Error Sensitivity to Environmental Noise in Quantum Circuits for Chemical State Preparation

Citation

Sawaya, Nicolas P. D., Mikhail Smelyanskiy, Jarrod R. McClean, and Alán Aspuru-Guzik. 2016. Error Sensitivity to Environmental Noise in Quantum Circuits for Chemical State Preparation. *Journal of Chemical Theory and Computation* 12: 3097–3108.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:42643041>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Error Sensitivity to Environmental Noise in Quantum Circuits for Chemical State Preparation.

Nicolas P. D. Sawaya,^{†,‡} Mikhail Smelyanskiy,[‡] Jarrod R. McClean,[¶] and Alán Aspuru-Guzik^{*,†}

Department of Chemistry and Chemical Biology, Harvard University, Parallel Computing Lab, Intel Corporation, and Computational Research Division, Lawrence Berkeley National Laboratory

E-mail: aspuru@chemistry.harvard.edu

Abstract

Calculating molecular energies is likely to be one of the first useful applications to achieve quantum supremacy, performing faster on a quantum than a classical computer. However, if future quantum devices are to produce accurate calculations, errors due to environmental noise and algorithmic approximations need to be characterized and reduced. In this study, we use the high performance *qHiPSTER* software to investigate the effects of environmental noise on the preparation of quantum chemistry states. We simulated eighteen 16-qubit quantum circuits under environmental noise, each corresponding to a unitary coupled cluster state preparation of a different molecule or molecular configuration. Additionally, we analyze the nature of simple gate errors in noise-free circuits of up to 40 qubits. We find that, in most cases, the Jordan–Wigner (JW) encoding produces smaller errors under a noisy environment as compared to the Bravyi–Kitaev (BK) encoding. For the JW encoding, pure-

dephasing noise is shown to produce substantially smaller errors than pure relaxation noise of the same magnitude. We report error trends in both molecular energy and electron particle number within a unitary coupled cluster state preparation scheme, against changes in nuclear charge, bond length, number of electrons, noise types, and noise magnitude. These trends may prove to be useful in making algorithmic and hardware-related choices for quantum simulation of molecular energies.

1 Introduction

For certain classes of problems, quantum computation promises to provide polynomial or exponential speedups compared to classical computers^{1,2}. Despite the rapid progress in obtaining longer qubit coherence times^{3–5}, current experimental efforts to demonstrate quantum circuits employ pre-threshold quantum information processors^{2,6,7}. Hence it is essential to determine the effects that noise has both on error-correction codes and on specific algorithms. By characterizing and understanding error trends, one can determine which aspects of a particular quantum device or algorithm need to be improved.

Here, we use the recently developed *qHiPSTER* software⁸ to simulate quantum

*To whom correspondence should be addressed

[†]Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA 02138, USA

[‡]Parallel Computing Lab, Intel Corporation, Santa Clara, CA 95054, USA

[¶]Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

state preparation circuits for the molecular electronic structure problem. We use circuits that prepare unitary coupled cluster states of a set of seven diatomic molecules, as well as multiple bond lengths along the dissociation coordinate of the nitrogen molecule. We also study and compare three isoelectronic diatomic species. Finally, we simulate short chemistry-inspired circuits on up to 40 qubits, approximating single-electron excitations from a reference state. In order to determine the effects of decoherence, the molecular quantum circuits are simulated under the presence of environmental noise. Studying the results reveals trends in error rates for molecular energy and electron particle number, which may help guide future algorithmic and hardware developments.

Two classes of quantum algorithms have emerged for calculating chemical energies. The first uses the quantum phase estimation (QPE) algorithm^{9–13} to evolve an initial state toward the molecular ground state. The second approach is the variational quantum eigensolver (VQE)^{14–16}. Similar ideas to the VQE have been independently introduced and experimentally realized within the context of matrix product states^{17,18}.

In VQE, a quantum circuit first prepares a parametrized approximation to the quantum state, which is then used to estimate the energy by direct measurement of the expectation value of the terms of the Hamiltonian. Then, the variationally-optimal state is found by modifying the state parametrization and corresponding preparation circuit after each energy calculation, to find the state(s) which minimize the energy. For chemistry applications, VQE generally requires fewer coherent gate operations per circuit than QPE does, making the VQE approach a more promising candidate for near-term quantum hardware.

In the current study, we focus on the preparation of molecular states using the unitary coupled cluster (UCC) ansatz^{15,19,20}. Quantum chemical state preparation is important in both the VQE approach, where it ultimately determines the state, energy, and properties measured, and in QPE, where it determines the success probability of the energy measure-

ment^{9,21,22}. Hence an understanding of the general effect of errors on state preparation is crucial to the success of both algorithms.

In addition to the choice between QPE and VQE, there are two common methods for mapping the electronic structure problem to a set of qubits: the Jordan–Wigner^{12,23} (JW) and Bravyi–Kitaev^{24,25} (BK) transformations. It has been shown for the hydrogen molecule that the BK mapping requires fewer gate operations than JW, in order to achieve a given chemical precision²⁵. Moreover, the reduced locality of the BK transformed Hamiltonian has many benefits for adiabatic schemes and state preparation methods. Each of these methods seems to display practical benefits and drawbacks that lend them best to different situations. These and other studies have investigated gate ordering, gate cancellation, and the effects of various Trotterization approximations^{26,27}.

The article is organized as follows. In Sec. 2 we present the physical background and methodology used in the study, including the noise model and the method for mapping the molecular electronic structure problem onto a set of qubits. In Sec. 3 we study the effects of environmental noise on the molecular energy of the prepared state. Sec. 4 presents results for the effects of both noise and gate errors on electron number preservation. Finally, Sec. 5 summarizes the implementation and performance details of *qHiPSTER*, before concluding remarks are given in Sec. 6.

2 Background and Methods

2.1 Environmental Noise in Quantum Circuits

Various noise models have been developed to study environmental effects and general errors in noisy quantum systems and nonideal quantum circuits, varying in accuracy and computational complexity. In this section, we briefly introduce the models of noise relevant to our study. These methods can be roughly split into two categories: the direct use of the Schrödinger

or quantum master equations, and the use of heuristic error operators.

In the first category, time is explicitly considered and the Schrödinger equation is used to propagate the state of the system. Highly accurate open quantum systems methods such as the quasiadiabatic propagator path integral method (QUAPI)^{28,29} are too computationally expensive to be used on more than a few qubits. Commonly used master equations such as the Redfield formalism³⁰ propagate the density matrix directly. Finally, state vector methods include the quantum jump formalism³¹ and various forms of the stochastic Schrödinger equation³². All of these methods often make use of a detailed bath spectral density, which describes the frequency-dependent system-bath interactions.

The second noise simulation category involves the use of one or more heuristic parameters that encapsulates the relevant noise effects. This strategy is used when the bath spectral density is unknown or its precise form is unimportant. Especially common in the simulation of quantum circuits for quantum computation, this simulation category is chosen when the goal is to elucidate general trends relating to a specific algorithm in its abstract circuit representation, as opposed to faithfully reproducing the results of a specific hardware setup. The operator-sum representation is often used for this type of noise modeling, in which the density matrix ρ evolves as

$$\rho \rightarrow \mathcal{E}(\rho) = \sum_i E_i \rho E_i^\dagger \quad (1)$$

where E_i are a set of Kraus operators. One example of such a quantum operation is the asymmetric depolarizing channel, defined by

$$\begin{aligned} \rho \rightarrow & (1 - p_x - p_y - p_z)\rho \\ & + p_x X \rho X + p_y Y \rho Y + p_z Z \rho Z \end{aligned} \quad (2)$$

where X , Y , and Z are the Pauli matrices, and the $p_{\{x,y,z\}}$ correspond to the probabilities of a Pauli operator acting on the system during one iteration of the channel. When multiple qubits are considered, Pauli matrices for each qubit are applied independently according to

the formula above.

