

Elucidating Reaction Mechanisms on Quantum Computers

Markus Reiher,¹ Nathan Wiebe,² Krysta M. Svore,² Dave Wecker,² and Matthias Troyer^{3,2,4}

¹*Laboratorium für Physikalische Chemie, ETH Zurich,
Valdimir-Prelog-Weg 2, 8093 Zurich, Switzerland*

²*Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA*

³*Theoretische Physik and Station Q Zurich, ETH Zurich, 8093 Zurich, Switzerland*

⁴*Station Q, Microsoft Research, Santa Barbara, CA 93106-6105, USA*

We show how a quantum computer can be employed to elucidate reaction mechanisms in complex chemical systems, using the open problem of biological nitrogen fixation in nitrogenase as an example. We discuss how quantum computers can augment classical-computer simulations for such problems, to significantly increase their accuracy and enable hitherto intractable simulations. Detailed resource estimates show that, even when taking into account the substantial overhead of quantum error correction, and the need to compile into discrete gate sets, the necessary computations can be performed in reasonable time on small quantum computers. This demonstrates that quantum computers will realistically be able to tackle important problems in chemistry that are both scientifically and economically significant.

Chemical reaction mechanisms are networks of molecular structures representing short- or long-lived intermediates connected by transition structures. At its core, this surface is determined by the nuclear coordinate dependent electronic energy, i.e., by the Born–Oppenheimer surface, which can be supplemented by nuclear-motion corrections. The relative energies of all stable structures on this surface determine the relative thermodynamical stability. Differences of the energies of these local minima to those of the connecting transition structures determine the rates of interconversion, i.e., the chemical kinetics of the process. Reaction rates depend on rate constants that are evaluated from these latter energy differences entering the argument of exponential functions. Hence, very accurate energies are required for the reliable evaluation of the rate constants.

The detailed understanding and prediction of complex reaction mechanisms such as transition-metal catalyzed chemical transformations therefore requires highly accurate electronic structure methods. However, the electron correlation problem remains, despite decades of progress [1], one of the most vexing problems in quantum chemistry. While approximate approaches, such as density functional theory [2], are very popular and have successfully been applied to all kinds of correlated electronic structures, their accuracy is too low for truly reliable quantitative predictions (see, e.g., Refs. [3, 4]). On classical computers, molecules with much less than a hundred strongly correlated electrons are already out of reach for systematically improvable ab initio methods that could achieve the required accuracy.

The apparent intractability of accurate simulations for such quantum systems led Richard Feynmann to propose quantum computers (see Appendix A for a brief review of quantum computing). The promise of exponential speedups for quantum simulation on quantum computers was first investigated by Lloyd [5] and Zalka [6] and was directly applied to quantum chemistry by Lidar, Aspuru–Guzik and others [7–11]. Quantum chemistry simulation has remained an active area within quantum algorithm

development, with ever more sophisticated methods being employed to reduce the costs of quantum chemistry simulation [12–20].

The promise of exponential speedups for the electronic structure problem has led many to suspect that quantum computers will one day revolutionize chemistry and materials science. However, a number of important questions remain. Not the least of these is the question of how exactly to use a quantum computer to solve an important problem in chemistry. The inability to point to a clear use case complete with resource and cost estimates is a major drawback; after all, even an exponential speedup may not lead to a useful algorithm if a typical, practical application requires an amount of time and memory so large that it would still be beyond the reach of even a quantum computer.

Here, we demonstrate for an important prototypical chemical system how a quantum computer would be used in practice to address an open problem and to estimate how large and how fast a quantum computer would have to be to perform such calculations within a reasonable amount of time. This sets a target for the type and size of quantum device that we would like to emerge from existing research and further gives confidence that quantum simulation will be able to provide answers to problems that are both scientifically and economically impactful.

The chemical process that we consider in this work is that of biological nitrogen fixation by the enzyme nitrogenase [22]. This enzyme accomplishes the remarkable transformation of turning a dinitrogen molecule into two ammonia molecules (and one dihydrogen molecule) under ambient conditions. Whereas the industrial Haber–Bosch catalyst requires high temperatures and pressures and is therefore energy intensive, the active site of Mo-dependent nitrogenase, the iron molybdenum cofactor (FeMoco) [23, 24], can split the dinitrogen triple bond at room temperature and standard pressure. Mo-dependent nitrogenase consists of two subunits, the Fe protein, a homodimer, and the MoFe protein, an $\alpha_2\beta_2$ tetramer. Figure 1 shows the MoFe protein of nitrogenase on the left




FIG. 1. X-ray crystal structure 4WES [21] of the nitrogenase MoFe protein from *Clostridium pasteurianum* taken from the protein data base (left; the backbone is colored in green and hydrogen atoms are not shown), the close protein environment of the FeMoco (center), and the structural model of FeMoco considered in this work (right; C in gray, O in red, H in white, S in yellow, N in blue, Fe in brown, and Mo in cyan).

and the FeMoco buried in this protein on the right. Despite the importance of this process for fertilizer production that makes nitrogen from air accessible to plants, the mechanism of nitrogen fixation at FeMoco is not known. Experiment has not yet been able to provide sufficient details on the chemical mechanism and theoretical attempts are hampered by intrinsic methodological limitations of traditional quantum chemical methods. Nitrogen fixation also turned out to be a true challenge for synthetic approaches and only three homogeneous catalysts working under ambient conditions have been discovered so far [25–27]. All of them decompose under reaction conditions as highlighted by an embarrassingly low turnover number of less than a dozen.

I. QUANTUM CHEMICAL METHODS FOR MECHANISTIC STUDIES

Standard concepts

At the heart of any chemical process is its mechanism, the elucidation of which requires the identification of all relevant stable intermediates and transition states and the calculation of their properties. The latter allow for experimental verification, whereas energies associated with these structures provide access to thermodynamic stability and kinetic accessibility. In general, a multitude of charge and spin states need to be explicitly calculated in search for the relevant ones that make the whole chemical process viable. This can lead to thousands of elementary reaction steps [28] whose reaction energies must be reliably calculated.

In the case of nitrogenase, numerous protonated intermediates of dinitrogen-coordinating FeMoco and subsequently reduced intermediates in different charge and spin states are feasible and must be assessed with respect to their relative energy. Especially kinetic mod-

elling poses tight limits on the accuracy of activation energies entering the argument of exponentials in rate expressions.

Although most chemical processes are local and therefore involve only a rather small number of atoms (usually one or two bonds are broken or formed at a time), they can be modulated by environmental effects (e.g., a protein matrix or a solvent). For the consideration of such environment effects, many different embedding methods [29–33] are available. They incorporate environment terms, ranging from classical electrostatics to full quantum descriptions, into the Hamiltonian to be diagonalized.

For nitrogenase, an electrostatic quantum-mechanical/molecular-mechanical (QM/MM) model that captures the embedding of FeMoco into the protein pocket of nitrogenase can properly account for the energy modulation of the chemical transformation steps at FeMoco. Accordingly, we consider a structural model for the active site of nitrogenase (Fig. 1 left) carrying only models of the anchoring groups of the protein, which represents a suitable QM part in such calculations. To study this bare model is no limitation as it does not at all affect our feasibility analysis (because electrostatic QM/MM embedding does not affect the number of orbitals considered for the wave function construction). We carried out (full) molecular structure optimizations with density functional theory (DFT) methods of this FeMoco model in different charge and spin states in order to not base our analysis on a single electronic structure. While our FeMoco model describes the resting state, binding of a small molecule such as dinitrogen, dihydrogen, diazene, or ammonia will not decisively increase or reduce the complexity of its electronic structure. See Appendix H for more details.

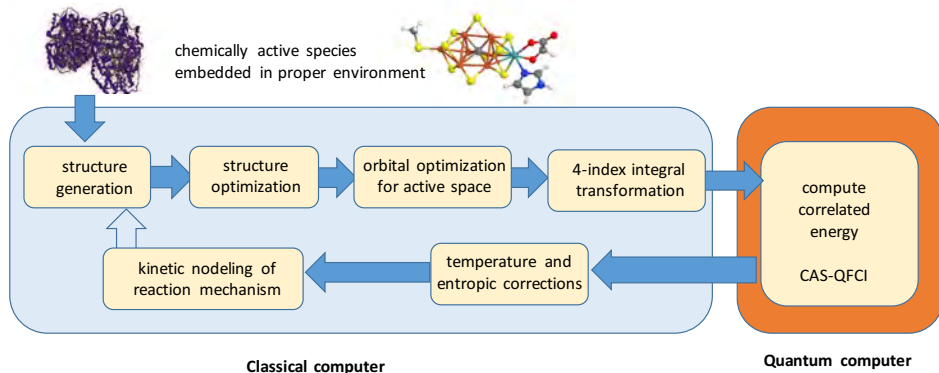


FIG. 2. Generic flowchart of a computational reaction-mechanism elucidation with a quantum computer part that delivers a quantum full configuration interaction (QFCI) energy in a (restricted) complete active orbital space (CAS). Once a structural model of the active chemical species (here FeMoco, top right) embedded in a suitable environment (the metalloprotein, top left) is chosen, structures of potential intermediates can be set up and optimized. Molecular orbitals are then optimized for a suitably chosen Fock operator. A four-index transformation from the atomic orbital to the molecular basis produces all integrals required for the second-quantized Hamiltonian. Once the quantum computer produces the (ground-state) energy of this Hamiltonian, it can be supplemented by corrections that consider nuclear motion effects to yield enthalpic and entropic quantities at a given temperature according to standard protocols (e.g., from DFT calculations). The temperature-corrected energy differences between stable intermediates and transition structures then enter rate expressions for kinetic modeling. For complex chemical mechanisms, this modeling might point to the exploration of additional structures.

Elucidating reaction mechanisms

Understanding a complex chemical mechanism requires the exploration of a potential energy landscape that results from the Born–Oppenheimer approximation. This approximation assigns an electronic energy to every molecular structure. The accurate calculation of this energy is the pivotal challenge, here considered by quantum computing. Characteristic molecular structures are optimized to provide local minimum structures indicating stable intermediates and first-order saddle points representing transition structures. The electronic energy differences for elementary steps that connect two minima through a transition structure enter expressions for rate constants by virtue of Eyring’s absolute rate theory (ART) which is a detailed quantum- and statistical-mechanical interpretation of the phenomenological rate expression by Arrhenius. Although more information on the potential energy surface as well as dynamic and quantum effects may be taken into account, ART is accurate even for large molecular systems such as enzymes [34, 35]. These rate constants then enter a kinetic description of all elementary steps that ultimately provide a complete picture of the chemical mechanism under consideration.

Exact diagonalization methods in chemistry

If the frontier orbital region around the Fermi level of a given molecular structure is dense, as is the case in π -conjugated molecules and open-shell transition metal complexes, then so-called strong *static* electron cor-

relation plays a decisive role already in the ground state. This is even more pronounced for electronically excited states relevant in photophysical and photochemical processes such as light harvesting for clean energy applications. Such situations require multi-configurational methods of which the complete-active-space self-consistent-field (CASSCF) approach has been established as a well-defined model that also serves as the basis for more advanced approaches [36].

CAS-type approaches require the selection of orbitals for the CAS, usually from those around the Fermi energy. This CAS selection is an art that, however, may be automatized [37, 38]. While CAS-type methods well account for static electron correlation, the remaining dynamic correlation is decisive for quantitative results. A remaining major drawback of exact-diagonalization schemes therefore is to include the contribution of all neglected virtual orbitals.

CASSCF is traditionally implemented as an exact diagonalization method, which limits its applicability to 18 electrons in 18 (spatial) orbitals because of the steep scaling of many-electron basis states with the number of electrons and orbitals [39]. The polynomially scaling density matrix renormalization group (DMRG) algorithm [40] can push this limit to about 100 spatial orbitals. This, however, also comes at the cost of an iterative procedure whose convergence for strongly correlated molecules is, due to the matrix product state representation of the electronic wave function, neither easy to achieve nor is it guaranteed. The groundstate energies output by quantum computer simulations are not affected by these issues, but present other accuracy considerations that we quantify below.

Ways quantum computers will help solve these problems

As long as molecular structure optimizations cannot be efficiently implemented on a quantum computer, this task can be accomplished with standard DFT approaches. DFT optimized molecular structures are in general reliable, even if the corresponding energies are affected by large uncontrollable errors. The latter problem can then be solved by a quantum computer that implements a multi-configurational (CAS-CI) wave function model to access truly large active orbital spaces. The orbitals for this model do not necessarily need to be optimized as natural orbitals can be taken from an unrestricted Hartree-Fock [41] or small-CAS CASSCF calculation. Clearly, the missing dynamic correlation needs to be implemented. This can be done in a 'perturb-then-diagonalize' fashion before the quantum computations start or in a 'diagonalize-then-perturb' fashion, where the quantum computer is used to produce the higher-order reduced density matrices required. The former approach, i.e., built-in dynamic electron correlation, is considerably more advantageous as no wave-function-derived quantities need to be calculated. One option for this approach is, for example, to consider dynamic correlation through DFT that avoids any double counting effects by virtue of range separation as has already been successfully studied for CASSCF and DMRG [42, 43]. Fig. 2 presents a flowchart that describes the steps of a quantum-computer assisted chemical mechanism exploration. Moreover, the quantum computer results can be used for the validation and improvement of parametrized approaches such as DFT in order to improve on the latter for the massive pre-screening of structures and energies.

II. QUANTUM SIMULATION OF QUANTUM CHEMICAL SYSTEMS

From among a number of different approaches to determine energies of ground states on a quantum computer [13, 44], we focus on quantum phase estimation (QPE). If we take the time evolution of an eigenstate to be $e^{-iHt}|n\rangle = e^{-iE_n t}|n\rangle$ then the task of QPE is to learn the phase $\phi = E_n t$ in a manner analogous to a Mach-Zehnder interferometer. If QPE is applied to a trial state that is a superposition of eigenvectors, it collapses the state to one of the eigenstates and will collapse it to the ground state with high probability if the trial state has a large overlap with the ground state. See Appendix F for a discussion of state preparation approaches in quantum simulation.

In practice, e^{-iHt} must be decomposed in QPE into elementary operations. There are several methods known for achieving this [45–47]; here we focus on the Trotter–Suzuki (TS) decompositions and specifically the second-order TS formula, which for the second-quantized Coulomb Hamiltonian takes the form

$$e^{-iHt} = \left(\prod_{j=1}^M e^{-iH_j t/2r} \prod_{j=M}^1 e^{-iH_j t/2r} \right)^r + O(t^3/r^2).$$

In the limit $r \rightarrow \infty$ this approximation becomes exact. Well-known methods exist for implementing each $e^{-iH_j t/2r}$ arising in a second-quantized formulation using a sequence of elementary gates and single-qubit rotations [11]. These rotations dominate the cost of the quantum simulation.

Circuit synthesis and quantum error correction

While existing estimates of the costs of quantum simulation have taken these rotations to be intrinsic quantum operations, doing so belies the cost of performing the computation fault tolerantly. To achieve reliable results, fault tolerant implementation of the algorithm is crucial. Fault tolerance is achieved by encoding a single *logical qubit* in a number of *physical qubits* with a quantum error-correcting code, such as the surface code [48]. This redundancy protects the logical qubit against decoherence and other experimental imperfections.

Quantum error correction cannot directly protect any arbitrary quantum operation, such as arbitrary rotations. Fortunately, however, it can protect a discrete set of gates, from which any continuous quantum operation can be approximated to within arbitrarily small error [49].

Approximation takes two steps. First the exponentials in the Trotter formula are decomposed into single-qubit rotations and so-called Clifford gates. In the surface code, which we consider here, Clifford gates can be implemented fault tolerantly. The single-qubit rotations, however, require approximation by a discrete set of gates consisting of Clifford operations and at least one non-Clifford operation, usually taken to be the T gate, a rotation by $\pi/8$ about the z -axis. Each T gate requires a procedure called magic state distillation, which consumes a host of noisy quantum states to output a single accurate magic state. The magic state is then used to teleport a T gate into the computation [50]. The space and time overheads of state distillation render it by far the most costly aspect of quantum error correction, leading to a large multiplicative overhead for each single-qubit rotation. The number of T gates therefore typically dominates the cost when implementing a fault tolerant algorithm.

