand of the computation being cyclic is critical to energy conservation; i.e., in the limit of closing this thermodynamic cycle ideally without error², the energy conservation demands

$$\mathcal{W} = \mathcal{Q}. \quad (3)$$

The energy consumption \mathcal{W} is the sum of the non-dissipative energy exchanges between external systems and the computer, which can be understood as the work cost of performing the computation. On one hand, we will analyze all the contributions to \mathcal{W} to identify an upper bound of energy consumption of quantum computation. On the other hand, the formulation (3) of energy conservation will be the basis for proving the implementation-independent lower bound on the energy consumption of classical computation by analyzing \mathcal{Q} .

² One might think that with a nonzero error ε in reinitialization after a round of computation, it would be possible that the final state after the reinitialization might have lower energy than the initial state of the computer's memory since the initial state in our framework is not assumed to be a ground state; thus, one could extract energy using the ε -difference between the initial and final states after one round of the computation. This reasoning, however, does not capture the fact that this energy is borrowed from the computer's memory only temporarily in this single round. After all, the initial state of the computer in the second round of the computation is then no longer exactly equal to the initial state of the first round, but also, in this second round, we still need to close the thermodynamic cycle by approximately reinitializing the computer's memory state. When averaging over many rounds of the computation using the same computer's memory, the energy extraction must perish. To capture this, our analysis deals with the energy conservation on average in the limit of $\varepsilon \rightarrow 0$.

B. Computational framework

Based on the model in the query-complexity setting (Fig. 1) that we conceptualized in the previous section, we here proceed to describe the explicit computational framework that represents the computation in terms of (quantum and classical) circuits, so that we can analyze all the contributions to the energy consumption therein. The significance of the computational framework that we will formulate here is that it makes it possible to quantitatively study bounds of energy consumption in the computation, especially, achievable upper bounds for quantum computation and fundamental lower bounds for classical computation. In the following, we discuss our core idea to obtain such bounds by presenting the definition of classical and quantum oracles as in the conventional query-complexity setting (Fig. 2), followed by formulating the full detail of the computational framework (Fig. 3) for our analysis.

Oracle in the classical case. To derive a fundamental nonzero lower bound of energy consumption of classical computation, our idea here is to formulate the framework in such a way that we can study the energy-consumption requirement for solving the computational task as a whole rather than implementing each gate. Conventionally, any classical computation can be rewritten as a classical circuit in a reversible way [37]. Formulating a computational task as a decision problem reduces the number of bits to output at the end of the computation to one bit. All other bits are to be reset to their initial state at the end of the computation. But if the original computation is written as a reversible circuit, it is possible in principle to uncompute the original circuit and get all the energy back that was invested into the computer to obtain the output, up to negligible $O(1)$ energy consumption to output the single bit, as established by Bennett in Refs. [37, 38]. A nonzero lower bound on the energy consumption can, therefore, arise only if uncomputation is impossible due to some additional restriction in the setting. A setting where this happens is when the computation is structured with oracles as in the query-complexity setting.

In this setting, the input to the computation is given via queries to an oracle that computes some function

$$f : \{0, 1\}^N \rightarrow \{0, 1\}^{W_f}, \quad (4)$$

where we let N be the number of bits of f 's input (which is the size of the computational problems in the query-complexity setting), and

$$W_f = O(\text{poly}(N)) \quad (5)$$

is the number of bits of f 's output. For example, Simon's problem [12, 13] is defined for $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$, i.e.,

$$W_f = N. \quad (6)$$

In the classical case, for an input string $x \in \{0, 1\}^N$, the classical oracle lets the computer know the corresponding output string $f(x) \in \{0, 1\}^{W_f}$ of the function

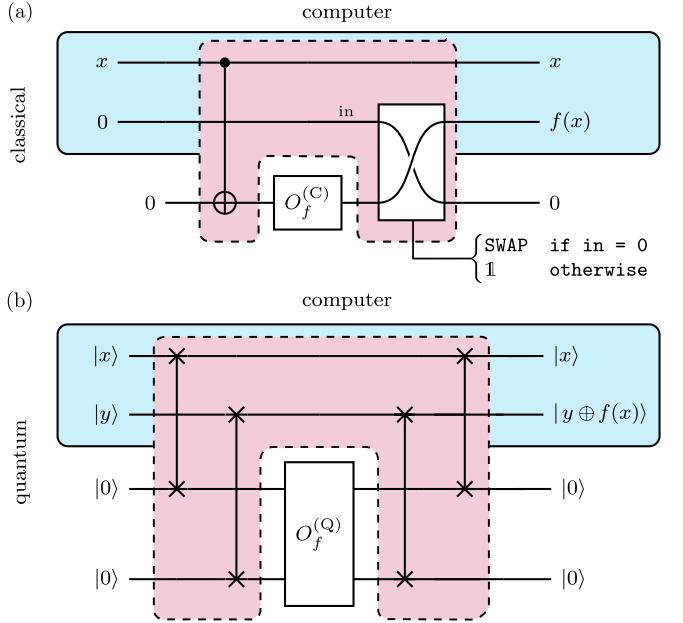


FIG. 2. A classical oracle (a) and a quantum oracle (b) in our computational framework of classical and quantum computation, respectively. (a) In the classical case, the computer sends a copy of the input bitstring x to the oracle. The oracle then allocates an output memory region in the computer (initialized in $00 \dots 0$ by the computer). In the allocated memory, the oracle writes out the corresponding output bitstring $f(x)$. As a result, each query to the oracle increases the information about the input-output relation $x \mapsto f(x)$ of the function f , which is to be stored in the computer until erased irreversibly. The oracle does nothing if there is no more memory in the computer to allocate. (b) In the quantum case, by convention, the oracle is implemented as a unitary map as defined in (8). The computer performs swap gates to set the input state into the input register for the oracle. After the oracle transforms the input state into the output state, the computer performs swap gates to receive this output state.

f , where the oracle internally computes a map

$$x \xrightarrow{O_f^{(C)}} f(x). \quad (7)$$

More precisely, as visualized in Fig. 2(a), the classical oracle receives a copy of the N -bit input string x from the computer, allocates an output memory region in the computer (which is initialized by the computer as $00 \dots 0$), and writes out the corresponding output bitstring $f(x)$ in the allocated memory region. The classical algorithm is then structured in such a way that the oracle queries are carried out outside of the computer, and the agent carrying out the algorithm on the computer is ignorant about the internals of the oracle so as to compute and learn a property of f from the information obtained via the queries. Since the oracle is a black box, the computer's memory initially has no information about f . Each call of the oracle for new x replaces 0 in an initialized memory region of the computer with $f(x)$ and thus adds the information about the input-output relation of f to the

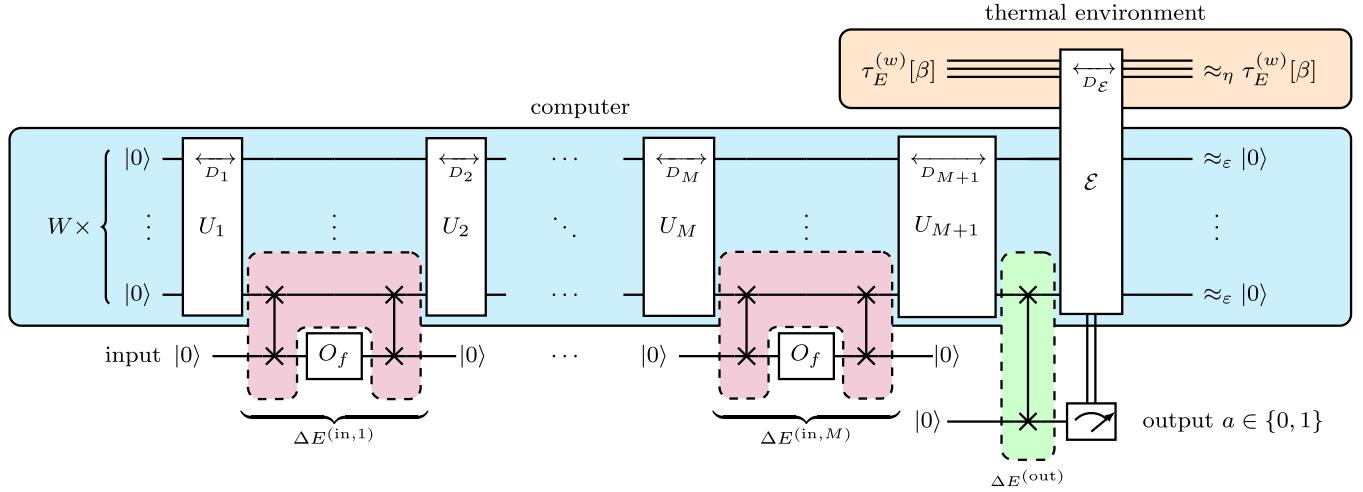


FIG. 3. A circuit representing our framework for implementing the thermodynamic cycle of computation in Fig. 1. The circuit in the figure represents, ordered from top to bottom, the thermal environment, the W (qu)bits of the (quantum or classical) computer, the outsourced oracle operation O_f for input, and finally the output (qu)bit. The computation starts with all W (qu)bits initialized in a fixed state $|0\rangle^{\otimes W}$. Then a sequence of $M + 1$ operations U_k is performed with oracle operations O_f spaced in-between. Each operation U_k has a depth of D_k with respect to some gate set. Interaction between the computer and the oracle (and that between the computer and the output) is performed as described in Fig. 2 (the figure describes the quantum oracle, i.e., the case of Fig. 2(b)), while the internal implementation of the oracle is ignored by conversion of the query-complexity setting. A single-(qu)bit measurement in the basis $\{|a\rangle : a = 0, 1\}$ of the output register is performed outside the computer to obtain the output $a \in \{0, 1\}$ of the computation for a decision problem. At the end of the computation, the W (qu)bits are reinitialized through the interaction \mathcal{E} with an environment E at inverse temperature β . In an idealized scenario, this would return each (qu)bit in the computer perfectly back to their initial state $|0\rangle$ so as to make the thermodynamic cycle closed. To capture a more realistic scenario, our analysis may allow for some non-zero infidelity $0 < \varepsilon \ll 1$, by which each (qu)bit's state after reinitialization can deviate from $|0\rangle$ by ε at the end of a single round of the computation. The required depth for this reinitialization is denoted by D_E . We also require that each set of qubits in the environment for initializing the w th qubit ($1 \leq w \leq W$) should be disturbed at most by η from its initial thermal state $\tau_E^{(w)}[\beta]$, as analyzed in Sec. III A.

computer's memory. To remove this information about f from the computer's memory, we have to erase the information in an irreversible way independently of f , using the Landauer erasure.³ This non-invertibility is what breaks Bennett's uncomputing argument [4, 37–39] in the classical case, which makes it possible to establish a nonzero lower bound on the energy consumption for any classical algorithm to solve a computational problem, as will be shown in Sec. III B.

Oracle in the quantum case. As for the quantum case, as described in Fig. 2(b), the quantum version of the oracle for f in (4) conventionally implements a unitary map that transforms each computational-basis state

as [4, 9, 10]

$$|x, y\rangle \xrightarrow{O_f^{(Q)}} |x, y \oplus f(x)\rangle, \quad (8)$$

for all $x \in \{0, 1\}^N$ and $y, f(x) \in \{0, 1\}^{W_f}$ with W_f in (5), where \oplus is the bit-wise XOR operation. For the superposition states $\sum_{x,y} \alpha_{x,y} |x, y\rangle$, $O_f^{(Q)}$ acts linearly on each term of $|x, y\rangle$ as (8). Note that we do not need the non-invertibility of the quantum oracle to derive the achievable upper bound of energy consumption for quantum computation, as we will show in Sec. III A.⁴

³ Once the computer knows $f(x)$, it is impossible to completely uncompute the information of $f(x)$ unless the computer irreversibly erases $f(x)$ by the Landauer erasure. If we were to allow for calling the classical oracle for the same input x twice, then the computer would obtain two copies of $f(x)$ in the computer's memory, and in this case, by taking the bit-wise XOR operation, one of the copies could be uncomputed; however, even in this case, it is still impossible to reversibly uncompute both copies completely as the circuit for uncomputing f requires knowledge about f . For simplicity of analysis, we assume that the different queries to the classical oracle use different inputs x since it is useless to make more than one query for the same x .

⁴ Since we follow the conventional definition (8) of quantum oracles in the query-complexity setting, unlike the classical case, the input $f(x)$ can be uncomputed by calling the quantum oracle twice; in particular, since $f(x) \oplus f(x) = 0$, the quantum oracle is its own inverse (i.e., $O_f^{(Q)} \circ O_f^{(Q)}$ is an identity map) and can, in principle, be used to uncompute all the oracle outputs. Due to this invertibility of the quantum oracle, it is a priori not clear whether there exists a fundamental lower bound on the energy consumption of the quantum computation in this setting. In practice, however, one may encounter another situation where the action of the quantum oracle except for $y = 0$ in (8) is undefined. Simon's problem can still be solved with such a quantum

Formulation of our computational framework based on query-complexity setting. Using these oracles in both classical and quantum cases, we here present the formulation of our computational framework, as shown in Fig. 3. While it is a standard setting in computational complexity theory [4], oracle-based computation comes with important assumptions that are worth being pointed out. For one, even if the algorithm can be written as a polynomial-size circuit with a polynomial number of oracle queries in terms of the problem size N , assuming a black-box oracle generally does not guarantee that the internals of the oracle itself can be implemented efficiently in polynomial time in N , which may be a usual criticism in computational complexity theory for oracle-based computation [10]. For another, similar to the time of implementing the oracle, the required energy consumption for implementing the internals of the oracle may not be polynomial in N in general either. For an actual implementation of the oracle, the bounds on the runtime and the energy consumption eventually depend on its physical details; however, we will propose how to make this additional cost negligible in practice, which we will discuss along with explaining the relevant energy consumption of the oracles in Sec. II C.

In the most general form, an oracle-based algorithm based on Fig. 1 can be structured as the circuit shown in Fig. 3, where the reinitialization of the computer's state is performed right after the output in an irreversible way based on Landauer erasure [11]. In particular, given a family of decision problems with problem size N for a function $f : \{0, 1\}^N \rightarrow \{0, 1\}^{W_f}$ in (4), the algorithm is described by the circuit with the number of oracle queries $M(N)$, the width $W(N)$, and the depth $D_k(N)$ for the k th part ($1 \leq k \leq M + 1$). The function f is unknown at the beginning of the computation, and the knowledge on f is input to the computer via multiple queries to the oracle. Usually, the algorithm's width and depth, as well as the number of queries, monotonically increase as the problem size N grows. For Simon's problem, we will explicitly write out their dependence on N in Sec. IV. In the following, we may omit the N dependency and write $M(N)$, $W(N)$, and $D_k(N)$ as M , W , and D_k , respectively, for simplicity of presentation.

In Fig. 3, the computer is initialized in a fixed pure state $|0\rangle^{\otimes W}$ of W (qu)bits, where the number of (qu)bits should be at least the number of bits required for storing the input and the output of the N -bit function f in (4),

oracle queried with $y = 0$. Then, it may be operationally infeasible to implement the inverse $(O_f^{(Q)})^{-1}$ by only calling $O_f^{(Q)}$ with $y = 0$ with a reasonable number of queries. Still, in light of recent findings showing protocols for reversing unknown unitary operations [40–43], it is currently unknown how to rigorously impose the non-invertibility to quantum oracles used in quantum algorithms. Finding a provable lower bound for the energy consumption of quantum computation in a variant of the query-complexity setting is therefore left for future work.

i.e., with W_f in (5),

$$W \geq N + W_f. \quad (9)$$

Note that in the classical case, storing information obtained from M queries to the classical oracle requires $W \geq (N + W_f)M$, but our analysis of the lower bound of energy consumption of classical computation does not explicitly use such requirements since the lower bound will be derived from the amount of information (i.e. the entropy) rather than the number of qubits, as will be shown in Sec. III B. In both classical and quantum cases, (9) holds true as the lower bound of W .

Then, a sequence of $M + 1$ unitary circuits U_k of computational depth

$$D_k \geq 0 \quad (10)$$

are performed, with the computational depth defined with respect to the decomposition of those unitaries into a product of gates selected from a finite gate set, such as Clifford and T gates, where each gate in the gate set acts on at most a constant number of (qu)bits.⁵

Between each operation U_k , the oracle $O_f^{(C)}$ in (7) for the classical and $O_f^{(Q)}$ in (8) for the quantum case is called in total M times. For simplicity of notation, we may omit the superscripts of the classical and quantum definitions of the oracles to simplify write O_f when it is clear from the context. As shown in Fig. 2, upon each query to O_f , the input state into O_f is pulled out of the computer to an input register to the oracle; then, this state of the input register goes through the oracle O_f , and the state output from O_f is put into the computer. In our analysis, the runtime of O_f is assumed to be negligible by the convention of the query-complexity setting, so the computer does not have to perform identity gates to wait for the runtime of O_f . We let

$$D_{\text{in}} = O(1) \quad (11)$$

denote the depth of each part of the circuit to make a query to the oracle for the input (i.e., each of the red parts in Figs. 2 and 3).

At the end of the computation, another swap gate (i.e., a green part in Fig. 3) is also performed to pull out an output state from the computer to an output register, which is a single (qu)bit for a decision problem. The

⁵ The particular choice of the gate set for quantum computation is irrelevant for our main results due to the Solovay-Kitaev Theorem [9, 44], as it would, at worst, contribute with a polylogarithmic correction to the circuit depth. For simplicity of the analysis, we assume that we use a gate set such that the swap gate is implementable within a constant depth; for example, if the gate set includes the CNOT gate, the swap gate is implementable with three CNOT gates [9]. For universal classical computation, the classical Toffoli gate can be used.

depth of a circuit of this output part (i.e., the green part of Fig. 3) is denoted by

$$D_{\text{out}} = O(1). \quad (12)$$

Reading the output register by the measurement in basis $\{|a\rangle : a = 0, 1\}$ yields the single-bit output $a \in \{0, 1\}$ of the algorithm.

Finally, the computer's memory is reinitialized in the part we label \mathcal{E} in Fig. 3, using an environment in a thermal state at inverse temperature

$$\beta > 0. \quad (13)$$

In particular, let ρ' denote the W -qubit state in the computer just after the reinitialization; then, for each $w = 1, \dots, W$, we require that each single-qubit reduced state for ρ' of the w th qubit should be ε -close to the initial state $|0\rangle$ of the computer, i.e.,

$$\langle 0 | \text{tr}_{\overline{w}}[\rho'] | 0 \rangle \geq 1 - \varepsilon, \quad (14)$$

where $\text{tr}_{\overline{w}}$ is the partial trace over all the W qubits but the w th. We also require that during reinitializing the state of the w th qubit, the environment, initially in the Gibbs state $\tau_E^{(w)}[\beta]$, should also be disturbed at most by a small constant

$$\eta > 0, \quad (15)$$

as will be detailed more in (69) of Sec. III A. The depth of this reinitialization part of the circuit is denoted by $D_{\mathcal{E}}$.

As a whole, the circuit depth of the circuit represented in Fig. 3 is denoted by

$$D := \sum_{k=1}^{M+1} D_k + M D_{\text{in}} + D_{\text{out}} + D_{\mathcal{E}} = D_{\text{comp}} + D_{\mathcal{E}}, \quad (16)$$

where we collectively represent the circuit depths (10), (11), and (12) for the operations U_1, \dots, U_{M+1} , the input, and the output as

$$D_{\text{comp}} := \sum_{k=1}^{M+1} D_k + M D_{\text{in}} + D_{\text{out}}. \quad (17)$$

For quantum computation, the operations U_k are unitaries acting on the W qubits in the computer, and the state of the qubits can be in a generic superposition. For classical computation, the operations U_k may only be reversible classical logic operations, and the states of bits are always dephased in their energy-eigenbasis; in other words, the classical computer's state is in a probabilistic mixture of bit strings, which we may represent as diagonal density operators to use the same notation as those in quantum computation.

C. Energy consumption

Accounting for the energy consumption of quantum computation in full depth requires considering both microscopic and macroscopic contributions. Macroscopic contributions would be the energetic costs coming from refrigerators, readout amplifiers, and other electronics not directly involved in qubit control, as studied in, e.g., Ref. [8]. Under microscopic contributions, we understand all energetics directly relevant to performing operations on the qubits, which is what we are interested in for the scope of this article. In Definition 1, we have already given a high-level definition of what we refer to as the energy consumption of the computation in our framework. Let us subdivide it further and explain the individual contributions:

(i) **Energetic cost:** Any non-trivial operation changing the state of a memory system with non-degenerate energy levels in the computer may cost some energy due to the change of the energy of the computer's state. Per gate U , this cost is given, in expectation, by

$$\Delta E^{(U)} = \text{tr}[H(U\rho U^\dagger - \rho)], \quad (18)$$

where H is the Hamiltonian of the system, and ρ is the state of the system before performing U .

(ii) **Control cost:** Implementing a desired unitary operation U on a target quantum system by some Hamiltonian dynamics generally requires an auxiliary system for the control to make the implementation energy-preserving as a whole. In general, this implementation may change the state of the auxiliary system, and compensating for this change has to require some thermodynamic control cost that we label as

$$E_{\text{ctrl.}}^{(U)} \geq 0, \quad (19)$$

which can be considered heat dissipation and is by definition a nonnegative quantity. Note that we here may not analyze model-specific control costs but just write the control cost as $E_{\text{ctrl.}}^{(U)}$ by abstracting the details, to maintain the general applicability of our analysis. See also Appendix B for a discussion on model-specific instantiation of the control cost.

(iii) **Initialization cost:** Landauer erasure (\mathcal{E} in Fig. 3) to reinitialize the computer's memory comes at an additional work cost, bounded by the heat dissipation Q_E into the environment necessary for erasure, as will be defined in more detail as (56). Alternatively, we can think of this initialization cost as the one from type (i), but not on the memory register of the computer but on the environment.

When it comes to deriving an achievability result for quantum computation, we account for the three types of

energy costs listed above: (i) energetic costs on the memory qubits, (ii) control costs, and (iii) initialization costs that arise as heat dissipation into a thermal environment when states in the computer are irreversibly erased into a fixed pure state. In the following, we write out the decomposition into these three costs individually for all the gates in the framework of Fig. 3.

The energy $\mathcal{W}_{\text{gates}}$ that the agent invests into performing the computation in Fig. 1 can be decomposed into contributions from all the gates listed in Fig. 3; that is, all the operations for U_k , the operations for the input from the oracles O_f and the output (the red and green parts of Fig. 3, respectively), and the erasure \mathcal{E} come at some work cost, which we write as

$$\mathcal{W}_{\text{gates}} = \sum_{k=1}^{M+1} \mathcal{W}_{\text{gate}}^{(U_k)} + \sum_{k=1}^M \mathcal{W}_{\text{gate}}^{(\text{in},k)} + \mathcal{W}_{\text{gate}}^{(\text{out})} + \mathcal{W}_{\text{gate}}^{(\mathcal{E})}. \quad (20)$$

In particular, for each $k \in \{1, \dots, M\}$, the work cost of U_k is given by

$$\mathcal{W}_{\text{gate}}^{(U_k)} = \Delta E^{(U_k)} + E_{\text{ctrl}}^{(U_k)}, \quad (21)$$

where $\Delta E^{(U_k)} = \text{tr}[H(U_k \rho_{k-1} U_k^\dagger - \rho_{k-1})]$ is the change of energy in the memory of the computer, H is the Hamiltonian of the W qubits in the computer, and ρ_{k-1} is the memory state just before application of U_k .

As for the energy consumption for the k th query to the oracle for the input, we let $\Delta E^{(\text{in},k)}$ denote the change of energy of states of the computer's memory before and after making the k th query, where the states before and after each query are shown in Fig. 2. Note that the states (and thus the energy) outside the computer in Fig. 2 remain unchanged. Apart from this energetic cost, the rest of the cost for the k th query is

$$\mathcal{W}_{\text{gate}}^{(\text{in},k)} = E_{\text{ctrl}}^{(\text{in},k)}, \quad (22)$$

where $E_{\text{ctrl}}^{(\text{in},k)}$ is the control cost of performing a red part of Figs. 2 and 3 for the k th query.⁶ The total energy transfer from the oracle into the computer's memory amounts to the contribution $\Delta E^{(\text{in})}$ in Fig. 1, given by

⁶ Due to the conceptual separation of computer and oracle, our theoretical analysis does not count the control cost of performing the oracle's internal operations as part of the computer's energy consumption; however, in practical cases, it is indeed possible to make this additional control cost negligible. For example, while the computer is a universal machine, the oracle is a single-purpose device for the input and thus may be made up of specialized parts with negligible energy loss compared to the computer. In Sec. V, we also explicitly propose an experimental setup with clarification on how the oracle can be implemented efficiently in the case of Simon's problem [12, 13], using a cryptographic primitive. In this setup, the non-invertibility of the classical oracle is justified by the computational hardness assumption on the security of the cryptographic primitive in practice, in place of the black-box assumption in our theoretical analysis.

the sum of the energy costs $\Delta E^{(\text{in},k)}$ of all M queries, i.e.,

$$\Delta E^{(\text{in})} = \sum_{k=1}^M \Delta E^{(\text{in},k)}. \quad (23)$$

Similarly, we let

$$-\Delta E^{(\text{out})} \in \mathbb{R} \quad (24)$$

denote the change of energy in the computer's memory before and after the output (which can be negative if the output state has higher energy than the initial state $|0\rangle$). The negative sign in (24) follows from that in (2), respecting the direction of the arrow in Fig. 1. Apart from this energetic cost, the rest of the cost of outputting the result of the algorithm is

$$\mathcal{W}_{\text{gate}}^{(\text{out})} = E_{\text{ctrl}}^{(\text{out})}, \quad (25)$$

where $E_{\text{ctrl}}^{(\text{out})}$ is the control cost of performing a green part of Fig. 3 for the output.⁷

Reinitializing the memory at the end of the computation also comes at a work cost given by

$$\mathcal{W}_{\text{gate}}^{(\mathcal{E})} = \Delta E^{(\mathcal{E})} + E_{\text{ctrl}}^{(\mathcal{E})} + \mathcal{Q}_E. \quad (26)$$

The three terms account for the energy change of the qubits $\Delta E^{(\mathcal{E})} = \text{tr}[H(\rho' - \rho)]$, the control cost $E_{\text{ctrl}}^{(\mathcal{E})}$, and the initialization cost \mathcal{Q}_E of the heat dissipation into the environment, where ρ' is the W -qubit state after erasure and should satisfy (14), and ρ the state just before erasure.

Control costs in (quantum) computation arise because performing a unitary operation U on a system usually requires external control. The system itself generically evolves according to its Hamiltonian, but if one desires to perform some non-trivial operation U on this target system, additional degrees of freedom of an auxiliary system should be needed. For example, by means of some external control field, an interaction Hamiltonian H_{int} with the target system can be turned on for a duration of time t , which then generates the unitary $U = e^{-iH_{\text{int}}t}$. In a simple model that we give in Appendix B as an illustration, the state of the auxiliary system for the control degrades through the interaction. Reinitializing the state of the auxiliary system then comes at some positive thermodynamic cost, whose precise form depends on

⁷ In our framework of Fig. 3, where the computer never performs measurements, the measurement to read the output is performed outside the computer. For decision problems, only a single (qu)bit of information is output. The energy consumption of performing the measurement of this (qu)bit with finite precision is $O(1)$ and does not grow with the problem size. Thus, the energy consumption of realizing this single-qubit measurement to a finite precision is negligible, which is not counted in our analysis to simplify the bounds.

the physical model of the actual system of interest. This energy is dissipated into some environment E' , which we do not specify further for the sake of the generality of our analysis. Rather than specifying the model, we work with generic positive constants for the control costs that give the overall control dissipation after summing up, i.e.,

$$\mathcal{Q}_{E'} = \sum_{k=1}^{M+1} E_{\text{ctrl.}}^{(U_k)} + \sum_{k=1}^M E_{\text{ctrl.}}^{(\text{in},k)} + E_{\text{ctrl.}}^{(\text{out})} + E_{\text{ctrl.}}^{\mathcal{E}} \geq 0. \quad (27)$$

The environment E' into which this energy is dissipated is not explicitly illustrated in our computational frame-

work of Fig. 3; there, only the initialization cost \mathcal{Q}_E appears. Adding these two terms together gives

$$\mathcal{Q} = \mathcal{Q}_E + \mathcal{Q}_{E'}, \quad (28)$$

as in the thermodynamic diagram of Fig. 1.

