

Towards Quantum Simulations

Elucidating Reaction Mechanisms on Quantum Computers

Markus Reiher, Nathan Wiebe, Krysta M. Svore, Dave Wecker, and Matthias Troyer
arXiv:1605.03590

Nitrogenase: A Draft Mechanism

Brian M. Hoffman, Dmitriy Lukoyanov, Dennis R. Dean, and Lance C. Seefeldt, *Acc. Chem. Res.* 2013 Feb 19; 46(2): 587595.

Recommended by David DiVincenzo, RWTH Aachen, FZ Juelich.

A quantum computer can effectively simulate other quantum systems: this is the idea that ushered in the quantum computer in Richard Feynman's 1981 Endicott House lecture. However, a distinct picture of how this new computer would actually function as a simulator has taken very long to emerge, while quantum computation's role in the solution of number theoretic problems, completely unanticipated in 1981 and for many years after, became very clear and consequential with Shor's work. Quantum simulation has definitely progressed in recent years, and the paper featured here does not represent the only promising step towards making it a reality. But in my opinion, the work of Reiher *et al.* gives the most concrete picture to date of how quantum simulation would actually work, in close conjunction with traditional quantum chemical and quantum materials modelling techniques, to solve a central scientific and technological problem in nitrogen biochemistry.

The subject taken up is nitrogen fixation. This is accomplished in nature by the enzyme nitrogenase, whose action is arguably the most central biochemical reaction in the living world (but which enzymes would you really be happy to do without?). A billion years of evolution have produced a remarkably complex active center in the nitrogenase protein; this center contains seven Fe atoms and (usually) one Mo atom, bonded together by nine bridging sulfurs. This makes the "FeMoco" (iron-molybdenum cofactor). In nature Mo is sometimes replaced by V or another Fe, but at the price of much lower activity – the action of this cofactor is remarkably specific. Six of the Fe atoms define an interstitial site, and it was only uncovered recently that this interstitial is definitely occupied, by one C atom.

This must be the most un-organic carbon in living matter, it is the kind of carbon site that we know best from austenitic steel.

This cluster manages something unattainable in the chemical engineering world; with the greatest efficiency, under ambient conditions, it receives an N_2 molecule one at a time, rips apart its triple bond, shepherds a set of protons to the reaction site, and releases two ammonia molecules for biological use (and, gratuitously, an extra H_2). I don't think that Hoffman *et al.* will be of detailed interest to most of the CMP readers, but a perusal of it can give an impression of the vast number of surmises that are needed for a detailed understanding of how the steps of the full reaction take place, each of which are a subject for study using computational simulation. The best available quantum-chemical or first-principles techniques (for the classical computer) have not proven to be up to the task.

As detailed by Reiher *et al.*, this is where the quantum computer can give a crucial boost to the predictive power of simulation. The basic algorithm to be employed, known as "Phase Estimation", has been well known since the initial wave of quantum algorithms in the 1990s. In Phase Estimation, if an eigenstate of a Hamiltonian is "loaded" into a quantum memory, and the time evolution under the quantum-chemical Hamiltonian is "enacted" by quantum gates, then an accurate estimate of the eigenenergy is obtained by measurement. Reiher *et al.* discuss a set of new strategies for the "enacting", showing that orders of magnitude of quantum computer running time are saved by nested or parallel strategies for this part of the algorithm. They also analyse the question of "loading"; first, there is the observation that even the most sophisticated classical quantum chemical (e.g., configuration interaction) representations of the ground-state wavefunction have a simple enough form that their "loading" into the state of a qubit register is feasible.

Second, they study the question of whether these approximate wavefunctions are close enough to the exact ground state. Phase Estimation does not improve the wavefunction; if the loaded wavefunction is a superposition of exact eigenstates, it will, with the corresponding probability, give back an accurate value of one of the eigenenergies in the superposition. Reiher *et al.* provide evidence that the best classical wavefunctions do have a sizable (perhaps on the order of percents) overlap with the exact ground state wavefunction, so that a sufficient number of runs of Phase Estimation will hit on the exact eigenenergy.

The paper provides a concrete estimate of how big a quantum computer will be needed to make real advances on the nitrogenase problem possible. The numbers found will surely be

a disappointment to some – the hundred-qubit computers that are in sight in laboratories today are definitely not capable of anything useful in this direction. But with perhaps “only” millions of qubits, quantum computers will have a chance to play a role, in conjunction with state-of-the-art classical simulation, and combined with the full savvy and intuition of the theoretical biochemist, to make an attack on this vital basic scientific problem.

Finally, there is also the prospect of a future technological angle. With a reliable classical/quantum catalytic-enzyme simulation tool, why stop at understanding what nature has made? The synthetic biochemical engineer would like to have something to replace the energy-hungry Haber-Bosch synthetic nitrogen fixing process. Existing efforts to emulate the prowess of the lowly rhizobia have had only modest success. If fully predictive quantum/classical simulation engines can aid in these efforts, perhaps a new technology could be forged.