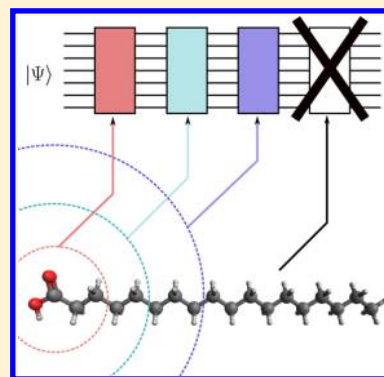


Exploiting Locality in Quantum Computation for Quantum Chemistry

Jarrod R. McClean,[†] Ryan Babbush,[†] Peter J. Love,[‡] and Alán Aspuru-Guzik^{*,†}[†]Department of Chemistry and Chemical Biology, Harvard University, Cambridge, Massachusetts 02138, United States[‡]Department of Physics, Haverford College, Haverford, Pennsylvania 19041, United States

ABSTRACT: Accurate prediction of chemical and material properties from first-principles quantum chemistry is a challenging task on traditional computers. Recent developments in quantum computation offer a route toward highly accurate solutions with polynomial cost; however, this solution still carries a large overhead. In this Perspective, we aim to bring together known results about the locality of physical interactions from quantum chemistry with ideas from quantum computation. We show that the utilization of spatial locality combined with the Bravyi–Kitaev transformation offers an improvement in the scaling of known quantum algorithms for quantum chemistry and provides numerical examples to help illustrate this point. We combine these developments to improve the outlook for the future of quantum chemistry on quantum computers.



Within chemistry, the Schrödinger equation encodes all information required to predict chemical properties ranging from reactivity in catalysis to light absorption in photovoltaics. Unfortunately, the exact solution of the Schrödinger equation is thought to require exponential resources on a classical computer due to the exponential growth of the dimensionality of the Hilbert space as a function of molecular size. This makes exact methods intractable for more than a few atoms.¹

Richard Feynman first suggested that this scaling problem might be overcome if a more natural approach was taken.² Specifically, instead of painstakingly encoding quantum information into a classical computer, one may be able to use a quantum system to naturally represent another quantum system and bypass the seemingly insurmountable storage requirements. This idea eventually developed into the field of quantum computation, which is now believed to hold promise for the solution of problems ranging from factoring numbers³ to image recognition^{4,5} and protein folding.^{6,7}

Initial studies by Aspuru-Guzik et al. showed that these approaches might be particularly promising for quantum chemistry.⁸ There have been many developments in both theory^{9–11} and experimental realization^{12–15} of quantum chemistry on quantum computers. The original gate construction for quantum chemistry introduced by Whitfield et al.¹⁶ was recently challenged as too expensive by Wecker et al.¹⁷ The pessimistic assessment was due mostly to the extrapolation of the Trotter error for artificial rather than realistic molecular systems, as was analyzed in detail in a followup study by many of the same authors.¹⁸ They subsequently improved the scaling by means of several circuit enhancements.¹⁹ The analysis of the Trotter error on realistic molecules in combination with their improvements led to a recent study where an estimate of the calculation time of Fe₂S₂ was reduced by orders of magnitude.¹⁸ In this Perspective, we further reduce the scaling by exploiting

the locality of physical interactions with local basis sets, as has been done routinely now in quantum chemistry for 2 decades.^{20,21} These improvements in combination with others make quantum chemistry on a quantum computer a very attractive application for early quantum devices. We describe the scaling under two prominent measurement strategies, quantum phase estimation and Hamiltonian averaging, which is a simple subroutine of the recently introduced Variational Quantum Eigensolver approach.¹⁴

Additionally, recent progress in accurate and scalable solutions of the Schrödinger equation on classical computers has also been significant.^{20–25} Some of these results have already appeared in the quantum computation literature in the context of in-depth studies of state preparation.^{26,27} A general review of quantum simulation^{28,29} and one on quantum computation for chemistry³⁰ cover these topics in more depth. A collection covering several aspects of quantum information and chemistry recently appeared.³¹ However, many developments that utilize fundamental physical properties of the systems being studied to enable scalability have not yet been exploited.

In this study, we hope to bring to light results from quantum chemistry as well as their scalable implementation on quantum computers. We begin by reviewing the standard electronic structure problem. Results based on the locality of physical interactions from linear scaling methods in quantum chemistry are then introduced with numerical studies to provide quantification of these effects. A discussion of the resulting impact on the most common quantum algorithms for quantum chemistry follows. We also investigate instances where a perfect oracle is not available to provide input states, demonstrating the

Received: August 6, 2014

Accepted: November 25, 2014

Published: November 25, 2014

need for advances in state preparation technology. Finally, we conclude with an outlook for the future of quantum chemistry on quantum computers.

To frame the problem and set the notation, we first briefly introduce the electronic structure problem of quantum chemistry.²³ Given a set of nuclei with associated charges $\{Z_i\}$ and a total charge (determining the number of electrons), the physical states of the system can be completely characterized by the eigenstates $\{|\Psi_i\rangle\}$ and corresponding eigenvalues (energies) $\{E_i\}$ of the Hamiltonian H

$$H = -\sum_i \frac{\nabla_{R_i}^2}{2M_i} - \sum_i \frac{\nabla_{r_i}^2}{2} - \sum_{i,j} \frac{Z_i}{|R_i - r_j|} + \sum_{j>i} \frac{Z_i Z_j}{|R_i - R_j|} + \sum_{j>i} \frac{1}{|r_i - r_j|} \quad (1)$$

where we have used atomic units; $\{R_i\}$ denote nuclear coordinates, $\{r_i\}$ electronic coordinates, $\{Z_i\}$ nuclear charges, and $\{M_i\}$ nuclear masses. Owing to the large difference in masses between the electrons and nuclei, typically the Born–Oppenheimer approximation is used to mitigate computational cost, and the nuclei are treated as stationary, classical point charges with fixed positions $\{R_i\}$. Within this framework, the parametric dependence of the eigenvalues on $\{R_i\}$, denoted by $\{E(\{R_i\})\}$, determines almost all chemical properties, such as bond strengths, reactivity, vibrational frequencies, and so forth. Work has been done in the determination of these physical properties directly on a quantum computer.³²

Due to the large energy gaps between electronic levels with respect to the thermal energy scale $k_B T$, it typically suffices to study a small subset of the eigenstates corresponding to the lowest energies. Moreover, for this reason, in many molecules, the lowest-energy eigenstate $|\Psi_0\rangle$, or ground state, is of primary importance, and for that reason, it is the focus of many methods, including those discussed here.

Direct computation in a positional basis accounting for antisymmetry in the wave function while using the Hamiltonian described is referred to as a first quantization approach and has been explored in the context of quantum computation.^{33–35} The first quantized approach has also been realized in experiment.³⁶ One may also perform first quantized calculations in a basis of Slater determinants. This was introduced as a representation of the electronic wave function by qubits in ref 8 (the compact mapping), and the efficiency of time evolution in this basis was recently shown.^{37,38} The second quantized approach places the antisymmetry requirements on the operators. After choosing some orthogonal spin–orbital basis $\{\varphi_i\}$ with a number of terms M , the second quantized Hamiltonian may be written as²³

$$\hat{H} = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s \quad (2)$$

with coefficients determined by

$$h_{pq} = \int d\sigma \varphi_p^*(\sigma) \left(-\frac{\nabla_r^2}{2} - \sum_i \frac{Z_i}{|R_i - r|} \right) \varphi_q(\sigma) \quad (3)$$

$$h_{pqrs} = \int d\sigma_1 d\sigma_2 \frac{\varphi_p^*(\sigma_1) \varphi_q^*(\sigma_2) \varphi_s(\sigma_1) \varphi_r(\sigma_2)}{|r_1 - r_2|} \quad (4)$$

where σ_i now contains the spatial and spin components of the electron, $\sigma_i = (r_{ip}, s_i)$. The operators a_p^\dagger and a_r obey the Fermionic anticommutation relations

$$\{a_p^\dagger, a_r\} = \delta_{p,r} \quad (5)$$

$$\{a_p^\dagger, a_r^\dagger\} = \{a_p, a_r\} = 0 \quad (6)$$

For clarity, we note that the basis functions used in quantum chemistry (such as atom-centered Gaussians) are frequently parametrized on the nuclear coordinates $\{R_i\}$, which can result in a dependence on the nuclear positions of the electronic integral terms $\{h_{pqrs}\}$. For notational simplicity, the dependence of the integrals on the nuclear positions in this work will remain implied.

It is clear by inspection that the maximum number of terms in the second-quantized Hamiltonian scales as $O(M^4)$. M can be quite large to reach chemical accuracy for systems of interest, and the number of terms present in the Hamiltonian is a dominant cost factor for almost all quantum computation algorithms for chemistry. However, due to the locality of physical interactions, one might imagine that many of the terms in the Hamiltonian are negligible relative to some finite precision ϵ . While this depends on the basis, it is this observation that forms the foundation for the linear scaling methods of electronic structure such as linear scaling density functional theory or quantum Monte Carlo.^{21,22,39–43} That is, in a local basis, the number of non-negligible terms scales more like $O(M^2)$, and advanced techniques such as fast multipole methods techniques can evaluate their contribution in $O(M)$ time.

These scaling properties are common knowledge within the domain of traditional quantum chemistry; however, they have not yet been exploited within the context of quantum computation.

These scaling properties are common knowledge within the domain of traditional quantum chemistry; however, they have not yet been exploited within the context of quantum computation. They are clearly vitally important for the correct estimate of the asymptotic scaling of any method.^{8,9,16,17} For that reason, we review the origin of that scaling here for the most common and readily available local basis, the Gaussian atomic orbital basis. We follow loosely the explanation presented by Helgaker, Jørgensen, and Olsen,²³ and refer readers to this text for additional detail on the evaluation of molecular integrals in local basis sets. The two elements that we will consider here are the cutoffs due to exponentially vanishing overlaps between Gaussians basis functions and a bound on the value of the largest integral.

