ARTICLE OPEN

Check for updates

Sign-problem free quantum stochastic series expansion algorithm on a quantum computer

Kok Chuan Tan[1,2]✉, Dhiman Bhowmick[2] and Pinaki Sengupta[2]

A quantum implementation of the Stochastic Series Expansion (SSE) Monte Carlo method is proposed, and is shown to offer significant advantages over classical implementations of SSE. In particular, for problems where classical SSE encounters the sign problem, the cost of implementing a Monte Carlo iteration scales only linearly with system size in quantum SSE, while it may scale exponentially with system size in classical SSE. In cases where classical SSE can be efficiently implemented, quantum SSE still offers an advantage by allowing for more general observables to be measured.

npj Quantum Information (2022)8:44 ; https://doi.org/10.1038/s41534-022-00555-x

INTRODUCTION

Quantum Monte Carlo methods have evolved to be indispensable in the study of strongly correlated many-body systems where the interplay between competing interactions result in quantum phases that are not observed in their non-interacting counterparts. In the study of strongly correlated systems, computational methods are vital, as standard analytical techniques based on single particle picture or perturbation theory are often rendered ineffective. Among the different numerical techniques, quantum Monte Carlo methods have proven to be very powerful for their ability to simulate a wide range of realistic microscopic Hamiltonians on relatively large system sizes at all temperatures and in all dimensions.

The phenomenal development of quantum information theory over the past two decades and the advent of quantum computers in the past couple of years have significantly broadened the potential for simulating strongly interacting quantum many-body systems. The ability to represent superposition of states directly on a quantum computer promises exponential speedup of quantum algorithms over their classical counterparts. Algorithms that exploit quantum hardware to speed up simulations of the thermal Gibbs state of many-body systems have previously been explored in refs. [1–8], but we are still at a nascent stage of harnessing the power of quantum computation in studying correlated systems.

The Stochastic Series Expansion (SSE)[9–12] method is a widely used Quantum Monte Carlo (QMC) method for simulating models of quantum many-body systems. It is based on sampling the series expansion of $\exp(-\beta H)$ up to a sufficiently high order. A significant advantage of SSE is that expectation values that are obtained via this method are exact, up to statistical errors. Other approaches include the world line method[13–16], and the Density Matrix Renormalization Group (DMRG) method[16,17,18]. In this article, we propose an implementation of SSE on a quantum computer and compare its efficiency to conventional SSE on a classical computer. We refer to the former as quantum SSE and the latter as classical SSE and demonstrate several advantages that quantum SSE has over classical SSE. In particular, it will be argued that the no-branching requirement[19] from classical SSE can be relaxed in quantum SSE, which leads to important consequences for the simulation of many-body systems. This quantum advantage stems from the ability of quantum processors to prepare nontrivial superpositions of quantum states. In the subsequent discussion, it is assumed the simulated object is a quantum many-body system with N particles and k-local interactions.

First, lifting the no-branching requirement in quantum SSE means that we are no longer limited to basis states that permit a diagonal representation. This has the effect of allowing more general quantum observables to be measured in quantum SSE.

The second consequence is that quantum SSE always leads to nonnegative weights, which are directly sampled from the measurement probabilities of quantum circuits that scale linearly with system size. By comparison, classical SSE by itself cannot guarantee nonnegative weights without imposing strict conditions such as no-branching. This implies that quantum computers may be able to simulate many-body systems currently inaccessible to classical SSE methods due to the famous sign problem[20,21]. Notably, the Quantum Metropolis Sampling (QMS)[3] algorithm also avoids the sign problem by repeated use of the quantum phase estimation algorithm[22]. However, quantum phase estimation requires deep quantum circuits, and approximates the unitary operation $U = \exp(iHt)$ via the Suzuki-Trotter decomposition[23]. This necessarily introduces a systematic error, unlike exact QMC methods such as SSE, which does not involve Trotterization. The same argument applies to any putative quantum version of Trotter-based algorithms such as the World Line QMC. Although the Continuous Time (CT) World Line method[16,24] avoids Trotterization, the simpler structure of the SSE operator string compared to CT QMC makes it better suited for implementation on quantum architectures, especially on near term, noisy quantum processors which require simpler quantum circuits in general.

This article is structured as follows: First, we introduce the broad ideas underlying the SSE QMC method. Second, we will describe a possible SSE implementation on a quantum computer, first for a simpler special case, then for the more general case. Third, we simulate the quantum SSE algorithm for one dimensional spin chains and compare it with exact results. Finally, we discuss how the no-branching condition and the sign problem affects classical SSE, and evaluate the advantages that quantum SSE offers over classical SSE.

[1]Institute of Fundamental and Frontier Sciences, University of Electronic Science and Technology of China, Chengdu, China. [2]School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore, Republic of Singapore. ✉email: bbtankc@gmail.com

RESULTS

Stochastic series expansion

The canonical SSE is a finite temperature quantum Monte Carlo algorithm based on the stochastic sampling of diagonal matrix elements in the Taylor series expansion of the density matrix in a suitably chosen basis[19]. For efficient implementation, it is necessary to decompose the Hamiltonian as $H=-\sum_b H_b$. The partition function is written as

$$\begin{aligned} Z &:= \mathrm{Tr}(\mathrm{e}^{-\beta H}) = \sum_\alpha \langle\alpha|\mathrm{e}^{-\beta H}|\alpha\rangle \\ &= \sum_\alpha \sum_{n=0}^{\infty} \frac{\beta^n}{n!}\langle\alpha|\sum_{b_n} H_{b_n}\Bigg)\ldots\Bigg(\sum_{b_1} H_{b_1}\Bigg)|\alpha\rangle \\ &\approx \sum_{n=0}^{M}\sum_b\sum_\alpha \frac{\beta^n}{n!}\langle\alpha|H_{b_n}\ldots H_{b_1}|\alpha\rangle, \end{aligned} \tag{1}$$