With these definitions at hand, we can make the sanity check to verify that the energy conservation (3) is true in the limit where the error ε in (14) is small enough to close the thermodynamic cycle of the computation. In particular, writing out (2) explicitly gives

$$\mathcal{W} = \left(\sum_{k=1}^{M+1} \mathcal{W}_{\text{gate}}^{(U_k)} + \sum_{k=1}^M \mathcal{W}_{\text{gate}}^{(\text{in},k)} + \mathcal{W}_{\text{gate}}^{(\text{out})} + \mathcal{W}_{\text{gate}}^{(\mathcal{E})} \right) + \left(\sum_{k=1}^M \Delta E^{(\text{in},k)} \right) - \Delta E^{(\text{out})} \quad (29)$$

$$= \left(\sum_{k=1}^{M+1} \Delta E^{(U_k)} + \sum_{k=1}^M \Delta E^{(\text{in},k)} - \Delta E^{(\text{out})} + \Delta E^{(\mathcal{E})} \right) + \left(\sum_{k=1}^{M+1} E_{\text{ctrl.}}^{(U_k)} + \sum_{k=1}^M E_{\text{ctrl.}}^{(\text{in},k)} + E_{\text{ctrl.}}^{(\text{out})} + E_{\text{ctrl.}}^{(\mathcal{E})} \right) + \mathcal{Q}_E \quad (30)$$

$$\rightarrow 0 + \mathcal{Q}_{E'} + \mathcal{Q}_E = \mathcal{Q}, \quad (31)$$

in the limit of $\varepsilon \rightarrow 0$. In (29), following the definition (2) of the energy consumption in Definition 1, we have represented the computer's energy consumption \mathcal{W} as the sum of all work costs $\mathcal{W}_{\text{gate}}^{(\dots)}$ and the energy exchanges with the input $\Delta E^{(\text{in})}$ and output $-\Delta E^{(\text{out})}$. Then, in the first term of (30), given $\varepsilon \rightarrow 0$, the qubit energy changes of type (i) together sum to zero because the computation together with reinitialization is a cyclic process on the computer. What is left is the second term, i.e., a sum of all control costs (ii), and the third term, i.e., the initialization cost (iii). Together, they yield energy conservation $\mathcal{W} = \mathcal{Q}$. This energy conservation is crucial for our analysis of the upper and lower bounds of energy consumption of computation in Sec. III B.

We also remark that, along the way of the analysis of the achievable upper bound of the energy consumption of quantum computation in Sec. III A, it is also important to check that the work costs (which may include both the energetic cost and the control cost as in (21) and (26)) are feasibly bounded in every step of quantum computation, whereas the energetic costs cancel out in evaluating the net energy consumption \mathcal{W} as shown in (30). To see the importance, we recall, for example, that the work cost $\mathcal{W}_{\text{gate}}^{(U_k)}$ is the cost that the agent needs to invest at each step U_k of the computation. For the implementability of each step of the computation, the work cost for each step should be bounded; otherwise, the agent would need to invest too much energy to run an intermediate step of the quantum algorithm. For this reason, in Sec. III A, we will provide achievable upper bounds of both the work costs and the control costs along the way of our analysis while using the control costs to present theorems on the upper

bounds of the net energy consumption.

III. GENERAL BOUNDS ON ENERGY CONSUMPTION OF COMPUTATION

In this section, we derive general upper and lower bounds of energy consumption of conducting quantum and classical computation, respectively, in the framework formulated in Sec. II. First, in Sec. III A, we provide a detailed analysis of the achievable energy consumption, accounting in particular for finite-fidelity and finite-step Landauer erasure for reinitialization and quantum error correction. Second, in Sec. III B, we derive a fundamental lower bound on the energy consumption.

A. Upper bound on energy consumption for quantum computation

In this section, we derive a general upper bound

$$\mathcal{W} \leq \mathcal{W}^{(Q)} \quad (32)$$

of energy consumption that is achievable by quantum computation within the framework of Fig. 3. The general idea for estimating $\mathcal{W}^{(Q)}$ is to break down the budget (29) of energy consumption into the contributions from the individual elementary gates and bound each contribution from above. The achievable upper bound of the energy consumption of quantum computation is determined not only by the query complexity; the total volume (width

times depth) of the quantum circuit for implementing the quantum algorithm is also crucial.⁸

In the following analysis, we begin with introducing the maximum work cost $\mathcal{W}_{\max, \text{el.}}$ per qubit for implementing an elementary gate in the gate set. Using $\mathcal{W}_{\max, \text{el.}}$, we will bound the work costs and the energy transfers between the oracle and the computer, i.e., each term in (29), as $\mathcal{W}_{\max, \text{el.}}$ times the volume (i.e., width times depth) of the part of the circuit in Fig. 3 corresponding to each term. At the same time, to estimate the overall energy consumption \mathcal{W} , we will also bound the control costs and the initialization cost appearing in (30). Our analysis of these bounds consisted of the following three steps, where we do a preliminary idealized estimate without considering gate errors in the first two steps, and then we account *a posteriori* for the overheads of quantum error correction in the third step.

- (1) *Contributions from gates, input, and output.* We estimate the contributions $\mathcal{W}_{\text{gate}}^{(U_k)}$, $\mathcal{W}_{\text{gate}}^{(\text{in},k)}$, $\mathcal{W}_{\text{gate}}^{(\text{out})}$, $\Delta E^{(\text{in},k)}$, and $\Delta E^{(\text{out})}$ for generic input sizes, widths, and depths.
- (2) *Contribution from reinitialization by finite Landauer erasure.* We bound the cost $\mathcal{W}_{\text{gate}}^{\mathcal{E}}$ of reinitializing the qubits, for which we derive a bound for finite Landauer erasure.
- (3) *Contribution from quantum error correction.* Lastly, we correct for the space and time overhead coming from quantum error correction.

Maximum work cost for implementing an elementary gate. We here analyze the work cost per qubit for performing each elementary quantum gate in the gate set, which is to be bounded by a constant independent of the size of the problem the gates are used to solve.

The work cost of each gate is determined as follows. In degenerate quantum computing, that is, when all computational basis states are energy-degenerate, unitary operations would be free in terms of the energetic cost [45–47]. Any implementation of a quantum computer on a physical platform inherently requires some energy splitting of the qubits in order to control their state. Several works have been investigating the fundamental lower and achievable energy cost of performing unitary operations [32, 48, 49], often with the goal to estimate *control costs* to achieve a certain target fidelity in implementing a given unitary gate. Physical realizations of the quantum gate are never perfect, but the requirement of the threshold theorem for fault-tolerant quantum computation (FTQC) is that the physical error rate of the

elementary gates should be below a certain threshold constant [50–61]. The existence of a finite threshold for FTQC implies that the work cost of performing a unitary operation within a better fidelity than the threshold can also be bounded. We can, therefore, take the maximum work cost per qubit over the gates in the universal gate set as

$$\mathcal{W}_{\max, \text{el.}} := E_{\text{qubit}} + E_{\text{ctrl.}}, \quad (33)$$

which splits into a contribution from the single-qubit energy change E_{qubit} and the remaining contribution $E_{\text{ctrl.}}$ representing the control cost per qubit. The universal gate set includes multiqubit gates such as CNOT gates, and for a multiqubit gate, the per-qubit work cost in (33) represents the total work cost divided by the number of qubits involved in the gate. We will explain E_{qubit} and $E_{\text{ctrl.}}$ in the following.

In (33), E_{qubit} is the maximum single-(qu)bit energy over all the (qu)bits in the computer, where

$$E_{\text{qubit}} \geq 0. \quad (34)$$

To be more specific, for the w th (qu)bit out of the W (qu)bits ($w \in \{1, 2, \dots, W\}$), let $H_{\text{qubit}}^{(w)} > 0$ be the single-(qu)bit Hamiltonian characterizing the energy of the w th (qu)bit, scaled appropriately so that its ground-state energy should be non-negative; in this case, we can take

$$E_{\text{qubit}} = \max_w \left\{ \left\| H_{\text{qubit}}^{(w)} \right\|_{\infty} \right\}, \quad (35)$$

where $\|\cdot\|_{\infty}$ is the operator norm. To ensure the generality of our analysis, we will keep E_{qubit} as a general variable, which may depend on the circuit size and scale, at most, polynomially with N . As for the control cost $E_{\text{ctrl.}}$ in (33), our analysis also keeps $E_{\text{ctrl.}}$ as a general variable, which may also grow, at most, polynomially with N . To summarize, the assumption on E_{qubit} and $E_{\text{ctrl.}}$ for our analysis is

$$E_{\text{qubit}} = O(\text{poly}(N)), \quad (36)$$

$$E_{\text{ctrl.}} = O(\text{poly}(N)), \quad (37)$$

while these may have much smaller, constant order $O(1)$ in conventional settings. This general setting allows us to account for the possibility that as the problem size grows, qubit crosstalk in a large quantum device may lead to a higher control cost per qubit. As shown in Ref. [32] (also see Appendix B), the variable $E_{\text{ctrl.}}$ is conventionally related to E_{qubit} as

$$E_{\text{ctrl.}} = O(E_{\text{qubit}}). \quad (38)$$

As a whole, we have a maximal work cost $\mathcal{W}_{\max, \text{el.}}$ per qubit in (33) to perform any elementary gate in the gate set, regardless of the initial state of the qubits before applying the gate. This upper bound $\mathcal{W}_{\max, \text{el.}}$ may be a loose worst-case estimate to cover the case where an elementary gate rotates the qubits it acts on from the

⁸ Indeed, there are edge cases where even if the query complexity scales polynomially, the quantum algorithm may require an exponential runtime to process the polynomial amount of information obtained from the queries. In such a case, the quantum algorithm may require exponentially large energy consumption due to the exponential runtime.

ground state to the excited state with the energy change of E_{qubit} . Note that some operations, e.g., a gate transforming the excited state to the ground state, may allow us to extract energy; also, the overall energetic cost sums up to zero as in (30). Still, it is important to bound the worst-case contributions of the work cost in terms of $\mathcal{W}_{\text{max., el.}}$, rather than only bounding the contributions of $E_{\text{ctrl.}}$ in (33), for implementability of every step of the computation.

Step (1): Contributions from gates, input, and output. Using $\mathcal{W}_{\text{max., el.}}$ in (33), we estimate the energy consumption for each part of the circuit in Fig. 3.

To estimate the cost of the unitaries U_k , we use the fact that the unitary acts on W qubits, and also each of these unitaries has depth D_k . Thus, the cost of $\mathcal{W}_{\text{gate}}^{(U_k)}$ applying U_k in (21) is bounded by

$$\mathcal{W}_{\text{gate}}^{(U_k)} \leq \mathcal{W}_{\text{max., el.}} \times WD_k \quad (39)$$

$$= O((E_{\text{qubit}} + E_{\text{ctrl.}}) \times WD_k), \quad (40)$$

where WD_k is the volume of this part of the circuit, and the second line follows from (33). Due to (21), $\mathcal{W}_{\text{gate}}^{(U_k)}$ in (40) is decomposed into

$$\Delta E^{(U_k)} = O(E_{\text{qubit}} \times WD_k), \quad (41)$$

$$E_{\text{ctrl.}}^{(U_k)} = O(E_{\text{ctrl.}} \times WD_k). \quad (42)$$

In each red part of calling the oracle between U_k and U_{k+1} ($k = 1, \dots, M$) in Fig. 3, we account for the control cost in swapping the states between the computer and the input register for the quantum oracle back and forth as in Fig. 2(b). In this part, i.e., the k th query to the

As a whole, due to (40), (44), (45), (47), and (48), the work cost and the energy transfer between the computer and the oracle in the part of the circuit in Fig. 3 for the computation (i.e., the parts except for \mathcal{E} in Fig. 3) is at most

$$\begin{aligned} & \sum_{k=1}^{M+1} \mathcal{W}_{\text{gate}}^{(U_k)} + \sum_{k=1}^M \mathcal{W}_{\text{gate}}^{(\text{in},k)} + \mathcal{W}_{\text{gate}}^{(\text{out})} + \sum_{k=1}^M \Delta E^{(\text{in},k)} - \Delta E^{(\text{out},k)} \\ &= O\left((E_{\text{qubit}} + E_{\text{ctrl.}}) \times W \left(\sum_{k=1}^{M+1} D_k + MD_{\text{in}} + D_{\text{out}}\right)\right) \end{aligned} \quad (49)$$

$$= O(\mathcal{W}_{\text{max., el.}} \times WD_{\text{comp}}), \quad (50)$$

where D_{comp} is given by (17), and $\mathcal{W}_{\text{max., el.}}$ by (33). In words, the work cost scales at most with the *volume* of the circuit for implementing the algorithm, i.e., the product of total width and depth of the circuit for implementing the algorithm (complexity-theoretic contribution) times the parameter $\mathcal{W}_{\text{max., el.}}$ in (33) representing the maximum work cost per qubit of gates (energetic contribution). Also, as shown in (30), the relevant contribution to the energy consumption \mathcal{W} is the control cost, which is a part of the work cost and bounded, due to (42), (44), and (47), by

$$\sum_{k=1}^{M+1} E_{\text{ctrl.}}^{(U_k)} + \sum_{k=1}^M E_{\text{ctrl.}}^{(\text{in},k)} + E_{\text{ctrl.}}^{(\text{out})} = O\left(E_{\text{ctrl.}} \times W \left(\sum_{k=1}^{M+1} D_k + MD_{\text{in}} + D_{\text{out}}\right)\right) \quad (51)$$

$$= O(E_{\text{ctrl.}} \times WD_{\text{comp}}). \quad (52)$$

Step (2): Contribution from reinitialization by finite Laudauer erasure. To close the thermodynamic cycle on

oracle with input size $N + W_f$ as in (8), we have $N + W_f$ auxiliary qubits for the input register in addition to the W qubits in the computer; i.e., the width of this part of the circuit in Fig. 3 is $W + N + W_f$. Therefore, the control cost $\mathcal{W}_{\text{gate}}^{(\text{in},k)}$ of the k th-query part for the input in (22) is given by

$$\mathcal{W}_{\text{gate}}^{(\text{in},k)} \leq E_{\text{ctrl.}} \times (W + N + W_f) D_{\text{in}} \quad (43)$$

$$= O(E_{\text{ctrl.}} \times WD_{\text{in}}), \quad (44)$$

where D_{in} is given by (11), and we use (9) in the last line. Also, using E_{qubit} in (34), we bound the energetic cost for the k th query by

$$\Delta E^{(\text{in},k)} \leq E_{\text{qubit}}(N + W_f) = O(E_{\text{qubit}} \times W), \quad (45)$$

where we use (9).

In the same way, in the green part of Fig. 3, the control cost $\mathcal{W}_{\text{gate}}^{(\text{out})}$ of outputting the result of the computation in (25) is at most

$$\mathcal{W}_{\text{gate}}^{(\text{out})} \leq E_{\text{ctrl.}} \times (W + 1) D_{\text{out}} \quad (46)$$

$$= O(E_{\text{ctrl.}} \times WD_{\text{out}}), \quad (47)$$

where D_{out} is given by (12), the factor $W + 1$ is composed of the W qubits in the computer and another single qubit for the output register. Moreover, using the upper bound E_{qubit} of the single-qubit energy in (34), we bound the energetic cost of the output by

$$|\Delta E^{(\text{out})}| \leq E_{\text{qubit}}. \quad (48)$$

the computer's memory, the qubits in the computer are

reinitialized by the end of the computation as shown by the part \mathcal{E} of the circuit in Fig. 3. The research field of algorithmic cooling [62–65] studies algorithms that cool down a mixed state of qubits in a quantum computer into a pure state to reset the qubits’ state, but these results do not straightforwardly apply to our framework due to the difference in settings from ours. Protocols for Landauer erasure [11] can be used for the reinitialization in our framework. Existing results on Landauer erasure use either a sequence of unitary operations as in Refs. [16, 23] or finite-time relaxation [24–26], but these works mostly provide lower bounds of the heat dissipation. Reference [23] also provides an upper bound of heat dissipation, but the protocol to achieve this upper bound requires an infinite number of steps. By contrast, what we need is an achievable upper bound of the energy consumption for erasing a given mixed state to a pure state up to finite precision ε within a finite number of steps. To obtain such a bound, we work on protocols using a sequence of unitary operations as in Ref. [23] (also called a collision model [66]), which are more suitable to obtain a finite bound of steps than those using finite-time relaxation.

Our goal here is to find an upper bound of the achievable cost $\mathcal{W}_{\text{gate}}^{\mathcal{E}}$ in this finite setting of reinitialization for our framework. We note that in Sec. III B, we will also clarify a lower bound of the required energy for this initialization by a Landauer-erasure protocol, which will be given by the required amount of heat \mathcal{Q}_E dissipated into the environment in the initialization (i.e., the initialization cost as defined in Sec. II C). By contrast, we here bound $\mathcal{W}_{\text{gate}}^{\mathcal{E}}$ from above, not from below, especially in the setting where we take into account all possible contributions including the control cost as well as the energetic cost. In principle, a part of the achievable cost of the Landauer-erasure protocol may be made as close to its lower bound \mathcal{Q}_E as possible [23], which is in part true, but only if we focus on the energetic cost and ignore the control cost. Problematically, exactly achieving \mathcal{Q}_E may require infinite time steps, and the control cost required for these infinite time steps may also diverge; indeed, based on the third law of thermodynamics in the formulation of Nernst’s unattainability principle [15], it is generally agreed upon that cooling a (quantum) state to absolute zero and thereby erasing its previous state inevitably comes at divergent resource costs in some form [16, 67–69].

To obtain the upper bound including the control cost in addition to the energetic cost, we need a finite upper bound on the number of steps (depths) for achieving finite infidelity ε in the Landauer-erasure protocol, which is where Theorem 2 below steps in. The setting of the Landauer-erasure protocol starts with the initial state ρ_S of the target system S of interest and an environment available in a thermal state at inverse temperature $\beta > 0$

$$\tau_E[\beta] := \frac{e^{-\beta H_E}}{\text{tr}[e^{-\beta H_E}]}, \quad (53)$$

where H_E is the Hamiltonian of the environment. Note that S will be each qubit in the computer in our analysis, but we here present a general result on a d -dimensional target system. The goal of the Landauer-erasure protocol is to transform the state

$$\rho_S \mapsto \rho'_S = \text{tr}_E[\rho'_{SE}] = \text{tr}_E[U(\rho_S \otimes \tau_E[\beta])U^\dagger] \quad (54)$$

by some unitary U , so that it should hold that

$$\langle 0 | \rho'_S | 0 \rangle \geq 1 - \varepsilon, \quad (55)$$

where $|0\rangle$ can be any pure state of the target system in the same way as the requirement (14) in our framework. In the Landauer-erasure protocol, energy is dissipated into the environment E as heat

$$\mathcal{Q}_E := \text{tr}[H_E(\rho'_E - \rho_E)]. \quad (56)$$

where we write

$$\rho'_E = \text{tr}_S[\rho'_{SE}]. \quad (57)$$

One can always write this increase in the environment’s energy as [23]

$$\beta \mathcal{Q}_E = \Delta S + D(E' \| E) + I(E' : S'), \quad (58)$$

where E, E' are the environment before and after the Landauer erasure, S' is the target system after the erasure,

$$\Delta S := S(\rho_S) - S(\rho'_S) \quad (59)$$

is called Landauer’s bound [11],

$$S(\rho) := -\text{tr}[\rho \ln \rho] \quad (60)$$

is the von Neumann entropy,

$$D(E' \| E) := \text{tr}[\rho'_E(\ln \rho'_E - \ln \tau_E[\beta])] \quad (61)$$

is the quantum relative entropy, and

$$I(E' : S') := S(\rho'_E) + S(\rho'_S) - S(\rho'_{SE}) \quad (62)$$

is the quantum mutual information. Using (58) together with $I(E' : S') \geq 0$, we can bound the relative entropy for the states of the environment after the erasure versus before as

$$D(E' \| E) \leq \beta \mathcal{Q}_E - \Delta S. \quad (63)$$

Due to $D(E' \| E) \geq 0$, it always holds that

$$\mathcal{Q}_E \geq \frac{\Delta S}{\beta}. \quad (64)$$

The quantity $\beta \mathcal{Q}_E - \Delta S \geq 0$ on the right-hand side of (63) represents the excess of the protocol’s heat \mathcal{Q}_E dissipated to the environment over that given by the Landauer’s bound ΔS , which is also an upper bound of how different the environment’s initial state $\tau_E[\beta]$ is from its

final state ρ'_E after the erasure; as the condition on the environment to be specified by the constant η in (15), we require

$$\beta \mathcal{Q}_E - \Delta S \leq \eta. \quad (65)$$

With these notations, we obtain the following finite bounds.

Theorem 2 (Finite Landauer-erasure bound). *Given a quantum state ρ_S of a d -dimensional target system for any $d \geq 2$ and any constants $\varepsilon \in (0, 1/2]$ and $\eta \in (0, 1]$, let T be*

$$T = \left\lceil \frac{(e+1)}{\varepsilon\eta} \ln \left(\frac{(e+1)(d-1)^2}{\varepsilon\eta} \right) \right\rceil = O\left(\frac{1}{\eta} \log \frac{d^2}{\varepsilon\eta} \right), \quad (66)$$

where $\lceil \dots \rceil$ is the ceiling function, and e is the Euler constant. Then, a T -step Landauer-erasure protocol can transform ρ_S into a final state ρ'_S with infidelity to a pure state $|0\rangle$ satisfying $\langle 0 | \rho'_S | 0 \rangle \geq 1 - \varepsilon$ in (55) by using an environment at inverse temperature $\beta > 0$, in such a way that the heat \mathcal{Q}_E in (56) dissipated into the environment differs from Landauer's bound ΔS in (59) at most by

$$\beta \mathcal{Q}_E - \Delta S \leq \eta, \quad (67)$$

as required in (65).

Proof. The proof can be found in Appendix A 1. \square

For our analysis of the framework in Fig. 3, this result on the finite erasure is now applied qubit-wise, i.e., $d = 2$ in Theorem 2, to the computer's state after the agent has obtained output from the computer. In particular, for each qubit of the computer, we run the erasure protocol in Theorem 2 in total W times in parallel to erase all the W qubits, using W sets of T auxiliary qubits of the environment E for an appropriate choice of H_E as shown in (A2) of Appendix A 1, where each of the W sets is used for erasing one of the W qubits in the computer. In this case, as described in Appendix A 1, the erasure protocol is composed of T steps (labeled $t \in \{1, \dots, T\}$) of swap gates, where the t th swap gate is applied between the qubit to be erased in the computer and the t th auxiliary qubit in the set of T auxiliary qubits of the environment to erase this qubit of the computer. Note that applying the qubit-wise erasure to each of the W qubits in total W times may be more costly as a whole than erasing the state of all W qubits in the computer simultaneously [70, 71] but provides a simpler upper bound of the work cost; remarkably, it suffices to use this potentially suboptimal Landauer-erasure protocol to prove the exponential quantum advantage in Sec. IV. On the other hand, for the analysis of the lower bound of the work cost in Sec. III B, we will analyze the efficient erasure protocol applied to all qubits simultaneously at the end of the computation to obtain a general lower bound. We also remark that since the analysis here deals with

the qubit-wise erasure, we indeed do not have to apply the erasure collectively at the end of the computation but can also initialize each qubit in the middle of the computation as long as the initialized qubit is no longer used for the rest of the computation. Later in our analysis of step (3), we will make a minor modification to the protocol in Fig. 3, so that we may initialize each qubit in the middle for quantum error correction rather than erasing all the qubits at the end. The following analysis of the upper bound of the work cost for the qubit-wise erasure is still applicable with this modification.

For the erasure protocol in Theorem 2 applied to our framework, the constant $\varepsilon > 0$ is to be determined by the threshold for FTQC, and $\eta > 0$ by the amount of thermal fluctuation of the environment. In particular, FTQC can be performed if every qubit is reinitialized to a fixed pure state $|0\rangle$ within constant infidelity ε satisfying the threshold theorem, i.e., (14). Moreover, we demand that the environment changes little in this process of reinitializing each qubit in the computer compared to a constant amount of inherent thermal fluctuation of the environment. To state this littleness quantitatively, we require for each single-qubit erasure that

$$\beta \mathcal{Q}_E^{(w)} - \Delta S^{(w)} \leq \eta, \quad (68)$$

where $\mathcal{Q}_E^{(w)}$ and $\Delta S^{(w)}$ are \mathcal{Q}_E and ΔS in Theorem 2 applied to erase the w th qubit in the computer ($1 \leq w \leq W$). Due to (63), this inequality implies that, for each qubit of the computer that is erased separately, the quantum relative entropy between the states of the corresponding environment $E'^{(w)}$ after the erasure protocol and $E^{(w)}$ before (in the thermal state $\tau_E^{(w)}[\beta]$ as shown in Fig. 3) is bounded by

$$D(E'^{(w)} \| E^{(w)}) \leq \eta \quad (69)$$

for some small constant $\eta > 0$. The choice of η is independent of that of ε .

For each qubit, both requirements for ε and η , i.e., (14) and (68), can be satisfied simultaneously if we choose the number of steps T large enough, but how large T should be was unclear in the existing asymptotic results [16, 23–26]. By contrast, Theorem 2 shows in more detail that if we choose (for fixed dimension $d = 2$)

$$T = O\left(\frac{1}{\eta} \log \frac{1}{\varepsilon\eta} \right), \quad (70)$$

we can satisfy the above requirements at the same time for any $\varepsilon \in (0, 1/2]$ and $\eta \in (0, 1]$. Using a gate set that can implement a swap gate within a constant depth, the depth of the erasure part \mathcal{E} in Fig. 3 is bounded by

$$D_{\mathcal{E}} \leq O(T). \quad (71)$$

Also, due to (68), if we sum together the heat dissipation for all W qubits, the protocol achieves

$$\mathcal{Q}_E \leq \sum_{w=1}^W \frac{\Delta S^{(w)} + \eta}{\beta} \leq \frac{W}{\beta} (\ln 2 + \eta), \quad (72)$$

where the last inequality follows from the upper-bound estimate $\Delta S^{(w)} \leq \ln 2$ that holds for each qubit.

Apart from the energetic cost \mathcal{Q}_E in the environment, for each of the W qubits in the computer, the energetic cost per qubit per single-depth part of the circuit for the erasure is bounded by the maximum energetic cost E_{qubit} in (34). Thus, similar to step (1), the energy change in the computer's memory due to the erasure is upper bounded by

$$\Delta E^{(\mathcal{E})} \leq E_{\text{qubit}} \times WD_{\mathcal{E}}. \quad (73)$$

As for the control cost, we can divide the control cost of erasure into a product of the maximum control cost per single-depth part of the circuit for the erasure and the depth $D_{\mathcal{E}}$ of the circuit in (71), in the same way as step (1). To bound the maximum control cost per single-depth part of the erasure, as in (53), let H_E denote the overall Hamiltonian H_E of the environment used for the single-qubit erasure protocol, which decomposes into the sum over the single-qubit Hamiltonians $H_E^{(t)}$ of the t th auxiliary qubit of the environment used at the t th time step in the protocol ($1 \leq t \leq T$). The control cost per single-depth part of the circuit depends on the energy scale of the qubit to be controlled, i.e., E_{qubit} for a qubit in the computer and $\|H_E\|_\infty$ for the environment to reinitialize the qubit by the Landauer erasure (note that $H_E \geq 0$ holds for the Landauer-erasure protocol here, as shown in (A2) of Appendix A 1). To avoid fixing the specifics of the control system, we here recall the conventional scaling of the control cost $E_{\text{ctrl.}}$ of the qubit in the computer in (38); following this conventional scaling,

we here assume that the control cost in the environment per single-depth part of the circuit depends linearly on the energy scale of the environment Hamiltonian, i.e.,

$$E_{\text{ctrl.}}^{(E)} = O(\|H_E\|_\infty), \quad (74)$$

as shown for the conventional models in, e.g., Ref. [32] (see also Appendix B). We calculate in Proposition 11 (Appendix A 1) how the largest energy eigenvalue $\|H_E\|_\infty$ of the environment Hamiltonian scales as a function of $T \rightarrow \infty$ and $\varepsilon \rightarrow 0$, for fixed inverse temperature β . This calculation allows us to estimate the per-depth control cost of all the T auxiliary qubits in the environment for erasing each of the W qubits in the computer (for fixed dimension $d = 2$) by

$$E_{\text{ctrl.}}^{(E)} = O\left(\frac{T}{\beta} \log \frac{1}{\varepsilon}\right). \quad (75)$$

Consequently, the control cost of erasure for the W qubits in the computer and the corresponding W sets of the auxiliary qubits in the environments for the Landauer erasure protocol is bounded by

$$E_{\text{ctrl.}}^{(\mathcal{E})} \leq \left(E_{\text{ctrl.}} + E_{\text{ctrl.}}^{(E)}\right) \times WD_{\mathcal{E}}. \quad (76)$$

If $E_{\text{ctrl.}}$ is independent of N , then $E_{\text{ctrl.}}$ may be negligible compared to the control cost $E_{\text{ctrl.}}^{(E)}$ of the environment, in the limit of $\eta, \varepsilon \rightarrow 0$ due to (70) and (75); however, in the general case where the control cost $E_{\text{ctrl.}}$ of the computer may grow in the problem size N and is not negligible, we keep both $E_{\text{ctrl.}}$ and $E_{\text{ctrl.}}^{(E)}$ separately in (76).