By far the most common basis used in electronic structure calculations is a set of atom-centered Gaussian (either Cartesian or “pure” spherical) functions. While the precise result can depend on the angular momentum associated with the basis function, for simplicity, consider only Gaussian S functions, which are defined by

$$|G_a\rangle = \exp(-ar_A^2) \quad (7)$$

where r_A is the vector from a point A that defines the center of the Gaussian. One property of Gaussian functions that turns out to be useful in the evaluation of molecular integrals is the Gaussian product rule. This rule states simply that the product of two spherical Gaussian functions may be written in terms of a single spherical Gaussian function on the line segment connecting the two centers. Consider two spherical Gaussian functions, $|G_a\rangle$ and $|G_b\rangle$ separated along the x -axis

$$\exp(-ax_A^2) \exp(-bx_B^2) = K_{ab}^x \exp(-px_p^2) \quad (8)$$

where K_{ab}^x is now a constant pre-exponential factor

$$K_{ab}^x = \exp(-\mu X_{AB}^2) \quad (9)$$

and the total exponent p , the reduced exponent μ , and the Gaussian separation X_{AB} are given by

$$p = a + b \quad (10)$$

$$\mu = \frac{ab}{a + b} \quad (11)$$

$$X_{AB} = A_x - B_x \quad (12)$$

That is, the product of two spherical Gaussians is a third Gaussian centered between the original two that decays faster than the original two functions, as given by the total exponent p . The overlap integral of two spherical Gaussian S functions may be obtained through application of the Gaussian product rule after factorizing into the three Cartesian dimensions followed by Gaussian integration and is given by

$$S_{ab} = \langle G_a | G_b \rangle = \left(\frac{\pi}{a + b} \right)^{3/2} \exp\left(-\frac{ab}{a + b} R_{AB}^2 \right) \quad (13)$$

where R_{AB} is the distance between the Gaussian centers A and B . Clearly, this integral decays exponentially with the square of the distance between centers, and one may determine a distance d_s such that beyond that distance, the integrals will be smaller than 10^{-k} in magnitude

$$d_s = \sqrt{a_{\min}^{-1} \log \left[\left(\frac{\pi}{2a_{\min}} \right)^3 10^{2k} \right]} \quad (14)$$

where a_{\min} is the minimal Gaussian exponent a (most diffuse function) in the set of Gaussian basis functions $\{|G_a\rangle\}$. While the exact decay parameters will depend on the basis set, it is generally true from this line of reasoning that there is a characteristic distance beyond which all overlap integrals are negligible. This means that the number of interactions per basis function becomes fixed, resulting in a linear number of significant overlap integrals. As kinetic energy integrals are just fixed linear combinations of overlap integrals of higher angular momentum, the same argument holds for them as well.

For S orbitals, the two-electron Coulomb integral may be written as

$$h_{abcd} = \frac{S_{ab} S_{cd}}{R_{PQ}} \operatorname{erf}(\sqrt{\alpha} R_{PQ}) \quad (15)$$

where erf is the error function and P and Q are Gaussian centers formed through application of the Gaussian product rule to $|G_a\rangle|G_b\rangle$ and $|G_c\rangle|G_d\rangle$, respectively. R_{PQ} is the distance between the two Gaussian centers P and Q , and α is the

reduced exponent derived from P and Q . For clarity, this may be bounded by the simpler expression

$$h_{abcd} \leq \min \left(\frac{4\alpha}{\pi} S_{ab} S_{cd}, \frac{S_{ab} S_{cd}}{R_{PQ}} \right) \quad (16)$$

The first of these two expressions in the min function comes from the short-range bound and the latter from the long-range bound of the error function. These bounds show that the integrals are determined by products of overlap terms, such that in the regime where overlap integrals scale linearly, we expect $O(M^2)$ as significant two-electron terms. Moreover, as seen in the long-range bound of the two-electron integral (TEI), there is some further asymptotic distance beyond which these interactions may be completely neglected, yielding an effectively linear scaling number of significant integrals. This limit can be quite large, however; thus, practically one expects to observe a quadratic scaling in the number of TEIs.

Additionally, we note from the form of the integrals that the maximal values that the TEIs will attain are determined by the basis set parameters, such as the width of the Gaussian basis functions or their angular momentum. The implication of this is that the maximal integral magnitude for the four index TEIs, $|h_{\max}^{\text{TEI}}|$ will be independent of the molecular size for standard atom-centered Gaussian basis sets and may be treated as a constant for scaling analysis that examines cost as a function of physical system size with fixed chemical composition. The overlap and kinetic energy integrals will similarly have a maximum independent of molecular size past a very small length scale. However, the nuclear attraction integrals must also be considered.

While not typically considered a primary source of difficulty due to the relative ease of evaluation with respect to TEIs, we separate the nuclear attraction integrals here due to the fact that the maximal norm of the elements may change as well. The nuclear attraction matrix element between S functions may be written as

$$h_{ab}^{\text{nuc}} = - \sum_i \frac{Z_i S_{ab}}{R_{pi}} \operatorname{erf}(\sqrt{p} R_{pi}) \quad (17)$$

where Z_i is the nuclear charge and R_{pi} refers to the distance between the Gaussian center P with total exponent p formed from the product $|G_a\rangle|G_b\rangle$ and the position of the i th nuclei. Following from the logic above, from the exponentially vanishing overlap S_{ab} at some distance, we expect only a linear number of these integrals to be significant. However, each of the integrals considers the sum over all nuclei, which can be related linearly to the number of basis functions in atom-centered Gaussian basis sets. Thus, the maximal one-electron integral (OEI) is not a constant but rather can be expected to scale with the Coulomb sum over distant nuclear charges. A conservative bound can be placed on such a maximal element as follows.

The temperature and pressure that a molecule resides in will typically determine the minimal allowed separation of two distinct nuclei and will thus define a maximum nuclear density ρ_{\max} . Denote the maximum nuclear charge in the systems under consideration as Z_{\max} . The maximal density and the number of nuclei will also define a minimal radius that a sphere of charge may occupy, r_{\max}

$$r_{\max}^3 = \frac{3Z_{\max}N_{\text{nuc}}}{4\pi\rho_{\max}} \quad (18)$$

where N_{nuc} is the number of nuclei in the system. Modeling the charge as spread uniformly within this minimal volume and using the maximum of the error function to find a bound on the maximum for the nuclear attraction matrix element, we find

$$\begin{aligned} |h_{ab}^{\text{nuc}}| &< 4\pi\rho_{\max}S_{ab}\left|\int_0^{r_{\max}} r^2 dr \frac{1}{r}\right| \\ &= 2\pi\rho_{\max}S_{ab}r_{\max}^2 \\ &= \beta_{ab}N_{\text{nuc}}^{2/3} \end{aligned} \quad (19)$$

where β_{ab} is now a system-size-independent quantity determined only by basis set parameters at nuclei a and b , and the size dependence is bounded as $O(N_{\text{nuc}}^{2/3})$. Atom-centered Gaussian basis sets will have a number of basis functions that are a linear multiple of the number of nuclei, and as such, we may now bound the maximal OEI element as

$$|h_{\max}^{\text{OEI}}| < \beta_{\max}^{\text{OEI}}M^{2/3} \quad (20)$$

The above analysis demonstrates that given some integral magnitude threshold, δ , there exists a characteristic distance d between atomic centers, beyond which integrals may be neglected. If one is interested in a total precision ϵ in the energy E_i , it is important to know how choosing δ will impact the solution and what choice of δ allows one to retain a precision ϵ .

By specification, the discarded integrals are small with respect to the rest of the Hamiltonian (sometimes as much as 10 orders of magnitude smaller in standard calculations). As such, one expects a perturbation analysis to be accurate. Consider the new, truncated Hamiltonian $H_t = H + V$, where V is the negation of the sum of all removed terms, each of which has a magnitude less than δ .

Assuming a nondegenerate spectrum for H , from perturbation theory, we expect the leading order change in eigenvalue E_i to be given by

$$\Delta E_i = \langle \Psi_i | V | \Psi_i \rangle \quad (21)$$

If the number of terms removed from the sum is given by N_r , a worst case bound on the magnitude of this deviation follows from the spectrum of the creation and annihilation operators and is given by

$$|\Delta E_i| \leq \sum_{\{h_i:|h_i|<\delta\}} |h_i| \leq N_r\delta \quad (22)$$

where $\{h_i:|h_i|<\delta\}$ is simply the set of Hamiltonian elements with a norm less than δ and the first inequality follows directly from the triangle inequality. We emphasize that this is a worst case bound, and generically, one expects at least some cancellation between terms, such as kinetic and potential terms, when the Hamiltonian is considered as a whole. Some numerical studies of these cancellation effects have been performed,¹⁸ but additional studies are required. Regardless, under this maximal error assumption, by choosing a value

$$\delta \leq \frac{\epsilon}{N_r} \quad (23)$$

one retains an accuracy ϵ in the final answer with respect to the exact answer when measuring the eigenvalue of the truncated

Hamiltonian H_t . Alternative, one may use the tighter bound based on the triangle inequality and remove the maximum number of elements such that the total magnitude of removed terms is less than ϵ . From the looser but simpler bound, we see that a reduction of scaling from M^4 to M^2 would require removal of the order of M^4 terms from the Hamiltonian; this constraint on δ can be rewritten in terms of M as

$$\delta \leq \frac{\epsilon}{M^4} \quad (24)$$

While the perturbation of the eigenvalue will have a direct influence on energy projective measurement methods such as quantum phase estimation, other methods evaluate the energy by averaging. In this case, we do not need to appeal to perturbation theory, and the δ required to achieve a desired ϵ can be found directly

$$\langle H_t \rangle = \langle \Psi_i | H_t | \Psi_i \rangle \quad (25)$$

$$= E_i + \langle \Psi_i | V | \Psi_i \rangle \quad (26)$$

We find that under our assumption of worst case error for averaging, the result is identical to that of the first-order perturbation of the eigenvalue E_i

$$|\Delta \langle H_t \rangle| \leq \sum_{\{h_i:|h_i|<\delta\}} |h_i| \leq N_r\delta \quad (27)$$

In summary, we find that for both the consideration of the ground-state eigenvalue and the average energy of the ground-state eigenvector, there is a simple formula for the value of δ , which scales polynomially in the system size, below which one may safely truncate to be guaranteed an accuracy ϵ in the final answer. Moreover, it suggests a simple strategy that one may utilize to achieve the desired accuracy. That is, sort the integrals in order of magnitude and remove the maximum number of integrals such that the total magnitude of removed integrals is less than ϵ .

