Efficient Mean-Field Simulation of Quantum Circuits Inspired by Density Functional Theory

Marco Bernardi[*,†]

†*Department of Applied Physics and Materials Science, California Institute of Technology, Pasadena, CA 91125, USA.*

‡*Department of Physics, California Institute of Technology, Pasadena, CA 91125, USA.*

E-mail: bmarco@caltech.edu

Exact simulations of quantum circuits (QCs) are currently limited to ∼50 qubits because the memory and computational cost required to store the QC wave function scale exponentially with qubit number. Therefore, developing efficient schemes for approximate QC simulations is a current research focus. Here we show simulations of QCs with a method inspired by density functional theory (DFT), a widely used approach to study many-electron systems. Our calculations can predict marginal single-qubit probabilities (SQPs) with over 90% accuracy in several classes of QCs with universal gate sets, using memory and computational resources linear in qubit number despite the formal exponential cost of the SQPs. This is achieved by developing a mean-field description of QCs and formulating optimal single- and two-qubit gate functionals − analogs of exchange-correlation functionals in DFT − to evolve the SQPs without computing the QC wave function. Current limitations and future extensions of this formalism are discussed.

1. Introduction

Noisy intermediate-scale quantum devices promise exciting advances in quantum algorithms with no classical counterpart[1–3]. Classical simulations remain essential to understand the physics of these quantum devices, improve their design, and accelerate their progress[4–7]. An important direction is the development of approximate schemes that are both accurate and computationally efficient, enabling simulations of generic QCs with arbitrary depth and degree of entanglement, ideally with favorable computational scaling. Work in this area has focused on tensor network matrix product states to simulate QCs with a range of structures, gate types, entanglement, and noise[8–14], and more recently on simulations of generic QCs using neural-network quantum states[15]. Despite these notable advances, approximate QC simulations remain an area of active investigation.

There is an intriguing parallel between many-electron and many-qubit systems. In the many-electron problem − a grand challenge in chemistry and materials physics[16] − exact solutions are possible only for systems with one electron (the hydrogen atom). Therefore, unlike QC simulations, electronic structure calculations of molecules and materials are dominated by approximate methods[16–23], among which density functional theory (DFT) is the main workhorse. Leveraging a mean-field description centered on the electron density, DFT achieves low-polynomial scaling with system size, enabling studies of matter with thousands of interacting electrons[17,24]. Methods to study QCs with a similar trade-off of cost and accuracy would be expedient. Early work on relating DFT to QCs focused on formal mapping of QCs onto lattice fermions[25] or connecting time-dependent DFT and spin Hamiltonians[26]. These notable efforts differ in method and scope from this work.

Here we show a DFT-inspired approach for QCs − in short, QC-DFT − able to accurately simulate single-qubit probabilities (SQPs) in QCs with computational cost scaling linearly with qubit number and depth, despite the formal exponential cost of the SQPs. We present results for various random QCs using two different universal gate sets, and demonstrate the formulation and optimization of QC-DFT gate functionals. We also apply this formalism to nonrandom QCs, and study how the SQP distribution changes with QC size as well as solve a simple model Hamiltonian. These results show that even though the exact QC wave function is exponentially complex, marginal probability distributions such as the SQPs can be obtained with a favorable trade-off of cost and accuracy without computing the QC wave function. Although the current formulation is limited to QCs with low entanglement, extensions based on reduced density matrices may enable further progress.

2. Theory: QC-DFT and Gate Functionals

The QC wave function for N qubits can be expanded in the computational basis as

$$\Psi = \sum_{i_1 i_2 \ldots i_N} c_{i_1 i_2 \ldots i_N} |i_1 i_2 \ldots i_N\rangle = \sum_{x}^{2^N-1} c_x |x\rangle \tag{1}$$

where $i_n = 0, 1$ are basis states for a single qubit, x are binary numbers from 0 to 2^{N-1}, c_x are state-vector amplitudes, namely expansion coefficients of the QC wave function, and $|x\rangle = |i_1 i_2 \ldots i_N\rangle$ are N-qubit states in the computational basis (N-bit long bitstrings). For N qubits, accessing this wave function requires storing and manipulating 2^N complex numbers, which is out of reach for modern computers for $N > 50$. (A laptop can handle $N \approx 25$ qubits, and a small computer cluster $N \approx 30$ on a single core; parallelization is needed beyond $N = 30$.) In a gate-based QC, the wave function evolves at each cycle (or step) via a unitary transformation, and it can be computed exactly with a classical algorithm by applying single- and two-qubit gates as 2×2 unitary matrices and updating pairs of amplitudes in place[4,5]. From the exact wave function at step s, one can obtain the N-qubit probability distribution $\tilde{P}_s(x) = |\langle x|\Psi_s\rangle|^2$, which can be measured experimentally but is exponentially hard to compute[3].

Here we take a different approach and focus on the evolution of each individual qubit as a result of mean-field interactions with single- and two-qubit gates. We define the *single-qubit probability (SQP)* for qubit n, with values between 0 and 1, as the probability of measuring qubit n in the excited state $|1\rangle$ at step s, regardless of the state of the other qubits:

$$p_s^{(n)} = \sum_{\{i_q,\, q \neq n\}} |\langle i_1, i_2, \ldots, i_n = 1, \ldots, i_N | \Psi_s \rangle|^2. \tag{2}$$

The *exact* SQPs are marginals of the N-qubit probability distribution $\tilde{P}_s(x)$, and are also exponentially hard to compute because they require knowledge of the QC wave function. We define the SQP vector at step s, $\mathbf{p}_s = (p^{(1)}, p^{(2)}, \ldots, p^{(N)})_s$, as the set of SQPs for all qubits in the QC. Note that the SQP vector has N components, and thus it can be stored with memory resources linear in qubit number N. Experimentally, the SQPs can be accessed by measuring the state of each single qubit.

We model the evolution of the SQP vector $\mathbf{p}_s$ under the effect of single- and two-qubit gates, using an approximate mean-field approach inspired by DFT. In a general QC, single- and two-qubit gates are applied to a set of qubits at each step s. As a result, the SQP vector evolves to a new value at step $s+1$:

$$\mathbf{p}_{s+1} = f_{\mathrm{G}}(\mathbf{p}_s) \tag{3}$$

where we define the map f_{G} as the exact gate functional. Here we derive approximate gate functionals which evolve independently the SQPs of qubits acted on by single-qubit gates, and couple qubits acted on by two-qubit gates (here, CZ and CNOT). Analogous to DFT, where the electron interactions depend on the density, here we derive gate-qubit interactions that depend only on the SQPs, and use them to evolve the SQP vector. Recall that p is the probability of measuring a single qubit in state $|1\rangle$. We define a single-qubit mean-field state consistent with this SQP:

$$|p\pm\rangle = \sqrt{1-p}\ |0\rangle \pm \sqrt{p}\ |1\rangle \tag{4}$$

where we use $\pm$ to take into account two opposite phases between the $|0\rangle$ and $|1\rangle$ states.

