Quantum computation of thermal averages in the presence of a sign problem

Giuseppe Clemente[,1,*] Marco Cardinali[,1,†] Claudio Bonati[,1,‡] Enrico Calore[,2,§] Leonardo Cosmai[,3,‖]
Massimo D'Elia,[1,¶] Alessandro Gabbana[,2,**] Davide Rossini[,1,††] Fabio Sebastiano Schifano[,2,‡‡]
Raffaele Tripiccione[,2,§§] and Davide Vadacchino[1,‖‖]

(QuBiPF Collaboration)

[1]*Dipartimento di Fisica dell'Università di Pisa and INFN—Sezione di Pisa,
Largo Pontecorvo 3, I-56127 Pisa, Italy*
[2]*Università degli Studi di Ferrara and INFN—Sezione di Ferrara, Via Saragat 1, I-44122 Ferrara, Italy*
[3]*INFN—Sezione di Bari, I-70126 Bari, Italy*

We illustrate the application of quantum computing techniques to the investigation of the thermodynamical properties of a simple system, made up of three quantum spins with frustrated pair interactions and affected by a hard sign problem when treated within classical computational schemes. We show how quantum algorithms completely solve the problem, and discuss how this can apply to more complex systems of physical interest, with emphasis on the possible systematics and on their control.

I. INTRODUCTION

The computation of thermal averages for a quantum mechanical system is a common challenge to many different fields of modern physics, in diverse areas such as condensed matter or fundamental interactions. In most cases, the issue is connected to the determination of the phase diagram of the system: this is the case, for instance, of quantum chromodynamics (QCD), the theory that describes strong interactions within the standard model of particle physics. In this context, the QCD phase diagram is related to fundamental questions regarding the properties of the Universe in its early stages, or the internal structure and composition of compact astrophysical objects, such as neutron stars [1,2].

*giuseppe.clemente@pi.infn.it
†marco.cardinali@pi.infn.it
‡claudio.bonati@unipi.it
§enrico.calore@fe.infn.it
‖leonardo.cosmai@ba.infn.it
¶massimo.delia@unipi.it
**gbblsn@unife.it
††davide.rossini@unipi.it
‡‡schifano@fe.infn.it
§§tripiccione@fe.infn.it
‖‖davide.vadacchino@pi.infn.it

In the absence of an exact solution, the strongly interacting nature of the system generally makes analytical tools, such as perturbation theory, unfeasible, and demands for a numerical approach. Monte Carlo (MC) algorithms represent a possible solution to the problem: in a nutshell, the computation of quantum thermal averages is brought to the form of a classical computation. Namely, after properly choosing the basis of the Hilbert space (the so-called computational basis), one performs a classical computation by means of classical MC methods, when possible. The path-integral formulation of quantum mechanics provides a suitable framework where this can be done systematically. The partition function of the quantum thermal system is rewritten in terms of a Euclidean path integral, which resembles the partition function of a classical statistical system in one more dimension (Euclidean time); Euclidean paths, or field configurations, are assigned a Boltzmann-like weight, corresponding to the exponential of minus the action. The path integral, in particular thermal averages, is then computed using standard tools, such as MC importance sampling.

Unfortunately, the prescription depicted above fails dramatically when the weight turns out to be nonpositive defined or, even worse, complex. This constitutes the infamous *sign problem*, which prevents progress in many fields, such as the accurate determination of the QCD phase diagram at finite baryon density [3–5], and can be shown to be a nondeterministic polynomial (NP) hard problem [6]. The problem is usually related to the path-integral formulation itself, or, in other words, to the choice of the

computational basis. Indeed, in the canonical formulation of the partition function, all the Hamiltonian eigenstates have a positive weight, as it stems from the hermiticity of the Hamiltonian itself, or of other possible conserved charges coupled to chemical potentials. However, determining the spectrum and the eigenstates of the Hamiltonian usually implies a similar or even greater level of computational difficulty.

The sign problem is often taken as one of the compelling motivations for the development of quantum computing (QC). Moving forward from the original proposal of Feynman [7], QC has been nowadays exploiting the possibility to operate algorithms based on quantum logic gates, aimed at either solving certain tasks more efficiently than classically, or in general reproducing the real-time evolution of a system on a digital, or analog, platform [8,9]. Cornerstone examples are the Lloyd's proposal for the quantum simulation of the Schrödinger equation [10], the Shor's algorithm for integer factorization [11], and the Grover's algorithm for searching from an unstructured database [12]. Our current investigation revolves around the question on how one can make use of a quantum computer in order to reproduce the quantum thermal averages and then explore the phase diagram, i.e., on how the sign problem can be solved in practice.

The reconstruction of quantum thermal averages via a quantum computer has been object of a rather intense theoretical activity in the past few years [13–25]. Here we focus on a specific implementation, proposed in Ref. [16], of the so-called quantum Metropolis sampling (QMS), where the main idea is to exploit quantum parallelism in order to sample directly in the eigenbasis of the Hamiltonian. The purpose of this study is to implement the algorithm for a prototype system where one finds a sign problem in the path-integral formulation, showing that this approach completely solves the problem and investigating the possible systematics of the method. In particular, we consider a system of three quantum spins coupled to each other via antiferromagnetic interactions [26,27], a model that has already been used as a QC testbed in several works, and in particular in quantum simulators using trapped ions [28,29] and nuclear magnetic resonance [30].

While it is almost trivial to diagonalize the system Hamiltonian and perform a classical sign problem-free computation, the way in which the quantum computer solves it is general and applies, at least in principle, to quantum systems of arbitrary complexity. Indeed, as we will better explain in the following, everything relies on the possibility of performing the quantum phase estimation (QPE) algorithm for the system at hand: it is at this point that the quantum advantage becomes manifest, since the algorithm is able to do that with no need for an explicit diagonalization of the Hamiltonian.

This paper is organized as follows. In Sec. II, we give a brief description of the model, showing how a discrete path integral (Trotter-Suzuki decomposition [31,32]) performed in the standard computational basis leads to a sign problem. Section III provides a short summary of the QMS algorithm adopted in this study. Section IV contains a summary of the results obtained with our model, while the possible systematics related to the algorithm are discussed in Sec. V. We conclude with Sec. VI, summarizing our results and discussing possible extensions to more complex systems.

The quantum algorithm illustrated in this study has been implemented on a quantum simulator (simulator for universal quantum algorithms) developed on purpose by one of the authors (G. C.), which is inspired to the well-known Qiskit simulator[1] [33].