Other noise channels include the symmetric depolarizing, bit flip, phase flip, amplitude damping, and phase damping channels. The limits of two common noise channels are instructive: the symmetric depolarizing channel converges to a completely mixed state, and the amplitude damping channel causes decay toward the ground state. We note that, though we do not do so in this study, it is also possible include the effects of correlated noise^{33,34}, which can lead to a faster total decay known as superdecoherence.

Determining noise-induced behavior is especially important if a physical implementation of a quantum circuit is to use gates—such as the precise Z rotations in our circuits—which require careful decomposition in an error correction scheme. The trade-off between imperfect arbitrary rotations and error correction schemes, which can be costly both in terms of gates and qubits, needs to be carefully studied in near-term devices where one may have sufficient resources to provide some, but not perfect, error correction³⁵. Additionally, we note that studying noise effects is important even if an implementation were to use error correction, because there will also be an effective decoherence time associated the error-corrected circuit, which may not be infinite on near-term devices.

2.2 Pauli Twirling Approximation

One of the challenging aspects of modeling environmental noise in quantum systems is that the effects become more dramatic as the system size increases. As a result, the eventual aim is to be able to treat the largest number of qubits feasible with given computational resources. In order to scale up our simulations, it is preferable to propagate a state vector instead of the density matrix, because the state vector contains 2^n complex coefficients while the density matrix requires 2^{2n} terms, where n is the number of qubits. This is the strategy we employ in this work, and describe in more detail in this section.

It is possible to map any Lindblad-type mas-

ter equation onto a linear stochastic differential equation (SDE) operating on the state vector^{36–38}. With such a mapping, many stochastic trajectories of the state vector are summed in order to converge to the same physical results as if propagating the full density matrix³⁸. More specifically, here we use the Pauli twirling approximation (PTA), in which a noise channel is mapped onto a set of Pauli operators. Specifically, we use PTA to approximate decoherence by the dephasing and amplitude damping channels. This noise model approximation has been previously used in simulating error-correcting circuits^{39–41}, with the assumption that the circuit time is much shorter than the coherence time. This assumption is allowable since once the circuit time nears the coherence time, the algorithm has likely long ceased to be useful. In order to use PTA for propagating a state vector with a SDE, one draws ν_X , ν_Y , and ν_Z from three independent Gaussian distributions, and performs three “noise gates” in sequence, resulting in a unitary operator

$$U_{noise} = \exp(-i\nu_X X) \exp(-i\nu_Y Y) \times \exp(-i\nu_Z Z) \quad (3)$$

based on the sampled values³⁸.

For small values of ν , these three rotations can be approximated as a single rotation:

$$U_{noise} = R_{\vec{n}}(\theta) \approx \exp(-i\frac{\theta}{2}\vec{n} \cdot \vec{\sigma}) \quad (4)$$

where

$$\vec{n} = \frac{\vec{\nu}}{||\vec{\nu}||} \quad (5)$$

$$\theta = 2||\vec{\nu}|| \quad (6)$$

Independent operators $R_{\vec{n}}(\theta)$ are applied to each qubit after every time step. For a particular qubit, long strings of these noise gates may be fused together for time steps during which there is no logic gate performed on the said qubit. PTA converges with fewer iterations than performing rare probabilistic bit-flips and phase-flips on the state vector, which is one common algorithm used to model noise chan-

nels through state vector propagation⁴².

The Gaussian distributions from which ν_X , ν_Y , and ν_Z are drawn have 0 mean and standard deviations s_x , s_y , and s_z defined by

$$s_\alpha = \sqrt{-\ln(1 - p_\alpha)} \quad (7)$$

with $\alpha \in \{x, y, z\}$ and

$$p_x = p_y = \frac{1 - e^{-1/M_1}}{4} \quad (8)$$

$$p_z = \frac{1 - e^{-1/M_2}}{2} - \frac{1 - e^{-1/M_1}}{4} \quad (9)$$

where we denote $M_\beta = T_\beta/\Delta t$ the “coherence parameter” and $\beta \in \{1, 2, \phi\}$. The time T_ϕ is related to the standard T_1 and T_2 times as

$$\frac{1}{T_\phi} = \frac{1}{T_2} - \frac{1}{2T_1} \quad (10)$$

and Δt is the time step between applications of noise, usually approximated as the average length of a single gate operation. These Gaussian distributions reproduce the T_1 , T_2 , and T_ϕ times from which they are derived, a fact we numerically verified in our implementation.

2.3 Molecular State Preparation

We simulate the quantum chemical state preparations within the Unitary Coupled Cluster (UCC) ansatz. Our molecular state preparation circuits are constructed from the UCC ansatz. The defining equation of the UCC state is given by

$$|\Psi_{CC}^{(k)}(\vec{\xi})\rangle = \exp\left(\sum_k T_k(\vec{\xi})\right) |\Phi_{HF}\rangle \quad (11)$$

where $|\Psi_{HF}\rangle$ is the Hartree–Fock reference state, the vector of state parameters $\vec{\xi}$ are the cluster amplitudes, and k is the order of the expansion. Up to second order, these operators are given by

$$T_1(\vec{\xi}) = \sum_{i_1 p_1} \xi_{i_1 p_1} (a_{i_1}^\dagger a_{p_1} - a_{p_1}^\dagger a_{i_1}) \quad (12)$$

$$T_2(\vec{\xi}) = \sum_{i_1 i_2 p_1 p_2} \xi_{i_1 i_2 p_1 p_2} (a_{i_1}^\dagger a_{p_1} a_{i_2}^\dagger a_{p_2} - a_{p_2}^\dagger a_{i_2} a_{p_1}^\dagger a_{i_1}) \quad (13)$$

where we index spin-orbitals occupied in the Hartree-Fock reference using i_j and p_j for those unoccupied.

In order to decompose Eq. 11 into one- and two-qubit gates, the Trotter-Suzuki decomposition is used, defined as

$$e^{\sum_i B_i} = \left(\prod_i e^{B_i/\eta} \right)^\eta \quad (14)$$

where the operator to be exponentiated is given as a sum of Pauli strings B_i , and η is an integer called the Trotter number. The approximation becomes exact as η approaches infinity. The effects of Trotterization order have been studied elsewhere^{12,25–27}. As the focus of the present study is the effects of decoherence, most of our quantum circuits estimate the molecular energy using a Trotter number of one.

Note that we restrict ourselves to spin-conserving excitations in the above operators (hence we are restricted to the lowest energy state corresponding to the spin of the reference state). This allows us to perform our simulations at reasonable values of the cluster amplitudes separately from the consideration of VQE optimization of these amplitudes. However, we note that the cluster amplitudes obtained by the traditional coupled cluster calculation are expected to differ from those variationally obtained in the UCC ansatz.