On the subject of general truncation, we note that while there may exist Hamiltonians with the same structure as the second quantized electronic structure Hamiltonian that have the property that removal of small elements will cause a drastic shift in the character of the ground state, this has not been seen for physical systems in quantum chemistry. Moreover, from the perturbation theory analysis given, such Hamiltonians would likely need to exhibit degenerate ground electronic states, which are not common in physical systems. In practice, it is observed that removing elements on the order of $\delta = 10^{-10}$ and smaller is more than sufficient to retain both qualitative and quantitative accuracy in systems of many atoms.^{22,23,39,40} Moreover, the convergence with respect to this value may be tested easily for any systems under consideration.

While the above analysis shows that locality of interactions in local basis sets provides a promise that, beyond a certain length scale, the number of non-negligible integrals will scale quadratically in the number of basis functions, it does not provide good intuition for the size of that length scale in physical systems of interest. Here, we provide numerical examples for chemical systems in basis sets used so far in quantum computation for quantum chemistry. The precise distance at which locality starts to reduce the number of significant integrals depends, of course, on the physical system and the basis set used. In particular, larger, more diffuse basis sets are known to exhibit these effects at comparatively larger

length scales than minimal, compact basis sets. However, the general scaling arguments given above hold for all systems of sufficient size.

An additional consideration that must be made for quantum computation is that, as of yet, no general technology has been developed for direct simulation in nonorthogonal basis sets. This prohibits direct simulation in the bare atomic orbital basis; however, the use of Löwdin symmetric orthogonalization yields the orthogonal basis set closest to the original atomic orbital basis set in an l^2 sense.^{44,45} We find that this is sufficient for the systems that we consider, but note that there have been a number of advances in orthogonal basis sets that are local in both the occupied and virtual spaces and may find utility in quantum computation.⁴⁶ Moreover, there has been recent work in the use of multiresolution wavelet basis sets that have natural sparsity and orthogonality while providing provable error bounds on the choice of basis.⁴⁷ Such a basis also allows one to avoid costly integral transformations related to orthogonality, which are known to scale as $O(M^5)$ when performed exactly. Further research is needed to explore the implications for quantum computation with these basis sets, and we focus here on the more common atom-centered Gaussian basis sets.

As a prototype system, we consider chains of hydrogen atoms separated by 1 Bohr (a_0) in the STO-3G basis set, an artificial system that can exhibit a transition to a strongly correlated wave function.⁴⁸ We count the total number of significant integrals for values of δ given by 10^{-15} and 10^{-7} for the symmetrically orthogonalized atomic orbital (OAO) basis and the canonical Hartree–Fock molecular orbital (MO) basis. The results are displayed in Figure 1 and demonstrate that with




Figure 1. Number of significant (magnitude $> 10^{-15}$) spin–orbital integrals in the STO-3G basis set as a function of the number of hydrogens in a linear hydrogen chain with a separation of $1 a_0$ for the Hartree–Fock canonical MO basis and the symmetrical OAO basis. The sMO and sOAO show the same quantity with a sharper cutoff (10^{-7}) and demonstrate the advantage to localized atomic basis functions at length scales as small as 10 \AA .

a cutoff of $\delta = 10^{-7}$, the localized character of the OAOs allows for a savings of on the order of 6×10^6 integrals with respect to the more delocalized canonical MOs. The s in the labeling of the orbital bases simply differentiates between two possible cutoffs. These dramatic differences begin to present with atomic chains as small as 10 \AA in length in this system with this basis set.

As an additional example, we consider linear alkane chains of increasing length. The results are displayed in Figure 2 and again display the dramatic advantages of preserving locality in the basis set. By the point one reaches 10 carbon atoms, a

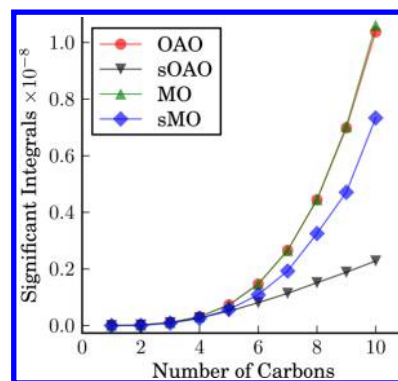


Figure 2. Number of significant (magnitude $> 10^{-15}$) spin–orbital integrals in the STO-3G basis set as a function of the number of carbons in a linear alkane chain for the Hartree–Fock canonical MO basis and the symmetrical OAO basis. The sMO and sOAO show the same quantity with a sharper cutoff (10^{-7}) and demonstrate the dramatic advantage to localized atomic basis even at this small atomic size.

savings of almost 10^8 integrals can be achieved at a truncation level of 10^{-7} .

Although localized basis sets provide a definitive scaling advantage in the medium–large size limit for molecules, one often finds that in the small-molecule limit, canonical MOs, the orbitals from the solution of the Hartree–Fock equations under the canonical condition, provide a more sparse representation. This is observed in the hydrogen and alkane chains studied here for the smallest molecule sizes, and the transition for this behavior will generally be basis-set-dependent. For example, in the alkane chains smaller than C_4H_{10} studied here, such as C_3H_8 , the number of significant integrals in the MO basis at a threshold of 10^{-7} is roughly 80% of that in the atomic orbital basis. The reason is that at smaller length scales, the “delocalized” canonical molecule orbitals have similar size to the more localized atomic orbitals but with the additional constraint of the canonical condition, a sufficient but not necessary condition for the solution of the Hartree–Fock equations that demands the Fock matrix be diagonal (as opposed to the looser variational condition of block-diagonal between the occupied and virtual spaces). A side effect of the canonical condition is that in the canonical MO basis, many of the h_{pqrs} terms for distinct indices are reduced in magnitude. However, there are not enough degrees of freedom present in the orbital rotations for this effect to persist to larger length scales, and as a result, local basis sets eventually become more advantageous. Moreover, it is known that at larger length scales, the canonical conditions tend to favor maximally delocalized orbitals, which can reduce the advantages of locality. These effects have been studied in some detail in the context of better orbital localizations by relaxing the canonical condition in Hartree–Fock and the so-called least-change Hartree–Fock method coupled with fourth-moment minimization.⁴⁶

Almost all algorithms designed for the study of quantum chemistry eigenstates on a quantum computer can be separated into two distinct parts, (1) state preparation and (2) energy estimation. For the purposes of analysis, it is helpful to treat the two issues separately, and in this section, we make the standard assumption in doing so that an oracle capable of producing good approximations to the desired eigenstates $|\Psi_i\rangle$ at unit cost is available. Under this assumption, energy estimation for a fixed desired precision ϵ is known to scale polynomially in the

size of the system for quantum chemistry; however, the exact scaling costs and trade-offs depend on the details of the method used. Here, we compare the costs and benefits of two prominent methods of energy estimation used in quantum computation for chemistry, quantum phase estimation and Hamiltonian averaging.

The first method used for the energy estimation of quantum chemical states on a quantum computer was quantum phase estimation.^{8,49,50} The method works by evolving the given quantum eigenstate $|\Psi_i\rangle$ forward under the system Hamiltonian H for a time T and reading out the accumulated phase, which can be easily mapped to the associated eigenenergy E_i . While the basic algorithm and its variations can have many different components, the cost is universally dominated by the coherent evolution of the system.

To evolve the system under the Hamiltonian, one must find a scalable way to implement the unitary operator $U = e^{-iHT}$. The standard procedure for accomplishing this task is the use of Suzuki–Trotter splitting,^{51,52} which approximates the unitary operator (at first order) as

$$\begin{aligned} U &= e^{-iHT} = (e^{-iH(T/m)})^m \\ &= (e^{-i(\sum_i H_i)\Delta t})^m \approx \left(\prod_i e^{-iH_i\Delta t}\right)^m \end{aligned} \quad (28)$$

where $\Delta t = T/m$ and H_i is a single term from the Bravyi–Kitaev transformed system Hamiltonian. Higher-order Suzuki–Trotter operator splittings and their benefits have been studied in the context of quantum simulation by Berry et al.⁵³ and by Whitfield et al.,⁵⁴ but we largely focus on the first-order formula in this work. If each of the simpler unitary operators $e^{-iH_i\Delta t}$ has a known gate decomposition, the total time evolution can be performed by chaining these sequences together.

The use of the Suzuki–Trotter splitting can be thought of as an evolution under an approximate Hamiltonian \tilde{H} , given by $e^{-i\tilde{H}T}$, whose eigenspectrum deviates from the original Hamiltonian by a factor depending on time step Δt . The precise dependence of this bias depends on the order of the Suzuki–Trotter expansion used. The total resolution, ϵ , in the energies of the approximate Hamiltonian \tilde{H} is determined by the total evolution time T . Thus, to achieve an accuracy of ϵ in the final energy, one must utilize a time step Δt small enough that the total bias is less than ϵ and a total run time T such that the resolution is better than ϵ . If the number of gates required to implement a single time step Δt is given by N_g , then the dominant cost of simulation (all of which must be done coherently) is given by

$$N_c = N_g \frac{T}{\Delta t} \quad (29)$$

The total evolution time T required to extract an eigenvalue to chemical precision $\epsilon_{\text{chem}} = 10^{-3}$ is typically set at the Fourier limit independent of molecular size and thus can be considered a constant for scaling analysis. We then focus on the number of gates per Suzuki–Trotter time step, N_g , and the time step Δt required to achieve the desired precision.