III. RESOURCE ESTIMATES

We now show that the resources required to elucidate reaction mechanisms on quantum computers not only scale polynomially, but also that the calculations are feasible on a relatively small quantum computer in reasonable time. These resource estimates are based on extrapolations from the costs found from simulations of smaller molecules (for loose upper bounds see Appendix C) and focus on two prototypical structures of FeMoco that are

Quantitatively accurate simulation (0.1 mHa)				
Struct. 1	T-Gates	Clifford Gates	Time	Log. Qubits
Serial	1.1×10^{15}	1.7×10^{15}	130 days	111
Nesting	3.5×10^{15}	5.7×10^{15}	15 days	135
PAR	3.1×10^{16}	3.1×10^{16}	110 hours	1982
Struct. 2	T-Gates	Clifford Gates	Time	Log. Qubits
Serial	2.0×10^{15}	3.1×10^{15}	240 days	117
Nesting	6.5×10^{15}	1.0×10^{16}	27 days	142
PAR	6.0×10^{16}	6.0×10^{16}	204 hours	2024
Qualitatively accurate simulation (1 mHa)				
Struct. 1	T-Gates	Clifford Gates	Time	Log. Qubits
Serial	1.0×10^{14}	1.6×10^{14}	12 days	111
Nesting	3.3×10^{14}	5.6×10^{14}	1.4 days	135
PAR	3.0×10^{15}	3.0×10^{15}	11 hours	1982
Struct. 2	T-Gates	Clifford Gates	Time	Log. Qubits
Serial	1.9×10^{14}	3.0×10^{14}	22 days	117
Nesting	6.0×10^{14}	9.9×10^{14}	2.5 days	142
PAR	5.5×10^{15}	5.5×10^{15}	20 hours	2024

TABLE I. Simulation time estimates. Listed are the number of logical qubits and gate operations and an estimate of the runtime required to obtain energies within 0.1 mHa or 1 mHa for two different structures of FeMoco on a quantum computer and the number of logical qubits needed for the simulations (width). Structure 1 is for spin state $S = 0$ and charge +3 elementary charges with 54 electrons in 54 spatial orbitals. Structure 2 is for spin state $S = 1/2$ and charge 0 with 65 electrons in 57 spatial orbitals (see Supporting Material for further details). These runtimes and gate counts are likely to exceed the actual requirements.

typical of the complexity of those that naturally would arise when probing the potential energy landscape of the complex to find the reaction rates. Our work differs from previous studies in that we not only look at an important target for simulation but we also examine it within a basis set that is reasonable match for the target accuracy required (see Appendix J).

We first estimate the runtime of the computation assuming a quantum computer can perform a *logical* T gate every 10 ns, Clifford circuits require negligible time and that a good trial state for the ground state is available. We then determine the cost of performing this simulation fault tolerantly using the surface code, such that each *physical* gate takes 10 ns.

For chemical significance, we aim to compute the energies with a total error of at most 0.1 mHartree, adding up all sources of systematic and statistical error (see Appendix E). These errors include the error in phase estimation, the error in the Trotter–Suzuki decomposition and errors from decomposing the Trotter–Suzuki approximation into H and T gates. We then optimize over these three costs to find the most efficient way to divide the error budget up over these three sources. We also consider a larger error of 1 mHartree that would still put

our approach into the accuracy range of standard state-of-the-art quantum chemical methods for simple (mononuclear) transition metal complexes.

Three concrete implementations of the quantum algorithm are considered here. In the first approach, called Serial, the single-qubit rotations are constrained to occur serially. In a second approach, called Nesting, Hamiltonian terms that affect disjoint sets of spin orbitals are executed in parallel, which substantially reduces the runtime of the calculations. In the third approach, programmable ancilla rotations (PAR), rotations are cached in qubits, which are then teleported into the circuit as needed [12]. For each approach, the overall cost is found by decomposing the rotations into Clifford and T gates using [51] for the serial case and [52] in the other two cases.

The costs at a logical level, shown in Table I for simulation of two different structures of FeMoco in different charge and spin states (see Appendix J for details) are within reach of a small scale quantum computer. If all gates are executed in series then we estimate that the simulation will complete in under a year and use a small number of logical qubits. PAR can reduce the time required to several days at the price of requiring nearly 20 times as many logical qubits. Nesting gives a reasonable tradeoff between the two and may be the preferred approach. Finally, if we focus on learning qualitatively accurate reaction rates then the costs drop by roughly a factor of 10 in all examples. Further details can be found in the appendix.

Resource requirements with quantum error correction

We next add the overheads required to perform the simulation fault tolerantly and summarize the costs incurred by performing the simulation of structure 1 using the surface code in Table VI. The underlying resource calculations follow [48]. These costs naturally depend on the physical error rates assumed in the hardware, and we consider three cases that correspond to (a) a near-term error rate of 10^{-3} , (b) an error rate of 10^{-6} that might one day be feasible in superconducting qubits, but is expected early on for topological quantum bits [53] and (c) an error rate of 10^{-9} that may be achievable for topological quantum computers.

Rather than focusing solely on the number of physical qubits we consider a fine-grained modular architecture for a quantum computer, sketched in Fig. 3, consisting of a classical supercomputer interfacing with a quantum computer that consists of a main quantum processor and dedicated separate T factories and rotation factories to produce T gates and to synthesize rotations, respectively. Only the main quantum processor is a general purpose device, while the other factories are special purpose hardware that is expected to be easier to construct.

We first observe in Table VI that the number of logical qubits in the *main quantum processor* is only of the order




FIG. 3. We show the architecture of a hybrid classical/quantum computer for quantum chemistry type calculations. The quantum computer acts as an accelerator to the classical supercomputer. It consists of a classical control frontend, a main quantum processor and a number of auxiliary processing units. The devices labeled QRot build single qubit rotations using $\pi/8$ rotations created in the T gate factories labeled TFac. Red arrows denote quantum communication and blue represent classical communication.

of a hundred, which translates into tens of thousands

IV. DISCUSSION

While at present a quantitative understanding of chemical processes involving complex open-shell species such as FeMoco in biological nitrogen fixation remains beyond the capability of classical-computer simulations, our work shows that quantum computers used as accelerators to classical computers could be used to elucidate this mechanism using a manageable amount of memory and time. In this context a quantum computer would be used to obtain, validate, or correct the energies of intermediates and transition states and thus give accurate activation energies for various transitions.

We note that the required quantum computing resources are comparable to that needed for Shor’s factoring algorithm for interesting 4096 bit numbers; both in terms of number of gates and also physical qubits [48]. The complexity of these simulations is thus typical of that required for other targets for quantum computing.

We also observe that the resources required depend

to millions of physical qubits – which is challenging but not out of reach for early devices. Most of the qubits are used in the T factories, which each need even fewer physical qubits than the main quantum processor. Even the number of T factories needed to perform the serial calculation is reasonable, with only 202 such factories required in this case for an error rate of 10^{-3} , and fewer for better qubits. After factoring in the number of qubits needed per T factory, we arrive at the conclusion that only 10^5 to 10^6 physical qubits are needed for the 10^{-6} and 10^{-9} error rates.

If we consider the parallelizing with the PAR approach then these costs become much more daunting. The number of T factories required increases by a factor of roughly 1000 although the number of qubits required in each factory remains comparable. This explosion is due to the fact that PAR perform about one thousand times as many rotations as the serial code. The costs required by the nesting approach, on the other hand, are much more modest and lead to only an order of magnitude increase in the number of qubits with a comparable reduction in runtimes.

At low error rates, the total number of qubits required in both the serial and nested cases are of the order of a million, which could even be envisioned as a single quantum processor. The case of 10^{-3} errors or the PAR approach require hundreds of millions to nearly a trillion physical qubits, which seems more daunting. However, considering a modular design where each factory requires about one million physical qubits this is nevertheless certainly within the realm of possibility for future generations of quantum devices and the potential economic benefits of solving chemical catalysis problems of industrial importance may offset the cost.

sensitively on gate error rates. Therefore robust qubits, such as topologically protected qubits [53], may be invaluable for using a quantum computer to elucidate reaction mechanisms for important chemical processes such as the mode of action of nitrogenase.

Chemical reactions that involve strongly correlated species that are hard to describe by traditional multi-configuration approaches are not just limited to nitrogen fixation: they are ubiquitous. They range from small molecule activating catalysts for C-H bond breaking, to those for hydrogen and oxygen production, carbon dioxide fixation and transformation to industrially useful compounds, and to photochemical processes. Given the economic and societal impact of chemical processes ranging from fertilizer production to polymerization catalysis and clean energy processes, the importance of a versatile, reliable, and fast quantum chemical approach powered by quantum computing can hardly be overemphasized. Parallelising the quantum computation of the energy landscape will be crucial to providing answers within a time-

	Serial rotations			PAR rotations			Nested rotations		
Error Rate	10^{-3}	10^{-6}	10^{-9}	10^{-3}	10^{-6}	10^{-9}	10^{-3}	10^{-6}	10^{-9}
Required code distance	35,17	9	5	37,19	9,5	5	37,17	9	5
	Quantum processor								
Logical qubits	111			110			109		
Physical qubits per logical qubit	15313	1013	313	17113	1013	313	17113	1013	313
Total physical qubits for processor	1.7×10^6	1.1×10^5	3.5×10^4	1.9×10^6	1.1×10^5	3.4×10^4	1.9×10^6	1.1×10^5	3.4×10^4
	Discrete Rotation factories								
Number	0			1872			26		
Physical qubits per factory	–	–	–	17113	1013	313	17113	1013	313
Total physical qubits for rotations	–	–	–	3.2×10^7	1.9×10^6	5.9×10^5	4.5×10^5	2.6×10^4	8.1×10^3
	T factories								
Number	202	68	38	166462	41110	29659	5845	1842	1029
Physical qubits per factory	8.7×10^5	1.7×10^4	5.0×10^3	1.1×10^6	7.5×10^4	5.0×10^3	8.7×10^5	1.7×10^4	5.0×10^3
Total physical qubits for T factories	1.8×10^8	1.1×10^6	1.9×10^5	1.8×10^{11}	3.1×10^9	1.5×10^8	5.1×10^9	3.0×10^7	5.2×10^6
Total physical qubits	1.8×10^8	1.2×10^6	2.3×10^5	1.8×10^{11}	3.1×10^9	1.5×10^8	5.1×10^9	3.0×10^7	5.2×10^6

TABLE II. This table gives the resource requirements including error correction for simulations of nitrogenase’s FeMoco in structure 1 within the times quoted in Table I using physical gates operating at 100 MHz and taking the error target to be 0.1 mHa. While the total number of qubits seems large, the quantum computer is composed of many small and weakly coupled devices that can be mass produced.

frame of several days instead of several years. Quantum computers therefore must be designed such that the components can be mass produced and clusters of quantum computers can be built in order to scan over the many structures that need to be examined to identify and estimate all important reaction rates accurately [28].

V. ACKNOWLEDGEMENTS

We acknowledge support by the Swiss National Science Foundation, the Swiss National Competence Center in Research NCCR QSIT, and hospitality of the Aspen Center for Physics, supported by NSF grant # PHY-1066293.

Appendix A: Introduction to Quantum Computing

Here we provide a brief review of quantum computing for those that are not experts in quantum computing. The aim of this review is to introduce the basic concepts and notations that are used in quantum computing as well as introduce quantum algorithms for simulating chemistry and also quantum error correction. For a more detailed exposition see [49]. These concepts will be used in section B 1 where we discuss quantum processes such as phase estimation.

1. Qubits and Quantum Gates

In quantum computation, quantum information is stored in a quantum bit, or *qubit*. Whereas a classical bit has a state value $s \in \{0, 1\}$, a qubit state $|\psi\rangle$ is a *linear superposition* of states:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}, \quad (\text{A1})$$

where the $\{0, 1\}$ basis state vectors are represented in Dirac notation (called *ket* vectors) as $|0\rangle = \begin{bmatrix} 1 & 0 \end{bmatrix}^T$, and $|1\rangle = \begin{bmatrix} 0 & 1 \end{bmatrix}^T$, respectively. The *amplitudes* α and β are complex numbers that satisfy the normalization condition: $|\alpha|^2 + |\beta|^2 = 1$. Upon *measurement* of the quantum state $|\psi\rangle$, either state $|0\rangle$ or $|1\rangle$ is observed with probability $|\alpha|^2$ or $|\beta|^2$, respectively.

Dirac notation is used largely because it contains an implicit tensor product structure that makes expressing qubit states much easier. For example, that the four-qubit state $|0000\rangle$ is equivalent to writing the tensor product of the four states: $|0\rangle \otimes |0\rangle \otimes |0\rangle \otimes |0\rangle = |0\rangle^{\otimes 4} = [1\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0]^T$.

A general n -qubit quantum state lives in a 2^n -dimensional Hilbert space and is represented by a $2^n \times 1$ -dimensional state vector whose entries represent the amplitudes of the basis states. A superposition over 2^n states is given by:

$$|\psi\rangle = \sum_{i=0}^{2^n-1} \alpha_i |i\rangle, \text{ such that } \sum_i |\alpha_i|^2 = 1, \quad (\text{A2})$$

where α_i are complex amplitudes and i is the binary representation of integer i . The ability to represent a superposition over exponentially many states with only a linear number of qubits is one of the essential ingredients of a quantum algorithm — an innate massive parallelism.

In a quantum computation, unitary transformations are used to transform quantum states into other quantum states. In particular, the quantum state $|\psi_1\rangle$ of the system at time t_1 is related to the quantum state $|\psi_2\rangle$ at time t_2 by a unitary $|\psi_2\rangle = U|\psi_1\rangle$.

In general we cannot expect that a quantum computer can implement every unitary transformation on n qubits exactly because there are an infinite number of such transformations. Instead, gate model quantum computers use a discrete set of quantum gates (which can be represented as $2^n \times 2^n$ unitary matrices) to approximate these continuous transformations. Any such set of gates is known as universal if any unitary transformation can be expressed, within arbitrarily small error, as a sequence of such gates.

The gate set that we consider in this work includes the following gates:

- the **X** gate which is the classical NOT gate that maps $|0\rangle \rightarrow |1\rangle$ and $|1\rangle \rightarrow |0\rangle$.
- the *Hadamard* gate **H** maps $|0\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|1\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$.
- the **Z** gate maps $|1\rangle \rightarrow -|1\rangle$ and **T** is the fourth-root of **Z**, and can be used interconvert **Z** and **X** gates via $\text{HZH} = \text{X}$.
- the **Y** gate maps $|1\rangle \rightarrow -i|0\rangle$ and $|0\rangle \rightarrow i|1\rangle$.
- the identity gate is represented by $\mathbb{1}$.
- the two-qubit *controlled-NOT* gate, **CNOT**, maps $|x, y\rangle \rightarrow |x, x \oplus y\rangle$. The corresponding unitary matrices are:

$$\text{H} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \text{X} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \text{Y} = \begin{bmatrix} 0 & i \\ -i & 0 \end{bmatrix}, \text{Z} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

$$\text{T} = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}, \mathbb{1} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

Measurement is an exception to this rule; it collapses the quantum state to the observed value, thereby erasing any information about the amplitudes α and β . This means that extracting information from a quantum system irreversibly damages a state. Consequently, although the exponential parallelism that quantum computers naturally possess is mitigated by the fact that most quantum computations will have to be repeated many times to extract the necessary information about the output of the quantum algorithm.

2. Quantum Circuit Synthesis

In order to understand how the cost estimates of our algorithms are found it is important to understand how arbitrary unitaries can be converted into discrete gate sequences. For our purposes, we consider the Clifford + T gate library discussed above, but compilation into other gate sets is also possible [54, 55]. The simplest case to consider is that of single qubit unitary synthesis where $U \in \mathbb{C}^{2 \times 2}$. In such cases, an Euler angle decomposition exists such that, up to an irrelevant overall phase,

$$U = e^{i\alpha Z} \text{H} e^{i\beta Z} \text{H} e^{i\gamma Z}. \quad (\text{A3})$$

Thus the problem of implementing a single qubit transformation reduces to the problem of performing the rotation $e^{i\theta Z} = \cos(\theta)\mathbb{1} + i\sin(\theta)Z$. The general case of $U \in \mathbb{C}^{2^n \times 2^n}$ similarly reduces to instances of single qubit synthesis interspersed with entangling gates such as **CNOT**. Thus implementing single qubit rotations can be seen as an atomic operation on which compilation process for U is based.

There are a host of methods known for decomposing rotations into Clifford and T gates. Perhaps the earliest such algorithm is the Solovay–Kitaev algorithm [56], which provided an efficient method for performing this decomposition based on Lie–algebraic techniques. In recent years, much more efficient methods for decomposition have been discovered that are based on number theory. These methods require near–quartically fewer operations than the Solovay–Kitaev algorithm would need [56] and are useful, if not necessary, for reducing the gate counts of the simulation algorithm to a palatable level.