II. DESCRIPTION OF THE MODEL

In this section, we introduce the system to be used as a playground for testing the QMS algorithm. It is made up of three quantum spin-1/2 variables, with frustrated pair interactions. Its Hamiltonian reads

$$H = J(\sigma_x \otimes \sigma_x \otimes \mathbb{1} + \sigma_x \otimes \mathbb{1} \otimes \sigma_x + \mathbb{1} \otimes \sigma_x \otimes \sigma_x), \quad (1)$$

where $\mathbb{1}$ stands for the 2×2 identity operator, σ_j (with $j = x, y, z$) denotes the usual spin-1/2 Pauli matrices, and J is a positive parameter (in order to make the couplings antiferromagnetic). Each of the $\mathbb{1}$ or σ_x operators acts on the spin corresponding to its position in the tensor product, and the full Hilbert space of the system is spanned by eight basis states. A possible choice, which makes the problem trivial, is to select as a basis the eight products of eigenstates (for each spin) of the σ_x operator: this is also a basis of eigenstates of the Hamiltonian, which has only two energy levels, the lower one with energy $E_0 = -J$ and degeneracy 6 (two spins aligned in either direction and opposite to the third spin), the higher one with energy $E_1 = 3J$ and degeneracy 2 (three spins aligned in either direction).

We are however interested in working with the standard computational basis, made up of eigenstates of the σ_z operators for all spins, which is the one where the sign problem appears. In the following, we will label, for simplicity, each of the eight states of this basis by its decimal correspondent, i.e., starting from $|\mathbf{0}\rangle \equiv |000\rangle$ to $|\mathbf{7}\rangle \equiv |111\rangle$. In this basis, the partition function of the system reads

$$Z = \mathrm{Tr}[\mathrm{e}^{-\beta H}] = \sum_{\alpha=0}^{7} \langle \boldsymbol{\alpha} | \mathrm{e}^{-\beta H} | \boldsymbol{\alpha} \rangle, \quad (2)$$

where β is the inverse temperature (we adopt natural units, $\hbar = k_B = 1$). The path-integral approach consists in rewriting the expression above by partitioning the exponential

---

[1]We used our own simulator to add some *ad hoc* features regarding the implementation of hybrid quantum-classical operations.

into the product of N smaller "evolution steps" (Trotter-Suzuki decomposition [31,32]) and inserting $N-1$ projectors over a complete set of states between each couple of evolution steps. Doing so, we arrive to the well-known expression

$$Z = \sum_{(\{\alpha_i\})} \prod_{i=1}^{N} \langle \alpha_{i+1} | e^{-\frac{\beta H}{N}} | \alpha_i \rangle, \qquad (3)$$

where $\alpha_{N+1} \equiv \alpha_1$ and $(\{\alpha_i\})$ is the space of all possible configurations, which we can interpret as periodic paths labeled by a discrete Euclidean time variable i. If the product of matrix elements turns out to be positive for each path, then the sum over all possible Euclidean paths appearing in Eq. (3) can be sampled by a classical MC algorithm. However, it is easy to show that this is not the case for this system, i.e., it is always possible to find paths for which the product of matrix elements turns out to be negative.

In order to show this fact, let us rewrite the elementary evolution step in the following form:

$$e^{-\frac{\beta H}{N}} = \frac{1}{4} \left[(e^{-3\varepsilon} + 3e^{\varepsilon}) \mathbb{1} + (e^{-3\varepsilon} - e^{\varepsilon}) \frac{H}{J} \right], \qquad (4)$$

where $\varepsilon \equiv \beta J / N$ [details about the derivation of Eq. (4) are reported in Appendix]. Since the Hamiltonian is composed of a symmetric combination of pair-products of σ_x operators, which flip the σ_z-component of the corresponding spin, it should be clear that the only nonzero nondiagonal matrix elements of $\exp(-\beta H/N)$ are those between states differing by exactly two spin flips. This fact enables to split the computational basis into two disjoint sets which are as follows:

$$\text{Set 1}: \{|0\rangle, |3\rangle, |5\rangle, |6\rangle\}, \qquad (5a)$$

$$\text{Set 2}: \{|1\rangle, |2\rangle, |4\rangle, |7\rangle\}. \qquad (5b)$$

Nondiagonal elements within the same set are all equal and negative, and in particular their common value is $-e^{\varepsilon}(1 - e^{-4\varepsilon})/4$, while nondiagonal elements across the two sets are exactly zero. In contrast, diagonal elements are all equal to $e^{\varepsilon}(3 + e^{-4\varepsilon})/4$.

To summarize, the discrete path-integral sum in Eq. (3) simplifies by decoupling into two distinct and simpler sums, where paths run either within *Set 1* or within *Set 2*. However, within each sum, paths presenting an odd number[2] of flips between different states carry a negative weight. The contribution of these paths cannot be rewritten in terms of a positive Boltzmann weight, i.e., in the form

----

[2]It is always possible to find such paths, since one can also insert an arbitrary number of diagonal matrix elements along the Euclidean time.

$\sum \exp(-S[\alpha_i])$, where $S[\alpha_i]$ is a real Euclidean action. This invalidates the possibility of using classical methods which approximate the sum by means of an MC sampling of its terms.

In the QMS algorithm described in the next section, the focus is brought again to the basis of eigenstates of the Hamiltonian, which are sampled according to their thermal weight. However, this is not done by exploiting the exact solution of the problem (that, in this case, would be trivial), but in a way which can be easily generalized to more complex cases. As we shall see below, the key ingredient to achieve this task will be the QPE algorithm [34,35].

III. THE QUANTUM METROPOLIS SAMPLING ALGORITHM

This section is a brief, yet self-contained, overview of the main steps of the QMS algorithm, to be operated on a quantum computer. Further details can be found in the original publication, Ref. [16].

The key ingredient of QC, which potentially enables an exponential speed-up over any known classical algorithm analogue, is quantum parallelism. Namely, the possibility, for a set of N two-level quantum systems (the so-called qubits, quantum analogues of classical bits), to be in a superposition of states, thus simultaneously acting on an exponential number 2^N of classical configurations. Such an ensemble of qubits can be suitably manipulated through a sequence of invertible single- and two-qubit controlled-NOT (CNOT) quantum logic gates (all mathematically representable as unitary operators); this constitutes a universal set of quantum gates [8,9].

The QMS algorithm is a procedure which enables to perform a quantum MC sampling of thermal averages, using a sequence of invertible quantum logic gates on a register of qubits and classical measures on appropriate subportions of it. More specifically, QMS makes use of QPE in order to record the eigenvalues of a given quantum Hamiltonian without explicitly knowing the detailed structure of its eigenvectors. Before entering the overview of QMS, we discuss the working principle of QPE, lying at the heart of it.

