Research Statement

Scott Aaronson

January 6, 2007

Most of my research deals with two questions: first, what are the ultimate limits on what can feasibly be computed in the physical world? Second, how can studying those limits shed light on basic issues in physics and cosmology? The first question involves bringing physics into computational complexity theory; the second, bringing computational complexity theory into physics. *A priori*, one might wonder whether there is any useful bridge to be built between these two subjects, one sturdy enough to carry not just metaphors but nontrivial technical results. As it happens, such a bridge has existed for thirteen years.

Ever since I learned how to program, I had imagined that the physical world consists of "bits all the way down." It seemed obvious to me that, if only we could probe nature at the Planck scale, we would find nothing but a vast array of bits getting updated by simple local rules: Conway's Game of Life writ large. The specific form of the rules was of no great consequence, since according to the Extended Church-Turing Thesis, any "reasonable" set of rules can simulate any other with at most polynomial slowdown.

After Peter Shor discovered his factoring algorithm [50], I and others who thought similarly were faced with a choice. Either

(1) the Extended Church-Turing Thesis is false,

(2) quantum mechanics as conventionally understood is false, or

(3) the factoring problem is solvable in polynomial time on a classical computer.

All three of these possibilities seem like wild, crackpot speculations, but at least one of them is correct! As I never tire of pointing out, this is a dilemma that confronts all of us—whether we choose to work on quantum computing or a different field, whether we think that practical quantum computers will be built in 20 years, 2000 years, or never. Personally, I follow current experimental work in quantum computing with admiration and interest, and I try to contribute to it when I can. But the prospect of breaking the RSA cryptosystem, or simulating quark-gluon plasmas, is not what keeps me awake at night. I chose to work on quantum computing because I want to know what the world is *like*—and because a world that allows quantum computers would be fundamentally different from a world that doesn't.

To me, Shor's algorithm represented a promise: that from now on, the study of the feasibly computable was going to be inextricably linked to the central conceptual problems in physics. My work, over the last seven years, can be seen as an attempt to make good on that promise.

1 Previous Work

The following capsules highlight eight ways in which my research has forged connections between computational complexity theory and physics. I've tried to provide self-contained explanations of my results, for the benefit of readers who have only a passing familiarity with quantum computing. This is what resulted in my research statement becoming rather long.

1.1 The Limits of Quantum Computers

If quantum computers can factor integers in polynomial time, then what *can't* they do? Contrary to widespread misconception, today we have evidence that quantum computers would face many of the same limitations as their classical counterparts. In particular, we strongly believe there is no quantum algorithm to solve NP-complete problems in polynomial time. Of course we can't *prove* that NP-complete problems are hard for quantum computers, since we can't even prove they're hard for classical computers! But in the early days of quantum computing, Bennett, Bernstein, Brassard, and Vazirani [31] showed the following fundamental result. Suppose we throw away the structure of an NP-complete problem, and consider only a "landscape" of 2^n possible solutions, together with a black box that tells us whether a given solution is correct or not. Then even if we can query the black box in quantum superposition, we still need $\Omega\left(2^{n/2}\right)$ queries to find a solution. This bound is tight, as evidenced by Grover's algorithm [36]. In other words, at least for "unstructured" search problems, the quadratic speedup of Grover's algorithm is optimal.

Bennett et al.'s result was only the beginning of a long, highly-successful quest to understand the black-box limitations of quantum computers. Among my contributions to that quest, the most important was a quantum lower bound for the so-called *collision problem* [3]. In this problem, we are given black-box access to a two-to-one function $f : [N] \rightarrow [N]$. Our goal is to find a *collision*—that is, a distinct x, y pair such that $f(x) = f(y)$—using as few queries to f as possible. It's not hard to see that, if we could find a collision after $O(\operatorname{polylog} N)$ queries, then we could efficiently break "collision-resistant hash functions" such as MD5, and thereby undermine much of modern cryptography. We could also efficiently solve the Graph Isomorphism problem.

Classically, finding a collision takes $\Theta\left(\sqrt{N}\right)$ queries, by the famous "Birthday Paradox." Brassard, Høyer, and Tapp [33] gave a quantum algorithm that does slightly better, finding a collision after $O\left(N^{1/3}\right)$ queries. The obvious question was whether or not this was optimal. Surprisingly, this question resisted attack for five years, with no lower bound better than constant (!) known. Finally, in 2001, I showed [3] that any quantum algorithm for finding a collision must make $\Omega\left(N^{1/5}\right)$ queries. Subsequently Shi [49], Kutin [41], and Ambainis [24] improved my lower bound to the optimal $\Omega\left(N^{1/3}\right)$.

What made proving a lower bound so difficult was the collision problem's "global" nature—the fact that we're no longer looking for a needle in a haystack, but merely for two pieces of hay of the same length! Another way of putting the point is that a quantum computer can "almost" find a collision after just *one* query to f. For we can easily prepare a state of the form $\frac{1}{\sqrt{2}}(|x\rangle + |y\rangle)$, for some random x, y pair such that $f(x) = f(y)$. The only problem is that, if we measure this state, then we'll obtain either $|x\rangle$ *or* $|y\rangle$ but not both of them.