where $\{|\alpha\rangle\}$ is a complete set of states, b denotes the operator string $b_n\ldots b_1$ and M is some sufficiently large cutoff in the expansion power. Assuming that each term $\langle\alpha|H_{b_n}\ldots H_{b_1}|\alpha\rangle$ is nonnegative, the algorithm proceeds by randomly sampling the configuration space $\mathcal{C} := \{(n,b,\alpha)\ \forall\ n,b,\alpha\}$. The thermal expectation value of any operator, O, is given by

$$\langle O\rangle = \mathrm{Tr}(O\mathrm{e}^{-\beta H})/Z \equiv \sum_C p_C f(O,C) \tag{2}$$

where $p_C := \frac{\beta^n}{n!}\langle\alpha|H_{b_n}\ldots H_{b_1}|\alpha\rangle/Z$. Finding $\langle f(O,C)\rangle$ for any O is not necessarily trivial, but when O is diagonal in the chosen basis, $f(O,C) := \langle\alpha|O|\alpha\rangle$ is an unbiased estimator that is readily evaluated.

The SSE method is numerically exact up to statistical errors because the truncation of the Taylor series expansion of the density matrix does not introduce any systematic error. While in principle, operator strings for every expansion power n should be considered, their weight decreases rapidly with $\frac{\beta^n}{n!}$ for large values of $n\gg\beta$, and is not reachable in practice when a finite number of MC steps are performed. A finite simulation always has some maximum operator string length that can be found empirically via the Monte Carlo updating scheme described as follows: The maximum expansion power depends on the values of the inverse temperature, β, and the Hamiltonian parameters, and is determined empirically by dynamically adjusting the operator string length during the equilibration stage of the simulation. M is then set as some higher number that is never reached by the finite simulation. Importantly, because no operator string length sampled ever reaches M, the truncated terms do not contribute to the final statistics of the simulation so there is no systematic contribution to the error, i.e., the only errors are statistical. In contrast, Trotterization with a finite number of steps always contribute errors regardless of the number of samples obtained, so the error contribution is inherently systematic.

SSE on a quantum computer

We now propose a method of implementing the SSE Monte Carlo algorithm on a quantum computer. The classical implementation of the SSE method requires that H_b satisfy a no-branching condition in order for the algorithm to be efficient (see Discussion). However, this requirement is no longer necessary on quantum computers as they naturally allow for superpositions of a large number of states. We can therefore choose a more convenient decomposition. In general, it is always possible to decompose any Hamiltonian as a sum of products of Pauli matrices:

$$H=\sum_b h_b \bigotimes_{i=1}^{N}\sigma_{b^i}^{(i)}, \tag{3}$$

where in general $b^i=0,1,2,3$ and $\sigma_0 := \mathbb{1}$, $\sigma_1 := \sigma_x$, $\sigma_2 := \sigma_y$ and $\sigma_3 := \sigma_z$. Note that in this notation, we used the upper index to label the ith Pauli matrix in the product. This is different from the lower index used to label the operator string $b=b_n\ldots b_1$ in $H_{b_n}\ldots H_{b_1}$.

In order to illustrate the quantum SSE method, we first consider a special case where the operators h_{b_i} mutually commute, for example, $b^i=0,1$ ($\mathbb{1}$ and σ_x). Specifically, we consider the Hamiltonian

$$H = -\sum_{\langle j,k\rangle} J_{jk}\sigma_x^{(j)}\sigma_x^{(k)}, \tag{4}$$

Such problems can already be nontrivial. For instance, in ref. [25], the Hamiltonian (4) was considered as an example of a many-body system that is NP hard to simulate for certain lattice configurations.

We first define $H_b := h_b\otimes_{i=1}^N\sigma_{b^i}^{(i)} + |h_b|\mathbb{1}$, which ensures that H_b is always positive semidefinite. We can readily verify that $\{H_b\}$ forms a set of mutually commuting semidefinite operators:

$$\frac{H_b H_{b'}}{|h_b h_{b'}|} = \frac{H_{b'} H_b}{|h_{b'} h_b|} \tag{5}$$

where we used the fact that $b_i=0,1$ and $\sigma_{b^i}^{(i)}$ can only be either be $\mathbb{1}$ or σ_x and they mutually commute. The positive semidefiniteness and commutativity of H_b guarantees that a product of such operators $H_{b_n}\ldots H_{b_1}$ is also positive semidefinite (see Supplementary Information under "Positive semidefiniteness: special case"). Making H_b positive semidefinite is equivalent to adding a constant to the Hamiltonian $H\to H+k\mathbb{1}$ where $k := \sum_b|h_b|$, such that the total Hamiltonian is also positive semidefinite. With the positivity of $\langle\alpha|H_{b_n}\ldots H_{b_1}|\alpha\rangle$ assured, we need a method of sampling the relative weight of a given configuration (n,b,α). In classical SSE, positive weights cannot be assured except for some special choices of $|\alpha\rangle$ (see Discussion). Here, we describe a quantum implementation which ensures positivity for arbitrary $|\alpha\rangle$.

A quantum algorithm relies on defining appropriate states and unitary operations that can be implemented efficiently on a quantum computer. In the following, we outline how to develop an efficient estimator of the relative weight of a configuration. Let us consider a state with $(N+n)$ qubits of the form:

$$|\psi_{\mathrm{in}}\rangle = |\alpha_{\mathrm{A}}\rangle|+_{\mathrm{B}_1}\rangle\ldots|+_{\mathrm{B}_n}\rangle, \tag{6}$$

where $|+\rangle := (|0\rangle+|1\rangle)/\sqrt{2}$, N is the number of particles in the system we are simulating, n is the expansion power in SSE, $\mathrm{A}=\mathrm{A}_1\ldots\mathrm{A}_N$ denotes the simulated system, and A_i with $i\in\{1,2,\ldots N\}$ denotes the ith particle in the simulated system. Observe that $H_b = |h_b|\left[\mathrm{sgn}\,(h_b)\otimes_{i=1}^N\sigma_{b^i}^{(\mathrm{A}_i)} + \mathbb{1}_{\mathrm{A}}\right]$ is a superposition of 2 unitary operators $\mathrm{sgn}\,(h_b)\otimes_{i=1}^N\sigma_{b^i}^{(\mathrm{A}_i)}$ and $\mathbb{1}_{\mathrm{A}}$. We define the following controlled unitary operation:

