

Quantum amplitude damping for solving homogeneous linear differential equations: A noninterferometric algorithm

João H. Romeiro ^{1,*} and Frederico Brito ^{2,†}

¹*Department of Electrical and Computer Engineering, São Carlos School of Engineering, University of São Paulo, São Carlos, Brazil*

²*Instituto de Física de São Carlos, Universidade de São Paulo, CP 369, 13560-970, São Carlos, São Paulo, Brazil*



(Received 12 May 2022; accepted 17 January 2023; published 27 January 2023)

In contexts where relevant problems can easily attain configuration spaces of enormous sizes, solving linear differential equations (LDEs) can become a hard achievement for classical computers; on the other hand, the rise of quantum hardware can conceptually enable such high-dimensional problems to be solved with a foreseeable number of qubits, while also yielding quantum advantage in terms of time complexity. Nevertheless, to bridge towards experimental realizations with several qubits and harvest such potential in a short-term basis, one must dispose of efficient quantum algorithms that are compatible with near-term projections of state-of-the-art hardware, in terms of both techniques and limitations. As the conception of such algorithms is no trivial task, insights on new heuristics are welcomed. This work proposes an approach by using the quantum amplitude damping operation as a resource to construct an efficient quantum algorithm for solving homogeneous LDEs. As the intended implementation involves performing amplitude damping exclusively via a simple equivalent quantum circuit, our algorithm shall be given by a gate-level quantum circuit (predominantly composed of elementary two-qubit gates) and is particularly nonrestrictive in terms of connectivity within and between some of its main quantum registers. We show that such an open quantum-system-inspired circuitry allows for constructing the real exponential terms in the solution in a noninterferometric way; we also provide a guideline for guaranteeing a lower bound on the probability of success for each realization, by exploring the decay properties of the underlying quantum operation.

DOI: [10.1103/PhysRevA.107.012431](https://doi.org/10.1103/PhysRevA.107.012431)

I. INTRODUCTION

The foreseeable reality of quantum technologies operating in a fully coherent manner imposes the challenge of being able to explore the underlying quantum phenomena as resources to enhance processes of interest effectively. Results ranging from quantum information processing, with the famous Shor's quantum factoring algorithm [1], to quantum heat engines [2] help foster those expectations. Nevertheless, identifying the quantum protocol that can conceptually surpass its classical counterparts (while also being adapted to near-term hardware in terms of employable paradigms and experimentally achievable techniques) has proven to be a generally nontrivial task, and the development of new strategies is welcomed.

Solving linear differential equations is one of the important examples for which classical computers can struggle to perform the task, particularly in contexts where one can quite easily reach very large configuration spaces, as the size of the problem increases, and a quantum computer can offer a real advantage. Indeed, several efficient quantum algorithms were already proposed for nonlinear differential equations [3–5] and linear differential equations [6]. Whereas the first set of examples allow for solving wide-ranging classes of equations, they lack a straightforward quantum circuit form, which could make its translation to an experimental implementation especially difficult. The second example, on the other

hand, was experimentally verified for a few qubits and is given by an equivalent quantum circuit, although it is still described in a high-level manner with abstract, potentially many-qubits quantum gates, which could hinder scaling to larger experimental setups with connectivity and other hardware limitations.

Here we introduce an efficient quantum algorithm for solving homogeneous linear differential equations (HLDEs), based on a strategy to construct the correct answer through the heuristic use of amplitude Damping (via an equivalent quantum circuit). Indeed, instead of using interference to suppress the wrong outcomes, we explore the decay properties of said quantum operation to favor the correct answer. This approach allows for an equivalent quantum circuit which (to the exception of a Hamiltonian simulation submodule) is given at a decomposed gate-level description, hence being significantly more friendly towards a potential experimental implementation, as well as towards application to use-cases of interest in foreseeable scalable hardware. The algorithm will solely use a quantum phase estimation (QPE) module followed by a section exclusively composed of elementary two-qubits gates (CNOTs and parameterized controlled one-qubit rotations) with an extra auxiliary register, responsible for implementing the nonunitary aspect of the calculations via a subsequent set of measurements. Also notably, this chain of controlled gates only requires entanglement to be created within pairs of qubits (one for each register), which can be less restrictive in terms of qubit connectivity on a quantum chip and facilitate the underlying mapping, as well as scalability.

*joao.henrique.alves@alumni.usp.br

†fbb@ifsc.usp.br

The use of quantum operations as direct resources to a specific protocol was already explored in [7] for classical verification of quantum computations. Here we exemplify their potential as building blocks to quantum algorithms as well, specifically as a versatile strategy for realizing nonunitary manipulations. Indeed, this could potentially be only an example from a family of circuit-defined algorithms inspired by open quantum system dynamics.

II. PROBLEM DELIMITATION

Our goal will be to solve HLDE, i.e., for a given set of initial conditions $x_0 \in \mathbb{C}^N$, time value $t \in \mathbb{R}_+$, and $A \in M_N(\mathbb{C})$, we would like to calculate an unknown vector $x(t) \in \mathbb{C}^N$ that obeys the following:

$$\begin{aligned} \frac{dx(t)}{dt} &= Ax(t), \\ x(0) &= x_0. \end{aligned} \tag{1}$$

In addition, a couple of extra conditions are imposed on matrix A . The primary extra condition is that A must be Hermitian, specifically because the existence of a basis of eigenvectors of A on \mathbb{C}^N (yielded by the spectral theorem) is a chore aspect on which our framework stands, as in [8]. The generalization to non-Hermitian cases can be achieved by iterating the presented algorithm and a Hamiltonian simulation module for the anti-Hermitian component via the Lie-Trotter formula [9].

As a second restriction, we will initially assume that all eigenvalues of A are strictly positive, as we deem this particular case to be illustrative of how our algorithm operates. This constraint will be promptly released in Appendix B, where it is shown that the more general case can be achieved by the same quantum circuit layout and size by simply redefining a few gate parameters.

Finally, we will hereby let the dimension N be a power of 2 for simplicity while describing the algorithm, given that the generalization to other cases should be trivial.

III. ALGORITHM'S OVERVIEW

A. Algorithm's framework

The HDLE problem from Eq. (1) has an analytical solution given by $x(t) = e^{At}x_0$. One way of going about solving this equation would be to calculate the matrix exponential e^{At} , for instance by finding the decomposition $A = VDV^{-1}$, where D is diagonal, and calculating $e^{At} = Ve^{Dt}V^{-1}$, which can be achieved efficiently if one is able to calculate the exponential of a scalar efficiently. Nevertheless, any method for a classical computer to approximate e^{At} in general requires time that scales at least as $O(N^3)$ [10], which dominates the total complexity since the subsequent matrix-vector multiplication $e^{At}x_0$ only takes quadratic time.