As a whole, by combining (70) – (76) together to bound the work cost for the erasure, we obtain

$$\mathcal{W}_{\text{gate}}^{(\mathcal{E})} = \Delta E^{(\mathcal{E})} + E_{\text{ctrl.}}^{(\mathcal{E})} + \mathcal{Q}_E \quad (77)$$

$$\leq \left(E_{\text{qubit}} + E_{\text{ctrl.}} + E_{\text{ctrl.}}^{(E)}\right) \times WD_{\mathcal{E}} + \frac{W}{\beta}(\ln 2 + \eta) \quad (78)$$

$$= O\left(\left(E_{\text{qubit}} + E_{\text{ctrl.}} + \frac{1}{\beta\eta}\right) \times \frac{W}{\eta} \text{polylog} \frac{1}{\varepsilon\eta}\right) \quad (79)$$

$$= O\left(\left(\mathcal{W}_{\text{max,el.}} + \frac{1}{\beta\eta}\right) \times \frac{W}{\eta} \text{polylog} \frac{1}{\varepsilon\eta}\right), \quad (80)$$

where $\mathcal{W}_{\text{max,el.}}$ is given by (33). Also, as shown in (30), the relevant contribution to the energy consumption \mathcal{W} is the control cost, which is bounded, due to (70), (72), (75), and (76), by

$$E_{\text{ctrl.}}^{(\mathcal{E})} + \mathcal{Q}_E \leq \left(E_{\text{ctrl.}} + E_{\text{ctrl.}}^{(E)}\right) \times WD_{\mathcal{E}} + \frac{W}{\beta}(\ln 2 + \eta) \quad (81)$$

$$= O\left(\left(E_{\text{ctrl.}} + \frac{1}{\beta\eta}\right) \times \frac{W}{\eta} \text{polylog} \frac{1}{\varepsilon\eta}\right). \quad (82)$$

To summarize the argument on the achievable upper bound of the energy consumption of the quantum computation in Fig. 3, in the case without gate errors, the energy consumption can be given by Theorem 3 shown in the following. The derivation is based on the expressions in (52) and (82); in particular, we use (30) to make all the energetic costs (i.e., the terms involving E_{qubit} in (34)) cancel out, and thus, only the control and initialization costs remain.

Theorem 3 (An upper bound of energy consumption in ideal case without gate errors). *Given constants $E_{\text{ctrl.}}$ in (38) and β in (13), working in an idealized case without gate errors, computation in the framework of Fig. 3 can be performed with an energy consumption \mathcal{W} (Definition 1) bounded, as the problem size N goes to infinity, and the infidelities $\varepsilon, \eta > 0$ vanish, by*

$$\mathcal{W} \leq \mathcal{W}^{(Q)} = O \left(E_{\text{ctrl.}} \times WD_{\text{comp}} + \left(E_{\text{ctrl.}} + \frac{1}{\beta\eta} \right) \times \frac{W}{\eta} \text{polylog} \left(\frac{1}{\varepsilon\eta} \right) \right), \quad (83)$$

where the width W of the circuit in (9) and the depth D_{comp} given by (17) are the parameters depending on the problem size N .

Proof. The proof can be found in Appendix A 2. \square

Step (3): Contribution from quantum error correction. The discussion so far has not taken into account imperfections in carrying out the unitary operations U_k and swap gates for input and output, which are necessary for realizing quantum computation. Any unitary gate carried out on an actual physical platform will inherently be imperfect, be it through unavoidable noise [72], non-ideal control parameters [73], or timing errors [74, 75]. Quantum error correction is a way to suppress the effect of these errors given that the rate at which these errors occur is below a certain threshold value, a result that is known as the threshold theorem [50–61]. The main principle of quantum error correction is to use a quantum error-correcting code of multiple physical qubits to redundantly represent a logical qubit on which the unitary operations are carried out. In this way, we, in principle, have protocols for FTQC. However, such fault-tolerant protocols to simulate the given original circuit by running the corresponding fault-tolerant circuit may incur space and time overheads, i.e., the overheads in terms of the number of qubits and the circuit depth, respectively, of the fault-tolerant circuit compared to the original circuit.

In the computational model herein presented, the space is quantified through the number of qubits W in (9), and the time through the circuit depth D in (90), as shown in Fig. 3. With a fault-tolerant protocol using measurements and classical post-processing, one could convert the original circuit in Fig. 3 into a fault-tolerant circuit with only polylogarithmic overheads in width and depth; in particular, to simulate a given W -width D -depth circuit, the fault-tolerant circuit can have [72]

$$W_{\text{FT}} = W \times \text{polylog}(WD), \quad (84)$$

$$D_{\text{FT}} = D \times \text{polylog}(WD), \quad (85)$$

where W_{FT} and D_{FT} are the width and depth of the fault-tolerant circuit, respectively. This choice is not unique but depends on the construction of the fault-tolerant protocol to achieve FTQC. For example, advances have been made to achieve a constant space overhead $W_{\text{FT}} = W \times O(1)$ and a quasi-polylogarithmic time overhead $D_{\text{FT}} = D \times \exp(\text{polylog}(\log(WD)))$ [61]; more recently, Ref. [76] has proven that a fault-tolerant protocol can achieve a constant space overhead $W_{\text{FT}} = W \times O(1)$ and a polylogarithmic time overhead $D_{\text{FT}} = D \times \text{polylog}(WD)$.

To analyze the upper bound of the achievable quantum cost $\mathcal{W}^{(Q)}$ of energy consumption with the fault-tolerant protocol in (84) and (85), one has to modify the analysis in steps (1) and (2), so as to allow initializing qubits in the middle of the computation rather than at the end and to include costs of measurements and classical post-processing present in the fault-tolerant protocol. The qubit initialization in the middle of computation is indispensable for performing quantum error correction to achieve FTQC; in the presence of noise, if the initialization were allowed only at the beginning or at the end, the class of problems that noisy quantum computation can solve would be significantly limited [77]. Regarding the initialization, as discussed in step (2), the same upper bound on the work cost of each qubit-wise erasure remains to hold even if we initialize each qubit in the middle rather than at the end. As for the measurement, an ideal projective measurement would come at a divergent resource cost as explored in Ref. [36], but for the quantum error correction, measurements with a finite fidelity above a certain constant threshold are sufficient. The costs of these finite-fidelity measurements may still be non-negligible, and their exact value is still subject to research [31, 78, 79]; however, in principle, this cost can also be bounded by a constant that is independent of the algorithm's other parameters such as the problem size N , similar to the cost of each gate in (33). The cost of classical post-processing can also be evaluated by writing the classical computation in terms of reversible classical logic circuits and bounding the cost of each classical logic gate by a constant, similar to (33). In this way, the increase of the work cost arising from quantum error correction can be taken into account in principle by setting ε in step (2) as a constant below the threshold and multiplying polylogarithmic factors to W , D_k , and D_{ε} as shown in (84) and (85).

On top of this, to solidify our analysis, we further employ the fact that measurements and classical post-processing are not fundamentally necessary for FTQC; in particular, instead of performing the syndrome measurement and classical post-processing for decoding in FTQC, it is possible to implement the decoding by quantum circuits and perform the correction by the controlled Pauli gates in place of applying Pauli gates conditioned on the output of the classical post-processing [80, 81]. In the fault-tolerant protocol with measurements, each

measured qubit can be reinitialized and reused in the rest of the computation, where a single qubit may be initialized multiple times in implementing the computation. By contrast, we here allow initializing each qubit only once in the middle of the computation, so as to apply the same upper bound of the work cost as that obtained in step (2); that is, the fault-tolerant protocol without measurements here never traces out each of these qubits, and furthermore, instead of reinitializing the same qubit multiple times, another auxiliary qubit is initialized by the finite erasure protocol in the middle of the computation. The fault-tolerant circuit with measurements has the depth of $D \times \text{polylog}(WD)$, and thus, each of the $W \times \text{polylog}(WD)$ qubits in this fault-tolerant circuit may have at most $D \times \text{polylog}(WD)$ measurements. To simulate these measurements and classical post-processing in the above way, it suffices to use at most $D \times \text{polylog}(WD)$ auxiliary qubits for each of the $W \times \text{polylog}(WD)$ qubits in the protocol without the measurements. Consequently, sacrificing the efficiency in the circuit width W (i.e., introducing a polynomial factor in D to W), we implement the original circuit in Fig. 3 by a fault-tolerant protocol without measurements, which achieves

$$W_{\text{FT}} = WD \times \text{polylog}(WD), \quad (86)$$

$$D_{\text{FT}} = D \times \text{polylog}(WD), \quad (87)$$

in simulating a W -width D -depth original circuit. While whether we should completely avoid the measurements in implementing quantum computation may be of questionable practical relevance, these results show that the measurement and post-processing costs are not a fundamental restriction toward bounding the energy consumption of quantum computing.

As a whole, we conclude this section by summarizing all the costs in one theorem including those of quantum error correction. Recall that the depth of the circuit in Fig. 3 is, according to (16),

$$D = D_{\text{comp}} + D_{\mathcal{E}} \quad (88)$$

$$= O\left(D_{\text{comp}} + \frac{1}{\eta} \log\left(\frac{1}{\varepsilon\eta}\right)\right) \quad (89)$$

$$= \tilde{O}\left(D_{\text{comp}} + \frac{1}{\eta}\right), \quad (90)$$

where the second line follows from (70) and (71), and with the big- \tilde{O} notation, we may ignore polylogarithmic factors. Thus, with D in (90), adding D auxiliary qubits for each of the W qubits in Fig. 3, up to polylogarithmic factors, we can use the fault-tolerant protocol without measurements in (86) and (87) to obtain the following theorem.

Theorem 4 (An upper bound of energy consumption including the cost of quantum error correction). *Given constants $E_{\text{ctrl.}}$ in (38) and β in (13), quantum computation in the framework of Fig. 3 can be performed in a fault-tolerant way with an energy consumption $\mathcal{W}^{(\text{Q})}$ (Definition 1) bounded, as the problem size N goes to infinity, and the infidelities $\varepsilon, \eta > 0$ vanish, by*

$$\mathcal{W} \leq \mathcal{W}^{(\text{Q})} = \tilde{O}\left(E_{\text{ctrl.}} \times W \left(D_{\text{comp}} + \frac{1}{\eta}\right) D_{\text{comp}} + \left(E_{\text{ctrl.}} + \frac{1}{\beta\eta}\right) \times \frac{W \left(D_{\text{comp}} + \frac{1}{\eta}\right)}{\eta}\right), \quad (91)$$

where the width W of the circuit in (9) and the depth D_{comp} given by (17) are the parameters depending on the problem size N , and with \tilde{O} , we may ignore polylogarithmic factors in W , D , $1/\varepsilon$, and $1/\eta$.

B. Lower bound on energy consumption of classical computation

In this section, we derive a general lower bound

$$\mathcal{W} \geq \mathcal{W}^{(\text{C})} \quad (92)$$

of energy consumption required for classical computation in the framework of Fig. 3. The techniques for finding an achievable upper bound in Sec. III A are not applicable to finding a fundamental lower bound of energy consumption. After all, in the limit of energy-efficient implementation, a single elementary gate may be performed at as close to zero control cost as possible, as we have pointed out in Sec. II A. Rather, the essential technique here for

deriving the nonzero fundamental lower bound is to use the non-invertibility of the classical oracle. Due to the non-invertibility of the classical oracle, to reset the computer's memory at the end of the computation, we need to perform the Landauer erasure of information obtained from the queries to the oracle, which inevitably consumes energy.

Using energy conservation in (3) together with the heat dissipation in the Landauer erasure, we obtain a general lower bound on the energy consumption of classical computation in the framework of Fig. 3, as shown by Theorem 5 below. In this theorem, we identify an entropic quantity $\frac{1}{\beta}S(C|a)$ as a lower bound of energy consumption, which falls into the category of the initialization cost

that is derived in an information-theoretic way and is thus implementation-independent.⁹ Even more notably, since we work on the query-complexity setting with the classical oracle, we can employ the techniques for deriving the lower bound of query complexity to prove rigorously that the required energy consumption for any classical algorithm to solve the problem can also be strictly gapped away from zero, even without complexity-theoretical assumptions on the hardness of the problems, such as the hardness of integer factoring for classical computation. Yet we remark that, as we will see in Corollary 9 of Sec. III B, the query complexity and the entropic lower bound of the energy consumption do not necessarily coincide in general up to constant factors.

Theorem 5 (A lower bound of energy consumption). *If the classical computer in Fig. 3 is allowed to interact with an environment at inverse temperature $\beta > 0$ for memory reinitialization, any computation in Fig. 3 requires an average energy consumption of at least*

$$\mathcal{W} \geq \mathcal{W}^{(C)} = \frac{S(C|a)}{\beta}, \quad (93)$$

in the limit of closing the thermodynamic cycle $\varepsilon \rightarrow 0$, where $S(C|a) := \sum_{a=0,1} p(a)S(\rho_a^C)$ is the conditional entropy of the classical-quantum state for the computer's state ρ_a^C before the erasure \mathcal{E} conditioned on the measurement outcome $a \in \{0,1\}$ of the output register, $p(a)$ is the probability distribution of the measurement outcome, and $S(\rho) := -\text{tr}[\rho \ln \rho]$.

To prove the lower bound in Theorem 5, we have already introduced most of the essential techniques in formulating the framework of computation in Sec. II. What remains for completing the proof is to assemble these techniques, which yields the following concise proof.

Proof of Theorem 5. We prove Theorem 5 by starting with some general energetic considerations along the lines of the thermodynamic diagram in Fig. 1. As introduced in Definition 1, \mathcal{W} is the energy consumption of the computer including reinitialization at the end. Together with the heat dissipation into the environment \mathcal{Q} , we have energy conservation in (3), i.e.,

$$\mathcal{W} = \mathcal{Q} = \mathcal{Q}_E + \mathcal{Q}_{E'}, \quad (94)$$

where the second equality is given by (28). Recall that the heat dissipation $\mathcal{Q}_{E'}$ due to control costs is by definition non-negative as shown in (27), and thus, the initialization cost \mathcal{Q}_E provides a lower bound

$$\mathcal{W} \geq \mathcal{Q}_E. \quad (95)$$

As shown in (58), Landauer's principle [11, 23, 82] guarantees that, for reinitialization of the computer's memory using a heat bath at inverse temperature $\beta > 0$, the heat dissipation \mathcal{Q}_E into the environment is bounded by

$$\beta \mathcal{Q}_E = \Delta S + D(E'||E) + I(E' : S') \geq \Delta S, \quad (96)$$

where we use $D(E'||E) \geq 0$ and $I(E' : S') \geq 0$. Note that ΔS is the change in entropy of the state of all the W bits in the computer to be erased; as discussed in Sec. III A, repeating the single-bit Landauer-erasure protocol W times in total (analyzed for deriving the upper bound) may be more costly as a whole than the erasure of all the W bits collectively [70, 71]. Thus, the simultaneous erasure of all W bits is analyzed here to derive the general lower bound based on Landauer's principle.

To write down ΔS more explicitly, we recall the fact that the agent has access to the output $a \in \{0,1\}$ stored in the output register before performing the erasure \mathcal{E} in Fig. 3, where a may, in general, be correlated with the computer's state ρ_a^C before the erasure conditioned on a . Even though the single-bit information in the output may not affect the asymptotic scaling, it is still necessary to take it into account to ensure that we bound the required energy consumption from below. Thus, ΔS on average is written as the conditional entropy

$$\Delta S \geq S(C|a) - S(C'|a) \rightarrow S(C|a) \quad \text{as } \varepsilon \rightarrow 0, \quad (97)$$

where $S(C'|a)$ is the conditional entropy of the classical-quantum state for the computer's state after the erasure \mathcal{E} conditioned on a , and we have $S(C'|a) \rightarrow 0$ as $\varepsilon \rightarrow 0$ since the computer's state after the erasure in (55) is a fixed pure state in the limit. These bounds yield

$$\mathcal{Q}_E \geq \frac{\Delta S}{\beta} \geq \frac{S(C|a)}{\beta}, \quad (98)$$

which, together with (95), implies the energy-consumption inequality

$$\mathcal{W} \geq \frac{S(C|a)}{\beta} \quad (99)$$

to prove Theorem 5. \square

One may wonder under which circumstances the bound provided by Theorem 5 on the energy consumption is strictly gapped away from zero, i.e., where the non-zero entropy comes from. The origin lies in the query-complexity setting, where usually, the oracle O_f is drawn randomly from some set of functions f specified by the problem to be solved. Then, the agent performs computation to learn some property of f via the queries to O_f ,

⁹ In our analysis of the lower bound of energy consumption of classical computation, we do not only analyze the scaling but provide the lower bound $\frac{1}{\beta}S(C|a)$ explicitly, i.e., including the constant factors, which is important for explicitly providing the goals of the experimental demonstrations shown by Table I of Sec. V. The agent's knowledge of the 1-bit information a of the result of the computation may be asymptotically negligible, but our analysis explicitly takes a into account to provide the lower bound including the constant factors.

but it is not necessarily required to specify the function f itself in full detail. Conditioned on the algorithm's output a , we have a probability distribution $p(f|a)$ according to which the function f (i.e., the oracle O_f) is chosen. A nonzero entropy of this probability distribution $p(f|a)$ leaves a nonzero entropy $S(C|a)$ in the computer's memory. Then, to erase this information obtained from the (randomly chosen) oracle, nonzero energy consumption is inevitably required due to Laudauer's principle. This analysis indicates that the entropic lower bound $S(C|a)/\beta$ of energy consumption in Theorem 5 is determined not only by the lower bound of query complexity but also by the probability distribution $p(f|a)$ for the oracle.

In Sec. IV, we will derive how the entropy arises and where the connection to $p(f|a)$ lies. Then, we employ the bound in Theorem 5 to prove the exponential energy-consumption advantage of quantum computation over classical computation in solving a computational task, in particular, Simon's problem [12, 13].

IV. EXPONENTIAL QUANTUM ADVANTAGE IN ENERGY CONSUMPTION

In this section, using the upper and lower bounds on energy consumption shown in Sec. III, we prove that the energy consumption of quantum computation can be exponentially smaller than that of classical computation within our framework. In Sec. IV A, we recall the standard formulation for Simon's problem given in Refs. [12, 13]. In Sec. IV B, we apply Theorems 3 and 4 to show that the energy consumption of quantum computation to solve Simon's problem scales polynomially with the problem size, and in Sec. IV C, we use Theorem 5 to prove that any classical algorithm for solving the same problem requires an energy consumption that scales exponentially in the problem size.

A. Definition of Simon's problem

Simon's problem [12, 13] is an exemplary case where the query complexity in quantum computation is exponentially separated from any classical computation. The structure of the quantum algorithm to solve Simon's problem is closely related to the ones to solve the non-oracle-based (i.e., non-relativized) problems of the discrete logarithm and the integer factorization [83]. The issue with the latter non-relativized problems is that lower bounds on the complexity of classical algorithms to solve them are notoriously hard to show in general because they mainly boil down to the open question of whether there exists a classical polynomial-time algorithm to solve the integer-factorization problem [84, 85]. Studying Simon's problem allows for a robust analysis, where lower bounds on the complexity of classical algorithms do not depend on the conjectured difficulty of a computational

task. Our analysis, therefore, will lay the groundwork for studying the other problems on a case-by-case basis.

Simon's problem [12, 13] may be posed as the computational task to search for an unknown bit string or, in the way we will define it in Problem 6, as the task of deciding whether an unknown function $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$ with N -bit input and N -bit output is 1-to-1 or 2-to-1 mapping. The computer solving the problem has access to an oracle O_f , to obtain the output of f for a given input x . For the *classical algorithm*, the oracle is defined as a map in (7) and Fig. 2(a), i.e.,

$$O_f^{(C)}(x, 0) = (x, f(x)), \quad (100)$$

for all inputs $x \in \{0, 1\}^N$ on the first N bits, with the second N bits initialized in 0. When it comes to the *quantum oracle*, on the other hand, the conventional prescription of the oracle is a unitary map that transforms each computational-basis state $|x, y\rangle$ as

$$O_f^{(Q)}(|x, y\rangle) = |x, y \oplus f(x)\rangle, \quad (101)$$

as shown in (8) and Fig. 2(b). For superposition states of $\sum_{x,y} \alpha_{x,y} |x, y\rangle$, the quantum oracle acts linearly on each term of $|x, y\rangle$.

Without further ado, we present the definition of Simon's problem as follows.

Problem 6 (Simon's problem [12, 13]). *Let $N > 0$ be the problem size and O_f the oracle for a function $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$ constructed in the following way: choose an N -bit string $s \in \{0, 1\}^N \setminus \{0\}$ and a single-bit state $b \in \{0, 1\}$ uniformly at random. If $b = 0$, then f is a uniformly randomly chosen 1-to-1 function. If $b = 1$, then f is a uniformly randomly chosen 2-to-1 function satisfying, for all $x \in \{0, 1\}^N$,*

$$f(x) = f(x \oplus s), \quad (102)$$

where \oplus is the bitwise exclusive OR. The goal is to output $a \in \{0, 1\}$ that estimates b correctly, i.e., $a = b$, with a high probability greater than $2/3$ on average, using as few calls to the oracle O_f as possible, where the average success probability is taken over the choices of s , b , and f .

In the following, we will apply the techniques developed in Sec. III to show that there exists a quantum algorithm for solving Problem 6 whose energy consumption is exponentially smaller than that required for any classical algorithm. In particular, we use Theorem 4 to show that a quantum computer can solve Simon's problem at an energy consumption $\mathcal{W} \leq \mathcal{W}^{(Q)} = O(\text{poly}(N))$ as the problem size N grows (Corollary 7), and we apply Theorem 5 to show that all classical algorithms solving the same problem consume the energy of at least $\mathcal{W} \geq \mathcal{W}^{(C)} = \Omega(2^{N/2}N)$ (Corollary 9).

B. Quantum upper bound on energy consumption for solving Simon's problem

In this section, we analyze the energy consumption of a quantum algorithm that solves Simon's problem as defined in Problem 6 in Sec. IV A. In particular, as a corollary of Theorem 4, we show in Corollary 7 by summing up all the contributions of the upper bound of energy consumption in Theorem 4, that a fault-tolerant quantum computer can solve Simon's problem with only a polynomial amount of energy consumption in the problem size N . We also remark that the result shows that the polynomial amount of energy consumption in N is achievable for quantum computation solving Simon's problem even if the parameter $E_{\text{ctrl.}}$ for the control cost in (38) grows polynomially in N .

Corollary 7 (Upper bound of energy consumption for quantum algorithm to solve Simon's problem). *Let $N > 0$ be the problem size and the oracle O_f be chosen according to the probability distribution for Simon's problem in Problem 6. Then, there exists a fault-tolerant quantum algorithm for solving Simon's problem with energy consumption bounded by*

$$\mathcal{W} \leq \mathcal{W}^{(Q)} = O(E_{\text{ctrl.}} \times \text{poly}(N)). \quad (103)$$

Proof. In deriving the general upper bound of energy consumption of quantum computation in Theorem 4 of Sec. III A, we have accounted for all the contributions to the work cost of performing a general quantum computation according to the framework in Fig. 3. For the proof here, the remaining point to be done is to translate the algorithm presented in Refs. [12, 13] (for introductory explanations, see also Refs. [86, 87]) to our computational framework, so as to calculate the bound in W and D_{comp} (i.e., D_k , and M in the definition (17) of D_{comp}) accordingly.

In the quantum algorithm for solving Simon's problem in Refs. [12, 13], a quantum subroutine *Fourier-twice* is repeatedly used. This quantum subroutine acts on $2N$ qubits initialized as $|0\rangle^{\otimes 2N}$ with the operations in the following order: (i) N parallel Hadamard gates act on the first N bits, i.e., $V = H^{\otimes N} \otimes \mathbb{1}^{\otimes N}$ with circuit depth $O(1)$, then, (ii) the oracle O_f acts on all $2N$ qubits simultaneously and finally, (iii) N parallel Hadamard gates V are again applied on the first N qubits. In the overall quantum algorithm, this subroutine is carried out $O(N)$ times; that is, the oracle O_f (called once in (ii) per performing this subroutine) is called in total M times with

$$M = O(N). \quad (104)$$

In fact, it has been shown that this is also the best a quantum algorithm can do; i.e., any quantum algorithm solving Simon's problem has to query O_f at least $O(N)$ times [88].

In the algorithm as structured in Fig. 3, M sets of $2N$ qubits are initialized at first, i.e.,

$$W = 2N \times M = O(N^2), \quad (105)$$

and the subroutine is sequentially applied to the m th set of the $2N$ qubits for $1 \leq m \leq M$. In this case, U_1 is of the form

$$U_1 = V \otimes \mathbb{1}^{\otimes 2N(M-1)}, \quad (106)$$

and all remaining U_k for $2 \leq k \leq M$ are of the form

$$U_k = \mathbb{1}^{\otimes 2N(k-2)} \otimes V \otimes V \otimes \mathbb{1}^{\otimes 2N(M-k)}, \quad (107)$$

where $\mathbb{1}^{\otimes 0} = 1$. The last unitary U_{M+1} equals the operation $\mathbb{1}^{\otimes 2N(M-1)} \otimes V$ followed by a quantum implementation $U_{\text{post-processing}}$ of the classical post-processing of the resulting state obtained from M repetitions of the subroutine, i.e.,

$$U_{M+1} = U_{\text{post-processing}}(\mathbb{1}^{\otimes 2N(M-1)} \otimes V), \quad (108)$$

where U_{M+1} can be performed within time complexity $O(N^3)$ [87]. Therefore, using, e.g., the Clifford+ T universal gate set, each of these unitaries can be implemented efficiently with circuit depths, for all $1 \leq k \leq M$,

$$D_k = O(1), \quad (109)$$

and

$$D_{M+1} = O(N^3). \quad (110)$$

Note that the classical post-processing in U_{M+1} can be represented by a reversible classical logic circuit, e.g., written in terms of the classical Toffoli gate, and by replacing each classical Toffoli gate with the quantum Toffoli gate (and further rewrite each quantum Toffoli gate in Clifford and T gates if we use the Clifford+ T universal gate set), we obtain the quantum circuit of $U_{\text{post-processing}}$ of depth $O(N^3)$.

Therefore, in the ideal case without gate error, inserting the explicit numbers in (104), (105), (109), and (110) into the upper bound of the energy consumption in Theorem 3, we find that the energy consumption in the quantum algorithm in Fig. 3 to solve Simon's problem is bounded by

$$\mathcal{W} \leq \mathcal{W}^{(Q)} = O(E_{\text{ctrl.}} N^5), \quad (111)$$

where the dominant part is the one for implementing U_{M+1} by a circuit of volume $WD_{M+1} = O(N^5)$, and β , η , and ε are set as constants. Even in the case where we implement this quantum algorithm by a fault-tolerant protocol without measurement in Theorem 4, the energy consumption in implementing the quantum algorithm to solve Simon's problem in a fault-tolerant way is bounded by

$$\mathcal{W} \leq \mathcal{W}^{(Q)} = \tilde{O}(E_{\text{ctrl.}} N^8), \quad (112)$$

where the dominant part is the one for implementing U_{M+1} by a circuit of volume $\tilde{O}(WD_{M+1} D_{M+1}) = \tilde{O}(N^8)$, and \tilde{O} may ignore polylogarithmic factors. In any of these cases, we obtain

$$\mathcal{W} \leq \mathcal{W}^{(Q)} = O(E_{\text{ctrl.}} \times \text{poly}(N)), \quad (113)$$

which scales polynomially in N so long as the parameter $E_{\text{ctrl.}}$ for the control cost also scales polynomially in N . \square

C. Classical lower bound on energy consumption for solving Simon's problem

In this section, we use the lower bound of energy consumption in Theorem 5 to prove that any classical computation that can solve Simon's problem (Problem 6) should require an exponentially large energy consumption in the problem size N , no matter how energy-efficiently the algorithm is implemented, as shown in Corollary 9 below.

The essential idea for deriving this lower bound of the energy consumption is to establish a connection between the conventional argument for a lower bound of query complexity of classical algorithms for solving Simon's problem [12, 13] and the entropic lower bound of the energy consumption that we derived in Theorem 5. In Problem 6, $s \in \{0, 1\}^N \setminus \{0\}$, $b \in \{0, 1\}$, and f are chosen uniformly at random, and the classical algorithm makes M queries to the oracle O_f using M different classical inputs $x_1, \dots, x_M \in \{0, 1\}^N$ to estimate b , i.e., to output $a \in \{0, 1\}$ satisfying $a = b$, with a high probability greater than $2/3$. We write these M inputs to O_f collectively as

$$\vec{x} := (x_1, \dots, x_M). \quad (114)$$

Note that a random guess of b would only achieve the average success probability of $1/2$, but Problem 6 requires a strictly better average success probability $2/3 = 1/2 + \Delta$ for $\Delta = 1/6$.

