



Elucidating reaction mechanisms on quantum computers

Markus Reiher^{a,1}, Nathan Wiebe^{b,1}, Krysta M. Svore^b, Dave Wecker^b, and Matthias Troyer^{b,c,2}

^aLaboratorium für Physikalische Chemie, ETH Zurich, 8093 Zurich, Switzerland; ^bStation Q Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052; and ^cTheoretische Physik and Station Q Zurich, ETH Zurich, 8093 Zurich, Switzerland

Edited by David P. DiVincenzo, Institute for Quantum Information, RWTH Aachen University, and accepted by Editorial Board Member Evelyn L. Hu May 24, 2017 (received for review December 31, 2016)

With rapid recent advances in quantum technology, we are close to the threshold of quantum devices whose computational powers can exceed those of classical supercomputers. Here, we show that a quantum computer can be used 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 used to probe these reaction mechanisms, to significantly increase their accuracy and enable hitherto intractable simulations. Our 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. Our results demonstrate that quantum computers will be able to tackle important problems in chemistry without requiring exorbitant resources.

quantum computing | quantum algorithms | reaction mechanisms

Chemical reaction mechanisms are networks of molecular structures representing short- or long-lived intermediates connected by transition structures. The relative energies of all stable structures determine the relative thermodynamical stability. Differences of the energies of local minima and connecting transition structures determine the rates of interconversion, i.e., the chemical kinetics of the process. As they enter exponential expressions, very accurate energy differences are required for the reliable evaluation of the rate constants. At its core, the detailed understanding and prediction of complex reaction mechanisms then 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. Although approximate approaches, such as density functional theory (DFT) (2), are very popular, their accuracy is often too low for quantitative predictions (see, e.g., refs. 3 and 4); this holds particularly true for molecules with many energetically close-lying orbitals. For such problems on classical computers, 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. 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 used 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 that is 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 we estimate how large and how fast a quantum computer would have to be to perform such calculations within a reasonable amount of time. Our findings set 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 dinitrogen into two ammonia molecules under ambient conditions. Whereas the industrial Haber–Bosch catalyst requires high temperature and pressure 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. Fig. 1 shows the MoFe protein of nitrogenase (Fig. 1, *Left*) and the FeMoco buried in this protein (Fig. 1, *Middle*). Despite the importance of this process for fertilizer production that makes nitrogen

Significance

Our work addresses the question of compelling killer applications for quantum computers. Although quantum chemistry is a strong candidate, the lack of details of how quantum computers can be used for specific applications makes it difficult to assess whether they will be able to deliver on the promises. Here, we show how quantum computers can be used to elucidate the reaction mechanism for biological nitrogen fixation in nitrogenase, by augmenting classical calculation of reaction mechanisms with reliable estimates for relative and activation energies that are beyond the reach of traditional methods. We also show that, taking into account overheads of quantum error correction and gate synthesis, a modular architecture for parallel quantum computers can perform such calculations with components of reasonable complexity.

Author contributions: M.R., N.W., K.M.S., D.W., and M.T. designed research; M.R., N.W., K.M.S., D.W., and M.T. performed research; M.R., N.W., K.M.S., D.W., and M.T. contributed new tools; N.W., K.M.S., D.W., and M.T. analyzed data; and M.R., N.W., K.M.S., D.W., and M.T. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. D.P.D. is a guest editor invited by the Editorial Board.

¹M.R. and N.W. contributed equally to this work.

²To whom correspondence should be addressed. Email: mtroyer@microsoft.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1619152114/-DCSupplemental.