In a first order Suzuki–Trotter splitting, the number of gates per Trotter time step is given by the number of terms in the Hamiltonian multiplied by the number of gates required to implement a single elementary term for the form $e^{-iH_i\Delta t}$. The gates per elementary term can vary based on the particular integral; however, for simplicity in developing bounds, we consider this as constant here. The number of terms is known

from previous analysis in this work to scale as $O(M^2)$ or in the truly macroscopic limit $O(M)$. The number of gates required to implement a single elementary term depends on the transformation used from Fermionic to qubit operators. The Jordan–Wigner transformation⁵⁵ results in nonlocal terms that carry with them an overhead that scales as the number of qubits, which in this case will be $O(M)$. Although there have been developments in methods to use teleportation to perform these nonlocal operations in parallel⁹ and by improving the efficiency of the circuits computing the phases in the Jordan–Wigner transformation,¹⁹ these issues can also be alleviated by choosing the Bravyi–Kitaev transformation that carries an overhead only logarithmic in the number of qubits, $O(\log M)$.^{10,56} As a result, one expects the number of gates per Suzuki–Trotter time step N_g to scale as $O(M^2 \log M)$ or in a truly macroscopic limit $O(M \log M)$.

To complete the cost estimate with fixed total time T , one must determine how the required time step Δt scales with the size of the system. As mentioned above, the use of the Suzuki–Trotter decomposition for the time evolution of H is equivalent to evolution under an effective Hamiltonian $\tilde{H} = H + V$, where the size of the perturbation is determined by the order of the Suzuki–Trotter formula used and the size of the time step. Once the order of the Suzuki–Trotter expansion to be used has been determined, the requirement on the time step is such that the effect of V on the eigenvalue of interest is less than the desired accuracy in the final answer ϵ .

This has been explored previously,^{18,19} but we now examine this scaling in our context. To find V , one may expand the k th-order Suzuki–Trotter expansion of the evolution of \tilde{H} into a power series as well as the power series of the evolution operator $\exp[-i(H + V)\Delta t]$ and find the leading order term V . As a first result, we demonstrate that for a k th-order propagator, the leading perturbation on the ground-state eigenvalue for a nondegenerate system is $O(\Delta t)^{k+1}$.

Recall the power series expansion for the propagator

$$\exp[-i(H + V)\Delta t] = \sum_{j=0}^{\infty} \frac{(-i)^j}{j!} (H + V)^j (\Delta t)^j \quad (30)$$

The definition of a k th-order propagator is one that is correct through order k in the power series expansion. As such, when this power series is expanded, V must make no contribution in the terms until $O((\Delta t)^{k+1})$. For this to be possible, it is clear that V must depend on Δt . In order for it to vanish for the first k terms, V must be proportional to $(\Delta t)^k$. Moreover, due to the alternation of terms between imaginary and real at each order in the power series with the first term being imaginary, the first possible contribution is order $(\Delta t)^k$ and imaginary. As is common in quantum chemistry, we assume a nondegenerate and real ground state, and thus, the contribution to the ground-state eigenvalue is well approximated by first-order perturbation theory as

$$E^{(1)} = \langle \Psi_g | V | \Psi_g \rangle \quad (31)$$

However, as V is an imaginary Hermitian and the ground state is known to be real in quantum chemistry, this expectation value must vanish. Thus, the leading order perturbation to the ground-state eigenvalue is at worst the real term depending on $(\Delta t)^{k+1}$.

To get a more precise representation of V for a concrete example, we now consider the first-order ($k = 1$) Suzuki–

Trotter expansion. As expected, the leading order imaginary error term is found to be

$$V^{(0)} = \frac{\Delta t}{2} \sum_{j < k} i[H_j, H_k] \quad (32)$$

whose contribution must vanish due to it being an imaginary Hermitian term. Thus, we look to the leading contributing error depending on $(\Delta t)^2$, which has been obtained previously¹⁸ from a Baker–Campbell–Hausdorff (BCH) expansion to read

$$V^{(1)} = \frac{(\Delta t)^2}{12} \sum_{i \leq j} \sum_j \sum_{k < j} \left[H_i \left(1 - \frac{\delta_{ij}}{2} \right), [H_j, H_k] \right] \quad (33)$$

Thus, the leading order perturbation is given by third powers of the Hamiltonian operators. To proceed, we count the number of OEIs and TEIs separately as $N_{\text{int}}^{\text{OEI}}$ and $N_{\text{int}}^{\text{TEI}}$, respectively. Their maximal norm elements are similarly denoted by $h_{\text{max}}^{\text{OEI}}$ and $h_{\text{max}}^{\text{TEI}}$. From this, we can draw a worst case error bound on the perturbation of the eigenvalue given by

$$\begin{aligned} E^{(1)} &\leq \frac{(\Delta t)^2}{12} \sum_{i \leq j} \sum_j \sum_{k < j} \left| H_i \left(1 - \frac{\delta_{ij}}{2} \right), [H_j, H_k] \right| \\ &\leq (|h_{\text{max}}^{\text{OEI}}| N_{\text{int}}^{\text{OEI}} + |h_{\text{max}}^{\text{TEI}}| N_{\text{int}}^{\text{TEI}})^3 (\Delta t)^2 \\ &\leq (|\beta_{\text{max}}^{\text{OEI}}| M^{2/3} N_{\text{int}}^{\text{OEI}} + |h_{\text{max}}^{\text{TEI}}| N_{\text{int}}^{\text{TEI}})^3 (\Delta t)^2 \end{aligned} \quad (34)$$

where the first inequality follows from the triangle inequality and the second is a looser but simpler bound that may be used to elucidate the scaling behavior. Holding the looser bound to the desired precision in the final answer ϵ , this yields

$$\Delta t \leq \left[\frac{\epsilon}{(|\beta_{\text{max}}^{\text{OEI}}| M^{2/3} N_{\text{int}}^{\text{OEI}} + |h_{\text{max}}^{\text{TEI}}| N_{\text{int}}^{\text{TEI}})^3} \right]^{1/2} \quad (35)$$

We emphasize that this is a worst case bound from first-order perturbation theory, including no possible cancellation between Hamiltonian terms. Some preliminary work has been done numerically in establishing average cancellation between terms that shows that these worst case bounds are too pessimistic.¹⁸ Additionally, a rigorous bound not depending on perturbation theory has been previously derived.^{17,18} Continuing, we expect the total scaling under a first-order Suzuki–Trotter expansion using a Bravyi–Kitaev encoding to be bounded by

$$\begin{aligned} N_c &= N_g \frac{T}{\Delta t} \leq \frac{N_g}{\epsilon \Delta t} \\ &\leq \frac{(|\beta_{\text{max}}^{\text{OEI}}| M^{2/3} N_{\text{int}}^{\text{OEI}} + |h_{\text{max}}^{\text{TEI}}| N_{\text{int}}^{\text{TEI}})^{3/2} N_{\text{int}} \log M}{\epsilon^{3/2}} \end{aligned} \quad (36)$$

In the large size limit where the number of significant TEIs in a local basis set scales quadratically and the number of significant OEIs scales linearly, this may be bounded by

$$N_c \leq \kappa \frac{(|\beta_{\text{max}}^{\text{OEI}}| M^{5/3} + |h_{\text{max}}^{\text{TEI}}| M^2)^{3/2} (M^2 + M)}{\epsilon^{3/2} (\log M)^{-1}} \quad (37)$$

where κ is a positive constant that will depend on the basis set and this expression scales as $O(M^5 \log M)$ in the number of spin–orbital basis functions. The quantum phase estimation algorithm has been central in almost all algorithms for energy

estimation in quantum simulation. However, it has a significant practical drawback in that after state preparation, all of the desired operations must be performed coherently. A different algorithm for energy estimation has recently been introduced^{11,14} that lifts all but an $O(1)$ coherence time requirement after state preparation, making it amenable to implementation on quantum devices in the near future. We briefly review this approach, which we will call Hamiltonian averaging, and bound its costs in applications for quantum chemistry.

As in quantum phase estimation, in Hamiltonian averaging, one assumes the eigenstates $|\Psi_i\rangle$ are provided by some oracle. By use of either the Jordan–Wigner or Bravyi–Kitaev transformation, the Hamiltonian may be written as a sum of tensor products of Pauli operators. These transformations at worst conserve the number of independent terms in the Hamiltonian; thus, we may assume for our worst case analysis that the number of terms is fixed by N_{int} and the coefficients remain unchanged. From the provided copy of the state and transformed Hamiltonian, to obtain the energy, one simply performs the average

$$\langle \hat{H} \rangle = \sum_{i,j,k,\dots \in x,y,z} h_{ijk\dots} \langle \sigma_1^i \otimes \sigma_2^j \otimes \sigma_3^k \dots \rangle \quad (38)$$

by independent Pauli measurements on the provided state $|\Psi_i\rangle$ weighted by the coefficients $h_{ijk\dots}$ which are simply a relabeling of the previous TEIs for convenience with the transformed operators. As $|\Psi_i\rangle$ is an eigenstate, this average will correspond to the desired eigenvalue E_i with some error related to sampling that we now quantify.

