

Are Quantum States Exponentially Long Vectors?

Scott Aaronson*

I'm grateful to Oded Goldreich for inviting me to the 2005 Oberwolfach Meeting on Complexity Theory. In this extended abstract, which is based on a talk that I gave there, I demonstrate that gratitude by explaining why Goldreich's views about quantum computing are wrong.

Why should anyone care? Because in my opinion, Goldreich, along with Leonid Levin [10] and other "extreme" quantum computing skeptics, deserves credit for focusing attention on the key issues, the ones that ought to motivate quantum computing research in the first place. Personally, I have never lain awake at night yearning for the factors of a 1024-bit RSA integer, let alone the class group of a number field. The real reason to study quantum computing is not to learn other people's secrets, but to unravel the ultimate Secret of Secrets: *is our universe a polynomial or an exponential place?*

Last year Goldreich [7] came down firmly on the "polynomial" side, in a short essay expressing his belief that quantum computing is impossible not only in practice but also in principle:

As far as I am concern[ed], the QC model consists of exponentially-long vectors (possible configurations) and some "uniform" (or "simple") operations (computation steps) on such vectors . . . The key point is that the associated complexity measure postulates that each such operation can be effected at unit cost (or unit time). My main concern is with this postulate. My own intuition is that the cost of such an operation or of maintaining such vectors should be linearly related to the amount of "non-degeneracy" of these vectors, where the "non-degeneracy" may vary from a constant to linear in the length of the vector (depending on the vector). Needless to say, I am not suggesting a concrete definition of "non-degeneracy," I am merely conjecturing that such exists and that it capture[s] the inherent cost of the computation.

My response consists of two theorem-encrusted prongs:¹ first, that you'd have trouble explaining even current experiments, if you didn't think that quantum states really *were* exponentially long vectors; and second, that for most complexity-theoretic purposes, the exponentiality of quantum states is not that much "worse" than the exponentiality of classical probability distributions, which nobody complains about. The first prong is based on my paper "Multilinear Formulas and Skepticism of Quantum Computing" [1]; the second is based on my paper "Limitations of Quantum Advice and One-Way Communication" [2].

Prong 1: Quantum States *Are* Exponential

For me, the main weakness in the arguments of quantum computing skeptics has always been their failure to suggest an answer to the following question: *what criterion separates the quantum states we're sure we can prepare, from the states that arise in Shor's factoring algorithm?* I call such a criterion a "Sure/Shor separator." To be clear, I'm not asking for a red line partitioning Hilbert space into two regions, "accessible" and "inaccessible." But a skeptic could at least propose a complexity measure for quantum states, and then declare that a state of n qubits is "efficiently accessible" only if its complexity is upper-bounded by a small polynomial in n .

In his essay [7], Goldreich agrees that such a Sure/Shor separator would be desirable, but avers that it's not his job to propose one:

*Currently at the University of Waterloo and supported by ARDA. Email: aaronson@ias.edu. This abstract is mostly based on work done while I was a student at UC Berkeley, supported by an NSF Graduate Fellowship.

¹Sanjeev Arora asked why I don't have *three* prongs, thereby forming a ψ -shaped pitchfork.

My main disagreement with Scott is conceptual: He says that it is up to the “skeptics” to make a [concrete] suggestion (of such a “complexity”) and views their [arguments] as weak without such a suggestion. In contrast, I think it is enough for the “skeptics” to point out that there is no basis to the (over-simplified and counter-intuitive to my taste) speculation by which a QC can manipulate or maintain such huge objects “free of cost” (i.e., at unit cost).

Motivated by the “hands-off” approach of Goldreich and other skeptics, in [1] I tried to carry out the skeptics’ research program for them, by proposing and analyzing possible Sure/Shor separators. The goal was to illustrate what a scientific argument against quantum computing might look like.

For starters, such an argument would be careful to assert the impossibility only of *future* experiments, not experiments that have already been done. As an example, it would not dismiss exponentially-small amplitudes as physically meaningless, since one can easily produce such amplitudes by polarizing n photons each at 45° . Nor would it appeal to the “absurd” number of particles that a quantum computer would need to maintain in coherent superposition—since among other examples, the Zeilinger group’s C_{60} double-slit experiments [4] have already demonstrated “Schrödinger cat states,” of the form $\frac{|0\rangle^{\otimes n} + |1\rangle^{\otimes n}}{\sqrt{2}}$, for n large enough to be interesting for quantum computation.