We here provide a necessary condition on M for achieving the average success probability $2/3$ in Simon's problem as required in Problem 6. In Refs. [12, 13], it has been shown that, in order to have the average success probability $1/2 + 2^{-N/2}$, it is necessary to make M queries with at least $M \geq 2^{N/4}$; later in Refs. [86, 87], this lower bound has been improved by proving that a larger lower bound $M = \Omega(2^{N/2})$ should be necessary for achieving the average success probability $1/2 + \Omega(1)$, but the constant factors of the lower bounds were not explicitly presented. Refining these analyses, we here show that, for any problem size N and required average success probability $1/2 + \Delta$, we have a lower bound of the query complexity of Simon's problem with an explicit constant factor given by $\left\lceil \sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2} \right\rceil$, where $\Delta \in (0, 1/2)$ is an arbitrary constant in this parameter region and is suitably tunable for our analysis. In particular, we prove the following proposition, where the proof is given in Appendix A 3.

Proposition 8 (Refined analysis of lower bound of query complexity of Simon's problem for classical algorithms). *For any parameter $\Delta \in (0, 1/2)$, if we have*

$$M < M_{N,\Delta} := \left\lceil \sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2} \right\rceil, \quad (115)$$

the average success probability of any probabilistic classical algorithm with at most M queries for Simon's problem

in Problem 6 is upper bounded by

$$\frac{1}{2} + \Delta. \quad (116)$$

Due to this proposition, the required number M of the oracle queries for the classical algorithms to solve Simon's problem is lower bounded, for any N and any choice of $\Delta \in (0, 1/2)$ in Proposition 8, by

$$M \geq M_{N,\Delta} = \Omega(2^{N/2}) \quad (117)$$

with a high probability, where we quantify this probability more precisely in Lemma 12 of Appendix A 3.

In the following corollary on the lower bound of energy consumption, we argue that the above query-complexity bound can be used for evaluating a dominant contribution to the entropy of the state of the computer. Thus, as the corollary of Theorem 5, we have the lower bound of the energy consumption of classical computation. Our lower bound of the energy consumption for any classical algorithm for solving Simon's problem scales $\Omega(2^{N/2}N)$ (with explicit constant factors derived in (146) and (149) in the proof), which is different from that of the query complexity $\Omega(2^{N/2})$ for the same problem (i.e., those in Refs. [86, 87] and Proposition 8). Indeed, as we will show later in Proposition 10, the upper bound of query complexity of Simon's problem for classical computation is also $O(2^{N/2})$ and thus is strictly smaller than the lower bound of energy consumption $\Omega(2^{N/2}N)$. As we have argued in Sec. III B, this difference arises because the query complexity only counts the number of queries, but the lower bound of the energy consumption of classical computation is determined by the amount of information (i.e., entropy) obtained from the queries, depending also on the probability distribution of the oracle.¹⁰ In this regard, the difference in the scaling of the energy consumption and the query complexity of classical computation solving Simon's problem quantitatively illustrates the fact that the energy consumption and the number of queries are different computational resources despite the resemblance in the proof techniques.

Corollary 9 (Lower bound of energy consumption for all classical algorithms to solve Simon's problem). *Let $N > 0$ be the problem size of Simon's problem in Problem 6.*

¹⁰ If one divides the entropy by the number of queries, one may be able to define the entropy per query, but it still remains unclear whether one could directly evaluate a lower bound of the entropy per query without analyzing the overall entropy for all the queries. In the case of Simon's problem, the entropy per query scales linearly with the problem size N ; however, this scaling depends on the probability distributions of the oracles used in the computational problems and thus may differ in general. Therefore, we need to evaluate a lower bound of the overall entropy, not just the query-complexity lower bound, to derive a lower bound of energy consumption.

Any probabilistic and adaptive classical algorithm to solve Simon's problem has an energy consumption of at least

$$\mathcal{W} \geq \mathcal{W}^{(C)} = \Omega\left(\frac{2^{N/2}N}{\beta}\right). \quad (118)$$

Proof. As formulated in Sec. II, the crucial assumption for our analysis is that the knowledge on the input-output relations for f obtained from the oracle queries in the classical computation remains in the computer, up to possible reversible transformations U_k , until erased with the thermodynamical work cost. As for the choice of the input sequence \vec{x} in (114), the classical algorithm can be probabilistic and adaptive in general. For our analysis to cover these general cases, we consider the classical algorithms where the classical input x_{k+1} to the oracle O_f is sampled probabilistically and can depend on the previous inputs x_1, \dots, x_k and the oracle outputs $f(x_1), \dots, f(x_k)$. Moreover, after each query to the oracle, the algorithm may probabilistically decide to stop or to continue further querying the oracle. In particular, let $q_1(x_1)$ denote the probability distribution on the choice of x_1 , and for each $k > 1$, conditioned on the previous oracle calls with input x_ℓ and output function value $f(x_\ell)$ for $1 \leq \ell < k$, we write the conditional probability distribution on the next choice of x_k as $q_k(x_k | \{(x_\ell, f(x_\ell)) : 1 \leq \ell < k\})$. The algorithm uses M different inputs, and thus, if x_k has already appeared

previously in x_1, \dots, x_{k-1} , i.e., $x_k \in \{x_\ell : 1 \leq \ell < k\}$, then we have $q_k(x_k | \{(x_\ell, f(x_\ell)) : 1 \leq \ell < k\}) = 0$. For a given sequence $((x_1, f(x_1)), \dots, (x_k, f(x_k)))$ of k queries, the algorithm stops with probability

$$q(\text{stop} | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq k\}), \quad (119)$$

or continues for the $(k+1)$ th query with probability

$$\begin{aligned} q(\text{continue} | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq k\}) \\ := 1 - q(\text{stop} | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq k\}). \end{aligned} \quad (120)$$

Since we require that $x_1, \dots, x_k \in \{0, 1\}^N$ should be k different inputs among 2^N elements, it holds by construction that

$$q(\text{continue} | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq 2^N\}) = 0, \quad (121)$$

which is the requirement that the algorithm must eventually terminate with at most 2^N queries. The overall strategy of the classical algorithm is determined by the strategy for choosing \vec{x} , i.e., q_1, q_2, \dots , and that for stopping, i.e., $q(\text{stop} | \dots)$ and $q(\text{continue} | \dots)$. Given this general formulation of strategies in choosing the input sequence \vec{x} to the oracle, we can write the computer's memory state before outputting a in Fig. 3, conditioned on having queried the function oracle exactly M times, up to reversible transformation, as

$$\begin{aligned} \rho_{f,M} = \frac{1}{p_f(M)} \sum_{\vec{x}} \left(q(\text{stop} | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq M\}) q_M(x_M | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq M-1\}) \right. \\ \left. \prod_{k=1}^{M-1} q(\text{continue} | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq k\}) q_k(x_k | \{(x_\ell, f(x_\ell)) : 1 \leq \ell \leq k-1\}) \right) \rho_{f,\vec{x}}. \end{aligned} \quad (122)$$

where $p_f(M)$ is the probability of the strategy yielding exactly M queries defined via normalization $\text{tr}[\rho_{f,M}] = 1$, we may write $q_1(x_1)$ as $q_k(x_k | \dots)$ with $k = 1$ for simplicity of notation,

$$\rho_{f,\vec{x}} := \bigotimes_{k=1}^M |x_k, f(x_k)\rangle \langle x_k, f(x_k)| \quad (123)$$

is a $(2N \times M)$ -bit classical state, i.e., a diagonal density operator, representing the M different input-output relations, and we here write only the relevant $(2N \times M)$ -bit part of the state depending on f and \vec{x} out of the W bits in the computer.

In general, for fixed s , b , and f in Problem 6, the required number of queries M for the algorithm to succeed in estimating b may change probabilistically, depending on the strategy of the algorithm, i.e., q_1, q_2, \dots , $q(\text{stop} | \dots)$, and $q(\text{continue} | \dots)$. We write the average

state before the output, in expectation taken over all possible number of queries $M \geq 1$, as

$$\rho_f = \sum_{M \geq 1} p_f(M) \rho_{f,M}. \quad (124)$$

For the agent before seeing any output from the computer, this state is further averaged over the uniform probability distributions $p(s)$, $p(b)$, $p(f|s, b)$ as described in Problem 6, i.e., up to reversible transformation,

$$\rho^C = \sum_{s,b} \sum_f p(s)p(b)p(f|s, b)\rho_f \quad (125)$$

$$= \sum_{M \geq 1} \sum_{s,b} \sum_f p(s)p(b)p(f|s, b)p_f(M)\rho_{f,M} \quad (126)$$

$$= \sum_{M \geq 1} p(M)\rho_M, \quad (127)$$

where we write in the last line

$$p(M) := \sum_{s,b} \sum_f p(s)p(b)p(f|s,b)p_f(M), \quad (128)$$

$$\rho_M := \frac{1}{p(M)} \sum_{s,b} \sum_f p(s)p(b)p(f|s,b)p_f(M)\rho_{f,M}. \quad (129)$$

In this notation, $p(M)$ represents the probability of the algorithm stopping exactly after M queries on average over the choices of s , b , and f , and ρ_M is the state of the computer's memory before the output, conditioned on the algorithm stopping exactly after M queries, on average over s , b , and f .

To output a in Fig. 3, the agent uses the output register prepared in $|0\rangle$ and performs a swap gate to pull out one of the W bits from the computer to the output register, where the state of the bit in the computer on which the swap gate acts becomes $|0\rangle$ after these output operations. After the swap gate, the $(W+1)$ -bit state of the computer and the output register is equivalent to $\rho^C \otimes |0\rangle\langle 0|$ up to reversible transformation. A measurement of the output register is then performed to obtain the single-bit output $a \in \{0, 1\}$ with probability $p_{\text{out}}(a)$. For the agent who knows the output a of the computation, before performing the erasure \mathcal{E} conditioned on a , the computer's state is in a W -qubit mixed state ρ_a^C , which is a classical state represented by a diagonal density operator obtained by measuring one of the $W+1$ qubits in the basis $\{|a\rangle : a = 0, 1\}$. As a result, for any agent erasing the computer's state, the probabilistic state of the computer's memory has the conditional entropy $S(C|a)$ of the state

$$\rho^{Ca} := \sum_a p_{\text{out}}(a) |a\rangle\langle a| \otimes \rho_a^C. \quad (130)$$

Since the entropy is invariant under the reversible transformation, we have

$$S(\rho^C \otimes |0\rangle\langle 0|) = S(\rho^{Ca}). \quad (131)$$

Thus, by definition of the conditional entropy, it holds that

$$S(C|a) = S(\rho^C \otimes |0\rangle\langle 0|) - S(a) \quad (132)$$

$$\geq S(\rho^C) - \ln 2 \quad (133)$$

$$\geq \sum_{M \geq 1} p(M)S(\rho_M) - \ln 2, \quad (134)$$

where the $\ln 2$ comes from $S(a) \leq \ln 2$ for a single-bit state, and the last line is due to (127) and the concavity of the entropy [9].

With the bound (134), we evaluate the entropic lower bound in Theorem 5, by taking into account the assumption that the algorithm should stop successfully. We are proving the entropic lower bound for *all* classical algorithms with success probability greater than $2/3$, and

importantly, our analysis takes into account all possible classical algorithms that try to minimize energy consumption at their best. For example, if the input sequence $\vec{x} = (x_1, \dots, x_M)$ obtained from M queries includes some pair (x_k, x_m) satisfying $f(x_k) = f(x_m)$, then the algorithm can know that f is a 2-to-1 function with $s = x_k \oplus x_m$. But even in such a case, the algorithm does not have to stop immediately but is still allowed to continue queries and computation, e.g., for trying to uncompute some part of the computer's memory so as to reduce the overall energy consumption at best. On the other hand, as long as the overall success probability is greater than $2/3$, the algorithm may stop even before making the $M_{N,\Delta}$ queries with some small probabilities to save the energy consumption (we allow such stops even if the input sequence \vec{x} with $M < M_{N,\Delta}$ does not contain any pair satisfying $f(x_k) = f(x_m)$). Our proof covers all these cases, to find the fundamental lower bound on energy consumption for all possible classical algorithms to solve Simon's problem 6.

To arrive at a lower bound on the entropy of the computer's state, we have to constrain the possible probability distributions $p(M)$ and the entropy $S(\rho_M)$ in (134). To bound $p(M)$ in (134), we use the requirement in Problem 6 that the overall success probability of the algorithm should be greater than $2/3$. Indeed, as we show in Lemma 12 in Appendix A 3, if the success probability is greater than or equal to $1/2 + \Delta$ for any parameter $\Delta \in (0, 1/2)$, the probability of making at least $M_{N,\Delta}$ queries is bounded from Proposition 8 by

$$\sum_{M \geq M_{N,\Delta}} p(M) \geq \frac{1 - 6\Delta}{3 - 6\Delta}. \quad (135)$$

The lower bound on the right-hand side depends on the choice of Δ . In Problem 6, we demand that the algorithm for solving this problem must have a success probability greater than $2/3$. Any such algorithm will also have a success probability greater than $1/2 + \Delta$ if the choice of Δ is in the range

$$\Delta \in \left(0, \frac{1}{6}\right), \quad (136)$$

i.e.,

$$\frac{1}{2} + \Delta < \frac{2}{3}. \quad (137)$$

With Δ in this range, the lower bound on the right-hand side of (135) applies to all the algorithms solving Problem 6 with a success probability greater than $2/3$; also, in this range, the right-hand side of (135) indeed becomes positive.

As for the bound of $S(\rho_M)$ in (134), in Proposition 13 of Appendix A 3, we show explicitly that, for any M , the entropy of ρ_M is bounded by

$$S(\rho_M) \geq \frac{1}{2} \ln \frac{2^N!}{(2^N - M)!}, \quad (138)$$

regardless of whether ρ_M consists of good or bad input sequences \vec{x} ; in particular, since the right-hand side of (138) is monotonously increasing in $M = 1, 2, \dots$, we can write, for all $M \geq M_{N,\Delta}$,

$$S(\rho_M) \geq \frac{1}{2} \ln \frac{2^N!}{(2^N - M_{N,\Delta})!}. \quad (139)$$

Inserting (135) and (138) into (134), we obtain

$$\sum_{M \geq 1} p(M)S(\rho_M) = \sum_{M < M_{N,\Delta}} p(M)S(\rho_M) + \sum_{M \geq M_{N,\Delta}} p(M)S(\rho_M) \quad (140)$$

$$\geq 0 + \frac{1}{2} \ln \frac{2^N!}{(2^N - M_{N,\Delta})!} \times \sum_{M \geq M_{N,\Delta}} p(M) \quad (141)$$

$$\geq \frac{1}{2} \ln \frac{2^N!}{(2^N - M_{N,\Delta})!} \times \frac{1 - 6\Delta}{3 - 6\Delta} \quad (142)$$

$$= \frac{1 - 6\Delta}{6 - 12\Delta} \ln \frac{2^N!}{(2^N - M_{N,\Delta})!}. \quad (143)$$

Furthermore, as we show in Lemma 14 of Appendix A 3, for $M_{N,\Delta}$ in (115), a variant of Stirling's approximation [89] yields

$$\ln \frac{2^N!}{(2^N - M_{N,\Delta})!} \geq \sqrt{\frac{2\Delta}{1 + \Delta}} 2^{N/2} N \ln 2 - 3. \quad (144)$$

Combining the results from (134), (143) and (144), for all classical algorithms that solve Simon's problem in our framework, we have

$$S(C|a) \geq \frac{1 - 6\Delta}{6 - 12\Delta} \left(\sqrt{\frac{2\Delta}{1 + \Delta}} 2^{N/2} N \ln 2 - 3 \right) - \ln 2 = \Omega(2^{N/2} N), \quad (145)$$

which is gapped away from zero independently of how energy-efficiently the classical algorithm is implemented, as long as the algorithm solves Simon's problem.

Therefore, for any choice of the constant parameter $\Delta \in (0, 1/6)$ in (136), applying Theorem 5 yields a lower bound of the energy consumption

$$\mathcal{W} \geq \frac{1}{\beta} S(C|a) \geq \frac{1}{\beta} \times \left(\frac{1 - 6\Delta}{6 - 12\Delta} \left(\sqrt{\frac{2\Delta}{1 + \Delta}} 2^{N/2} N \ln 2 - 3 \right) - \ln 2 \right) = \Omega\left(\frac{2^{N/2} N}{\beta}\right). \quad (146)$$

In particular, the maximum of the constant factor

$$\frac{1 - 6\Delta}{6 - 12\Delta} \sqrt{\frac{2\Delta}{1 + \Delta}} \quad (147)$$

of the leading order in (146) is achieved with the choice of

$$\Delta = 2 - \frac{\sqrt{15}}{2} = 0.0635 \dots \in \left(0, \frac{1}{6}\right), \quad (148)$$

and thus, we can explicitly obtain the tightest lower bound among the choices of $\Delta \in (0, 1/6)$ as

$$\mathcal{W} \geq \frac{1}{\beta} \times \left(\frac{3\sqrt{15} - 11}{6\sqrt{15} - 18} \left(\sqrt{\frac{8 - 2\sqrt{15}}{6 - \sqrt{15}}} \times 2^{N/2} N \ln 2 - 3 \right) - \ln 2 \right). \quad (149)$$

These lower bounds conclude the proof. \square

V. PROPOSAL ON EXPERIMENTAL DEMONSTRATION

In this section, we propose an experimental scenario for demonstrating the energy-consumption advantage of

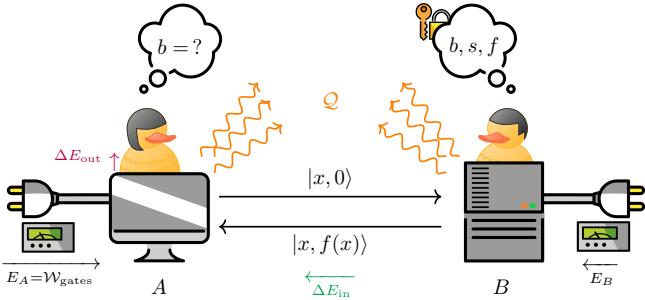


FIG. 4. A schematic depiction of the setup to experimentally demonstrate the exponential energy-consumption advantage of quantum computation over classical computation based on solving Simon’s problem. The setup is comprised of two parties: B who keeps a bit $b \in \{0, 1\}$ and an N -bit string $s \in \{0, 1\}^N \setminus \{0\}$ as B ’s secret, and A who tries to estimate B ’s secret b , as in Simon’s problem in Problem 6. With b and s chosen uniformly at random, B initially constructs a function f in such a way that B can implement the oracle O_f feasibly within a polynomial time for quantum algorithms to solve Simon’s problem, yet A cannot distinguish f from the uniformly random choice as required for Simon’s problem in Problem 6, using a cryptographic primitive as described in the main text. To solve Simon’s problem, A runs the quantum algorithm for solving Simon’s problem on A ’s quantum computer, and B acts as the oracle O_f . Throughout the computation, B ’s internal implementation of O_f is never directly revealed to A ; to estimate B ’s secret b , A can query B by sending a superposition of states $|x, 0\rangle$ to B , and B internally performs O_f and sends back the corresponding superposition of $|x, f(x)\rangle$ to A . These queries provide the only information A can access for guessing b , as in Simon’s problem. What we care about in this experiment is the energy consumption. By measuring A ’s energy consumption E_A and B ’s energy consumption E_B in solving Simon’s problem, we obtain an achievable energy consumption for the implemented quantum computation from their sum $E_A + E_B$, as shown in (153). In the main text, we also clarify how to estimate the energy consumption for classical computation to solve the same problem, using extrapolation or the fundamental lower bound (Table I). The experimental demonstration of the energy-consumption advantage of quantum computation will be successful if $E_A + E_B$ is smaller than the estimate of energy consumption for classical computation.

quantum computation over classical computation in solving Simon’s problem analyzed in Sec. IV. In the following, we begin by presenting a schematic setup for measuring the energy consumption of quantum computation in experiments. Then, we clarify the explicit polynomial-time way of instantiating the oracle for Simon’s problem to implement this setup. Finally, we also show approaches to estimate the energy consumption for classical computation to solve the same problem, which yields criteria to be compared with the results of the quantum experiments.

Setup for measuring the energy consumption of quantum computation in experiments. We propose a schematic setup depicted in Fig. 4 to capture our frame-

work in Fig. 3 and measure the energy consumption of quantum computation therein. A challenge in the experimental realization of the framework in Fig. 3 is the implementation of the oracle as required in Sec. II B. To address this challenge, the setup in Fig. 4 consists of two parties A and B with their quantum computers, in a setup similar to that of a client and a server: A is where the main computation is carried out, and B is where the oracle is instantiated without revealing its internal to A . Quantum communication between A and B in Fig. 4 implements the query to the oracle in Fig. 3. For Simon’s problem, i.e., Problem 6, B selects a bit $b \in \{0, 1\}$ and a bit string $s \in \{0, 1\}^N \setminus \{0\}$ of length N uniformly at random. Based on b and s , B constructs an N -bit function $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$ in such a way that B can internally implement the oracle O_f for f . The particular choice of b , s , and f is kept as B ’s secret so that A should not be able to distinguish B ’s construction of f from the uniformly random choice of f as in Problem 6.

In this two-party setup, the task is for A to estimate the bit b that B has internally. To achieve this task, A performs the quantum algorithm for Simon’s algorithm, i.e., that analyzed in Sec. IV B. One way to implement the initialization \mathcal{E} in Fig. 3 is to use our finite-step Laudauer-erasure protocol for Theorem 2, but one can also use other initialization techniques available in the experiments. Without the finite-step Laudauer-erasure protocol, the upper bound of the achievable initialization and control costs may not be theoretically guaranteed, but the experiment will still be successful as long as the energy consumption of the implementation of the quantum algorithm is below an estimate of that of the classical algorithms. Each query to the oracle in the quantum algorithm, implemented by a swap gate in Figs. 2(b) and 3, is replaced with that implemented by (quantum) communication between A and B , as shown in Fig. 4. Throughout the experiment, B never directly reveals b to A , but B only receives the input state for the query from A and returns the corresponding output state of the oracle O_f to A by using B ’s internal implementation of O_f . In other words, in implementing this experiment in a lab, imagine that the party A has some program to run its computation, and also the party B has some program to compute the function f of the oracle; then, in the lab, we require that A ’s program should not directly read b in B ’s internal program but should estimate b through the queries. Thus, in the setup of Fig. 4, A can estimate B ’s secret b only from the input-output relations of O_f obtained from the queries implemented by communication, in the same way as estimating b from those obtained from the oracle queries in Fig. 3.

To measure the achievable energy consumption of quantum computation in this setup, i.e., $\mathcal{W}_{\text{gates}} + \Delta E^{(\text{in})} - \Delta E^{(\text{out})}$ in Definition 1, it suffices to measure upper bounds of $\mathcal{W}_{\text{gates}} - \Delta E^{(\text{out})}$ and $\Delta E^{(\text{in})}$, so as to sum them up. For this purpose, we keep track of the amount of energy supply that is used for the implementation of A ’s computation, and that for the implementation

on B 's side. For example, these amounts can be obtained by reading the difference of the electricity meters before and after conducting the computation, which shows the electricity consumption. Also, on the output system in the experiment, we set the state 0 as a low-energy state and 1 as a high-energy state, so that it should hold that

$$\Delta E^{(\text{out})} \geq 0 \quad (150)$$

in the experiment. With this setup, the energy consumption E_A on A 's side is, by construction,

$$E_A = \mathcal{W}_{\text{gates}} \geq \mathcal{W}_{\text{gates}} - \Delta E^{(\text{out})}. \quad (151)$$

Similarly, the energy consumption E_B on B 's side is the sum of the energetic cost $\Delta E^{(\text{in})}$ and the control cost of implementing the oracle O_f . As explained in Sec. II B, our theoretical analysis has ignored the control cost of implementing the oracle by convention, but the control cost is nonnegative by definition. Therefore, we have

$$E_B \geq \Delta E^{(\text{in})}. \quad (152)$$

Note that the implementation of communication between A and B may also incur a non-negative control cost, but we can include this control cost in E_A or E_B (e.g., by using the energy supply from A and B as the power source for the communication devices). As a whole, the sum of the measured energy consumption

$$E_A + E_B \quad (153)$$

serves as an upper bound of the energy consumption $\mathcal{W}_{\text{gates}} + \Delta E^{(\text{in})} - \Delta E^{(\text{out})}$.

The experimental demonstration of the advantage of quantum computation will be successful if the measured upper bound of the energy consumption $E_A + E_B$ in (153) is smaller than an estimate of the energy consumption for classical computation. Based on our results in Corollaries 7 and 9, a separation of $\mathcal{W}^{(Q)}$ and $\mathcal{W}^{(C)}$ will emerge. If the experiment achieves

$$E_A + E_B = O(\text{poly}(N)), \quad (154)$$

then the advantage can be demonstrated for large N due to the exponentially large energy consumption for the classical computation to solve the same problem. In the following, we will propose an explicit scheme according to which B can construct O_f as required in our framework yet still within a polynomial energy consumption to achieve (154), followed by also clarifying how to explicitly evaluate the lower bound of the energy consumption for the classical computation.

Construction of oracle for Simon's problem. To realize (154), it is necessary that the control cost of B 's implementation of O_f should be bounded by $O(\text{poly}(N))$ so that we should have $E_B = O(\text{poly}(N))$. Then, it is essential that B can compute f within a polynomial time (i.e., by a polynomial-size circuit) so that the oracle O_f can also be implemented in a polynomial time (and a polynomial control cost). The formulation of Simon's problem

in Problem 6, however, allows for random 1-to-1 and 2-to-1 functions f on N bits, which may not be computable within a polynomial time but take an exponential time in general. Note that in Simon's problem, the choice of f is random, but once f is chosen, the oracle needs to output the same $f(x)$ for fixed input x , which is different from producing random outputs conditioned on x . The challenge here arises from the fact that, in general, implementation of randomly chosen Boolean functions almost always requires an exponential-size circuit [90]. In constructing f in the experiment, we, therefore, have to accept an approximation to the uniformly random choice of f required for Simon's problem.

To overcome the challenge, we propose a construction of f that can be implemented within a polynomial time by B yet cannot be distinguished from the uniformly random choice of f for A . To instantiate the random functions $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$ in Simon's problem, we use a pseudorandom permutation, a well-established primitive of cryptography [91]. For some key bit string $k \in \{0, 1\}^L$ of length L , a pseudorandom permutation f_k can be computed by a block-cipher algorithm such as Advanced Encryption Standard (AES) [92, 93], using a classical circuit of a polynomial size in N and L . For a sufficiently large fixed L , the key $k \in \{0, 1\}^L$ is selected uniformly at random by B and kept as B 's secret in the same way as b and s , so that k should be kept unknown to A . A crucial property of the pseudorandom permutation is that whatever polynomial-time algorithm A performs in the experiment, A cannot distinguish f_k from a uniformly randomly drawn 1-to-1 function unless A knows k . Conventionally, this computational hardness assumption is justified for a large key length L because no known polynomial-time algorithm can distinguish f_k from a truly random choice of 1-to-1 functions. Note that multiple block-cipher algorithms such as AES are believed to be secure against attacks even by quantum computers as well as classical computers as we increase the key length L [94–96], while the detailed security analysis of the choice of L for the computational hardness assumption is beyond the scope of this paper and is left for future work.

In place of the uniformly random 1-to-1 function $f_{1\text{-to-1}}$, we propose to use the one-way permutation f_k as its replacement $\tilde{f}_{1\text{-to-1}}$ in the experiment, i.e.,

$$\tilde{f}_{1\text{-to-1}}(x) := f_k(x). \quad (155)$$

To construct a replacement $\tilde{f}_{2\text{-to-1}}$ for the uniformly random 2-to-1 function $f_{2\text{-to-1}}$ satisfying $f_{2\text{-to-1}}(x) = f_{2\text{-to-1}}(x \oplus s)$, we use the fact that for any $a < b$, the probability of having

$$f_{1\text{-to-1}}(x) = a, f_{1\text{-to-1}}(x \oplus s) = b \quad (156)$$

and that of

$$f_{1\text{-to-1}}(x \oplus s) = a, f_{1\text{-to-1}}(x) = b \quad (157)$$

are exactly the same for the uniformly random choice of $f_{1\text{-to-1}}$. Thus, using a uniformly random choice of $f_{1\text{-to-1}}$,

we can also construct a uniformly random 2-to-1 function $f_{2\text{-to-1}}$ by

$$f_{2\text{-to-1}}(x) := f_{1\text{-to-1}}(\min\{x, x \oplus s\}), \quad (158)$$

where the order between x and $x \oplus s$ is defined as that in terms of the binary integers. By replacing $f_{1\text{-to-1}}$ with $\tilde{f}_{1\text{-to-1}} = f_k$ in (155), we propose to use

$$\tilde{f}_{2\text{-to-1}}(x) := f_k(\min\{x, x \oplus s\}). \quad (159)$$

With the definitions in (155) and (159), we have an instantiation of the function f for Simon's problem that is indistinguishable from the uniformly random choice for A under the computational hardness assumption on the block cipher and is implementable within a polynomial time for B . Note that a polynomial-size quantum circuit for implementing the oracle O_f for the quantum computation can be obtained from the reversible classical logic circuit for the instantiation of f , e.g., written in terms of the classical Toffoli gate, by replacing each classical Toffoli gate with the quantum Toffoli gate (and further rewrite each quantum Toffoli gate into Clifford and T gates if we use the Clifford+ T universal gate set).