$$U_{\mathrm{A,B}_i}|\alpha_{\mathrm{A}}\rangle|0_{\mathrm{B}_i}\rangle = \mathbb{1}_{\mathrm{A}}|\alpha_{\mathrm{A}}\rangle|0_{\mathrm{B}_i}\rangle \tag{7}$$

$$U_{\mathrm{A,B}_i}|\alpha_{\mathrm{A}}\rangle|1_{\mathrm{B}_i}\rangle = \left[\mathrm{sgn}\,(h_b)\bigotimes_{j=1}^{N}\sigma_{b^j}^{(\mathrm{A}_j)}\right]|\alpha_{\mathrm{A}}\rangle|1_{\mathrm{B}_i}\rangle. \tag{8}$$

By calculating the expectation value $\langle\psi_{\mathrm{in}}|U_{\mathrm{A,B}_i}|\psi_{\mathrm{in}}\rangle$, one can evaluate an estimator for the relative weight

$$q(n,b,\alpha) := |\langle\psi_{\mathrm{in}}|U_{\mathrm{A,B}_i}|\psi_{\mathrm{in}}\rangle|^2 \equiv \left|\frac{\langle\alpha_{\mathrm{A}}|H_{b_n}\ldots H_{b_1}|\alpha_{\mathrm{A}}\rangle}{2^n h_{b_n}\ldots h_{b_1}}\right|^2. \tag{9}$$

The spectrum of $H_{b_i}/|h_{b_i}|$ lies in the range $[0,2]$ and so the spectrum of $H_{b_n}\ldots H_{b_1}/|h_{b_n}\ldots h_{b_1}|$ is within $[0,2^n]$. The projected amplitude is therefore not necessarily exponentially small due to the $1/2^n$ factor even for relatively large expansion orders n. In order to find a bounded error estimate for $q(n,b,\alpha)$, quantum circuits are repeatedly prepared and then measured in the computational basis t times. An example of the type of quantum circuit being

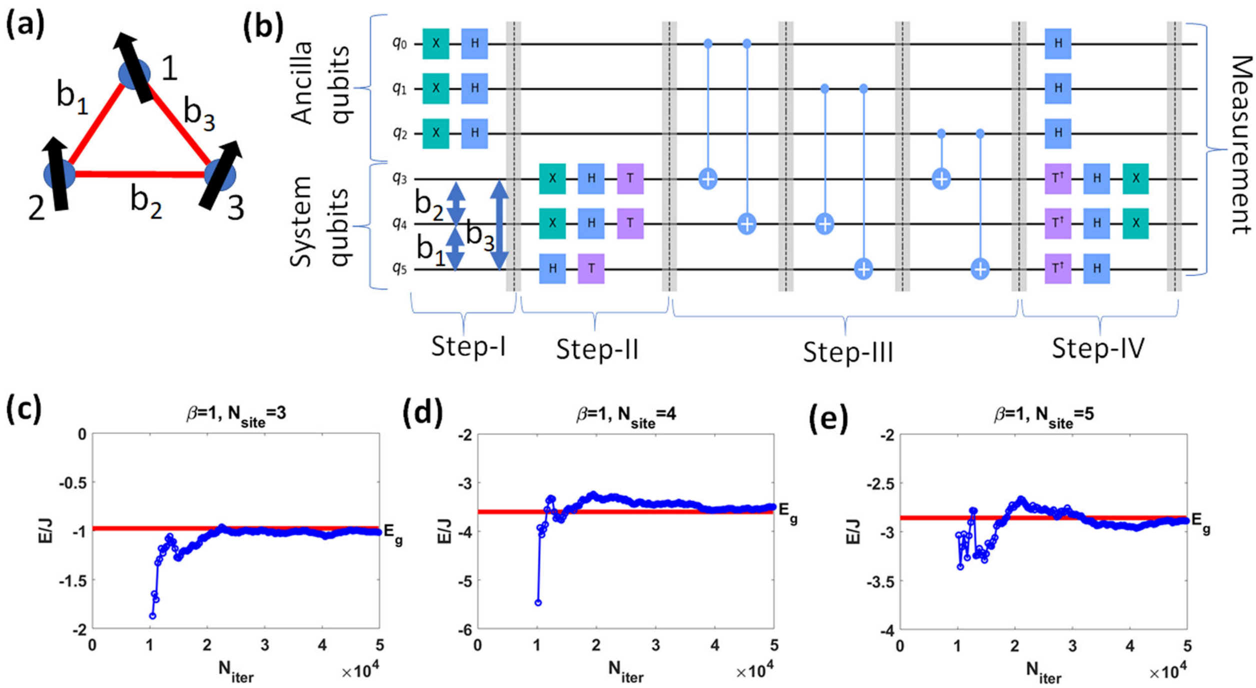

Fig. 1 Quantum SSE simulation of 1D antiferromagnetic spin-1/2 chain. a 1D spin-1/2 chain with antiferromagnetic interaction and periodic boundary condition. For $N = 3$, the sites are labeled 1, 2, 3 and the corresponding bonds b_1, b_2, b_3. **b** An example schematic of the quantum circuit calculating the expectation value of string of unitary operators $U^{b1}_{A,q_1} U^{b2}_{A,q_2} U^{b3}_{A,q_3}$. Further details are given in the main text. **c**, **d**, **e** Illustrates the convergence of the mean energy (blue-line with circles) determined by quantum SSE at $\beta = 1$. The finite temperature energy of the system represented by the red horizontal line is obtained via exact diagonalization. The x-axis indicates the number of Metropolis iterations $N_{\rm iter}$ for $N = 3$, 4, 5 respectively.

measured is shown in Fig. 1b. Each of these measurements yield a classical bitstring that is a series of 0s or 1s and out of these t bitstrings, we count the number of bitstrings that are all 0s. The ratio to the total number of samples t estimates $q(n, b, \alpha)$. Since this is a binary distribution, the sample variance scales with ~$1/t$, so estimating $q(n, b, \alpha)$ to target precision ε requires $t \sim 1/\varepsilon^2$ samples in general. In this way, the configuration weights can be estimated to any target degree of numerical precision.