Of course, the real problem is that, once we accept $|\psi\rangle$ and $|\varphi\rangle$ into our set of possible states, consistency almost *forces* us to accept $\alpha|\psi\rangle + \beta|\varphi\rangle$ and $|\psi\rangle \otimes |\varphi\rangle$ as well. So is there any defensible place to draw a line? This conundrum is what led me to investigate “tree states”: the class of n -qubit pure states that are expressible by polynomial-size *trees* of linear combinations and tensor products. As an example, the state

$$\left(\frac{|0\rangle + |1\rangle}{\sqrt{2}}\right) \otimes \dots \otimes \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}}\right)$$

is a tree state. For that matter, so is *any* state that can be written succinctly in the Dirac ket notation, using only the symbols $|0\rangle, |1\rangle, +, \otimes, (,)$ together with constants (no \sum ’s are allowed). In evaluating tree states as a possible Sure/Shor separator, we need to address two questions: first, should all quantum states that arise in present-day experiments be seen as tree states? And second, would a quantum computer permit the creation of non-tree states?

My results imply a positive answer to the second question: not only could a quantum computer efficiently generate non-tree states, but such states arise naturally in several quantum algorithms.² In particular, let C be a random linear code over \mathbb{GF}_2 . Then with overwhelming probability, a uniform superposition over the codewords of C cannot be represented by any tree of size $n^{\varepsilon \log n}$, for some fixed $\varepsilon > 0$.³ Indeed, $n^{\varepsilon \log n}$ symbols would be needed even to *approximate* such a state well in L_2 -distance, and even if we replaced the random linear code by a certain explicit code (obtained by concatenating the Reed-Solomon and Hadamard codes). I also showed an $n^{\varepsilon \log n}$ lower bound for the states arising in Shor’s algorithm, assuming an “obviously true” but apparently deep number-theoretic conjecture: basically, that the multiples of a large prime number, when written in binary, constitute a decent erasure code. All of these results rely on a spectacular recent advance in classical computer science: the first superpolynomial lower bounds on “multilinear formula size,” which were proven by Ran Raz [12] about a month before I needed them for my quantum application. Incidentally, in all of the cases discussed above, I conjecture that the actual tree sizes are exponential in n ; currently, though, Raz’s method can only prove lower bounds of the form $n^{\varepsilon \log n}$.⁴

Perhaps more relevant to physics, I also conjecture that 2-D and 3-D “cluster states” (informally, 2-D and 3-D lattices of qubits with pairwise nearest-neighbor interactions) have exponential tree sizes.⁵ If true, this conjecture suggests that states with enormous tree sizes might have already been observed in condensed-matter experiments—for example, those of Ghosh et al. [6] on long-range entanglement in magnetic salts.

²On the other hand, I do not know whether a quantum computer restricted to tree states always has an efficient classical simulation. All I can show is that such a computer would be simulable in $\Sigma_3^P \cap \Pi_3^P$, the third level of the polynomial-time hierarchy.

³In Raz’s proof, ε was about 10^{-6} , but what’s a constant between friends (or more precisely, between theoretical computer scientists)?

⁴I did manage to prove an exponential lower bound, provided we restrict ourselves to linear combinations $\alpha|\psi\rangle + \beta|\varphi\rangle$ that are “manifestly orthogonal”—which means that for all computational basis states $|x\rangle$, either $\langle\psi|x\rangle = 0$ or $\langle\varphi|x\rangle = 0$.

⁵By contrast, I can show that 1-D cluster states have tree size $O(n^4)$.

In my personal fantasy land, once the evidence characterizing the ground states of these condensed-matter systems became overwhelming, the skeptics would come back with a *new* Sure/Shor separator. Then the experimentalists would try to refute *that* separator, and so on. As a result, what started out as a philosophical debate would gradually evolve into a scientific one—on which progress not only can be made, but is.

Prong 2: It’s Not That Bad

To describe a state of n particles, we need to write down an exponentially long vector of exponentially small numbers, which themselves vary continuously. Moreover, the instant we measure a particle, we “collapse” the vector that describes its state—and not only that, but possibly the state of another particle on the opposite side of the universe. Quick, what theory have I just described?

The answer is classical probability theory. The moral is that, before we throw up our hands over the “extravagance” of the quantum worldview, we ought to ask: is it so much *more* extravagant than the classical probabilistic worldview? After all, both involve linear transformations of exponentially long vectors that are not directly observable. Both allow fault tolerance, in stark contrast with analog computing. Neither lets us reliably pack $n + 1$ bits into an n -bit state. And neither (apparently!) would provide enough power to solve NP-complete problems in polynomial time.