In the absence of ancillae, the cost required to approximate an axial rotation within error ϵ for the worst possible input angle, $C(\epsilon)$, is known to lie within the interval [52]

$$4 \log_2(1/\epsilon) - 9 \leq C(\epsilon) \leq 4 \log_2(1/\epsilon) + 11. \quad (\text{A4})$$

Subsequent work [57] has provided a method that has an average case complexity roughly 3/4 of this worst case bound and also has made the compilation process efficient [58].

More recent methods have introduced the use of ancillae to reduce the costs of synthesis [51]. These methods can substantially reduce the number of T gates required to perform the synthesis. This approach can reduce the scaling of the average number of T gates needed to implement an axial rotation

$$C(\epsilon) \approx 1.15 \log_2(1/\epsilon) + 9.2. \quad (\text{A5})$$

This approach requires one additional ancilla qubit.

Asymptotically better methods exist, such as repeat-until-success (RUS) synthesis with fallback [51], but these methods do not outperform this method given the range of ϵ that we require for the simulation to reach chemical accuracy. Further methods such as PAR rotations [12] or

gearbox circuits [59] can be used to substantially reduce the T-depth of the simulation circuits at the price of requiring greater parallel width. While we do not consider the impact of gearbox synthesis here, we will examine the costs and benefits of PAR rotations.

An important issue arises when using such methods for parallel execution. If several rotations are implemented simultaneously on different qubits and RUS synthesis is used then the dominant contribution to the cost is given by the longest sequence of gates used in each block of parallelized rotations. In particular, while RUS synthesis reduces the expectation value by a factor of roughly 4 from the worst case bounds, it is clear that when performing many of them in parallel it is very likely that at least one of them will saturate the bound. For this reason we use upper bounds on deterministic synthesis given by (A4), which have a better worst case scaling.

3. Simulating Quantum Chemistry

Our approach to quantum chemistry simulations closely follows the strategy proposed in Ref. [9]. We begin by representing the quantum chemistry Hamiltonian in a second quantized form, keeping track of the locations of the electrons by using the occupations of a discrete basis of spin-orbitals. Since a spin orbital can only contain one electron, such states are very natural to express in a quantum computer. Each spin orbital is assigned a qubit where the state $|1\rangle$ corresponds to an occupied orbital and $|0\rangle$ an unoccupied orbital.

These states are often described using creation and annihilation operators which obey $a^\dagger|0\rangle = |1\rangle$, $a^\dagger|1\rangle = 0$, $a|1\rangle = |0\rangle$ and $a|0\rangle = 0$. Since these operators correspond to creating electrons, they must respect the symmetries appropriate for Fermions. The most important property is the anti-commutation property $\{a_i^\dagger, a_j\} = \delta_{i,j}$. This means that while it may be tempting to identify $a^\dagger = (\mathbf{X} - i\mathbf{Y})/2$, such a representation does not satisfy the anti-commutation relation for Fermionic operators. Instead, these creation and annihilation operators can be converted into Pauli operators using the Jordan–Wigner or Bravyi–Kitaev [60] transformations.

In order to simulate the dynamics of the system we need to emulate the unitary dynamics that the system undergoes using a sequence of gate operations. This dynamics is of the form e^{-iHt} where

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s, \quad (\text{A6})$$

Once these fermionic creation and annihilation operators are translated into Pauli operators using, for example, the Jordan–Wigner transformation then the Hamiltonian can be translated into a sequence of operations that individually could be performed on a quantum computer. Well known circuits exist for performing the exponentials yielded by the Trotter–Suzuki formulas [11, 14, 49].

The most direct method for estimating the ground-state energy is phase estimation (see Section B 1 for more detail). Phase estimation uses quantum interference between controlled executions of $\mathbb{1}, e^{-iHt_1}, e^{-iHt_2}, \dots$, to infer the eigenvalues of e^{-iHt} . If t is smaller than $\pi/\|H\|$ this also directly yields the eigenvalues of H . Apart from the ability to directly sample from the eigenvalues, a further advantage to phase estimation is that it requires $O(1/\epsilon)$ applications of e^{-iHt} to learn its eigenphase within error ϵ with high probability. This is quadratically better than bounds on the variance that would be seen from optimal classical sampling methods using an unbiased estimator.

It is necessary to apply phase estimation on an initial state that has large overlap with the ground state in order to find the ground-state energy with high probability. Specifically, if applied to an initial state $|\psi\rangle = a_g|E_G\rangle + \sqrt{1 - |a_g|^2}|E_G^\perp\rangle$ then PE returns an estimate of ground-state energy E_G with probability $|a_g|^2$. The simplest state to use is the Hartree–Fock state, which only requires applying a sequence of NOT gates to the state $|0\rangle^n$, however configuration interaction with sufficiently high excitations may be required to achieve high overlap for systems that have strong correlations in their ground states. We do not consider the cost of preparing such a state here, since such a cost of preparing a sufficiently accurate approximation to the ground state of FeMoco is difficult to determine in absentia of large scale quantum computers. We discuss this issue in more detail in Section F.

4. Quantum error correction

Quantum hardware is far less robust to errors as classical hardware. Quantum error correction provides a way to reduce the errors in the device without sacrificing the quantum nature of the system. Quantum error correcting codes require that the physical error rates, say of qubits, quantum gates, and measurements, are less than a given threshold value. This threshold value depends on the error correcting code being used and the type of noise of the system. If error rates are below the threshold, then errors in the computation, referred to as the logical gates and logical qubits, can be made arbitrarily small at only a polylogarithmic overhead.

The surface code is currently the most popular code due in part to its relatively high threshold of 1% [61]. Much of this enthusiasm has arisen because of evidence that existing superconducting quantum computers may already have error rates near this threshold [62]. This raises the hope that a fault-tolerant, scalable quantum computer may be just over the horizon.

While quantum error correction promises the ability to perform arbitrarily long quantum computations on a noisy device, the resources required to execute such a fault-tolerant computation can be large if the system operates close to the threshold. The majority of the cost

of quantum error correction arises from the need to perform a universal set of quantum gates. While protecting so-called Clifford gates requires very little overhead in the surface code, in comparison protecting a non-Clifford gate, such as a T or Toffoli gate, requires substantial resources. While several techniques for producing fault-tolerant non-Clifford gates exist [63–66], here we focus on the use of magic state distillation in conjunction with the surface code [61]. Magic state distillation [50] takes as input a set of noisy resource states and outputs a cleaner resource state. For the surface code, we must distill the T state as it cannot be implemented directly in the code. Here we consider the 15 – 1 distillation scheme of Bravyi and Kitaev [50], where 15 noisy input states with error rate p produce a single magic resource state with error rate roughly $35p^3$.

Magic state distillation consists only of Clifford operations, which can be implemented easily within the surface code.

Appendix B: Implementing Phase Estimation

In this section, we discuss how to implement the phase estimation protocol in quantum computers. This is important to our subsequent estimates of the complexity of simulating nitrogenase’s FeMoco because phase estimation constitutes the outer most loop of the quantum simulation and hence is a major driver of the cost of the simulation.

1. Phase Estimation

Phase estimation is one of the most critical components of the quantum simulation algorithm. Without phase estimation, the amount of simulation time needed to estimate the ground-state energy would grow quadratically as ϵ^{-2} with the precision ϵ required. Since ϵ is on the order of 0.1mHartree, the phase must be estimated for one time step of the evolution operator within an error 0.1/ r mHartree where r is the number of Trotter steps required. If statistical sampling, rather than phase estimation, were used to estimate the phase then on the order of 10^{12} experiments would be needed to make the variance in the estimate sufficiently small. This would be impractical and so phase estimation is crucial for most large scale applications in quantum chemistry.

The standard quantum phase estimation algorithm [49] allows eigenvalues to be learned within error ϵ with probability at least 1/2 uses a number of applications of the unitary circuit, $M(\epsilon)$, that is bounded above by

$$M(\epsilon) \leq \frac{16\pi}{\epsilon}. \quad (\text{B1})$$

This value is far from optimal time scaling of π/ϵ , but has the advantage of requiring neither measurements of the quantum system nor a classical computer to infer the




FIG. 4. This circuit is used in iterative phase estimation algorithms wherein the eigenvalues are inferred from measurement statistics.

most likely eigenvalue. We will also use this result below because it provides an upper bound on the cost of the optimal phase estimation algorithm.

An alternative approach is to use iterative phase estimation. Iterative phase estimation forgoes storing the phase in a quantum register and instead uses a classical inference algorithm to learn the eigenphase from measurements of an auxiliary probe qubit that is iteratively measured and re-entangled with the system. Kitaev proposed the first variant of iterative phase estimation. The circuit used in this process is given in Fig 4.

The ultimate limit that can be achieved, in terms of the number of times the unitary is applied as a function of desired error tolerance, is given by [67–69]

$$M(\epsilon) \approx \frac{\pi}{\epsilon}. \quad (\text{B2})$$

Furthermore, Gaussian strategies such as RFPE [70] yield Bayesian Cramer–Rao bounds that come close to saturating this: $3.3/\epsilon$. As these bounds are often saturated for likelihood functions of this form [71] and because the Gaussian assumption causes the user to throw away substantial information from higher moments in the posterior distribution, we expect that (B2) represents the ultimate limit for phase estimation and that based on previous studies the use of optimized policies [72–74] will enable performance that comes close to saturating this limit. As a result we use $M(\epsilon) \approx \pi/(2\epsilon)$ as a surrogate for the expected performance of these optimized approaches, where the factor of 2 difference follows from an important optimization that holds for the case of quantum simulation algorithms that we discuss in detail in the following section.

2. Controlled rotations

While the previous discussion only showed how to perform a Z-rotation, we need to perform controlled Z-rotations to perform the phase estimation algorithm. Fortunately, there are well known methods that can be used to perform such controlled rotations using a pair of rotations and two CNOT gates. A better approach, previously shown in unpublished work by Tsuyoshi Ito in 2012, is given in Figure 5 and the validity of the circuit is proven in the following lemma.




FIG. 5. Low depth circuit for controlled Z-rotations where the θ gate represents $e^{i\theta Z}$.

Lemma 1. *The circuit of Fig. 5 implements the controlled operation $\Lambda(R_z(2\theta))$ where the top most qubit is the control.*

Proof. Assume that $|c\rangle = |0\rangle$ or $|1\rangle$ then the CNOT gate performs

$$|c\rangle|\psi\rangle \mapsto (\mathbf{X}^c \otimes \mathbf{1})(a|00\rangle + b|11\rangle). \quad (\text{B3})$$

The rotation gates then prepare the state

$$(ae^{i(1-(-1)^c)\theta/2}|c\rangle|0\rangle + be^{-i(1-(-1)^c)\theta/2}|c \oplus 1\rangle|1\rangle). \quad (\text{B4})$$

Finally the CNOT gate yields the state

$$\begin{aligned} |c\rangle(ae^{i(1-(-1)^c)\theta/2}|0\rangle + be^{-i(1-(-1)^c)\theta/2}|1\rangle) \\ = (\Lambda(R_z(2\theta))|c\rangle|\psi\rangle). \end{aligned} \quad (\text{B5})$$

Therefore the circuit functions properly for either $|c\rangle = |0\rangle$ or $|c\rangle = |1\rangle$. By linearity it is also valid for any input. \square

This form of a controlled rotation is well suited for many quantum simulation applications because it allows the controlled evolutions to be replaced with evolutions that control only on the single qubit rotations in the circuits for implementing the individual terms in the Hamiltonian. However, it is not optimal here because we can implement a variant of a controlled rotation gate to in effect double the impact that the eigenphase has on the measurement probabilities.

Lemma 2. *Let $U = \prod_{j=1}^{N_{\text{exp}}} B_j(\mathbf{1} \otimes R_z(\phi_j))B_j^\dagger$ where B_j are unitary transformations and assume that it is chosen such that $U^\dagger = \prod_{j=1}^{N_{\text{exp}}} B_j(\mathbf{1} \otimes R_z(-\phi_j))B_j^\dagger$ then if $U|\phi\rangle = e^{i\phi}$ then there exists a quantum protocol that uses an M rotations and samples from a Bernoulli distribution such that the probability of the protocol outputting 0 is*

$$P(0|\phi; \theta, M) = \frac{1 + \cos(2M[\phi + \theta])}{2}.$$

Proof. Our proof is constructive and follows directly from the arguments presented informally in [75]. The idea behind the protocol is to replace every controlled rotation used in the circuit in Fig. 4 with the circuit in Fig. 6. Formally we denote this by replacing the controlled operation $\Lambda(U^M)$ with W^M , which is defined such that $W|0\rangle|\psi\rangle = |0\rangle U|\psi\rangle$ and $W|1\rangle|\psi\rangle = |1\rangle U^\dagger|\psi\rangle$. Since

$X R_z(\theta_j) X = R_z(-\theta_j)$, the operation W can be written as

$$W = \prod_{j=1}^{N_{\text{exp}}} B_j \Lambda(X) (\mathbf{1} \otimes R_z(-\phi_j)) \Lambda(X) B_j^\dagger, \quad (\text{B6})$$

where the controlled-not operation $\Lambda(X)$ uses the control qubit as its control and the qubit that R_z acts on as its target. This circuit is shown in Fig. 6.

Applying $(\mathbf{H} \otimes \mathbf{1}) W (\mathbf{H} \otimes \mathbf{1})$ to the eigenstate $|0\rangle|\phi\rangle$, such that $U|\phi\rangle = e^{i\phi}|\phi\rangle$, yields

$$\frac{1}{2} \left(|0\rangle|\phi\rangle (1 + e^{2i(\theta-\phi)M}) + |1\rangle|\phi\rangle (1 - e^{2i(\theta-\phi)M}) \right), \quad (\text{B7})$$

up to a global phase. The probability of measuring the ancilla qubit to be zero is $(1 + \cos(2M(\phi - \theta)))/2$ as claimed. \square

Lemma 2 shows that the number of rotations needed to perform phase estimation is 1/4 the value that would be expected if circuits such as that of Fig. 5 were used (or 1/2 the cost in parallel settings). As an example, for the case of RFPE numerical experiments show that we can learn the eigenphase within error ϵ with probability 1/2 using

$$M(\epsilon) \approx \frac{2.3}{\epsilon}. \quad (\text{B8})$$

In comparison, the ultimate lower limit given by previous studies becomes $\pi/(2\epsilon)$ and the Bayesian Cramer-Rao bound gives a lower limit of approximately $1.6/\epsilon$ under the assumptions of Gaussian priors made in RFPE [70]. Similarly the upper bounds on the cost of traditional QPE in [49] become $8\pi/\epsilon$ rather than $16\pi/\epsilon$ if these circuits are used.

While the assumption that the underlying Trotter-Suzuki formula can be inverted by simply inverting the sign of the evolution time. Specifically, this assumption applies for the $(2k+1)$ th order Trotter-Suzuki formulas for all $k \geq 1$ but it does not apply to the second order formula unless further assumptions are made. We can see this from the fact that

$$\left(\prod_{j=1}^N e^{-iH_j t} \right) \left(\prod_{j=1}^N e^{iH_j t} \right) = 1 + O(t^2).$$

If we assume that we are interested only in the ground-state energy and the Hamiltonian is real valued (like in the quantum chemistry applications that we consider) then the error in assuming that U^\dagger can be formed by simply flipping the signs is $O(t^3)$ [16], which is by no means fatal but it could potentially contribute to the error in Trotter-Suzuki decompositions. Whereas if the third (or higher) order Trotter-Suzuki formula is used then no such danger exists. This, along with the superior bounds proven for the error in the third order formula, provide the justification for using this formula in preference to the asymptotically equivalent second-order formula.




FIG. 6. Circuit used to implement analogue of controlled Z-rotations used in Lemma 2, which is a rotation whose sign is conditioned on the control qubit.

Appendix C: Trotter errors

Here we discuss the issue of how errors from the use of Trotter Suzuki formulas lead to systematic errors in the ground-state energy estimates output by phase estimation. We further discuss the methodology that we use to upper bound these errors and also give empirical estimates of the scaling of such errors.

1. Rigorous bounds

A major source of error in most quantum simulation algorithms arises from the use of Trotter formula expansions. Such errors can be made arbitrarily small but the need to make such errors smaller than chemical accuracy mean that the cost of doing so can substantially impact the time required to perform the simulation. In this appendix we will provide a detailed discussion about how to bound the error in low-order Trotter formulas.