Therefore, combining this instantiation of f with the quantum algorithm analyzed in Corollary 7, we have an explicit construction of the setup to demonstrate the polynomial amount of energy consumption in quantum computation for solving Simon's problem, as described in (154).

Estimation of energy consumption of classical computation. Finally, we clarify how to estimate the energy consumption of classical computation for solving Simon's problem to be compared with the results of the quantum experiments. With the above implementation, the quantum experiment can demonstrate the energy-consumption advantage of quantum computation if the measured upper bound (153) of the energy consumption for the quantum computation is smaller than an estimate for the classical computation. However, classical algorithms for solving the same problem cannot be realized due to the exponential runtime and the exponential energy consumption, and thus, the energy consumption of the classical computation cannot be measured directly unlike that of quantum computation; after all, the main point of our analysis is that it is fundamentally infeasible for any classical algorithm to solve this problem while the quantum algorithm can. To overcome this difficulty, we here propose two approaches for the estimation: one is an extrapolation-based approach, which is more suitable for the near-term demonstration in comparison with the existing best-effort classical algorithm, and the other is a fundamental approach based on our fundamental lower bound in Corollary 9, which may be more demanding yet can demonstrate the advantage for any possible classical algorithm.

The extrapolation-based approach works for a particular choice of classical algorithm. For example, for fixed N , we consider a classical algorithm with querying M

different inputs $\vec{x} = (x_1, \dots, x_M)$ uniformly at random to the oracle for Simon's problem. The algorithm works as follows:

- if at least one pair (x_k, x_m) satisfying $f(x_k) = f(x_m)$ for some $x_k \neq x_m$ is found in the M queries, the algorithm has an evidence that f is a 2-to-1 function and thus outputs $a = 1$;
- otherwise, the algorithm considers f to be 1-to-1 and outputs $a = 0$.

For fixed N , we can conduct this classical algorithm for any $M \in \{1, \dots, 2^N\}$, but the success probability of finding (x_k, x_m) satisfying $f(x_k) = f(x_m)$ may be low if M is set to be too small. By contrast, Refs. [86, 87] have shown that there exists $M = O(2^{N/2})$ such that this algorithm solves Simon's problem from the uniformly random M queries, i.e., achieves $a = b$ with a high probability as required in Problem 6. However, these existing analyses only show the existence of $M = O(2^{N/2})$ without explicitly clarifying the constant factors therein, and thus, it was unknown how large we should set M in this classical algorithm explicitly. To establish an explicit criterion on the quantum advantage, we here provide the following proposition with the explicit constant factor.

Proposition 10 (Refined analysis of upper bound of query complexity of Simon's problem for a classical algorithm). *Fix any N . For any $\delta > 0$, if we have*

$$M = \left\lceil \frac{1}{2} \sqrt{\left(8 \ln \frac{1}{\delta}\right) 2^N + 1} + \frac{1}{2} \right\rceil = O(2^{N/2}), \quad (160)$$

the classical algorithm that makes the uniformly random M (or more) queries for solving Simon's problem in Problem 6 achieves the following with a high probability greater than $1 - \delta$:

- if there exists any $s \in \{0, 1\}^N \setminus \{0\}$ such that f is a 2-to-1 function satisfying $f(x) = f(x \oplus s)$, then this M -query algorithm outputs $a = 1$;
- if f is a 1-to-1 function, this M -query algorithm outputs $a = 0$.

Proof. The proof is presented in Appendix A 4. \square

For this classical algorithm, the extrapolation for estimating the energy consumption can be performed in the following procedure. In the setup of Fig. 4, we replace the quantum computer of A with a classical computer to perform the classical algorithm and the communication between A and B with classical communication. For a fixed choice of N , conduct the classical algorithm for several different choices of M in the order of

$$M = O(\text{poly}(N)), \quad (161)$$

and measure the energy consumption as in (153). From the measured energy consumption for several M , we estimate the leading order of the energy consumption and

its constant factor as M increases. Using this estimation of the function of M representing the energy consumption, perform an extrapolation to obtain an estimate of the energy consumption for M in (160) of Proposition 10 with setting $\delta = 1/3$ as in Problem 6. For each fixed choice of N , use this estimate as the energy consumption of this classical algorithm for solving Simon's problem, so we can compare it with the quantum case.

We note that the goal of this approach is to estimate the scaling and the constant factor of the energy consumption of the classical algorithm as M grows. In the proposed experiment with polynomial large M in (161), the classical algorithm may not return the correct output; still, one should not try to conduct the classical algorithm for exponentially large M in (160) to estimate the energy consumption directly, which is infeasible. After all, Proposition 10 already guarantees that the algorithm works for M in (160), and the aim of the experiment is not the reproduction of this theoretical result on query complexity. Rather, our proposal here is to perform only polynomial-time experiments to estimate the energy consumption for M in (160) by extrapolation, to obtain a criterion on the advantage in the energy consumption.

Next, the fundamental approach is based on the fundamental lower bound of energy consumption for any classical algorithm shown in Corollary 9. Problematically, the above extrapolation-based approach can be used for demonstration of the advantage of quantum computation only over the existing classical algorithm used in the experiment, rather than any possible classical algorithm. For example, apart from the probabilistic classical algorithm analyzed here, Ref. [97] provides a deterministic classical algorithm using $M = O(2^{N/2})$ queries to solve Simon's problem, but the bound that we have shown in Proposition 10 may not directly be applicable to this different algorithm; thus, as long as the extrapolation-based approach is used, one may need a case-by-case analysis to determine the constant factor for the extrapolation depending on the choice of classical algorithms. By contrast, in the fundamental approach here, we will address this issue by using the fundamental lower bound applicable to any classical algorithm.

In particular, for any classical algorithm to solve Simon's problem, we have the explicit lower bound in (149) in the proof of Corollary 9, i.e.,

$$\begin{aligned} \mathcal{W} &\geq \mathcal{W}^{(C)} = \frac{1}{\beta} S(C|a) \\ &\geq \frac{1}{\beta} \times \left(\frac{3\sqrt{15} - 11}{6\sqrt{15} - 18} \times \right. \\ &\quad \left. \left(\sqrt{\frac{8 - 2\sqrt{15}}{6 - \sqrt{15}}} \times 2^{N/2} N \ln 2 - 3 \right) - \ln 2 \right), \quad (162) \end{aligned}$$

where we have chosen the Boltzmann constant to be $k_B = 1$. With the room temperature in mind, we set

$$\frac{1}{\beta} = 300 \text{ K}, \quad (163)$$

N	$\mathcal{W}^{(C)} \text{ (J)}$
50	2×10^{-13}
100	1×10^{-5}
150	7×10^2
200	3×10^{10}
250	1×10^{18}
300	5×10^{25}

TABLE I. Fundamental lower bounds on the energy consumption of any classical algorithm for solving Simon's problem in Problem 6 with various problem sizes N . For small values of $N \lesssim 100$ and room temperature $1/\beta = 300 \text{ K}$, the Boltzmann constant $k_B \approx 1 \times 10^{-23} \text{ J/K}$ is dominant, keeping the lower bound of the energy consumption below 1 J. For larger N , the lower bound of energy consumption rapidly becomes too large to realize in practice for any classical algorithm. Note that apart from this fundamental approach, in the main text, we also discuss an extrapolation-based approach to estimate the lower bound of energy consumption more tightly for some specific choice of the classical algorithm, which is expected to be more suitable for near-term experiments.

and substituting $1/\beta$ in (162) with k_B/β for $k_B \approx 1 \times 10^{-23} \text{ J/K}$, we obtain the energy consumption in joule. With this calculation, typical values of the fundamental lower bounds in (162) are summarized in Table I. For comparison, we have the following facts [98].

1. Energy of using a 10-watt flashlight for 1 minute is $6 \times 10^2 \text{ J}$.
2. Magnetic stored energy in the world's largest toroidal superconducting magnet for the ATLAS experiment at CERN, Geneva, is $1 \times 10^9 \text{ J}$.
3. Yearly electricity consumption in the U.S. as of 2009 is $1 \times 10^{19} \text{ J}$.
4. Total energy from the Sun that strikes the face of the Earth each year is $5 \times 10^{25} \text{ J}$.

With an appropriate choice of N , if the quantum experiment demonstrates an energy consumption below the corresponding value in Table I, our analysis shows that no classical algorithm can outperform such an implementation of the quantum algorithm in terms of energy consumption.

Although the values in the table may be demanding at least under the current technology, the fundamental approach here has a merit instead. The conventional experiments for demonstrating quantum computational supremacy in terms of time complexity compare the quantum algorithm with the existing classical algorithm; in this case, even if a quantum algorithm in the experiments outperforms a classical algorithm, another faster classical algorithm may again outperform the implemented quantum algorithm as technology develops further. In our case of energy-consumption advantage, the extrapolation-based approach also makes a similar type of comparison between the quantum algorithm and

the existing classical algorithm. By contrast, the fundamental approach here is applicable to any possible classical algorithm, which was hard to claim for the conventional analysis of quantum advantage in time complexity. Therefore, our analysis establishes an ultimate goal of quantum technologies from a new perspective: outperform the fundamental classical lower bound of energy consumption by an experimental realization of a quantum algorithm.

Consequently, our proposal for these experimental demonstrations clarifies how to realize the query-complexity setting (formulated as a learning problem, where a party A is to learn a property of the function f through queries realized as quantum communication with another party B , and B is to implement the oracle of f). Our proposal and analysis also set quantitative goals of quantum technologies for the demonstrations (e.g., the values in Table I). On top of these theoretical developments, accomplishing a demonstration of the energy-consumption advantage of quantum computation over classical computation would be a major milestone in the experiments.

VI. DISCUSSION AND OUTLOOK

In this work, we have formulated a framework for studying the energy consumption of quantum and classical computation. With this framework, we have derived general upper and lower bounds of the energy consumption of quantum and classical computation, respectively, in the query-complexity setting.

To derive the upper bound of the energy consumption of quantum computation, our analysis provides a comprehensive account of all the contributions arising from the cost of each operation used for implementing the computation, of initializing the state of the computer at the end of the computer into the initial state, and of the overhead of quantum error correction, with imperfections of the implementation taken into account. In particular, to obtain an upper bound of the achievable initialization cost, we have developed a finite-step Landauer-erasure protocol to initialize a given state into a pure state within a finite number of time steps up to finite infidelity, progressing over the existing infinite-step protocols [16, 23]. Our analysis indicates that an upper bound of energy consumption of quantum computation is determined by the gate and query complexities and the heat dissipation achievable by the finite-step Landauer erasure in our framework. These upper bounds rigorously clarify that, even though quantum computation might have larger per-gate energy consumption than classical computation, it is indeed true that polynomial-time quantum computation can be physically realized with a polynomial amount of energy consumption including all the above costs, serving as a widely applicable (yet potentially non-tight) estimate of the achievable energy consumption of quantum computation.

As for the lower bound, we developed new techniques to prove a fundamentally nonzero lower bound that works for any classical algorithm regardless of how energy-efficiently the classical computation is implemented. To establish this fundamental lower bound, we have employed two fundamental physical insights: the energy-conservation law and Landauer's principle. In particular, working on a query-complexity setting, we develop a technique to argue that a lower bound on the query complexity of any classical algorithm solving a computational problem can be used to establish a lower bound on the Landauer cost—that is, the minimum heat dissipation required to erase the information obtained from the queries, measured in terms of entropy. Consequently, due to the energy-conservation law, the energy consumption of any classical algorithm solving such a computational problem must be at least the Landauer cost. Our analysis suggests that the lower bound on energy consumption is governed by this entropic quantity, which depends not only on the query complexity but also on the probability distribution of the oracle in our framework. Owing to the generality of these techniques, our lower bound holds regardless of technological advances in the energy efficiency of the physical implementation of classical computation.

Using these upper and lower bounds, we rigorously prove that quantum computation can be exponentially advantageous in terms of energy consumption over classical computation for solving a computational problem, in particular, Simon's problem, as the problem size increases. Since our formulation and analysis are based on the query-complexity setting of Simon's problem, our proof of this exponential quantum advantage is free from the complexity-theoretic assumptions on the hardness of the problem, such as the assumption on the conjectured hardness of integer factoring for classical computation. We have also proposed a schematic setup of experiments to demonstrate this energy-consumption advantage of quantum computation, using a cryptographic primitive to instantiate and implement the oracle for Simon's problem feasibly.

These results bridge the gap between the complexity-theoretical techniques for studying query complexity and the physical notion of energy that has been studied previously in the field of quantum thermodynamics, opening an alternative way of studying the advantages of quantum computation. Our framework and the general bounds on energy consumption can be used for studying a potentially broader class of problems; for example, similar to Simon's problem, it is interesting to extend our analysis to other problems in the query-complexity setting, such as the Abelian hidden subgroup problem [99], the Bernstein-Vazirani problem [100], and the problem with verifiable quantum advantage without structure [101]. Another relevant question would be whether such an energy-consumption advantage of quantum computation could still be proven even if the assumption of a black-box oracle is relaxed and whether the energetic contributions from the oracle can be included

in the overall bounds, unlike the conventional analysis in the query-complexity setting. Also, while our analysis provides an achievable upper bound of energy consumption for quantum computation and a fundamental lower bound for classical computation, it is interesting to develop techniques for proving a lower bound of energy consumption of quantum computation.

Finally, through our analysis, we have proposed a new type of experiment to demonstrate the energy-consumption advantages of quantum computation for Simon’s problem, offering an explicit criterion for implementing a quantum algorithm that outperforms any possible classical algorithm in energy consumption (Table I). As a result, our work elevates the significance of Simon’s problem—and potentially a broader class of query-complexity problems—from both theoretical and experimental perspectives; this aligns with the pioneering theoretical studies on quantum computational supremacy [102–108], which emphasized the importance of special sampling problems that now receive considerable attention in quantum experiments [109–111]. Remarkably, unlike the existing experiments on quantum computational supremacy that focus on time complexity, once the proposed quantum experiment satisfies our energy-consumption advantage criterion in Table I, our analysis guarantees that no classical algorithm will be able to surpass the quantum algorithm, regardless of technological advances in energy-efficient computation. To make this demonstration possible in principle, we have explicitly evaluated the lower bounds of query complexity and energy consumption for any classical algorithm to solve Simon’s problem, providing clarification on the constant factors involved. On the theoretical side, refining these constant factors in our analysis would further reduce the technological requirements for experimental demonstration. Experimentally demonstrating this fundamental energy-consumption advantage of quantum computation over any classical computation would represent a significant milestone, and our developments have opened a route toward further studies in this direction with a solid theoretical foundation.

ACKNOWLEDGMENTS

The authors thank Jake Xuereb, Parnam (Faraj) Bakhshinezhad, Philip Taranto, Satoshi Yoshida, Nat-suto Isogai, Mio Murao, Pauli Erker, and Marcus Huber for the discussion. Quantum circuits have been generated using the `quantikz` package [112]. Figure 4 has been created with emojis designed by OpenMoji under the license CC BY-SA 4.0. F.M. acknowledges funding by the European flagship on quantum technologies (‘ASPECTS’ consortium 101080167) and funding from the European Research Council (Consolidator grant ‘Coco-quest’ 101043705). H.Y. acknowledges JST PRESTO Grant Number JPMJPR201A, JPMJPR23FC, JSPS KAKENHI Grant Number JP23K19970, and MEXT

Quantum Leap Flagship Program (MEXT QLEAP) JP-MXS0118069605, JPMXS0120351339.

APPENDICES

Appendices of the article “Energy-Consumption Advantage of Quantum Computation” are organized as follows. In Appendix A, we provide technical details and proofs for the results in the main text. In Appendix B, we illustrate the physics of control cost for a specific example and highlight further references on this topic.

Appendix A: Proofs

In this appendix, we provide the proofs of the theorems, propositions, and lemmas present in the main text. In Appendix A 1, we prove the finite-step and finite-fidelity Landauer-erasure bound (Theorem 2 in the main text) together with an estimate of the relevant energy scales (Proposition 11). Then, in Appendix A 2, we provide a line-by-line decomposition and proof of the energy-consumption upper bound (Theorem 3 of the main text). In Appendix A 3, we provide technical details for the entropic lower bound in Simon’s problem (Proposition 8, Lemma 12, Proposition 13, and Lemma 14) used in the proof of Corollary 9 in the main text. Lastly, in Appendix A 4, we provide a proof of the achievable upper bound of query complexity of Simon’s problem for a classical algorithm that we use in our analysis (Proposition 10).

1. Proof on finite Landauer-erasure bound

In this appendix, we present the details of the Landauer-erasure protocol and prove a finite version of the work cost of erasing states of a finite-dimensional target quantum system for this protocol. In particular, we will prove Theorem 2 in the main text and Proposition 11 shown below.

Regarding the protocol for the Landauer erasure, we consider a protocol based on the one in Ref. [23]. Given a quantum state ρ_S of the d -dimensional target system to be erased as in Theorem 2 in the main text, the Landauer-erasure protocol transforms ρ_S into a final state ρ'_S satisfying the pure-state fidelity $1 - \varepsilon = \langle 0 | \rho'_S | 0 \rangle$; in particular, the protocol here uses

$$\rho'_S = (1 - \varepsilon) |0\rangle\langle 0| + \frac{\varepsilon}{d-1} (\mathbb{1} - |0\rangle\langle 0|), \quad (\text{A1})$$

where $\mathbb{1}$ is the identity operator acting on the d -dimensional space. For the sake of analysis, we require $\varepsilon \leq 1/2$ (to be used in (A46), (A62), and (A73)). Importantly, this final state has full rank and therefore satisfies $\text{supp}(\rho_S) \subseteq \text{supp}(\rho'_S)$ as required also in Ref. [23], where $\text{supp}(\rho)$ represents the support of the operator ρ . To define the protocol, one can consider any continuously differentiable path $\rho[u]$ through the space of density operators parameterized by $u \in [0, 1]$ with endpoints $\rho(0) = \rho_S$ and $\rho(1) = \rho'_S$. For any choice of T points

$0 < u_1 < u_2 < \dots < u_T = 1$ in $[0, 1]$, the environment is composed of T d -dimensional subsystems and has the total Hamiltonian

$$H_E := \sum_{t=1}^T H_E^{(t)} \otimes \mathbb{1}_{\text{rest} \neq t} \geq 0, \quad (\text{A2})$$

where the Hamiltonian of the t th subsystem ($t = 1, \dots, T$) is given by

$$H_E^{(t)} := -\frac{1}{\beta} \ln \rho[u_t] \geq 0, \quad (\text{A3})$$

and $\mathbb{1}_{\text{rest} \neq t}$ is the identity operator acting on all the subsystems but the t th. For each t , we write the thermal state of the t th subsystem of the environment as

$$\tau_E^{(t)}[\beta] := \frac{e^{-\beta H_E^{(t)}}}{\text{tr}[e^{-\beta H_E^{(t)}}]} = \rho[u_t], \quad (\text{A4})$$

where the last equality holds by definition (A3). Then, the thermal state of the environment $\tau_E[\beta]$ in (53) of the main text becomes

$$\tau_E[\beta] = \bigotimes_{t=1}^T \tau_E^{(t)}[\beta] = \bigotimes_{t=1}^T \rho[u_t]. \quad (\text{A5})$$

The protocol is composed of T steps; in the t th step for $t = 1, \dots, T$, the protocol swaps the target system’s state with the thermal state $\tau_E^{(t)}[\beta]$ of the t th subsystem of the environment, so that the target system’s state after the t th step should be $\rho[u_t]$ along the path $\rho[u]$. For this protocol, Ref. [23] shows that the heat dissipation into the environment is given by

$$\beta \mathcal{Q}_E = \sum_{t=1}^T \text{tr} \left[(\rho[u_t] - \rho[u_{t-1}]) \ln (\rho[u_t]) \right]. \quad (\text{A6})$$

For the sake of analysis, our protocol chooses a straight path $\rho[u]$ through the state space and equally spaced points on the path, i.e.,

$$\rho[u] = \rho_S + u(\rho'_S - \rho_S), \quad (\text{A7})$$

$$u_t = \frac{t}{T}. \quad (\text{A8})$$

With this protocol, we prove Theorem 2 in the main text, which is repeated again below. In particular, we progress beyond the asymptotic bound shown in Ref. [16] and analyze in detail how the number of steps an erasure protocol requires scales as a function of the final pure-state infidelity ε and the difference of the actual heat dissipation η to Landauer’s ideal bound.

Theorem 2 (Finite Landauer-erasure bound). *Given a quantum state ρ_S of a d -dimensional target system for*

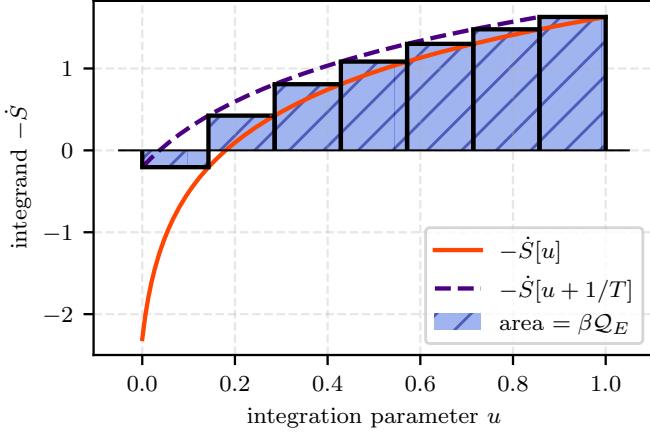


FIG. 5. Illustration visualizing the arguments used for deriving the upper and lower bound on βQ_E in (A23) and (A27) for $T = 7$. The sum on the right-hand side of (A6) making up βQ_E is represented through the area below the step function in blue. This step function is lower bounded by the integral $\int_0^1 du (-\dot{S}[u])$ due to the monotonicity of the integrand (solid red curve), which yields to (A23). An upper bound is given by the shifted integral given by the dashed violet curve, together with a correction of the last step on the right of the interval of the integral, leading to (A27).

any $d \geq 2$ and any constants $\varepsilon \in (0, 1/2]$ and $\eta \in (0, 1]$, let T be

$$T = \left\lceil \frac{(e+1)}{\varepsilon\eta} \ln \left(\frac{(e+1)(d-1)^2}{\varepsilon\eta} \right) \right\rceil = O\left(\frac{1}{\eta} \log \frac{d^2}{\varepsilon\eta} \right), \quad (66)$$

where $\lceil \dots \rceil$ is the ceiling function, and e is the Euler constant. Then, a T -step Landauer-erasure protocol can transform ρ_S into a final state ρ'_S with infidelity to a pure state $|0\rangle$ satisfying $\langle 0 | \rho'_S | 0 \rangle \geq 1 - \varepsilon$ in (55) by using an environment at inverse temperature $\beta > 0$, in such a way that the heat Q_E in (56) dissipated into the environment differs from Landauer's bound ΔS in (59) at most by

$$\beta Q_E - \Delta S \leq \eta, \quad (67)$$

as required in (65).

Proof. We will prove that, for a generic number of steps T , the protocol described above to transform ρ_S into ρ'_S in (A1), which is ε -close to the pure state $|0\rangle$, has a heat dissipation Q_E to the environment bounded by

$$\beta Q_E - \Delta S \leq \frac{\ln \left(\frac{e(d-1)^2 T}{\varepsilon} \right)}{T}, \quad (A9)$$

so that setting T as in the theorem statement should achieve $\beta Q_E - \Delta S \leq \eta$.

To bound Q_E in (A6), we write, for simplicity of notation,

$$S[u] := S(\rho[u]) = -\text{tr}[\rho[u] \ln(\rho[u])]. \quad (A10)$$

By using the shorthand $\dot{S}[u] := \frac{\partial}{\partial u} S[u]$, it holds that

$$\dot{S}[u] = -\text{tr}[\dot{\rho}[u] \times (\ln(\rho[u]) + \mathbb{1})] \quad (A11)$$

$$= -\text{tr}[(\rho'_S - \rho_S) \times (\ln(\rho[u]) + \mathbb{1})] \quad (A12)$$

$$= -\text{tr}[(\rho'_S - \rho_S) \ln(\rho[u])]. \quad (A13)$$

Using this notation, each term on the right-hand side of (A6) simplifies into

$$\text{tr}[(\rho[u_t] - \rho[u_{t-1}]) \ln(\rho[u_t])] = \text{tr}\left[\frac{\rho'_S - \rho_S}{T} \ln(\rho[u_t])\right] \quad (A14)$$

$$= -\frac{\dot{S}[u_t]}{T}. \quad (A15)$$

Moreover, $S[u]$ is a concave function; i.e., for any u_1, u_2 , and $\lambda \in [0, 1]$, we have

$$S[\lambda u_1 + (1 - \lambda)u_2] = S(\rho[\lambda u_1 + (1 - \lambda)u_2]) \quad (A16)$$

$$\stackrel{(A7)}{=} S(\lambda \rho[u_1] + (1 - \lambda) \rho[u_2]) \quad (A17)$$

$$\geq \lambda S(\rho[u_1]) + (1 - \lambda) S(\rho[u_2]) \quad (A18)$$

$$= \lambda S[u_1] + (1 - \lambda) S[u_2]. \quad (A19)$$

Thus, $-\dot{S}[u]$ is a monotonically non-decreasing function as u increases. In Fig. 5, we illustrate an example of a monotonically non-decreasing function and how we use this property to derive upper and lower bounds of βQ_E . By evaluating the integral according to this figure, (A6) can be bounded from below by

$$\beta Q_E = \sum_{t=1}^T \frac{-\dot{S}[u_t]}{T} \quad (A20)$$

$$\geq \int_0^1 du (-\dot{S}[u]) \quad (A21)$$

$$= S[0] - S[1] \quad (A22)$$

$$= \Delta S \quad (A23)$$

and from above by

$$\beta Q_E = \sum_{t=1}^T \frac{-\dot{S}[u_t]}{T} \quad (A24)$$

$$\leq \int_{1/T}^1 du (-\dot{S}[u]) + \frac{-\dot{S}[1]}{T} \quad (A25)$$

$$= \int_0^1 du (-\dot{S}[u]) - \int_0^{1/T} du (-\dot{S}[u]) + \frac{-\dot{S}[1]}{T} \quad (A26)$$

$$\leq \Delta S + \left| S[0] - S\left[\frac{1}{T}\right] \right| + \frac{-\dot{S}[1]}{T}. \quad (A27)$$

Therefore, we can rewrite (A6) as

$$0 \leq \beta Q_E - \Delta S \leq \left| S[0] - S\left[\frac{1}{T}\right] \right| - \frac{1}{T} \dot{S}[1]. \quad (A28)$$

We bound the right-hand side of (A28). The first term on the right-hand side can be estimated using the continuity bound of the quantum entropy in Refs. [113–116]. In particular, let

$$\delta := \frac{1}{2} \left\| \rho[0] - \rho \left[\frac{1}{T} \right] \right\|_1 \quad (\text{A29})$$

denote the trace distance, and

$$h(x) := \begin{cases} -x \ln(x) - (1-x) \ln(1-x), & x \in (0, 1), \\ 0 & x = 0, 1 \end{cases} \quad (\text{A30})$$

the binary entropy. Since we have

$$\rho[0] - \rho \left[\frac{1}{T} \right] = \frac{\rho'_S - \rho_S}{T}, \quad (\text{A31})$$

it holds that

$$\delta = \frac{1}{2T} \|\rho'_S - \rho_S\|_1 \leq \frac{1}{T}, \quad (\text{A32})$$

where the last inequality follows from $\frac{1}{2} \|\rho' - \rho\|_1 \in [0, 1]$ for all density operators ρ and ρ' . Thus, for any $T \geq 2$, it holds that

$$\delta \leq \frac{1}{2} \leq 1 - \frac{1}{d}, \quad (\text{A33})$$

where the last inequality holds due to $d \geq 2$. For δ satisfying (A33), the continuity bound in Refs. [113–116] leads to

$$\left| S[0] - S \left[\frac{1}{T} \right] \right| = \left| S(\rho[0]) - S \left(\rho \left[\frac{1}{T} \right] \right) \right| \quad (\text{A34})$$

$$\leq \delta \ln(d-1) + h(\delta). \quad (\text{A35})$$

For any $T \geq 2$, it also follows from (A32) that

$$h(\delta) \leq h \left(\frac{1}{T} \right) \quad (\text{A36})$$

$$\leq \frac{1 + \ln T}{T}, \quad (\text{A37})$$

where the last inequality follows from

$$h(x) \leq x \left(1 + \ln \frac{1}{x} \right) \quad (\text{A38})$$

for all $x > 0$. Therefore, for any $T \geq 2$, we obtain from (A32), (A35), and (A37)

$$\left| S[0] - S \left[\frac{1}{T} \right] \right| \leq \frac{\ln(d-1)}{T} + \frac{1 + \ln T}{T}. \quad (\text{A39})$$

To bound the second term of the right-hand side of (A28), we use ρ'_S in (A1) to evaluate

$$-\dot{S}[1] = \text{tr}[(\rho'_S - \rho_S) \ln \rho'_S] \quad (\text{A40})$$

$$= -S(\rho'_S) - \text{tr}[\rho_S \ln \rho'_S] \quad (\text{A41})$$

$$= -S(\rho'_S) - \langle 0 | \rho_S | 0 \rangle \ln(1 - \varepsilon)$$

$$- \sum_{j=1}^{d-1} \langle j | \rho_S | j \rangle \ln \left(\frac{\varepsilon}{d-1} \right) \quad (\text{A42})$$

$$\leq 0 - \langle 0 | \rho_S | 0 \rangle \ln(1 - \varepsilon) - \sum_{j=1}^{d-1} \langle j | \rho_S | j \rangle \ln \left(\frac{\varepsilon}{d-1} \right) \quad (\text{A43})$$

$$\leq - \sum_{j=0}^{d-1} \langle j | \rho_S | j \rangle \ln \left(\frac{\varepsilon}{d-1} \right) \quad (\text{A44})$$

$$= \ln \left(\frac{d-1}{\varepsilon} \right), \quad (\text{A45})$$

where (A44) follows from

$$-\ln(1 - \varepsilon) \leq -\ln \left(\frac{\varepsilon}{d-1} \right) \quad (\text{A46})$$

for all $\varepsilon \leq 1/2$ and $d \geq 2$, and (A45) uses normalization of $\text{tr}[\rho_S] = 1$ and the identity $-\log(x) = \log(1/x)$ for $x > 0$. As a whole, applying (A39) and (A45) to (A28) yields

$$\beta Q_E - \Delta S \leq \frac{\ln(d-1)}{T} + \frac{1 + \ln T}{T} + \frac{\ln \left(\frac{d-1}{\varepsilon} \right)}{T} \quad (\text{A47})$$

$$= \frac{\ln \left(\frac{e(d-1)^2 T}{\varepsilon} \right)}{T}, \quad (\text{A48})$$

which provides the bound (A9) on the heat dissipation Q_E for a generic number of steps $T \geq 2$.