But none of this addresses the central complexity-theoretic question: if someone gives you a polynomial-size quantum state, *how much more useful* is that than being given a sample from a classical probability distribution? In their textbook [11], Nielsen and Chuang were getting at this question when they made the following intriguing speculation:

[W]e know that many systems in Nature “prefer” to sit in highly entangled states of many systems; might it be possible to exploit this preference to obtain extra computational power? It might be that having access to certain states allows particular computations to be done much more easily than if we are constrained to start in the computational basis.

To a complexity theorist like me, them’s fightin’ words—or at least, complexity-class-definin’ words. In particular, let’s consider the class BQP/qpoly, which consists of all problems solvable in polynomial time on a quantum computer, if the quantum computer has access to a polynomial-size “quantum advice state” $|\psi_n\rangle$ that depends only on the input length n . (For the uninitiated, BQP stands for “Bounded-Error Quantum Polynomial-Time,” and /qpoly means “with polynomial-size quantum advice.”) Note that $|\psi_n\rangle$ might be arbitrarily hard to prepare; for example, it might have the form $2^{-n/2} \sum_x |x\rangle |f(x)\rangle$ for an arbitrarily hard function f . We can imagine that $|\psi_n\rangle$ is given to us by a benevolent wizard; the only downside is that the wizard doesn’t know which input $x \in \{0, 1\}^n$ we’re going to get, and therefore needs to give us a single advice state that works for all x .

The obvious question is this: *is quantum advice more powerful than classical advice?* In other words, does $\text{BQP/qpoly} = \text{BQP/poly}$, where BQP/poly is the class of problems solvable in quantum polynomial time with the aid of polynomial-size *classical* advice? As usual in complexity theory, the answer is that we don’t know. This raises a disturbing possibility: could quantum advice be similar in power to *exponential-size* classical advice, which would let us solve any problem whatsoever (since we’d simply have to store every possible answer in a giant lookup table)? In particular, could BQP/qpoly contain the NP-complete problems, or the halting problem, or even *all* problems?

If you know me, you know I’d sooner accept that pigs can fly. But how to support that conviction? In [2], I did so with the help of yet another complexity class: PostBQP, or BQP with postselection. This is the class of problems solvable in quantum polynomial time, if at any stage you could measure a qubit and then *postselect* on the measurement outcome being $|1\rangle$ (in other words, if you could kill yourself if the outcome was $|0\rangle$, and then condition on remaining alive). My main result was that BQP/qpoly is contained in PostBQP/poly. Loosely speaking, anything you can do with polynomial-size quantum advice, you can

also do with polynomial-size *classical* advice, provided you’re willing to use exponentially more computation time (or settle for an exponentially small probability of success).⁶

On the other hand, since $\text{NP} \subseteq \text{PP}$, this result still says nothing about whether a quantum computer with quantum advice could solve NP-complete problems in polynomial time. To address that question, in [2] I also created a “relativized world” where $\text{NP} \not\subseteq \text{BQP}/\text{qpoly}$. This means, roughly, that there is no “brute-force” method to solve NP-complete problems in quantum polynomial time, even with the help of quantum advice: any proof that $\text{NP} \subset \text{BQP}/\text{qpoly}$ would have to use techniques radically unlike any we know today.

In my view, these results support the intuition that quantum states are “more like” probability distributions over n -bit strings than like exponentially-long strings to which one has random access. If exponentially-long strings were rocket fuel, and probability distributions were grape juice, then quantum states would be wine—the alcoholic “kick” in this analogy being the minus signs. I can imagine someone objecting: “What a load of nonsense! Whether quantum states are more like grape juice or rocket fuel is not a mathematical question, about which theorems could say anything!” To which I’d respond: if results such as $\text{BQP}/\text{qpoly} \subseteq \text{PostBQP}/\text{poly}$, which sharply limit the power of quantum advice, do not count as evidence against Goldreich’s view of quantum states, then what *would* count as evidence? And if nothing would count, then how scientifically meaningful is that view in the first place?