To map the chemical problem to qubits from Fermions, the Jordan-Wigner (JW) or Bravyi-Kitaev (BK) mapping is used. The JW mapping uses creation and annihilation operators defined as

$$a_p^\dagger = \left(\prod_{m < p} \sigma_m^z \right) \sigma_p^+ \quad (15)$$

$$a_p = \left(\prod_{m < p} \sigma_m^z \right) \sigma_p^- \quad (16)$$

$$\sigma^\pm \equiv (\sigma^x \mp i\sigma^y)/2 \quad (17)$$

where p is the index of a spin-orbital and

σ_j^α are the Pauli matrices that act on qubit j . The occupation is stored *locally*, i.e. each qubit stores the occupation number of a single spin-orbital. The BK transformation involves a more intricate recursive procedure that is presented elsewhere⁴³. Unlike the JW mapping, which is $O(n)$ -local, the BK mapping yields an $O(\log n)$ -local Hamiltonian, at the expense of the spin-orbital occupancy being $O(\log n)$ -local. It is also notable that parity storage is $O(n)$ -local and $O(\log n)$ -local in JW and BK, respectively.

2.4 Simulation Details

We considered the elemental diatomic molecules of the first row of the periodic table (excluding neon) at their equilibrium geometries, as well as nitrogen at several bond lengths and the isoelectronic series C_2^{2-} , N_2 , and O_2^{2+} .

Using the PTA, we simulated quantum chemical state preparations, within the Unitary Coupled Cluster (UCC) ansatz, under three different noise scenarios: relaxation-dephasing ($T_1 = T_\phi$), pure relaxation ($T_1 \ll T_\phi$), and pure dephasing ($T_\phi \ll T_1$). Note that T_1 and T_ϕ produce the same magnitude of noise (same probability of error) for a given coherence parameter. Note also that there is an energy error and particle number error resulting from our use of only a single Trotter step in approximating the full exponential. However, our aim is to characterize the effects of noise, so all errors are reported as the difference in error between the pristine (noiseless) and noisy circuits. As noted in the introduction, Trotter errors have been analyzed in several previous studies^{12,25–27}.

Nonetheless, we performed one set of calculations to estimate the trade-off between noise error and Trotterization error. In a noise-free circuit, using a Trotter number $\eta=1$ will be less accurate than $\eta=2$. However, the latter case produces a circuit that is twice as long, yielding a greater noise-induced error. We define M_{Trot} as the coherence parameter (eqs. 8 & 9) at which the total error (Trotter error and noise error) is equal for $\eta=1$ and $\eta=2$.

In the present simulations we obtain the values of the cluster amplitudes from traditional

(nonunitary) coupled cluster with single and double excitations (CCSD) from the quantum chemistry package Psi4⁴⁴. We used the minimal STO-3G basis set with frozen $1s^2$ core electrons and a spin-restricted reference state. Cluster amplitudes with magnitude less than 10^{-5} were truncated in the experiments performed in this work. Experimental molecular geometries were taken from the NIST database⁴⁵.

To obtain explicit quantum circuits, the above Fermionic cluster operators were translated into Pauli operators via the JW and BK mappings. These terms were then ordered numerically by qubit index. Finally, this sorted list of Pauli operators was translated into gate sequences of basis changes Hadamard and Yb (basis change to Y), CNOTs, and Z-rotations with the standard circuit synthesis for exponentials of tensor products of Pauli operators². A useful introduction to these circuits for the JW mapping can be found in Whitfield et al.¹²

Figure 1: For pristine (noise-free) circuits, the bars show one standard deviation of the chemical energy estimator (square root of the variance of the estimator). A larger standard deviation implies that a larger number of measurements is required. BK and JW denote the Bravyi–Kitaev and Jordan–Wigner mappings, respectively.

One consideration that is especially important for VQE is the variance of the objective function estimator. Even in a pristine circuit, many measurements are required before converging on the correct answer. As an illustration, in Fig. 1 we show the square root of

the variance for noiseless quantum circuits that simulate the first-row dimers. Hence, in addition to the mean noise-induced error, we will also report results for the noise-induced increase in variance for the chemical energy estimator, as summarized below. For VQE, an increased variance requires a larger number of quantum measurements to arrive at a particular answer. Thus trends in the variance are an important consideration, as this variance determines how many times the quantum circuit must be run.

We define operators

$$\langle \cdot \rangle_p = \langle \psi_p | \cdot | \psi_p \rangle \quad (18)$$

and

$$\langle \cdot \rangle_n = \sum_j \langle \psi_{n,j} | \cdot | \psi_{n,j} \rangle \quad (19)$$

where subscripts p and n denote the state vectors that result from pristine (noise-free) and noisy quantum circuits, respectively, and the index j runs over all noisy iterations of the simulated circuit. The estimator for the chemical energy, determined by a set of quantum measurements, is

$$\widehat{\langle H \rangle} = \sum_\gamma \widehat{\langle H_\gamma \rangle} \quad (20)$$

where each H_γ consists of a Pauli string, O_γ , multiplied by a constant:

$$H_\gamma = w_\gamma \times O_\gamma. \quad (21)$$

Because variances of independent random variables are additive, we can estimate the variance in a pristine quantum circuit as

$$\sigma_p^2 = \text{Var}[\widehat{\langle H_\gamma \rangle}] = \sum_\gamma (\langle H_\gamma^2 \rangle_p - \langle H_\gamma \rangle_p^2) \quad (22)$$

Because O_γ is simply a Pauli string, $\langle O_\gamma^2 \rangle$ is unity. Hence, we use the expression

$$\sigma_p^2 = \sum_\gamma (w_\gamma^2 - \langle H_\gamma \rangle_p^2). \quad (23)$$

Similarly,

$$\sigma_{noisy}^2 = \sum_\gamma (w_\gamma^2 - \langle H_\gamma \rangle_n^2). \quad (24)$$

Finally, we define the noise-induced uncertainty or “error width” as the quantity

$$\sigma_{noisy} - \sigma_p \quad (25)$$

For all simulated molecules, we performed 10,000 noisy iterations of the quantum circuit.

3 Errors in Molecular Energy

3.1 First-Row Diatomics

To begin studying the correlation between chemical properties and the effects of errors on quantum chemical state preparation, we first simulated a set of first-row diatomics of the periodic table. Though they are stable enough to be experimentally studied in the gas phase, several of the diatomics are unstable in air under standard laboratory conditions. However, we chose this molecular set in order to study chemical trends in susceptibility to error for the UCC ansatz.

Fig. 2 shows M_{Trot} (defined in Sec. 2.4) for both pure dephasing and pure relaxation, in the JW and BK mappings. M_{Trot} is defined such that when $M = M_{Trot}$, one cannot improve accuracy by increasing the Trotter number, because any decrease in Trotter error is counteracted by a gain in noise-induced error. Though there is no clear molecular trend, M_{Trot} tends to be larger for the BK than the JW mapping. This implies that for BK there is a larger range of decoherence rates for which increasing the Trotter number is beneficial, though this analysis does not yet allow one to draw conclusions concerning which mapping produces smaller errors in noisy circuits.

The bottom plot of Fig. 2 shows the distribution in chemical energy error for both JW and BK mappings, under an arbitrary coherence parameter of $M = 10^8$. Recall that energy error here is defined as the difference between the measured energy and the energy of the same state preparation under noiseless conditions. We plot results for this order of magnitude noise estimate because errors are within

Figure 2: *Top*: First row diatomic molecules. JW and BK data are plotted as squares and circles, respectively, using the color scheme of the bottom plot. The plot gives coherence parameter M_{Trot} , defined as the noise magnitude at which a single Trotter step produces the same error as two Trotter steps. *Bottom*: Mean error in the chemical energy relative to a noiseless state preparation for first row diatomics with coherence parameter $M = 10^8$, in both the Jordan–Wigner and Bravyi–Kitaev encodings, for a single Trotter step. The bars represent the increase in standard deviation of the chemical energy estimator, relative to the noise-free simulation. The JW circuits exhibit less error on average than the BK mapped molecules in the cases studied. Note the multiplier corresponding to the vertical axis ($1e-3$).

chemical accuracy (1 kcal/mol or 1.6×10^{-3} Hartrees) for the majority of our molecular set. The JW mapping results in smaller mean error and smaller error width for a given coherence parameter, often less than half compared to the BK result. As mentioned above, error widths are related to the increase in variance of the VQE molecular energy calculation. The superior robustness of JW is especially noteworthy considering that, on average, there tend to be fewer gates in our BK circuits.