To complete the proof, we set T in (A48) in such a way that we should obtain $\beta Q_E - \Delta S \leq \eta$. We write (A48) as

$$f(T) := \frac{\ln \left(\frac{e(d-1)^2 T}{\varepsilon} \right)}{T}. \quad (\text{A49})$$

Given the fact that $\ln x/x$ does not increase for all $x \geq e$, we see that $f(T)$ is a monotonically non-increasing function for all $T \geq \varepsilon/(d-1)^2$, which always holds here due to $T \geq 1$, $\varepsilon \leq 1/2$, and $d \geq 2$. An exact solution of

$$f(T) = \eta \quad (\text{A50})$$

with respect to T would require the Lambert W function, which cannot be expressed in a closed form through elementary functions [117, 118]. But for our analysis, finding T satisfying

$$f(T) \leq \eta \quad (\text{A51})$$

within the order of

$$T = O \left(\frac{1}{\eta} \log \frac{d^2}{\varepsilon \eta} \right) \quad (\text{A52})$$

is sufficient; to provide such T , let us choose T as

$$T = \left\lceil \frac{(e+1)}{e\eta} \ln \left(\frac{(e+1)(d-1)^2}{\varepsilon \eta} \right) \right\rceil. \quad (\text{A53})$$

Then, we have

$$f(T) \quad (\text{A54})$$

$$\leq f\left(\frac{e+1}{e\eta} \ln\left(\frac{(e+1)(d-1)^2}{\varepsilon\eta}\right)\right) \quad (\text{A55})$$

$$= \frac{e\eta}{e+1} \left(\frac{\ln\left(\frac{(e+1)(d-1)^2}{\varepsilon\eta}\right)}{\ln\left(\frac{(e+1)(d-1)^2}{\varepsilon\eta}\right)} + \frac{\ln\left(\ln\left(\frac{(e+1)(d-1)^2}{\varepsilon\eta}\right)\right)}{\ln\left(\frac{(e+1)(d-1)^2}{\varepsilon\eta}\right)} \right) \quad (\text{A56})$$

$$\leq \frac{e\eta}{e+1} \left(1 + \frac{1}{e}\right) \quad (\text{A57})$$

$$= \eta, \quad (\text{A58})$$

where the first inequality follows from the monotonicity of f , the second inequality follows from

$$\frac{\ln x}{x} \leq \frac{1}{e} \quad (\text{A59})$$

for all $x > 0$. For all $\varepsilon \leq 1/2$, $\eta \leq 1$, and $d \geq 2$, the choice of T in (A53) indeed satisfies

$$T \geq \frac{e+1}{e \times 1} \ln\left(\frac{(e+1)(2-1)^2}{\frac{1}{2} \times 1}\right) \quad (\text{A60})$$

$$= \frac{e+1}{e} \ln(2(e+1)) \quad (\text{A61})$$

$$\geq 2. \quad (\text{A62})$$

Note that $\eta \leq 1$ is not necessarily the tightest upper bound to derive the last line and thus to show the theorem here. Therefore, we obtain from (A48) and (A58)

$$\beta \mathcal{Q}_E - \Delta S \leq \eta, \quad (\text{A63})$$

which yields the conclusion. \square

In Proposition 11, we analyze the energetics involved in the finite Landauer-erasure protocol in Theorem 2. In particular, we are interested how the energy scale of the Hamiltonian H_E of the environment in (A2) behaves as a function of the required infidelity ε between the final state and the pure state and the required steps T of the protocol in Theorem 2. Note that we have $T \geq 2$ in Theorem 2 due to (A62).

Proposition 11 (Maximum energy of environment in finite Laudauer-eraser protocol). *Let H_E be the Hamiltonian in (A2). The maximum energy eigenvalue of H_E is bounded by*

$$\|H_E\|_\infty \leq \frac{T}{\beta} \left(\ln \frac{d-1}{\varepsilon} + 1 \right) - \frac{1}{\beta} \left(\frac{\ln 2\pi T}{2} + \frac{1}{12T+1} \right) \quad (\text{A64})$$

$$= O\left(\frac{T}{\beta} \ln \frac{d}{\varepsilon}\right), \quad (\text{A65})$$

where $T \geq 2$ is the number of steps of the Landauer-erasure protocol in Theorem 2, $d \geq 2$ is the dimension of the target system for the Landauer-erasure protocol, and $\varepsilon \in (0, 1/2]$ is the infidelity between the final state and the pure state in (A1).

Proof. Since H_E is a sum over terms $H_E^{(t)} \otimes \mathbb{1}_{\text{rest} \neq t}$ as shown in (A2), the maximum eigenvalue of H_E is bounded by the sum of the maxima over the Hamiltonian $H_E^{(t)}$ of each individual subsystem, i.e.,

$$\|H_E\|_\infty \leq \sum_{t=1}^T \|H_E^{(t)} \otimes \mathbb{1}_{\text{rest} \neq t}\|_\infty = \sum_{t=1}^T \|H_E^{(t)}\|_\infty. \quad (\text{A66})$$

Due to the definition of $H_E^{(t)}$ in (A3), we have

$$\|H_E^{(t)}\|_\infty = \left\| -\frac{1}{\beta} \ln \rho[u_t] \right\|_\infty = -\frac{1}{\beta} \ln \lambda_{\min}(\rho[u_t]), \quad (\text{A67})$$

where we write the minimum eigenvalue of a density operator ρ as $\lambda_{\min}(\rho)$. Thus, we have

$$\|H_E\|_\infty \leq -\frac{1}{\beta} \sum_{t=1}^T \ln \lambda_{\min}(\rho[u_t]). \quad (\text{A68})$$

The problem reduces to bounding $\lambda_{\min}(\rho[u_t])$. As shown in (A7) and (A8), we have

$$\rho[u_t] = \left(1 - \frac{t}{T}\right) \rho_S + \frac{t}{T} \rho'_S. \quad (\text{A69})$$

It generally holds that the minimum eigenvalue of the sum of two positive semi-definite Hermitian matrices is bounded from below by the sum of the smallest two eigenvalues of the summands [119–121]. Since $\rho \geq 0$ is not further specified, we here bound the minimum eigenvalue of $\rho[u_t]$ from below by that of ρ'_S , i.e.,

$$\lambda_{\min}(\rho[u_t]) \geq \left(1 - \frac{t}{T}\right) \lambda_{\min}(\rho_S) + \frac{t}{T} \lambda_{\min}(\rho'_S) \quad (\text{A70})$$

$$\geq 0 + \frac{t}{T} \lambda_{\min}(\rho'_S) \quad (\text{A71})$$

$$= \frac{t}{T} \frac{\varepsilon}{d-1}, \quad (\text{A72})$$

where the last line follows from the fact that the eigenvalues of ρ'_S in (A1) is $(1 - \varepsilon)$ and $\varepsilon/(d-1)$, and we have

$$1 - \varepsilon \geq \frac{\varepsilon}{d-1} \quad (\text{A73})$$

for all $\varepsilon \in (0, 1/2]$ and $d \geq 2$. As a result, summing over all the contributions, we arrive at the following upper bound for the maximum eigenvalue of H_E

$$\|H_E\|_\infty \leq -\frac{1}{\beta} \sum_{t=1}^T \ln \frac{t\varepsilon}{T(d-1)} \quad (\text{A74})$$

$$= \frac{T}{\beta} \ln \frac{T(d-1)}{\varepsilon} - \frac{1}{\beta} \ln T! \quad (\text{A75})$$

$$\leq \frac{T}{\beta} \ln \frac{T(d-1)}{\varepsilon}$$

$$-\frac{1}{\beta} \left(T \ln T - T + \frac{\ln 2\pi T}{2} + \frac{1}{12T+1} \right) \quad (A76)$$

$$= \frac{T}{\beta} \left(\ln \frac{d-1}{\varepsilon} + 1 \right) - \frac{1}{\beta} \left(\frac{\ln 2\pi T}{2} + \frac{1}{12T+1} \right) \quad (A77)$$

$$= O \left(\frac{T}{\beta} \ln \frac{d}{\varepsilon} \right), \quad (A78)$$

ling's approximation [89]

$$T! \geq \sqrt{2\pi T} \left(\frac{T}{e} \right)^T e^{\frac{1}{12T+1}}. \quad (A79)$$

This concludes the proof. \square

2. Detail of proof of general upper bound on energy consumption

In this appendix, we expand upon the expressions for the energy consumption of quantum algorithms derived in Sec. IIIA. In particular, we prove Theorem 3 repeated again below.

where the second inequality follows from a variant of Stir-

Theorem 3 (An upper bound of energy consumption in ideal case without gate errors). *Given constants $E_{\text{ctrl.}}$ in (38) and β in (13), working in an idealized case without gate errors, computation in the framework of Fig. 3 can be performed with an energy consumption \mathcal{W} (Definition 1) bounded, as the problem size N goes to infinity, and the infidelities $\varepsilon, \eta > 0$ vanish, by*

$$\mathcal{W} \leq \mathcal{W}^{(Q)} = O \left(E_{\text{ctrl.}} \times WD_{\text{comp}} + \left(E_{\text{ctrl.}} + \frac{1}{\beta\eta} \right) \times \frac{W}{\eta} \text{polylog} \left(\frac{1}{\varepsilon\eta} \right) \right), \quad (83)$$

where the width W of the circuit in (9) and the depth D_{comp} given by (17) are the parameters depending on the problem size N .

Proof. Proving the statement essentially boils down to summing up all the terms in (30). To this end, we use the expressions derived in steps (1) and (2) for the gate and erasure costs and replace them line-by-line. We also note that the contributions from energetic costs sum up to zero, and hence, in the overall energy consumption, the terms that matter are the control cost and the initialization cost.

As a prerequisite, we give some more detailed expressions for the cost of the erasure in (77). We start by recalling that the number of steps T in the Landauer-erasure protocol for each qubit (i.e., for dimension $d = 2$) is given, due to Theorem 2, by

$$T = O \left(\frac{1}{\eta} \log \frac{1}{\varepsilon\eta} \right), \quad (A80)$$

as $\eta, \varepsilon \rightarrow 0$. This expression will be used for expressing all the erasure-related costs in terms of ε and η . First of all, the circuit depth (71) of the erasure part \mathcal{E} in Fig. 3 can be expressed as

$$D_{\mathcal{E}} \leq O(T) = O \left(\frac{1}{\eta} \log \frac{1}{\varepsilon\eta} \right), \quad (A81)$$

where the swap gate is implemented within a constant depth in our setting. In addition, the environment control cost in (75) can be expressed as

$$E_{\text{ctrl.}}^{(E)} \stackrel{(75)}{=} O \left(\frac{T}{\beta} \log \frac{1}{\varepsilon} \right) \stackrel{(A80)}{=} O \left(\frac{1}{\beta\eta} \text{polylog} \frac{1}{\varepsilon\eta} \right). \quad (A82)$$

All these expressions together can then be inserted into the control cost for the erasure, $E_{\text{ctrl.}}^{(\mathcal{E})}$ to give

$$\mathcal{W}_{\text{gate}}^{(\mathcal{E})} - \Delta E^{(\mathcal{E})} = E_{\text{ctrl.}}^{(\mathcal{E})} + \mathcal{Q}_E \quad (A83)$$

$$\stackrel{(72), (76)}{\leq} \left(E_{\text{ctrl.}} + E_{\text{ctrl.}}^{(E)} \right) \times WD_{\mathcal{E}} + \frac{W}{\beta} (\ln 2 + \eta) \quad (A84)$$

$$\stackrel{(A81), (A82)}{\leq} \left(E_{\text{ctrl.}} + O \left(\frac{1}{\beta\eta} \text{polylog} \frac{1}{\varepsilon\eta} \right) \right) \times W \cdot O \left(\frac{1}{\eta} \log \frac{1}{\varepsilon\eta} \right) + \frac{W}{\beta} (\ln 2 + \eta) \quad (A85)$$

$$= O \left(\left(E_{\text{ctrl.}} + \frac{1}{\beta\eta} \right) \times \frac{W}{\eta} \text{polylog} \frac{1}{\varepsilon\eta} \right), \quad (A86)$$

which leads to the bound in (82). Finally, using (30), we obtain

$$\mathcal{W} = \left(\sum_{k=1}^{M+1} E_{\text{ctrl.}}^{(U_k)} + \sum_{k=1}^M E_{\text{ctrl.}}^{(\text{in},k)} + E_{\text{ctrl.}}^{(\text{out})} + E_{\text{ctrl.}}^{(\mathcal{E})} \right) + \mathcal{Q}_E \quad (\text{A87})$$

$$\stackrel{(40),(44),(47)}{=} \sum_{k=1}^{M+1} O(E_{\text{ctrl.}} \times WD_k) + \sum_{k=1}^M O(E_{\text{ctrl.}} \times WD_{\text{in}}) + O(E_{\text{ctrl.}} \times WD_{\text{out}}) + E_{\text{ctrl.}}^{(\mathcal{E})} + \mathcal{Q}^{(\mathcal{E})} \quad (\text{A88})$$

$$= \underbrace{O\left(E_{\text{ctrl.}} \times W \left(\sum_{k=1}^{M+1} D_k + MD_{\text{in}} + D_{\text{out}} \right) \right)}_{\text{all costs for all } U_k, \text{ input, and output}} + E_{\text{ctrl.}}^{(\mathcal{E})} + \mathcal{Q}^{(\mathcal{E})} \quad (\text{A89})$$

$$\stackrel{(\text{A86})}{=} O\left(E_{\text{ctrl.}} \times WD_{\text{comp}} + \left(E_{\text{ctrl.}} + \frac{1}{\beta\eta}\right) \times \frac{W}{\eta} \text{polylog} \frac{1}{\varepsilon\eta}\right), \quad (\text{A90})$$

where the last line follows from the definition of D_{comp} in (17). This concludes the proof. \square

3. Proof on the classical lower bound on energy consumption for solving Simon's problem

In this appendix, we prove Proposition 8, Lemma 12, Proposition 13, and Lemma 14 used for establishing a lower bound on the energy consumption of all classical algorithms for solving Simon's problem analyzed in Sec. IV C.

We start with providing the proof of Proposition 8 repeated again below for readability. The proposition shows that a lower bound of the query complexity of Simon's problem in Problem 6 is given for problem size N and average success probability $1/2 + \Delta$ by (115), i.e.,

$$M_{N,\Delta} := \left\lceil \sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2} \right\rceil, \quad (\text{A91})$$

where $\Delta \in (0, 1/2)$ is any constant in this parameter region and is suitably tunable for our analysis. Note that Ref. [86] have also shown a bound on M to achieve the average success probability $3/4$ in Simon's problem, but for our analysis of Problem 6, especially for the proof of Lemma 12, it is critical to obtain the bound on M for an average success probability between $1/2$ and $2/3$, as with Proposition 8 with parameter chosen as $\Delta < 1/6$.

Proposition 8 (Refined analysis of lower bound of query complexity of Simon's problem for classical algorithms). *For any parameter $\Delta \in (0, 1/2)$, if we have*

$$M < M_{N,\Delta} := \left\lceil \sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2} \right\rceil, \quad (115)$$

the average success probability of any probabilistic classical algorithm with at most M queries for Simon's problem in Problem 6 is upper bounded by

$$\frac{1}{2} + \Delta. \quad (116)$$

Proof. As in Refs. [12, 13], given some fixed input sequence $\vec{x} = (x_1, \dots, x_M)$ of length M in (114), we call \vec{x} *good* if $x_k = x_m \oplus s$ for some $1 \leq k < m \leq M$; otherwise, we call \vec{x} *bad*. For a fixed input sequence \vec{x} , let

$$p_{\text{good}} \quad (\text{A92})$$

denote the probability of \vec{x} being good; i.e., that being bad is $1 - p_{\text{good}}$. To succeed in estimating b with probability greater than $2/3$ as required in Problem 6, any classical algorithm in our framework of Fig. 3 needs to query the function oracle O_f repeatedly at least until the input sequence \vec{x} becomes good with some probability sufficiently larger than zero. When \vec{x} is good and $b = 1$, i.e., f is 2-to-1, the classical algorithm can find elements of \vec{x} satisfying $x_k = x_m \oplus s$ for some $1 \leq k < m \leq M$ to conclude that f is 2-to-1. But we note that even if \vec{x} is good, in the case of $b = 0$, i.e., the case where f is 1-to-1, the M different inputs $x \in \{0, 1\}^N$ still yield M distinct outputs $y = f(x)$. Conditioned on having a good input sequence, the success probability of the classical algorithm estimating b correctly has a trivial upper bound 1. For a bad input sequence \vec{x} , the M input-output relations $(x_1, f(x_1)), \dots, (x_M, f(x_M))$ are identical regardless of $b = 0$ or $b = 1$ (i.e., regardless of f being 1-to-1 or 2-to-1). Thus, conditioned on having a bad input sequence, the success probability of the classical algorithm estimating b correctly cannot exceed $1/2$, i.e., a completely random guess. Therefore, the overall success probability of the classical algorithm estimating b correctly is upper bounded by

$$p_{\text{good}} \times 1 + (1 - p_{\text{good}}) \times \frac{1}{2} = \frac{1}{2} + \frac{p_{\text{good}}}{2}. \quad (\text{A93})$$

As shown in Ref. [86] (see also Refs. [12, 13]), the probability of having a good input sequence $\vec{x} = (x_1, \dots, x_M)$ at the M th query, conditioned on the case where the input sequence at the $(M-1)$ th query is bad, is at most

$$\frac{M-1}{2^N - 1 - \frac{(M-2)(M-1)}{2}}. \quad (\text{A94})$$

Following Ref. [86], for any M , we bound

$$\frac{M-1}{2^N-1-\frac{(M-2)(M-1)}{2}} \leq \frac{2M}{2^{N+1}-M^2}. \quad (\text{A95})$$

Thus, for any sequence \vec{x} obtained from M queries, the probability of \vec{x} being good is at most

$$p_{\text{good}} \leq \sum_{m=1}^M \frac{m-1}{2^N-1-\frac{(m-2)(m-1)}{2}} \quad (\text{A96})$$

$$\leq \sum_{m=1}^M \frac{2M}{2^{N+1}-M^2} \quad (\text{A97})$$

$$= \frac{2M^2}{2^{N+1}-M^2}. \quad (\text{A98})$$

Since M is an integer, the bound (115), i.e.,

$$M < \left\lceil \sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2} \right\rceil \quad (\text{A99})$$

holds if and only if we have

$$M < \sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2}. \quad (\text{A100})$$

Therefore, for any M satisfying (115), we have

$$p_{\text{good}} < \frac{2\left(\sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2}\right)^2}{2^{N+1}-\left(\sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2}\right)^2} = 2\Delta. \quad (\text{A101})$$

Due to (A93), the average success probability of the classical algorithm is at most

$$\frac{1}{2} + \frac{p_{\text{good}}}{2} < \frac{1}{2} + \Delta. \quad (\text{A102})$$

□

Lemma 12 below aims to provide a high-probability lower bound on the required number of queries to the oracle for all classical algorithms to solve Simon's problem. We have shown in Proposition 8 that, for any classical algorithm with the number of queries smaller than $M_{N,\Delta}$, its success probability in solving Simon's problem

in Problem 6 is upper bounded by $1/2 + \Delta$. For our purpose, however, we need to invert this statement; in words, what we need for our analysis is that to achieve a fixed success probability $2/3$, the probability that the oracle is queried at least $M \geq M_{N,\Delta}$ times should be reasonably high. Intuitively, if the algorithm stops with $M < M_{N,\Delta}$ too often, the algorithm cannot gain enough information from the queries to successfully solve Simon's problem with probability $2/3$. In particular, let us write the probability of the algorithm performing strictly less than $M_{N,\Delta}$ queries as

$$p(M < M_{N,\Delta}) = \sum_{M < M_{N,\Delta}} p(M), \quad (\text{A103})$$

and that of more than or equal to $M_{N,\Delta}$ queries

$$p(M \geq M_{N,\Delta}) = \sum_{M \geq M_{N,\Delta}} p(M), \quad (\text{A104})$$

where $p(M)$ is defined in (128) as the probability distribution of the algorithm querying the oracle exactly M times. Then, we need to show that a fair fraction of the volume of the probability distribution $p(M)$ lies in the region of $M \geq M_{N,\Delta}$; that is, $p(M \geq M_{N,\Delta}) = \Omega(1)$ for large N . We prove such a probabilistic bound quantitatively in the following lemma.

Lemma 12 (High-probability lower bound on the required number of queries for all classical algorithms to solve Simon's problem). *Let $N > 0$ be a problem size and the oracle O_f be chosen according to Simon's problem as described in Problem 6. For any probabilistic and adaptive classical algorithm for solving Simon's problem with probability greater than or equal to $2/3$, the probability of the algorithm querying the oracle O_f greater than or equal to $\sqrt{2\Delta/(1+\Delta)} \times 2^{N/2}$ times is bounded by*

$$p(M \geq M_{N,\Delta}) \geq \frac{1-6\Delta}{3-6\Delta}, \quad (\text{A105})$$

where $M_{N,\Delta}$ is given by (115), $\Delta \in (0, 1/2)$ is the parameter in Proposition 8, and $p(M \geq M_{N,\Delta})$ is defined as (A104).

Proof. We here more rigorously express the ideas sketched above. To this end, let us write the algorithm's success probability as

$$p(\text{success}) = p(\text{success}|M < M_{N,\Delta}) p(M < M_{N,\Delta}) + p(\text{success}|M \geq M_{N,\Delta}) p(M \geq M_{N,\Delta}) \geq \frac{2}{3}, \quad (\text{A106})$$

where the $p(\text{success}|M < M_{N,\Delta})$ is the success probability of the algorithm conditioned on the cases where the algorithm has queried the oracle strictly less than $M_{N,\Delta}$ times, and $p(\text{success}|M \geq M_{N,\Delta})$ is that conditioned on the cases where the oracle has queried more than or equal to $M_{N,\Delta}$ times. To (A106), we apply the result in Proposition 8 showing that the success probability conditioned on the case where the oracle has queried strictly less than $M_{N,\Delta}$

times is at most $1/2 + \Delta$; i.e., it holds that

$$p(\text{success}|M < M_{N,\Delta}) \leq \frac{1}{2} + \Delta. \quad (\text{A107})$$

Moreover, due to the normalization of probability distributions, we have

$$p(M < M_{N,\Delta}) = 1 - p(M \geq M_{N,\Delta}), \quad (\text{A108})$$

$$p(\text{success}|M \geq M_{N,\Delta}) \leq 1. \quad (\text{A109})$$

Combining (A106), (A107) (A108), and (A109), we find

$$\left(\frac{1}{2} + \Delta\right)(1 - p(M \geq M_{N,\Delta})) + p(M \geq M_{N,\Delta}) \geq \frac{2}{3}. \quad (\text{A110})$$

Therefore, we have

$$p(M \geq M_{N,\Delta}) \geq \frac{1 - 6\Delta}{3 - 6\Delta}, \quad (\text{A111})$$

and thereby conclude the proof. \square

The next ingredient for proving the entropic lower bound on the energy consumption for classical algorithms to solve Simon's problem in Sec. IV C is Proposition 13 below. For a given number of queries to the oracle, this proposition gives a lower bound on how much entropy the algorithm produces.

Proposition 13 (Lower bound on entropy produced by all classical algorithms with a fixed number of queries for Simon's problem). *For fixed $M \geq 0$, any classical probabilistic and adaptive algorithm structured as in Fig. 3 calling the oracle O_f for Simon's problem in Problem 6 on M different inputs $x \in \{0, 1\}^N$ generates the computer's state with entropy*

$$S(\rho_M) \geq \frac{1}{2} \ln \frac{2^N!}{(2^N - M)!}, \quad (\text{A112})$$

where ρ_M is given by (129) in the main text, and S is the entropy defined according to $S(\rho) = -\text{tr}[\rho \ln \rho]$.

Proof. To start the proof, we clarify an explicit form of ρ_M in (129). Problem 6 is defined using s , b , and f sampled according to the uniform distributions $p(s)$, $p(b)$, and $p(f|s, b)$, respectively, as shown by (125) in the main text and detailed in the following. The single-bit state $b \in \{0, 1\}$ is given by a fair coin toss, i.e., chosen uniformly at random from $\{0, 1\}$, giving us, for each b ,

$$p(b) = \frac{1}{2}. \quad (\text{A113})$$

According to the prescription for Simon's problem, $s \in \{0, 1\}^N \setminus \{0\}$ is chosen uniformly at random, giving us

$$p(s) = \frac{1}{2^N - 1}, \quad (\text{A114})$$

because there exist $2^N - 1$ elements in $\{0, 1\}^N \setminus \{0\}$. Furthermore, for the sake of our analysis, we here write

$$p(f|b) := \sum_{s \in \{0, 1\}^N \setminus \{0\}} p(s)p(f|s, b). \quad (\text{A115})$$

Then, we have

$$p(f|b) = \begin{cases} \frac{1}{2^N!} & \text{if } b = 0 \\ \frac{1}{2^N - 1} \frac{1}{\binom{2^N}{2^N - 1} 2^{N-1}!} & \text{if } b = 1. \end{cases} \quad (\text{A116})$$

The probabilities on the right-hand side can be derived from normalization. In the case of $b = 0$, we choose 1-to-1 functions as f , and $p(f|b)$ comes from $|S_{2^N}| = 2^N!$, where S_{2^N} is the set of permutations, i.e., 1-to-1 functions, over 2^N elements in $\{0, 1\}^N$. In the case of $b = 1$, we choose 2-to-1 functions as f , and we can see $p(f|b)$ as follows. The factor $1/(2^N - 1)$ comes from $p(s)$. Moreover, to fully define the function f satisfying $f(x) = f(x \oplus s)$, half of all inputs of f can be assigned the unique corresponding outputs; that is, we have to choose $2^N/2$ ordered elements from the set $\{0, 1\}^N$ of 2^N elements. There are $\binom{2^N}{2^N - 1} 2^{N-1}!$ such choices. The product of these numbers constitutes $p(f|b)$ in (A116).