Trotter errors arise from the fact that the terms in the Hamiltonian used in the expansion do not commute. In principle, the Zassenhaus formula provides everything that is needed in order to understand the scaling of these errors:

$$e^{(A+B)t} = e^{At}e^{Bt}e^{\frac{1}{2}[A,B]t^2} + O(t^3), \quad (\text{C1})$$

and thus

$$\|e^{(A+B)t} - e^{At}e^{Bt}\| \in O(\|[A, B]\|t^2). \quad (\text{C2})$$

Although this expression is useful in estimating the scaling of simulation errors, the question that we're interested in is somewhat orthogonal to this. Instead, we are interested in the errors in estimated eigenvalues. The Baker–Campbell–Hausdorff formula, which is the dual to the Zassenhaus formula, can be used to estimate these errors. If $H = \sum_{\alpha=1}^L H_{\alpha}$ where the H_{α} correspond to terms in (A6) then the second order Trotter–Suzuki formula (also known as the Strang splitting) gives

$$\prod_{\alpha=1}^L e^{-iH_{\alpha}t/2} \prod_{\alpha=L}^1 e^{-iH_{\alpha}t/2} = e^{-iH_{\text{eff}}t}, \quad (\text{C3})$$

where $H - H_{\text{eff}}$ is

$$-\frac{1}{12} \sum_{\alpha \leq \beta} \sum_{\beta} \sum_{\alpha' < \beta} [H_{\alpha}(1 - \frac{\delta_{\alpha,\beta}}{2}), [H_{\beta}, H_{\alpha'}]]t^2 + O(t^3). \quad (\text{C4})$$

This shows that, to leading order, the error in the Strang splitting can be estimated by the ground-state expectation value of a double commutator sum. Furthermore, since $L \in O(N^4)$ it is clear from this form that the error in the Trotter formula must scale at most as $O(N^{10}t^2)$. It is not $O(N^{12}t^2)$ because the commutator structure restricts two of the orbitals H_{α} and H_{β} act on. Although this estimate can be computed in polynomial time, there are too many terms for this to be reliably estimated for molecules on the scale of nitrogenase even using Monte Carlo methods [19].

An upper bound on the asymptotic scaling can be trivially found by applying the triangle inequality to (C4). This approach is not appropriate for our purposes because we do not know a priori how large t must be for the leading order term to be dominant. As a result, we use the following result, which is provably an upper bound on the error in the energy of the evolution:

$$\begin{aligned} \Delta E_{\text{TS}}/t^2 \leq & 4 \sum_{\alpha, \beta, \alpha'} \|H_{\alpha}\| \|H_{\beta}\| \|H_{\alpha'}\| \\ & \times (\delta_{\alpha > \beta} \delta_{\alpha' > \beta} + \delta_{\beta > \alpha} \delta_{\alpha', \alpha}) W(\alpha, \beta, \alpha'), \end{aligned} \quad (\text{C5})$$

where $W(\alpha, \beta, \alpha')$ is an indicator function that takes the value 1 if and only if the corresponding double commutator is non-zero. Note that this is slightly tighter than the bound used in (16) of [17].

There are several criteria that we know a priori lead to a double commutator vanishing:

1. $H_{\alpha'}$ and H_{β} act on disjoint sets of qubits.
2. H_{α} acts on a disjoint set of qubits from the set of qubits that $H_{\alpha'}$ and H_{β} act on.
3. $H_{\alpha'}$ and H_{β} correspond to PP or PQQP terms.
4. $H_{\alpha'}$ and H_{β} correspond to PR and PQQR terms with the same P and R.
5. Only one of $[H_{\alpha}, [H_{\beta}, H_{\alpha'}]]$, $[H_{\beta}, [H_{\alpha'}, H_{\alpha}]]$ and $[H_{\alpha'}, [H_{\alpha}, H_{\beta}]]$ is non-zero according to the prior rules (Jacobi identity).

There are other symmetries to the terms that can be used to argue that even more terms are necessarily zero. Since we do not consider these properties, we overcount the contribution of any such terms and so our estimate remains an upper bound.

This bound is much more computable than the original expression, but is computationally challenging to compute exactly owing to the $O(N^{10})$ terms in the double commutator sum. The Cauchy–Schwarz inequality can be used to convert this expression into one that can be

computed in $O(N^4)$ operations, but doing so can over-represent the influence of large terms in the Hamiltonian [17].

2. Empirical estimates of Trotter Error

Given that computation of the matrix elements of the error operator in (C4) is computationally challenging, we rely on Monte Carlo sampling to estimate the upper bound in (C5). Monte Carlo sampling is much more effective here because the use of the triangle inequality removes the alternating sign that we see when summing the original series. We achieve this by drawing M samples uniformly at random for $\{\alpha_j : j = 1 \dots M\}$, $\{\beta_j : j = 1 \dots M\}$ and $\{\beta'_j : j = 1 \dots M\}$. We then reject the sample if $(\delta_{\alpha>\beta}\delta_{\alpha'>\beta} + \delta_{\beta>\alpha}\delta_{\alpha',\alpha})W(\alpha, \beta, \alpha') = 0$, and otherwise compute the product of the products of the norms of the three corresponding Hamiltonian terms. If there are L terms in the Hamiltonian and define

$$\Gamma(\alpha, \beta, \alpha') := 4\|H_\alpha\|\|H_\beta\|\|H_{\alpha'}\| \times (\delta_{\alpha>\beta}\delta_{\alpha'>\beta} + \delta_{\beta>\alpha}\delta_{\alpha',\alpha})W(\alpha, \beta, \alpha') \quad (\text{C6})$$

then

$$\Delta E_{\text{TS}} \lesssim \frac{L^3}{M} \sum_{j=1}^M \Gamma(\alpha_j, \beta_j, \alpha'_j) t^2 := h t^2. \quad (\text{C7})$$

This implies that, for a fixed h , if we wish to achieve error ϵ in the eigenvalues of the Trotter–Suzuki expansion then it suffices to pick

$$t = \sqrt{\frac{\epsilon}{h}}. \quad (\text{C8})$$

The variance in this estimator is

$$\frac{L^6 \mathbb{V}_{\alpha, \beta, \alpha'}(\Gamma(\alpha, \beta, \alpha')) t^4}{M}, \quad (\text{C9})$$

which implies using Chebyshev’s inequality that with probability greater than 75% the sample error is less than 2ϵ if

$$M \geq \frac{L^6 \mathbb{V}_{\alpha, \beta, \alpha'}(\Gamma(\alpha, \beta, \alpha')) t^4}{\epsilon^2} = \frac{L^6 \mathbb{V}_{\alpha, \beta, \alpha'}(\Gamma(\alpha, \beta, \alpha'))}{h^2}. \quad (\text{C10})$$

In practice, uniform sampling is not necessarily the best option because the importance of the different terms can vary wildly within a class. For example, the one-body terms tend to be much larger than the two-body terms but the two-body terms are far more numerous. This means that uniform sampling can underestimate the contributions of such terms because of their relative scarcity in the sample space.

We combat this by sampling from each *type* of double commutator. In particular, we sample over a class

of double commutators such as [PQ, [PP, PQ]] and uniformly draw PQ, PP and PQ terms to estimate those terms contribution to the overall error. The total error is then the sum of the estimates over all such classes. We use a minimum of 10^8 samples per class, which renders the sample standard deviation in our estimates of the error less than 1%.

The resultant bounds can be seen in Fig. 7 wherein we examine the predicted Trotter numbers and the empirically observed Trotter numbers for small molecules. The rough scaling of the upper bound on the Trotter number that we see corresponds to $N^{2.5}$ which was noted in previous numerical studies [17], but owing to the scatter of the data due to the widely varying chemical properties of the molecules, this scaling should not be seen as definitive. We observe that the upper bounds seem to be roughly a factor of 10 000 times too loose for small molecules. For this reason we plot three reasonable extrapolations of the scaling based on the numerical results for small molecules. The most pessimistic bound rescales the scaling extracted from the upper bound such that all the data remains beneath the curve. The middle one rescales the upper bound data by the average discrepancy between the upper bound and the numerically computed examples. The most optimistic curve is simply a polynomial fit to the numerical data that ignores the upper bound. We expect the middle curve to be the most realistic estimates, but provide resource estimates for these three cases below as well as results that follow from using the upper bound.

Appendix D: Error propagation

Here we provide proofs of some basic results that we will use to propagate these errors through the quantum simulation. These results are crucial for the cost estimates in the subsequent section because they show how large the worst case errors can be in the eigenvalue estimation given errors of these magnitudes. It is worth noting that we expect these results to yield substantial overestimates of the error because they do not consider the natural cancellations that are likely to occur in practical eigenvalue estimation problems.

Lemma 3. *Let A and B be Hermitian operators acting on finite dimensional Hilbert spaces such that $\|A - B\|_2 \leq \epsilon$ and $A|\psi_A\rangle = E_A|\psi_A\rangle$ and $B|\psi_B\rangle = E_B|\psi_B\rangle$ where E_A and E_B are the smallest eigenvalues of either operator then $|E_A - E_B| \leq \epsilon$.*

Proof. Because A and B are Hermitian they satisfy the variational property meaning that

$$E_B \leq \langle \psi_A | B | \psi_A \rangle. \quad (\text{D1})$$

Since $\|A - B\|_2 \leq \epsilon$ it follows that there exists C such that $\|C\|_2 \leq 1$ and $A = B + \epsilon C$. This implies that

$$E_B \leq \langle \psi_A | A + \epsilon C | \psi_A \rangle = E_A + \epsilon \langle \psi_A | C | \psi_A \rangle. \quad (\text{D2})$$




FIG. 7. Trotter number ($1/dt$) needed to reach 0.1 mHartree of accuracy assuming no errors from synthesis or phase estimation. Lines represent projected upper bounds based on current data.

Thus

$$E_B - E_A \leq \epsilon \langle \psi_A | C | \psi_A \rangle. \quad (\text{D3})$$

If $E_A \leq E_B$ then we have from the definition of $\|\cdot\|_2$ that

$$|E_B - E_A| \leq \epsilon |\langle \psi_A | C | \psi_A \rangle| \leq \epsilon. \quad (\text{D4})$$

Now assume that $E_B > E_A$. We then have

$$E_A \leq \langle \psi_B | A | \psi_B \rangle. \quad (\text{D5})$$

Then by repeating the same argument we conclude that $|E_B - E_A| \leq \epsilon$ regardless of the sign of $E_A - E_B$. \square

Now we will go beyond this bound to show that the error scaling in the eigenvalues of the unitary evolutions generated by two similar Hamiltonians is no more pathological than the scaling of errors in the ground-state energies.

Lemma 4. *Assume that for Hermitian bounded operators A and B acting on a finite dimensional Hilbert space $\|e^{-iAt} - e^{-iBt}\|_2 \leq t\gamma(t)$ for $\gamma(t)$ a non-decreasing continuous function of t on $[0, \infty)$ then $\|A - B\|_2 \leq \gamma(t)$.*

Proof. Using standard bounds [49], we have that $\|e^{-iAt} - e^{-iBt}\|_2 \leq \|A - B\|_2 t$ and furthermore from Taylor's theorem $\|e^{-iAt} - e^{-iBt}\|_2 = \|A - B\|_2 t + O(\|A - B\|_2^2 t^2)$.

Therefore the former upper bound is tight in the limit as $t \rightarrow 0$. By assumption $\|e^{-iAt} - e^{-iBt}\|_2 \leq t\gamma(t)$ for all t in a compact subinterval containing 0. Assume that $\lim_{t \rightarrow 0} \gamma(t)/\|A - B\|_2 < 1$. This implies that there exists a compact interval containing 0 such that for all t in this interval $\|e^{-iAt} - e^{-iBt}\|_2 > \|A - B\|_2 t$, which leads to a contradiction because we have already demonstrated that $\|A - B\|_2 t$ is a tight bound on the error in this limit. Therefore $\gamma(t) \geq \lim_{t \rightarrow 0} \gamma(t) \geq \|A - B\|_2$ under the assumptions of the lemma. \square

Theorem 1. *Let $H = \sum_{j=1}^M H_j$ where each H_j is a bounded Hermitian operator acting on a finite dimensional Hilbert space. Furthermore, let $\tilde{H} = \sum_{j=1}^M \tilde{H}_j$ be a similar sum of bounded operators such that $\|H_j - \tilde{H}_j\|_2 \leq \delta$. Finally, let $\|e^{-iHt} - \prod_{j=1}^M e^{-iH_j t/2} \prod_{j=M}^1 e^{-iH_j t/2}\|_2 \leq \Delta E_{TS}(t)t$. Then the difference in ground-state energies between $H(t)$ and $\tilde{H}(t) := i \log(\prod_{j=1}^M e^{-i\tilde{H}_j t/2} \prod_{j=M}^1 e^{-i\tilde{H}_j t/2})/t$ is at most $\Delta E_{TS}(t) + (2M - 1)\delta$.*

Proof. First, since \tilde{H}_j is Hermitian

$$\prod_{j=1}^M e^{-i\tilde{H}_j t/2} \prod_{j=M}^1 e^{-i\tilde{H}_j t/2}$$

Upper Bound			LIQUi)		
Molecule	Spin Orbitals	Basis	Molecule	Spin Orbitals	Basis
HF	12	sto6g	H2O (frozen core)	12	sto6g
FeMoco	16	tzvp	BeH2	14	sto6g
NH3	16	sto6g	H2O	14	sto6g
CH4	18	sto6g	CH4 (frozen core)	16	sto6g
F2	20	sto6g	FeMoco	16	tzvp
HCl	20	sto6g	NH3	16	sto6g
H2S	22	sto6g	Li2	20	sto6g
FeMoco	24	tzvp	HCl	20	sto6g
H2CO	24	tzvp	F2	20	sto6g
H2O	26	p321	CH4	20	sto6g
CO2	30	sto3g	H2S	22	sto6g
O3	30	sto6g			
H2O	38	dzvp			
H2O	50	p6311ss			
CO2	54	p321			
H2O	62	p6311ss			
CO2	90	dzvp			
FeMoco	108	tzvp			
Fe2S2	112	sto3g			
FeMoco	114	tzvp			

TABLE III. Table contains the identities of each molecule sorted first by the number of *spin* orbitals, which is twice the number of spatial orbitals, and then by the actual, or upper bounded, Trotter number.

is a unitary operator. The matrix logarithm is defined if and only if the matrix in question is invertible and hence the matrix logarithm exists because unitary matrices are invertible. The logarithm is then clearly an anti-Hermitian operator and hence $\tilde{H}(t)$ is Hermitian. This implies that $\prod_{j=1}^M e^{-iH_j t/2} \prod_{j=M}^1 e^{-iH_j t/2} \equiv e^{-i\tilde{H}(t)t}$.

The triangle inequality implies that

$$\begin{aligned} & \|e^{-iHt} - e^{-i\tilde{H}(t)t}\|_2 \leq \\ & \left\| e^{-iHt} - \prod_{j=1}^M e^{-iH_j t/2} \prod_{j=M}^1 e^{-iH_j t/2} \right\|_2 \\ & + \left\| \prod_{j=1}^M e^{-iH_j t/2} \prod_{j=M}^1 e^{-iH_j t/2} - e^{-i\tilde{H}(t)t} \right\|_2. \end{aligned} \quad (\text{D6})$$

Then using our assumptions and standard inequalities for the errors in unitary operations [49] this error is at most

$$[\Delta E_{TS}(t) + (2M - 1)\delta]t. \quad (\text{D7})$$

Then applying Lemma 4 we have that $\|H - \tilde{H}(t)\|_2 \leq \Delta E_{TS}(t) + (2M - 1)\delta$. The result then follows from applying Lemma 3. \square

These error bounds apply for generic quantum simulation based on the second order Trotter formula (also

known as the Strang splitting), however they apply in particular to the case of simulating quantum chemistry. We have already discussed the Trotter–Suzuki error for quantum chemistry in (C5). Even in the absence of decoherence, a source of error inevitably comes into our analysis from approximating the single qubit rotations. Now we need to show a result that bounds the impact of such eigenvalue estimation.