A. Quantum phase estimation

The QPE algorithm applies to the following general problem. Suppose a given unitary operator U has an eigenvector $|u\rangle$ with eigenvalue $e^{i\theta}$ ($0 \leq \theta \leq 2\pi$). Also, assume that one is able to prepare the state $|u\rangle$ and to perform controlled-U^{2^j} gates ($j \in \mathbb{N}$). The latter operation needs a set of two quantum registers, the first one (1) containing the n qubits necessary to store $|u\rangle$; the second one (2) containing an ancillary qubit, such that

$$\frac{|0\rangle_2 + |1\rangle_2}{\sqrt{2}} \otimes |u\rangle_1 \xrightarrow{\text{c-}U^{2^j}} \frac{|0\rangle_2 + e^{i2^j\theta}|1\rangle_2}{\sqrt{2}} \otimes |u\rangle_1. \qquad (6)$$

The purpose of QPE is to obtain the best l-bit estimate of θ, for a given integer value of l [34,35].

To achieve this goal, one discretizes the possible values of θ choosing a suitable encoding in binary notation with r bits [up to $O(2^{-(r+1)})$ accuracy]. A sequence of the above controlled-U^{2^j} gates needs to be applied (varying $j = 0, 1, 2, \ldots$) on r different ancillary qubits, here referred to as the control register. Since the quantum register that stores $|u\rangle$ is prepared in an eigenstate of the operators $U, U^2, U^4, \cdots$, the state of this register never changes, while the phase factors $e^{i\theta}, e^{2i\theta}, e^{4i\theta}, \ldots$ are propagated backward in the control register. Finally, a quantum Fourier transform on the control register can be shown to provide an approximate binary encoding of θ, where each digit is represented by the state of a specific qubit state of the register. In particular, the algorithm outputs the best l-bit estimate of θ with probability $> 1 - \epsilon$, if the control register contains $r = l + O(\log(1/\epsilon))$ qubits and the obtained result is rounded off to its most significant l bits [34]. For further details on the QPE algorithm, see, e.g., Ref. [9].

In the specific implementation of QMS, the QPE procedure is applied to the unitary operator $U = e^{-iHt}$, which, for a large class of physically relevant Hamiltonians, can be efficiently simulated on a quantum computer [10], after an appropriate discretization of the real-time evolution (Trotterization) [31,32]. Since an eigenvector $|\psi_\alpha\rangle$ of H with eigenvalue E_α is also an eigenvector of $U(t)$ with eigenvalue $e^{-iE_\alpha t}$, it is thus possible to encode, in the control register, the eigenvalues of H without explicitly knowing the detailed structure of its eigenvectors [35].

B. Implementation of quantum Metropolis sampling

We consider a generic quantum system that can be represented (or approximated) by a finite number n of qubits and is described by the Hamiltonian H. In principle, the eigenvalues of H could take any real value, but for the sake of simplicity, let us assume that they can be exactly represented by a certain number r of qubits. This is always possible for commensurate energy spacings, by discretizing their possible values and using the binary representation of the integers 0 through $2^r - 1$. As emerging from the discussion above in Sec. III A, and will be clarified later in Sec. V B, for systems with energy levels that are not multiple of a certain energy unit this introduces a systematic error, which however can be made arbitrarily small by increasing r (at the cost of increasing the complexity of the algorithm).

The idea of the QMS algorithm is to adapt to the QC framework the classical Markov Chain MC method [36], by generating a sequence of eigenstates of H distributed according to the Boltzmann statistical weight. While some steps of the original classical algorithm can be extended to the quantum case straightforwardly, the accept-reject step

requires more care since, due to the no-cloning theorem [37,38], it is not possible to have a backup copy of the original configuration to be used if the proposed update is rejected.

To overcome this difficulty, the QMS algorithm uses four quantum registers: a register representing the system wave function $|\psi\rangle$ with n qubits, two r-qubits registers holding information about the energy before ($|E^{\text{old}}\rangle$) and after ($|E^{\text{new}}\rangle$) a proposed update, and a single-qubit register with information about the acceptance ($|\text{acc}\rangle$). All these registers can be arranged as

$$|\text{acc}\rangle_4 \otimes |E^{\text{new}}\rangle_3 \otimes |E^{\text{old}}\rangle_2 \otimes |\psi\rangle_1, \qquad (7)$$

in the following, when omitted, the subscripts corresponding to the various registers are intended as in (7). The QMS algorithm then proceeds as follows:

Step 0.—At the beginning, the input state (7) has to be initialized to zero for all registers, except the one holding the system state, which must be set to an eigenstate $|\psi_k\rangle$ of the Hamiltonian: $|0, 0, 0, \psi_k\rangle$. Any eigenstate will work, so the simulation can be initialized, e.g., by making the wave function collapse by means of an energy measure.

Step 1.—A QPE procedure $\Phi^{(\text{old})}$ is performed between register 1 and register 2,

$$|0\rangle_2 \otimes |\psi_k\rangle_1 \xrightarrow{\Phi^{(\text{old})}} |E_k\rangle_2 \otimes |\psi_k\rangle_1, \qquad (8)$$

E_k being the energy of the eigenstate $|\psi_k\rangle$. Thus, the global output state after this step is $|0, 0, E_k, \psi_k\rangle$.

Step 2.—The choice of the trial update is implemented by randomly selecting a unitary operator C from a certain set $\mathcal{C} \subseteq \mathcal{U}(2^n)$ and applying it to the first register,[3]

$$|E_k\rangle_2 \otimes |\psi_k\rangle_1 \xrightarrow{C} \sum_p x_{k,p}^{(C)} |E_k\rangle_2 \otimes |\psi_p\rangle_1, \qquad (9)$$

where $x_{k,p}^{(C)}$ are the matrix elements of C. This operation is followed by a QPE $\Phi^{(\text{new})}$ between register 1 and register 3,

$$|0\rangle_3 \otimes |\psi_p\rangle_1 \xrightarrow{\Phi^{(\text{new})}} |E_p\rangle_3 \otimes |\psi_p\rangle_1. \qquad (10)$$

In summary, this brings the output state of step 1 into a new (formal) superposition of eigenstates,

$$|0, 0, E_k, \psi_k\rangle \xrightarrow{\text{Step 2}} \sum_p x_{k,p}^{(C)} |0, E_p, E_k, \psi_p\rangle. \qquad (11)$$

---

[3]There is a lot of arbitrariness in the choice of the set $\mathcal{C}$, but with two requirements: for each element of $\mathcal{C}$, its inverse must be in $\mathcal{C}$ as well, and the set must be large enough to ensure ergodicity in the sampling of energy eigenstates.