The probability distribution $p(f|b)$ allows us to describe the computer's state at the end of the computation before the output in Fig. 3, up to the reversible transformation, as

$$\rho_M = \frac{1}{p(M)} \sum_{b \in \{0, 1\}} p(b) \sum_f p(f|b) p_f(M) \rho_{f,M}, \quad (\text{A117})$$

as shown in (129) of the main text. We can furthermore split the state ρ_M into a convex sum of the case where f is 1-to-1, $\rho_M^{(b=0)}$, and the case where f is 2-to-1, $\rho_M^{(b=1)}$,

$$\rho_M = \frac{1}{2} \rho_M^{(b=0)} + \frac{1}{2} \rho_M^{(b=1)}, \quad (\text{A118})$$

where the prefactor $1/2$ comes from $p(b)$ in (A113), and the two contributions for $b = 0, 1$ are defined as

$$\rho_M^{(b)} := \frac{1}{p(M)} \sum_f p(f|b) p_f(M) \rho_{f,M}. \quad (\text{A119})$$

We use the concavity of the entropy together with the decomposition (A118) to give a lower bound of $S(\rho_M)$ by

$$S(\rho_M) \geq \frac{1}{2}S\left(\rho_M^{(b=0)}\right) + \frac{1}{2}S\left(\rho_M^{(b=1)}\right) \quad (\text{A120})$$

$$\geq \frac{1}{2}S\left(\rho_M^{(b=0)}\right), \quad (\text{A121})$$

where the last line follows from $S\left(\rho_M^{(b=1)}\right) \geq 0$. The strategy for the remainder of the proof is to find a lower bound on $S\left(\rho_M^{(b=0)}\right)$ corresponding to the 1-to-1 functions. Although this strategy may ignore the contribution from the 2-to-1 functions, it will turn out that the contribution from the 1-to-1 functions is already sufficient for our analysis.

To bound (A121) further, we cast the quantum-state picture in (A117) (i.e., the diagonal density operator) into the corresponding probability distribution to calculate its entropy. To describe the input-output relation of $y = f(x)$, let us introduce a rank-1 computational-basis projector $\Pi_{\vec{x}, \vec{y}}$ acting on the space of the $2N \times M$ bits for ρ_M in (A117) as

$$\Pi_{\vec{x}, \vec{y}} = \bigotimes_{k=1}^M |x_k, y_k\rangle \langle x_k, y_k|, \quad (\text{A122})$$

]

where \vec{x} and \vec{y} are any ordered choices of M elements in the set $\{0, 1\}^N$. In particular, we have $\vec{x} = (x_1, \dots, x_M)$ and $\vec{y} = (y_1, \dots, y_M)$, where $x_k, y_k \in \{0, 1\}^N$ for all $1 \leq k \leq M$. Note that the algorithm uses M different inputs as \vec{x} by construction, but we do not assume the outputs \vec{y} to be distinct from each other in general. To calculate the entropy $S\left(\rho_M^{(b=0)}\right)$, we consider a probability distribution obtained from the diagonal operator $\rho_M^{(b=0)}$ by

$$p(\vec{x}, \vec{y}) := \text{tr} \left[\Pi_{\vec{x}, \vec{y}} \rho_M^{(b=0)} \right], \quad (\text{A123})$$

so that it holds that

$$S\left(\rho_M^{(b=0)}\right) = - \sum_{\vec{x}, \vec{y}} p(\vec{x}, \vec{y}) \ln p(\vec{x}, \vec{y}). \quad (\text{A124})$$

Substituting $\rho_{f,M}$ in (A119) with (122) in the main text, we can express the probability $p(\vec{x}, \vec{y})$ of the $2N \times M$ bits in (A123) as

$$p(\vec{x}, \vec{y}) = \sum_f \left(\frac{1}{p(M)} q(\text{stop} | \{(x_\ell, f(x_\ell))\}_{1 \leq \ell \leq M}) q(x_M | \{(x_\ell, f(x_\ell))\}_{1 \leq \ell \leq M-1}) \right. \quad (\text{A125})$$

$$\begin{aligned} & \left. \prod_{k=1}^{M-1} q(\text{continue} | \{(x_\ell, f(x_\ell))\}_{1 \leq \ell \leq k}) q(x_k | \{(x_\ell, f(x_\ell))\}_{1 \leq \ell \leq k-1}) \times p(f | b=0) \delta_{f(x_1), y_1} \cdots \delta_{f(x_M), y_M} \right) \\ & = \overbrace{\left(\frac{1}{p(M)} q(\text{stop} | \{(x_\ell, y_\ell)\}_{1 \leq \ell \leq M}) q(x_M | \{(x_\ell, f(x_\ell))\}_{1 \leq \ell \leq M-1}) \right.}^{=: q(\vec{x} | \vec{y})} \\ & \quad \times \overbrace{\left. \prod_{k=1}^{M-1} q(\text{continue} | \{(x_\ell, y_\ell)\}_{1 \leq \ell \leq k}) q(x_k | \{(x_\ell, y_\ell)\}_{1 \leq \ell \leq k-1}) \right)}^{=: p(\vec{y})} \end{aligned}$$

$$\times \underbrace{\sum_f p(f | b=0) \delta_{f(x_1), y_1} \cdots \delta_{f(x_M), y_M}}_{=: p(\vec{y})} \quad (\text{A126})$$

$$= q(\vec{x} | \vec{y}) \times p(\vec{y}), \quad (\text{A127})$$

where $q(\vec{x} | \vec{y})$ and $p(\vec{y})$ are defined as shown in the second line, and $\delta_{a,b}$ is a delta function, i.e.,

$$\delta_{a,b} = \begin{cases} 1, & \text{if } a = b, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A128})$$

The notations of $q(\vec{x} | \vec{y})$ and $p(\vec{y})$ are chosen suggestively to distinguish between probabilities arising from the al-

gorithm's probabilistic strategy (written in q) and the setting of Simon's problem (written in p), as explained in more detail below. Note that in general, $p(\vec{y})$ may depend on \vec{y} , which can be any ordered sequence of M elements from $\{0, 1\}^N$ and may include the same elements $y_k = y_\ell$ for $1 \leq k < \ell \leq M$; however, $p(\vec{y})$ does not depend on the choice of M different inputs x_k due to uniform randomness of the choice of 1-to-1 functions

$f : \{0, 1\}^N \rightarrow \{0, 1\}^N$ in Problem 6 in the case of $b = 0$. By construction, the probability distribution $p(\vec{x}, \vec{y})$ is normalized with respect to a sum over all \vec{x} and \vec{y} , i.e.,

$$\sum_{\vec{x}, \vec{y}} p(\vec{x}, \vec{y}) = 1. \quad (\text{A129})$$

The quantity $p(\vec{y})$ can be interpreted as the probability of measuring the output sequence \vec{y} , given the computer's state ρ^C . Summing over all values of \vec{y} lets us recover the normalization, i.e.,

$$\sum_{\vec{y}} p(\vec{y}) = \sum_f p(f|b=0) \sum_{\vec{y}} \delta_{f(x_1), y_1} \cdots \delta_{f(x_M), y_M} \quad (\text{A130})$$

$$= \sum_f p(f|b=0) \times 1 = 1. \quad (\text{A131})$$

On the other hand, $q(\vec{x}|\vec{y})$ can be interpreted as the probability of querying the input sequence \vec{x} conditioned on having exactly M queries and on measuring the output sequence \vec{y} . For any choice of $M \geq 1$ and \vec{y} , the function $q(\vec{x}|\vec{y})$ is a probability distribution over \vec{x} by construction (as shown in (122) and (128) in the main text) and normalized by

$$\sum_{\vec{x}} q(\vec{x}|\vec{y}) = 1. \quad (\text{A132})$$

These interpretations also show that we can understand $p(\vec{y})$ as the distribution obtained by marginalizing $p(\vec{x}, \vec{y})$ over \vec{x} ,

$$\sum_{\vec{x}} p(\vec{x}, \vec{y}) = \left(\sum_{\vec{x}} q(\vec{x}|\vec{y}) \right) p(\vec{y}) = p(\vec{y}). \quad (\text{A133})$$

Using the fact that entropy does not increase under marginalization [9], we can bound the entropy of $p(\vec{x}, \vec{y})$ in (A124) by that of the marginal distribution $p(\vec{y})$, i.e.,

$$-\sum_{\vec{x}, \vec{y}} p(\vec{x}, \vec{y}) \ln p(\vec{x}, \vec{y}) \geq -\sum_{\vec{y}} p(\vec{y}) \ln p(\vec{y}). \quad (\text{A134})$$

This inequality allows us to analyze a lower bound of the entropy of ρ_M solely through the contribution coming from properties of Simon's problem (captured by $p(\vec{y})$), irrespectively of the contribution from the algorithm's strategy to solve it (captured by $q(\vec{x}|\vec{y})$).

In the following, we estimate the entropy $-\sum_{\vec{y}} p(\vec{y}) \ln p(\vec{y})$ in (A134). To this end, we further split the sum over all possible output sequences \vec{y} in (A134) into one for M distinct outputs in $\{0, 1\}^N$ and the rest. In particular, we write the set of M distinct elements in $\{0, 1\}^N$ as

$$Y_{\text{distinct}} = \{ \vec{y} = (y_1, \dots, y_M) \in (\{0, 1\}^N)^M : y_k \neq y_\ell \text{ for all } 1 \leq k < \ell \leq M \}. \quad (\text{A135})$$

and its complement as

$$Y_{\text{indistinct}} = \{ \vec{y} = (y_1, \dots, y_M) : y_i \in \{0, 1\}^N \} \setminus Y_{\text{distinct}}. \quad (\text{A136})$$

For any $\vec{y} \in Y_{\text{indistinct}}$ that includes indistinct pairs of $y_k = y_\ell$, we have

$$p(\vec{y}) = 0. \quad (\text{A137})$$

To see this, recall that $p(\vec{y})$ in (A125) is defined as a sum over all 1-to-1 functions, but the input \vec{x} to each 1-to-1 function is a sequence of distinct values, i.e., $x_k \neq x_\ell \in \{0, 1\}^N$ for all $1 \leq k < \ell \leq M$. Thus, the sequence of outputs $f(x_1), \dots, f(x_M)$ of any 1-to-1 function is also a sequence of pairwise distinct values in $\{0, 1\}^N$. As a consequence, we have (A137) if $\vec{y} \notin Y_{\text{distinct}}$. On the other hand, for any choice of M distinct values $\vec{y} \in Y_{\text{distinct}}$, the probability $p(\vec{y})$ does not depend on \vec{y} ; i.e., it holds for some constant $p_{\text{distinct}} > 0$ that

$$p(\vec{y}) = p_{\text{distinct}}. \quad (\text{A138})$$

This holds due to the uniform randomness of the choice of 1-to-1 functions $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$ in Problem 6, i.e.,

$$p(\vec{y}) = p(\pi(\vec{y})), \quad (\text{A139})$$

for any permuted vector $\pi(\vec{y}) = (\pi(y_1), \dots, \pi(y_M))$. Using the normalization (A131), we further have

$$1 = \sum_{\vec{y} \in Y_{\text{distinct}}} p(\vec{y}) + \sum_{\vec{y} \in Y_{\text{indistinct}}} p(\vec{y}) \quad (\text{A140})$$

$$= \sum_{\vec{y} \in Y_{\text{distinct}}} p(\vec{y}) + 0 \quad (\text{A141})$$

$$= |Y_{\text{distinct}}| p_{\text{distinct}}. \quad (\text{A142})$$

A combinatorial argument yields

$$|Y_{\text{distinct}}| = \binom{2^N}{M} M! = \frac{2^N!}{(2^N - M)!}, \quad (\text{A143})$$

which is the number of ordered choices of M distinct elements in the set $\{0, 1\}^N$ of 2^N elements. Thus, we have

$$p_{\text{distinct}} = \frac{(2^N - M)!}{2^N!}. \quad (\text{A144})$$

Therefore, it holds that

$$-\sum_{\vec{y}} p(\vec{y}) \ln p(\vec{y}) \quad (\text{A145})$$

$$= -\sum_{\vec{y} \in Y_{\text{distinct}}} p(\vec{y}) \ln p(\vec{y}) - \sum_{\vec{y} \in Y_{\text{indistinct}}} p(\vec{y}) \ln p(\vec{y}) \quad (\text{A146})$$

$$= -\sum_{\vec{y} \in Y_{\text{distinct}}} p(\vec{y}) \ln p(\vec{y}) + 0 \quad (\text{A147})$$

$$= \ln \frac{2^N!}{(2^N - M)!}. \quad (\text{A148})$$

Consequently, it follows from (A121), (A124), (A134), and (A148) that

$$S(\rho_M) \geq \frac{1}{2} \ln \frac{2^N!}{(2^N - M)!}, \quad (\text{A149})$$

which leads to the conclusion. \square

Finally, we show Lemma 14 below to evaluate the fraction of factorials appearing in the lower bound of the entropy in the main text.

Lemma 14 (Estimation of fraction of factorials). *For*

any $N \in \{1, 2, \dots\}$ and $\Delta \in (0, 1/6)$, it holds that

$$\ln \frac{2^N!}{(2^N - M_{N,\Delta})!} \geq \sqrt{\frac{2\Delta}{1+\Delta}} 2^{N/2} N \ln 2 - 3, \quad (\text{A150})$$

where $M_{N,\Delta}$ is defined as (115).

Proof. We bound the factorials using a variant of Stirling's approximation [89]

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\frac{1}{12n+1}} \leq n! \leq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\frac{1}{12n}}, \quad (\text{A151})$$

which holds for arbitrary positive integers n . In particular, with $M_{N,\Delta}$ denoting $\left\lceil \sqrt{\frac{2\Delta}{1+\Delta}} \times 2^{N/2} \right\rceil$ as in (115), we have

$$\begin{aligned} & \ln \frac{2^N!}{(2^N - M_{N,\Delta})!} \\ & \geq \ln \frac{\sqrt{2\pi 2^N} \left(\frac{2^N}{e}\right)^{2^N} e^{\frac{1}{12 \times 2^N + 1}}}{\sqrt{2\pi(2^N - M_{N,\Delta})} \left(\frac{2^N - M_{N,\Delta}}{e}\right)^{2^N - M_{N,\Delta}} e^{\frac{1}{12(2^N - M_{N,\Delta})}}} \end{aligned} \quad (\text{A152})$$

$$= \ln \left(\left(\frac{2^N}{2^N - M_{N,\Delta}} \right)^{2^N + \frac{1}{2}} (2^N - M_{N,\Delta})^{M_{N,\Delta}} e^{-M_{N,\Delta} + \frac{1}{12 \times 2^N + 1} - \frac{1}{12(2^N - M_{N,\Delta})}} \right) \quad (\text{A153})$$

$$= - \left(2^N + \frac{1}{2} \right) \ln \left(1 - \frac{M_{N,\Delta}}{2^N} \right) + M_{N,\Delta} (N \ln 2 - 1) - \left(-M_{N,\Delta} \ln \left(1 - \frac{M_{N,\Delta}}{2^N} \right) + \frac{1}{12} \left(\frac{1}{2^N - M_{N,\Delta}} - \frac{1}{2^N + \frac{1}{12}} \right) \right) \quad (\text{A154})$$

$$\geq \left(2^N + \frac{1}{2} \right) \frac{M_{N,\Delta}}{2^N} + M_{N,\Delta} (N \ln 2 - 1) - \left(\frac{\frac{M_{N,\Delta}^2}{2^N}}{1 - \frac{M_{N,\Delta}}{2^N}} + \frac{1}{12} \left(\frac{1}{2^N - M_{N,\Delta}} - \frac{1}{2^N + \frac{1}{12}} \right) \right) \quad (\text{A155})$$

$$\geq M_{N,\Delta} + M_{N,\Delta} (N \ln 2 - 1) - 3 \quad (\text{A156})$$

$$= M_{N,\Delta} N \ln 2 - 3 \quad (\text{A157})$$

$$\geq \sqrt{\frac{2\Delta}{1+\Delta}} 2^{N/2} N \ln 2 - 3, \quad (\text{A158})$$

where (A155) follows from

$$\frac{x}{1+x} \leq \ln(1+x) \leq x \quad (\text{A159})$$

for all $x > -1$ with $x = -M_{N,\Delta}/2^N$, and (A156) holds due to

$$2^N + \frac{1}{2} \geq 2^N, \quad (\text{A160})$$

$$\frac{\frac{M_{N,\Delta}^2}{2^N}}{1 - \frac{M_{N,\Delta}}{2^N}} + \frac{1}{12} \left(\frac{1}{2^N - M_{N,\Delta}} - \frac{1}{2^N + \frac{1}{12}} \right) \leq 3 \quad (\text{A161})$$

for all $N \in \{1, 2, \dots\}$ and $\Delta \in (0, 1/6)$. To see that (A161) holds true, we define a function representing

the left-hand side of (A161), i.e.,

$$F(N, \Delta) := \frac{\frac{M_{N,\Delta}^2}{2^N}}{1 - \frac{M_{N,\Delta}}{2^N}} + \frac{1}{12} \left(\frac{1}{2^N - M_{N,\Delta}} - \frac{1}{2^N + \frac{1}{12}} \right), \quad (\text{A162})$$

and observe that this function has a constant order $F(N, \Delta) = O(1)$ as $N \rightarrow \infty$, in particular,

$$\lim_{N \rightarrow \infty} F(N, \Delta) = \lim_{N \rightarrow \infty} \frac{M_{N,\Delta}^2}{2^N} = \frac{2\Delta}{1+\Delta}; \quad (\text{A163})$$

then, for every N , due to the monotonicity of $F(N, \Delta)$

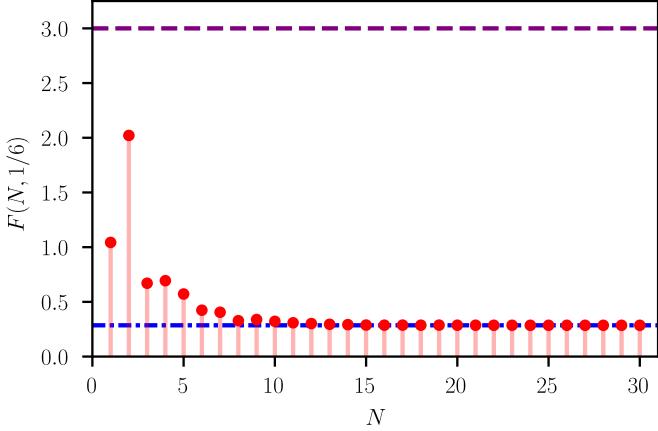


FIG. 6. A plot of $F(N, 1/6)$ in (A162) for $N \in \{1, 2, \dots, 30\}$ (red dots). As discussed in the main text, we have $\lim_{N \rightarrow \infty} F(N, 1/6) = 2\Delta/(1 + \Delta)|_{\Delta=1/6} = 2/7$ (blue dashed-dotted line). As the plot shows, the function $F(N, 1/6)$ is upper bounded by 3 (purple dashed line) since it takes the maximum at $N = 2$ with $F(2, 1/6) = 2.02 \dots \leq 3$.

in terms of Δ , we have

$$F(N, \Delta) \leq F(N, 1/6) \in \left[\frac{2}{7}, 3 \right], \quad (\text{A164})$$

where the range can be confirmed also from Fig. 6. \square

4. Proof on classical upper bound for solving Simon's problem

In this appendix, we present a proof of the achievable upper bound of query complexity of Simon's problem in Problem 6 by the classical algorithm with M uniformly random queries, as described in Proposition 10 of Sec. V. The statement of Proposition 10 is repeated again below for readability.

Proposition 10 (Refined analysis of upper bound of query complexity of Simon's problem for a classical algorithm). *Fix any N . For any $\delta > 0$, if we have*

$$M = \left\lceil \frac{1}{2} \sqrt{\left(8 \ln \frac{1}{\delta} \right) 2^N + 1} + \frac{1}{2} \right\rceil = O(2^{N/2}), \quad (\text{160})$$

the classical algorithm that makes the uniformly random M (or more) queries for solving Simon's problem

in Problem 6 achieves the following with a high probability greater than $1 - \delta$:

- if there exists any $s \in \{0, 1\}^N \setminus \{0\}$ such that f is a 2-to-1 function satisfying $f(x) = f(x \oplus s)$, then this M -query algorithm outputs $a = 1$;
- if f is a 1-to-1 function, this M -query algorithm outputs $a = 0$.

Proof. Recall that, as described in Sec. V, with the M different inputs $\vec{x} = (x_1, \dots, x_M)$ chosen uniformly at random, the algorithm works as follows:

- if at least one pair (x_k, x_m) satisfying $f(x_k) = f(x_m)$ for some $x_k \neq x_m$ is found in the M queries, the algorithm has an evidence that f is a 2-to-1 function and thus outputs $a = 1$;
- otherwise, the algorithm considers f to be 1-to-1 and outputs $a = 0$.

If f is a 1-to-1 function, the algorithm always returns the correct output $a = 0$; thus, it suffices to analyze the probability of outputting $a = 0$ in the case of f being 2-to-1. For any $s \in \{0, 1\}^N \setminus \{0\}$, the set $\{0, 1\}^N$ is divided into $2^N/2$ equivalence classes with an equivalence relation $x \sim x \oplus s$. If the uniformly random M queries of the algorithm include a pair of x and $x \oplus s$ in the same equivalence class for some x , then the algorithm finds $f(x) = f(x \oplus s)$ and thus outputs $a = 1$ correctly. The false output, i.e., $a = 0$ for the 2-to-1 function f , arises when all the M queries x_1, \dots, x_M are in different equivalence classes.

The failure probability is given by a fraction of the number of the sequences that give a false output, divided by the number of possible sequences. The number of sequences that return the false output is given by $\binom{2^N/2}{M} 2^M$, where M unordered inputs are chosen among the $2^N/2$ equivalence classes, which amounts to the factor $\binom{2^N/2}{M}$, but for each of the M equivalence classes, there are 2 possible choices x or $x \oplus s$, giving the other factor 2^M . On the other hand, the number of possible sequences is given by $\binom{2^N}{M}$, the number of unordered choices of M elements in 2^N . All-in-all, the failure probability is given by

$$\frac{\binom{2^N/2}{M} 2^M}{\binom{2^N}{M}}, \quad (\text{A165})$$

and we have a bound on this failure probability by

$$\frac{\binom{2^N/2}{M} 2^M}{\binom{2^N}{M}} = \frac{2^{N-1}(2^{N-1}-1)(2^{N-1}-2) \dots (2^{N-1}-(M-1)) \times 2^M}{2^N(2^N-1)(2^N-2) \dots (2^N-(M-1))} \quad (\text{A166})$$

$$= \frac{2^N(2^N-2)(2^N-4) \dots (2^N-2(M-1))}{2^N(2^N-1)(2^N-2) \dots (2^N-(M-1))} \quad (\text{A167})$$

$$= 1 \left(1 - \frac{1}{2^N - 1}\right) \left(1 - \frac{2}{2^N - 2}\right) \cdots \left(1 - \frac{M-1}{2^N - (M-1)}\right) \quad (\text{A168})$$

$$\leq 1 \left(1 - \frac{1}{2^N}\right) \left(1 - \frac{2}{2^N}\right) \cdots \left(1 - \frac{M-1}{2^N}\right) \quad (\text{A169})$$

$$\leq e^{-0} e^{-1/2^N} e^{-2/2^N} \cdots e^{-(M-1)/2^N} \quad (\text{A170})$$

$$= e^{-M(M-1)/2^{N+1}}, \quad (\text{A171})$$

where the last inequality follows from $e^x \geq 1 + x$ for all x .

Therefore, for any M satisfying

$$M > \frac{1}{2} \sqrt{\left(8 \ln \frac{1}{\delta}\right) 2^N + 1} + \frac{1}{2}, \quad (\text{A172})$$

we have

$$e^{-M(M-1)/2^{N+1}} < \delta. \quad (\text{A173})$$

Since M is an integer, with M chosen as (160), the classical algorithm with the M or more uniformly random queries achieves the success probability in solving Simon's problem greater than $1 - \delta$. \square

Appendix B: Details of control-cost model

In this appendix, we aim to provide a clearer explanation of the concept of control cost discussed in Sec. II C of the main text. In the main text, we have argued that when performing a unitary operation on a computer, it can require more energy than just the change of energy in the memory alone. We call this cost the control cost. We have introduced a generic constant to account for the control cost so that our analysis should not depend on a specific architecture of the computer, whereas the exact value of control costs may depend on specific models being used. Crucially, we assume that the control cost is non-negative yet at most in the order of the energy scale of the Hamiltonian, as shown in (19) and (38). Here, we present a toy example that exhibits these properties. To this end, we sketch a model that can implement arbitrary unitaries to a single-qubit target system assisted by an auxiliary system but may incur the control cost, based on the model shown, e.g., in Ref. [122].

Interactions between systems are described by interaction Hamiltonians, and the time evolution of the quantum states, according to the Hamiltonians, conserves energy. If we work with non-energy-degenerate qubits as our computational platforms, such as trapped ions, superconducting qubits, and optics, then the energy conservation restricts the set of possible operations implementable by the Hamiltonian dynamics. In particular, if the quantum computer's memory is described by the Hamiltonian H_S , then unitary operations U imple-

mentable within the energy-conservation law would satisfy a condition of

$$[H_S, U] = 0. \quad (\text{B1})$$

In more general cases, a unitary operation to be applied to a target system may not satisfy this energy-conservation condition; however, the implementation can be made to conserve energy by enlarging the system it acts on to include an auxiliary system. Such an auxiliary system for the control may also be called a catalyst or battery. For a selection of references where this setting is discussed in further detail, see Refs. [122–125]. The local operation on the target system is then only the effective operation that arises by tracing out the auxiliary system.

Toy example based on Ref. [122]. Let us model the target system $\mathcal{H}_S = \mathbb{C}^2$ as a qubit and the auxiliary system \mathcal{H}_C for the control as a harmonic oscillator, i.e., a half-infinite ladder. The two systems are governed by the Hamiltonian

$$H = H_S + H_C = \omega |1\rangle\langle 1|_S \otimes \mathbb{1}_C + \mathbb{1}_S \otimes \sum_{n \geq 0} n\omega |n\rangle\langle n|_C, \quad (\text{B2})$$

where $\{|0\rangle_S, |1\rangle_S\}$ and $\{|n\rangle_C : n = 0, 1, \dots\}$ are orthonormal bases of \mathcal{H}_S and \mathcal{H}_C , and $\mathbb{1}_C$ and $\mathbb{1}_S$ are the identity operators acting on \mathcal{H}_S and \mathcal{H}_C , respectively. If we desire to perform a unitary U on our target qubit, we can do so in an energy-preserving way by defining the dilation on the composite system $\mathcal{H}_S \otimes \mathcal{H}_C$ as

$$V[U] := |0\rangle\langle 0|_S \otimes |0\rangle\langle 0|_C + \sum_{n \geq 1} \sum_{i,j=0}^1 U_{i,j} |i\rangle\langle j|_S \otimes |n-i\rangle\langle n-j|_C, \quad (\text{B3})$$

where we write $U_{i,j} := \langle i|U|j\rangle$. The operator $V[U]$ commutes with $H = H_S + H_C$ in (B2) because it commutes termwise. In particular, the term in the first line of (B3) commutes with H because it is diagonal in terms of the same basis. The same is true for all terms in the second line with $i = j$. All other terms are of the form

$$(\text{constant}) \times |0\rangle\langle 1|_S \otimes |n\rangle\langle n-1|_C + (\text{h.c.}) \quad (\text{B4})$$

where (h.c.) represents the hermitian conjugate of the first term, and $n > 0$ is some integer. These terms also commute since they preserve the degenerate subspaces

$\text{span}_{\mathbb{C}}\{|0\rangle_S \otimes |n\rangle_C, |1\rangle_S \otimes |n-1\rangle_C\}$ of H (recall, two diagonalizable matrices commute if and only if they are simultaneously diagonalizable, i.e., if and only if they preserve each other's eigenspace).