Alternatively, one could rewrite the right-hand side of the expression as $x(t) = \sum_{i=1}^N e^{a_i t} \tilde{c}_i v_i$, where $\{v_i\}_{1 \leq i \leq N}$ and $\{a_i\}_{1 \leq i \leq N}$ are the normalized eigenvectors and associated eigenvalues of A , respectively, and $\tilde{c}_i \equiv \langle x_0, v_i \rangle_{\mathbb{C}^N}$; this decomposed analytical expression provides a guideline for our quantum algorithm. In what follows, we will consider that

$(v_i)_{1 \leq i \leq N}$ and $(a_i)_{1 \leq i \leq N}$ are labeled in such a way that the eigenvectors are sorted in ascending order, i.e., a_N is an instance of the largest eigenvalue and thus $\|A\|_2 = a_N$. For now on, whenever we evoke $\|A\|$, it will correspond to the 2-norm, unless an alternative subindex is given.

For the sake of translating the basics of our problem into quantum, we now define the following normalized state vectors, which are, respectively, proportional to the initial conditions and the solution

$$\begin{aligned} |x_0\rangle &\equiv \frac{1}{\|x_0\|} \sum_{i=1}^N (x_0)_i |i\rangle, \\ |x(t)\rangle &\equiv \frac{1}{\sqrt{\gamma(t)}} \sum_{i=1}^N e^{-(a_N - a_i)t} c_i |v_i\rangle. \end{aligned} \tag{2}$$

For the $|x(t)\rangle$ expression, $\gamma(t)$ is a normalisation factor [with $\gamma(0) = 1$ and $\gamma(t) \leq 1$], $\forall i \in \{1, 2, \dots, N\}$ $c_i \equiv \tilde{c}_i / \|x_0\|$ and $|v_i\rangle = \sum_{j=1}^N (v_i)_j |j\rangle$. Note that the expression for $|x(t)\rangle$ is chosen such that it includes a global factor $e^{-a_N t}$ in such a way that the explicit exponential terms in Eq. (2) are either vanishing or constant in t . This choice is motivated by the gamut of operations that are more intuitively implemented via amplitude damping (the algorithm's main resource) due to its generally dissipative nature.

Our global goal will be to initialize a quantum register composed of $n \equiv \log_2(N)$ qubits (the work register) with the state $|x_0\rangle$ and to apply a series of operations to approximate $|x(t)\rangle$. We hereby identify the need to perform a nonunitary operation locally to the associated quantum register given the scaling by real exponential terms that must be performed. For equivalent reasons, we will allow for a potentially nonunit **probability of success** $p_S \in]0, 1]$, i.e., the algorithm might fail for a finite set of repetitions as long as it also outputs classical information telling if the algorithm succeeds, which will come in the form of a measure to a set of auxiliary qubits at the end of the calculations.

We also explore the following constraint: for some given A and x_0 , one should be able to affirm that the algorithm succeeds for any t with at least some nonzero probability, i.e., we should be able to put a t -independent positive lower bound on p_S . Generally, for all the instances that do not inherently meet the aforementioned criterion, we propose a very simple alternative method in Sec. IV that can also be used to boost the probability of success p_S even in cases where it already holds a positive lower-bound. One could also aim at adapting our algorithm to an amplitude amplification [11] framework to optimize the number of required repetitions.

Finally, if the algorithm succeeds, then some global characteristic $F(x) = \langle x(t) | M | x(t) \rangle$ can be extracted through the application of the corresponding observable M . Alternatively, the entire state can be approximated through several repetitions and quantum tomography [12,13]. Under any of these circumstances, all findings are linked to the original problem through the explicit relation $x(t) \doteq \sqrt{\gamma(t)} \|x_0\| e^{\|A\|t} |x(t)\rangle$.

B. Equivalent quantum circuit

The overall intuition behind the algorithm's functioning can be understood as follows: after preparing the work register in the initial condition state $|x(0)\rangle$, we let it nonunitarily

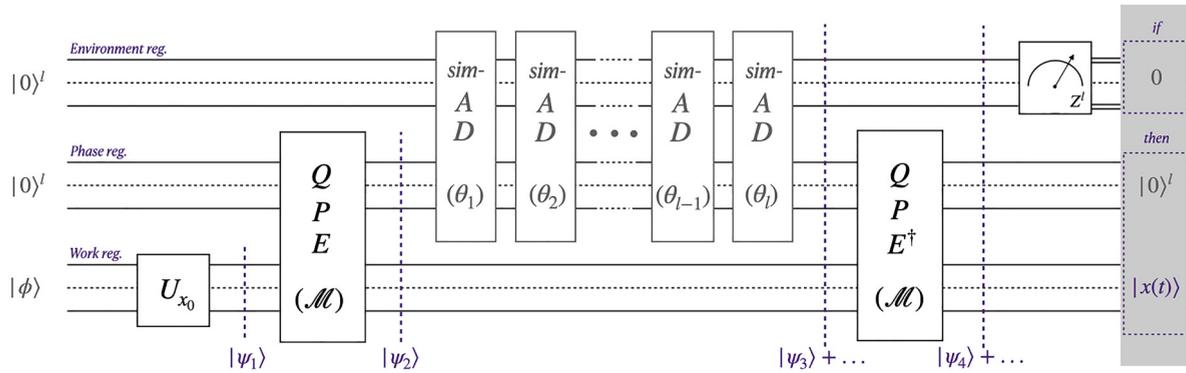


FIG. 1. Proposed quantum circuit for solving HDLEs when A is Hermitian. It includes three main quantum registers: the work register [of size $n \equiv \log_2(N)$], the phase register and the environment register [both of size $l \equiv \log_2(L)$, stipulated by the intended precision]; the circuit is divided into four modules plus measurements at the end.

evolve to $|x(t)\rangle$ by creating a set of decay effects to the qubits on an auxiliary register. The probability of such decays are conditioned on the eigenvalues $(a_i)_{1 \leq i \leq N}$ and on t and can thus be used to create the $(e^{-(a_N - a_i)t})_{1 \leq i \leq N}$ terms. Measuring the auxiliary register and verifying that no decay occurred from the other registers (i.e., measurement yields $|0\rangle$) will then project the state of the work register into the desired (potentially approximated) $|x(t)\rangle$ state.

The resulting implementation is given by the quantum circuit in Fig. 1 and is composed of four main modules, which will be covered by order of appearance in the following subsections.

1. IC preparation

As previously stated, the first stage of our algorithm corresponds to loading the initial conditions (IC) vector into the work register, resulting in $|\psi_1\rangle \equiv |x_0\rangle$ as given by Eq. (2). This is a reoccurring starting module for several algorithms aimed at solving similar problems [6,8] and also in quantum machine learning applications [14].