Step 3.—This step is the quantum analogue of the accept-reject procedure: a one-qubit operator $W(E_p, E_k)$ is applied to the acceptance register 4, conditioned (by controlled gates) on both the old and new energy registers (2 and 3), such that

$$|0, E_p, E_k, \psi_p\rangle$$
$$\xrightarrow{W} \left(\sqrt{f(p,k)}|1\rangle + \sqrt{1-f(p,k)}|0\rangle \right) \otimes |E_p, E_k, \psi_p\rangle,$$
$$(12)$$

where

$$f(p,k) = \min\left(1, e^{-\beta(E_p - E_k)}\right). \qquad (13)$$

Now, a classical measure is done on the acceptance register, which can only take the value 1 or 0. The outcome of this measure is 1 with probability

$$\sum_p f(p,k)|x_{k,p}^{(C)}|^2, \qquad (14)$$

and in this case we proceed by performing a classical measure of energy on register 3, making the system qubits collapse onto a given energy eigenstate $|1, E_p, E_k, \psi_p\rangle$. From this state, one resets[4] all the nonsystem qubits (registers 2, 3, and 4) to zero and then starts again the algorithm from step 1. Things are slightly more complex if the readout of the acceptance register is 0: in this case, the trial state has to be rejected, and therefore the update step needs to be reverted, using the procedure described in step 4.

Step 4.—First, apply the inverse of the operator

$$U \equiv W \cdot \Phi^{(\text{new})} \cdot C, \qquad (15)$$

which represents the sequence of unitary operators applied in steps 2 and 3. Notice that register 3 will be set to zero by $U^\dagger$, while register 2 is kept untouched from step 1, giving us a reference value for the initial energy. A QPE $\Phi^{(\text{new})}$ followed by a classical measure on register 3 can return two possible results: if $E^{\text{new}} = E_k$, one has succeeded in reverting back the system state into an energy eigenstate with energy[5] E_k; it is thus possible to reset all nonsystem qubits and start again from step 1. If instead $E^{\text{new}} \neq E_k$, one has to apply $[\Phi^{(\text{new})}]^{-1}$, followed by U and by a measure in

---

[4]Resetting a qubit can be implemented, for example, by taking a classical measure on that qubit, followed by a σ_x-gate if the measure returns 1 (i.e., a CNOT controlled by a classical bit).

[5]The system state does not have to be equal to $|\psi_k\rangle$; indeed, in case of degeneracies, invertibility and ergodicity of the set of updates $\mathcal{C}$ anyhow guarantee the exactness of the algorithm. In particular, going to a different state can be viewed as reverting to the original state followed by a microcanonical step.

the acceptance register (in order to reach a configuration analogous to the one at the end of step 3), and repeat step 4 until E^{new} is equal to E_k. From a practical point of view, it is reasonable to set a maximum number of iterations for this procedure, after which the whole update is aborted and one has to restart from step 0.

The QMC algorithm is thus completely determined once the set $\mathcal{C}$ of unitary operators needed in the step 3 is known. This set has to satisfy some requirements for the QMC procedure to correctly sample the Boltzmann weight, which however do not fix completely its form: it is indeed sufficient that if $C \in \mathcal{C}$ also $C^\dagger \in \mathcal{C}$, moreover $\mathcal{C}$ has to be large enough to enable any energy eigenstate to be reached in a finite number of steps. For complex systems, the numerical effectiveness of the QMC algorithm could however significantly depend on the choice of the set $\mathcal{C}$, and it would be of great help to dispose of a systematic procedure to select a convenient set of unitary transformations. For the few-level systems investigated in this work, this is however not a big problem and thus, in the following, we will not specifically investigate the issue of optimizing the choice of the set $\mathcal{C}$.

We finally emphasize that the possibility of using quantum gates controlled by classical measure is a key ingredient for the implementation of the QMS algorithm. This feature is used at the end of step 3 when, depending on the result of the classical measure on the acceptance register, one accepts (go to step 1) or rejects (go to step 4) the proposed configuration. Moreover, in case of reject, another classical control is needed to check whether E^{new} is equal to E^{old} or not.

C. Measuring observables

Using the QMS algorithm, we can generate a sequence of energy eigenstates whose energy is distributed according to the Boltzmann weight and, in particular, we can easily compute thermal averages of operators commuting with the Hamiltonian H. One is also interested, however, in computing thermal averages of a generic observable O, which does not commute with H. The measure of such an observable on the system register would however spoil the stochastic exactness of the QMS algorithm, since the state register would collapse away from the given energy eigenstate.

This problem is obviously absent in the classical algorithm; a possible way out is to measure the observable O every n_{th} steps of the QMS algorithm, with n_{th} large enough to let the system rethermalize between consecutive measures. A different strategy could be to start from step 0 of the QMS algorithm, perform n_{th} updates to reach thermalization, measure the observable O, and finally abort the simulation, repeating this procedure several times. In this way, we have however less or no variability in the initial configuration at step 0 (depending on the initialization procedure adopted), and by performing n_{th} rethermalization

updates after every measure we expect to have a more efficient sampling of the energies.

Note that, in both cases (i.e., using rethermalization or different independent simulations), we have to fix the number n_{th}, and this introduces a systematic error in the algorithm. By selecting n_{th} large enough, this systematic can be made negligible with respect to the statistical errors that are present in every MC simulation; however, this means that we cannot improve arbitrarily the accuracy by simply increasing the statistics. A detailed discussion of this point will be presented in Sec. V. Finally, let us explicitly remark that, if we are interested in computing the quantum thermal average of several operators, rethermalization updates are required after the measure of each of these observables, the only exception being the case of observables commuting with each other, which can be measured (in any order) with no need for in-between rethermalization.

IV. NUMERICAL RESULTS

In this section, we present the results obtained by applying the QMS algorithm to the frustrated triangle system described by the Hamiltonian in Eq. (1).

States of the frustrated triangle can be represented using $n = 3$ qubits, while the energy of the system can be exactly encoded by using only one qubit ($r = 1$). Indeed, since the energy takes only two values ($-J$ and $3J$), by shifting and rescaling H, we can define the new Hamiltonian $\tilde{H} = (\mathbb{1} + H/J)/4$, whose spectrum is only made of eigenvalues equal to 0 or 1. The total number of required qubits is therefore $n + 2r + 1 = 6$.

To proceed with the QMS algorithm, we need two basic ingredients: to be able to perform a QPE and to determine the set $\mathcal{C}$ of unitary operators to be used for generating trial states. As previously discussed, the set $\mathcal{C}$ has to ensure ergodicity (in this context, by ergodicity we do not mean the possibility to actually explore the whole Hilbert space, but only the whole set of energy eigenstates) and if $U \in \mathcal{C}$, also $U^{\dagger}$ has to belong to the same set (reversibility).