For single-qubit gates, we apply the gate U to this mean-field state, and then compute the probability of measuring $|1\rangle$ while taking the phase average over the $\pm$ states. This approach provides explicit rules to update the SQPs at each step:

$$p_{s+1} = \frac{1}{2} \sum_{\pm} |\langle 1|U|p_s \pm\rangle|^2 \,. \tag{5}$$

This equation defines the local-probability approximation (LPA) gate functional. Using eq 5, we derive the following LPA update rules for common single-qubit gates:

$$\begin{aligned} &\text{Pauli X and Y:} \quad p_{s+1} = 1 - p_s \\ &\text{Pauli Z, S and T:} \quad p_{s+1} = p_s \\ &\text{H, } \sqrt{X} \text{ and } \sqrt{Y}\text{:} \quad p_{s+1} = 0.5. \end{aligned} \tag{6}$$

These results imply that the Pauli X and Y gates flip the SQP, the Pauli Z, S and T gates leave the SQP unchanged as they act only on the phase, and the Hadamard, Pauli $\sqrt{X}$ and $\sqrt{Y}$ gates set the SQP to 1/2.

For the two-qubit gates considered here, CZ and CNOT, we use our intuition combined with the LPA rules to approximate the SQP evolution. The probability $p^{(c)}$ of the control qubit is left unchanged, while the probability $p^{(t)}$ of the target qubit is evolved according to:

$$\begin{aligned} &\text{CZ:} \qquad p^{(t)}_{s+1} = p^{(t)}_s \\ &\text{CNOT: if } \; p^{(c)} < 0.5\,, \quad p^{(t)}_{s+1} = p^{(t)}_s \\ &\qquad\quad \text{if } p^{(c)} > 0.5\,, \quad p^{(t)}_{s+1} = 1 - p^{(t)}_s \;. \end{aligned} \tag{7}$$

The CZ result follows from the LPA rule for the Pauli Z gate, which leaves the SQP unchanged. The CNOT gate uses a $p^{(c)} = 0.5$ threshold for controlling the target qubit, but the case $p^{(c)} = 0.5$ is more subtle and needs a separate update rule:

$$\text{CNOT: if } \; p^{(c)} = 0.5 \text{ and } p^{(t)} = 0 \text{ or } 1\,, \quad p^{(t)}_{s+1} = 0.5 \qquad (8)$$
$$\text{if } \; p^{(c)} = 0.5 \text{ and } p^{(t)} \neq 0 \text{ or } 1\,, \quad p^{(t)}_{s+1} = 1 - p^{(t)}_{s} \;.$$

This choice allows us to address the important case of a Hadamard gate acting on the control qubit of a CNOT gate, as in the Bell-state preparation QC[27], a key building block in the random QCs discussed below. In particular, our CNOT and Hadamard update rules give the correct SQPs for all possible two-qubit initial basis states in the Bell-state preparation QC[27] (see Table 1).

Table 1 | LPA functional applied to the Bell-state preparation QC. Exact wave function Ψ_s, and the corresponding SQP vector $\mathbf{p}_s$, given as a function of step s for the Bell-state preparation two-qubit QC[27]. This QC consists of H applied to qubit 0 (step 1) followed by CNOT with control qubit 0 and target qubit 1 (step 2). As one can verify, the LPA rules in eqs 5−7 give the same SQPs as the exact ones shown in the table, at all steps and for all initial states in the computational basis.

Initial state $\Psi_{s=0}$	$\mathbf{p}_{s=0}$	$\Psi_{s=1}$	$\mathbf{p}_{s=1}$	$\Psi_{s=2}$	$\mathbf{p}_{s=2}$
$\lvert 00\rangle$	(0,0)	$\frac{1}{\sqrt{2}}(\lvert 00\rangle + \lvert 10\rangle)$	(0.5, 0)	$\frac{1}{\sqrt{2}}(\lvert 00\rangle + \lvert 11\rangle)$	(0.5, 0.5)
$\lvert 01\rangle$	(0,1)	$\frac{1}{\sqrt{2}}(\lvert 01\rangle + \lvert 11\rangle)$	(0.5, 1)	$\frac{1}{\sqrt{2}}(\lvert 01\rangle + \lvert 10\rangle)$	(0.5, 0.5)
$\lvert 10\rangle$	(1,0)	$\frac{1}{\sqrt{2}}(\lvert 00\rangle - \lvert 10\rangle)$	(0.5, 0)	$\frac{1}{\sqrt{2}}(\lvert 00\rangle - \lvert 11\rangle)$	(0.5, 0.5)
$\lvert 11\rangle$	(1,1)	$\frac{1}{\sqrt{2}}(\lvert 01\rangle - \lvert 11\rangle)$	(0.5, 1)	$\frac{1}{\sqrt{2}}(\lvert 01\rangle - \lvert 10\rangle)$	(0.5, 0.5)

We implement these QC-DFT simulations using an in-house code (see Data Availability), and apply them to random and nonrandom QCs ranging from small to large, with up to a billion interacting qubits. For small QCs with less than ∼30 qubits, we compare the approximate SQPs with exact values obtained from wave-function (also known as state-vector) QC simulations carried out using the QuEST code[6] (see Appendix). For this comparison, we define the *SQP accuracy* A_s as the fraction of SQPs predicted correctly by QC-DFT at step s (equivalently, the SQP error $1 - A_s$ is the fraction of qubits with incorrect SQPs).

3. Random QC Simulations with the LPA Functional

First, we discuss results for random QCs with a universal Clifford+T gate set [28], a moderate depth (20 steps), and QC sizes ranging from 20 to 32 qubits. In these QCs, at each step half of the qubits, chosen at random, are acted on by a randomly chosen single-qubit gate in the set, while the other half are acted on by CNOT gates that couple randomly-selected control and target qubits (Fig. 1a). In the exact simulations, the state-vector of the QC is initially set to $|00\ldots0\rangle_N$ and then evolved according to the gate sequence, while the exact SQPs are computed at each step via eq 2. The exact SQPs, with initial value of $\mathbf{p} = 0$, evolve nontrivially for $10-15$ steps, after which in our random QCs they reach a fully randomized value of $\mathbf{p} = 0.5$. Our approximate QC-DFT simulations aim to capture the nontrivial SQP dynamics in the first $10-15$ steps before randomization occurs.

The results of our QC-DFT simulations for these Clifford+T QCs are shown in Fig. 1a. We find that our approach can predict the SQPs with an accuracy greater than ~90% at all steps. The highest error occurs near steps $5-7$, where the qubits become nontrivially correlated, and then decreases to zero when the QC becomes fully randomized, with all SQPs trivially equal to 0.5. Throughout the dynamics, the exact SQP values for most qubits are 0, 0.5, or 1 due to the combined action of the Hadamard and CNOT gates, and thus the main challenge for the approximate simulations is to capture the transitions between these values, as further discussed below. We also simulate Clifford+T QCs with a different structure, which alternates one step where all qubits are acted on by randomly chosen single-qubit gates, and one step where all qubits are acted on by CNOT gates with randomly selected control and target qubits. The results, given in Fig. S1 in the Supplementary Information, show a similar SQP accuracy of 90% or higher at all steps.