Across our full set of eighteen molecules, BK mean errors are 1.2 to 3.2 times larger for a given molecule. Error widths are up to 2.0 times larger, with the sole case of a larger JW error width occurring in the case of Li_2 .

It is plausible that the superior robustness of JW over BK is due to the occupation numbers being stored locally. That is, a local single-qubit error in the BK mapping will affect several occupation numbers, whereas in the JW mapping it will affect just one. The corollary to this hypothesis is that occupation number errors may yield greater errors in the energy than parity errors do (see Sec. 2.3). However, further study is needed to definitively determine the source of the observed superior resistance of the JW encoding to noise, and to determine which of the two transformations is asymptotically more robust to noise as the size of the problem increases. In the remainder of this article, we use the JW mapping.

It is apparent from Fig. 3 that pure dephasing produces consistently smaller errors than pure relaxation. This is straight-forwardly explained by noting that the larger magnitude terms in the Hamiltonian are those corresponding to expectation values of the Z operator. Consequently, the change in energy result is more sensitive to a change in reference orbital occupation number (directly related to the Z operator) because of the strength of coulombic interactions, and because particle number errors produce states outside of the manifold of physical states (see Section 4). On the other hand, the main effect of dephasing (related to X and Y operators in molecular energy formula) is simply to evolve the state towards a product state. However, many of these trends in er-

Figure 3: Mean error in the chemical energy relative to a noiseless state preparation for first row diatomics in the Jordan–Wigner encoding, under three types of noise with coherence parameter $M = 10^8$. The bars again represent the increase in standard deviation of the chemical energy estimator, relative to the noise-free simulation. In all cases, the error is more sensitive to relaxation than to dephasing. Note the multiplier corresponding to the vertical axis ($1e-3$).

rors can become conflated with trends related to the number of gates in each quantum circuit, as longer gate sequences suffer more from decoherence.

To help clarify the nature of these errors, we then plot energy errors per gate (equivalent to error per time step) in Fig. 4, and see a clearer trend emerge. Going to the right across the periodic table produces a roughly consistent increase in mean error. We attribute this error largely to the increase in nuclear charge, but note that many factors will combine to produce this trend, including the magnitude of terms in the Hamiltonian, the fill factor, and the distance between nuclei. The error width, on the other hand, increases as one moves in either direction away from carbon, a trend that may be attributable to the fact that C_2 ’s particle number is the best-preserved out of this set of dimers (see Sec. 4). That is, it may be that states which are further from the manifold of physical states (physical states being those with the exactly correct electron number) have a larger error width.

Finally, we observe that, for moderate noise

Figure 4: Mean error *per gate* in the chemical energy relative to a noiseless state preparation for first row diatomics in the Jordan–Wigner encoding. The bars (error width) represent the increase in standard deviation relative to the noise-free simulation. The rough trend is for mean errors to increase with increasing nuclear charge. Error width per gate, on the other hand, increases as one moves in either direction away from carbon.

Figure 5: Product of mean observed error (with respect to a noiseless preparation) and coherence parameter for Be_2 (red squares), B_2 (green circles), and N_2 (blue triangles) in the Jordan–Wigner mapping. Note that this plotted value is equal to the ratio of mean error to decoherence rate. Dashed lines: T_1 type noise only; dotted lines: T_ϕ type noise only. The ratio asymptotes as the noise magnitude decreases. For lower values of M , the final state approaches a completely mixed state, which eventually yields a constant error and therefore a decreasing error to noise ratio.

magnitudes, the mean energy error to noise ratio (equal to the product of the mean energy error and the coherence parameter) is roughly constant above a certain threshold (Fig. 5). We compare Be_2 , B_2 , and N_2 because their circuits contain similar total gate numbers, from around 6800 to 7600 gates. This trivial result implies that, within a large range of noise magnitudes, error decreases roughly linearly with coherence parameter, assuming uncorrelated errors. The lower values of M (higher noise magnitudes) produce smaller ratios because the state approaches the completely mixed state as M decreases. Since the error itself approaches a constant value as the completely mixed state is approached, the error to noise ratio decreases.

3.2 Nitrogen Dissociation

Next we study the dissociation of nitrogen, the primary component of Earth’s atmosphere. Its dissociation is one required step of nitrogen fixation⁴⁶, necessary for producing fertilizer for the world’s agricultural needs. Comparing different conformations of a particular molecule allows us to elucidate trends based on attributes such as atomic orbital overlap or nuclear separation.

Figure 6: Mean error in the energy with respect to a noiseless state preparation for several bond lengths of the N_2 molecule in the JW encoding. The bars enclose one standard deviation of the observed errors. Sampling lengths are not evenly spaced. Note the multiplier corresponding to the vertical axis (10^{-3}).

The larger error in internuclear separations of

0.95 and 0.74 Å are simply a result of the larger number of gates. The varying number of gates is due in part to our using a cutoff of 10^{-5} for the cluster amplitudes. However, separations of 1.04 through 2.65 Å have nearly identical gate counts (within 3% of each other), which allows for more direct comparisons. We see little change in error for bond lengths greater than 1 Å, which suggests that error behavior is similar for a broad range of bond lengths.

Figure 7: In order to isolate nuclear charge, we plot the mean error in the energy with respect to a noisy state preparation for a series of isoelectronic species with equal bond lengths, in the JW mapping. The bars encompass one standard deviation of the observed errors. Increasing nuclear charge leads to increased error under all simulated types of noise.

3.3 Isoelectronic Species

As a final demonstration of a simple chemical-based trend, we compare a series of compounds with an identical number of electrons but varying charge and nuclei, so-called isoelectronic species. We simulated noisy circuits for C_2^{2-} , N_2 , and O_2^{2+} , each at the same bond length of 1.25 Å. Because the bond length, electron count, and spin remain constant, this allows us to isolate the nuclear charge as a variable. Errors increase with increasing nuclear charge, again despite nearly identical gate counts. As before, we attribute this trend to larger magnitudes in the electronic Hamiltonian, a direct result of the increasing nuclear charges. Further

connections to the previous work relating the error of gate-model quantum chemistry simulations with respect to nuclear charge²⁷ are left for future work.

We also simulated a set of molecules in which bond length and nuclear charge were constant while the total number of electrons increased. These calculations (omitted from the article) did not produce a clear trend in molecular energy error, partly because of a large variation in the number of gates required to implement the correct cluster amplitudes. However, for all simulated molecules there is a clear trend in particle number error as electron filling varies (independent of bond length and nuclear charge), as shown in Sec. 4.

4 Error in Electron Number Preservation

As discussed above, VQE and QPE are robust to errors assuming that we are concerned only with the final ground-state energy, not with knowing the true values of the parameters ξ themselves. There may of course be some cases in which it is important to know the correct cluster amplitudes that correspond to the ground state. For instance, one may want to prepare the state on a different quantum device prone a different set of systematic errors, or to calculate additional state properties for which the cluster amplitudes are explicitly needed.