The naive choice of single spin flips $\mathbb{1} \otimes \mathbb{1} \otimes \sigma_x$, $\mathbb{1} \otimes \sigma_x \otimes \mathbb{1}$, and $\sigma_x \otimes \mathbb{1} \otimes \mathbb{1}$ does not ensure ergodicity, as evident from the observation that all these operators commute with H. An ergodic choice can be obtained instead by replacing the spin flip operator σ_x in the previous expressions by the Hadamard gate $\frac{1}{\sqrt{2}}(\sigma_z + \sigma_x)$ [8]. Since these transformations are involutions, reversibility is also automatically satisfied. With this choice, we obtain an average acceptance probability in the Metropolis step that goes from 99% for $\beta = 0.1$ to 91% for $\beta = 1.0$.

For what concerns the QPE, this actually boils down to a real-time evolution (see Sec. III A) and, for the frustrated triangle, time evolution is particularly simple: since all terms in H commute pairwise and their square is proportional to the identity, we can easily construct the time evolution operator as a product of terms like

$$e^{-itJ\mathbb{1}\otimes\sigma_x\otimes\sigma_x} = (\mathbb{1} \otimes \mathbb{1} \otimes \mathbb{1})\cos(Jt) - i(\mathbb{1} \otimes \sigma_x \otimes \sigma_x)\sin(Jt). \quad (16)$$

We will come back later to the systematics associated to the QPE when such an analytical expression is not available.

Now we have all the ingredients required for the application of the QMS algorithm and, in order to test its effectiveness, we start by studying the dependence on the temperature of the internal energy $\langle\tilde{H}\rangle$, comparing numerical data with exact analytical results. Such a comparison is performed in Fig. 1, and we can see that there is excellent agreement between the two determinations.

We now move to the computation of thermal averages of operators not commuting with the Hamiltonian. We have tested several Hermitian operators, but here we only report the results obtained for

$$A = \sigma_x \otimes \sigma_x \otimes (\mathbb{1} + \sigma_y). \quad (17)$$

The outcomes obtained in all other cases are equivalent. After some preliminary test (see Sec. V A for a thorough analysis), we decided to use 200 rethermalization steps after each measure of A, and in Fig. 1, we show the results obtained with a statistics of about 10^7 measures. Even in this case, the agreement between the numerical estimates and the exact results is remarkable; this is a direct proof of the fact that the algorithm properly samples the configurations and effectively solves the sign problem of this model.

After having checked that the QMS algorithm correctly samples the states of the system, we can investigate some aspects related to the efficiency of the method. In particular, a possible weak point of the QMS algorithm is the iterative procedure required in step 4 (see Sec. III B) to project back a rejected update: if a very large number of iterations are required, the algorithm obviously loses its effectiveness. However, we found that, at least for the frustrated triangle studied in this work, the number of reverting steps required for each iteration is relatively small, and its distribution can be well approximated by an exponential, as shown in the histogram of Fig. 2 for some values of β. The distribution of the number of reverting steps gets broader for small β values, signaling that the algorithm is less efficient in this regime; on the other hand, this is mitigated by the fact that at low β the acceptance is higher, so that reverting steps occur less frequently. For this system and setup, we observe $\sim$10 reverting steps on average at $\beta = 0.1$ and a lower number at larger β. In general, a possible way of reducing the impact of the reverting steps on the simulation time would be to look for a set $\mathcal{C}$ of moves with larger acceptance probability, so to reduce the occurrence of the iterative step. We note however that the dependence of the average number of reversing steps on β is quite mild, being to a good accuracy proportional to $1/\beta$, to be compared with the exponential behavior typically encountered in systems characterized by sign problems.
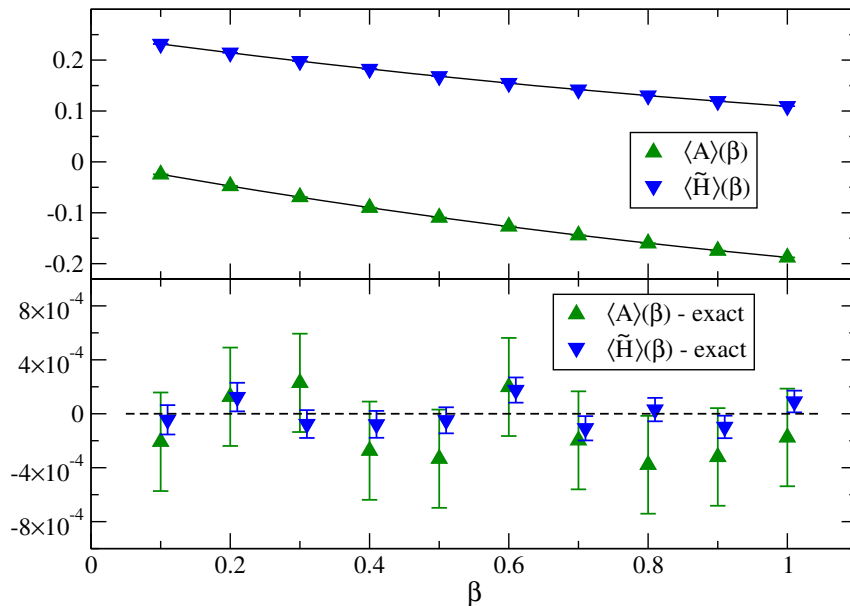
FIG. 1.   Expectation values of the energy $\tilde{H}$ and the observable A, defined in Eq. (17), for several values of β. In the top panel, solid lines represent the exact analytical values of the observables; the bottom panel shows residuals (data are slightly shifted horizontally, to improve readability) in order to better appreciate the consistency between analytical and numerical results.

One last comment must be done on the number of quantum gates required to perform a single Metropolis step. In order to perform the QPE and apply the operator U, i.e., before measuring the acceptance register, we need ~ 100 gates. If we accept the new configuration, we need additional ~ 200 gates to go back to step 1; otherwise, in case of rejection, we need ~ 100 gates in order to perform a single reverse attempt. Beyond these gates, six gates controlled by classical measures are needed for each iteration of the algorithm and one for each reverse attempt. We see that, even for a very simple system like the frustrated triangle, the number of quantum gates used is large; therefore, in order to run the algorithm on a real quantum computer, a very high fidelity on the quantum gates is required.
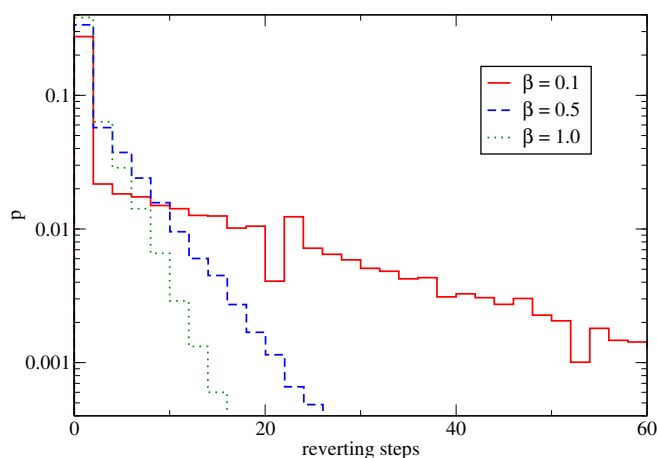

FIG. 2.   Histogram of the number of times step 4 of the QMS algorithm had to be repeated for $\beta = 0.1$, 0.5, and 1.0.