Alternatively, we can also perform a quantum subroutine called amplitude estimation[26] (see Supplementary Information under "Amplitude estimation"). In general, to estimate the probability p to any desired precision ϵ with success probability $1 - \delta$, the subroutine needs to be invoke certain unitary operations a total of $t = t(\epsilon, \delta)$ times, where $t(\epsilon, \delta)$ only depends on the desired precision ϵ and success probability $1 - \delta$. In this case, the variance scales with ~$1/t^2$, where t is now the number of times the unitary operations are applied rather than the number of independent samples.

Stochastic sampling of operator space

Once the relative weight of a configuration C is sampled, the Monte Carlo simulation proceeds by stochastically sampling the operator space via the Metropolis method. This consists of randomly selecting some new configuration C', and then accepting the newly chosen configuration with probability $P_{\rm accept}(C \to C') := \min[W(C')/W(C), 1]$ where $W(C)$ is the weight of a configuration $C = (n, b, \alpha)$. It is given by the following expression

$$\begin{aligned} W(C) &:= \tfrac{\beta^n}{n!} \langle \alpha | H_{b_n} \dots H_{b_1} | \alpha \rangle \\ &= \tfrac{\beta^n}{n!} |2^n h_{b_n} \dots h_{b_1}| \sqrt{q(n, b, \alpha)}, \end{aligned} \tag{10}$$

where $q(n, b, \alpha)$ is the probability sampled in Eq. 9. In the Metropolis algorithm, it is implicitly assumed that the probability of selecting C' when the current configuration is C is the same as the probability of selecting C when the current configuration is C', i.e., $P_{\rm select}(C \to C') = P_{\rm select}(C' \to C)$. Each of the variables n, b, α is updated separately using the ratio $W(C')/W(C)$ as the probability of accepting an update. In quantum SSE, the update probabilities are determined by $\sqrt{q(n', b', \alpha')/q(n, b, \alpha)}$, which are obtained by sampling the probabilities $q(n, b, \alpha)$ from the quantum circuit (see Eq. 9). Explicit expressions for the update probabilities are provided in the Supplementary Information under "Metropolis acceptance weights". The calculation of the update probabilities does not involve the 2^n multiplicative factor in Eq. 10, so an exponential blowup in the sampling error does not occur.

Quantum implementation of SSE for general Hamiltonians

In the section "SSE on a quantum computer" we demonstrated a special case implementation of quantum SSE where the quantity $\langle \alpha | H_{b_n} \dots H_{b_1} | \alpha \rangle$ is guaranteed to be non-negative. For general Hamiltonians, this may not always be possible because the operator $H_{b_n} \dots H_{b_1}$ may not be Hermitian. In this section, we discuss an implementation for more general Hamiltonians.

Recall the expression for the expectation value of an operator:

$$\langle O \rangle = \sum_{n,b,\alpha} \frac{\beta^n}{n!} \langle \alpha | H_{b_n} \dots H_{b_1} | \alpha \rangle \langle \alpha | O | \alpha \rangle / Z.$$

Note that the summation over all possible strings b contain $\langle \alpha | H_{b_n} \dots H_{b_1} | \alpha \rangle$, as well as its complex conjugate $\langle \alpha | H_{b_1} \dots H_{b_n} | \alpha \rangle$. Since $\langle \alpha | H_{b_n} \dots H_{b_1} | \alpha \rangle + \langle \alpha | H_{b_1} \dots H_{b_n} | \alpha \rangle = 2\mathrm{Re} \langle \alpha | H_{b_n} \dots H_{b_1} | \alpha \rangle$, only the real part of each term contributes to the expectation value. Hence we can equivalently write

$$\langle O \rangle = \sum_{n,b,\alpha} \frac{\beta^n}{n!} \mathrm{Re} \langle \alpha | \mathrm{H}_{b_n} \dots \mathrm{H}_{b_1} | \alpha \rangle \langle \alpha | \mathrm{O} | \alpha \rangle / \mathrm{Z}. \tag{11}$$

Therefore, in order to implement quantum SSE, we only need to sample the real portion of $\langle \alpha|H_{b_n}\dots H_{b_1}|\alpha\rangle$ and ensure that it is nonnegative. We show that this can be done by adding a sufficiently large constant to the Hamiltonian (see Supplementary Information under "Positive semidefiniteness: general case").

Suppose $M \geq n$ is the cutoff in the expansion power (see Eq. 1). For a fixed M, let $H_b := |h_b|\left[\operatorname{sgn}(h_b)\otimes_{i=1}^{N}\sigma_{b^i}^{(A_i)} + 2M\mathbb{1}_A\right]$. We note that this is an unequal superposition of 2 unitary operations that depends on the cutoff value M.

We introduce the state

$$|\psi_{\text{in}}\rangle := |\alpha_A\rangle|\phi_{B_1}\rangle\dots|\phi_{B_n}\rangle|+_C\rangle, \tag{12}$$

where

$$|\phi_{B_i}\rangle := \sqrt{(2M)/(2M+1)}|0_{B_i}\rangle + \sqrt{1/(2M+1)}|1_{B_i}\rangle$$

and $|+_C\rangle := \frac{1}{\sqrt{2}}(|0_C\rangle + |1_C\rangle)$. As before, by introducing appropriate unitary operators, the relative weight of the configurations can be evaluated. In addition to U_{A,B_i}, we further define the unitary $V_{AB,C}$, which is controlled by qubit C:

$$\begin{aligned} &V_{AB,C}|\alpha_A\rangle|\phi_{B_1}\rangle\dots|\phi_{B_n}\rangle|0_C\rangle \\ &:= U_{A,B_1}\dots U_{A,B_n}|\alpha_A\rangle|\phi_{B_1}\rangle\dots|\phi_{B_n}\rangle|0_C\rangle \\ &V_{AB,C}|\alpha_A\rangle|\phi_{B_1}\rangle\dots|\phi_{B_n}\rangle|1_C\rangle \\ &:= U_{A,B_n}\dots U_{A,B_1}|\alpha_A\rangle|\phi_{B_1}\rangle\dots|\phi_{B_n}\rangle|1_C\rangle \end{aligned} \tag{13}$$