To demonstrate the versatility of our approach, we also simulate a different family of random QCs, introduced by Boixo et al. [29], which use a T, Pauli $\sqrt{X}$ and $\sqrt{Y}$, and CZ gate set (Fig. 1b). (We removed the Hadamard gate layer from the QCs in Ref. [29] because it would make the SQPs trivially equal to 0.5 at all steps). The initial conditions are the

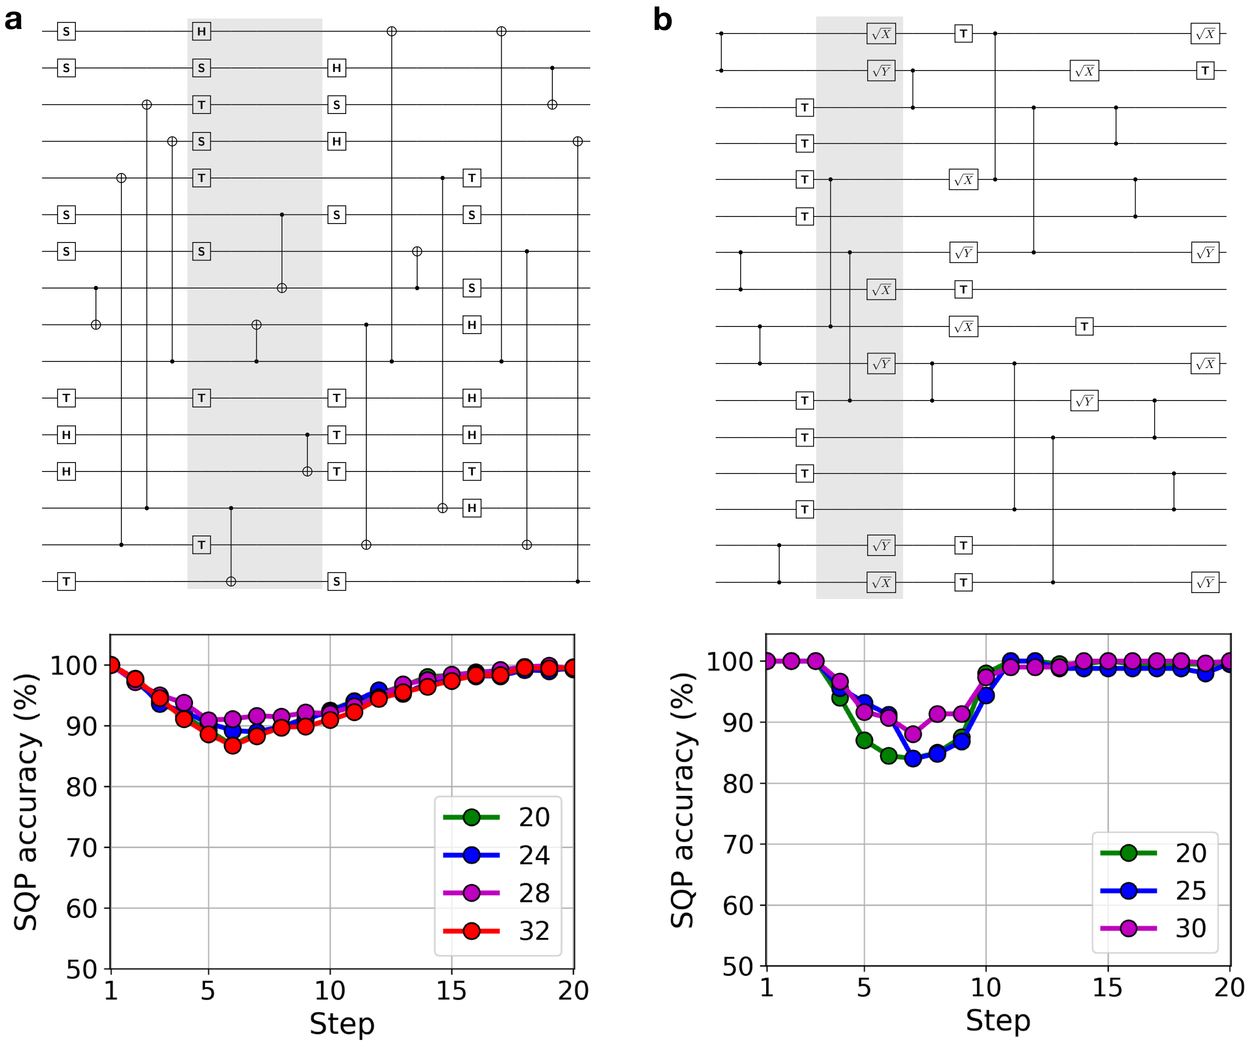

Figure 1 | QC-DFT simulations of random quantum circuits using LPA gate functionals. a, Random QC using a universal Clifford+T gate set. Each step consists of randomly chosen single-qubit gates applied to half of the qubits and CNOT gates applied to the other half (top). The accuracy of the simulated SQPs at each step is shown for this type of random QCs with different numbers of qubits (bottom). **b,** Random QC with T, $\sqrt{X}$, $\sqrt{Y}$ and CZ gates, taken from Ref.[29] but with the H gates removed. Similar to (**a**), we plot the SQP accuracy for different QC sizes (bottom). Both types of QCs have a depth of 20 steps, with a single step shown in shaded gray. The SQPs are obtained by averaging results from 20 distinct random QC instances in (**a**) and 10 instances in (**b**).

same as in the Clifford+T circuits discussed above, but the quantum dynamics is richer, with more possible SQP values than in the Clifford+T case due to the combined effects of the T and square-root Pauli gates. Despite this greater complexity, our QC-DFT approach can simulate the dynamics of these QCs with an SQP accuracy greater than 85−90% for sizes ranging from 20 to 30 qubits (Fig. 1b). These results demonstrate that our QC-DFT approach, combined with the LPA rules, can accurately predict the SQPs for various random QCs without computing the exponentially complex QC wave function.

4. Improved Functionals: Multi-Gate Approximation

We study whether the accuracy of QC-DFT can be improved by fine-tuning the gate functionals, in a spirit similar to improving electronic exchange-correlation functionals in DFT [24]. Many SQP errors in the LPA simulations derive from applying two consecutive times the same gate to a given qubit − a situation analogous to a strong local interaction in the many-electron problem − or from specific gate sequences acting on a qubit. Because our LPA focuses on local gate-qubit interactions at the current step, it lacks memory effects and cannot accurately describe such multi-gate correlations. To address this problem, we formulate multi-gate approximation (MGA) functionals encoding the effects of gate sequences, and apply them as a *correction* to the LPA in the first ∼10 steps, where multi-gate correlations are important for predicting the SQP dynamics in our random QCs.