V. SYSTEMATICS OF THE QMS ALGORITHM

The QMS algorithm generally presents a certain number of systematic errors. In the previous section, we already discussed the rethermalization updates to be performed after each measure of an observable not commuting with the Hamiltonian. For the frustrated triangle, this was the only source of systematic error in the QMS algorithm. Here we discuss the systematics that have to be taken into account, beyond rethermalization, when using QMS to simulate a generic quantum system.

Digitalization of the energy.—In the QMS update, energy has to be stored in two quantum registers, each composed of r qubits, and if the spectrum of the Hamiltonian cannot be exactly mapped to an $r-$bit binary number, one introduces a truncation error in the energies and thus a systematic in the sampling. Even for toy models, this is clearly unavoidable whenever energy eigenvalues are incommensurable with each other, like for a three-state system with energy levels $0, 1/\sqrt{2}$, and 1.

Digitalization of the state.—This problem is present for systems with continuous degrees of freedom (like gauge theories associated to continuous Lie groups). In such cases, we need to choose a proper digitalization of the state, i.e., a way of approximating the state by using only n qubits. For more details, see Refs. [39–42].

QPE implementation.—As briefly recalled in Sec. III A, the implementation of the QPE algorithm requires a real-time unitary evolution on the system register using e^{-iHt}. In practice, such an evolution cannot be performed exactly, but has to be approximated by a Trotterization, thus introducing a systematic error related to the finite

time-step (see, e.g., Refs. [43–48] for recent discussions about this systematic effect in similar contexts).

The effect of all these systematics can be reduced at the expense of increasing the computational complexity of the algorithm: by increasing the number of used qubits, the digitalization of both energies and states can be made arbitrarily accurate and similarly the Trotter time step can be reduced at will, making the QPE more and more accurate. While this is true in theory, in practice one is interested in understanding how these systematics scale with the computational complexity, in order to understand whether the algorithm is useful or not for the problem at hand.

The last two systematics of the list (i.e., digitalization of the state and QPE implementation) are not specific of the QMS, but are common to all QC algorithms. For this reason, they have already been discussed in the literature and we will not further investigate them, thus concentrating on the rethermalization and the digitalization of energy.

A. Rethermalization

We have performed several simulations with different numbers of rethermalization steps to quantify the number of needed rethermalization updates. As an example, the results obtained for the average values $\langle \tilde{H} \rangle$ and $\langle A \rangle$ [defined in Eq. (17)] are shown in Fig. 3 for $\beta = 1$. The observed behavior is analogous to that seen in the thermalization of classical MC simulations: average values depend on the number of rethermalization steps n_{th} adopted, but they converge to an asymptotic value when n_{th} increases.

This asymptotic value is the correct thermal average, which can however be reached at different values of n_{th} for different observables, as clearly seen in Fig. 3; this is just a consequence of the fact that different observables have different thermalization times. In order not to introduce systematics in the final estimates, n_{th} has to be chosen large enough that all the observables have reached their plateau values with an accuracy that is larger than the target precision of the simulation. We explicitly note that, in order not to introduce biases, the maximum number of iterations in step 4 of QMS have to be large enough to guarantee that a complete thermalization cycle can be performed without aborting the algorithm.

B. Digitalization of the energy

In order to assess how much the QMS algorithm relies on an accurate estimation of the energies, we considered a simple system made up of two qubits representing an Hamiltonian H with the following four energy levels: $\{0, \frac{1}{2}, \frac{1}{\sqrt{2}}, \frac{3}{4}\}$. With this setup, energy levels can never be exactly digitalized, and we can also investigate the effect of having two nearby levels ($1/\sqrt{2}$ and $3/4$) that cannot be resolved if the number of used qubits is too small.

As a preliminary test, we isolated the QPE step: in Fig. 4, we show the results of a single application of the QPE procedure Φ to the eigenstate of H associated to the eigenvalue $\lambda_1 = \frac{1}{\sqrt{2}}$, using a variable number r of qubits for the energy register. Since λ_1 cannot be represented by a fraction of powers of 2, a measure in the energy register does not always collapse to the same state, but has a certain
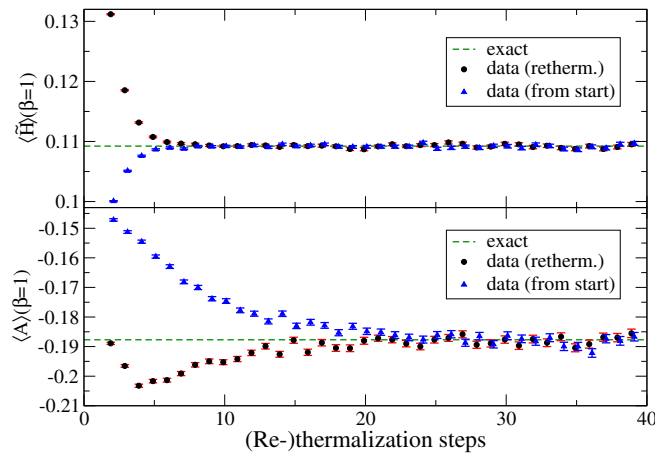

FIG. 3. Effects of different numbers of (re-)thermalization steps at $\beta = 1$ on the expectation values of the energy and of the non-H-commuting observable A defined in Eq. (17). The points with a circle refer to measures taken during the Metropolis chain after each rethermalization sequence, while the ones with a triangle refer to measures taken at the end of a single Metropolis chain after a certain number of thermalization steps, starting always from the initial eigenstate $\frac{1}{\sqrt{2}}(|\mathbf{0}\rangle - |\mathbf{3}\rangle)$.
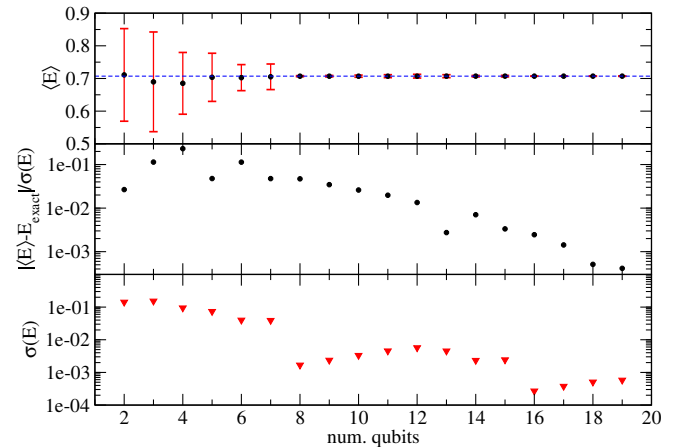