2. Quantum phase estimation

To create the dependency of the upcoming decay effects on the eigenvalues of A , representations of $(a_i)_{1 \leq i \leq N}$ must be loaded into an auxiliary register. First, we define the operators $\bar{A} \equiv A/||A||$ and the unitary operator $\mathcal{M} \equiv e^{-i2\pi\bar{A}}$. The second stage of our algorithm is the quantum phase estimation (QPE) [15] of \mathcal{M} , i.e., the phase estimation module is constructed via controlled- \mathcal{M} operations. The register used to code the phases (hereby called phase register) is of size $l \in \mathbb{N}^*$, thus individually corresponding to a Hilbert space of size $L \equiv 2^l$. The output of the quantum circuitry presented so far is

$$|\psi_2\rangle \equiv \sum_{i=1}^N c_i |\tilde{\phi}_i\rangle |v_i\rangle, \quad (3)$$

where the state vectors $(|\tilde{\phi}_i\rangle)_{1 \leq i \leq N}$ code the phases $(e^{i2\pi\phi_i})_{1 \leq i \leq N} \equiv (e^{-i2\pi a_i/||A||})_{1 \leq i \leq N}$, the spectrum of \mathcal{M} , in a potentially approximated manner, and are explicitly given by [16]

$$|\tilde{\phi}_i\rangle = \frac{1}{L} \sum_{k,j=0}^{L-1} e^{-\frac{i2\pi kj}{L}} e^{i2\pi k\phi_i} |j\rangle,$$

with $\phi_i \equiv 1 - a_i/||A||$. It should be noted that at least $\tilde{\phi}_N$ is an exact approximation, that is, $|\tilde{\phi}_N\rangle = |0\rangle^l$.

Indeed, it can be verified that by rewriting $|\tilde{\phi}_i\rangle$ as a superposition $\sum_{j=0}^{L-1} \beta_{i,j} |j\rangle$, then $|\beta_{i,j}\rangle$ takes a high value iff $j \approx L(1 - a_i/||A||) \pmod{L}$. If this approximation is valid for all $i \in \{1, 2, \dots, N-1\}$ and for some $j \in \{0, L-1\}$ each, then the right-hand side of Eq. (3) can be simplified as

$$|\psi_2\rangle \approx \sum_{i=1}^N c_i |d_i\rangle |v_i\rangle, \quad (4)$$

where $d_i \equiv \lfloor L(1 - a_i/||A||) \rfloor \pmod{L}$, with the $\lfloor \cdot \rfloor$ notation indicating the closest integer value. The validity of this approximation depends on the distribution of eigenvalues of \mathcal{M} around the unitary circle, but is also a fair assumption if the phase register is sufficiently large, i.e., allowing for more possible values of j . This is intuitively depicted in Fig. 2.

We shall continue by using Eq. (4) as the expression for $|\psi_2\rangle$, given that much more substantial analysis over the associated error will be provided in Appendix C.

3. sim-AD modules

The subsequent section in our quantum circuit concentrates the heuristic use of the quantum operation known as amplitude damping (AD). This operation is an important tool for describing processes that dissipate energy in open quantum systems, where a (ground) state amplitude is favored against others due to their decay. Here we explore and combine several of these amplitude decay effects to construct the exponential terms $(e^{-(a_N - a_i)t})_{1 \leq i \leq N}$ that appear on the solution. As shown in Sec. IV, such a feature can also be used to single out the correct answer through the amplitude decay of the wrong outcomes, differing from the standard approaches which rely on interferometric schemes.

Before proceeding, it should be made clear that we do not demand the manipulation of an actual open quantum system to implement our algorithm; such a requirement would not be desirable in several physical implementations since it could be hard ensuring that only such a process would be present. Nevertheless, we deem the AD operation to be heuristic in our task, as we shall use the quantum circuit shown in Fig 3, which is capable of *simulating* this effect [16], particularly if

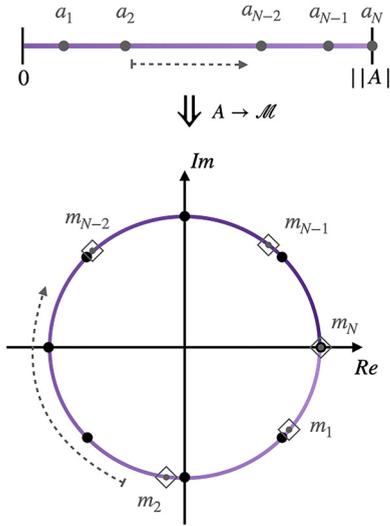


FIG. 2. Illustrative example of the eigenvalues mapping between A and \mathcal{M} , given by $(m_i)_{1 \leq i \leq N} \equiv (e^{-i2\pi a_i/|A|})_{1 \leq i \leq N}$. The eigenvalues of A are wrapped in a clockwise manner around the unit circle, as indicated by the dashed arrow. The black filled dots around the circle represent all phases that can be exactly represented by a single computational state in the phase register, in this example for $L = 8$. These can also be regarded as the allowed levels of damping that can later be performed by the *sim-AD* chain. The largest eigenvalue of A is always set to $m_N = 0$. For all other eigenvalues, a mismatch with relation to the black dots will engender a superposition of computational states in the phase register, which is a mechanism of error in the algorithm.

it were happening under completely controlled circumstances: hence our usage of the term *sim-AD*.

If the measurement result from Fig. 3 is unknown, then the output for an input density operator ρ follows the operator-sum expression given by $\rho' = E_0^k \rho E_0^{k\dagger} + E_1^k \rho E_1^{k\dagger}$, where

$$\begin{aligned} E_0^k &= \begin{bmatrix} 1 & 0 \\ 0 & \cos(\theta_k/2) \end{bmatrix}, \\ E_1^k &= \begin{bmatrix} 0 & \sin(\theta_k/2) \\ 0 & 0 \end{bmatrix}, \end{aligned} \quad (5)$$

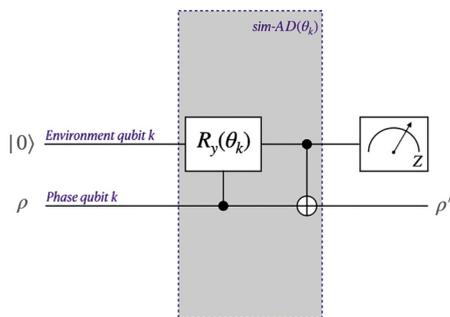


FIG. 3. White-box representation of a *sim-AD* module. This simple quantum circuitry realizes the amplitude damping quantum operation on ρ , whereas the auxiliary upper qubit plays the role of an environment. $R_y(\theta_k) \equiv e^{-i\theta_k \sigma_y/2}$, where σ_y is the complex Pauli matrix.

correspond to the AD's Kraus operators [17]. However, if the measurement result is known and equal to 0, then the final state is proportional to $E_0^k \rho E_0^{k\dagger}$, i.e., if the input is a pure state $|\psi\rangle$, the output is $\frac{E_0^k |\psi\rangle}{\sqrt{\langle \psi | E_0^{k\dagger} E_0^k | \psi \rangle}}$.