Consider an individual term

$$X_{ijkl\dots} = h_{ijkl\dots} \sigma_1^i \otimes \sigma_2^j \otimes \sigma_3^k \dots \quad (39)$$

It is clear from the properties of qubit measurements that the full range of values that this quantity can take on is $[-h_{ijkl\dots}, h_{ijkl\dots}]$. As a result, we expect that the variance associated with this term can be bounded by

$$\text{Var}[X_{ijkl\dots}] \leq |h_{ijkl\dots}|^2 \quad (40)$$

Considering a representative element, namely, the maximum magnitude integral element h_{max} , we can bound the variance of \hat{H} as

$$\text{Var}[\hat{H}] \leq N_{\text{int}}^2 |h_{\text{max}}|^2 \quad (41)$$

The variance of the mean, which is the relevant term for our sampling error, comes from the central limit theorem and is bounded by

$$\text{Var}[\langle \hat{H} \rangle] \leq \frac{\text{Var}[\hat{H}]}{N} \quad (42)$$

where N is the number of independent samples taken of $\langle \hat{H} \rangle$. Collecting these results, we find

$$\begin{aligned} \text{Var}[\langle \hat{H} \rangle] &\leq \frac{\sum |h_{ijkl\dots}|^2}{N} \\ &\leq \frac{(|\beta_{\text{max}}^{\text{OEI}}| M^{2/3} N_{\text{int}}^{\text{OEI}} + |h_{\text{max}}^{\text{TEI}}| N_{\text{int}}^{\text{TEI}})^2}{N} \end{aligned} \quad (43)$$

Now setting the variance to the desired statistical accuracy ϵ^2 (which corresponds to a standard error of ϵ at a 68% confidence interval), we find that the number of independent samples expected, N_s , is bounded by

$$N_s \leq \frac{(|\beta_{\max}^{\text{OEI}}| M^{2/3} N_{\text{int}}^{\text{OEI}} + |h_{\max}^{\text{TEI}}| N_{\text{int}}^{\text{TEI}})^2}{\epsilon^2} \quad (44)$$

If a single independent sample of $\langle \hat{H} \rangle$ requires the measurement of each of the N_{int} quantities, then the bound on the total cost in the number of state preparations and measurements, N_{mv} , is

$$N_m \leq \frac{N_{\text{int}} (|\beta_{\max}^{\text{OEI}}| M^{2/3} N_{\text{int}}^{\text{OEI}} + |h_{\max}^{\text{TEI}}| N_{\text{int}}^{\text{TEI}})^2}{\epsilon^2} \quad (45)$$

which if one considers the large size limit, such that the number of TEIs scales quadratically and the number of OEIs scales linearly, we find

$$N_m \leq \kappa \frac{(M + M^2) (|\beta_{\max}^{\text{OEI}}| M^{5/3} + |h_{\max}^{\text{TEI}}| M^2)^2}{\epsilon^2} \quad (46)$$

where κ is a positive constant that depends upon the basis set. It is clear that this expression scales as $O(M^6)$ in the number of spin-orbital basis functions. We see from this that under the same maximum error assumptions, Hamiltonian averaging scales only marginally worse in the number of integrals and precision as compared to quantum phase estimation performed with a first-order Suzuki-Trotter expansion but has a coherence time requirement of $O(1)$ after each state preparation. Note that each measurement is expected to require single qubit rotations that scale as either $O(M)$ for the Jordan-Wigner transformation or $O(\log M)$ for the Bravyi-Kitaev transformation. However, we assume that these trivial single qubit rotations can be performed in parallel independent of the size of the system without great difficulty, and we thus do not consider this in our cost estimate. This method is a suitable replacement for quantum phase estimation in situations where coherence time resources are limited and good approximations to the eigenstates are readily available. Additional studies are needed to quantify the precise performance of the two methods beyond worst case bounds.

A central assumption for successful quantum phase estimation and typically any energy evaluation scheme is access to some oracle capable of producing good approximations to the eigenstate of interest, where a “good” approximation is typically meant to imply an overlap that is polynomial in the size of the system. Additionally, a supposed benefit of phase estimation over Hamiltonian averaging is that given such a good (but not perfect) guess, by projective measurement in the energy basis, in principle, one may avoid any bias in the final energy related to the initial state. Here, we examine this assumption in light of the van Vleck catastrophe,⁵⁷ which we review below, and examine the consequences for measurements of the energy by QPE and Hamiltonian averaging.

The van Vleck catastrophe⁵⁷ refers to an expected exponential decline in the quality of trial wave functions, as measured by overlap with the true wave function of a system, as a function of size. We study a simple case of the catastrophe here in order to frame the consequences for quantum computation. Consider a model quantum system consisting of a collection of N noninteracting two-level subsystems with subsystem Hamiltonians given by H_i . These subsystems have ground and excited eigenstates $|\psi_g^i\rangle$ and $|\psi_e^i\rangle$ with eigenenergies $E_g < E_e$, such that the total Hamiltonian is given by

$$H = \sum_i H_i \quad (47)$$

and eigenstates of the total Hamiltonian are formed from tensor products of the eigenstates of the subsystems. As such, the ground state of the full system is given by

$$|\Psi_g\rangle = \bigotimes_{i=0}^{N-1} |\psi_g^i\rangle \quad (48)$$

Now, suppose we want to measure the ground-state energy of the total system, but the oracle is only capable of producing trial states for each subsystem $|\psi_t^i\rangle$ such that $\langle \psi_t^i | \psi_g^i \rangle = \Delta$, where $|\Delta| < 1$. The resulting trial state for the whole system is

$$|\Psi_t\rangle = \bigotimes_{i=0}^{N-1} |\psi_t^i\rangle \quad (49)$$

From normalization of the two-level system, we may also write the trial state as

$$|\psi_t^i\rangle = \Delta |\psi_g^i\rangle + e^{-i\theta} \sqrt{1 - \Delta^2} |\psi_e^i\rangle \quad (50)$$

where $\theta \in [0, 2\pi)$. Moreover, from knowledge of the gap, one can find the expected energy on each subsystem, which is given by

$$\langle \psi_t^i | H_i | \psi_t^i \rangle = \Delta^2 E_g + (1 - \Delta^2) E_e \quad (51)$$

For the case of Hamiltonian averaging on the total system, the expected answer is given by

$$\begin{aligned} E &= \langle \Psi_t | H | \Psi_t \rangle \\ &= \sum_{i=0}^{N-1} \langle \psi_t^i | H_i | \psi_t^i \rangle \\ &= N(\Delta^2 E_g + (1 - \Delta^2) E_e) \end{aligned} \quad (52)$$

which yields an energy bias from the true ground state, ϵ_b , given by

$$\begin{aligned} \epsilon_b &= N(\Delta^2 E_g + (1 - \Delta^2) E_e) - N E_g \\ &= N(1 - \Delta^2)(E_e - E_g) \\ &= N(1 - \Delta^2)\omega \end{aligned} \quad (53)$$

where we denote the gap for each subsystem as $\omega = (E_e - E_g)$. As such, it is clear that the resulting bias is only linear in the size of the total system N .

Quantum phase estimation promises to remove this bias by projecting into the exact ground state. However, this occurs with a probability proportional to the square of the overlap of the input trial state with the target state. In this example, this is given by

$$|\langle \Psi_t | \Psi_g \rangle|^2 = |\Delta|^{2N} \quad (54)$$

which is exponentially small in the size of the system. That is, quantum phase estimation is capable of removing the bias exactly in this example noninteracting system but at a cost that is exponential in the size of the system. The expected cost of removing some portion of the bias may be calculated by considering the distribution of states and corresponding energies.

Consider first the probability of measuring an energy with a bias of $\epsilon(M) = M(1 - \Delta^2)\omega$. For this to happen, it is clear that exactly M of the subsystems in the measured state are in the excited state. It is clear that this is true for $\binom{N}{M}$ eigenstates, and

the square of the overlap with such an eigenstate is $(\Delta^2)^{N-M}(1 - \Delta^2)^M$ or

$$P(\epsilon(M)) = \binom{N}{M} (\Delta^2)^{N-M} (1 - \Delta^2)^M \quad (55)$$

which is clearly a binomial distribution. As a result, in the large N limit, this distribution is well-approximated by a Gaussian, and we may write

$$P(\epsilon) \approx \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2}\left(\frac{M - \bar{N}}{\sigma}\right)^2\right] \quad (56)$$

$$\bar{N} = N(1 - \Delta^2) \quad (57)$$

$$\sigma^2 = N\Delta^2(1 - \Delta^2) \quad (58)$$

Bringing this together, we find that the probability of measuring a bias of less than $\epsilon(M)$ is given by

$$\begin{aligned} P(\langle \epsilon \rangle) &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^M dM' \exp\left[-\frac{1}{2}\left(\frac{M' - \bar{N}}{\sigma}\right)^2\right] \\ &= \frac{1}{2} \left[\operatorname{erf}\left(\frac{M - \bar{N}}{\sqrt{2\sigma^2}}\right) + \operatorname{erf}\left(\frac{\bar{N}}{\sqrt{2\sigma^2}}\right) \right] \end{aligned} \quad (59)$$

where erf is again the error function.

Thus, the expected cost in terms of number of repetitions of the full phase estimation procedure to remove a bias of at least $\epsilon(M)$ from this model system is

$$\begin{aligned} C(\langle \epsilon(M) \rangle) &= \frac{1}{P(\langle \epsilon(M) \rangle)} \\ &= 2 \left[\operatorname{erf}\left(\frac{M - \bar{N}}{\sqrt{2\sigma^2}}\right) + \operatorname{erf}\left(\frac{\bar{N}}{\sqrt{2\sigma^2}}\right) \right]^{-1} \end{aligned} \quad (60)$$

We plot the expected cost function for a range of oracle guess qualities Δ on a modest system of $N = 100$ in Figure 3. From




Figure 3. A log–log plot of the expected cost in number of repetitions of measuring an energy with a bias $\epsilon(M)$ as a function of M in quantum phase estimation for different values of the oracle quality Δ . A system of $N = 100$ noninteracting subsystems is considered. A perfect, unbiased answer corresponds to $M = 0$ with expected cost $O(\Delta^{2N})$; however, to aid in visualization, this plot is provided only beyond $M = 1$. In general, one sees that depending on the oracle quality Δ , different fractions of the bias may be removed with ease, but there is always some threshold for imperfect guesses ($|\Delta| < 1$) such that there is an exponential growth in cost.

this, we see that the amount of bias that can feasibly be removed depends strongly on the quality of the oracle guess. Generically, we see that for any fixed imperfect guess on the subsystem level ($|\Delta| < 1$), there will be an exponential cost in phase estimation related to perfect removal of the bias.