For any given expansion power n, we can verify the expression:

$$\begin{aligned} \langle\psi_{\text{in}}|V_{AB,C}|\psi_{\text{in}}\rangle &= \frac{\langle\alpha_A|H_{b_1}\dots H_{b_n}|\alpha_A\rangle + \langle\alpha_A|H_{b_n}\dots H_{b_1}|\alpha_A\rangle}{2(2M+1)^n|h_{b_n}\dots h_{b_1}|} \\ &= \frac{\text{Re}\langle\alpha_A|H_{b_n}\dots H_{b_1}|\alpha_A\rangle}{(2M+1)^n|h_{b_n}\dots h_{b_1}|}, \end{aligned} \tag{14}$$

which allows us to define an estimator for the relative weight:

$$\begin{aligned} q(n,b,\alpha) &:= |\langle\psi_{\text{in}}|V_{AB,C}|\psi_{\text{in}}\rangle|^2 \\ &\equiv \left|\frac{\text{Re}\langle\alpha_A|H_{b_n}\dots H_{b_1}|\alpha_A\rangle}{(2M+1)^n|h_{b_n}\dots h_{b_1}|}\right|^2 \end{aligned} \tag{15}$$

Note that the spectrum of $H_{b_i}/|h_{b_i}|$ is in the range $[0, 2M+1]$ so the absolute value of $\text{Re}\langle\alpha_A|H_{b_n}\dots H_{b_1}|\alpha_A\rangle/|h_{b_n}\dots h_{b_1}|$ is within the range $[0,(2M+1)^n]$. The configuration weights, and hence $\text{Re}\langle\alpha_A|H_{b_n}\dots H_{b_1}|\alpha_A\rangle$, are always nonnegative.

From the above arguments, we see that the configurations can be sampled by measuring the probability $q(n,b,\alpha)$. The Metropolis portion of the simulation then proceeds as before, where the acceptance probability depends on the ratio $\sqrt{q(n',b',\alpha')/q(n,b,\alpha)}$. Note that the above argument finds a sufficiently large constant to add to the Hamiltonian to avoid negative weights. This constant may be too large for many specific problems. We expect that the minimum constant that is required can be optimized on a case by case basis. We also highlight that adding a constant to the Hamiltonian is only necessary when the initial Hamiltonian contains negatively weighted configurations. The above arguments show that quantum SSE is able to avoid negative weights for general Hamiltonians and arbitrary basis states $|\alpha\rangle$, which is not possible in general for classical SSE except for special cases.

Example: 1D antiferromagnetic spin-1/2 chain

To demonstrate the proposed algorithm, we simulated the Hamiltonian of one dimensional periodic spin chains with $N = 3$, 4, 5 sites. The algorithm was implemented using the quantum simulation toolkit Qiskit[27] and the measured observables were compared with results from exact diagonalization. The Hamiltonian with antiferromagnetic exchange is given by

$$H' = J\sum_b \sigma_x^{b(1)}\sigma_x^{b(2)}, \tag{16}$$

where $J > 0$ and $b(i)$ is the i-th site of the b-th bond (see Fig. 1a). The classical SSE implementation violates the no-branching condition and suffers from the sign problem if we do not use the eigenstates of σ_x to construct the basis states, $|\alpha\rangle$. In quantum SSE this constraint no longer exists as the no-branching requirement is lifted and the string of bond operators have positive-semidefinite weights. To illustrate this, we choose the basis states $|\alpha\rangle$ to be product states of σ_z eigenvectors (i.e., products of $|\uparrow\rangle, |\downarrow\rangle$) rotated by Hadamard gates and non-Clifford phase gates $T = \begin{pmatrix}1 & 0 \\ 0 & e^{i\pi/4}\end{pmatrix}$. In general, quantum circuits with non-Clifford gates are known to be difficult to simulate classically[28,29].

The energy calculations from quantum SSE as a function of the number of Metropolis iterations are shown Fig. 1c–e for site numbers $N = 3$, 4, 5 respectively at $\beta = 1$. For all the cases considered, it was observed that the mean energy computed via quantum SSE converges towards the exact finite temperature energy of the system obtained from exact diagonalization, thus demonstrating the validity of the algorithm.

DISCUSSION

Recall that implementing SSE Monte Carlo requires each term $\langle\alpha|H_{b_n}\dots H_{b_1}|\alpha\rangle$ in the expansion to be nonnegative. In general, this cannot be always guaranteed except for special cases. This results in the infamous *sign problem*[20,21] which severely restricts the applicability of QMC methods.

In the classical implementation of SSE, the sign problem arises from the so-called no-branching condition. This is the requirement that $H_b|\alpha\rangle \propto |\alpha'\rangle$, where $|\alpha'\rangle$ is also a basis vector. In other words, we always have to use a decomposition of $H = \Sigma_b H_b$ such that H_b does not create superpositions of basis states. For any given basis, this means that every H_b satisfying the no-branching requirement can be classified as a diagonal operator satisfying $H_b|\alpha\rangle \propto |\alpha\rangle$ for every α, or an off-diagonal update satisfying $H_b|\alpha\rangle \propto |\alpha'\rangle$ where $\alpha \neq \alpha'$ for some α.

A diagonal update can always be made positive by adding a sufficiently large constant. This is because if H_b is a diagonal update, then $H_b'|\alpha\rangle := (H_b + k\mathbb{1})|\alpha\rangle \propto |\alpha\rangle$ is also a diagonal update.