The question we now have is: for a given input state $|\psi\rangle_S = \sum_j \psi_j |j\rangle$ of the target system and an initially pure state $|\phi\rangle_C$ of the auxiliary system given by

$$|\phi\rangle_C = \sum_n \phi_n |n\rangle = \frac{1}{\sqrt{L}} \sum_{n=0}^{L-1} |n + \ell_0\rangle_C \quad (\text{B5})$$

with some constant $\ell_0 \geq 1$, how close are $U|\psi\rangle_S$ and $\rho'_S = \text{tr}_C\{V[U]|\psi\rangle\langle\psi|_S \otimes |\phi\rangle\langle\phi|_C V[U]^\dagger\}$ in general? We write the approximate implementation (in (B3)) of U acting on the subsystem S as a channel

$$\Lambda_U(\psi_S) = \text{tr}_C [V[U](|\psi\rangle\langle\psi|_S \otimes |\phi\rangle\langle\phi|_C)V[U]^\dagger]. \quad (\text{B6})$$

To estimate the quality of Λ_U in approximating U , we calculate how $V[U]$ acts on the two spaces $\mathcal{H}_S \otimes \mathcal{H}_C$. For this purpose, we define another state of the auxiliary system

$$|\phi'\rangle_C = \frac{1}{\sqrt{L-2}} \sum_{n=1}^{L-2} |n + \ell_0\rangle_C, \quad (\text{B7})$$

which overlaps with $|\phi\rangle_C$ up to the two boundary terms $|\ell_0\rangle_C$ and $|\ell_0 + L - 1\rangle_C$. With the help of this state, we can concisely write how $V[U]$ acts on $|\psi\rangle_S \otimes |\phi\rangle_C$, i.e.,

$$\begin{aligned} V[U](|\psi\rangle_S \otimes |\phi\rangle_C) &= \sqrt{\frac{L-2}{L}} (U|\psi\rangle_S) \otimes |\phi'\rangle_C \\ &\quad + \sqrt{\frac{2}{L}} |\chi\rangle_{SC}, \end{aligned} \quad (\text{B8})$$

where $|\chi\rangle_{SC}$ is some state orthogonal to $(U|\psi\rangle_S) \otimes |\phi'\rangle_C$. This state moreover satisfies

$$\langle\chi|_{SC} (\mathbb{1}_S \otimes |\phi'\rangle_C) = 0. \quad (\text{B9})$$

Calculating the reduced density matrix $\rho'_S = \Lambda_U(\psi_S)$ yields, for any initial state $|\psi\rangle$,

$$\rho'_S = \left(1 - \frac{2}{L}\right) U|\psi\rangle\langle\psi|_S U^\dagger + \frac{2}{L} \chi_S, \quad (\text{B10})$$

where $\chi_S = \text{tr}_C[|\chi\rangle\langle\chi|_{SC}]$ is the remainder from the partial trace, and no cross terms from $(U|\psi\rangle_S) \otimes |\phi'\rangle_C \langle\chi|_{SC} + (\text{h.c.})$ appear because of orthogonality shown in (B9).

Now, we can prove that in the limit of $L \rightarrow \infty$, the channel Λ_U converges to the original unitary channel for U . In terms of channel fidelity $f(\Lambda_U, U)$ (see, e.g., Refs. [126–128]), we can quantify how well Λ_U approximates the unitary channel U on average over the input states. In our case, we have

$$f(\Lambda_U, U) = \int d\mu(\psi) \langle\psi|U^\dagger \Lambda_U(\psi) U|\psi\rangle_S \quad (\text{B11})$$

$$\stackrel{(\text{B10})}{=} \int d\mu(\psi) \left(\langle\psi|U^\dagger U|\psi\rangle \langle\psi|U^\dagger U|\psi\rangle - O\left(\frac{1}{L}\right) \right)$$

$$= \int d\mu(\psi) \left(1 - O\left(\frac{1}{L}\right) \right) \quad (\text{B12})$$

$$= 1 - O\left(\frac{1}{L}\right), \quad (\text{B13})$$

where the integral is over the normalized and uniform measure $d\mu(\psi)$ of the projective space $P(\mathcal{H}_S) \cong \mathbb{C}P^1$ of the single-qubit Hilbert space $\mathcal{H}_S \cong \mathbb{C}^2$. The bound (B13) follows from the approximation using (B10) and the negative sign in front of the big- O term comes from $|\langle\psi|\rho'_S|\psi\rangle_S| < 1$ for all states $\rho'_S \neq |\psi\rangle\langle\psi|_S$. This shows that in the limit of large L , the two unitary channels corresponding to U and Λ_U coincide.

Emergence of the control cost. In Ref. [122], it is argued that the channel on the auxiliary system for the control, defined by $\rho_C \mapsto \rho'_C = \text{tr}_S\{V(U)\rho_S \otimes \rho_C V(U)^\dagger\}$, always increases the entropy of the auxiliary system. As a consequence, we can argue that the cost of resetting the state of the auxiliary system for the control will always be positive within this framework since the Landauer-erasure cost of reinitializing the auxiliary system is bounded by the positive entropy difference $S(\rho'_C) - S(\rho_C) \geq 0$ [23]. This fact also justifies the assumption that $E_{\text{ctrl.}} \geq 0$ in (19).

To show this fact quantitatively, we here bound the increase of the entropy of the auxiliary system on average over the initial states $|\psi\rangle_S$ of the target system in performing $V[U]$. To this end, we compute the average entropy change for the auxiliary system as

$$\langle\Delta S_C\rangle := \int d\mu(\psi) |S(\rho'_C) - S(|\phi\rangle\langle\phi|_C)| \quad (\text{B14})$$

$$= \int d\mu(\psi) S(\rho'_C), \quad (\text{B15})$$

where the integral is over all possible initial states on the target system. A lower bound of $S(\rho'_C)$ can be established by using concavity of the entropy as

$$S(\rho'_C) = S(\rho'_S) \quad (\text{B16})$$

$$= S\left(\left(1 - \frac{2}{L}\right) U|\psi\rangle\langle\psi|_S U^\dagger + \frac{2}{L} \chi_S\right) \quad (\text{B17})$$

$$\geq \frac{2}{L} S(\chi_S), \quad (\text{B18})$$

where the first equality follows from the fact that the state $V[U](|\psi\rangle_S \otimes |\phi\rangle_C)$ in (B8) is pure, the second equality from (B10), and we write $\chi_S = \text{tr}_C[|\chi\rangle\langle\chi|_{SC}]$. The entropy $S(\chi_S) \geq 0$ may depend on U and $|\psi\rangle_S$ but is independent of the value of $L \geq 2$, which can be seen by explicitly calculating (B8) and observing that no factor of L appears in $|\chi\rangle_{SC}$. For some particular choices of unitary U , it may hold that $S(\chi_S) = 0$ or even $S(\rho'_C) = 0$ (for example, for $U = \mathbb{1}$). However, for general unitaries, the entropy does not vanish. Under the assumption that the entropy $S(\chi_S) > 0$ is strictly positive for at least a

constant fraction of the initial state $|\psi\rangle_S$ among all possible choices, we asymptotically have

$$\langle \Delta S_C \rangle = \Omega\left(\frac{1}{L}\right), \quad (\text{B19})$$

which provides a lower bound of the control cost arising from reinitializing the auxiliary system by the Laudauer-erasure protocol per implementing U with $V[U]$.

Apart from this control cost of reinitialization, in practice, one may also need to account for the energetic cost required for preparing the state (B5) of the auxiliary system. We also note that when connecting these results to actual energy requirements in an experimental implementation of a quantum computer, other control costs like the ones to control the timing of the operations in implementing $V[U]$ should also be considered [75]. As for the energetic cost, by definition of $|\phi\rangle_C$ and H_C , we have

$$\langle \phi | H_C | \phi \rangle = \Omega(\omega L). \quad (\text{B20})$$

This scaling shows that preparing the state (B5) of the auxiliary system comes at an energetic cost that grows with the target channel fidelity for the implementation Λ_U to approximate the gate U . Letting $\varepsilon = 1 - f(\Lambda_U, U)$ denote the infidelity, we can combine (B13) and (B20) to qualitatively recover the scaling derived in Ref. [32], i.e.,

$$\langle \phi | H_C | \phi \rangle = O\left(\frac{\|H_S\|_\infty}{\varepsilon}\right). \quad (\text{B21})$$

A detailed analysis and the conception of a more optimal auxiliary system for the control are also given in Ref. [32], where the linear dependency of the control cost $E_{\text{ctrl.}} = O(E_{\text{qubit}})$ is also derived for a certain model, as in (38) of the main text. Note that a similar bound can also be obtained from the results in Refs. [34, 35, 129–131]. These bounds show that the assumptions on the control cost in our analysis indeed hold for the conventional theoretical models that we believe capture the essence of physics in practice.

[1] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).

[2] A. Auffèves, Quantum Technologies Need a Quantum Energy Initiative, *PRX Quantum* **3**, 020101 (2022).

[3] D. Jaschke and S. Montangero, Is quantum computing green? an estimate for an energy-efficiency quantum advantage, *Quantum Science and Technology* **8**, 025001 (2023).

[4] J. Watrous, Quantum Computational Complexity, in *Encyclopedia of Complexity and Systems Science*, edited by R. A. Meyers (Springer New York, New York, NY, 2009) pp. 7174–7201.

[5] J. Watrous, On the complexity of simulating space-bounded quantum computations, *Comput. Complex.* **12**, 48–84 (2004).

[6] P. Benioff, Quantum mechanical models of turing machines that dissipate no energy, *Phys. Rev. Lett.* **48**, 1581 (1982).

[7] R. P. Feynman, Quantum mechanical computers, *Foundations of Physics* **16**, 507 (1986).

[8] M. Fellous-Asiani, J. H. Chai, Y. Thonnart, H. K. Ng, R. S. Whitney, and A. Auffèves, Optimizing resource efficiencies for scalable full-stack quantum computers, *PRX Quantum* **4**, 040319 (2023).

[9] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, 2010).

[10] H. Buhrman and R. de Wolf, Complexity measures and decision tree complexity: a survey, *Theoretical Computer Science* **288**, 21 (2002), complexity and Logic.

[11] R. Landauer, Irreversibility and Heat Generation in the Computing Process, *IBM Journal of Research and Development* **5**, 183 (1961).

[12] D. Simon, On the power of quantum computation, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994) pp. 116–123.

[13] D. R. Simon, On the power of quantum computation, *SIAM Journal on Computing* **26**, 1474 (1997).

[14] O. A. Mukhanov, Energy-efficient single flux quantum technology, *IEEE Transactions on Applied Superconductivity* **21**, 760 (2011).

[15] W. Nernst, Ueber die Berechnung chemischer Gleichgewichte aus thermischen Messungen, *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* **1906**, 1 (1906).

[16] P. Taranto, F. Bakhshinezhad, A. Bluhm, R. Silva, N. Friis, M. P. E. Lock, G. Vitagliano, F. C. Binder, T. Debarba, E. Schwarzhans, F. Clivaz, and M. Huber, Landauer versus nernst: What is the true cost of cooling a quantum system?, *PRX Quantum* **4**, 010332 (2023).

[17] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, Challenges and opportunities in quantum machine learning, *Nature Computational Science* **2**, 567 (2022).

[18] H.-Y. Huang, R. Kueng, and J. Preskill, Information-theoretic bounds on quantum advantage in machine learning, *Phys. Rev. Lett.* **126**, 190505 (2021).

[19] D. Aharonov, J. Cotler, and X.-L. Qi, Quantum algorithmic measurement, *Nature communications* **13**, 887 (2022).

[20] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, Exponential separations between learning with and without quantum memory, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (2022) pp. 574–585.

[21] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Quantum advantage in learning from experiments, *Science* **376**, 1182 (2022).

[22] H.-Y. Huang, S. Chen, and J. Preskill, Learning to predict arbitrary quantum processes, *PRX Quantum* **4**, 040337 (2023).

[23] D. Reeb and M. M. Wolf, An improved Landauer principle with finite-size corrections, *New Journal of Physics* **16**, 103011 (2014).

[24] K. Proesmans, J. Ehrlich, and J. Bechhoefer, Finite-time

landauer principle, *Phys. Rev. Lett.* **125**, 100602 (2020).

[25] T. Van Vu and K. Saito, Finite-time quantum landauer principle and quantum coherence, *Phys. Rev. Lett.* **128**, 010602 (2022).

[26] A. Rolandi and M. Perarnau-Llobet, Finite-time Landauer principle beyond weak coupling, *Quantum* **7**, 1161 (2023).

[27] S. Bravyi, D. Gosset, and R. König, Quantum advantage with shallow circuits, *Science* **362**, 308 (2018).

[28] M. Ozawa, Conservative Quantum Computing, *Phys. Rev. Lett.* **89**, 057902 (2002).

[29] J. Gea-Banacloche and M. Ozawa, Minimum-energy pulses for quantum logic cannot be shared, *Phys. Rev. A* **74**, 060301(R) (2006).

[30] J. Ikonen, J. Salmilehto, and M. Möttönen, Energy-efficient quantum computing, *npj Quantum Information* **3**, 17 (2017).

[31] J. Stevens, D. Szombati, M. Maffei, C. Elouard, R. Assouly, N. Cottet, R. Dassonneville, Q. Ficheux, S. Zeppezauer, A. Bienfait, A. N. Jordan, A. Auffèves, and B. Huard, Energetics of a Single Qubit Gate, *Phys. Rev. Lett.* **129**, 110601 (2022).

[32] G. Chiribella, Y. Yang, and R. Renner, Fundamental Energy Requirement of Reversible Quantum Operations, *Phys. Rev. X* **11**, 021014 (2021).

[33] Y. Yang, R. Renner, and G. Chiribella, Energy requirement for implementing unitary gates on energy-unbounded systems, *Journal of Physics A: Mathematical and Theoretical* **55**, 494003 (2022).

[34] H. Tajima, N. Shiraishi, and K. Saito, Uncertainty relations in implementation of unitary operations, *Phys. Rev. Lett.* **121**, 110403 (2018).

[35] H. Tajima, N. Shiraishi, and K. Saito, Coherence cost for violating conservation laws, *Phys. Rev. Res.* **2**, 043374 (2020).

[36] Y. Guryanova, N. Friis, and M. Huber, Ideal Projective Measurements Have Infinite Resource Costs, *Quantum* **4**, 222 (2020).

[37] C. H. Bennett, Logical reversibility of computation, *IBM journal of Research and Development* **17**, 525 (1973).

[38] C. H. Bennett, The thermodynamics of computation—a review, *International Journal of Theoretical Physics* **21**, 905 (1982).

[39] A. Baumeler and S. Wolf, Free energy of a general computation, *Phys. Rev. E* **100**, 052115 (2019).

[40] Q. Dong, M. T. Quintino, A. Soeda, and M. Murao, Success-or-draw: A strategy allowing repeat-until-success in quantum computation, *Phys. Rev. Lett.* **126**, 150504 (2021).

[41] S. Yoshida, A. Soeda, and M. Murao, Universal construction of decoders from encoding black boxes, *Quantum* **7**, 957 (2023).

[42] S. Yoshida, A. Soeda, and M. Murao, Reversing unknown qubit-unitary operation, deterministically and exactly, *Phys. Rev. Lett.* **131**, 120602 (2023).

[43] Y.-A. Chen, Y. Mo, Y. Liu, L. Zhang, and X. Wang, Quantum advantage in reversing unknown unitary evolutions, [arXiv:2403.04704 \[quant-ph\]](https://arxiv.org/abs/2403.04704) (2024).

[44] A. Y. Kitaev, Quantum computations: algorithms and error correction, *Russian Mathematical Surveys* **52**, 1191 (1997).

[45] F. G. S. L. Brandão, M. Horodecki, J. Oppenheim, J. M. Renes, and R. W. Spekkens, Resource Theory of Quantum States Out of Thermal Equilibrium, *Phys. Rev. Lett.* **111**, 250404 (2013).

[46] M. Lostaglio, An introductory review of the resource theory approach to thermodynamics, *Reports on Progress in Physics* **82**, 114001 (2019).

[47] Z.-W. Liu, K. Bu, and R. Takagi, One-Shot Operational Quantum Resource Theory, *Phys. Rev. Lett.* **123**, 020401 (2019).

[48] A. Misra, U. Singh, S. Bhattacharya, and A. K. Pati, Energy cost of creating quantum coherence, *Phys. Rev. A* **93**, 052335 (2016).

[49] F. Bakhshinezhad, F. Clivaz, G. Vitagliano, P. Erker, A. Rezakhani, M. Huber, and N. Friis, Thermodynamically optimal creation of correlations, *Journal of Physics A: Mathematical and Theoretical* **52**, 465303 (2019).

[50] D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error, in *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '97 (Association for Computing Machinery, New York, NY, USA, 1997) p. 176.

[51] D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error rate, *SIAM J. Comput.* **38**, 1207 (2008).

[52] P. W. Shor, Fault-tolerant quantum computation, in *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, FOCS '96 (IEEE Computer Society, USA, 1996) p. 56.

[53] A. Y. Kitaev, Quantum computations: algorithms and error correction, *Russian Mathematical Surveys* **52**, 1191 (1997).

[54] E. Knill, R. Laflamme, and W. H. Zurek, Resilient quantum computation: error models and thresholds, *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **454**, 365 (1998).

[55] P. Aliferis, D. Gottesman, and J. Preskill, Quantum accuracy threshold for concatenated distance-3 codes, *Quantum Info. Comput.* **6**, 97 (2006).

[56] B. W. Reichardt, Fault-tolerance threshold for a distance-three quantum code, in *Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part I*, ICALP'06 (Springer-Verlag, Berlin, Heidelberg, 2006) p. 50–61.

[57] B. M. Terhal and G. Burkard, Fault-tolerant quantum computation for local non-markovian noise, *Phys. Rev. A* **71**, 012336 (2005).

[58] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *Journal of Mathematical Physics* **43**, 4452 (2002).

[59] A. G. Fowler, Proof of finite surface code threshold for matching, *Phys. Rev. Lett.* **109**, 180502 (2012).

[60] D. Gottesman, Fault-tolerant quantum computation with constant overhead, *Quantum Info. Comput.* **14**, 1338 (2014).

[61] H. Yamasaki and M. Koashi, Time-efficient constant-space-overhead fault-tolerant quantum computation, *Nature Physics* **20**, 247 (2024).

[62] D. K. Park, N. A. Rodriguez-Briones, G. Feng, R. Rahimi, J. Baugh, and R. Laflamme, Heat Bath Algorithmic Cooling with Spins: Review and Prospects, in *Electron Spin Resonance (ESR) Based Quantum Computing*, edited by T. Takui, L. Berliner, and G. Hanson (Springer New York, New York, NY, 2016) pp. 227–255.

[63] Á. M. Alhambra, M. Lostaglio, and C. Perry, Heat-Bath Algorithmic Cooling with optimal thermalization strategies, *Quantum* **3**, 188 (2019).

[64] A. Serafini, M. Lostaglio, S. Longden, U. Shackerley-Bennett, C.-Y. Hsieh, and G. Adesso, Gaussian Thermal Operations and The Limits of Algorithmic Cooling, *Phys. Rev. Lett.* **124**, 010602 (2020).

[65] R. R. Soldati, D. B. R. Dasari, J. Wrachtrup, and E. Lutz, Thermodynamics of a minimal algorithmic cooling refrigerator, *Phys. Rev. Lett.* **129**, 030601 (2022).

[66] P. Taranto, P. Lipka-Bartosik, N. A. Rodríguez-Briones, M. Perarnau-Llobet, N. Friis, M. Huber, and P. Bakhshinezhad, Efficiently cooling quantum systems with finite resources: Insights from thermodynamic geometry, [arXiv:2404.06649 \[quant-ph\]](https://arxiv.org/abs/2404.06649) (2024).

[67] H. Wilming and R. Gallego, Third Law of Thermodynamics as a Single Inequality, *Phys. Rev. X* **7**, 041033 (2017).

[68] L. Masanes and J. Oppenheim, A general derivation and quantification of the third law of thermodynamics, *Nature Communications* **8**, 14538 (2017).

[69] N. Freitas, R. Gallego, L. Masanes, and J. P. Paz, Cooling to Absolute Zero: The Unattainability Principle, *Thermodynamics in the Quantum Regime* , 597–622 (2018).

[70] L. Buffoni and M. Campisi, Cooperative quantum information erasure, *Quantum* **7**, 961 (2023).

[71] A. Rolandi, P. Abiuso, and M. Perarnau-Llobet, Collective advantages in finite-time thermodynamics, *Phys. Rev. Lett.* **131**, 210401 (2023).

[72] D. Gottesman, An introduction to quantum error correction and fault-tolerant quantum computation, in *Quantum information science and its contributions to mathematics*, Proceedings of Symposia in Applied Mathematics, Vol. 68 (American Mathematical Society, Providence, Rhode Island, 2010) pp. 13–58.

[73] X. Jiang, J. Scott, M. Friesen, and M. Saffman, Sensitivity of quantum gate fidelity to laser phase and intensity noise, *Phys. Rev. A* **107**, 042611 (2023).

[74] H. Ball, W. D. Oliver, and M. J. Biercuk, The role of master clock stability in quantum information processing, *npj Quantum Information* **2**, 16033 (2016).

[75] J. Xuereb, P. Erker, F. Meier, M. T. Mitchison, and M. Huber, The impact of imperfect timekeeping on quantum control, [arXiv:2301.10767 \[quant-ph\]](https://arxiv.org/abs/2301.10767) (2023).

[76] S. Tamiya, M. Koashi, and H. Yamasaki, Polylog-time- and constant-space-overhead fault-tolerant quantum computation with quantum low-density parity-check codes (2024), [arXiv:2411.03683 \[quant-ph\]](https://arxiv.org/abs/2411.03683).

[77] D. Aharonov, M. Ben-Or, R. Impagliazzo, and N. Nisan, Limitations of noisy reversible computation (1996), [arXiv:quant-ph/9611028 \[quant-ph\]](https://arxiv.org/abs/quant-ph/9611028).

[78] S. Deffner, J. P. Paz, and W. H. Zurek, Quantum work and the thermodynamic cost of quantum measurements, *Phys. Rev. E* **94**, 010103(R) (2016).

[79] T. Debarba, G. Manzano, Y. Guryanova, M. Huber, and N. Friis, Work estimation and work fluctuations in the presence of non-ideal measurements, *New Journal of Physics* **21**, 113002 (2019).

[80] D. Aharonov and M. Ben-Or, Fault-Tolerant Quantum Computation with Constant Error, in *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '97 (Association for Computing Machinery, New York, NY, USA, 1997) p. 176–188.

[81] D. Aharonov and M. Ben-Or, Fault-Tolerant Quantum Computation with Constant Error Rate, *SIAM J. Comput.* **38**, 1207 (2008).

[82] L. d. Rio, J. Åberg, R. Renner, O. Dahlsten, and V. Vedral, The thermodynamic meaning of negative entropy, *Nature* **474**, 61 (2011).

[83] P. W. Shor, Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer, *SIAM Journal on Computing* **26**, 1484 (1997).

[84] S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, 2009).

[85] F. Boudot, P. Gaudry, A. Guillevic, N. Heninger, E. Thomé, and P. Zimmermann, *IEEE Security & Privacy* **20**, 80 (2022).

[86] R. Cleve, Lecture notes: Introduction to quantum information processing, part 2, quantum algorithms (i) (2022).

[87] R. de Wolf, Quantum Computing: Lecture Notes, [arXiv:1907.09415 \[quant-ph\]](https://arxiv.org/abs/1907.09415) (2019).

[88] P. Koiran, V. Nesme, and N. Portier, A Quantum Lower Bound for the Query Complexity of Simon’s Problem, in *Automata, Languages and Programming*, edited by L. Caires, G. F. Italiano, L. Monteiro, C. Palamidessi, and M. Yung (Springer Berlin Heidelberg, Berlin, Heidelberg, 2005) pp. 1287–1298.

[89] H. Robbins, A remark on stirling’s formula, *The American Mathematical Monthly* **62**, 26 (1955).

[90] C. E. Shannon, The synthesis of two-terminal switching circuits, *The Bell System Technical Journal* **28**, 59 (1949).

[91] S. Goldwasser and M. Bellare, Lecture notes on cryptography, Summer course “Cryptography and computer security” at MIT (2008).

[92] M. J. Dworkin, *Advanced encryption standard (aes)*, Federal Information Processing Standards Publication 197 (2001).

[93] J. Daemen and V. Rijmen, *The design of Rijndael: AES — the Advanced Encryption Standard* (Springer-Verlag, 2002) p. 238.

[94] X. Bonnecain, M. Naya-Plasencia, and A. Schrottenloher, Quantum security analysis of AES, *IACR Transactions on Symmetric Cryptology* , 55 (2019).

[95] S. Jaques, M. Naehrig, M. Roetteler, and F. Virdia, Implementing grover oracles for quantum key search on AES and LowMC, in *Advances in Cryptology – EUROCRYPT 2020* (Springer International Publishing, 2020) pp. 280–310.

[96] K. Jang, A. Baksi, H. Kim, G. Song, H. Seo, and A. Chattopadhyay, Quantum analysis of aes, Cryptology ePrint Archive, Paper 2022/683 (2022), <https://eprint.iacr.org/2022/683>.

[97] G. Cai and D. Qiu, Optimal separation in exact query complexities for simon’s problem, *Journal of Computer and System Sciences* **97**, 83 (2018).

[98] Wikipedia, *Orders of magnitude (energy)* (2004), [Online; accessed 10-June-2024].

[99] M. Ettinger, P. Høyer, and E. Knill, The quantum query complexity of the hidden subgroup problem is polynomial, *Information Processing Letters* **91**, 43 (2004).

[100] E. Bernstein and U. Vazirani, Quantum complexity theory, *SIAM Journal on Computing* **26**, 1411 (1997).

[101] T. Yamakawa and M. Zhandry, Verifiable quantum advantage without structure, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science* (IEEE, 2022) pp. 1–12.

nual Symposium on Foundations of Computer Science (FOCS) (IEEE Computer Society, Los Alamitos, CA, USA, 2022) pp. 69–74.

[102] S. Aaronson and A. Arkhipov, The computational complexity of linear optics, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11 (Association for Computing Machinery, New York, NY, USA, 2011) p. 333–342.

[103] S. Aaronson and L. Chen, Complexity-theoretic foundations of quantum supremacy experiments, in *Proceedings of the 32nd Computational Complexity Conference*, CCC '17 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, DEU, 2017).

[104] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nature Physics* **14**, 595 (2018).

[105] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, *Nature Physics* **15**, 159 (2019).

[106] D. Shepherd and M. J. Bremner, Temporally unstructured quantum computation, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **465**, 1413 (2009).

[107] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **467**, 459 (2011).

[108] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Average-case complexity versus approximate simulation of commuting quantum computations, *Phys. Rev. Lett.* **117**, 080501 (2016).

[109] Google Quantum AI, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).

[110] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).

[111] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, Logical quantum processor based on reconfigurable atom arrays, *Nature* **626**, 58 (2024).

[112] A. Kay, Tutorial on the Quantikz Package, *arXiv:1809.03842 [quant-ph]* (2018).

[113] Z. Zhang, Estimating mutual information via kolmogorov distance, *IEEE Transactions on Information Theory* **53**, 3280 (2007).

[114] K. M. R. Audenaert, A sharp continuity estimate for the von neumann entropy, *Journal of Physics A: Mathematical and Theoretical* **40**, 8127 (2007).

[115] D. Petz, *Quantum Information Theory and Quantum Statistics*, 1st ed. (Springer Berlin Heidelberg, 2008).

[116] A. Winter, Tight uniform continuity bounds for quantum entropies: Conditional entropy, relative entropy distance and energy constraints, *Communications in Mathematical Physics* **347**, 291 (2016).

[117] T. Y. Chow, What is a Closed-Form Number?, *The American Mathematical Monthly* **106**, 440 (1999).

[118] S. R. Valluri, D. J. Jeffrey, and R. M. Corless, Some applications of the Lambert W function to physics, *Canadian Journal of Physics* **78**, 823 (2000).

[119] H. Weyl, Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen, *Mathematische Annalen* **71**, 441 (1912).

[120] C. R. Johnson, Precise intervals for specific eigenvalues of a product of a positive definite and a Hermitian matrix, in *Linear Algebra and its Applications*, Vol. 117, edited by R. A. Brualdi and H. Schneider (1989) pp. 159–164.

[121] W. Fulton, Eigenvalues, invariant factors, highest weights, and Schubert calculus, *Bulletin of the American Mathematical Society* **37**, 209 (2000).

[122] J. Åberg, Catalytic coherence, *Phys. Rev. Lett.* **113**, 150402 (2014).

[123] P. Skrzypczyk, A. J. Short, and S. Popescu, Extracting work from quantum systems, *arXiv:1302.2811 [quant-ph]* (2013).

[124] P. Skrzypczyk, A. J. Short, and S. Popescu, Work extraction and thermodynamics for individual quantum systems, *Nature Communications* **5**, 4185 (2014).

[125] M. Navascués and S. Popescu, How energy conservation limits our measurements, *Phys. Rev. Lett.* **112**, 140502 (2014).

[126] M. Horodecki, P. Horodecki, and R. Horodecki, General teleportation channel, singlet fraction, and quasidistillation, *Phys. Rev. A* **60**, 1888 (1999).

[127] M. A. Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, *Physics Letters A* **303**, 249 (2002).

[128] E. Magesan, R. Blume-Kohout, and J. Emerson, Gate fidelity fluctuations and quantum process invariants, *Phys. Rev. A* **84**, 012309 (2011).

[129] H. Tajima and K. Saito, Universal limitation of quantum information recovery: symmetry versus coherence (2022), *arXiv:2103.01876 [quant-ph]*.

[130] H. Tajima, R. Takagi, and Y. Kuramochi, Universal trade-off structure between symmetry, irreversibility, and quantum coherence in quantum processes (2022), *arXiv:2206.11086 [quant-ph]*.

[131] H. Tajima, private communication (2023).