Lemma 5. *For Hermitian H and $t \geq 0$ let $U(t)$ be a family of unitary operations such that $\|e^{-iHt} - U(t)\|_2 \leq \epsilon t$ for all $t \geq 0$, then for every t there exists a Hermitian operator $\tilde{H}(t)$ such that $U = e^{-i\tilde{H}t}$ and $\|H - \tilde{H}(t)\|_2 \leq \epsilon$.*

Proof. By taking the principal matrix logarithm of $U(t)$ it is clear that there exists $\tilde{H}(t)$ such that $U = e^{-i\tilde{H}(t)t}$. Furthermore, using standard inequalities it is easy to see that $\|e^{-iHt} - e^{-i\tilde{H}(t)t}\|_2 \leq \|H - \tilde{H}(t)\|_2 t$. Because this bound is tight in the limit as $t \rightarrow 0$, if $\epsilon < \|H - \tilde{H}(t)\|_2$ then a contradiction is reached for t sufficiently small. Therefore since we require that the inequality hold for all $t \geq 0$, $\|H - \tilde{H}(t)\|_2 \leq \epsilon$. \square

It is easy to see from the fact that there is a one to one mapping between the Hamiltonian terms and the rotations performed in the circuit. Thus if the error in each such rotation is δt then the error in each term of the effective Hamiltonian is at most δ . This means that the total shift in energy from circuit synthesis is at most $M\delta$.

Corollary 1. Assume $H = \sum_{j=1}^M h_j P_j$ for $h_j \in \mathbb{R}$, Pauli operators P_j and $\|e^{-iHt} - \prod_{j=1}^M e^{-ih_j P_j t/2} \prod_{j=M}^1 e^{-ih_j P_j t/2}\|_2 \leq \Delta E_{TS}(t)t$. Also let each $e^{-iH_j t/2}$ be implemented using a unitary $U_j(t)$ such that $\|U_j(t) - e^{-iH_j t/2}\|_2 \leq \Delta_{\text{synth}}$ for all t . The error in the ground-state energy that results from such a simulation is at most $\Delta E_{TS} + (2M - 1)\Delta_{\text{synth}}/t$.

Proof. The result then follows directly from Theorem 1 and Lemma 5. \square

This shows that it suffices to choose synthesis error that shrinks linearly with the timestep used in the Trotter decomposition. Since the cost of circuit synthesis of rotations in the the Clifford + T gate library scales logarithmically with Δ_{synth}/t [52].

Appendix E: Cost estimates for nitrogenase

Fundamentally, two factors contribute to the cost of the quantum simulation (assuming that the user can prepare an exact copy of the ground state at negligible cost). The first is the cost of implementing the Trotter decomposition of the Hamiltonian and the second is the number of times that the Trotter circuit must be repeated in the phase estimation algorithm.

One might object that the number of time steps required in the Trotter–Suzuki decomposition also is a driving factor in the cost. Of course the Trotter decomposition is a major driver of the cost, but it comes in only indirectly through the cost of phase estimation. This is because, in principle, the phase estimation algorithm learns the eigenphases of a single Trotter step. The Trotter error can be made arbitrarily small by choosing shorter evolution times, but this in turn requires the phase estimation algorithm to take more steps. As the cost of phase estimation scales inversely with the desired uncertainty, this causes the cost to scale inversely with the time step used in the Trotter–Suzuki decomposition. Thus the cost of the simulation can be thought of as arising from only two sources, the cost of each depending on the error tolerances allowed for all three contributions to the error.

If we then define ϵ_1 to be the error in phase estimation, $\epsilon_2 := \Delta E_{TS}$ to be the error in the Trotter–Suzuki expansion and $\epsilon_3 := (2M - 1)\Delta_{\text{synth}}/t$ to be the error in circuit synthesis then it follows from the triangle inequality and Corollary 1 that the error in the ground-state energy is at most 0.1 mHartree if

$$\epsilon_1 + \epsilon_2 + \epsilon_3 \leq \epsilon := 10^{-4} \text{Ha}. \quad (\text{E1})$$

In this section we will focus on the target of 0.1 mHa level of accuracy, which is appropriate for quantitative calculations of the reaction rates.

We estimate the cost of the circuit using the number of T gates required in the algorithm, which is a function

of the form

$$C = 2M \left[\frac{\alpha}{\epsilon_1} \right] \left[\beta \sqrt{\frac{\epsilon}{\epsilon_2}} \right] \left(\gamma \log_2 \left(\frac{2M}{\epsilon_3} \left[\beta \sqrt{\frac{\epsilon}{\epsilon_2}} \right] \right) + \delta \right), \quad (\text{E2})$$

and then optimize over ϵ_1, ϵ_2 and ϵ_3 to minimize C subject to the constraint in (E1). Here α is the scaling constant for the phase estimation algorithm used, β is the Trotter number (or multiplicative factor by which t is decreased from 1 Ha⁻¹) needed to achieve an error of 0.1 mHartree in the ground-state energy estimate and γ and δ are the constants used in the quantum circuit synthesis algorithm. Here $M = 6.1 \times 10^6$ for nitrogenase in the 54 orbital basis and $M = 8.2 \times 10^6$ for nitrogenase in the 57 orbital basis using the circuits of [9].

The true optima of (E2) are difficult to find because of the factor of $\sqrt{\epsilon/\epsilon_2}$ in the logarithm. In order to simplify our optimization we instead choose ϵ_1, ϵ_2 and ϵ_3 to minimize

$$\tilde{C} = 2M \left(\frac{\alpha}{\epsilon_1} \right) \left(\beta \sqrt{\frac{\epsilon}{\epsilon_2}} \right) \left(\gamma \log_2 \left(\frac{2M}{\epsilon_3} \beta \right) + \delta \right), \quad (\text{E3})$$

subject to the same constraint. The global optimum of (E3) can be found directly from calculus, which allows near optimal parameters for (E2) to be found easily. The value of (E2) at these parameters is then an upper bound on the minimum of (E2) and so the estimates provided remain upper bounds (modulo assumptions about the Trotter error).

The “worst” case assumptions in Table IV correspond to only using rigorously proven upper bounds on the cost. These lead to estimates that are clearly extremely pessimistic. Even given our optimistic assumptions about the target computer, the worst case bounds suggest between millions and tens of thousands of years depending on whether parallelism is used.

The “pessimistic” assumptions use empirical scalings for circuit synthesis and phase estimation and use the worst scaling supported from our bounds on the Trotter error, but rescaled by the ratios observed between the actual Trotter numbers required and the theoretically predicted ones.

The “rescaled” case takes the rigorous upper bound for nitrogenase and divides it by the average ratio observed between the exact Trotter numbers and their upper bounds for the tractable molecules molecules. Rescaling the upper bound by a constant and applying least squares fitting to find the most consistent constant yields similar results. We suspect these rescaled estimates may provide the most realistic estimate of the Trotter number required to simulate nitrogenase.

The “optimistic” assumptions again use the same empirical scalings for PE and synthesis, but instead extrapolate the average scaling observed for the ensemble of molecules whose Trotter numbers we can compute and scaling up the result so that all of the data lies beneath the curve. This scaling is optimistic, as there is little evidence for a clear trend in the empirical data and the

Case	Gates	Time (100 MHz T gates)
Rigorous bound	1.0×10^{21}	3.2×10^5 years.
Clifford	1.4×10^{21}	–
Rigorous + PAR	3.2×10^{22}	8300 years.
Clifford	3.3×10^{22}	–
Pessimistic bound	7.9×10^{15}	30 months
Clifford	1.2×10^{16}	–
Pessimistic + PAR	2.3×10^{17}	31 days
Clifford	2.3×10^{17}	–
Rescaled bound	1.2×10^{15}	135 days
Clifford	1.8×10^{15}	–
Rescaled + PAR	3.5×10^{16}	120 hours
Clifford	3.5×10^{16}	–
Optimistic bound	1.6×10^{14}	19 days
Clifford	2.4×10^{14}	–
Optimistic + PAR	4.7×10^{15}	17 hours
Clifford	4.7×10^{15}	–

Case	α	β	γ	δ
Rigorous	8π	7×10^6	4	11
Pessimistic	$\pi/2$	1075	1.15	9.2
Rescaled	$\pi/2$	166	1.15	9.2
Optimistic	$\pi/2$	24	1.15	9.2

TABLE IV. Resource estimates for simulation of nitrogenase’s FeMoco in structure 1 which requires a small basis consisting of 108 spin orbitals. PAR uses $\gamma = 4$ and $\delta = 11$.

range provided is insufficient to meaningfully extrapolate out to 108 spin orbitals (54 spatial orbitals) or more. We provide this estimate because it provides the best scaling that could reasonably be claimed to be supported by the data.

1. Variance-based estimates

Such errors arise from three sources: the systematic error in the TS decomposition ϵ_{TS} , the statistical error tolerance for phase estimation ϵ_{QPE} , and the statistical error in synthesizing rotations from Clifford and T gates ϵ_{Rot} . We then require the total error to be $\epsilon_{\text{TS}} + \sqrt{\epsilon_{\text{QPE}}^2 + \epsilon_{\text{Rot}}^2} = 0.1$ mHa. The three uncertainties are then chosen such that the number of T gates required for the simulation is minimized given the target accuracy.

The previous analysis for the estimates in the error can be used within this expression for the error under the assumptions that the errors in QPE and synthesis are not adversarial. This approach was taken with the estimates in the main body, wherein the three dominant costs are optimized against each other to minimize the resources needed to achieve the 0.1 mHartree target. The

Case	Gates	Time (100 MHz T gates)
Rigorous bound	1.8×10^{21}	5.5×10^5 years.
Clifford	2.5×10^{21}	–
Rigorous + PAR	5.9×10^{22}	1.5×10^4 years.
Clifford	6.0×10^{22}	–
Pessimistic bound	1.2×10^{16}	3.8 years
Clifford	1.8×10^{16}	–
Pessimistic + PAR	3.5×10^{17}	48 days
Clifford	3.5×10^{17}	–
Rescaled bound	2.2×10^{15}	250 days
Clifford	3.1×10^{15}	–
Rescaled + PAR	6.3×10^{16}	8.8 days
Clifford	6.3×10^{16}	–
Optimistic bound	2.3×10^{14}	27 days
Clifford	3.5×10^{14}	–
Optimistic + PAR	6.6×10^{15}	23 hours
Clifford	6.6×10^{15}	–

Case	α	β	γ	δ
Rigorous	8π	9.5×10^6	4	11
Pessimistic	$\pi/2$	1233	1.15	9.2
Rescaled	$\pi/2$	225	1.15	9.2
Optimistic	$\pi/2$	25	1.15	9.2

TABLE V. Resource estimates for simulation of nitrogenase’s FeMoco in structure 2 which requires a small basis consisting of 114 spin orbitals. PAR uses $\gamma = 4$ and $\delta = 11$.

optimization process is exactly the same as that used to minimize the cost given in (E2), however a different constraint linking the three errors is used. This leads to modest reductions in the costs relative to the worst case bounds, which we provide in Tables IV and V.

2. PAR circuits

There are several approaches that can be taken to parallelize rotations. The first, often coined nesting, is discussed in [16, 76]. It involves taking terms that commute with each other in the Hamiltonian and grouping them together so that they can be executed simultaneously. In principle, this can lead to substantial reductions in the depth but in practice it is difficult to assess the performance of these schemes here because of the size of the molecule and the fact that we have chosen to restrict ourselves to lexicographic ordering. This means that if we are to estimate the impact that parallelization can bring to these calculations we need to introduce a method that can reduce the T -depth without changing the ordering of terms in the Trotter–Suzuki decomposition.

The PAR method gives a way to achieve this goal [12]. It works by teleporting a rotation into a state with probability 1/2 using only Clifford operations and a pre-

rotated ancilla. In the event that this method fails then instead of performing $R_z(\theta)$ it performs $R_z(-\theta)$. This can be corrected by teleporting a rotation $R_z(2\theta)$ into the qubit in question. Should this fail (and it will half the time) the rotation can be corrected by teleporting a rotation of $R_z(4\theta)$ and so on. This creates a geometric distribution of the number of pre-cached qubits needed to perform a given rotation. These rotation angles are known before hand and so can be prepared offline in parallel. Hence a logarithmic multiplicative overhead in space is needed to guarantee that enough ancilla qubits are prepared to perform the rotations with such high probability that it is unlikely that the cache of ancillae will ever be depleted.

In order to bound the number of ancillae needed in order to parallelize a factor of M rotations with high probability consider the following protocol.

- Divide the terms in the Trotter expansion of the Hamiltonian into blocks consisting of M sequential terms.
- For each rotation angle θ_j in a given block and each $j = 0, \dots, n-1$, prepare the states $R_z(2^j \theta_j)|+\rangle$ in parallel for a predetermined value of n .
- Implement each of the M (possibly sequential operations) using programmable ancilla rotations using at most n attempts, if the PAR circuit fails in each attempt then a failure is said to have occurred.
- If a failure occurs then implement the correct rotation and proceed to the next precached rotation.

This protocol can be used to perform the desired rotation and its performance is summarized below.

Theorem 2. *There exists a protocol for implementing PAR rotations that caches M rotations of the form $R_z(\theta_j), R_z(2\theta_j), \dots, R_z(2^{n-1}\theta_j)$ for $j = 1, \dots, M$ such that the expected number of rotations performed before a failure occurs is*

$$2^n(1 - (1 - 2^{-n})^M).$$

Proof. The proof is constructive. To see this consider the protocol discussed above. Such a protocol fails when all n PAR circuits fail for any of the M rotations cached, and a failure occurs when all n attempts at the rotation that have been precomputed are expended. The probability of such a failure is clearly 2^{-n} because the PAR circuit's success probabilities are independent and each attempt has success probability $1/2$ [12].

From the geometric distribution, the probability that no failure occurs in M trials is then simply $(1 - 2^{-n})^M$. Similarly, the probability that a failure occurs after precisely k attempts is $2^{-n}(1 - 2^{-n})^{k-1}$. Since there are only M rotations in the cache any branch that has more than

M successes can only yield M rotations. This implies that the mean is

$$\sum_{k=1}^M k 2^{-n} (1 - 2^{-n})^{k-1} + M (1 - 2^{-n})^M = 2^n (1 - (1 - 2^{-n})^M). \quad (\text{E4})$$

□

A consequence of Theorem 2 is that a PAR cache of M rotations that further caches the correction operations for n failures can reduce the T -depth by a factor of $2^n(1 - (1 - 2^{-n})^M)$. We can therefore adjust these parameters to substantially reduce the T -depth without adding a prohibitive number of rotations.

Note that as $n \rightarrow \infty$ (E4) approaches M as expected. This suggests that taking large n allows the PAR rotations to be parallelized more efficiently, but this comes at the price of requiring more T gates and hence increases the overheads of quantum error correction.

a. Factory approach

A major challenge with PAR facing costing it in a fault tolerant setting is that all the non-Clifford operations can be prepared simultaneously and offline. This means that if we take the cost model where only T gates are considered then we come to the absurd conclusion that the costs of all quantum simulation algorithms can be reduced to that of synthesizing a single rotation. In order to prevent such absurd tradeoffs we assume here that a delay of time equal to 1 T gate is included to model the measurement and feed-forward step that is needed for the programmable ancilla rotation. This fixed cost means that even if all of the T-gates are precached before hand then the time required for the remainder of the simulation will never be zero.

The above assumptions lead to an alternative approach to implementing PAR rotations, which we follow in the PAR costs in the following as well as the main body. Rather than constructing a large cache of rotations offline, it makes sense to produce the rotations just in time. Specifically if the cost of synthesizing a rotation is C T gates then it makes sense to have nC factories constantly producing rotation states of the form $R_z(2^j \theta_j)|+\rangle$ for $j = 1, \dots, n$. Each of these nC factories is staggered such that one set of factories finishes with their n states at least 1 cycle before the next rotation is needed. As soon as a set of factories finishes it then proceeds on to the next set of rotations needed (excluding those that are currently being generated). If a failure occurs, then the factory approach halts just like the traditional approach and waits for a rotation to be synthesized that applies the correct rotation online.

Imagine for the moment that the PAR circuit succeeds in C consecutive attempts. Since each attempt requires time equivalent to a single invocation of a T gate and the cost of synthesis is C T gates, the first set of factories

will have finished producing their states before the last set applies theirs. This means that with C such factories rotation states can be continuously generated even in the “worst case scenario” where each PAR circuit succeeds on the first attempt. This is why in this setting it does not make sense to use more than nC rotation factories given the assumption that the cost of each PAR attempt is 1 T gate and that the rotations for at most n failures are to be pre-cached.