We define MGA-n gate-functionals which treat explicitly single-qubit gate sequences with length $l \leq n$. Their SQP update rules can be written as

$$p_{s+1} = |\langle 1| \prod_{i=s-l}^{s} U_i |0\rangle|^2, \tag{9}$$

where U_i is a single-qubit gate acting at step i. This approach captures the effect of sequences of l gates, from step $s-l$ to the current step s, and focuses on early multi-gate corrections in the QC by assuming that the gate sequences act on the initial single-qubit state $|0\rangle$. When using these MGA-n gate-functionals, the SQPs are evolved at each step using the LPA, but gate-sequences with length up to n are searched at each step; if a sequence included in the functional is found, then the SQP is updated according to eq 9. For example, for an MGA functional encoding a sequence of two Hadamard gates U_H, the first gate gives $p_s = 0.5$ due to the LPA rules, and the second gives $p_{s+1} = \langle 1|U_H^2|0\rangle = 0$, thus correcting the erroneous LPA value $p_{s+1} = 0.5$. Similarly, an MGA treating a sequence of two square-root of Pauli X gates gives $p_s = 0.5$ after the first and $p_{s+1} = \langle 1|(\sqrt{\sigma_X})^2|0\rangle = 1$ after the second $\sqrt{X}$ gate.

We develop several MGA functionals (see Appendix) aimed at improving the SQP accuracy by addressing the shortcomings of the LPA in our random QCs. For the Clifford+T QCs in Fig. 1a, our analysis of the LPA results reveals that two main gate sequences lead to SQP errors: the H−H sequence consisting of two consecutive Hadamard gates applied to the same qubit, which leads to $p_{s+1} = 0.5$ in the LPA instead of the exact $p_{s+1} = 0$, and the three-gate sequence H−T−H, which gives $p_{s+1} = 0.5$ in the LPA instead of the exact result $p_{s+1} = 0.146447$. Accordingly, we develop a simple MGA-3 functional addressing these two sequences, and apply it to our Clifford+T random QCs. Figure 2a compares the accuracy of this MGA-3 functional with the LPA for the random Clifford+T QCs in Fig. 1a. We apply the multi-gate corrections in the first 7 steps, and find a significant improvement of SQP accuracy, by roughly 5−8%, during those steps. Beyond step ∼10, the QC state randomizes and the SQP accuracy becomes nearly identical for the two functionals. We find similar accuracy improvements for Clifford+T random QCs with a different structure, as shown in Fig. S2 in Supplementary Information.

For the second family of random QCs discussed above, which employ T, Pauli $\sqrt{X}$ and $\sqrt{Y}$, and CZ gates, we develop two types of MGA functionals: MGA-2 addressing only sequences of two consecutive $\sqrt{X}$ or $\sqrt{Y}$ Pauli gates, and a systematically improved MGA-6 functional encoding sequences of up to six single- and two-qubit gates. Both of these MGA functionals lead to accuracy improvements over the LPA, with the MGA-6 further improving over the simpler MGA-2 (Fig. 2b). For both types of random QCs studied here, we analyze the SQP dynamics for selected qubits (bottom panels in Fig. 2). We find that multi-gate corrections can have several different effects: the MGA can leave the LPA results unchanged, correct SQP errors in the LPA, fail to correct the LPA errors, or occasionally introduce errors not present in the LPA. When the MGA correction is successful, the SQP accuracy improve-

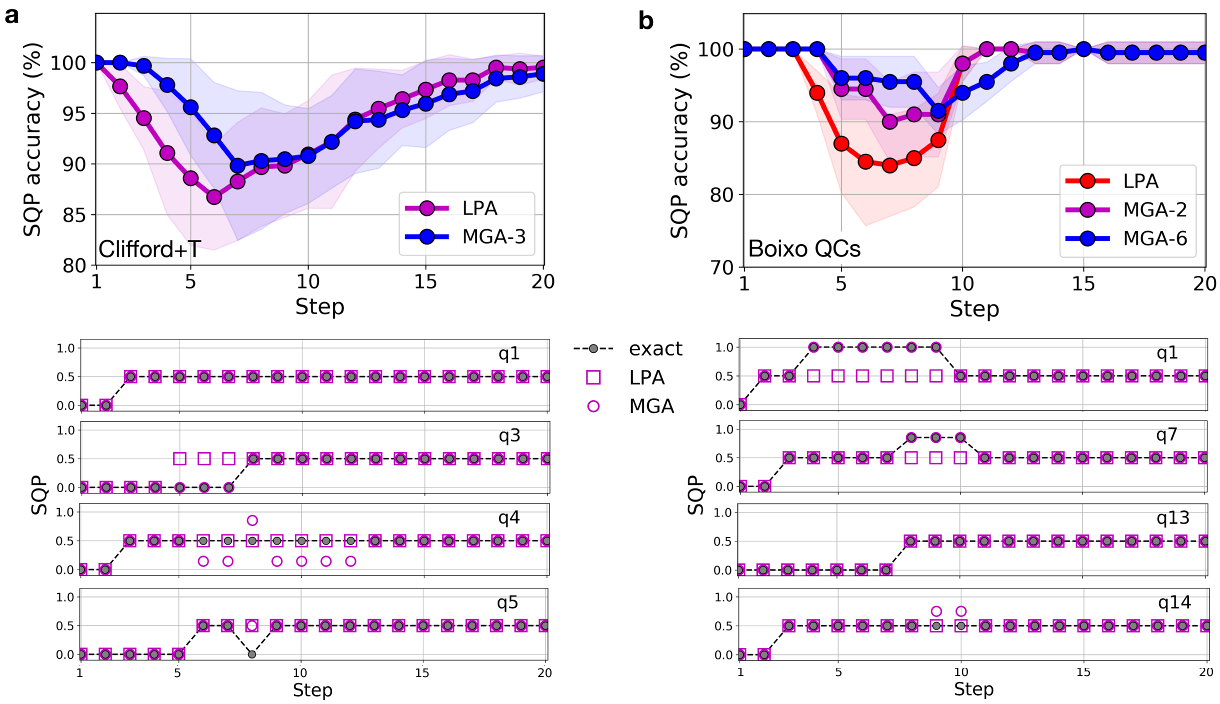

Figure 2 | Optimized MGA gate functionals. a, Accuracy comparison between the LPA and MGA-3 functionals applied to the Clifford+T circuits in Fig. 1a (top). The SQPs at each step are plotted for selected qubits for both functionals and compared with exact results (bottom). **b,** Accuracy comparison between the LPA and two different MGAs, MGA-2 and MGA-6, encoding respectively up to two- and six-gate sequences (top), shown together with the SQPs at each step for selected qubits (bottom). The results in (**a**) are for QCs with 32 qubits and in (**b**) for QCs with 20 qubits. The standard deviation of the SQP accuracy is shown for each curve using shaded colors. These results are obtained by averaging over the same number of QCs as in Fig. 1.

ment typically lingers for several steps, leading to sizable accuracy improvements relative to the LPA. In addition, the multi-gate corrections allow us to capture SQP values different from from 0, 0.5 and 1, the only possible values in the LPA. These results demonstrate a systematic approach for improving QC-DFT gate functionals by explicitly addressing multi-gate correlations.