This problem can be circumvented by improving the quality of the subsystem guesses as a function of system size. In particular, one can see that if $|\Delta|$ is improved as $(1 - (1/2N))$, then $|\Delta|^{2N}$ is $O(1)$. However, as the subsystems in a general case could be of arbitrary size, classical determination of a subsystem state of sufficient quality may scale exponentially in the required precision and thus system size. Moreover, one would not expect the problem to be easier in general cases where interactions between subsystems are allowed. As a result, further developments in variational methods,¹⁴ quantum cooling,⁵⁸ and adiabatic state preparation^{8,27,59} will be of key importance in this area. Moreover improvements in the ansatz used to prepare the wave function such as multiconfigurational self-consistent field (MCSCF) calculations^{26,27} or unitary coupled cluster (UCC)¹¹ will be integral parts of any practical quantum computing for quantum chemistry effort.

A complementary solution for the problem of molecular simulation on quantum computers is that of adiabatic quantum computation. It is not known to show the same direct dependence on the overlap of the initial guess state as QPE, which may allow it to solve different problems than the quantum phase estimation or variational quantum eigensolver in practice. In ref 59, Babbush et al. show how to scalably embed the eigenspectra of molecular Hamiltonians in a programmable physical system so that the adiabatic algorithm can be applied directly. In this scheme, the molecular Hamiltonian is first written in second quantization using Fermionic operators. This Hamiltonian is then mapped to a qubit Hamiltonian using the Bravyi–Kitaev transformation.^{10,56} The authors show that the more typical Jordan–Wigner transformation cannot be used to scalably reduce molecular Hamiltonians to two-local qubit interactions as the Jordan–Wigner transformation introduces linear locality overhead that translates to an exponential requirement in the precision of the couplings when perturbative gadgets are applied. Perturbative gadgets are used to reduce the Bravyi–Kitaev transformed Hamiltonian to a two-local programmable system with a restricted set of physical couplings. Finally, tunneling spectroscopy of a probe qubit⁶⁰ can be used to measure eigenvalues of the prepared state directly.

While the exact length of time that one must adiabatically evolve is generally unknown, Babbush et al. argue that the excited-state gap could shrink polynomially with the number of spin–orbitals when interpolating between exactly preparable noninteracting subsystems and the exact molecular Hamiltonian in which those subsystems interact. This would imply that adiabatic state preparation is efficient. Their argument is based on the observation that molecular systems are typically stable in their electronic ground states and that the natural processes that produce these states should be efficient to simulate with a quantum device. Subsequently, Veis and Pittner analyzed adiabatic state preparation for a set of small chemical systems and observed that for all configurations of these systems, the minimum gap occurs at the very end of the evolution when the state preparation is initialized in an eigenstate given by a CAS (complete active space) ground state.²⁷ The notion that the minimum gap could be bounded by the physical HOMO (highest occupied molecular orbital)–

LUMO (lowest unoccupied molecular orbital) gap lends support to the hypothesis put forward by Babbush et al.

In the adiabatic model of quantum computation, the structure of the final problem Hamiltonian (encoding the molecular eigenspectrum) determines experimental resource requirements. Because programmable many-body interactions are generally unavailable, we will assume that any experimentally viable problem Hamiltonian must be two-local. Any two-local Hamiltonian on n qubits can be expressed as

$$H = \alpha \cdot 1 + \sum_{i=1}^n \vec{\beta}_i \cdot \vec{\sigma}_i + \sum_{i=1}^{n-1} \sum_{j=i+1}^n \vec{\gamma}_{ij} \cdot (\vec{\sigma}_i \otimes \vec{\sigma}_j) \quad (61)$$

where $\vec{\sigma}_i = \langle \sigma_i^x, \sigma_i^y, \sigma_i^z \rangle$ is the vector of Pauli matrices on the i th qubit, $\alpha \in \mathbb{R}$ is a scalar, and $\vec{\beta}_i \in \mathbb{R}^3$ and $\vec{\gamma}_{ij} \in \mathbb{R}^9$ are vectors of coefficients for each possible term.

In addition to the number of qubits, the most important resources are the number of qubit couplings and the range of field values needed to accurately implement the Hamiltonian. Because local fields are relatively straightforward to implement, we are concerned with the number of two-local couplings

$$\sum_{i=1}^{n-1} \sum_{j=i+1}^n \text{card}(\vec{\gamma}_{ij}) \quad (62)$$

where $\text{card}(\vec{\gamma})$ is the number of nonzero terms in vector $\vec{\gamma}$. Because the effective molecular electronic structure Hamiltonian is realized perturbatively, there is a trade-off between the error in the eigenspectrum of the effective Hamiltonian, ϵ , and the strength of couplings that must be implemented experimentally. The magnitude of the perturbation is inversely related to the gadget spectral gap Δ , which is directly proportional to the largest term in the Hamiltonian

$$\max_{ij} \{ \|\vec{\gamma}_{ij}(\epsilon)\|_{\infty} \} \propto \Delta(\epsilon) \quad (63)$$

Thus, the smaller Δ is, the easier the Hamiltonian is to implement but the greater the error in the effective Hamiltonian. In general, there are other important resource considerations, but these are typically scale-invariant, for instance, the geometric locality of a graph or the set of allowed interaction terms. The Hamiltonian can be modified to fit such constraints using additional perturbative gadgets but typically at the cost of using more ancilla qubits that require greater coupling strength magnitudes.

The Bravyi–Kitaev transformation is crucial when embedding molecular electronic structure in two-local spin Hamiltonians due to the fact that this approach guarantees a logarithmic upper bound on the locality of the Hamiltonian. A loose upper bound (i.e., overestimation) for the number of qubits needed to gadgetize the molecular electronic Hamiltonian can be obtained by assuming that all terms have the maximum possible locality of $O(\log(M))$, where M is the number of spin-orbitals.

In general, the number of terms produced by the Bravyi–Kitaev transformation scales the same as the number of integrals in the electronic structure problem, $O(M^4)$; however, as pointed out in earlier, this bound can be reduced to $O(M^2)$ if a local basis is used and small integrals are truncated. Using the “bit-flip” gadgets of refs 61 and 62 to reduce M^2 terms of locality $\log(M)$, we would need $M^2 \log(M)$ ancillae. Because the number of ancilla qubits is always more than the number of

logical qubits for this problem, an upper-bound on the total number of qubits needed is $O(M^2 \log(M))$.

The number of couplings needed will be dominated by the number of edges introduced by ancilla systems required as penalty terms by the bit-flip gadgets. Each of the $O(M^2)$ terms is associated with a different ancilla system that contains a number of qubits equal to the locality of that term. Furthermore, all qubits within an ancilla system are fully connected. Thus, if we again assume that all terms have maximum locality, an upper bound on the number of couplers is $O(M^2 \log^2(M))$. On the basis of this analysis, the adiabatic approach to quantum chemistry has rather modest qubit and coupler requirements.

In ref 59, Babbush et al. reduce the locality of interaction terms using perturbative gadgets from the bit-flip family first introduced in ref 61 and later generalized by ref 62. In the Supporting Information presented in a later paper analyzing the scaling of gadget constructions,⁶³ it is shown that for bit-flip gadgets, $\lambda^{k+1}/\Delta^k = O(\epsilon)$ and

$$\max_{ij} \{ \|\vec{\gamma}_{ij}(\epsilon)\|_{\infty} \} = O\left(\frac{\lambda^k}{\Delta^{k-1}}\right) \quad (64)$$

Here, λ is the perturbative parameter, Δ is the spectral gap, ϵ is the error in the eigenspectrum, and $\vec{\gamma}_{ij}$ is the coefficient of the term to be reduced. Putting this together and representing the largest coupler value as γ , we find that $\Delta = \Omega(\epsilon^{-k} \gamma^k)$, where Ω is the “Big Omega” lower bound. Due to the Bravyi–Kitaev transformation, the locality of terms is bounded by $k = O(\log(M))$; thus, $\Delta = \Omega(\epsilon^{-\log(M)} \gamma^{-\log(M)})$.

Prior analysis from this Perspective indicates that the maximum integral size is bounded by $\gamma \leq |\beta_{\max}^{\text{OEI}}| M^{2/3}$. This gives us the bound

$$\Delta = \Omega(\epsilon^{-\log(M)} \|\beta_{\max}^{\text{OEI}}\| M^{2/3} \|\epsilon\|^{\log(M)}) \quad (65)$$

However, Δ also depends polynomially on M^2 , the number of terms present. Though known to be polynomial, it is extremely difficult to predict exactly how Δ depends on M^2 as applying gadgets to terms “in parallel” leads to “cross-gadget contamination”, which contributes at high orders in the perturbative expansion of the self-energy used to analyze these gadgets.⁶³ Without a significantly deeper analysis, we can only conclude that

$$\Delta = \Omega\left(\text{poly}(M) \left\| \frac{\beta_{\max}^{\text{OEI}} M^{2/3}}{\epsilon} \right\|^{\log(M)}\right) \quad (66)$$

This analysis indicates that the most significant challenge to implementing the adiabatic approach to quantum chemistry is the required range of coupler values, which is certain to span at least several orders of magnitude for nontrivial systems.