Theorem 3. *The average time required per rotation to apply the factory-based PAR strategy, assuming each PAR application requires time at most equal to 1 T gate and all remaining Clifford operations are free is*

$$\left(2 - \frac{n+2}{2^n}\right) + \frac{C+n}{2^n}.$$

Proof. The probability of success in any PAR attempt is $1/2$ therefore the expected number of T gates that need to be applied before a solution is found is

$$\sum_{j=1}^n \frac{j}{2^n} + \frac{C+n}{2^n}. \quad (\text{E5})$$

The latter term gives the expected impact of a failure, which occurs with probability $1/2^n$, makes n PAR attempts and incurs a cost of C T gates in applying the fallback rotation. The result then follows from summing the geometric series. \square

We use the above theorem to compute the expected time required by PAR under our assumptions of the costs of feed forwarding. If such costs are neglected then the online cost of performing a PAR rotation is simply $C/2^n$ T gates.

b. Nesting estimates

Finally, we would like to reiterate that in both the serial and PAR cases, the costs are found by building the circuit corresponding to the TS formula and counting the gates that compose it. The estimates for nesting given in the main body are upper bounds based on empirical estimates of the number of terms that can be simultaneously executed. We estimate this by taking FeMoco and greedily grouping terms that act on distinct sets of spin orbitals. We specifically find that there exists a grouping that can simultaneously execute 26.43 terms simultaneously for structure 1 and 27.83 for structure 2. These values roughly correspond to the optimal scaling of $N/4$ that can be achieved through this nesting strategy. We then assume the quantum computer can simultaneously execute each of these commuting groups and find the corresponding time by dividing the T count by these factors.

Rather than executing the grouping in $\text{LIQU}i\rangle$ we use a simple upper bound on the number of Clifford gates that could arise from the grouping. We do this because

the grouping strategy breaks the lexicographic ordering that leads to dramatic cancelation of the CNOT strings that arise from the Jordan–Wigner decomposition.

For the nested data, the number of timesteps needed is assumed to be the same as for the other cases, which is reasonable given that operator ordering tends to not have a dramatic impact on the error. Subsequent work will investigate the precise interplay that operator ordering has in nesting.

Appendix F: State Preparation

We discuss in this section the issue of state preparation, which is a major unanswered question that impacts the cost of quantum simulations. While we do not discuss these costs in detail in the rest of the paper, we discuss below the issues that arise when using elementary state preparation methods based on coupled cluster, configuration interaction or Hartree–Fock states as well as adiabatic state preparation. We also provide numerical results showing that the ground state overlap with Hartree–Fock states scales for small molecules and discuss the costs involved in adiabatic state preparation.

1. Elementary state preparation methods

Although we cannot rigorously prove that elementary state preparation methods such as Hartree–Fock states, unitary coupled cluster or truncated configuration interaction states (such as CISD or difference dedicated CI (DDCI) [77] states) will suffice for preparing a state with large overlap with the ground state, it is still important to ask how good simple ansatzes perform for numerically tractable cases. We provide some numerical evidence for small molecules showing that the overlaps of the true ground state with the Hartree–Fock ground state is not necessarily small. We leave similar studies of the overlap for unitary coupled cluster and truncated configuration interaction ansatzes for subsequent work.

We see in Figure 8 that there is substantial overlap between the Hartree–Fock state and the true ground state of the molecules calculated using $\text{LIQU}i\rangle$. In particular, the smallest overlap that we see is 89%. Other studies that have looked at chains of hydrogen atoms that are near disassociation show very small overlaps with the Hartree–Fock states. Thus there are small molecules that can be constructed that are not well described by such ansatzes.

We find weak evidence for a downward trend in the overlap with the Hartree–Fock state as the size of the system grows. Specifically, we see that roughly 50% of the data points are well approximated by $107.75\% \times e^{-0.0076n}$ where n is the number of spin orbitals. This scaling would suggest that the overlap with the Hartree–Fock state for a molecule typical of this ensemble of the scale of FeMoco may be roughly 43%, we cannot say whether nitrogenase




FIG. 8. Percent overlap ($|\langle \psi | \psi_{\text{HF}} \rangle|^2 \times 100$) of the Hartree–Fock state with the electronic ground state computed by LIQUi). All integrals are computed in a sto6g basis except for H_2 and HeH^+ . The integrals for those molecules are computed using sto3g and 3-21g bases respectively. The blue dashed line shows a possible extrapolated trend from the data.

is indeed typical of this ensemble. Indeed, one may expect stronger correlations to be present for molecules that contain atoms with d -electrons than those examined in Figure 8. Such molecules are frequently outside of our ability to simulate classically so finding an appropriate ensemble of molecules to use for such benchmarks remains an open research problem; however, the success of methods such as CI dynamically extended active space (CI-DEAS) [78, 79] in DMRG suggests that elementary truncated configuration interaction states may also suffice for quantum simulation [80].

We further see no compelling evidence for scaling with the maximum nuclear charge of the constituent atoms for the molecules in this set. This can clearly be seen from BeH_2 which has substantially better fidelity with the ground state than Be does. Similarly, HF and HCl have nearly identical overlaps despite the fact that Cl has nearly twice the nuclear charge of F. More study is needed in order to understand how these overlaps scale for large molecules, however it is strongly suspected that multi-reference states will be required to accurately model highly correlated ground states.

2. Adiabatic State Preparation

Adiabatic state preparation [81] provides an alternative approach state preparation wherein the Hamiltonian to be simulated is replaced with a time-dependent Hamiltonian whose final ground state coincides with the target ground state and whose initial ground state is an easily

preparable state. An example of such a Hamiltonian is

$$H(s) = \sum_p h_{pp} a_p^\dagger a_p + \sum_{p,q} h_{pqqp} a_p^\dagger a_q^\dagger a_p a_q + s \left(H - \sum_p h_{pp} a_p^\dagger a_p - \sum_{p,q} h_{pqqp} a_p^\dagger a_q^\dagger a_p a_q \right), \quad (\text{F1})$$

where $s \in [0, 1]$ is a dimensionless time that is 0 at the beginning of the evolution and 1 at the end. It is then easy to see that the ground state of $H(0)$ is the Hartree–Fock state and the ground state of $H(1)$ is the full configuration interaction (FCI) ground state.

For such Hamiltonians (or more generally for those that are differentiable at least three times and whose resulting derivatives are $O(1)$ [82]) a sufficiently slow evolution under the Hamiltonian will cause the Hartree–Fock state to be transformed into the FCI ground state of H . In other words if \mathcal{T} is the time-ordering operator which is defined such that

$$\partial_s \mathcal{T}(e^{-i \int_0^s H(s') ds' T}) := -i H(s) \mathcal{T}(e^{-i \int_0^s H(s') ds' T}) \quad (\text{F2})$$

and if we define \mathcal{P} to be a projector onto the FCI ground state and $\Delta(s)$ to be the smallest eigenvalue gap between the ground state and the rest of the spectrum of $H(s)$ then

$$|(\mathbb{1} - \mathcal{P}) \mathcal{T}(e^{-i \int_0^1 H(s') ds' T}) | \psi_{\text{HF}} \rangle | \in O \left(\frac{\max_s \|\dot{H}(s)\|}{\min_s \Delta^2(s) T} \right). \quad (\text{F3})$$

Here we take $T \gg 1$ and treat the other parameters to be bounded above by a constant in this asymptotic expansion. For such Hamiltonians the triangle inequality clearly shows that $\|\dot{H}(s)\| \in O(N^4)$ (although on physical grounds we expect $\|\dot{H}(s)\|$ to be in $O(\eta^2) \in O(N^2)$ because the potential energy scales quadratically with the number of constituent particles). Thus if we want to have $O(1)$ probability of preparing the FCI ground state it suffices to simulate the time-dependent Hamiltonian $H(s)$ for time

$$T \in O \left(\frac{N^4}{\min_s \Delta^2(s)} \right). \quad (\text{F4})$$

Prima facie, the best known bounds on the costs of Trotter–Suzuki based simulation give the circuit size for such a simulation to be [83, 84]

$$N_{\text{operations}} \in \left(N^4 \left[\frac{h N^8}{\min_s \Delta^2(s)} \right]^{1+o(1)} \right), \quad (\text{F5})$$

where $h \geq \max\{|h_{pq}|, |h_{pqrs}|\}$. If we take $h \in O(1)$, this rigorous bound suggests that the cost of adiabatic state preparation may be prohibitively expensive even if the eigenvalue gap is constant, it is important to note that this upper bound on the norm of the derivative of H is

expected to be extremely loose and the scaling of the error in the Trotter–Suzuki formulas is expected to be much better than the scaling quoted above.

If we take the depth of second-order Trotter–Suzuki formula simulations of real molecules to scale as $O(N^{5.5})$, as observed in previous studies [17], and take the norm of the Hamiltonian to scale as $O(N^2)$ as anticipated asymptotically for a local basis [85] then the scaling that arises from using the second-order Trotter formula for time-dependent Hamiltonians [83] would be

$$O\left(\frac{h^{3/2}N^{8.5}}{\min_s \Delta^3(s)}\right). \quad (\text{F6})$$

If $h \in O(N^{-2})$ then this scaling reduces to $N^{5.5}/\min_s \Delta^3(s)$, which is expected if the two body terms dominate the cost of the Trotter–Suzuki decomposition. This means that even after making empirical assumptions, highly gapped adiabatic paths are likely to be necessary for adiabatic state preparation to be useful.

It is worth mentioning that adiabatic state preparation has been investigated for other systems and in these settings it has been found to be a highly practical method of state preparation [20, 75]. Although it should be noted that the paths used from the initial Hamiltonian to the final Hamiltonian are often non-trivial. These observations suggest that the above complexity analysis may be quite loose. Further work is needed to better estimate the cost of adiabatic state preparation for realistic molecules and also the cost of learning optimal adiabatic paths from easily preparable Hamiltonians to the FCI Hamiltonian.

Appendix H: FeMoco — the active site of nitrogenase

In this section we provide more background on the importance of biological dinitrogen fixation and on the active site model of nitrogenase prepared in different charge and spin states applied in the feasibility analysis of this work.

For decades, a Holy Grail in chemistry has been the catalytic fixation of molecular nitrogen under ambient conditions. Less than half a dozen synthetic catalysts have been developed for this purpose [25–27] (after decades of fruitless efforts). All of them suffer from low turnover numbers and the synthetic dinitrogen-fixation problem under ambient conditions can thus be considered largely unsolved. The process is of tremendous importance for society as fertilizers are produced from ammonia, the final product of dinitrogen fixation. Industrially,

Appendix G: Cost estimates for topological qubits

In this section we will examine the impact topological quantum computing may have on these numbers. This is important because the fault tolerant overheads quoted in the main body depend sensitively on the error rate. Topological quantum computing promises to provide physical error rates that are orders of magnitude beyond what is achievable in alternative platforms; however there is a catch. Much of the current research underway focuses on topological quantum computing using Ising anyons, which provide topological protection for Clifford operations but do not provide protection for T gates. This means that even if we achieve very low error rates using this technology then the costs of magic state distillation may not be reduced dramatically despite the quality of the Clifford gates that topological quantum computing affords. We examine this by modeling the errors in generating the T states to be 10^{-4} and then consider the costs of using the surface code to distill the necessary gates. We provide the data for this scenario in Table VI.

The data in Table VI shows that even in this setting assuming low quality magic states does not remove the benefit of topological protection for the Clifford operations. In particular, if we look at the savings in physical qubits that occur from going from error rate 10^{-6} to 10^{-9} we see in the data in the main body that roughly an order of magnitude separates the two numbers of physical qubits. In contrast, if we assume the lower quality magic states given above then the reductions are more modest: they are roughly a factor of 5. This shows that while having high fidelity magic states is ideal, even if such states are low quality then a topological quantum computer with protected Clifford operations can nonetheless see advantages from having topologically protected gates at the level of accuracy required to simulate nitrogenase.

ammonia is produced in the very efficient Haber–Bosch process, which, however, requires high temperature and pressure (and consumes up to 2% of the annual energy production) [86].

While it currently appears unrealistic that this simple heterogeneous process will be replaced by a sophisticated synthetic homogeneous catalyst, which is likely to be less stable and expensive to produce, a mono- or polynuclear iron-based catalyst working under ambient conditions and feeding on an easily accessible source for ‘hydrogen’ (and dinitrogen from air) could become important for local small-scale fertilizer-production concepts.

In any case, nitrogen fixation represents a tremendous chemical challenge to activate and break the strong triple bond in dinitrogen at room temperature and pressure. Nature found an efficient way to achieve this goal. It is accomplished by the enzyme nitrogenase whose active site, the iron–molybdenum cofactor FeMoco, consists of

	Serial rotations		PAR rotations		Nested rotations	
Clifford Error Rate	10^{-6}	10^{-9}	10^{-6}	10^{-9}	10^{-6}	10^{-9}
Required code distance	9,3	5,3	9,5	5,3	9,3	5,3
	Quantum processor					
Logical qubits	111		110		109	
Physical qubits per logical qubit	1013	313	1013	313	1013	313
Total physical qubits for processor	1.1×10^5	3.5×10^4	1.1×10^5	3.4×10^4	1.1×10^5	3.4×10^4
	Discrete Rotation factories					
Number	0		1872		26	
Physical qubits per factory	–	–	1013	313	1013	313
Total physical qubits for rotations	–	–	1.9×10^6	5.9×10^5	2.6×10^4	8.1×10^3
	T factories					
Number	64	30	41110	23248	1427	813
Physical qubits per factory	2.7×10^4	2.7×10^4	7.5×10^4	2.7×10^4	7.5×10^4	2.7×10^4
Total physical qubits for T factories	1.7×10^6	8.1×10^5	3.1×10^9	6.4×10^8	1.1×10^8	2.2×10^7
Total physical qubits	1.8×10^6	8.5×10^5	3.1×10^9	6.4×10^8	1.1×10^8	2.2×10^7

TABLE VI. This table gives the resource requirements including error correction for simulations of nitrogenase’s FeMoco in a 54 (spatial) orbital basis within the times quoted in Table I in the main body using physical gates operating at 100 MHz. Here we use error rates that are appropriate for quantum computing with Ising anyons, wherein topological protection is granted to Clifford operations but not to non-Clifford rotations. The error rate used in the production of the raw magic states is taken to be 10^{-4} in all of the above cases.

seven iron atoms and one molybdenum atom, which are clamped together by bridging sulfur atoms [87]. The complete structure of the FeMoco was solved only very recently when a central main-group atom was discovered [88] that, surprisingly, turned out to be a carbon atom [23, 24, 89]. The complex electronic structure of this cluster of open-shell iron atoms, the possible charge, spin, and protonation states as well as the different ligand and binding sites to be considered makes this active site a nightmare for electronic structure calculations, which is the basis of all theoretical approaches toward the elucidation of the mode of action of metalloenzymes such as nitrogenase.

It is thus no surprise that the specific mechanism of dinitrogen reduction at this active site has been elusive, especially given the fact that the mechanism of nitrogenase is difficult to study experimentally. Computational approaches suffer from the static electron correlation problem. For two ammonia molecules to be produced, the transfer of six protons and six electrons is required per dinitrogen molecule (in fact, eight protons and eight electrons are needed as one dihydrogen molecule is produced stoichiometrically). The transfer of these highly reactive agents leads to many stable intermediates and side products (see, for example, the analogous discussion of these steps in Ref. [28] for the first synthetic dinitrogen-fixating complex by Yandulov and Schrock). To elucidate the mechanism of nitrogenase, which is important for a better understanding of the activation of inert bonds by synthetic catalysts, therefore requires the consideration of a very large number of molecular struc-

tures.

While molecular structure of stable intermediates and transition states may be optimized within unrestricted Kohn–Sham DFT, the calculation of their energies demands an accurate wave-function-based approach. We therefore optimize molecular structures of a FeMoco model in the resting structure (Fig. 1 (right) in the main article) for varying charge and spin states in order to create different electronic situations that challenge the feasibility analysis presented in this work. For these structures, integrals in a molecular orbital basis have been obtained that parametrize the electronic structure of the cluster in the second-quantized quantum-chemical Hamiltonian of Eq. (A6).