On the other hand, we see that if H_b is an off-diagonal update, adding a constant will necessarily create a superposition of basis states, since $(H_b + k\mathbb{1})|\alpha\rangle \propto h_{b,\alpha}|\alpha'\rangle + k|\alpha\rangle$ where $\alpha \neq \alpha'$. This means that we cannot guarantee that H_b is always positive semidefinite for off-diagonal updates, without violating the no-branching condition. This in turn implies that $\langle\alpha|H_{b_n}\dots H_{b_1}|\alpha\rangle$ is not necessarily positive, which is the sign problem.

From the above, we see that the sign problem exists because of the no-branching requirement. If we avoid the sign problem by lifting no-branching requirement, one will have to keep track of all the off-diagonal elements of $H_{b_n}\dots H_{b_1}|\alpha\rangle$. In the worst case, the computational resources required to keep track of an arbitrary superposition of basis states is of the order $\mathcal{O}(\exp(N))$, where N is the number of particles.

The primary benefit of the quantum SSE method is that it does not require the no-branching condition, as quantum computers naturally allows for the creation of superpositions of quantum states. This allows us to sample the relative weights of a given configuration directly, without needing to keep track of all the off-diagonal elements. By lifting the no-branching requirement, we can always ensure that the relative weights are nonnegative, thus also avoiding the sign problem. We have shown this for the special case where the Hamiltonian can be decomposed into products of $\mathbb{1}$ or σ_x, as well as for more general Hamiltonians.

We now discuss several advantages that quantum SSE offers over classical SSE. The first advantage is that a much wider range of observables can be measured using quantum SSE. Both quantum and classical SSE compute statistical averages most easily when the observable O is diagonal in the basis $|\alpha\rangle$. Unlike classical SSE, however, quantum SSE does not require the no-branching condition so there is no longer any limitations on the choice of basis states $|\alpha\rangle$. For any given operator O, we can now choose any basis $\{|\alpha\rangle\}$ that diagonalizes the observable O, as long as the the basis state $|\alpha\rangle$ can be efficiently prepared on the quantum computer. The corresponding estimator for the observable O is then $f(O, C) := \langle\alpha|O|\alpha\rangle$. Consequently, quantum SSE allows more general quantum observables to be measured compared to classical SSE. An example of this is when $O = |\phi\rangle\langle\phi|$ for a known quantum state $|\phi\rangle$. In this case, O is the projector onto the state $|\phi\rangle$ and $\langle O\rangle = \langle\phi|e^{-\beta H'}/Z|\phi\rangle$ is the overlap between $|\phi\rangle$ and the thermal state $e^{-\beta H'}/Z$. In general, finding the state overlap is not easily implementable using classical SSE. In the Shastry-Sutherland model[30–32] for instance, this can be used to directly verify that the ground state is a product of singlet pairs. This is achieved by by letting $|\phi\rangle$ be a product of singlets and then sampling the expectation values using quantum SSE.

A second advantage of quantum SSE is the low circuit complexity required to estimate the configuration weight. Consider the computational cost of implementing quantum SSE for the special case discussed in the section "SSE on a quantum computer". Here, sampling $\langle\alpha_A|H_{b_n}\dots H_{b_1}|\alpha_A\rangle$ given some operator string b requires a total of n unitary operations to be performed, multiplied by the number of samples t for any target numerical precision. Combine this with the fact that $\langle n\rangle$, the average expansion power in SSE, is proportional to the system energy and scales with βN, and the end result is that the total number of quantum gates required to sample the configuration weight requires $\mathcal{O}(n) \sim \mathcal{O}(N)$ number of operations, i.e., it scales linearly with system size. A similar argument can also be made if we employ the amplitude estimation algorithm (see Supplementary Information under "Amplitude estimation"). An analysis for general Hamiltonians will lead to the same conclusion, since essentially the same set of unitary operations are performed, except with an additional control operation. We therefore expect the size of the quantum circuits in quantum SSE to scale with $\mathcal{O}(N)$ in general.

Compare this to the classical SSE algorithm. In classical SSE, the cost of sampling the configuration weight is $\mathcal{O}(N)$ only when the no-branching condition is satisfied and there is no sign problem. As previously discussed, a consequence of this is that the no-branching condition fixes the basis $|\alpha\rangle$, which limits the kinds of observables that classical SSE can measure. More generally, without the no-branching condition, classical SSE encounters the sign problem and the computational cost of avoiding negative weights is generally $\sim \mathcal{O}(\exp(N))$. By comparison, the complexity of the quantum circuit is $\mathcal{O}(N)$ if the basis state $|\alpha\rangle$ is generated by a circuit of size $\mathcal{O}(N)$, or poly(N) if $|\alpha\rangle$ is generated by a circuit of size poly(N). Taking into account all such possible basis states, we expect the improvement in terms of computational complexity to be exponential.

The third advantage of quantum SSE is that it can be efficiently parallelized to a short depth quantum circuit. Short depth circuits belong to the quantum complexity class QNC, which is the class of quantum circuits with polylogarithmic depth. In order to see this, observe that the unitary operations U_{A,B_i} and $V_{AB,C}$ are composed of layers of controlled-Pauli operations. In general, such operations are known to belong to the quantum complexity class QNC[133], which has depth complexity $\mathcal{O}(\log N)$. This means that the depth complexity of quantum SSE belongs to QNC[1] for any $|\alpha\rangle$ generated in logarithmic depth, or more generally QNC if $|\alpha\rangle$ is generated in polylogarithmic depth. In particular, the example in Fig. 1 belongs to QNC[1]. Note that this circuit contains non-Clifford T gates and that non-Clifford gates are not efficiently simulable on classical computers[28,29]. Such short depth quantum circuits are also especially well suited for implementation on NISQ processors[34] with limited coherence times. To the best of our knowledge, we are not aware of a similar result for classical SSE.