We now introduce a third and last register, called the environment register, which should also be composed of l qubits. Each AD_k module from Fig. 1 corresponds to the two-gates circuit from Fig. 3 applied to the k th phase qubit, while the state of all the other qubits is trivially operated upon. The parameters $(\theta_k)_{0 \leq k \leq l-1}$ are selected such that $\cos(\theta_k/2) \equiv e^{-|A|2^k t/L}$, leading to the following set of design equations:

$$\{\forall k \in \{0, 1, \dots, l-1\} \theta_k \equiv \arccos(2e^{-|A|2^{k+1}t/L} - 1). \quad (6)$$

Furthermore, if the measurements for each *sim-AD* $_k$ module outputs 0, the resulting (nonunitary) operator applied to the phase register is $G_0 \equiv \bigotimes_{k=0}^{l-1} E_0^k$. An expression for this operator is explicitly calculated in Appendix A, where it is shown to be equal to $\sum_{j=0}^{L-1} e^{-|A||tj/L|} |j\rangle \langle j|$.

Finally, we can look into the outcome of applying the chain of *sim-AD* modules to $|\psi_2\rangle$, if the measurement outcomes are all 0:

$$|\psi_3\rangle \approx \frac{1}{\sqrt{\gamma(t)}} \sum_{i=1}^N c_i e^{-(a_N - a_i)t} |d_i\rangle |v_i\rangle, \quad (7)$$

where, once again, a more thorough step-by-step calculation is given in Appendix A. The explicit value of the normalization constant is $\gamma(t) \equiv \sum_{i=1}^N |c_i|^2 e^{-2(a_N - a_i)t}$.

Intuitively, if each local $|1\rangle$ state is interpreted as an excitation [18], then the *sim-AD* chain models the qubit-wise decay of excitations to the environment register. This is set up such that the excitations in more significant phase qubits have a higher probability of decaying into their respective environment qubit. Since smaller eigenvalues are represented by larger phases, they are more likely to induce a decay and are thus more severely damped when the measurement of the environment register does not find any excitation.

4. Reverse operations and measurement

The final step is the uncomputation of the phase register through the reversely applied PPE module, engendering the product state $|\psi_4\rangle \approx |0\rangle^l |0\rangle^l |x(t)\rangle$. Moreover, in Fig. 1, the measurements corresponding to the *AD* stage are pushed back to the end as it is customary in quantum circuits. This does not, however, alter our calculations, nor does it temper with the probability of measuring all 0's (i.e., the probability of success p_s).

The final link to the original problem is guaranteed by noting $x(t) \doteq \sqrt{\gamma(t)} \|x_0\| e^{|A|t} |x(t)\rangle$. To calculate $\gamma(t)$ term-by-term by its expression given above, one should find the components $(c_i)_{1 \leq i \leq N}$ and the eigenvalues $(a_i)_{1 \leq i \leq N}$, which is just as hard as solving the original HLDE [10]. However, this apparent obstacle is suppressed if one notices that $p_s = \gamma(t)$ and thus that $\gamma(t)$ can be estimated with the quantum algorithm itself and does not need to be classically pre or postcalculated, much like the normalization constant in [8].

IV. SETTING THE LOWER-BOUND ON p_S

In cases where p_S goes to zero as t increases or in cases where the stationary value (lower-bound) of p_S is deemed too low for some given application, one can opt to solve the following alternative HDLE problem:

$$\begin{aligned} \frac{d}{dt} \begin{bmatrix} y(t) \\ x(t) \end{bmatrix} &= \begin{bmatrix} \|A\| \mathbb{I}_N & 0 \\ 0 & A \end{bmatrix}_{2N \times 2N} \begin{bmatrix} y(t) \\ x(t) \end{bmatrix}, \\ \begin{bmatrix} y(0) \\ x(0) \end{bmatrix} &= \begin{bmatrix} y_0 \\ x_0 \end{bmatrix}_{2N \times 1}. \end{aligned} \tag{8}$$

This is equivalent to calculating simultaneously the evolution of two noninteracting systems, one that follows $\frac{dx(t)}{dt} = Ax(t)$ and one that follows $\frac{dy(t)}{dt} = \|A\|y(t)$, the first being the one we are interested in and the second only being used to increase the lower bound on p_S .

The reasoning behind this method is to assure the existence of a (at least N -degenerate) subspace associated to the eigenvalue $\|A\|$, such that the weight decomposition of the normalized initial condition over this space is at least $\frac{\|y_0\|}{\sqrt{\|x_0\|^2 + \|y_0\|^2}}$. This implicates the existence of a non-anishing term, which also plays the role of a lower bound, in the expression of $\gamma(t)$.

Indeed, if one takes $y_0 \in \{r \in \mathbb{C}^N \mid \|r\| \geq \|x_0\|\}$, then it is guaranteed that $p_S \geq 0.5$ for any $t \in \mathbb{R}_+$. In terms of implementation, this is payed off by adding an extra qubit to the work register, doubling the dimension of the associated Hilbert space. If one aims to extract a global characteristic of $x(t)$ through an observable, the choice of y_0 must be made accordingly; for instance, if the chosen global characteristic is the absolute average $|\frac{1}{N} \sum_{i=1}^N (x(t))_i|$, then y_0 should be set such that its absolute average is 0, thus not interfering in the final result.

Moreover, since this new larger matrix inherits the very same eigenvalue distribution of A , the other registers' sizes can be chosen as if the original problem was being tackled, hence yielding the same result in both cases if all measurement outcomes are 0.

V. TIME COMPLEXITY ANALYSIS

As our algorithm makes use of a QPE module for decomposing the initial condition state vector $|x_0\rangle$ into the eigenbasis of A , one can observe parallelism of the nonunitary operations realised for all resulting subspaces, which are each conditioned on the associated entangled state $|\tilde{\phi}_i\rangle$ in the phase register. Under such a scheme, one needs not to explicitly calculate the eigenvectors and eigenvalues of A , as would be required by a general classical approach. Hence, we deem it natural to encounter a quantum speed-up on the problem's dimension. This will be explicitly shown in this section through the inspection of all underlying modules. We also conclude on the limitations imposed on the complexity for the time variable t .

A. Potential speed-up for scaling N

There are two main ways to address the IC preparation stage: either it hints at another section of a larger routine (in

which case the algorithm here described is simply a module) or an efficiently implemented unitary gate U_{x_0} should be applied to load $|x_0\rangle$ starting from some initial state $|\phi\rangle$, more closely in accordance with Fig. 1. In this second case, the loading of classical data can be thought of as some standard state-preparation method such as [19] or through a quantum random access memory (qRAM) approach [20]. Indeed, when quantum memory is established, the complexity for loading the data into a qRAM can be about $O[\log(N)]$ [21].