This calls attention to an important open question in the field of Hamiltonian gadgets: whether there exist “exact” gadgets that can embed the ground-state energy of arbitrary many-body target Hamiltonians without the use of perturbation theory. A positive answer to this conjecture would allow us to embed molecular electronic structure Hamiltonians without needing large spectral gaps. For entirely diagonal Hamiltonians, such gadgets are well-known in the literature^{64,65} but fail when terms do not commute.⁶³ Exact reductions have also been achieved for certain Hamiltonians. For instance, “frustration-free” gadgets have been used in proofs of the QMA-

Completeness of quantum satisfiability and in restricting the necessary terms for embedding quantum circuits in an instance of the local Hamiltonian problem.^{66–68}

In this work, we analyzed the impact on scaling for quantum chemistry on a quantum computer that results from consideration of locality of interactions and exploitation of local basis sets. The impact of locality has been exploited to great advantage for some time in traditional algorithms for quantum chemistry but has received relatively little attention in quantum computation thus far. From these considerations, we showed that in practical implementations of quantum phase estimation using a first-order Suzuki–Trotter approximation, one expects a scaling cost on the order of $O(M^5 \log M)$ with respect to number of spin-orbitals rather than more pessimistic estimates of $O(M^8)$ – $O(M^9)$ ^{17,19} or $O(M^{5.5})$ – $O(M^{6.5})$ ¹⁸ related to the use of unphysical random integral distributions or the restriction to molecules too small to observe the effects of physical locality. We believe that the combination of the algorithmic improvements suggested by Poulin and Hastings et al.^{18,19} with strategies that exploit locality presented here will result in even greater gains, and more work is needed in this area.

We also considered the cost of Hamiltonian averaging, an alternative to quantum phase estimation with minimal coherence time requirements beyond state preparation. This method has some overhead with respect to quantum phase estimation, scaling as $O(M^6)$ in the number of spin-orbitals, but has significant practical advantages in coherence time costs, as well as the ability to make all measurements in parallel. This method can at best give the energy of the state provided when oracle guesses are imperfect; however, it can easily be combined with a variational or adiabatic approach to improve the accuracy of the energy estimate. Moreover, while quantum phase estimation promises to be able to remove the bias of imperfect oracle guesses, we demonstrated how the cost of removal may strongly depend on how imperfect the guesses are. The consideration of physical locality greatly improves the outlook for quantum chemistry on a quantum computer, and in light of the goal of quantum chemistry to study physical systems rather than abstract constructs, it is correct to include this physical locality in any analysis pertaining to it.

The consideration of physical locality greatly improves the outlook for quantum chemistry on a quantum computer, and in light of the goal of quantum chemistry to study physical systems rather than abstract constructs, it is correct to include this physical locality in any analysis pertaining to it.

Finally, we analyzed the impact of locality on a complementary approach for quantum chemistry, namely, adiabatic quantum computation. This approach does not have a known direct dependence on the quality of guess states provided by an oracle and can in fact act as the state oracle for the other approaches discussed here.

In all cases, the consideration of physical locality greatly improves the outlook for quantum chemistry on a quantum computer, and in light of the goal of quantum chemistry to study physical systems rather than abstract constructs, it is correct to include this physical locality in any analysis pertaining to it. We believe that with these and other developments made in the area of quantum computation, quantum chemistry remains one of the most promising applications for exceeding the capabilities of current classical computers.

AUTHOR INFORMATION

Corresponding Author

*E-mail: aspuru@chemistry.harvard.edu.

Notes

The authors declare no competing financial interest.

Biographies

Jarrod R. McClean is a Department of Energy Computational Graduate Science Fellow currently studying with Alán Aspuru-Guzik as a Ph.D. student at Harvard University. He is interested in the intersection between traditional and quantum computation for applications to chemistry.

Ryan Babbush is a Ph.D. student in Chemical Physics studying quantum information with Alán Aspuru-Guzik at Harvard University. As an undergraduate, he doubled majored in chemistry and physics at Carleton College and will go on to join Google's quantum computing research team following the defense of his dissertation next year.

Peter J. Love is an Associate Professor of Physics at Haverford College.

Alán Aspuru-Guzik is Professor of Chemistry and Chemical Biology of Harvard University. His research group works at the intersection of quantum information and chemistry. He has been working on quantum computing for quantum chemistry since 2005. He has received numerous awards, including a Technology Review 35 Innovators under 35 award. <http://aspuru.chem.harvard.edu>

ACKNOWLEDGMENTS

J.M. is supported by the Department of Energy Computational Science Graduate Fellowship under Grant Number DE-FG02-97ER25308. P.J.L. acknowledges the National Science Foundation Grant Number PHY-0955518. A.A.-G. and P.J.L. appreciate the support of the Air Force Office of Scientific Research under Award No. FA9550-12-1-0046. R.B. and A.A.-G. are grateful to the National Science Foundation for Award No. CHE-1152291.

REFERENCES

- (1) Thogersen, L.; Olsen, J. A Coupled Cluster and Full Configuration Interaction Study of CN and CN⁻. *Chem. Phys. Lett.* **2004**, *393*, 36–43.
- (2) Feynman, R. P. Simulating Physics with Computers. *Int. J. Theor. Phys.* **1982**, *21*, 467–488.
- (3) Shor, P. W. Algorithms for Quantum Computation Discrete Logarithms and Factoring. . *Proceedings 35th Annual Symposium on Foundations of Computer Science*; Santa Fe, NM, 1994; pp 124–134.
- (4) Neven, H.; Rose, G.; Macready, W. G. Image Recognition with an Adiabatic Quantum Computer I. Mapping to Quadratic Unconstrained Binary Optimization. 2008, arXiv: quant-ph/0804.4457. arXiv.org e-Print archive.
- (5) Babbush, R.; Denchev, V.; Ding, N.; Isakov, S.; Neven, H. Construction of Non-Convex Polynomial Loss Functions for Training

a Binary Classifier with Quantum Annealing. 2014, arXiv: quant-ph/1406.4203. arXiv.org e-Print archive.

(6) Perdomo-Ortiz, A.; Dickson, N.; Drew-Brook, M.; Rose, G.; Aspuru-Guzik, A. Finding Low-Energy Conformations of Lattice Protein Models by Quantum Annealing. *Sci. Rep.* **2012**, *2*, 571/1–571/7.

(7) Babbush, R.; Perdomo-Ortiz, A.; O’Gorman, B.; Macready, W.; Aspuru-Guzik, A. *Advances in Chemical Physics*; Wiley-Blackwell: Hoboken, NJ, 2014; pp 201–244.

(8) Aspuru-Guzik, A.; Dutoi, A. D.; Love, P. J.; Head-Gordon, M. Simulated Quantum Computation of Molecular Energies. *Science* **2005**, *309*, 1704–1707.

(9) Jones, N. C.; Whitfield, J. D.; McMahon, P. L.; Yung, M.-H.; Van Meter, R.; Aspuru-Guzik, A.; Yamamoto, Y. Faster Quantum Chemistry Simulation on Fault-Tolerant Quantum Computers. *New J. Phys.* **2012**, *14*, 115023.

(10) Seeley, J. T.; Richard, M. J.; Love, P. J. The Bravyi–Kitaev Transformation for Quantum Computation of Electronic Structure. *J. Chem. Phys.* **2012**, *137*, 224109/1–224109/16.

(11) Yung, M.; Casanova, J.; Mezzacapo, A.; McClean, J.; Lamata, L.; Aspuru-Guzik, A.; Solano, E. From Transistor to Trapped-Ion Computers for Quantum Chemistry. *Sci. Rep.* **2014**, 3589/1–3589/7.

(12) Lanyon, B. P.; Whitfield, J. D.; Gillett, G. G.; Goggin, M. E.; Almeida, M. P.; Kassal, I.; Biamonte, J. D.; Mohseni, M.; Powell, B. J.; Barbieri, M. Towards Quantum Chemistry on a Quantum Computer. *Nat. Chem.* **2010**, *2*, 106–111.

(13) Aspuru-Guzik, A.; Walther, P. Photonic Quantum Simulators. *Nat. Phys.* **2012**, *8*, 285–291.

(14) 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 Quantum Processor. *Nat. Commun.* **2014**, *5*, 4213/1–4213/7.

(15) Wang, Y.; Dolde, F.; Biamonte, J.; Babbush, R.; Bergholm, V.; Yang, S.; Jakobi, I.; Neumann, P.; Aspuru-Guzik, A.; Whitfield, J. D. et al. Quantum Simulation of Helium Hydride in a Solid-State Spin Register. 2014, arXiv: quant-ph/1405.2696. arXiv.org e-Print archive.

(16) Whitfield, J. D.; Biamonte, J.; Aspuru-Guzik, A. Simulation of Electronic Structure Hamiltonians Using Quantum Computers. *Mol. Phys.* **2011**, *109*, 735–750.

(17) Wecker, D.; Bauer, B.; Clark, B. K.; Hastings, M. B.; Troyer, M. Gate-Count Estimates for Performing Quantum Chemistry on Small Quantum Computers. *Phys. Rev. A* **2014**, *90*, 022305.

(18) Poulin, D.; Hastings, M. B.; Wecker, D.; Wiebe, N.; Doherty, A. C.; Troyer, M. The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry. 2014, arXiv: quant-ph/1406.4920. arXiv.org e-Print archive.

(19) Hastings, M. B.; Wecker, D.; Bauer, B.; Troyer, M. Improving Quantum Algorithms for Quantum Chemistry. 2014, arXiv: quant-ph/1403.1539. arXiv.org e-Print archive.

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

(21) Goedecker, S. Linear Scaling Electronic Structure Methods. *Rev. Mod. Phys.* **1999**, *71*, 1085.

(22) Artacho, E.; Sánchez-Portal, D.; Ordejón, P.; García, A.; Soler, J. M. Linear-Scaling Ab-Initio Calculations for Large and Complex Systems. *Phys. Status Solidi B* **1999**, *215*, 809–817.

(23) Helgaker, T.; Jørgensen, P.; Olsen, J. *Molecular Electronic Structure Theory*; Wiley: Sussex, U.K., 2002.