In summary, we proposed a quantum implementation of the SSE Monte Carlo algorithm and compared it to its classical counterpart. It was shown that the cost of implementing a single Monte Carlo update in quantum SSE scales linearly with the number of particles N. We compare this to the classical implementation of SSE, where certain many-body systems exhibit the sign problem and incurs an additional cost that scales exponentially with N. The quantum algorithm was able to avoid this by ensuring that the weight of the configuration is always positive, regardless of the chosen basis. This suggests that quantum computers can significantly speed up the simulation of complex quantum many-body systems. Even when the sign problem is not present in classical SSE, quantum SSE can still be advantageous, since it allows for more general observables to be measured. This was demonstrated via a numerical simulation of a 1D spin-1/2 chain using the quantum SSE algorithm in combination with a basis that is typically hard to implement in classical SSE. In all cases considered, it was shown that quantum SSE converged to the correct results obtained from exact diagonalization.

It is known that estimating the ground state energy of a k-local Hamiltonian is QMA complete[25,35,36]. Here, we have shown that quantum SSE can perform bounded error estimates of the configuration weights in polynomial time, which is not always possible in classical SSE. While this quantum speedup removes a major bottleneck in classical SSE, it does not necessarily imply an exponentially fast rate of convergence to the ground state energy, nor that QMA complete problems can be solved in polynomial time in general. Nonetheless, the quantum SSE algorithm shows that quantum computers are promising tools for accelerating the SSE Monte Carlo simulation in many scenarios. This may provide a pathway for probing the quantum properties of many-body systems that are currently inaccessible to existing classical techniques.

METHODS

We describe in detail how the simulation in Section "Example: 1D antiferromagnetic spin chain" was performed. After adding identity operators to the bond operators to make them positive-semidefinite, the effective Hamiltonian for $J = 1$ in the quantum SSE simulation is

$$H = \sum_b H_b, \tag{17}$$

where $H_b = \mathbb{1} - \sigma_x^{b(1)}\sigma_x^{b(2)}$.

The controlled unitary operator U^b_{A,B_i} performs the map,

$$U^b_{A,B_i}|\alpha_A\rangle|+_{B_i}\rangle = \frac{1}{\sqrt{2}}\left(|\alpha_A\rangle|0_{B_i}\rangle - \sigma_x^{b(1)}\sigma_x^{b(2)}|\alpha_A\rangle|1_{B_i}\rangle\right). \tag{18}$$

that were introduced in the section "SSE on a quantum computer". The expectation value of operator strings $H_{b_1}H_{b_2}\dots H_{b_n}$—needed to evaluate the partition function—is related to U^b_{A,B_i} according to ($|h_b| = 1\ \forall\, b$ since $J = 1$):

$$\langle\alpha_A|\langle+_{B_1}|\dots\langle+_{B_n}|U^{b_n}_{A,B_n}\dots U^{b_1}_{A,B_1}|\alpha_A\rangle|+_{B_1}\rangle\dots|+_{B_n}\rangle = \frac{1}{2^n}\langle\alpha_A|H_{b_n}\dots H_{b_1}|\alpha_A\rangle. \tag{19}$$

An example quantum circuit performing this measurement is shown in Fig. 1b. In this example circuit, there are six physical qubits where qubits q_0, q_1, q_2 are ancillae, and qubits q_3, q_4, q_5 represent the system. The quantum circuit determines the expectation value of string of $U^{b_3}_{A,q_2}U^{b_2}_{A,q_1}U^{b_1}_{A,q_0}$ for a three site periodic system when $n = 3$. We will describe in detail the steps involved in the circuit in Fig. 1b.

In Step I, the ancilla qubits q_0, q_1 and q_3 are prepared in the states $|-_{q_0}\rangle$, $|-_{q_1}\rangle$, $|-_{q_2}\rangle$ respectively, using Hadamard and Pauli X gates.

In Step II the system qubits q_3, q_4 and q_5 representing the spin-1/2 sites of the physical spin chain are prepared in some initial state $|\alpha\rangle_A$ by applying either the identity or the Pauli X-gate followed by a Hadamard and the non-Clifford T-gate. This generates basis states via a constant depth, non-Clifford quantum circuit.

In Step III the unitary operations $U^{b_3}_{A,q_2}$, $U^{b_2}_{A,q_1}$ and $U^{b_1}_{A,q_0}$ are completed by sequentially applying CNOT operations onto $|\alpha\rangle_A|-_{q_0}\rangle|-_{q_1}\rangle|-_{q_2}\rangle$.

Finally in Step IV, the qubits are rotated using Hadamard, X gates, inverse T gates and then measured in the computational basis. The probability of measuring all qubits with the outcome 0 gives the square of the expectation value of $U^{b_3}_{A,q_2}U^{b_2}_{A,q_1}U^{b_1}_{A,q_0}$. All of the quantum gates used to construct the quantum circuit are standard quantum gates included as part of the Qiskit library.

After evaluating the expectation value for a given operator string and spin state, the relative weight is computed from $\sqrt{q(n', b', \alpha')/q(n, b, \alpha)}$. The Metropolis algorithm, as described in the section "Stochastic sampling of operator space", is then implemented accordingly to update the quantum state and the operator string. In SSE, the finite temperature energy of the system is evaluated using the expression[19],

$$E = -\frac{\langle n \rangle}{\beta} + N, \tag{20}$$

where $\langle n \rangle$ is the average length of operator string per Metropolis loop. Note that the contributing term N in Eq. 20 is due to adding a constant to the Hamiltonian to ensure positive semidefiniteness.

The Metropolis sampling is initialized with some arbitrary string of operators. The average operator string length $\langle n \rangle$ is then calculated after an initial 10^4 Metropolis steps, and the mean energy is evaluated using Eq. 20. The simulation was performed for site numbers $N = 3, 4, 5$ at a fixed inverse temperature $\beta = 1$.

DATA AVAILABILITY

Data are available from the authors on reasonable request.

CODE AVAILABILITY

Code used to perform the numerical examples are available from the authors on reasonable request.

