

Why Quantum Chemistry Is Hard

by Scott Aaronson

Appeared in *Nature Physics News & Views*, 5(10):707-708, 2009

As someone who studies quantum computing, I'm often asked whether I worry that useful quantum computers will never be built in our lifetimes—and hence, that all the research devoted to studying these hypothetical devices will have been wasted. I try to convey that the *real* reason to study quantum computing has nothing to do with the distant prospect of building a machine that could factor 10,000-digit numbers in an instant. Rather, it's the transformative insights that we get right now into physics and computation. And there are few better illustrations of such insights than that provided by the paper of Schuch and Verstraete.

A central problem of quantum chemistry is to compute the ground states—that is, the lowest-energy states—of systems of electrons interacting via the Coulomb force. In this way, one can work out, for example, the spatial configuration of a molecule. Schuch and Verstraete study the computational difficulty of finding electron ground states in the framework of Density Functional Theory (DFT). The achievement of DFT—for which one of its developers, Walter Kohn, received the 1998 Nobel Prize in Chemistry—was to split the problem of computing ground states into two parts. The first part is to compute or approximate a “universal functional” governing the interaction between electrons; this functional is the same for every molecule with the same number of electrons. The second part is to minimize the sum of that universal functional and a term that applies to each electron separately. Because of the counterintuitive nature of quantum mechanics—wherein the amount of information needed to describe a physical system increases exponentially with the number of particles—a minimization problem involving just a single electron can be exponentially easier than one involving many electrons. Thus, DFT made possible computing the properties of molecules previously considered beyond reach for even the most powerful computers.

However, DFT also left a crucial question unanswered: *how hard is it to compute the universal functional?* In practice, we know of fairly efficient ways to approximate this functional, but sometimes those methods break down. Is the breakdown necessary, or simply a result of insufficient cleverness? While not giving definitive answers, the modern theory of computational complexity offers a powerful toolkit for such questions, one that lets us go beyond intuition or trial-or-error. In particular, we can argue that a problem is hard by showing that it's "at least as hard" as hundreds of other problems that have resisted efficient solution for decades.

This is what Schuch and Verstraete do. Their main result is that minimizing the energy for a system of N interacting electrons is "QMA-complete." Intuitively, this means that the problem is at least as hard as any *other* problem involving minimizing energy of a quantum system. If you're wondering, the QMA stands for "Quantum Merlin Arthur": the terminology comes from a thought experiment wherein an omniscient prover, Merlin, has to provide a skeptical verifier, Arthur, with a short "proof" (in this case, the ground state of a quantum system) in order to convince Arthur of the truth of some proposition (in this case, that the ground state energy is at most some stated value). The proof can be quickly verified once found (at least by a quantum computer), but it might be fiendishly difficult to find.

Now, minimizing a single-electron energy, *given* the ability to compute the Universal Functional, turns out to be easy (in technical terms, it is a convex optimization problem). And thus, an implication of Schuch and Verstraete's result is that we shouldn't hope for a fast, general-purpose method to compute the Universal Functional. For if such a method existed, then all QMA problems would be efficiently solvable, which is considered vanishingly unlikely. Rather, we should expect that in the worst case, every algorithm to compute the Universal Functional will take time that scales exponentially with the number of electrons N .

Admittedly, the same conclusion also follows from a simpler and earlier-known fact: that minimizing electron energies belongs to the class NP-hard. Roughly speaking, that's the class of problems that are at least as hard as the notorious Traveling Salesman Problem, of finding the shortest route that visits a large

number of cities. NP stands for “Nondeterministic Polynomial-Time”: the class of problems that might be hard to solve, but for which a solution, once found, can be efficiently recognized by a classical computer. QMA can be seen as the quantum generalization of NP. All QMA-hard problems are also NP-hard, but there are conjectured to be NP-hard problems that are not QMA-hard.

Now, it’s reasonable to ask: once we know a problem is NP-hard, isn’t that hard enough? NP-hardness already establishes a problem as intractable in the worst case, at least under the famous “ $P \neq NP$ ” conjecture. So why go to the additional step, as Schuch and Verstraete do, of proving the problem QMA-hard?

Here is where things get interesting. Suppose we consider a slight variant of the electron ground state problem, where we want to minimize the energy over all *pure* states, but are not interested in mixed states. In that case, minimizing a *single*-electron energy could already be a difficult NP problem. And thus, if we found a fast algorithm to compute the Universal Functional, the consequence would be, not to solve QMA problems, but “merely” to make the class QMA equal to the class NP—which is again considered unlikely. Thus, here we can get evidence that a practical problem is hard, but only by reasoning about a hypothetical collapse of “higher-level” computational classes. The conclusion really does depend on the fine-toothed distinction between QMA and NP, between quantum proofs and classical proofs.

In important respects, Schuch and Verstraete’s result is illustrative of quantum information science as a whole. This field does nothing to challenge the laws of quantum mechanics—the framework for almost all of physics since the 1920s. But it does ask a new set of questions about those laws: in this case, “what is the complexity of computing the DFT Universal Functional?” Because those questions straddle disciplines, they can look strange at first both to physicists and to computer scientists. But often enough they’ve turned out to have illuminating and nontrivial answers.