Table 1. Simulation time estimates

Structure	T gates	Cl. gates	Δt (10 ns)	Δt (100 ns)	Qubits
Quantitatively accurate simulation (0.1 mHa)					
Structure 1					
Serial	1.1×10^{15}	1.7×10^{15}	130 d	3.6 y	111
Nesting	3.5×10^{15}	5.7×10^{15}	15 d	4.9 mo	135
PAR	3.1×10^{16}	3.1×10^{16}	110 h	1.5 mo	1,982
Structure 2					
Serial	2.0×10^{15}	3.1×10^{15}	240 d	6.6 y	117
Nesting	6.5×10^{15}	1.0×10^{16}	27 d	8.9 mo	142
PAR	6.0×10^{16}	6.0×10^{16}	210 h	2.9 mo	2,024
Qualitatively accurate simulation (1 mHa)					
Structure 1					
Serial	1.0×10^{14}	1.6×10^{14}	12 d	3.9 mo	111
Nesting	3.3×10^{14}	5.6×10^{14}	1.4 d	14 d	135
PAR	3.0×10^{15}	3.0×10^{15}	11 h	4.6 d	1,982
Structure 2					
Serial	1.9×10^{14}	3.0×10^{14}	22 d	7.2 mo	117
Nesting	6.0×10^{14}	9.9×10^{14}	2.5 d	25 d	142
PAR	5.5×10^{15}	5.5×10^{15}	20 h	8.3 d	2,024

Listed are the number of Clifford and T-gate operations, the estimate of the run time (Δt), and the number of logical qubits required to obtain energies within 0.1 mHa or 1 mHa for two different structures of FeMoco on a quantum computer. 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 [SI Appendix](#) for further details). These run times and gate counts are likely to be pessimistic.

rate theory (ART). 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 molecules such as enzymes (27, 28). 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 or open-shell transition metal complexes (such as FeMoco), then so-called strong static electron correlation plays a decisive role already in the ground state. Static electron correlations are even more pronounced for electronically excited states relevant in photophysical and photochemical processes such as light harvesting for clean energy applications. Such situations require multiconfigurational 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 (29). CAS-type approaches require the selection of orbitals for the CAS, usually from those around the Fermi energy, which can be automatized (30–33). Although 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 (34). The polynomially scaling density matrix renormalization group (DMRG) algorithm (35) 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 guaranteed.

Ways Quantum Computers Will Help Solve These Problems.

Molecular structure optimizations are commonly found 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 be solved by a quantum computer that implements a multiconfigurational 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 (36) or small-CAS CASSCF calculation. The missing dynamic correlation can then be implemented 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 compute 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 (37, 38). 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 to improve on the latter for the massive prescreening of structures and energies.

Quantum Simulation of Quantum Chemical Systems

Ground state energies on a quantum computer can be obtained by 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 QPE learns the phase $\phi = E_n t$ in a manner analogous to a Mach–Zehnder interferometer. Measuring the phase starting from an approximate ground state collapses the wave function into an exact eigenstate and gives its energy. Time evolution is implemented using the Trotter–Suzuki approach ([SI Appendix](#)), and efficient methods exist to implement each of the terms in a second quantized Hamiltonian (11). Although algorithms are known that can achieve better scaling than low-order Trotter–Suzuki methods for some problems (39–44), they are more challenging to lay out, as circuits and preliminary estimates suggest that they perform worse at this problem size. For these reasons, we focus on low-order Trotter–Suzuki methods here and leave the task of fully costing out these alternative methods for future work.

To achieve reliable results, a fault-tolerant implementation of the quantum algorithm is crucial. Encoding a single logical qubit in a number of physical qubits with a quantum error-correcting code, such as the surface code (45), protects the logical qubit against decoherence and other experimental imperfections.

Quantum error correction cannot directly protect any arbitrary quantum operation, but it can protect a discrete set of gates, from which any continuous quantum operation can be approximated to within arbitrarily small error (46). Approximation takes two steps. First, the exponentials in the time evolution 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 (47). The space and time overheads of state distillation render it the most costly aspect of quantum error correction, leading to a large overhead that is not

These estimates can be improved, if necessary, using the techniques we provide in *SI Appendix*. Specifically, we provide simulation circuits that reduce the depth of the quantum simulation by a factor of 4 and also give near-optimal phase estimation algorithms that use a cluster of quantum computers to estimate the phase. These methods not only can allow our algorithms to be run in much less time, if necessary, but also may be of independent interest to researchers.