As for the QPE module, the implementation of \mathcal{M} translates to the problem of the Hamiltonian simulation of A for the simulated interval of time $\tau \equiv 2\pi/\|A\|$, which can be done efficiently if A is efficiently row-computable and s -sparse (i.e., if each row has at most s nonzero entries). More explicitly, the QPE requires the implementation of the family of unitary operators $(e^{-i2\pi A\tau 2^j})_{0 \leq j \leq l-1}$, applied to the work register. Several approaches for the sparse Hamiltonian simulation were proposed in past years [22–25]. A known lower bound for the query complexity of this problem is $\Omega(s\|A\|_{\max}\tau + \log(1/\epsilon_H)/\log \log[1/\epsilon_H])$, and an optimal method in all parameters of interest was already introduced [26]. The error introduced by this step (ϵ_H) will not be considered, as the dominant source of error in the context of our calculations will be the nonexact phase estimation, hence the adoption of the soft complexity $\tilde{O}(s\|A\|_{\max}\tau)$ for subsequent considerations.

In our algorithm, the total simulated time interval for implementing all the controlled- \mathcal{M} gates within the QPE module amounts for $(L-1)\tau$ and thus $O(L\|A\|)$. The subsequent inverse quantum Fourier transform, only takes $\Theta(l^2)$ steps, which can be efficiently implemented [27]. Similarly, the *sim-AD* modules contribute with a total count of $\Theta(l)$ gates applied to two qubits each. Thus, the total run-time complexity of the algorithm can be expressed as

$$O[\log(N) + sL + \log^2(L) + \log(L)]. \tag{9}$$

We now derive, in a simple way, the complexity introduced by L , given the problem's parameters and some allowed additive error ϵ when approximating $x(t)$. This approach will only acknowledge the error introduced by the best possible approximation of each eigenvalue, given that a far more detailed approach (considering the totality of components in each $|\tilde{\phi}_i\rangle$) is given in Appendix C. This second approach yields the same result here depicted, albeit slightly more complete with the addition of the condition number of A as a new term in the complexity expression.

Let $b_i \equiv (1 - \frac{j}{L})\|A\|$, for some $j \in \{0, 1, \dots, L-1\}$, be the best valid approximation of some eigenvalue a_i . Then, the algorithm errs by at most $O(\|A\|/L)$ when estimating a_i as b_i . This error is propagated to the exponential terms ($\epsilon_i \equiv |e^{b_i t} - e^{a_i t}|$) as follows:

$$\epsilon_i \leq e^{a_i t} |e^{\frac{\|A\|t}{L}} - 1| \leq e^{\|A\|t} (e^{\frac{\|A\|t}{L}} - 1).$$

Solving for L and using the asymptotic behavior $1/\ln(1 + ae^{-bx}) = \Theta(e^{bx}/a)$, found through the first-order Taylor series approximation of $\ln(1 + y)$ around $y = 0$:

$$L = O\left(\frac{\|A\|t e^{\|A\|t}}{\epsilon_i}\right).$$

More globally, if ϵ is the global additive error allowed to $x(t)$, it suffices to take $\|\epsilon_i\| \equiv \|\epsilon\|/\|x_0\|$, for all $i \in \{1, 2, \dots, N\}$. Thus

$$L = O\left(\frac{\|A\|t e^{\|A\|t} \|x_0\|}{\epsilon}\right). \quad (10)$$

By combining Eqs. (9) and (10), if s , $\|x_0\|$ and $1/\epsilon$ are polylogarithmic in N and if $\|A\|t = O(1)$ [28], the algorithm reaches an exponential speedup compared to classical methods, in relation to the size of the problem N .

B. Limitations on the complexity for scaling t

Whereas granting the aforementioned speedup in relation to the problem's size N , Eqs. (9) and (10) imply an exponential complexity in $\|A\|t$. In this section, we present two complementary arguments for why this cannot be enhanced for any similar algorithm built from applying QPE with the Hamiltonian simulation.

First, since the QPE module implicates an additive error when approximating the eigenvalues and given that this error scales linearly in relation to $1/L$, the Hilbert space dimension L must scale exponentially with t to keep the exponential terms $(e^{a_i t})_{1 \leq i \leq N}$ under some error boundary.

Second, It should be noted that the only term in Eq. (9) that takes linear L is the one related to the Hamiltonian simulation, hence dominating the run-time when it comes to scaling t . This complexity, however, cannot be vanquished, coming as a direct result of the No-Fast-Forwarding Theorem [22], which yields that the simulation of a sparse Hamiltonian cannot be achieved in sublinear complexity in relation to the simulated time, which in our construction is itself $O(L)$ (when one accounts for the accumulated simulated time among all controlled- \mathcal{M} gates in the QPE module).

VI. CONCLUSION

In summary, the quantum circuitry presented here succeeds in solving HLDEs with a potentially exponential speedup while compared to any classical method, concerning the problem's dimension N . Although no such speedup is observed for other parameters (perhaps most importantly for t), we provided the underlying results that yield the impossibility to perform better with any similar algorithm based on QPE with the Hamiltonian simulation. Also importantly, features such as the dominance of at most two-qubits elementary gates and the low connectivity requirements between the phase and environment registers boost the algorithm's appeal for scalable experimental implementations.

The usage of a *sim-AD* chain also bears novelty by allowing the solution to be found through an open quantum system-inspired scheme, creating a subspace where each state of the computational basis is adequately decayed thanks to the tuning of a set of parametric gates. We stress that such a strategy does not demand real open system dynamics. Despite its noninterferometric nature, we showed that one can also foster the decaying property of AD to favor the correct output (Sec. IV), which ultimately imposes a t -independent minimal probability of success for every single realization of the algorithm, independently of the underlying HDLE problem.

Such a scheme may create a new perspective for designing quantum algorithms based on open-systems-inspired operations, particularly for applications where specific nonunitary transformations are needed.

ACKNOWLEDGMENTS

F.B. is supported by Instituto Nacional de Ciência e Tecnologia de Informação Quântica (CNPq INCT-IQ 465469/2014-0), Brazil. The authors would like to thank Professor Diogo Soares-Pinto for the fruitful discussions.

APPENDIX A: CLOSER LOOK INTO THE NONUNITARY SECTION

The chore of the algorithm here presented lies in the *sim-AD* stage, as it enables the nonunitary calculation using exclusively two common types of one-qubit controlled gates, and with a gate depth that only scales linearly with the register size 1 (hence logarithmically in the allowed error $1/\epsilon$). In this Appendix, we explicitly derive Eq. (7), while also taking the opportunity to justify our choice of operator \mathcal{M} for the previous QPE step.