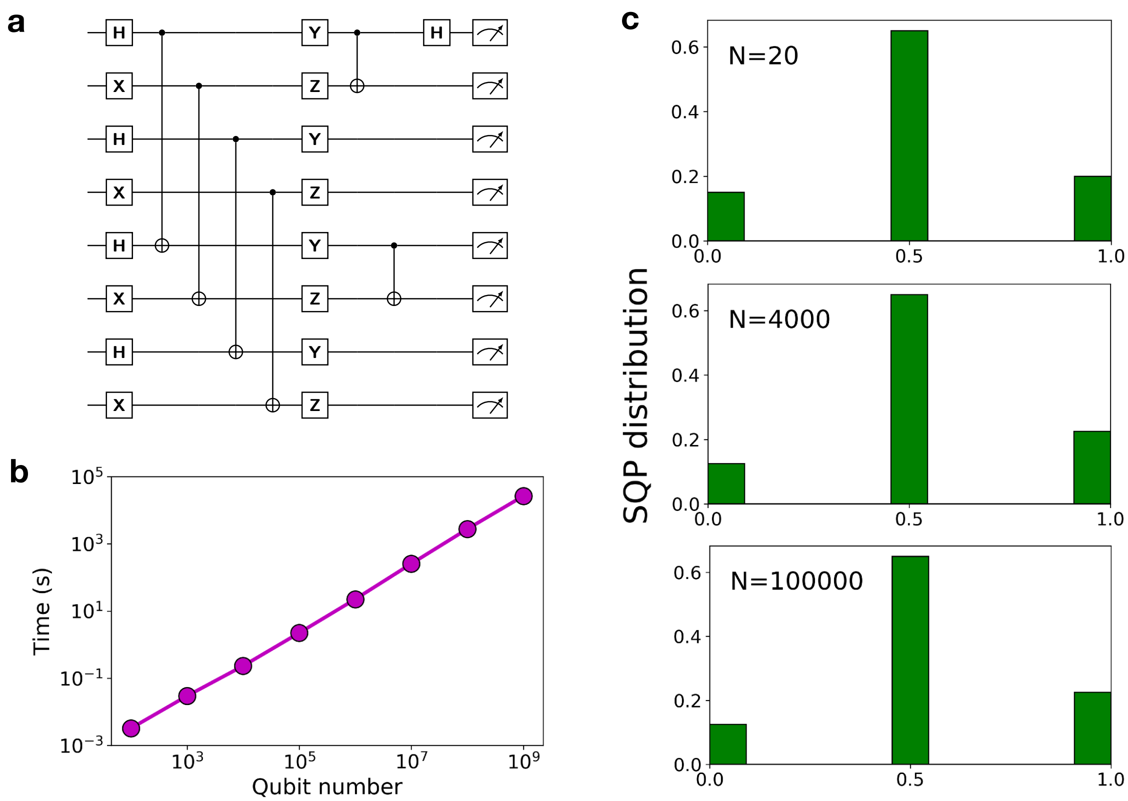

Figure 3 | Computational cost and SQP distribution scaling with quantum circuit size. a, Quantum circuit structure used to obtain the computational cost and SQP distribution as a function of QC size. The rules used to generate this type of QC are given in Appendix. **b,** Linear scaling of QC-DFT computation time with qubit number. **c,** Invariance of the SQP distribution with respect to QC size for QCs with the same structure, which is given in (**a**). Results are shown for QCs with sizes of 20, 4000, and 10^5 qubits.

5. Large QCs and SQP Scaling

The favorable scaling of our approach allows us to simulate very large QCs. We focus on a family of nonrandom QCs (with Hadamard, Pauli X, Y, Z, and CNOT gates), generated with a set of deterministic rules given in Appendix, whose circuit diagram for an example size of 8 qubits is shown in Fig. 3a. Using this class of QCs, with sizes ranging from 20 to 10^9 qubits, we demonstrate that the computational cost of QC-DFT has a linear scaling with number of qubits in the QC (Fig. 3b). We are able to complete the largest calculation, with size one billion qubits, using only a laptop computer for a few hours. Note that both the memory and computational cost to obtain accurate SQPs scale linearly with QC size and depth in our

QC-DFT approach, in clear contrast with the *exact* SQPs from state-vector simulations[6], for which memory and computational cost scale exponentially with qubit number.

Our analysis of the nonrandom QCs in Fig. 3a further reveals an intriguing physical result: for QCs with a given structure, the SQP distribution is independent of QC size, and thus is scale-invariant with respect to qubit number, as shown in Fig. 3c for three illustrative QC sizes. (Although the simulations for large N values cannot be validated against exact results, we have verified that simulations for $N < 30$ qubits achieve a 90% SQP accuracy, similar to the other LPA results shown in this work.) This finding shows that the SQP distribution is a fingerprint of the QC linked to its structure, a result reminiscent of the map between the electron density and the material structure shown in the celebrated Hohenberg-Kohn theorems of DFT[17]. Our analysis suggests that the SQP is a central quantity in mean-field simulations of QCs − similar to the electron density in DFT, which is also a one-body marginal probability − justifying the focus on SQPs in our approach.

6. Model Spin Hamiltonian

We conclude with an application of QC-DFT to obtain the ground-state energy of a simple spin Hamiltonian. Extending the variational quantum eigensolver (VQE)[30,31] discussion in Ref.[27], we model the N-qubit Hamiltonian $H = \bigotimes_{j=1}^{N} \sigma_z^{(j)}$, where $\sigma_z^{(j)}$ is the Pauli Z gate acting on qubit j, and search the energy minimum starting from the trial wave function $|\Psi(\boldsymbol{\theta})\rangle = \prod_{j=1}^{N} R_x^{(j)}(\theta_j)\,|00\ldots0\rangle_N$, where $R_x(\theta_j)$ is a rotation through angle θ_j around the x-axis applied to qubit j (starting from an initial state $|00\ldots0\rangle_N$), and $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_N)$ is the set of rotation angles parametrizing the wave function. For this example Hamiltonian, the energy $E(\boldsymbol{\theta})$ can be obtained analytically: the rotation $R_x(\theta)$ acting on each qubit gives a state $|\varphi(\theta)\rangle$ written as

$$R_x(\theta)\,|0\rangle \equiv |\varphi(\theta)\rangle = \cos(\theta/2)\,|0\rangle - i\,\sin(\theta/2)\,|1\rangle\,, \tag{10}$$

and thus we obtain:

$$E(\boldsymbol{\theta}) = \langle\Psi(\boldsymbol{\theta})|H|\Psi(\boldsymbol{\theta})\rangle = \prod_{j=1}^{N}\langle\varphi(\theta_j)|\sigma_z^{(j)}|\varphi(\theta_j)\rangle = \prod_{j=1}^{N}[\cos^2(\theta_j/2) - \sin^2(\theta_j/2)] = \prod_{j=1}^{N}\cos(\theta_j). \tag{11}$$

In a state-vector simulation, preparing the trial wave function $|\Psi(\boldsymbol{\theta})\rangle$ and computing the associated energy $E(\boldsymbol{\theta})$ for any set of angles $\boldsymbol{\theta}$ requires the application of N rotations about the x-axis, with a computational cost growing exponentially with qubit number N. Therefore, the search for the ground-state energy with state-vector simulations would require exponential resources.