Regardless of whether accurate cluster amplitude are desired, the energy calculation is useful as part of the VQE and QPE schemes only as long as the prepared state is chemically valid. In the case of the VQE approach this relates to the input state, whereas in the QPE approach this relates to the projected state at the energy measurement that could be chemically invalid as a result of Trotter or other errors. A necessary condition for validity is that the error in the electron number operator be small. Hence, it might be considered even more important to consider the errors in the number operator than in the energy calculation, since strictly speaking errors in chemical energy do not themselves invalidate the variational principle.

In the JW mapping, an orbital in the reference state is fully occupied for qubit state $|1\rangle$ and empty for state $|0\rangle$. Hence, the total electron number operator in the second quantization formalism is written as follows:

$$N = \sum_i N_i = \sum_i a_i^\dagger a_i \quad (26)$$

In this section we determine the error in the prepared state’s electron number expectation value for the molecules listed above. Error bars in the following plots enclose one standard deviation of the error in the particle number, relative to the noise-free simulation, over the set of noisy runs (note that this definition is different from the error bars defined in the plots of chemical energy error). We also show how $q\text{HiPSTER}$ may be useful for studying particle number errors resulting from systematic gate errors.

4.1 Effect of Noise on Electron Number

Figure 8: Mean error in the electron particle number relative to a noise-free state preparation for first row diatomic species in a JW encoding. The bars enclose one standard deviation in the observed data. Increased errors in molecules closer to the plot’s center are due in part to a larger number of time steps (gates) present in the circuit. Note the multiplier corresponding to the vertical axis (1e-4).

Here we report the robustness of the electron particle number for our molecular set, in

Figure 9: Mean error in the electron particle number relative to a noise-free state preparation *per gate* for first row diatomic species, in a JW encoding. The bars enclose one standard deviation in the observed data. Within our noise model, relaxation causes the particle number to approach $n/2$, where n is the number of spin-orbitals. Note the multiplier corresponding to the vertical axis (1e-7).

the Jordan–Wigner mapping. The simulations demonstrate that the electron particle number is significantly more robust to pure dephasing noise. This result is expected, as occupation is stored in the computational basis and all states studied in this section are close to the Hartree-Fock reference state, which is partially robust under the effects of dephasing. In other words: if no gate operations were performed on the initial reference state and pure-dephasing was present, the particle number would be preserved. It is the Hadamard and Y-basis gates that temporarily expose these Z amplitudes to error.

Insight can be gained more easily from the error per gate (Fig. 9) than from the total error (Fig. 8). The pure relaxation errors produce a clean trend, increasing based on how far the state is from the fully mixed state. This evolution towards the mixed state results from our decision to use the Pauli twirling approximation. We note that physical quantum hardware may behave asymptotically differently, depending on which quantum hardware is used. For example, qubits in ion traps⁵ often are implemented in such a way that they decay to the

ground state after dephasing, whereas transmon qubits³ decay to the vacuum state that is outside the computational manifold. Hence our use of PTA is meant to provide estimates for error magnitudes in relaxation noise relative to dephasing noise, and we emphasize that results in real hardware may differ qualitatively from the results shown in Figures 8 and 9.

Within our noise model, the smallest relaxation-induced errors are present in diatomics whose electron filling fractions are approximately half of the number of available spin-orbitals (B_2 , C_2 , N_2). We attribute this to the completely mixed state being largely composed of states corresponding to half-filling, and again emphasize that this trend will depend on the asymptotic behavior of the particular quantum hardware implementation.

Figure 10: Mean error in the electron particle number with respect to a noiseless state preparation for N_2 in the JW encoding at different internuclear separations. For bond distances larger than 0.95 Å, the circuits have nearly identical gate counts (equal to the number of time steps during which noise acts on the system), which results in similar occupation number errors.

In nitrogen dissociation, a slightly larger error from dephasing noise implies a slightly smaller error in relaxation noise. (Again, we consider only bond lengths 1.04 through 2.65 Å, because they have nearly identical numbers of gates.) Consider quantum circuits for bond lengths 1.46 and 1.85 Å. The 1.46 Å state’s larger departure from the original Z-aligned ref-

erence state leads to qubits whose directions are less aligned with Z eigenvalues on the Bloch sphere, making them comparably less robust to dephasing noise but more robust to relaxation noise. Though dephasing noise does not appreciably affect the particle number while the qubit is not being acted upon by an exponentiation operator, it does result in larger particle number errors when the qubit is rotated into the X basis with a Hadamard, exposed to additional environmental noise, and then rotated back to the Z basis.

4.2 Effect of Gate Errors on Electron Number

Systematic gate errors, if present, will also affect the particle number expectation value, even in the absence of appreciable environmental errors. As mentioned above, oftentimes a key concern for a VQE type approach is the preservation of the particle number. In this section, we demonstrate the use of $qHiPSTER$ for studying simple trends in the particle number error, for up to 40 qubits. We note that, in the minimal basis of spin-orbitals (with frozen $1s^2$ electrons for heavy atoms), 40 qubits (spin-orbitals) are sufficient for simulating propane, a hydrocarbon with the chemical formula C_3H_8 . For this article, we do not simulate a full quantum circuit for a specific molecule of 40 spin-orbitals, because of the long simulation times required.

As a first demonstration, we simulated two single-excitation operators, one after the other, assuming Jordan–Wigner mapping (circuit shown in reference¹²). In second quantization, this single-excitation operation is given by Eq. 12.

One such circuit consists of a mix of basis-change gates (Hadamard and its Y-basis analogue, which we denote Yb), CNOT gates, and two Z-rotation gates.

For these tests, all Hadamard and Yb gates were arbitrarily given systematic over-rotation errors of 10^{-4} radians. Simplistic CNOT errors, where present, were represented as controlled rotations of $\pi + 10^{-3}$ radians around the X-axis. At the beginning of each simulation, the $n/4$ lowest qubits were flipped so that the initial

particle number was $n/4$. All excitation magnitudes are in the range typical for quantum chemistry. Our purpose in this section is to observe the error trend, not to compare directly to gate errors in current quantum hardware. We note that our reported particle number errors are not due to Trotter errors, which our simulations showed were on the order of 10^{-14} or less for error-free circuits.

Figure 11 shows particle number errors that result when these systematic errors are assumed to be present on all gates. The first excitation is from qubit 0 to $n/4$, while the second is from qubit $n/4 - 1$ to $n - 1$. All Z-rotation gates are set to 0.1 radians. Errors increase approximately linearly with the number of qubits in the simulation. For these arbitrary error magnitudes, the particle number errors from a single exponential is too small to appreciably affect the result of a VQE or QPE calculation. However, assuming multiplicative behavior and thousands of gates, these particle number errors may degrade the overall accuracy of the chemical energy calculation.

A less predictable trend appears when we vary the magnitude of the single-particle excitation, while keeping the rest of the circuit constant. We again place two single excitations in a row, varying the magnitude of the first while keeping the second fixed. There are two Z-rotations of equal magnitude associated with each of the two excitations (four Z-rotations total). The first two Z-rotations (θ_1) are varied between 0 and 0.5 radians, and the second two Z-rotations (θ_2) are a constant 0.1 radians. Errors were present on all basis-change gates, while systematic CNOT errors were present only on the uninterrupted set of CNOT gates immediately preceding the fourth rotation operator. Somewhat counter-intuitively, larger Z-rotations can lead to smaller particle number errors. Starting with a θ_1 value of 0, the error increases slightly before declining at an increasing rate. As expected, the particle error increases with qubit number, simply because there are more total errors present in the circuit.