Appendix I: Exact diagonalization techniques in chemistry

The electronic structure of a molecular structure determines its reactivity. Predicting chemical reactions requires the solution of the electronic Schrödinger equation to obtain the electronic energy and wave function. Whereas the expansion of the many-electron wave function into a (quasi-) complete many-electron basis will produce the exact solution, called full configuration interaction in chemistry or exact diagonalization in physics, this approach is unfeasible for molecules of more than a few atoms. As all standard quantum-chemical solution approaches construct many-electron basis functions from orbitals, the number of the former is determined by the

number of the latter. Exact diagonalization is therefore limited to about 18 electrons distributed among 18 spatial orbitals due to the exponential scaling of the many-electron basis states with the number of orbitals [39]. Unfortunately, the size of an orbital basis is already very large for moderately sized molecules. As a consequence, a restricted orbital space must be chosen.

Of all approximations developed in quantum chemistry to overcome this problem [36] the complete-active-space self-consistent-field (CASSCF) approach (and related models) utilizes exact diagonalization, but, because of the exponential scaling, in a reduced orbital space, the so-called CAS, that selects orbitals around the Fermi energy. To compensate for this approximation, the orbitals are relaxed self consistently, hence the name. Still, the CAS is limited by the 18-orbital wall and by the neglect of most of the (virtual) orbitals for the construction of the wave function. Considering the fact that molecules of a hundred atoms or more quickly require much more than a thousand one-electron basis functions for an accurate description of their electronic structure, most of the orbitals constructed from these basis functions are omitted in the construction of a CASSCF wave function. Even iterative techniques such as the density matrix renormalization group (DMRG), which can be understood as a polynomially scaling CASSCF approach, can push this wall only to about a hundred spatial orbitals.

As a result, a CASSCF-type wave function solves only the so-called static electron correlation problem. It is therefore particularly well suited for molecular structures with near-degenerate orbitals. The resulting electronic structure is only qualitatively well described. However, this feature is maintained throughout a reaction coordinate, which makes a CASSCF-type approach an appealing universal approach. For a quantitative description,

What chemical problems will benefit from such an implementation? Clearly, these will be problems that are dominated by strong static electron correlation (rather than by weak dispersion interactions originating from dynamic electron correlation). The electronic Schrödinger equation assigns an energy to a given molecular structure in the Born–Oppenheimer approximation, and this electronic energy (evaluated at zero Kelvin without vibrational, temperature, and entropy corrections) should dominate the energy change of a chemical process. Transition-metal catalysis is a field that presents such situations.

Many important chemical transformations are mediated by complicated electronic structures featured by transition-metal complexes. Especially late 3*d* transition metals are of this kind, most importantly iron, which is cheap and ubiquitous, often yields nontoxic compounds, and therefore represents an ideal catalytic center. Moreover, stable intermediates and, in particular, transition

the many virtual orbitals not considered for the CAS make a nonnegligible contribution to electronic-energy differences. To account for this so-called dynamic correlation is mandatory and typically achieved by a subsequent multi-reference perturbation theory calculation. Such perturbation theories to second order require elements of the three- and four-electron reduced density matrices, which are difficult to calculate and to store.

Instead of this ‘diagonalize, then perturb’ approach, also ‘perturb, then diagonalize’ ideas have been studied in chemistry. These latter approaches produce a ‘dressed’ many-electron Hamiltonian which is then better conditioned for a CASSCF-type approach. A new development in this area is the combination of density functional theory (DFT), known to describe dynamic electron correlation well with a CASSCF-type approach by spatial range separation introduced by Savin (see Refs. [42, 43] and references cited therein). Range separation is accomplished for the electron–electron interaction matrix elements and allows one to apply DFT at short range, whereas at long range the full flexibility of a CASSCF-type wave function can be exploited. This approach is very efficient and does not compromise the efficiency of a CASSCF-type approach (by contrast to perturbation theory). In combination with the polynomially scaling DMRG, this approach has delivered very promising results [43] for benchmark reactions involving transition metals [4].

As a consequence, a chemically sensible exact diagonalization technique is a CASSCF-type approach that captures dynamic correlation effects in the one-electron states. Such an approach could be directly implemented on a quantum computer with the specific exact diagonalization technology developed for such a machine.

states of a reaction mechanism often represent typical static electron correlation problems as bonds are formed and broken on the way to stable products. To reliably predict a chemical transformation of this kind usually requires to study many more than one elementary reaction step — especially when reactive intermediates are involved. The number of stable intermediates also increases due to unwanted side reactions that need to be inspected. Therefore, the total number of molecular structures whose electronic energy is required for an understanding of a reaction mechanism is, in general, very large.

Clearly, all these structures must be optimized with an efficient quantum chemical method. While the accuracy of electronic energies obtained with present-day DFT approaches is often not satisfactory for predictive purposes (see, e.g., Refs. [4, 90]), molecular structures can be reproduced with remarkable accuracy (obviously, a geometry gradient would replace DFT structures, too).

Structure	Structure opt./B3LYP		small-CAS CASSCF orbitals			
	Total spin S	Charge	Act. electrons	Act. orbitals	Total spin S	Charge
1	0	3	54	54	0	3
2	1/2	0	65	57	3/2	0

TABLE VII. The structures optimized for FeMoco and the settings for the small-CAS CASSCF orbital optimization. S is the total spin quantum number and the charge is measured in units of the elementary charge.

Atom	Coordinates Structure 1		
S	0.193509	-1.756174	-6.077728
FE	-0.073029	0.147462	-7.060398
S	-0.155670	-0.304014	-9.055299
FE	-1.194342	-0.633669	-4.857848
C	0.097884	0.293121	-3.822079
FE	0.330922	1.682738	-2.639449
S	0.144858	3.427365	-3.756994
FE	0.000866	1.806617	-5.056623
S	-1.759709	1.058486	-6.097784
FE	1.393308	-0.388901	-4.927653
S	2.935549	-1.083502	-3.683191
FE	1.581418	-0.451806	-2.290305
S	1.643565	1.242708	-0.922047
MO	-0.046077	-0.172415	0.259925
O	-0.081763	1.100331	1.760010
FE	-0.867802	-0.634367	-2.449654
S	-1.333714	1.160052	-1.286619
S	-2.410505	-1.577384	-3.437041
S	1.630797	1.240870	-6.345822
S	0.353345	-1.960563	-1.213579
N	1.652155	-0.910661	1.409953
O	-1.541962	-0.693648	1.167835
C	-2.111338	-0.028674	2.280713
C	-0.152453	1.144086	-10.129068
C	-1.083149	1.057101	2.683832
H	-3.060803	0.437015	2.003587
H	-2.276775	-0.745389	3.087505
O	-1.167081	1.740174	3.642482
H	0.625836	0.971552	-10.878592
H	-1.119129	1.154099	-10.642807
H	0.015059	2.077284	-9.598730
C	2.339711	-0.178076	2.302234
N	3.244459	-0.944382	2.902881
C	3.160159	-2.227901	2.407388
C	2.176178	-2.202712	1.475167
H	1.794821	-3.019854	0.890693
H	3.781670	-3.031530	2.764468
H	3.882526	-0.630293	3.624464
H	2.182421	0.864873	2.517705

TABLE VIII. Coordinates for Structure 1 of FeMoco in Å.

Atom	Coordinates Structure 2		
S	0.032866	-2.093214	-6.099450
FE	0.017622	-0.056816	-7.223052
S	0.143644	-0.401241	-9.315852
FE	-1.471854	-0.721074	-4.983076
C	-0.152873	0.225677	-3.664235
FE	-0.118135	1.845485	-2.405214
S	-0.036661	3.550618	-3.771677
FE	-0.049780	1.771611	-5.047025
S	-1.760858	1.077437	-6.479794
FE	1.306370	-0.657731	-4.880018
S	2.846629	-1.294232	-3.368478
FE	1.196355	-0.482357	-2.281465
S	1.566535	1.288016	-0.916398
MO	-0.186653	-0.006921	0.115165
O	-0.140683	1.329614	1.698525
FE	-1.529490	-0.648712	-2.302651
S	-1.922731	1.225686	-0.982284
S	-3.001156	-1.537533	-3.665135
S	1.731662	1.109781	-6.342619
S	-0.092449	-1.984660	-1.157533
N	1.598971	-0.845206	1.346832
O	-1.365000	-0.810743	1.369699
C	-1.748263	-0.161994	2.555155
C	0.321553	1.195240	-10.180658
C	-0.891249	1.088353	2.753014
H	-2.802446	0.128482	2.480632
H	-1.641761	-0.837543	3.407640
O	-0.904893	1.741420	3.768510
H	0.360712	0.987604	-11.249659
H	-0.529456	1.838638	-9.965894
H	1.243686	1.686635	-9.874888
C	2.191192	-0.163147	2.313386
N	3.212741	-0.875946	2.820353
C	3.284255	-2.074995	2.147363
C	2.276098	-2.042456	1.234236
H	1.996684	-2.787079	0.511609
H	4.022709	-2.823292	2.371900
H	3.814462	-0.573873	3.568856
H	1.907357	0.820486	2.645710

TABLE IX. Coordinates for Structure 2 of FeMoco in Å.

Hence, CASSCF-type methods will mostly be required for the validation of electronic energies of DFT-optimized molecular structures, which is the essential piece of information for establishing a reaction mechanism.

Appendix J: Computational Methodology

We optimized the structure of a FeMoco resting-state model that takes those residues of the protein backbone into account which anchor the metal cluster in the enzyme (see Fig. 1 (right) in the main article). Note that different spin and charge states were considered for the structure optimization in order to obtain a variety of electronic structures for the assessment of a solution algorithm on a quantum computer. These spin and charge states do not necessarily match the one of the resting state of nitrogenase.

Structure **1** for three positive excess charges and an equal number of α - and β -spins, and structure **2** for an uncharged FeMoco model with one unpaired α -spin.

The Cartesian coordinates of these structures are collected in Tables VIII and IX. In these unrestricted DFT calculations, the spin symmetry was broken [91] and only S_z remains as a good quantum number. For these structure optimizations we chose the TURBOMOLE program package (V6.4) [92] and employed the B3LYP density

functional [93–96] with the def2-TZVP Ahlrichs triple-zeta basis set plus polarization functions on all atoms [97]. An effective core potential was chosen only for the molybdenum atom [98], which also takes care of all scalar-relativistic effects on this heavy atom.

Then, integrals in the molecular orbital (MO) basis were produced MO integrals for structure **1** were generated for a CAS of 54 electrons in 54 spatial CASSCF orbitals (108 spin orbitals), which were obtained from a singlet CASSCF calculation with 24 electrons in 16 orbitals. Accordingly, for structure **2** we generated MO integrals for a CAS of 65 electrons in 57 spatial CASSCF orbitals (114 spin orbitals) from a quartet CASSCF calculation with 21 electrons in 12 orbitals.

The choice of the large active spaces was based on Pulay’s UNO-CAS criterion [37, 99], in which the occupation number of the natural orbitals serves as a selection criterion. However, rather than unrestricted Hartree–Fock natural orbitals, we selected those small-CAS CASSCF natural orbitals in the occupation intervals [1.98,0.02] and [1.99,0.01], respectively. The molecular orbital integrals for the second-quantized electronic Hamiltonian in these small-CAS orbital bases were calculated with the MOLCAS program [39].

All settings for the calculations are summarized in Table VII.

-
- [1] Clifford Dykstra, Gernot Frenking, Kwang S. Kim, and Gustavo E. Scuseria, *Theory and Applications of Computational Chemistry: The First Forty Years* (Elsevier, Amsterdam, 2005).
- [2] Christopher J. Cramer and Donald G. Truhlar, “Density Functional Theory for Transition Metals and Transition Metal Chemistry,” *Phys. Chem. Chem. Phys.* **11**, 10757–10816 (2009).
- [3] Wanyi Jiang, Nathan J. DeYonker, John J. Determan, and Angela K. Wilson, “Toward accurate theoretical thermochemistry of first row transition metal complexes,” *J. Phys. Chem. A* **116**, 870–885 (2012).
- [4] Thomas Weymuth, Erik P. A. Couzijn, Peter Chen, and Markus Reiher, “New benchmark set of transition-metal coordination reactions for the assessment of density functionals,” *J. Chem. Theory Comput.* **10**, 3092–3103 (2014).
- [5] Seth Lloyd, “Universal quantum simulators,” *Science* **273**, 1073 (1996).
- [6] Christof Zalka, “Simulating quantum systems on a quantum computer,” in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 454 (The Royal Society, 1998) pp. 313–322.
- [7] Daniel A Lidar and Haobin Wang, “Calculating the thermal rate constant with exponential speedup on a quantum computer,” *Physical Review E* **59**, 2429 (1999).
- [8] Hefeng Wang, Sabre Kais, Alán Aspuru-Guzik, and Mark R Hoffmann, “Quantum algorithm for obtaining the energy spectrum of molecular systems,” *Physical Chemistry Chemical Physics* **10**, 5388–5393 (2008).
- [9] Benjamin P Lanyon, James D Whitfield, Geoff G Gillett, Michael E Goggin, Marcelo P Almeida, Ivan Kassal, Jacob D Biamonte, Masoud Mohseni, Ben J Powell, Marco Barbieri, and Andrew G White, “Towards quantum chemistry on a quantum computer,” *Nature Chemistry* **2**, 106–111 (2010).
- [10] Ivan Kassal, James D Whitfield, Alejandro Perdomo-Ortiz, Man-Hong Yung, and Alán Aspuru-Guzik, “Simulating chemistry using quantum computers,” *Annual Review of Physical Chemistry* **62**, 185–207 (2011).
- [11] James D Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik, “Simulation of electronic structure hamiltonians using quantum computers,” *Molecular Physics* **109**, 735–750 (2011).
- [12] N Cody Jones, James D Whitfield, Peter L McMahon, Man-Hong Yung, Rodney Van Meter, Alán Aspuru-Guzik, and Yoshihisa Yamamoto, “Faster quantum chemistry simulation on fault-tolerant quantum computers,” *New Journal of Physics* **14**, 115023 (2012).
- [13] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L OBrien, “A variational eigenvalue solver on a photonic quantum processor,” *Nature communications* **5** (2014).
- [14] Dave Wecker, Bela Bauer, Bryan K Clark, Matthew B Hastings, and Matthias Troyer, “Gate-count estimates for performing quantum chemistry on small quantum computers,” *Physical Review A* **90**, 022305 (2014).
- [15] Jarrod R McClean, Ryan Babbush, Peter J Love, and Alán Aspuru-Guzik, “Exploiting locality in quantum