Here we employ QC-DFT as an alternate route for efficient energy calculations. Using the LPA gate functional, the update rule for the $R_x(\theta)$ rotation is

$$p_{s+1} = \frac{1}{2}\sum_{\pm}\langle 1|R_x(\theta)|p_{s\,\pm}\rangle = p_s\cos^2(\theta/2) + (1-p_s)\sin^2(\theta/2)\,, \tag{12}$$

which becomes $p_{s+1} = \sin^2(\theta/2)$ for our initial state with $p_s = 0$. Using as trial wave function the resulting mean-field state obtained with eq 4, $|p_{s+1}\rangle = \cos(\theta/2)\,|0\rangle \pm \sin(\theta/2)\,|1\rangle$, we write the mean-field energy for a single-qubit Hamiltonian σ_z as:

$$\varepsilon_{\mathrm{MF}}(\theta, p_{s+1}) = \langle p_{s+1}|\sigma_z|p_{s+1}\rangle = \cos^2(\theta/2) - \sin^2(\theta/2) = \cos(\theta), \tag{13}$$

where $\varepsilon_{\mathrm{MF}}$ depends explicitly on p_{s+1}, the SQP obtained after applying the $R_x(\theta)$ rotation in QC-DFT. This result can be extended to N qubits, by applying rotations $R_x^{(j)}(\theta_j)$ to each qubit j to obtain the N-qubit mean-field state $|\mathbf{p}_{s+1}\rangle = |p_{s+1}^{(1)}, p_{s+1}^{(2)}, \ldots p_{s+1}^{(N)}\rangle$, with $|p_{s+1}^{(j)}\rangle = \cos(\theta_j/2)\,|0\rangle_j \pm \sin(\theta_j/2)\,|1\rangle_j$ as above. For our Hamiltonian $H = \bigotimes_{j=1}^{N}\sigma_z^{(j)}$, the N-qubit mean-field energy E_{MF} factors into a product of single-qubit mean-field energies, and depends explicitly on the SQPs, $\mathbf{p}_{s+1}$, obtained after applying the x-axis rotation gates

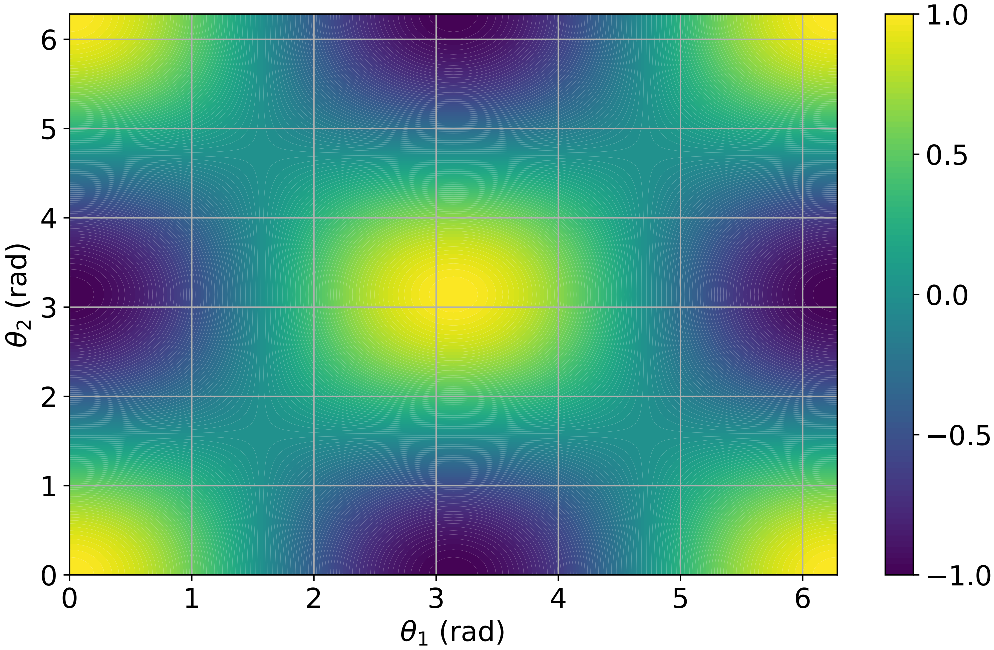

Figure 4 | Energy calculation with QC-DFT. Energy $E(\theta_1, \theta_2)$, color-coded in arbitrary units, for the two-qubit Hamiltonian $H = \sigma_z^{(1)}\sigma_z^{(2)}$ and the trial wave function $\Psi(\theta_1, \theta_2) = R_x^{(1)}(\theta_1)\, R_x^{(2)}(\theta_2)\, |00\rangle$. The energy is computed with QC-DFT on a fine grid of angles (θ_1, θ_2) using eq (14) together with the LPA gate functional for the rotations R_x.

in QC-DFT:

$$E_{\mathrm{MF}}(\boldsymbol{\theta}, \mathbf{p}_{s+1}) = \langle \mathbf{p}_{s+1} | H | \mathbf{p}_{s+1} \rangle = \prod_{j=1}^{N} \varepsilon_{\mathrm{MF}}^{(j)} = \prod_{j=1}^{N} \cos(\theta_j). \tag{14}$$

This mean-field energy is identical to the exact analytic result in eq 11, and it can be obtained in QC-DFT directly from the SQPs, without preparing the trial wave function. Figure 4 shows the energy $E(\theta_1, \theta_2)$ for the two-qubit case computed with QC-DFT on a fine grid of rotation angles θ_1 and θ_2. In this case, since $E(\theta_1, \theta_2) = \cos(\theta_1)\cos(\theta_2)$, the minima are found for $(\theta_1, \theta_2) = (\pi, 0)$ and $(0, \pi)$ (see Fig. 4). For the general case with N qubits, the energy $E(\boldsymbol{\theta})$ can be computed in QC-DFT using only $\mathcal{O}(N)$ memory and computational resources, by applying the x-axis rotation gates within the LPA. Conversely, the same calculation has $\mathcal{O}(2^N)$ memory and computational cost in state-vector simulations. Although the example examined here has a simple analytic solution, it illustrates the point that QC-DFT may enable efficient mean-field calculations of ground-state energies.