FIG. 4. Accuracy of the QPE procedure for the exact eigenvalue $\lambda_1 = \frac{1}{\sqrt{2}}$ using different numbers of qubits for the energy register. Here $\langle E \rangle$ and $\sigma(E)$ denote the average value and the standard deviation of the energy distribution, as obtained from QPE. In the upper panel, QPE estimates for $\langle E \rangle$ are compared with the exact value (dashed line); the middle panel shows the ratio between the absolute value of the discrepancy and the standard deviation; the bottom panel shows the scaling of $\sigma(E)$ with the number of qubits.
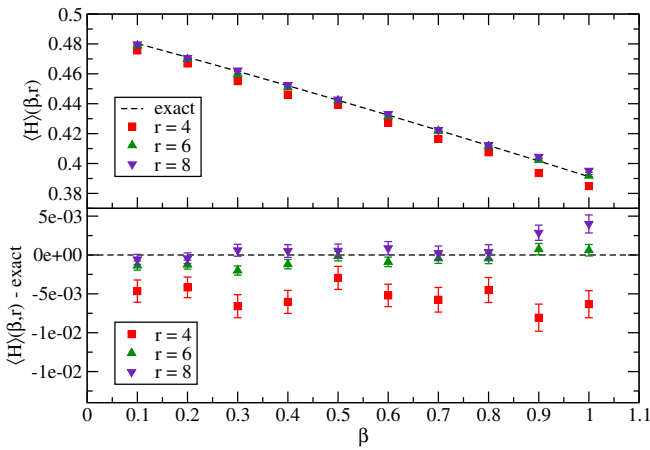
FIG. 5. Energy sampling for a system with energy levels $\{0, \frac{1}{2}, \frac{1}{\sqrt{2}}, \frac{3}{4}\}$, using a different number of qubits for the energy registers ($r = 4$, 6, and 8), and comparison with the exact expectation value $\langle H(\beta) \rangle$ for several values of β.

probability to collapse to a state with neighboring energy, so that each measure is affected by a sampling error, which however vanishes exponentially in r.

More interesting are the effects of errors in the estimates occurring during the full QMS algorithm, as can be observed from the average of the energies measured, shown in Fig. 5 for three different digitalizations of the energy ($r = 4$, 6, and 8). Using four qubits seems not to give a sufficiently accurate estimate along the explored range in β, while 6 and 8 qubits show reasonable accuracy, even if the convergence behavior is not always clear, as apparent, for example, at $\beta = 1$. Nevertheless, comparing the energy probability distribution of the samples with the exact one (Fig. 6), we can observe a clearer convergence to the exact weights at increasing number of qubits r. Notice that a
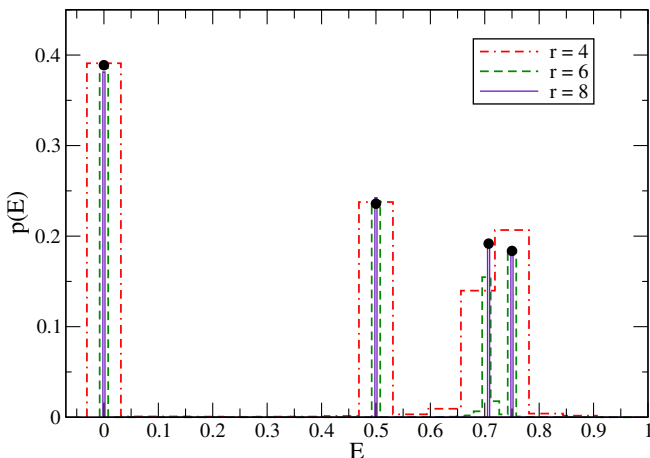

FIG. 6. Energy sampling at $\beta = 1.0$ with energy levels $\{0, \frac{1}{2}, \frac{1}{\sqrt{2}}, \frac{3}{4}\}$, using different numbers of qubits for the energy registers ($r = 4$, 6, and 8), and comparison with the exact probability distribution $p(E)$. For each choice of r, around $O(10^4)$ measures have been performed.

proper discrimination between neighboring eigenvalues in the spectra (as $\frac{1}{\sqrt{2}}$ and $\frac{3}{4}$) can occur only if the number of qubits is sufficient.

VI. CONCLUSIONS

In this paper we have described, implemented, and tested a quantum algorithm able to measure thermal averages for quantum mechanical systems. Our results are for a very simple system of frustrated spins, which exhibits the infamous sign problem when rewritten in terms of Euclidean path integrals, thus making the use of classical Monte Carlo simulations impossible. In this sense, our toy model shares the same problems of more complex and physically meaningful problems, such as the study of the phase structure of Lattice QCD at finite temperature and density. It therefore conceptually paves the way to quantum simulations of those systems, as soon as appropriately powerful quantum computers become available in the future.

The algorithm chosen for our study is the quantum Metropolis algorithm introduced in Ref. [16], which has been briefly reviewed in Sec. III. Even if inspired by the standard Metropolis algorithm, it cannot be considered a Markov chain in the classical sense, since the measurement of observables interferes with the Markov process itself, thus destroying equilibration and leading to the need for rethermalization after each measurement is taken (unless the observable commutes with the Hamiltonian). Of course, this is not a peculiar feature of the algorithm, but rather something which one expects in general, because of the quantum nature of the system and of the algorithm.

We have verified that our quantum algorithm reproduces, within statistical errors, exact results for the thermal averages of the energy and of an additional operator that does not commute with the Hamiltonian. We have quantified the size of the statistical and sampling errors and their behavior with respect to the number of simulation steps. We have also discussed and measured the systematic errors due to several different effects, such as the accuracy of the phase estimation process, the sampling accuracy as a function of the qubit register size, as well as some conceptually unavoidable performance losses associated to the reverting steps of the algorithm.

All these systematic errors, together with the computing depth of the complete algorithm, will ultimately play a key role to define the actual viability of this quantum algorithm as one moves from simple toy models to the study of real-life complex physical systems. This perspective suggests several avenues for further studies, ranging from the implementation of the present model onto actual QC machines, to the implementation of the algorithm for more complex (possibly gauge invariant) models, to the analysis of the optimal balance point between computational accuracy and efficiency versus register size and algorithm complexity: work is in progress along these directions.