- computation for quantum chemistry,” *The journal of physical chemistry letters* **5**, 4368–4380 (2014).
- [16] Matthew B Hastings, Dave Wecker, Bela Bauer, and Matthias Troyer, “Improving quantum algorithms for quantum chemistry,” *Quantum Information & Computation* **15**, 1–21 (2015).
- [17] David Poulin, Matthew B Hastings, Dave Wecker, Nathan Wiebe, Andrew C Doberty, and Matthias Troyer, “The trotter step size required for accurate quantum simulation of quantum chemistry,” *Quantum Information & Computation* **15**, 361–384 (2015).
- [18] Dave Wecker, Matthew B Hastings, and Matthias Troyer, “Progress towards practical quantum variational algorithms,” *Physical Review A* **92**, 042303 (2015).
- [19] Ryan Babbush, Jarrod McClean, Dave Wecker, Alán Aspuru-Guzik, and Nathan Wiebe, “Chemical basis of trotter-suzuki errors in quantum chemistry simulation,” *Physical Review A* **91**, 022311 (2015).
- [20] Bela Bauer, Dave Wecker, Andrew J Millis, Matthew B Hastings, and Matthias Troyer, “Hybrid quantum-classical approach to correlated materials,” arXiv preprint arXiv:1510.03859 (2015).
- [21] L. M. Zhang, C. N. Morrison, J. T. Kaiser, and D. C Rees, “Nitrogenase mofe protein from clostridium pasteurianum at 1.08 angstrom resolution: comparison with the azotobacter vinelandii mofe protein,” *Acta Crystallogr.* **D71**, 274–282 (2015).
- [22] Brian M. Hoffman, Dmitriy Lukoyanov, Zhi-Yong Yang, Dennis R. Dean, and Lance C. Seefeldt, “Mechanism of nitrogen fixation by nitrogenase: The next stage,” *Chem. Rev.* **114**, 4041–4062 (2014).
- [23] Thomas Spatzal, Müge Aksoyoglu, Limei Zhang, Susana L. A. Andrade, Erik Schleicher, Stefan Weber, Douglas C. Rees, and Oliver Einsle, “Evidence for interstitial carbon in nitrogenase fmo cofactor,” *Science* **334**, 940 (2011).
- [24] Kyle M. Lancaster, Michael Roemelt, Patrick Ettenhuber, Yilin Hu, Markus W. Ribbe, Frank Neese, Uwe Bergmann, and Serena DeBeer, “X-ray emission spectroscopy evidences a central carbon in the nitrogenase iron-molybdenum cofactor,” *Science* **334**, 974–977 (2011).
- [25] Dmitry V. Yandulov and Richard R. Schrock, “Catalytic reduction of dinitrogen to ammonia at a single molybdenum center,” *Science* **301**, 76–78 (2003).
- [26] K. Arashiba, Y. Miyake, and Y. Nishibayashi, “A molybdenum complex bearing pnp-type pincer ligands leads to the catalytic reduction of dinitrogen into ammonia,” *Nature Chem.* **3**, 120–125 (2011).
- [27] John S. Anderson, Jonathan Rittle, and Jonas C. Peters, “Catalytic conversion of nitrogen to ammonia by an iron model complex,” *Nature* **501**, 84–87 (2013).
- [28] Maike Bergeler, Gregor N. Simm, Jonny Proppe, and Markus Reiher, “Heuristics-guided exploration of reaction mechanisms,” *J. Chem. Theory Comput.* **11**, 5712–5722 (2015).
- [29] Jgván Magnus Haugaard Olsen and Jacob Kongsted, “Molecular properties through polarizable embedding,” *Adv. Quantum Chem.* **61**, 107–143 (2011).
- [30] Christoph R. Jacob and Johannes Neugebauer, “Subsystem density-functional theory,” *WIREs Comput. Molec. Sci.* **4**, 325–362 (2014).
- [31] Arieh Warshel, “Multiscale modeling of biological functions: From enzymes to molecular machines (nobel lecture),” *Angew. Chem. Int. Ed.* **53**, 10020–10031 (2014).
- [32] Tomasz A. Wesolowski Sapana Shedge and Xiuwen Zhou, “Frozen-density embedding strategy for multilevel simulations of electronic structure,” *Chem. Rev.* **115**, 5891–5928 (2015).
- [33] Sebastian Wouters, Carlos A. Jiménez-Hoyos, Qiming Sun, and Garnet Kin-Lic Chan, “A practical guide to density matrix embedding theory in quantum chemistry,” arXiv , 1603.08443 (2016).
- [34] M. H. Olsson, J. Mavri, and A. Warshel, “Transition state theory can be used in studies of enzyme catalysis: lessons from simulations of tunnelling and dynamical effects in lipoxxygenase and other systems,” *Philos. Trans. Roy. Soc. Lond. B* **361**, 1417–1432 (2006).
- [35] David R. Glowacki, Jeremy N. Harvey, and Adrian J. Mulholland, “Taking ockham’s razor to enzyme dynamics and catalysis,” *Nature Chem.* **4**, 169–176 (2012).
- [36] Trygve Helgaker, Poul Jørgensen, and Jeppe Olsen, *Molecular Electronic-Structure Theory* (John Wiley & Sons, 2000).
- [37] Peter Pulay and Tracy P. Hamilton, “Uhf natural orbitals for defining and starting mc-scf calculations,” *J. Chem. Phys.* **88**, 4926–4933 (1988).
- [38] Christopher J. Stein and Markus Reiher, “Automated selection of active orbital spaces,” *J. Chem. Theory Comput.* **12**, 1760–1771 (2016).
- [39] Francesco Aquilante, Jochen Autschbach, Rebecca K. Carlson, Liviu F. Chibotaru, Mickael G. Delcey, Luca De Vico, Ignacio Fdez. Galván, Nicolas Ferré, Luis Manuel Frutos, Laura Gagliardi, Marco Garavelli, Angelo Giussani, Chad E. Hoyer, Giovanni Li Manni, Hans Lischka, Dongxia Ma, Per Åke Malmqvist, Thomas Müller, Artur Nenov, Massimo Olivucci, Thomas Bondo Pedersen, Daoling Peng, Felix Plasser, Ben Pritchard, Markus Reiher, Ivan Rivalta, Igor Schapiro, Javier Segarra-Mart, Michael Stenrup, Donald G. Truhlar, Liviu Ungur, Alessio Valentini, Steven Vancoillie, Valera Veryazov, Victor P. Vysotskiy, Oliver Weingart, Felipe Zapata, and Roland Lindh, “Molcas 8: New capabilities for multiconfigurational quantum chemical calculations across the periodic table,” *J. Comput. Chem.* , 506–541 (2016).
- [40] Steven R. White, “Density matrix formulation for quantum renormalization groups,” *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
- [41] Josep M. Bofill and Peter Pulay, “The unrestricted natural orbital-complete active space (uno-cas) method: An inexpensive alternative to the complete active space-self-consistent-field (cas-scf) method,” *J. Chem. Phys.* **90**, 3637–3646 (1989).
- [42] E. Fromager, J. Toulouse, and H. J. Å. Jensen, “On the universality of the long-/short-range separation in multiconfigurational density-functional theory,” *J. Chem. Phys.* **126**, 074111 (2007).
- [43] Erik Donovan Hedegård, Stefan Knecht, Jesper Skau Kielberg, Hans Jörgen Aagaard Jensen, and Markus Reiher, “Density matrix renormalization group with efficient dynamical electron correlation through range separation,” *J. Chem. Phys.* **142**, 224108 (2015).
- [44] Daniel S Abrams and Seth Lloyd, “Simulation of many-body fermi systems on a universal quantum computer,” *Physical Review Letters* **79**, 2586 (1997).
- [45] Dominic W Berry, Graeme Ahokas, Richard Cleve, and Barry C Sanders, “Efficient quantum algorithms for sim-

- ulating sparse hamiltonians,” *Communications in Mathematical Physics* **270**, 359–371 (2007).
- [46] Andrew M Childs and Nathan Wiebe, “Hamiltonian simulation using linear combinations of unitary operations,” *Quantum Information & Computation* **12**, 901–924 (2012).
- [47] Dominic W Berry, Andrew M Childs, Richard Cleve, Robin Kothari, and Rolando D Somma, “Simulating hamiltonian dynamics with a truncated taylor series,” *Physical review letters* **114**, 090502 (2015).
- [48] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Physical Review A* **86**, 032324 (2012).
- [49] Michael A Nielsen and Isaac L Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
- [50] Sergey Bravyi and Alexei Kitaev, “Universal quantum computation with ideal Clifford gates and noisy ancillas,” *Physical Review A* **71**, 022316 (2005).
- [51] Alex Bocharov, Martin Roetteler, and Krysta M Svore, “Efficient synthesis of probabilistic quantum circuits with fallback,” *Physical Review A* **91**, 052317 (2015).
- [52] Peter Selinger, “Efficient Clifford+ t approximation of single-qubit operators,” *Quantum Information & Computation* **15**, 159–180 (2015).
- [53] Sankar Das Sarma, Michael Freedman, and Chetan Nayak, “Majorana zero modes and topological quantum computation,” arXiv preprint arXiv:1501.02813 (2015).
- [54] Alex Bocharov, Yuri Gurevich, and Krysta M Svore, “Efficient decomposition of single-qubit gates into v basis circuits,” *Physical Review A* **88**, 012313 (2013).
- [55] Simon Forest, David Gosset, Vadym Kliuchnikov, and David McKinnon, “Exact synthesis of single-qubit unitaries over Clifford-cyclotomic gate sets,” *Journal of Mathematical Physics* **56**, 082201 (2015).
- [56] Christopher M Dawson and Michael A Nielsen, “The solovay-kitaev algorithm,” arXiv preprint quant-ph/0505030 (2005).
- [57] Vadym Kliuchnikov, Dmitri Maslov, and Michele Mosca, “Fast and efficient exact synthesis of single-qubit unitaries generated by Clifford and t gates,” *Quantum Information & Computation* **13**, 607–630 (2013).
- [58] Neil J Ross and Peter Selinger, “Optimal ancilla-free Clifford+ t approximation of z-rotations,” arXiv preprint arXiv:1403.2975 (2014).
- [59] Nathan Wiebe and Vadym Kliuchnikov, “Floating point representations in quantum circuit synthesis,” *New Journal of Physics* **15**, 093041 (2013).
- [60] Jacob T Seeley, Martin J Richard, and Peter J Love, “The bravyi-kitaev transformation for quantum computation of electronic structure,” *The Journal of chemical physics* **137**, 224109 (2012).
- [61] Austin G Fowler, Ashley M Stephens, and Peter Groszkowski, “High-threshold universal quantum computation on the surface code,” *Physical Review A* **80**, 052312 (2009).
- [62] R Barends, J Kelly, A Megrant, A Veitia, D Sank, E Jeffrey, TC White, J Mutus, AG Fowler, B Campbell, *et al.*, “Superconducting quantum circuits at the surface code threshold for fault tolerance,” *Nature* **508**, 500–503 (2014).
- [63] Adam Paetznick and Ben W Reichardt, “Universal fault-tolerant quantum computation with only transversal gates and error correction,” *Physical review letters* **111**, 090505 (2013).
- [64] H Bombin and MA Martin-Delgado, “Quantum measurements and gates by code deformation,” *Journal of Physics A: Mathematical and Theoretical* **42**, 095302 (2009).
- [65] Sergey Bravyi and Andrew Cross, “Doubled color codes,” arXiv preprint arXiv:1509.03239 (2015).
- [66] Cody Jones, “Multilevel distillation of magic states for quantum computing,” *Physical Review A* **87**, 042305 (2013).
- [67] HM Wiseman and RB Killip, “Adaptive single-shot phase measurements: A semiclassical approach,” *Physical Review A* **56**, 944 (1997).
- [68] Wim Van Dam, G Mauro D’Ariano, Artur Ekert, Chiara Macchiavello, and Michele Mosca, “Optimal phase estimation in quantum networks,” *Journal of Physics A: Mathematical and Theoretical* **40**, 7971 (2007).
- [69] Dominic W Berry, Brendon L Higgins, Stephen D Bartlett, Morgan W Mitchell, Geoff J Pryde, and Howard M Wiseman, “How to perform the most accurate possible phase measurements,” *Physical Review A* **80**, 052114 (2009).
- [70] Nathan Wiebe and Christopher E Granade, “Efficient bayesian phase estimation,” arXiv preprint arXiv:1508.00869 (2015).
- [71] Nathan Wiebe, Christopher Granade, Christopher Ferrie, and DG Cory, “Hamiltonian learning and certification using quantum resources,” *Physical review letters* **112**, 190501 (2014).
- [72] Alexander Hentschel and Barry C Sanders, “Machine learning for precise quantum measurement,” *Physical review letters* **104**, 063603 (2010).
- [73] Christopher E Granade, Christopher Ferrie, Nathan Wiebe, and David G Cory, “Robust online hamiltonian learning,” *New Journal of Physics* **14**, 103013 (2012).
- [74] Cristian Bonato, Machiel S Blok, Hossein T Dinani, Dominic W Berry, Matthew L Markham, Daniel J Twitchen, and Ronald Hanson, “Optimized quantum sensing with a single electron spin using real-time adaptive measurements,” *Nature nanotechnology* (2015).
- [75] Dave Wecker, Matthew B Hastings, Nathan Wiebe, Bryan K Clark, Chetan Nayak, and Matthias Troyer, “Solving strongly correlated electron models on a quantum computer,” *Physical Review A* **92**, 062318 (2015).
- [76] Sadegh Raeisi, Nathan Wiebe, and Barry C Sanders, “Quantum-circuit design for efficient simulations of many-body quantum dynamics,” *New Journal of Physics* **14**, 103017 (2012).
- [77] V. M. Garcaa, O. Castell, R. Caballol, and J. P. Malrieu, “An iterative difference-dedicated configuration interaction. proposal and test studies,” *Chem. Phys. Lett.* **238**, 222–229 (1995).
- [78] Ö. Legeza and J. Sólyom, **International Workshop on Recent Progress and Prospects in Density-Matrix Renormalization**, lorentz Center, Leiden University, The Netherlands, **2004**.
- [79] Ö. Legeza, **CECAM workshop for tensor network methods for quantum chemistry**, eTH Zürich, **2010**.
- [80] Sebastian F Keller and Markus Reiher, “Determining factors for the accuracy of dmrg in chemistry,” *CHIMIA International Journal for Chemistry* **68**, 200–203 (2014).
- [81] L-A Wu, MS Byrd, and DA Lidar, “Polynomial-time

- simulation of pairing models on a quantum computer,” *Physical Review Letters* **89**, 057904 (2002).
- [82] Donny Cheung, Peter Høyer, and Nathan Wiebe, “Improved error bounds for the adiabatic approximation,” *Journal of Physics A: Mathematical and Theoretical* **44**, 415302 (2011).
- [83] Nathan Wiebe, Dominic Berry, Peter Høyer, and Barry C Sanders, “Higher order decompositions of ordered operator exponentials,” *Journal of Physics A: Mathematical and Theoretical* **43**, 065203 (2010).
- [84] Nathan Wiebe, Dominic W Berry, Peter Høyer, and Barry C Sanders, “Simulating quantum dynamics on a quantum computer,” *Journal of Physics A: Mathematical and Theoretical* **44**, 445308 (2011).
- [85] Jarrod R McClean, Ryan Babbush, Peter J Love, and Alán Aspuru-Guzik, “Exploiting locality in quantum computation for quantum chemistry,” *The journal of physical chemistry letters* **5**, 4368–4380 (2014).
- [86] G. Jeffery Leigh, ed., *Nitrogen Fixation at the Millenium* (Elsevier Science, Amsterdam, 2002).
- [87] J. Kim and D. C. Rees, “Structural Models for the Metal Centers in the Nitrogenase Molybdenum-Iron Protein,” *Science* **257**, 1677–1682 (1992).
- [88] Oliver Einsle, F. Akif Tezcan, Susana L. A. Andrade, Benedikt Schmid, Mika Yoshida, James B. Howard, and Douglas C. Rees, “Nitrogenase MoFe-Protein at 1.16 Å Resolution: A Central Ligand in the FeMo-Cofactor,” *Science* **297**, 1696–1700 (2002).
- [89] M. W. Ribbe, Y. Hu, K. O. Hodgson, and B Hedman, “Biosynthesis of nitrogenase metalloclusters,” *Chem. Rev.* **114**, 4063–4080 (2014).
- [90] Caiping Liu, Tianbiao Liu, and Michael B. Hall, “Influence of the Density Functional and Basis Set on the Relative Stabilities of Oxygenated Isomers of Diiron Models for the Active Site of [FeFe]-Hydrogenase,” *J. Chem. Theory Comput.* **11**, 205–214 (2015).
- [91] C. R. Jacob and M. Reiher, “Spin in density-functional theory,” *Int. J. Quantum Chem.* **112**, 3661–3684 (2012).
- [92] Reinhart Ahlrichs, Michael Bär, Marco Häser, Hans Horn, and Christoph Kölmel, “Electronic Structure Calculations on Workstation Computers: The Program System TURBOMOLE,” *Chem. Phys. Lett.* **162**, 165–169 (1989).
- [93] Chengteh Lee, Weitao Yang, and Robert G. Parr, “Development of the colle-salvetti correlation-energy formula into a functional of the electron density,” *Phys. Rev. B* **37**, 785–789 (1988).
- [94] A. D. Becke, “Density-Functional Exchange-Energy Approximation with Correct Asymptotic Behavior,” *Phys. Rev. A* **38**, 3098–3100 (1988).
- [95] Axel D. Becke, “Density-functional thermochemistry. iii. the role of exact exchange,” *J. Chem. Phys.* **98**, 5648–5652 (1993).
- [96] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch, “Ab initio calculation of vibrational absorption and circular dichroism spectra using density functional force fields,” *J. Phys. Chem.* **98**, 11623–11627 (1994).
- [97] F. Weigend and R. Ahlrichs, “Balanced Basis Sets of Split Valence, Triple Zeta Valence and Quadruple Zeta Valence Quality for H to Rn: Design and Assessment of Accuracy,” *Phys. Chem. Chem. Phys.* **7**, 3297–3305 (2005).
- [98] D. Andrae, U. Häußermann, M. Dolg, H. Stoll, and H. Preuß, “Energy-adjusted *ab initio* Pseudopotentials for the Second and Third Row Transition Elements,” *Theor. Chim. Acta* **77**, 123–141 (1990).
- [99] Sebastian Keller, Katharina Boguslawski, Tomasz Janowski, Markus Reiher, and Peter Pulay, “Selection of active spaces for multiconfigurational wavefunctions,” *J. Chem. Phys.* **142**, 244104 (2015).