6. Discussion

It is important to understand the limitations of the proposed QC-DFT approach. At present, our mean-field method cannot capture qubit phase and interference effects, which are essential in many quantum algorithms. Therefore this method cannot compete with more established techniques such as tensor networks. Satisfactory results are expected mainly in QCs with low entanglement. For example, we have verified explicitly that LPA simulations lead to inaccurate SQP predictions for the Deutsch-Jozsa, Bernstein-Vazirani, and Grover's search quantum algorithms[27]. Extensions of the QC-DFT formalism using one- or two-qubit reduced density matrices (RDMs) (instead of the SQPs), inspired by recent advances in electronic structure using RDMs[32,33], may enable an improved description of gates such as S, T, Pauli Z and controlled-Z, which act on the qubit phase and are currently ignored in our SQP-based approach. Preliminary results for QC-DFT using RDMs, to be presented elsewhere, show improved accuracy for some of the quantum algorithms mentioned above.

Our formulation of a DFT analog for QC simulations motivates several future research directions, including using machine learning to improve the QC-DFT gate functionals, as shown recently for exchange-correlation functionals in DFT[34,35], and applying QC-DFT and its future extensions to spin Hamiltonians and quantum algorithms.

7. Conclusion

In summary, we demonstrated mean-field simulations of QCs inspired by DFT. The approach shown in this work, called QC-DFT, can accurately predict the SQPs − marginals of the full QC probability distribution − with low computational cost (despite their formal exponential scaling) for various random and nonrandom QCs. Although the current approach is not generic and is limited to QCs with low entanglement, improvements to this formalism based on one- and two-qubit RDMs may enable simulations of more general classes of QCs.

Appendix A. Numerical Methods

Exact QC simulations. The exact QC simulations are carried out using the QuEST code[6]. All single- and two-qubit gates are used as provided in the code. We use appropriate rotation operations to implement the square root Pauli gates, and compute the exact SQPs from the state vector. Example input files for QuEST are available in the data sets accompanying this manuscript.

Multi-gate functionals. The MGA-n functionals are implemented in our QC-DFT code by looking for specific gate sequences in the QC. If a gate sequence encoded in the functional is found within cycle $s_{\max}$, the SQPs are updated using p_{s+1} from eq 9. These multi-gate corrections are applied only up to once for each qubit. For the Clifford+T QCs, the MGA-3 functional used in Fig. 2a corrects for the gate sequences H-H (using $p_{s+1} = 0$) and H-T-H ($p_{s+1} = 0.146447$) up to cycle $s_{\max} = 7$, including cases where CNOT gates act on the qubit within these sequences. This means that CNOT control and target operations are ignored when looking for these gate sequences − for example, the gate sequence H−CNOT−H acting on a qubit is treated as H−H and corrected. For the QCs in Fig. 2b, the MGA-2 functional corrects for the $\sqrt{X}$−$\sqrt{X}$ and $\sqrt{Y}$−$\sqrt{Y}$ sequences up to cycle $s_{\max} = 7$. In this case, the CZ control operations are ignored when looking for gate sequences, while CZ target operations are taken into account. For example, if the gate sequence $\sqrt{X}$−CZ−$\sqrt{X}$ involves the CZ control qubit, it is treated as $\sqrt{X}$−$\sqrt{X}$ and the multi-gate correction is applied. If the same sequence is found for the CZ target qubit, the multi-gate correction is not applied.

The MGA-6 functional includes several multi-gate corrections with up to 6-gate sequences, applied up to cycle $s_{\max}$ between 6 and 10 depending on the sequence. Some sequences take into account CZ gates, while others ignore them. Next we provide the full list of gate-sequences for our MGA-6 functional, using a naming convention for gate sequences where, for a given qubit, the rightmost gate acts at the current step, and the leftmost gate acts at the earliest step in the sequence; steps where no gates act on the qubit are ignored. This means that sequences are given in the same order as when reading the QC from left to right, ignoring steps with no gates. The CZ gates are explicitly taken into account, in the same way for control and target qubits, unless otherwise stated. With these conventions, the gate sequences treated in our MGA-6 functional are as follows: 2-gate sequences $\sqrt{X}-\sqrt{X}$ and $\sqrt{Y}-\sqrt{Y}$ (both with $p_{s+1}=1$ and $s_{\max}=6$); 3-gate sequences $\sqrt{X}-\mathrm{CZ}-\sqrt{X}$ and $\sqrt{Y}-\mathrm{CZ}-\sqrt{Y}$ ($p_{s+1}=1,\ s_{\max}=6$), $\sqrt{Y}-\mathrm{T}-\sqrt{X}$ ($p_{s+1}=0.14645,\ s_{\max}=8$, CZ gates ignored) and $\sqrt{Y}-\mathrm{T}-\sqrt{Y}$, $\sqrt{X}-\mathrm{T}-\sqrt{X}$, and $\sqrt{X}-\mathrm{T}-\sqrt{Y}$ ($p_{s+1}=0.85355,\ s_{\max}=8$, CZ gates ignored); 4-gate sequences $\mathrm{T}-\sqrt{Y}-\mathrm{T}-\sqrt{X}$ and $\mathrm{T}-\sqrt{X}-\mathrm{T}-\sqrt{Y}$ ($p_{s+1}=0.75,\ s_{\max}=10$, CZ gates ignored); 5-gate sequences $\mathrm{CZ}-\sqrt{X}-\mathrm{T}-\mathrm{CZ}-\sqrt{X}$, $\mathrm{CZ}-\sqrt{X}-\mathrm{T}-\mathrm{CZ}-\sqrt{Y}$, $\mathrm{CZ}-\sqrt{Y}-\mathrm{T}-\mathrm{CZ}-\sqrt{X}$, and $\mathrm{CZ}-\sqrt{Y}-\mathrm{T}-\mathrm{CZ}-\sqrt{Y}$ ($p_{s+1}=0.5,\ s_{\max}=10$); 6-gate sequences $\sqrt{Y}-\mathrm{CZ}-\sqrt{X}-\mathrm{T}-\mathrm{CZ}-\sqrt{X}$ ($p_{s+1}=0.14645,\ s_{\max}=10$), $\sqrt{X}-\mathrm{CZ}-\sqrt{Y}-\mathrm{T}-\mathrm{CZ}-\sqrt{X}$ and $\sqrt{X}-\mathrm{CZ}-\sqrt{Y}-\mathrm{T}-\mathrm{CZ}-\sqrt{Y}$ ($p_{s+1}=0.85355,\ s_{\max}=10$). These sequences can be found in the MGA-n QC-DFT codes provided in the data sets accompanying this manuscript.

Nonrandom QCs. The nonrandom QCs used for the scaling calculations in Fig. 3 are generated using deterministic rules. For a QC with size N qubits, step 1 consists of alternating H and Pauli X gates; in step 2, a CNOT gate connects each qubit $i < N/2$ (control) to qubit $i + N/2$ (target); step 3 consists of alternating Pauli Y and Z gates; step 4 has CNOT gates every 4 qubits, each with neighboring control and target qubits; step 5 applies H gates every 10 qubits. Only QCs with size N multiple of 4 and 10 have the same structure, and thus give the same SQP distribution as shown above. Codes for generating these QCs and reproducing the calculations in Fig. 3 are provided in the data sets accompanying this manuscript.