(24) Shao, Y.; Molnar, L. F.; Jung, Y.; Kussmann, J.; Ochsenfeld, C.; Brown, S. T.; Gilbert, A. T.; Slipchenko, L. V.; Levchenko, S. V.; O’Neill, D. P. Advances in Methods and Algorithms in a Modern Quantum Chemistry Program Package. *Phys. Chem. Chem. Phys.* **2006**, *8*, 3172–3191.

(25) Bowler, D.; Miyazaki, T. O(N) Methods in Electronic Structure Calculations. *Rep. Prog. Phys.* **2012**, *75*, 036503.

(26) Wang, H.; Kais, S.; Aspuru-Guzik, A.; Hoffmann, M. R. Quantum Algorithm for Obtaining the Energy Spectrum of Molecular Systems. *Phys. Chem. Chem. Phys.* **2008**, *10*, 5388–5393.

(27) Veis, L.; Pittner, J. Adiabatic State Preparation Study of Methylene. *J. Chem. Phys.* **2014**, *140*, 214111/1–214111/13.

(28) Buluta, I.; Nori, F. Quantum Simulators. *Science* **2009**, *326*, 108–111.

(29) Georgescu, I. M.; Ashhab, S.; Nori, F. Quantum Simulation. *Rev. Mod. Phys.* **2014**, *86*, 153–185.

(30) Kassal, I.; Whitfield, J. D.; Perdomo-Ortiz, A.; Yung, M.-H.; Aspuru-Guzik, A. Simulating Chemistry Using Quantum Computers. *Annu. Rev. Phys. Chem.* **2011**, *62*, 185–207.

(31) Kais, S. *Advances in Chemical Physics, Quantum Information and Computation for Chemistry*; John Wiley & Sons: New York, 2014; Vol. 154.

(32) Kassal, I.; Aspuru-Guzik, A. Quantum Algorithm for Molecular Properties and Geometry Optimization. *J. Chem. Phys.* **2009**, *131*, 224102/1–224102/6.

(33) Kassal, I.; Jordan, S. P.; Love, P. J.; Mohseni, M.; Aspuru-Guzik, A. Polynomial-Time Quantum Algorithm for the Simulation of Chemical Dynamics. *Proc. Natl. Acad. Sci. U.S.A.* **2008**, *105*, 18681–18686.

(34) Ward, N. J.; Kassal, I.; Aspuru-Guzik, A. Preparation of Many-Body States for Quantum Simulation. *J. Chem. Phys.* **2009**, *130*, 194105/1–194105/9.

(35) Welch, J.; Greenbaum, D.; Mostame, S.; Aspuru-Guzik, A. Efficient Quantum Circuits for Diagonal Unitaries without Ancillas. *New J. Phys.* **2014**, *16*, 033040.

(36) Lu, D.; Xu, N.; Xu, R.; Chen, H.; Gong, J.; Peng, X.; Du, J. Simulation of Chemical Isomerization Reaction Dynamics on a NMR Quantum Simulator. *Phys. Rev. Lett.* **2011**, *107*, 020501.

(37) Veis, L.; Višňák, J.; Fleig, T.; Knecht, S.; Saue, T.; Visscher, L.; Pittner, J. C. V. Relativistic Quantum Chemistry on Quantum Computers. *Phys. Rev. A* **2012**, *85*, 030304.

(38) Toloui, B.; Love, P. J. Quantum Algorithms for Quantum Chemistry based on the Sparsity of the CI-matrix. 2013, arXiv: quant-ph/1312.2579. arXiv.org e-Print archive.

(39) Williamson, A.; Hood, R. Q.; Grossman, J. Linear-Scaling Quantum Monte Carlo Calculations. *Phys. Rev. Lett.* **2001**, *87*, 246406.

(40) Aspuru-Guzik, A.; Salomón-Ferrer, R.; Austin, B.; Lester, W. A. A Sparse Algorithm for the Evaluation of the Local Energy in Quantum Monte Carlo. *J. Comput. Chem.* **2005**, *26*, 708–715.

(41) Werner, H.-J.; Pflger, K. *Annual Reports in Computational Chemistry*; Elsevier {BV}: Amsterdam, The Netherlands, 2006; pp 53–80.

(42) Ochsenfeld, C.; Kussmann, J.; Lambrecht, D. S. *Reviews in Computational Chemistry*; Wiley-Blackwell: Hoboken, NJ, 2007; pp 1–82.

(43) Zalesny, R.; Papadopoulos, M. G.; Mezey, P. G.; Leszczynski, J., Eds. *Linear-Scaling Techniques in Computational Chemistry and Physics*; Springer: The Netherlands, 2011.

(44) Löwdin, P. On the Nonorthogonality Problem Connected with the Use of Atomic Wave Functions in the Theory of Molecules and Crystals. *J. Chem. Phys.* **1950**, *18*, 365–375.

(45) Mayer, I. On Löwdin’s Method of Symmetric Orthogonalization. *Int. J. Quantum Chem.* **2002**, *90*, 63–65.

(46) Ziolkowski, M.; Jansik, B.; Jørgensen, P.; Olsen, J. Maximum Locality in Occupied and Virtual Orbital Spaces Using a Least-Change Strategy. *J. Chem. Phys.* **2009**, *131*, 124112.

(47) Harrison, R. J.; Fann, G. I.; Yanai, T.; Gan, Z.; Beylkin, G. Multiresolution Quantum Chemistry: Basic Theory and Initial Applications. *J. Chem. Phys.* **2004**, *121*, 11587.

(48) Hachmann, J.; Cardoen, W.; Chan, G. K.-L. Multireference Correlation in Long Molecules with the Quadratic Scaling Density Matrix Renormalization Group. *J. Chem. Phys.* **2006**, *125*, 144101/1–144101/12.

(49) Kitaev, A. Y. Quantum Measurements and the Abelian Stabilizer Problem. 1995, arXiv: quant-ph/9511026. arXiv.org e-Print archive.

(50) Abrams, D. S.; Lloyd, S. Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors. *Phys. Rev. Lett.* **1999**, *83*, 5162–5165.

- (51) Trotter, H. F. On the Product of Semi Groups of Operators. *Proc. Am. Math. Soc.* **1959**, *10*, 545–551.
- (52) Suzuki, M. General Decomposition Theory of Ordered Exponentials. *Proc. Jpn. Acad., Ser. B* **1993**, *69*, 161–166.
- (53) Berry, D.; Ahokas, G.; Cleve, R.; Sanders, B. Efficient Quantum Algorithms for Simulating Sparse Hamiltonians. *Commun. Math. Phys.* **2007**, *270*, 359–371.
- (54) Whitfield, J.; Biamonte, J.; Aspuru-Guzik, A. Quantum Computing Resource Estimate of Molecular Energy Simulation. 2010, arXiv: quant-ph/1001.3855v1. arXiv.org e-Print archive.
- (55) Jordan, P.; Wigner, E. Über das Paulische Äquivalenzverbot. *Z. Phys.* **1928**, *47*, 631–651.
- (56) Bravyi, S. B.; Kitaev, A. Y. Fermionic Quantum Computation. *Ann. Phys.* **2002**, *298*, 210–226.
- (57) van Vleck, J. H. Nonorthogonality and Ferromagnetism. *Phys. Rev.* **1936**, *49*, 232–240.
- (58) Xu, J.-S.; Yung, M.-H.; Xu, X.-Y.; Boixo, S.; Zhou, Z.-W.; Li, C.-F.; Aspuru-Guzik, A.; Guo, G.-C. Demon-Like Algorithmic Quantum Cooling and Its Realization with Quantum Optics. *Nat. Photonics* **2014**, *8*, 113–118.
- (59) Babbush, R.; Love, P. J.; Aspuru-Guzik, A. Adiabatic Quantum Simulation of Quantum Chemistry. 2013, arXiv: quant-ph/1311.3967. arXiv.org e-Print archive.
- (60) Berkley, A. J.; Przybysz, A. J.; Lanting, T.; Harris, R.; Dickson, N.; Altomare, F.; Amin, M. H.; Bunyk, P.; Enderud, C.; Hoskinson, E.; et al. Tunneling Spectroscopy Using a Probe Qubit. *Phys. Rev. B* **2013**, *87*, 020502.
- (61) Kempe, J.; Kitaev, A.; Regev, O. The Complexity of the Local Hamiltonian Problem. *SIAM J. Comput.* **2006**, *35*, 1070–1097.
- (62) Jordan, S. P.; Farhi, E. Perturbative Gadgets at Arbitrary Orders. *Phys. Rev. A* **2008**, *77*, 1–8.
- (63) Cao, Y.; Babbush, R.; Biamonte, J.; Kais, S. Towards Experimentally Realizable Hamiltonian Gadgets. 2013, arXiv: 1311.2555. arXiv.org e-Print archive.
- (64) Biamonte, J. D. Non-Perturbative k-Body to Two-Body Commuting Conversion Hamiltonians and Embedding Problem Instances into Ising Spins. *Phys. Rev. A* **2008**, *77*, 1–8.
- (65) Babbush, R.; O’Gorman, B.; Aspuru-Guzik, A. Resource Efficient Gadgets for Compiling Adiabatic Quantum Optimization Problems. *Ann. Phys. (Berlin)* **2013**, *525*, 877–888.
- (66) Nagaj, D. Fast Universal Quantum Computation with Railroad-Switch Local Hamiltonians. *J. Math. Phys.* **2010**, *51*, 2201.
- (67) Gosset, D.; Nagaj, D. Quantum 3-SAT is QMA1-complete. 2013, arXiv: quant-ph/1302.0290. arXiv.org e-Print archive.
- (68) Childs, A. M.; Gosset, D.; Webb, Z. The Bose-Hubbard model is QMA-complete. 2013, arXiv: quant-ph/1311.3297. arXiv.org e-Print archive.