These circuits demonstrate the type of intuition that may be quickly gleaned by testing simple circuit trends, and one potential use of

qHiPSTER. Alternatively, one can imagine a case in which a limited fraction of available quantum resources have high fidelities, for example if there are enough quantum resources to error-correct only some of the available qubits. A powerful classical simulator of quantum devices is useful for determining how calculation accuracies are affected by the placement of error-prone gates on a particular portion of the quantum circuit. Significant work would be needed to come to concrete conclusions concerning the optimal placement of error-prone hardware. Future studies could consider how the particle number changes after a substantial portion is already in the excited state. Finally, we note that the possibility exists to use knowledge of error trends to deliberately introduce gates into the circuit that correct the particle number, though we do not explore this possibility further here.

Figure 11: Error in the particle number resulting from systematic single-qubit and CNOT errors, in a gate sequence representing two sequential single-excitation operators. E-H and E-Y denote errors due to systematic over-rotations in the Hadamard and Y-basis gates, respectively. Errors increase approximately linearly with the number of qubits.

5 High-Performance Quantum Simulator

In this section we summarize the implementation and optimization of *qHiPSTER*, our high-

Figure 12: Error in the electron particle number as a function of the excitation magnitude in the presence of systematic gate errors, for circuits of 28 and 36 qubits, in a gate sequence representing two sequential single-excitation operators. As a general trend, particle number error decreases with increasing excitation magnitude, for this set of systematic gate errors. The trend is not monotonic; of these five simulated values, the largest error occurs at $\theta = 0.1$.

performance quantum simulator used to study the noisy circuits described in earlier sections. Additional details are available in .⁸

Currently, *qHiPSTER* propagates pure states, which we use in the PTA as described. Hence, we store a $2^n \times 1$ state vector instead of the $2^n \times 2^n$ density matrix. As illustrated above, this still allows us to reproduce arbitrary observables of the density matrix that results from a noisy simulation.

We focus on implementing general single-qubit gates as well as two-qubit controlled gates (including CNOT gates), which are known to be universal.⁴⁷ Let Q be a 2×2 unitary matrix that represents single qubit gate operation:

$$Q = \begin{pmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{pmatrix}$$

Figure 13a shows the vector representation of a quantum state. Each amplitude has a subscript index in the binary representation. Figure 13b shows single-qubit gate operations on qubits 0 and 1. To perform a single-qubit gate on qubit k of the n -qubit quantum register, we apply Q to pairs of amplitudes whose indices

differ in the k -th bits of their binary index:

$$\begin{aligned} \beta'_{*...0_k*...} &= q_{11} \cdot \beta_{*...0_k*...} + q_{12} \cdot \beta_{*...1_k*...} \\ \beta'_{*...1_k*...} &= q_{21} \cdot \beta_{*...0_k*...} + q_{22} \cdot \beta_{*...1_k*...} \end{aligned} \quad (27)$$

A generalized two-qubit controlled- Q gate, with a control qubit c and a target qubit t , works as follows: if c is set to $|1\rangle$, Q is applied to t ; otherwise t is left unmodified:

$$\begin{aligned} \beta'_{*1_c*0_t*...} &= q_{11} \cdot \beta_{*1_c*0_t*...} + q_{12} \cdot \beta_{*1_c*1_t*...} \\ \beta'_{*1_c*1_t*...} &= q_{21} \cdot \beta_{*1_c*0_t*...} + q_{22} \cdot \beta_{*1_c*1_t*...} \end{aligned} \quad (28)$$

5.1 Single-Node Implementation and Optimization

The single node implementation of a single- and two-qubit controlled gates is trivial, and directly follows from equations 27 and 28. Namely, we iterate over consecutive groups of amplitudes of length 2^{k+1} , while applying Q to every pair of amplitudes within the group, separated by a stride of 2^k . Here we describe several performance optimizations.

Vectorization. Modern Intel CPUs support SIMD (Single Instruction Multiple Data) instructions, such as AVX2,⁴⁸ which can perform 4 double-precision operations simultaneously on 4 elements of the input registers. We map computation of every two pairs of complex amplitudes to 4-wide SIMD instructions. Each pair, which operates on real and imaginary parts, uses half of the SIMD register.

Multithreading. Modern CPUs can execute many concurrent hardware threads. We parallelize single- and two-qubit controlled gate operations on these threads using OpenMP 4.0.⁴⁹ Groups of amplitudes are evenly divided among the threads. When there are not enough groups to utilize all available threads, we parallelize computations within a group.

Cache Blocking. Instead of bringing the entire state from slow memory for each gate operation, we block gate operations in fast Last Level Cache (LLC). Namely, we “fuse” groups of consecutive gates, where each gate operates

Figure 13: Example of (a) a two-qubit quantum state, (b) single-qubit gate operations, applied to qubits 0 and 1, respectively, and (c) an implementation of a distributed single-qubit gate operation. Communication occurs between pairs of processors, CPU_i and CPU_j . Processors exchange half of their states, place it into temporary storage, compute on exchanged halves, and then perform another exchange.

on some qubit k , $k < l_c$, where 2^l_c is LLC size. We iterate over blocks of 2^l_c amplitudes of the state vector, applying each of the fused gates to this block, while the block remains resident in LLC and therefore benefits from the LLC's high bandwidth.

5.2 Multinode Implementation and Optimization

In our distributed implementation, a state vector of 2^n amplitudes (2^{n+4} bytes) is distributed among 2^p nodes, such that each node stores a local state of 2^{n-p} amplitudes. Let $m = n - p$. Naturally, 2^{m+4} must be less than the total memory capacity of the node.

Given single- or two-qubit gate operation on qubit k , if $k < m$, the operation is fully contained within a node. If $k \geq m$, the first and second elements of the pair are located on two different nodes and communication is required. We implement the communication scheme described in⁴² and demonstrated in Figure 13c. The two nodes exchange half of their state vectors into each other's temporary storage, then compute on exchanged halves, which is followed by another pair-wise exchange.

To reduce the memory requirements of temporary storage, we divide the distributed phase into multiple steps. At each step we exchange and reserve temporary storage for only a small

portion of the state vector, as opposed to the entire half, as done in reference⁴². This also allows us to overlap communication and computation in step i with state exchange in steps $i - 1$ and $i + 2$, which helps to partially hide the overhead of communication.

5.3 qHiPSTER Performance

We evaluate the performance and scalability of qHiPSTER on the Stampede supercomputer. Stampede⁵⁰ at the Texas Advanced Computing Center (TACC)/University of Texas (# 10 in the current TOP500 list) consists of 6,400 compute nodes, each of which is equipped with two sockets of Xeon E5-2680 and 32GB of DDR4 memory per node (16GB per socket). Each socket has 8 cores. The nodes are connected via a Mellanox FDR 56 Gb/s InfiniBand interconnect. Achievable memory bandwidth $B_{mem} = 40$ GB/s, while achievable network bandwidth $B_{net} = 5.5$ GB/s (bidirectional), per socket.

We have used 1000 nodes (2000 sockets), the maximum available allocation. With aggregate memory capacity of 32 Tbytes across 1000 nodes, we are able to simulate a quantum system of up to 40 qubits.