Resource Requirements with Quantum Error Correction. We next add the overheads required to perform the simulation fault-tolerantly and summarize the costs for structure 1 in Table 2. The underlying resource calculations for the fault-tolerant overheads follow ref. 45. These costs depend on the physical error rates in the hardware, and we consider three cases corresponding to (i) a near-term error rate of 10^{-3} , (ii) error rates of 10^{-6} that may one day be achievable in superconducting qubits or ion traps, and (iii) 10^{-9} , which could be achievable for topological quantum computers (50).

We consider a high-level 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 intended for general purpose computing, whereas the other factories are intended for special purposes.

The subdivisions in our architecture need not be physical. Indeed, we implicitly assume that they are logical in our analysis, but thinking about the device from this perspective reveals that quantum computer architectures can be tailored to quantum simulation. In particular, such hardware can exploit the fact that such simulations only require magic states when the algorithm is performing single-qubit rotations. Optimizing against a fixed architecture designed for quantum simulation can furthermore reduce the bandwidth needed to control the device, make compilation easier, and simplify communication within the device.

We first observe in Table 2 that the number of logical qubits in the main quantum processor is only of the order of a hundred, which translates into tens of thousands to millions of physical qubits, which is challenging but not out of reach. Most of the qubits are used in the T factories, which each need fewer physical qubits than the main quantum processor. The number of T factories needed to perform the serial calculation is small, with 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 parallelize with the PAR approach, then these costs are more daunting. The number of T factories required increases by a factor of roughly 1,000. If we use the nesting approach, then the costs are an order of magnitude greater than the serial case in the number of qubits, with a comparable reduction in run times.

To summarize, our estimates suggest that a quantum computer that operates with 10-MHz gates with 10^{-3} error rates could be simulated within a reasonable time using either nesting or PAR, but the number of physical qubits required will likely

be prohibitive. On the other hand, the requirements are much more modest with the aspirational target of 100-MHz gates with 10^{-6} error rates; a large but manageable number of physical qubits suffices, and the need to parallelize is nowhere near as dire. Therefore, barring further developments in quantum simulation algorithms or quantum error correction, we should aim to see quantum computers that meet our aspirational goal emerge from quantum computing programs. New simulation methods such as in refs. 41–44 may one day provide decisive advantages for problems like FeMoco, but more work is needed to provide complete cost estimates for them and, in turn, a fair comparison between the costs that they incur relative to Trotter–Suzuki methods for problems of this size. Finally, these low gate errors could one day be provided by topological quantum computers (50), which underscores the value in developing such hardware.

Discussion

Although, 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 employed to elucidate this mechanism using a manageable amount of memory and time. The quantum computer is used here to obtain, validate, or correct the energies of intermediates and transition states and thus gives accurate activation energies for various transitions. The required space and time resources for simulating FeMoco using the 54-orbital basis and nesting are comparable to that of Shor's factoring algorithm for 4,096-bit numbers (45).

Parallelizing the quantum computation of the energy landscape will be crucial to providing answers within a timeframe of several days instead of several years. Bounding the number of repetitions of phase estimation needed to prepare the ground state from an initial ansatz remains an open problem (see *SI Appendix*), and parallelism may often be needed to allow us to tolerate low success probability. Quantum computers therefore must be designed with a scalable architecture in mind and also built with the realization that constructing a single quantum computer is insufficient to solve such tasks. Instead, we should aim to have quantum computers that can be built en masse, because clusters of quantum computers will be needed to scan over the many structures that need to be examined to identify and estimate all important reaction rates (25).

Finally, chemical reactions that involve strongly correlated species that are hard to describe by traditional multiconfiguration approaches are not just limited to nitrogen fixation: They are ubiquitous. They range from C–H bond activating catalysts; to those for hydrogen and oxygen production, carbon dioxide fixation, and transformation; to industrially useful compounds; 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.