As presented before, the correct transformation over the phase register is applied if and only if the measurement outcomes are all 0. In this case, we know that the output state of the QPE stage $|\psi_2\rangle$ has been taken into the state $\frac{G_0|\psi_2\rangle}{\langle\psi_2|G_0^\dagger G_0|\psi_2\rangle}$. An uncomplicated way to find the expression for G_0 , which was used in this article is to look into how it operates over a state of the computational basis $|j\rangle$. First of all, it should be noted that the E_0^k operators can be rewritten as $E_0^k = \sum_{n=0}^1 e^{-\|A\|2^k n t/L} |n\rangle\langle n|$. If j is expressed in its binary form $(j_{l-1} j_{l-2}, \dots, j_0)_2$, it becomes clear that

$$\begin{aligned} G_0|j\rangle &= \bigotimes_{k=0}^{l-1} \sum_{n=0}^1 e^{-\|A\|2^k n t/L} |n\rangle\langle n|_{j_k} \\ &= \bigotimes_{k=0}^{l-1} e^{-\|A\|2^k j_k t/L} |j_k\rangle \\ &= e^{-\|A\|(\sum_{k=0}^{l-1} j_k 2^k) t/L} |j\rangle \\ &= e^{-\|A\|j t/L} |j\rangle. \end{aligned}$$

Thus, G_0 can be generally expressed as the following sum of projectors in the computational basis:

$$G_0 = \sum_{j=0}^{L-1} e^{-\|A\|j t/L} |j\rangle\langle j|. \quad (A1)$$

Recalling that the states $|j\rangle_{0 \leq j \leq L-1}$ code the phases of the spectrum of \mathcal{M} , Eq. (A1) shows that the damping effect increases (i.e., the exponent's absolute value increases linearly) as we scan the complex unit circle in an anticlockwise manner starting from the rightmost extremity (recall Fig. 2). While $|0\rangle$ suffers no damping, the largest allowed damping is $e^{-\|A\|(L-1)t/L} \xrightarrow{L \rightarrow \infty} e^{-\|A\|t}$. This should be enough to justify our choice of unitary operator \mathcal{M} as it linearly wraps the eigenvalues of A around the complex unit circle in such a way that subspaces related to the smallest eigenvalues suffer

more damping (i.e., are associated to larger phases within the $[0, 2\pi[$ range.

The effect of G_0 on $|\psi_2\rangle$ from Eq. (4) can now be calculated:

$$\begin{aligned} G_0|\psi_2\rangle &= \sum_{i=1}^N c_i \sum_{j=0}^{L-1} e^{-\|A\|jt/L} |j\rangle \langle j|d_i\rangle |v_i\rangle \\ &= \sum_{i=1}^N c_i e^{-\|A\|((L(1-a_i/\|A\|)) \pmod{L})t/L} |d_i\rangle |v_i\rangle. \end{aligned}$$

We can dispose of the modular notation by assuming that there is no $i \in \{1, 2, \dots, N-1\}$ such that $\lfloor L(1 - a_i/\|A\|) \rfloor = L$; note that if the negation of this statement is true, then we could observe an additive error which grows as fast as $O(e^{\|A\|t})$ since at least one subspace associated to some $a_i < \|A\|$ would suffer no damping. This is particularly bound to happen if L is not sufficiently large for a particular high value of the condition number (κ) of A . This is the case because m_1 approximates 0 in Fig. 2 through the fourth quadrant as κ increases. This consideration will be further developed in Appendix C, where it will ultimately lead to the introduction of $1/\kappa$ in a revised form of Eq. (10).

Without the modular notation, we will now provide a guideline for proving that $G_0|\psi_2\rangle \xrightarrow{L \rightarrow \infty} \sum_{i=1}^N c_i e^{-(a_N - a_i)t} |d_i\rangle |v_i\rangle$ and thus that the precision can always be enhanced by augmenting the size of the phase register. Indeed, for any $z \in [0, 1[$:

$$\left| \frac{\lfloor Lz \rfloor}{L} - z \right| \leq \left| \frac{\lfloor Lz + 1 \rfloor}{L} - z \right| = 1/L,$$

which implies $\lfloor L(1 - a_i/\|A\|) \rfloor / L \xrightarrow{L \rightarrow \infty} 1 - a_i/\|A\|$. The convergence that we aim to prove then follows from the continuity of the associated exponential functions.

APPENDIX B: GENERALIZED ALGORITHM FOR ANY HERMITIAN MATRIX

As suggested in this article, our algorithm can be extended to solve HDLE problems with any Hermitian matrix A . In fact, this task can be achieved by simply modifying a few parameters of the equivalent quantum circuit: as the number and size of quantum registers and applied quantum gates remain untouched.

The first modification comes to the QPE stage: While the $\bar{A} \equiv A/\|A\|$ operator is still defined the same way, the new form of the unitary operator \mathcal{M} should be constructed such that the interval $[-\|A\|, \|A\|]$ is mapped in one single lap around the unit circle, without overlapping both ends. The revised operator is given by

$$\mathcal{M} \equiv e^{-i2\pi(\frac{L-1}{2L}\bar{A} + \frac{L+1}{2L}\mathbb{I}_N)}. \tag{B1}$$

One may refer to Fig. 4 for a visual interpretation of \mathcal{M} . A brief side-note should be made about the implementation for the equivalent circuit: \mathcal{M} can be achieved by Hamiltonian simulation of $\frac{L-1}{2L}\bar{A} + \frac{L+1}{2L}\mathbb{I}_N$, which, for a s -sparse matrix A , is at most $(s+1)$ -sparse. As a direct result, this stage keeps the same overall complexity as presented in Sec. V for the more restrict formulation.

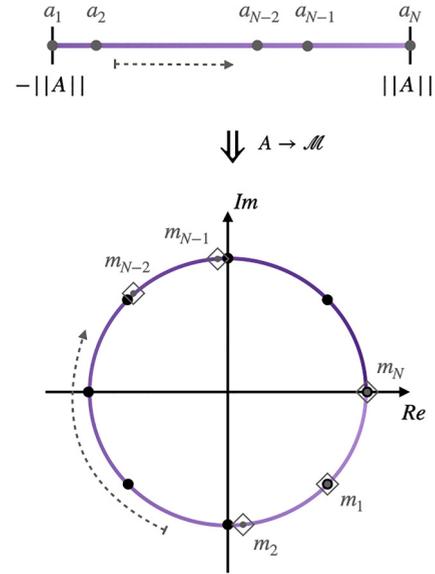


FIG. 4. Illustrative example of the eigenvalues mapping between A and \mathcal{M} , for the less restrictive Hermitian A formulation. The $[-\|A\|, \|A\|]$ interval is mapped to the complex phase interval $[0, 2\pi(L-1)/L]$. The black filled dots around the circle represent the allowed values of damping, in this case for $l = 3$.