Data Availability

The data sets generated and analyzed in this study, as well as the QC-DFT codes, will be made available in the CaltechDATA repository. Additional data and information are available upon reasonable request. The QuEST code[6] used for the exact QC simulations is an open source software, which can be downloaded at https://quest.qtechtheory.org. The QC drawings were prepared using the Quantikz LaTeX package[36], which can be downloaded at https://ctan.org/pkg/quantikz. The QC-DFT Python code will be made available in the CaltechDATA repository.

Supporting Information

The Supporting Information is available free of charge at [link]. Figure S1, additional LPA simulations of Clifford+T random QCs; Figure S2, additional results for optimized MGA functionals.

Acknowledgements

This work was supported by the National Science Foundation under Grant No. 1750613, which provided for method development. M.B. was also supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of Basic Energy Sciences, Scientific Discovery through Advanced Computing (SciDAC) program under Award Number DE-SC0022088, which supported code development.

1 Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2018**, *2*, 79.

2 Nielsen, M. A.; Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*; Cambridge University Press, 2010.

3 Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **2019**, *574*, 505–510.

4 De Raedt, H.; Jin, F.; Willsch, D.; Willsch, M.; Yoshioka, N.; Ito, N.; Yuan, S.; Michielsen, K. Massively parallel quantum computer simulator, eleven years later. *Comput. Phys. Commun.* **2019**, *237*, 47–61.

5 Smelyanskiy, M.; Sawaya, N. P. D.; Aspuru-Guzik, A. qHiPSTER: The Quantum High Performance Software Testing Environment. *arXiv 1601.07195* **2016**,

6 Jones, T.; Brown, A.; Bush, I.; Benjamin, S. C. QuEST and high performance simulation of quantum computers. *Sci. Rep.* **2019**, *9*, 1–11.

7 Chen, Z.-Y.; Zhou, Q.; Xue, C.; Yang, X.; Guo, G.-C.; Guo, G.-P. 64-qubit quantum circuit simulation. *Sci. Bull.* **2018**, *63*, 964–971.

8 Jozsa, R. On the simulation of quantum circuits. *arXiv quant-ph/0603163* **2006**,

9 Markov, I. L.; Shi, Y. Simulating Quantum Computation by Contracting Tensor Networks. *SIAM J. Comput.* **2008**, *38*, 963–981.

10 Vidal, G. Efficient Classical Simulation of Slightly Entangled Quantum Computations. *Phys. Rev. Lett.* **2003**, *91*, 147902.

11 Yoran, N.; Short, A. J. Classical Simulation of Limited-Width Cluster-State Quantum Computation. *Phys. Rev. Lett.* **2006**, *96*, 170503.

12 Jozsa, R.; Miyake, A. Matchgates and classical simulation of quantum circuits. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **2008**, *464*, 3089–3106.

13 Zhou, Y.; Stoudenmire, E. M.; Waintal, X. What Limits the Simulation of Quantum Computers? *Phys. Rev. X* **2020**, *10*, 041038.

14 Gray, J. quimb: A python package for quantum information and many-body calculations. *J. Open Source Softw.* **2018**, *3*, 819.

15 Jónsson, B.; Bauer, B.; Carleo, G. Neural-network states for the classical simulation of quantum computing. *arXiv 1808.05232* **2018**,

16 Martin, R. M.; Reining, L.; Ceperley, D. M. *Interacting Electrons: Theory and Computational Approaches*; Cambridge University Press, 2016.

17 Martin, R. M. *Electronic Structure: Basic Theory and Practical Methods*; Cambridge University Press, 2020.

18 Car, R.; Parrinello, M. Unified Approach for Molecular Dynamics and Density-Functional Theory. *Phys. Rev. Lett.* **1985**, *55*, 2471–2474.

19 Ceperley, D.; Alder, B. Quantum Monte Carlo. *Science* **1986**, *231*, 555–560.

20 White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **1992**, *69*, 2863–2866.

21 Kent, P. R. C.; Kotliar, G. Toward a predictive theory of correlated materials. *Science* **2018**, *361*, 348–354.

22 Motta, M. et al. Towards the solution of the many-electron problem in real materials: Equation of state of the hydrogen chain with state-of-the-art many-body methods. *Phys. Rev. X* **2017**, *7*, 031059.

23 Williams, K. T.; Yao, Y.; Li, J.; Chen, L.; Shi, H.; Motta, M.; Niu, C.; Ray, U.; Guo, S.; Anderson, R. J.; Authors, M.; Authors, M. Direct comparison of many-body methods for realistic electronic Hamiltonians. *Phys. Rev. X* **2020**, *10*, 011041.

24 Burke, K. Perspective on density functional theory. *J. Chem. Phys.* **2012**, *136*, 150901.

25 Gaitan, F.; Nori, F. Density functional theory and quantum computation. *Phys. Rev. B* **2009**, *79*, 205117.

26 Tempel, D. G.; Aspuru-Guzik, A. Quantum computing without wavefunctions: Time-dependent density functional theory for universal quantum computation. *Sci. Rep.* **2012**, *2*, 391.

27 Hidary, J. D. *Quantum Computing: an Applied Approach*; Springer, 2019.

28 Boykin, P. O.; Mor, T.; Pulver, M.; Roychowdhury, V.; Vatan, F. A new universal and fault-tolerant quantum basis. *Inf. Process. Lett.* **2000**, *75*, 101–107.

29 Boixo, S.; Isakov, S. V.; Smelyanskiy, V. N.; Babbush, R.; Ding, N.; Jiang, Z.; Bremner, M. J.; Martinis, J. M.; Neven, H. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **2018**, *14*, 595–600, The QCs used in this work were taken from the first author's Github account, last accessed in January 2021.

30 Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; O'brien, J. L. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **2014**, *5*, 1–7.

31 Fedorov, D. A.; Peng, B.; Govind, N.; Alexeev, Y. VQE method: A short survey and recent developments. *Mater. Theory* **2022**, *6*, 1–21.

32 Mazziotti, D. A. Two-electron reduced density matrix as the basic variable in many-electron quantum chemistry and physics. *Chem. Rev.* **2012**, *112*, 244–262.

33 Mazziotti, D. A. Anti-Hermitian contracted Schrödinger equation: Direct determination of the two-electron reduced density matrices of many-electron molecules. *Phys. Rev. Lett.* **2006**, *97*, 143002.

34 Snyder, J. C.; Rupp, M.; Hansen, K.; Müller, K.-R.; Burke, K. Finding Density Functionals with Machine Learning. *Phys. Rev. Lett.* **2012**, *108*, 253002.

35 Pederson, R.; Kalita, B.; Burke, K. Machine learning and density functional theory. *Nat. Rev. Phys.* **2022**, *4*, 357.

36 Kay, A. Tutorial on the Quantikz Package. *arXiv 1809.03842* **2020**,