- Dykstra C, Frenking G, Kim KS, Scuseria GE (2005) *Theory and Applications of Computational Chemistry: The First Forty Years* (Elsevier, Amsterdam).
- Cramer CJ, Truhlar DG (2009) Density functional theory for transition metals and transition metal chemistry. *Phys Chem Chem Phys* 11:10757–10816.
- Jiang W, DeYonker NJ, Determan JJ, Wilson AK (2012) Toward accurate theoretical thermochemistry of first row transition metal complexes. *J Phys Chem A* 116:870–885.
- Weymuth T, Couzijn EPA, Chen P, Reiher M (2014) New benchmark set of transition-metal coordination reactions for the assessment of density functionals. *J Chem Theor Comput* 10:3092–3103.
- Lloyd S (1996) Universal quantum simulators. *Science* 273:1073–1078.
- Zalka C (1998) Simulating quantum systems on a quantum computer. *Proc R Soc A Math Phys Eng Sci* 454:313–322.
- Lidar DA, Wang H (1999) Calculating the thermal rate constant with exponential speedup on a quantum computer. *Phys Rev E* 59:2429–2438.
- Wang H, Kais S, Aspuru-Guzik A, Hoffmann MR (2008) Quantum algorithm for obtaining the energy spectrum of molecular systems. *Phys Chem Chem Phys* 10:5388–5393.
- Lanyon BP, et al. (2010) Towards quantum chemistry on a quantum computer. *Nat Chem* 2:106–111.
- Kassal I, Whitfield JD, Perdomo-Ortiz A, Yung MH, Aspuru-Guzik A (2011) Simulating chemistry using quantum computers. *Annu Rev Phys Chem* 62:185–207.