The output of the QPE stage will still follow the same form depicted in Eq. (3), where if $|\tilde{\phi}_i\rangle = \sum_{j=0}^{L-1} \beta_{i,j}|j\rangle$, then $|\beta_{i,j}\rangle$ takes a high value if and only if $j \approx \frac{L-1}{2}(1 - a_i/\|A\|) \pmod{L}$. The mod notation is immediately dropped with no loss of generality since $\min_{1 \leq i \leq N} (a_i) \geq -\|A\|$. Thus, following the same assumptions described in Sec. III B 2:

$$|\psi_2\rangle \approx \sum_{i=1}^N c_i \left| \left[\frac{L-1}{2}(1 - a_i/\|A\|) \right] \right\rangle |v_i\rangle. \tag{B2}$$

The second and last modification to our algorithm comes to the rotation angles within the *sim*-AD modules:

$$\{\forall k \in \{0, 1, \dots, l-1\} \theta_k \equiv \arccos(2e^{-\|A\|2^{k+2}t/(L-1)} - 1)\}. \tag{B3}$$

By comparing Eq. (B3) to Eq. (6), the new expression for $G_0|\psi_2\rangle$ is

$$G_0|\psi_2\rangle = \sum_{i=1}^N c_i \left(\sum_{j=0}^{L-1} e^{-2\|A\|jt/(L-1)} |j\rangle \langle j| \right) |d_i\rangle |v_i\rangle.$$

If all measurements to the environment register result in 0, the global output is once again achieved as $|\psi_3\rangle \approx \frac{1}{\sqrt{\gamma(t)}} \sum_{i=1}^N c_i e^{-(\|A\|-a_i)t} |d_i\rangle |v_i\rangle$, with $\gamma(t) \equiv \sum_{i=1}^N |c_i|^2 e^{-2(\|A\|-a_i)t}$. However, a new subtlety compared to the more restrictive case is that a_N can now be strictly smaller than $\|A\|$. Under such circumstances, $p_S = \gamma(t) \xrightarrow{t \rightarrow \infty} 0$, i.e., the probability of success necessarily vanishes as t increases, regardless of the initial conditions. Nevertheless, this can be simply resolved by the approach presented in Sec. IV.

APPENDIX C: IN-DEPTH ANALYSIS OF THE PHASE REGISTER SIZE

In Sec. V, a simple calculation of the complexity of L was given [Eq. (10)]. In this Appendix, we present a more thorough analysis of the phase register size, which will ultimately lead to the same conclusion on the complexity, although with a new parameter, the condition number κ . We start by looking into the particular case of one single eigenvalue a_i , i.e., $|x_0\rangle \equiv |v_i\rangle$ and easily extend it to any initial condition. Indeed, one can revisit Eq. (3) and the exact expression for $|\tilde{\phi}_i\rangle$ given in Sec. III B 2. The latter should be modified upon relabeling the phase register as shown below [16]:

$$|\tilde{\phi}_i\rangle = \sum_{k=-L/2+1}^{L/2} \alpha_{j|i} |j\rangle,$$

$$\text{with } \alpha_{j|i} = \frac{1}{L} \sum_{k=-L/2+1}^{L/2} [e^{i2\pi(\phi_i - \frac{d_i+j}{L})}]^k,$$

where $\phi_i \equiv (1 - \bar{a}_i)$ and $d_i/L \in [0, 1[$ is the best l bits approximation of $\phi_i \pmod{1}$. Continuing with the procedure, the *sim*-AD modules and the inverse QPE, upon a successful final measurement and after the adequate rescaling, will result in

$$\tilde{x}(t) = \left| \sum_{k=-L/2+1}^{L/2} \alpha_{j|i} e^{b_{j|i}t} \right| v_i,$$

$$\text{with } b_{j|i} \equiv |A| \left(1 - \frac{d_i + j}{L} \right),$$

where the absolute value is introduced to discard a global phase that may appear. However, the analytical solution simply yields $x(t) = e^{a_i t} v_i$. We also introduce the overall additive error $\epsilon_i = \tilde{x}(t) - x(t)$ of the algorithm.

To proceed, the j subindexes shall be clustered into two groups, by verifying if they respect the following criteria: $|\phi_i - \frac{d_i+j}{L}| < P$, for some precision $P \in]0, 1[$ which will be later specified. Since $d_i/L \approx \phi_i$, this criteria can be simplified as $|j| \leq p$, with $p \equiv \lceil LP \rceil$, which stipulates the following upper bound on ϵ_i :

$$\|\epsilon_i\|^2 \leq \|\epsilon_{\text{QPE}}\|^2 + \|\epsilon_{P_r}\|^2,$$

$$\text{with } \|\epsilon_{\text{QPE}}\|^2 \equiv \left\| \sum_{j \notin \{-p, \dots, p-1, p\}} |\alpha_{j|i}| e^{b_{j|i}t} v_i \right\|^2,$$

$$\text{and } \|\epsilon_{P_r}\|^2 \equiv \left\| \left(\sum_{j=-p}^p |\alpha_{j|i}| e^{b_{j|i}t} - e^{a_i t} \right) v_i \right\|^2.$$

We will hereby refer to the first term as the QPE error $\|\epsilon_{\text{QPE}}\|^2$ and the second term as the precision error $\|\epsilon_{P_r}\|^2$.

1. QPE error

The QPE error can be bounded as follows:

$$\|\epsilon_{\text{QPE}}\|^2 \leq \left(\sum_{\substack{-L/2 < j < -p \\ p < j \leq L/2}} |\alpha_{j|i}|^2 \right) e^{2\|A\|t}.$$

This sum of squared probability has $\frac{1}{2(p-1)}$ as a valid upper bound [16]. Using $p \equiv \lceil 2^l P \rceil$:

$$\|\epsilon_{\text{QPE}}\|^2 \leq \frac{e^{2\|A\|t}}{2(2^l P - 1)}.$$

Solving for l , one may conclude that it suffices to use a phase register as large as

$$l \leq \log_2 \left(\frac{e^{2\|A\|t}}{2\|\epsilon_{\text{QPE}}\|^2} + 1 \right) + \log_2 \left(\frac{1}{P} \right),$$

and, thus

$$l = O \left[\|A\|t + \log \left(\frac{1}{\|\epsilon_{\text{QPE}}\|} \right) + \log \left(\frac{1}{P} \right) \right] \quad (\text{C1})$$

Equation (C1) will be revisited when we set a bound on P , which will come naturally as we delve into the precision error analysis.

2. Precision error

For every value of j within $\{-p, -p+1, \dots, p-1, p\}$, the individual error introduced can be bounded as follows:

$$|e^{b_{j|i}t} - e^{a_i t}|^2 \leq |e^{b_{-p|i}t} - e^{a_i t}|^2.$$

The right-hand side can then be rewritten as $|e^{b_{-p|i}t} - e^{b_{0|i}t}|^2 + |e^{b_{0|i}t} - e^{a_i t}|^2$ through triangular inequality. The first term in this sum translates to how close the most extreme accepted approximation ($e^{b_{-p|i}t}$) of the exact exponential term is from the best possible approximation ($e^{b_{0|i}t}$) of the exponential term and the second term corresponds to how close this best approximation meets the exact solution. While the first quadratic error will lead to a bound on P , the second will result in another bound on the phase register size l .