The collision lower bound is one of the main pieces of evidence that secure electronic commerce will still be possible in a world with quantum computers. Yet my own motivation for working on the problem came not from cryptography, but rather from a desire to understand the computational power of "hidden-variable theories" such as Bohmian mechanics. In these theories, one supplements quantum mechanics by a rule for calculating the "actual" trajectories of particles, in such a way that the measurement probabilities agree precisely with the standard quantum-mechanical ones. I was interested in the following question: *how hard is it to calculate the entire trajectory of a particle?* Or conversely, supposing that (contrary to experience) we could observe a particle's entire trajectory, what computational powers would that give us? I proved in [12] that, in *any* hidden-variable theory satisfying two reasonable axioms, sampling an entire trajectory would give us the ability to solve the collision problem in a constant number of queries. Combining this result with the collision lower bound, we obtain strong evidence that *sampling trajectories is hard for quantum computers*—or in other words, that "hidden-variable quantum computers" would be strictly more powerful than ordinary quantum computers. On the other hand, I also showed in [12] that, in the black-box or oracle setting, observing a trajectory would *still* not let us solve NP-complete problems in polynomial time. I therefore obtained the first reasonable model of computation that seems more powerful than quantum computing, but only *slightly* so.

In addition to my work on the collision problem, I also proved the first "direct product theorem" for quantum search [10], as well as quantum lower bounds for finding local minima [15], evaluating total Boolean functions [6], and Recursive Fourier Sampling [7]. These other limitations of quantum computers will be discussed later in the research statement where relevant.

1.2 Addressing Skepticism of Quantum Computing

Several prominent computer scientists and physicists—including Oded Goldreich, Leonid Levin, Gerard 't Hooft, and Stephen Wolfram—have argued that building large-scale quantum computers will be not merely difficult, but fundamentally impossible. The suggested reasons vary, but they often center around the idea that describing a state of, say, $10,000$ particles by a vector of $2^{10,000}$ amplitudes is "inherently extravagant"—a sign that quantum mechanics is being pushed beyond its domain of validity.

The obvious response is twofold: first, that I certainly *hope* we encounter such a breakdown of quantum mechanics, since that would be much more scientifically interesting than mere success in building a quantum computer! And second, that if quantum computing skeptics want to advance research, then they ought to propose specific, testable *ways* in which quantum mechanics might break down.

In a paper on "Multilinear Formulas and Skepticism of Quantum Computing" [8], I tried to go beyond these generalities, and actually *start* the research program that I advocate for skeptics. To do so I introduced the notion of a "Sure/Shor separator," which is a set of quantum states that includes all states that experimentalists have already demonstrated, but *not* the states that would arise in, say, Shor's factoring algorithm. I explained why the obvious ways of measuring a state's complexity fail to produce a Sure/Shor separator. I then proposed a new measure that I called *tree size*: basically, the minimum number of vertices needed to express a state by a "tree" of linear combinations and tensor products over the basis $\{|0\rangle, |1\rangle\}$. Most of the states that physicists tend to discuss—such as product states, "Schrödinger cats," and 1-dimensional spin chains—have tree size polynomial in the number of qubits n. On the other hand, by using a spectacular recent result of Raz [43], I was able to prove that certain states arising in quantum error-correction would have tree size $n^{\Omega(\log n)}$. Later, I was pleased to learn that experimental work on 2- and 3-dimensional spin lattices [35] might have already succeeded in demonstrating states with superpolynomial tree size.

1.3 The Power of Quantum Proofs and Quantum Advice

Even if they can't propose a Sure/Shor separator, many computer scientists are profoundly uneasy about the apparent "exponentiality" of quantum states. One way of addressing their unease is to ask: in what settings does the exponentiality *actually* manifest itself? So for example, are there theorems that admit a short quantum proof but not a short classical proof? Are there problems that can be solved with the help of short quantum advice, but not short classical advice? Four of my results are directly relevant to these questions.

A few years ago I studied the complexity class $\mathsf{BQP/qpoly}$, which consists of all problems solvable in polynomial time on a quantum computer, with the help of a polynomial-size "quantum advice state" $|\psi_n\rangle$ that depends only on the input length n but can otherwise be arbitrary [10]. *A priori*, it's not obvious that a quantum computer couldn't solve any problem whatsoever, were it only equipped with the right advice state! However, I showed that $\mathsf{BQP/qpoly}$ is contained in the classical class $\mathsf{PP/poly}$. Intuitively, this means that 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 to extract what the advice is telling you.