Multi-node results: Figure 14a shows time per single-qubit gate operation on multiple nodes. We report time per gate for 32, 256, 1K, and 2K sockets, which enable simulating quantum

(a) Weak scaling up to 40 qubits

(b) Quantum Fourier Transform on 40 qubits

Figure 14: Performance results: (a) Multinode time in seconds per single-qubit gate. Numbers next to several bars show exact measured time. Results of gate operation qubits 0 – 27 are similar to qubit 28 and are omitted. The large jump at 30 qubits is due to internode communication and is proportional to the ratio between network memory bandwidth. (b) Quantum Fourier Transform (QFT) up to 40 qubits. The upper curve shows total time per single QFT call. The lower curve shows average time per gate.

systems with 32, 37, 39, and 40 qubits, respectively. Note this is a “weak scaling” experiment. Specifically, we fix the local state vector to use the maximum amount of memory available on a socket. As we increase the number of qubits, we also use more sockets, and, as a result, the size of local state vector on a socket remains the same.

Gate operations applied to qubits 0 – 28 require no communication for all four quantum systems, and achieve the performance of $\sim 0.44s$ per gate (shown in Figure 14a, above the bar that corresponds to qubit position 28), limited only by memory bandwidth of the machine. Gates applied to higher-order qubits require communication. For the 32-node configuration we consistently see $\sim 3.53s$. This is $\sim 8\times$ higher than single qubit performance, and corresponds to the ratio between memory and network bandwidth of the system ($7.2 = 40/5.5$).

As the number of nodes increases, time per gate continues to increase. For example, a gate operation applied to qubit 39 on 1024 nodes takes 8.7 s (shown above the corresponding bar). This is a $2.5\times$ increase, compared to applying a gate operation to qubit 33 on 32 nodes, which takes 3.53 s (also shown above the cor-

responding bar). This is due to network contention and interference with other jobs running on the system at the same time. Controlled two-qubit gates show similar run-times and follow similar performance trends.

Note that compared to the JUMPIQCS distributed quantum simulator,⁴² *qHiPSTER* is $3\times$ to $10\times$ faster. The advantage is due to higher memory and network bandwidth, as well as the better interconnect of Stampede system, compared to the JUMP system of the Jülich Supercomputing Center on which JUMPIQCS was run¹.

Performance of QFT. Finally, we report the performance of the Quantum Fourier Transform (QFT). QFT is the fundamental kernel of many quantum algorithms, such as Shor’s algorithm for factoring,⁵² the quantum phase estimation (QPE) algorithm for estimating the eigenvalues of a unitary operator,⁵³ and algorithms for the hidden subgroup problem.⁵⁴ Its relevance to chemistry simulation is through QPE, which is used for one of the two major algorithmic approaches for chemical quantum circuits, as

¹In particular, Stampede uses 2-level Clos⁵¹ fat tree topology which has much higher path diversity, compared to the omega interconnect of the JUMP system

discussed in the introduction.

Figure 14b shows the performance of QFT as the number of qubits varies from 29 to 40. Note this is also a weak scaling experiment, as the size of the local state vector per node is fixed to be 2^{29} complex amplitudes. We see that total QFT time varies from 116 s for 29 qubits up to 997 s for 40 qubits. On average, for 40 qubits, each QFT gate operation takes ~ 1.22 s, as shown in the last column of Figure 14b.

Understanding run-time requirements of important quantum subroutines (such as QFT) is important, as it allows for estimating the simulation time of a quantum application that calls this kernel. For example, on the Stampede cluster, a single user application is limited to a maximum run-time of 24 hours. For a 40-qubit system, this would allow ~ 86 ($24 \times 3600/997$) calls to QFT for the total of $\sim 70,000$ quantum gates.

6 Conclusion

We have simulated noisy quantum circuits for preparing molecular electronic states in the unitary coupled cluster ansatz, in order to characterize errors in molecular energy and in electron particle number. It is necessary to study the effects of noise on the accuracy of quantum computation, because early quantum hardware will likely not be of high enough fidelity or will not contain sufficient resources to implement error correction codes. It is important to note that imperfectly error-corrected qubits will also have effective decoherence times.

For our set of eighteen molecules, we showed that the Jordan–Wigner mapping is less sensitive to noise than the Bravyi–Kitaev mapping. Isolating pure dephasing and pure relaxation noise demonstrated that relaxation noise produces larger errors than dephasing noise of the same magnitude, for both mappings. Additionally, these initial simulations suggest that there is a large range of relevant bond lengths for which error behavior will be similar.

The intuition gained from this study may be useful for guiding experimental and algorithmic choices, when implementing quantum circuits

for modeling chemistry. For example, the differing sensitivities to T_1 and T_ϕ times may help in choosing between similar quantum devices that are available for use. Finally, this article has demonstrated the utility of high performance simulators such as *qHiPSTER* for characterizing the effects of noise and gate errors in quantum circuits of 10 to 40 qubits, especially in circuits for which analytical characterization is difficult or impossible.

Acknowledgement The authors thank Jhonathan Romero, Alejandro Perdomo-Ortiz, Borja Peropadre, and Kostyantyn Kechedzhi for fruitful conversations, as well as Gian Giacomo Guerreschi for proofreading the manuscript. N.P.D.S. and A.A.-G. acknowledge support from the Air Force Office of Scientific Research under award FA9550-12-1-0046, the Army Research Office under award, W911NF-15-1-0256 and the Defense Security Science Engineering Fellowship managed by the Office of Naval Research under award N00014-16-1-2008. J.R.M. is a Luis W. Alvarez Fellow in Computing Sciences supported by Laboratory Directed Research and Development (LDRD) funding from Berkeley Lab, provided by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The authors also acknowledge the Texas Advanced Computing Center (TACC, <http://www.tacc.utexas.edu>) at the University of Texas at Austin for providing HPC resources.

References

- (1) Montanaro, A. *Npj Quantum Information* **2016**, 2, 15023.
- (2) Nielsen, M. A.; Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed.; Cambridge University Press: New York, NY, USA, 2011.
- (3) Kelly, J.; Barends, R.; Fowler, A. G.; Megrant, A.; Jeffrey, E.; White, T. C.; Sank, D.; Mutus, J. Y.; Campbell, B.; Chen, Y.; Chen, Z.; Chiaro, B.; Dunsworth, A.; Hoi, I.-C.; Neill, C.;