Indeed, the first aforementioned quadratic error can be rewritten as follows:

$$\|\epsilon_{P_r, l}\|^2 \leq e^{2b_{0|i}t} |e^{p/L} - 1|^2 \leq e^{2\|A\|t} |e^{\lceil LP \rceil / L} - 1|^2,$$

which, simplifying and solving for P , yields

$$\frac{1}{P} \leq \frac{1}{\ln(1 + \|\epsilon_{P_r, l}\| e^{-\|A\|t})},$$

and, once again using $1/\ln(1 + ae^{-bx}) = \Theta(e^{bx}/a)$:

$$\frac{1}{P} = O \left(\frac{e^{\|A\|t}}{\|\epsilon_{P_r, l}\|} \right). \quad (\text{C2})$$

One can also advocate for the inclusion of κ , the condition number of A , in the expression of $1/P$. Figure 2 gives some visual input to why this is the case. Indeed, we should not allow P to be bigger than $a_1/a_N \equiv 1/\kappa$ as this would imply that the region covered by the $\{-p, -p+1, \dots, p-1, p\}$ indexes (normally associated to significant amplitudes $|\alpha_{j|1}|$

according to our construction) would cross over to the first quadrant, where terms suffer much less damping, generating an error that could grow as fast as $O(e^{|A|t})$. Hence, it would also be more suitable to set $1/P = O(\kappa e^{|A|t}/|\epsilon_{Pr,t}|)$.

Moving to the second quadratic error term, it should be noted that the difference between the exact eigenvalue a_i and its best allowed approximation b_{0i} cannot be larger than $|A|/2L$. With that being said,

$$|\epsilon_{Pr,t}|^2 \leq e^{2a_i t} |e^{\frac{|A|t}{2L}} - 1|^2 \leq e^{2|A|t} |e^{\frac{|A|t}{2L}} - 1|^2,$$

and upon some very similar calculations to those that preceded Eq. (C2)

$$l = O\left[|A|t + \log\left(\frac{1}{|\epsilon_{Pr,t}|}\right)\right], \quad (C3)$$

where a slower growing logarithmic term on $|A|t$ was suppressed.

3. Concluding on the phase register size

Putting Eqs. (C1), (C2), and (C3) together, we stipulate that the overall expression for the phase register size so far is $l =$

$$O\left[|A|t + \log\left(\frac{1}{|\epsilon|}\right) + \log(\kappa)\right].$$

As we finally lift the constraint of looking into one single eigenvalue and move to any initial condition $x_0 = \sum_{i=1}^N \tilde{c}_i v_i$, it suffices to verify that the total error is bounded by $|\epsilon|^2 \leq \|x_0\|^2 \max_{1 \leq i \leq N} (|\epsilon_i|^2)$, and thus that one can choose $|\epsilon_i| \equiv |\epsilon|/|x_0|$ in our last expression to conclude about the asymptotic behavior imposed over l :

$$l = O\left[|A|t + \log(|x_0|) + \log\left(\frac{1}{|\epsilon|}\right) + \log(\kappa)\right]. \quad (C4)$$

[1] P. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE, New York, 1994), pp. 124–134.

[2] J. Klatzow, J. N. Becker, P. M. Ledingham, C. Weinzettl, K. T. Kaczmarek, D. J. Saunders, J. Nunn, I. A. Walmsley, R. Uzdin, and E. Poem, Experimental Demonstration of Quantum Effects in the Operation of Microscopic Heat Engines, *Phys. Rev. Lett.* **122**, 110601 (2019).

[3] S. K. Leyton and T. J. Osborne, A quantum algorithm to solve nonlinear differential equations, [arXiv:0812.4423](https://arxiv.org/abs/0812.4423).

[4] S. Lloyd, G. De Palma, C. Gokler, B. Kiani, Z.-W. Liu, M. Marvian, F. Tennie, and T. Palmer, Quantum algorithm for nonlinear differential equations, [arXiv:2011.06571](https://arxiv.org/abs/2011.06571).

[5] J.-P. Liu, H. Ø. Kolden, H. K. Krovi, N. F. Loureiro, K. Trivisa, and A. M. Childs, Efficient quantum algorithm for dissipative nonlinear differential equations, *Proc. Natl. Acad. Sci.* **118**, e2026805118 (2021).

[6] T. Xin, S. Wei, J. Cui, J. Xiao, I. Arrazola, L. Lamata, X. Kong, D. Lu, E. Solano, and G. Long, Quantum algorithm for solving linear differential equations: Theory and experiment, *Phys. Rev. A* **101**, 032307 (2020).

[7] U. Mahadev, Classical verification of quantum computations, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, 2018), pp. 259–267.

[8] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.* **103**, 150502 (2009).

[9] D. W. Berry, High-order quantum algorithm for solving linear differential equations, *J. Phys. A: Math. Theor.* **47**, 105301 (2014).

[10] C. Moler and C. Van Loan, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later, *SIAM Rev.* **45**, 3 (2003).

[11] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, *Contemporary Mathematics* **305**, 53 (2002).

[12] J. B. Altepeter, E. R. Jeffrey, and P. G. Kwiat, Photonic state tomography, *Adv. At. Mol. Opt. Phys.* **52**, 105 (2005).

[13] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).

[14] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).

[15] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Quantum algorithms revisited, *Proc. R. Soc. London A* **454**, 339 (1998).

[16] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000), pp. 380 and 383.

[17] M. M. Wilde, *Quantum Information Theory* (Cambridge University Press, Cambridge, 2017), pp. xi–xii.

[18] Note that this is merely an interpretation and does not rely on $|1\rangle$ physically being an excited state, as dictated by the underlying hardware.

[19] L. Grover and T. Rudolph, Creating superpositions that correspond to efficiently integrable probability distributions, [arXiv:quant-ph/0208112](https://arxiv.org/abs/quant-ph/0208112).

[20] V. Giovannetti, S. Lloyd, and L. Maccone, Architectures for a quantum random access memory, *Phys. Rev. A* **78**, 052310 (2008).

[21] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Random Access Memory, *Phys. Rev. Lett.* **100**, 160501 (2008).

[22] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Efficient quantum algorithms for simulating sparse hamiltonians, *Commun. Math. Phys.* **270**, 359 (2007).

[23] D. W. Berry and A. M. Childs, Black-box hamiltonian simulation and unitary implementation, *Quantum Inf. Comput.* **12**, 29 (2012).

[24] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian Dynamics with a Truncated Taylor Series, *Phys. Rev. Lett.* **114**, 090502 (2015).

[25] D. W. Berry, A. M. Childs, and R. Kothari, Hamiltonian simulation with nearly optimal dependence on all parameters, in

- 2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (IEEE, New York, 2015), pp. 792–809.
- [26] G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [27] T. Kim and B.-S. Choi, Efficient decomposition methods for controlled- r using a single ancillary qubit, *Sci. Rep.* **8**, 5445 (2018).
- [28] A less restrictive and still valid condition is $\|A\|_t = \mathcal{O}(\log^{k_1} \log^{k_2} N)$ for some exponents k_1 and k_2 .