Received: 22 February 2021; Accepted: 9 March 2022;
Published online: 26 April 2022

REFERENCES

1. Terhal, B. M. & DiVincenzo, D. P. Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A* **61**, 022301 (2000).
2. Bilgin, E. & Boixo, S. Preparing thermal states of quantum systems by dimension reduction. *Phys. Rev. Lett.* **105**, 170405 (2010).
3. Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D. & Verstrate, F. Quantum Metropolis sampling. *Nature* **471**, 87–90 (2011).
4. Riera, A., Gogolin, C. & Eisert, J. Thermalization in nature and on a quantum computer. *Phys. Rev. Lett.* **108**, 080402 (2012).
5. Yung, M.-H. & Aspuru-Guzik, A. A quantum–quantum Metropolis algorithm. *Proc. Natl Acad. Sci.* **109**, 754–759 (2012).
6. Montanaro, A. Quantum speedup of Monte Carlo methods. *Proc. R. Soc. A* **471**, 20150301 (2015).
7. Ge, Y., Molnár, A. & Cirac, J. I. Rapid adiabatic preparation of injective projected entangled pair states and Gibbs states. *Phys. Rev. Lett.* **116**, 080503 (2016).
8. Motta, M. et al. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nat. Phys.* **16**, 205 (2019).
9. Sandvik, A. W. & Kurkijärvi, J. Quantum Monte Carlo simulation method for spin systems. *Phys. Rev. B* **43**, 5950–5961 (1991).
10. Sandvik, A. W. A generalization of Handscomb's quantum Monte Carlo scheme-application to the 1D Hubbard model. *J. Phys. A* **25**, 3667–3682 (1992).
11. Sandvik, A. W. Finite-size scaling of the ground-state parameters of the two-dimensional Heisenberg model. *Phys. Rev. B* **56**, 11678–11690 (1997).
12. Sandvik, A. W. Stochastic series expansion method with operator-loop update. *Phys. Rev. B* **59**, R14157–R14160 (1999).
13. Hirsch, J. E., Sugar, R. L., Scalapino, D. J. & Blankenbecler, R. Monte Carlo simulations of one-dimensional fermion systems. *Phys. Rev. B* **26**, 5033–5055 (1982).
14. Suzuki, M. Relationship between d-dimensional quantal spin systems and (d+1)-dimensional Ising systems: Equivalence, critical exponents and systematic approximants of the partition function and spin correlations. *Prog. Theor. Phys.* **56**, 1454–1469 (1976).
15. Suzuki, M., Miyashita, S. & Kuroda, A. Monte Carlo simulation of quantum spin systems. *I. Prog. Theor. Phys.* **58**, 1377–1387 (1977).
16. Beard, B. B. & Wiese, U.-J. Simulations of discrete quantum systems in continuous Euclidean time. *Phys. Rev. Lett.* **77**, 5130–5133 (1996).
17. White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
18. Schollwöck, U. The density-matrix renormalization group. *Rev. Mod. Phys.* **77**, 259–315 (2005).
19. Sandvik, A. W. Computational studies of quantum spin systems. *AIP Conf. Proc.* **1297**, 135–338 (2010).
20. Foulkes, W. M. C., Mitas, L., Needs, R. J. & Rajagopal, G. Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.* **73**, 33–83 (2001).
21. Henelius, P. & Sandvik, A. W. Sign problem in Monte Carlo simulations of frustrated quantum spin systems. *Phys. Rev. B* **62**, 1102–1113 (2000).
22. Abrams, D. S. & Lloyd, S. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.* **83**, 5162–5165 (1999).
23. Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
24. Prokofev, N. V., Svistuniov, B. V. & Tupitsyn, I. S. Exact, complete, and universal continuous-time worldline Monte Carlo approach to the statistics of discrete quantum systems. *Sov. Phys. J.E.T.P.* **87**, 310–321 (1998).
25. Troyer, M. & Wiese, U.-J. Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Phys. Rev. Lett.* **94**, 170201 (2005).
26. Brassard, G., Mosca, M. & Tapp, A. Quantum amplitude amplification and estimation. *Quantum Comput. Quantum Inf. A Millennium* **305**, 53–74 (2002).
27. Abraham, H. et al. Qiskit: an Open-source Framework for Quantum Computing. https://doi.org/10.5281/ZENODO.2562111 (2019).
28. Gottesman, D. The Heisenberg representation of quantum computers. Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics (Cambridge, MA, International Press, 1999) pp. 32–43.
29. Aaronson, S. & Gottesman, D. Improved simulation of stabilizer circuits. *Phys. Rev. A* **70**, 052328 (2004).
30. Shastry, B. S. & Sutherland, B. Exact ground state of a quantum mechanical antiferromagnet. *Phys. B+C.* **108**, 1069–1070 (1981).
31. Richter, J., Ivanov, N. B. & Schulenburg, J. The antiferromagnetic spin-chain with competing dimers and plaquettes: numerical versus exact results. *J. Phys. Condens. Matter* **10**, 3635–3649 (1998).
32. Miyahara, S. & Ueda, K. Theory of the orthogonal dimer Heisenberg spin model for SrCu2(BO3)2. *J. Phys.: Condens. Matter* **15**, R327–R366 (2003).
33. Moore, C. & Nilsson, M. Parallel quantum computation and quantum codes. *SIAM J. Comput.* **31**, 799–815 (2001).
34. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
35. Kitaev, A. Yu., Shen, A. H. & Vyalyi, M. N. *Classical and quantum computation, volume 47 of Graduate Studies in Mathematics* (AMS, Providence, RI, 2002).
36. Kempe, J., Kitaev, A. & Regev, O. The complexity of the local Hamiltonian problem. *SIAM J. Comput.* **35**, 1070–1097 (2006).

ACKNOWLEDGEMENTS

It is a pleasure to thank Anders Sandvik for helpful discussions. K.C.T. acknowledges support by the University of Electronic Science and Technology of China and the NTU Presidential Postdoctoral Fellowship program funded by Nanyang Technological University. Financial support from the Ministry of Education, Singapore, in the form of Grant No. MOE2018-T1-1-021 is gratefully acknowledged.

AUTHOR CONTRIBUTIONS

K.C.T. conceived the research. D.B. performed the simulation of the quantum algorithm. K.C.T., D.B., and P.S. contributed equally to the analyses and wrote the paper.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41534-022-00555-x.

Correspondence and requests for materials should be addressed to Kok Chuan Tan.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.