- O'Malley, P. J. J.; Quintana, C.; Roushan, P.; Vainsencher, A.; Wenner, J.; Cleland, A. N.; Martinis, J. M. *Nature* **2015**, *519*, 66–69.
- (4) Senko, C.; Richerme, P.; Smith, J.; Lee, A.; Cohen, I.; Retzker, A.; Monroe, C. *Phys. Rev. X* **2015**, *5*, 1–9.
 - (5) Schindler, P.; Nigg, D.; Monz, T.; Barreiro, J. T.; Martinez, E.; Wang, S. X.; Quint, S.; Brandl, M. F.; Nebendahl, V.; Roos, C. F.; Chwalla, M.; Hennrich, M.; Blatt, R. *New J. Phys.* **2013**, *15*, 1–37.
 - (6) Rieffel, E. G.; Polak, W. H. *Quantum Computing: A Gentle Introduction*; The MIT Press: Cambridge, MA, USA, 2014.
 - (7) O'Malley, P. J. J.; Babbush, R.; Kivlichan, I. D.; Romero, J.; McClean, J. R.; Barends, R.; Kelly, J.; Roushan, P.; Tranter, A.; Ding, N.; Campbell, B.; Chen, Y.; Chen, Z.; Chiaro, B.; Dunsworth, A.; Fowler, A. G.; Jeffrey, E.; Megrant, A.; Mutus, J. Y.; Neill, C.; Quintana, C.; Sank, D.; Vainsencher, A.; Wenner, J.; White, T. C.; Coveney, P. V.; Love, P. J.; Neven, H.; Aspuru-Guzik, A.; Martinis, J. M. *ArXiv e-prints* **2015**, arXiv:1512.06860v1.
 - (8) Smelyanskiy, M.; Sawaya, N. P. D.; Aspuru-Guzik, A. *ArXiv e-prints* **2016**, quant-ph/1601.07195.
 - (9) Aspuru-Guzik, A.; Dutoi, A. D.; Love, P. J.; Head-Gordon, M. *Science* **2005**, *309*, 1704–1707.
 - (10) Abrams, D. S.; Lloyd, S. *Phys. Rev. Lett.* **1999**, *83*, 5162–5165.
 - (11) Veis, L.; Pittner, J. *J. Chem. Phys.* **2010**, *133*.
 - (12) Whitfield, J. D.; Biamonte, J.; Aspuru-Guzik, A. *Mol. Phys.* **2011**, *109*, 735–750.
 - (13) Reiher, M.; Wiebe, N.; Svore, K. M.; Wecker, D.; Troyer, M. *arXiv preprint* **2016**, quant-ph/1605.03590v2.
 - (14) Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; O'Brien, J. L. *Nat. Commun.* **2015**, *5*, 4213.
 - (15) McClean, J. R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. *New J. Phys.* **2016**, *18*, 023023.
 - (16) Wecker, D.; Hastings, M. B.; Troyer, M. *Phys. Rev. A* **2015**, *92*, 042303.
 - (17) Barrett, S.; Hammerer, K.; Harrison, S.; Northup, T. E.; Osborne, T. J. *Phys. Rev. Lett.* **2013**, *110*, 090501.
 - (18) Eichler, C.; Mlynec, J.; Butscher, J.; Kurpiers, P.; Hammerer, K.; Osborne, T. J.; Wallraff, A. *Phys. Rev. X* **2015**, *5*, 041044.
 - (19) Taube, A. G.; Bartlett, R. J. *Int. J. Quant. Chem.* **2006**, *106*, 3393–3401.
 - (20) Bartlett, R. J.; Musiał, M. *Rev. Mod. Phys.* **2007**, *79*, 291–352.
 - (21) Wang, H.; Kais, S.; Aspuru-Guzik, A.; Hoffmann, M. R. *Phys. Chem. Chem. Phys.* **2008**, *10*, 5388–5393.
 - (22) McClean, J. R.; Babbush, R.; Love, P. J.; Aspuru-Guzik, A. *J. Phys. Chem. Lett.* **2014**, *5*, 4368–4380.
 - (23) Jordan, P.; Wigner, E. *Zeitschrift für Physik* **1928**, *47*, 631–651.
 - (24) Bravyi, S. B.; Kitaev, A. Y. *Ann. Phys.* **2002**, *298*, 210 – 226.
 - (25) Seeley, J. T.; Richard, M. J.; Love, P. J. *J. Chem. Phys.* **2012**, *137*, 224109.
 - (26) Hastings, M. B.; Wecker, D.; Bauer, B.; Troyer, M. *Quantum Info. Comput.* **2015**, *15*, 1–21.
 - (27) Babbush, R.; McClean, J.; Wecker, D.; Aspuru-Guzik, A.; Wiebe, N. *Phys. Rev. A* **2015**, *91*, 022311.
 - (28) Makri, N.; Makarov, D. E. *J. Chem. Phys.* **1995**, *102*, 4600–4610.

- (29) Makri, N.; Makarov, D. E. *J. Chem. Phys.* **1995**, *102*, 4611–4618.
- (30) Redfield, A. G. In *Adv. Magn. Reson.*; Waugh, J. S., Ed.; Advances in Magnetic and Optical Resonance; Academic Press, 1965; Vol. 1; pp 1 – 32.
- (31) Plenio, M. B.; Knight, P. L. *Rev. Mod. Phys.* **1998**, *70*, 101–144.
- (32) Fox, R. F. *Phys. Rep.* **1978**, *48*, 179 – 283.
- (33) Aharonov, D.; Kitaev, A.; Preskill, J. *Phys. Rev. Lett.* **2006**, *96*, 050504.
- (34) Monz, T.; Schindler, P.; Barreiro, J. T.; Chwalla, M.; Nigg, D.; Coish, W. a.; Harlander, M.; Hänsel, W.; Hennrich, M.; Blatt, R. *Phys. Rev. Lett.* **2011**, *106*, 130506.
- (35) Kitaev, A. Y. *Russ. Math. Surv.* **1997**, *52*, 1191.
- (36) Jacobs, K.; Knight, P. L. *Phys. Rev. A* **1998**, *57*, 2301–2310.
- (37) Bassi, A. *Phys. Rev. A* **2003**, *67*, 062101.
- (38) Bassi, A.; Deckert, D.-A. *Phys. Rev. A* **2008**, *77*, 032323.
- (39) Silva, M.; Magesan, E.; Kribs, D. W.; Emerson, J. *Phys. Rev. A* **2008**, *78*, 012347.
- (40) Geller, M. R.; Zhou, Z. *Phys. Rev. A* **2013**, *88*, 012314.
- (41) Tomita, Y.; Svore, K. M. *Phys. Rev. A* **2014**, *90*, 062320.
- (42) Trieu, D. B. Large-scale simulations of error prone quantum computation devices. Ph.D. thesis, University of Wuppertal, 2010.
- (43) Tranter, A.; Sofia, S.; Seeley, J.; Kaicher, M.; McClean, J.; Babbush, R.; Coveney, P. V.; Mintert, F.; Wilhelm, F.; Love, P. J. *Int. J. Quant. Chem.* **2015**, *115*, 1431–1441.
- (44) Turney, J. M.; Simmonett, A. C.; Parrish, R. M.; Hohenstein, E. G.; Evangelista, F. A.; Fermann, J. T.; Mintz, B. J.; Burns, L. A.; Wilke, J. J.; Abrams, M. L.; Russ, N. J.; Leininger, M. L.; Janssen, C. L.; Seidl, E. T.; Allen, W. D.; Schaefer, H. F.; King, R. A.; Valeev, E. F.; Sherrill, C. D.; Crawford, T. D. *WIREs Comput. Mol. Sci.* **2012**, *2*, 556–565.
- (45) NIST Computational Chemistry Comparison and Benchmark Database, NIST Standard Reference Database Number 101 Release 17b. 2015; <http://cccbdb.nist.gov/>.
- (46) Erisman, J. W.; Sutton, M. A.; Galloway, J.; Klimont, Z.; Winiwarter, W. *Nat. Geosci.* **2008**, *1*, 636.
- (47) DiVincenzo, D. P. *Phys. Rev. A* **1995**, *51*, 1015–1022.
- (48) New AVX2 instruction descriptions available. software.intel.com, Retrieved: 2012-01-17.
- (49) OpenMP Architecture Review Board, OpenMP Application Program Interface Version 3.0. 2013; <http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf>.
- (50) Texas Advanced Computing Center (TACC). <https://www.tacc.utexas.edu/stampede/>, Retrieved: 2014-12-01.
- (51) Clos, C. *Bell Syst. Tech. J.* **1953**, *32*, 406–424.
- (52) Shor, P. W. *SIAM J. Comput.* **1997**, *26*, 1484–1509.
- (53) Kitaev, A. *Electronic Colloquium on Computational Complexity (ECCC)* **1996**, *3*.
- (54) Lomont, C. *ArXiv e-prints* **2004**, quant-ph/0411037v1.

Graphical TOC Entry