A meatier objection centers around a recent result of Raz [13], that a quantum interactive proof system where the verifier gets quantum advice can solve any problem whatsoever—or in complexity language, that QIP/qpoly equals ALL. (Here ALL is the class of all problems, which means, literally, the class of all problems.) I made a related observation in [2], where I pointed out that $\text{PostBQP}/\text{qpoly}$ equals ALL.⁷ However, a key point about both results is that *they have nothing to do with quantum computing*, and indeed, would work just as well with classical randomized advice. In other words, the classes IP/rpoly and PP/rpoly are also equal to ALL. Since I like making conjectures, I’ll conjecture more generally that quantum advice does not wreak much havoc in the complexity zoo that isn’t already wreaked by randomized advice. So for example, I’ll conjecture that just as the class MA/rpoly is strictly contained in ALL,⁸ so too its quantum analogue QMA/qpoly is strictly contained in ALL. I might be proven wrong, but that’s the whole point!

Conclusion

For almost a century, quantum mechanics was like a Kabbalistic secret that God revealed to Bohr, Bohr revealed to the physicists, and the physicists revealed (clearly) to no one. So long as the lasers and transistors worked, the rest of us shrugged at all the talk of complementarity and wave-particle duality, taking for granted that we’d never understand, or need to understand, what such things actually meant. But today—largely because of quantum computing—the Schrödinger’s cat is out of the bag, and all of us are being forced to confront the exponential Beast that lurks inside our current picture of the world. And as you’d expect, not everyone is happy about that, just as the physicists themselves weren’t all happy when *they* first had to confront it in the 1920’s.

Yet this unease has to contend with two traditions of technical results: the first showing that many of the obvious alternatives to quantum mechanics are nonstarters; and the second showing that quantum mechanics isn’t *quite* as strange as one would naïvely think. Both traditions are decades old: the first includes Bell’s theorem [5] and the Kochen-Specker theorem [9], while the second includes Holevo’s theorem [8] and the results of decoherence theory. But theoretical computer scientists come to quantum mechanics with their own set of assumptions (some would say prejudices), so in this abstract I’ve tried to indicate how they, too, might eventually be shoved into the vast quantum ocean, which isn’t *that* cold once one gets used to it.

⁶In [3] I characterized PostBQP exactly in terms of the classical complexity class PP (Probabilistic Polynomial-Time), which consists of all decision problems solvable in polynomial time by a randomized Turing machine, which accepts with probability greater than 1/2 if and only if the answer is “yes.” Thus, a more conventional way to state my result is $\text{BQP}/\text{qpoly} \subseteq \text{PP}/\text{poly}$.

⁷Here’s a one-sentence proof: given the advice state $2^{-n/2} \sum_x |x\rangle |f(x)\rangle$, to evaluate the Boolean function f on any given input x we simply need to measure in the standard basis, then postselect on seeing the $|x\rangle$ of interest.

⁸Indeed $\text{MA}/\text{rpoly} = \text{MA}/\text{poly}$; that is, we can replace the randomized advice by deterministic advice.

References

- [1] S. Aaronson. Multilinear formulas and skepticism of quantum computing. In *Proc. ACM STOC*, pages 118–127, 2004. quant-ph/0311039.
- [2] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005. quant-ph/0402095.
- [3] S. Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. Roy. Soc. London*, 2005. To appear. quant-ph/0412187.
- [4] M. Arndt, O. Nairz, J. Vos-Andreae, C. Keller, G. van der Zouw, and A. Zeilinger. Wave-particle duality of C_{60} molecules. *Nature*, 401:680–682, 1999.
- [5] J. S. Bell. *Speakable and Unspeakable in Quantum Mechanics*. Cambridge, 1987.
- [6] S. Ghosh, T. F. Rosenbaum, G. Aeppli, and S. N. Coppersmith. Entangled quantum state of magnetic dipoles. *Nature*, 425:48–51, 2003. cond-mat/0402456.
- [7] O. Goldreich. On quantum computing. www.wisdom.weizmann.ac.il/~oded/on-qc.html, 2004.
- [8] A. S. Holevo. Some estimates of the information transmitted by quantum communication channels. *Problems of Information Transmission*, 9:177–183, 1973. English translation.
- [9] S. Kochen and S. Specker. The problem of hidden variables in quantum mechanics. *J. of Math. and Mechanics*, 17:59–87, 1967.
- [10] L. A. Levin. Polynomial time and extravagant models, in The tale of one-way functions. *Problems of Information Transmission*, 39(1):92–103, 2003. cs.CR/0012023.
- [11] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [12] R. Raz. Multi-linear formulas for permanent and determinant are of super-polynomial size. In *Proc. ACM STOC*, pages 633–641, 2004. ECCC TR03-067.
- [13] R. Raz. Quantum information and the PCP theorem. In *Proc. IEEE FOCS*, 2005. To appear. quant-ph/0504075.