ACKNOWLEDGMENTS

APPENDIX: DERIVATION OF EQ. (4)

Here we present a derivation of Eq (4), which was used in Sec. II to better understand the structure of the path-integral representation of the frustrated triangle. We start from the Taylor expansion of the Euclidean evolution operator,

$$e^{-\frac{\beta H}{N}} = \sum_{n=0}^{\infty} \frac{1}{n!} \left(-\frac{\beta}{N}\right)^n H^n. \qquad (A1)$$

It is then easy to check that the Hamiltonian (1) satisfies the property $H^2 = (3J^2\mathbb{1} + 2JH)$; therefore we have, for any integer n, a relation of the form

$$H^n = \alpha_n \mathbb{1} + \beta_n H. \qquad (A2)$$

To fix the values of α_n and β_n, we can diagonalize both members of this equation and remember that the eigenvaules of H are just $-J$ and $3J$. In this way, we immediately get

$$H^n = \frac{J^n}{4} \left\{ [3(-1)^n + 3^n]\mathbb{1} + [3^n + (-1)^{n+1}]\frac{H}{J} \right\} \qquad (A3)$$

and, plugging this result into Eq. (A1), we finally obtain

$$e^{-\frac{\beta H}{N}} = \sum_{n=0}^{\infty} \frac{1}{n!} \left(-\frac{\beta}{N}\right)^n H^n$$
$$= \frac{1}{4} \left[(e^{-3\frac{\beta J}{N}} + 3e^{\frac{\beta J}{N}})\mathbb{1} + (e^{-3\frac{\beta J}{N}} - e^{\frac{\beta J}{N}})\frac{H}{J} \right]. \qquad (A4)$$

[1] S. L. Shapiro and S. A. Teukolsky, *Black Holes, White Dwarfs, and Neutron Stars* (Wiley-Interscience, New York, 1983).

[2] K. Rajagopal and F. Wilczek, *The Condensed Matter Physics of QCD*, edited by M. Shifman, At the Frontier of Particle Physics Vol. 3 (World Scientific, Singapore, 2001), pp. 2061–2151.

[3] O. Philipsen, arXiv:1009.4089.

[4] G. Aarts, J. Phys. Conf. Ser. **706**, 022004 (2016).

[5] H. T. Ding, F. Karsch, and S. Mukherjee, Int. J. Mod. Phys. E **24**, 1530007 (2015).

[6] M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).

[7] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).

[8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 2nd ed. (Cambridge University Press, Cambridge, United Kingdom, 2010).

[9] G. Benenti, G. Casati, D. Rossini, and G. Strini, *Principles of Quantum Computation and Information: A Comprehensive Textbook* (World Scientific, Singapore, 2018).

[10] S. Lloyd, Science **273**, 1073 (1996).

[11] P. W. Shor, SIAM J. Comput. **26**, 1484 (1997).

[12] L. K. Grover, Phys. Rev. Lett. **79**, 325 (1997).

[13] B. M. Terhal and D. P. DiVincenzo, Phys. Rev. A **61**, 022301 (2000).

[14] D. Poulin and P. Wocjan, Phys. Rev. Lett. **103**, 220502 (2009).

[15] E. Bilgin and S. Boixo, Phys. Rev. Lett. **105**, 170405 (2010).

[16] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, Nature (London) **471**, 87 (2011).

[17] M.-H. Yung and A. Aspuru-Guzik, Proc. Natl. Acad. Sci. U.S.A. **109**, 754 (2012).

[18] A. Riera, C. Gogolin, and J. Eisert, Phys. Rev. Lett. **108**, 080402 (2012).

[19] Y. Ge, A. Molnár, and J. I. Cirac, Phys. Rev. Lett. **116**, 080503 (2016).

[20] M. J. Kastoryano and F. G. S. L. Brandão, Commun. Math. Phys. **344**, 915 (2016).

[21] A. Shabani and H. Neven, Phys. Rev. A **94**, 052301 (2016).

[22] A. N. Chowdhury and R. D. Somma, Quantum Inf. Comput. **17**, 41 (2017).

[23] J. Cohn, K. S. Najafi, F. Yang, B. Jones, and J. K. Freericks, arXiv:1812.03607.

[24] E. R. Anschuetz and Y. Cao, arXiv:1903.01359.

[25] J. E. Moussa, arXiv:1903.01451.

[26] G. H. Wannier, Phys. Rev. **79**, 357 (1950); **7**, 5017(E) (1973).

[27] L. Balents, Nature (London) **464**, 199 (2010).

[28] K. Kim, M.-S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe, Nature (London) **465**, 590 (2010).

[29] E. E. Edwards, S. Korenblit, K. Kim, R. Islam, M.-S. Chang, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe, Phys. Rev. B **82**, 060412(R) (2010).

[30] J. Zhang, M.-H. Yung, R. Laflamme, A. Aspuru-Guzik, and J. Baugh, Nat. Commun. **3**, 880 (2012).

[31] H. F. Trotter, Proc. Am. Math. Soc. **10**, 545 (1959).

[32] M. Suzuki, Commun. Math. Phys. **51**, 183 (1976).

[33] https://qiskit.org.

[34] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Proc. R. Soc. A **454**, 339 (1998).

[35] D. S. Abrams and S. Lloyd, Phys. Rev. Lett. **83**, 5162 (1999).

[36] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).

[37] W. K. Wootters and W. H. Zurek, Nature (London) **299**, 802 (1982).

[38] D. Dieks, Phys. Lett. A **92**, 271 (1982).

[39] N. Klco and M. J. Savage, Phys. Rev. A **99**, 052335 (2019).

[40] A. Alexandru, P. F. Bedaque, H. Lamm, and S. Lawrence (NuQS Collaboration), Phys. Rev. Lett. **123**, 090501 (2019).

[41] A. Alexandru, P. F. Bedaque, S. Harmalkar, H. Lamm, S. Lawrence, and N. C. Warrington (NuQS Collaboration), Phys. Rev. D **100**, 114501 (2019).

[42] N. Klco, J. R. Stryker, and M. J. Savage, arXiv:1908.06935 [Phys. Rev. D (to be published)].

[43] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Commun. Math. Phys. **270**, 359 (2007).

[44] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, arXiv:1912.08854.

[45] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Proc. Natl. Acad. Sci. U.S.A. **115**, 9456 (2018).

[46] H. Lamm, S. Lawrence, and Y. Yamauchi (NuQS Collaboration), Phys. Rev. D **100**, 034518 (2019).

[47] B. Chakraborty, M. Honda, T. Izubuchi, Y. Kikuchi, and A. Tomiya, arXiv:2001.00485.

[48] D. E. Kharzeev and Y. Kikuchi, arXiv:2001.00698.