- Whitfield JD, Biamonte J, Aspuru-Guzik A (2011) Simulation of electronic structure Hamiltonians using quantum computers. *Mol Phys* 109:735–750.
- Jones NC, et al. (2012) Faster quantum chemistry simulation on fault-tolerant quantum computers. *New J Phys* 14:115023.
- Peruzzo A, et al. (2014) A variational eigenvalue solver on a photonic quantum processor. *Nat Commun* 5.
- Wecker D, Bauer B, Clark BK, Hastings MB, Troyer M (2014) Gate-count estimates for performing quantum chemistry on small quantum computers. *Phys Rev A* 90:022305.
- McClean JR, Babbush R, Love PJ, Aspuru-Guzik A (2014) Exploiting locality in quantum computation for quantum chemistry. *J Phys Chem Lett* 5:4368–4380.
- Hastings MB, Wecker D, Bauer B, Troyer M (2015) Improving quantum algorithms for quantum chemistry. *Quan Inf Comput* 15:1–21.
- Poulin D, et al. (2015) The Trotter step size required for accurate quantum simulation of quantum chemistry. *Quan Inf Comput* 15:361–384.
- Wecker D, Hastings MB, Troyer M (2015) Progress towards practical quantum variational algorithms. *Phys Rev A* 92:042303.
- Babbush R, McClean J, Wecker D, Aspuru-Guzik A, Wiebe N (2015) Chemical basis of Trotter–Suzuki errors in quantum chemistry simulation. *Phys Rev A* 91:022311.
- Bauer B, Wecker D, Millis AJ, Hastings MB, Troyer M (2015) Hybrid quantum-classical approach to correlated materials. arXiv:1510.03859.
- Zhang LM, Morrison CN, Kaiser JT, Rees DC (2015) Nitrogenase MoFe protein from *Clostridium pasteurianum* at 1.08 Å resolution: Comparison with the *Azotobacter vinelandii* MoFe protein. *Acta Crystallogr D* 71:274–282.
- Hoffman BM, Lukoyanov D, Yang ZY, Dean DR, Seefeldt LC (2014) Mechanism of nitrogen fixation by nitrogenase: The next stage. *Chem Rev* 114:4041–4062.
- Spatzal T, et al. (2011) Evidence for interstitial carbon in nitrogenase FeMo cofactor. *Science* 334:940.
- Lancaster KM, et al. (2011) X-ray emission spectroscopy evidences a central carbon in the nitrogenase iron-molybdenum cofactor. *Science* 334:974–977.
- Bergeler M, Simm GN, Proppe J, Reiher M (2015) Heuristics-guided exploration of reaction mechanisms. *J Chem Theor Comput* 11:5712–5722.
- Warshel A (2014) Multiscale modeling of biological functions: From enzymes to molecular machines (Nobel Lecture). *Angew Chem Int Ed* 53:10020–10031.
- Olsson MH, Mavri J, Warshel A (2006) Transition state theory can be used in studies of enzyme catalysis: Lessons from simulations of tunnelling and dynamical effects in lipoygenase and other systems. *Philos Trans R Soc Lond B* 361:1417–1432.
- Glowacki DR, Harvey JN, Mulholland AJ (2012) Taking Ockham's razor to enzyme dynamics and catalysis. *Nat Chem* 4:169–176.
- Helgaker T, Jørgensen P, Olsen J (2000) *Molecular Electronic-Structure Theory* (John Wiley, New York).
- Pulay P, Hamilton TP (1988) UHF natural orbitals for defining and starting MC-SCF calculations. *J Chem Phys* 88:4926–4933.
- Stein CJ, Reiher M (2016) Automated selection of active orbital spaces. *J Chem Theor Comput* 12:1760–1771.
- Stein CJ, Reiher M (2017) Automated identification of relevant frontier orbitals for chemical compounds and processes. *Chimia* 71:170–176.
- Stein CJ, Reiher M (February 22, 2017) Measuring multi-configurational character by orbital entanglement. *Mol Phys*, 10.1080/00268976.2017.1288934.
- Aquilante F, et al. (2016) Molcas 8: New capabilities for multiconfigurational quantum chemical calculations across the periodic table. *J Comput Chem* 37:506–541.
- White SR (1992) Density matrix formulation for quantum renormalization groups. *Phys Rev Lett* 69:2863–2866.
- Bofill JM, Pulay P (1989) 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.
- Fromager E, Toulouse J, Jensen HJA (2007) On the universality of the long-/short-range separation in multiconfigurational density-functional theory. *J Chem Phys* 126:074111.
- Hedegård ED, Knecht S, Kielberg JS, Jensen HJA, Reiher M (2015) Density matrix renormalization group with efficient dynamical electron correlation through range separation. *J Chem Phys* 142:224108.
- Berry DW, Ahokas G, Cleve R, Sanders BC (2007) Efficient quantum algorithms for simulating sparse Hamiltonians. *Commun Math Phys* 270:359–371.
- Childs AM, Wiebe N (2012) Hamiltonian simulation using linear combinations of unitary operations. *Quan Inf Comput* 12:901–924.
- Berry DW, Childs AM, Cleve R, Kothari R, Somma RD (2015) Simulating Hamiltonian dynamics with a truncated Taylor series. *Phys Rev Lett* 114:090502.
- Babbush R, et al. (2015) Exponentially more precise quantum simulation of fermions ii: Quantum chemistry in the CI matrix representation. arXiv:1506.01029.
- Babbush R, et al. (2016) Exponentially more precise quantum simulation of fermions in second quantization. *New J Phys* 18:033032.
- Kivlichan ID, Wiebe N, Babbush R, Aspuru-Guzik A (2016) Bounding the costs of quantum simulation of many-body physics in real space. arXiv:1608.05696.
- Fowler AG, Mariantoni M, Martinis JM, Cleland AN (2012) Surface codes: Towards practical large-scale quantum computation. *Phys Rev A* 86:032324.
- Nielsen MA, Chuang IL (2010) *Quantum Computation and Quantum Information* (Cambridge Univ Press, New York).
- Bravyi S, Kitaev A (2005) Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys Rev A* 71:022316.
- Bocharov A, Roetteler M, Svore KM (2015) Efficient synthesis of probabilistic quantum circuits with fallback. *Phys Rev A* 91:052317.
- Selinger P (2015) Efficient Clifford+T approximation of single-qubit operators. *Quan Inf Comput* 15:159–180.
- Sarma SD, Freedman M, Nayak C (2015) Majorana zero modes and topological quantum computation. arXiv:1501.02813.