In the same paper, I gave another limitation of quantum advice, by constructing an oracle relative to which $\mathsf{NP} \not\subset \mathsf{BQP/qpoly}$. This provided the first evidence that NP-complete problems are still hard for quantum computers even in the presence of quantum advice. The oracle result was actually just a corollary of a much stronger result: my so-called *direct product theorem for quantum search*. The direct product theorem states that if a quantum algorithm is searching a list of N items, K of which are marked, and if the algorithm doesn't have enough time to find even *one* marked item, then the probability that the algorithm will find all K of the marked items decreases exponentially with K. Even though this theorem seems intuitively obvious, proving it had been a notorious open problem for years, since one needs to rule out the possibility of subtle correlations between finding one marked item and finding another one. Shortly after

I obtained the direct product theorem, Klauck, Špalek, and de Wolf [39] improved it, and also applied it to prove the first *quantum time-space tradeoffs*: that is, tradeoffs between the number of steps used by a quantum algorithm and the number of qubits.

In recent work with Greg Kuperberg [20], we switched attention from quantum advice to quantum *proofs*. In particular, we studied the longstanding open problem of whether there are mathematical truths that you can prove to someone by handing them a quantum state, but *not* by handing them a classical string of comparable size. More formally, the question is whether $\mathsf{QMA} = \mathsf{QCMA}$, where QMA (Quantum Merlin Arthur) is the class of problems for which every 'yes' answer has a polynomial-size quantum proof that can be verified in quantum polynomial time, and QCMA (Quantum Classical Merlin Arthur) is the same except that now the proof has to be classical. Kuperberg and I gave a novel sort of evidence that quantum proofs really are more powerful than classical ones, by constructing a "quantum oracle" relative to which $\mathsf{QMA} \neq \mathsf{QCMA}$. On the other hand, we also showed that constructing a *classical* oracle relative to which $\mathsf{QMA} \neq \mathsf{QCMA}$ is likely to be much harder than previously supposed—since the central group-theoretic problems that were known to be in QMA are actually in QCMA as well.

Finally, last year I studied the power of quantum proofs *combined* with quantum advice [17]. I showed that, even when taken together, these resources would not make a quantum computer infinitely powerful. This contrasts with a counterintuitive recent result of Raz [44]: that *interactive* proofs combined with quantum advice *would* make a quantum computer infinitely powerful. More formally, I showed that the class $\mathsf{QMA/qpoly}$ is contained in $\mathsf{PSPACE/poly}$, whereas Raz showed that the class $\mathsf{QIP/qpoly}$ contains *every* computational problem.

1.4 The Learnability of Quantum States

One of the most basic tasks in current experimental physics is called *quantum state tomography*. Here we are given some physical process that reliably produces a quantum mixed state ρ; the goal is to learn an approximate description of ρ by repeatedly applying the process and then measuring the result. Quantum state tomography has been used to study chemical reactions [51] and to confirm the preparation of entangled states [37], among other applications. But there is a fundamental problem, which has thus far prevented tomography from being applied to states of more than about 8 qubits. This is that, if ρ is an n-qubit state, then the number of measurements needed to reconstruct it even approximately grows exponentially with n: in particular like 4^n, the number of independent parameters in a $2^n \times 2^n$ density matrix.

Recently I showed how to get around this problem, provided we are willing to relax the goal to what I call "pretty-good tomography" [14]. More formally, suppose we can repeatedly prepare an n-qubit state ρ, and then apply a measurement drawn from some (possibly-unknown) probability distribution $\mathcal{D}$. Suppose also that we only want to predict the approximate expectation values of *most* measurements drawn from $\mathcal{D}$, rather than all of them. I showed that we can do this, with high probability, using a number of sample measurements that increases only *linearly* with the number of qubits n, and inverse-polynomially with the relevant error parameters. The proof of this theorem relates quantum information theory to computational learning theory, and in particular to sample complexity bounds for "PAC" (Probably Approximately Correct) learning. I have already spoken to some experimental groups about doing a small-scale demonstration of pretty-good tomography, and I hope these efforts will bear fruit. What I've already been able to do is to use my learning theorem to prove several new results in quantum computing—including a classical simulation of quantum one-way communication protocols, and an approximate verification procedure for untrusted quantum advice.

1.5 Quantum Computing and Spacetime

A common criticism of Grover's search algorithm is that, while it might yield a quadratic speedup for solving, say, NP-complete optimization problems, it can't yield a speedup for searching a "physical" database. Benioff [30] boiled this criticism down to what I see as its essence: that since the speed of light is finite, any quantum search algorithm will necessarily be limited by the time needed for signals to propagate from one part of the database to another. In joint work with Andris Ambainis [18], we showed that this criticism fails for a nontrivial reason. By using a carefully-optimized recursive version of Grover search, Ambainis and I proved that a "quantum robot," moving at unit speed, could search a two-dimensional lattice of size $\sqrt{N} \times \sqrt{N}$ for a marked item using only $O\left(\sqrt{N}\log^{3/2} N\right)$ steps. Subsequently, Ambainis, Kempe, and Rivosh [25] showed how to get the running time down to $O\left(\sqrt{N}\log N\right)$ using quantum random walks. In hypercubes of dimension 3 or higher, the running time decreases further to $O\left(\sqrt{N}\right)$, which precisely matches the time needed for Grover search with no locality constraints at all.

Using our spatial search algorithm, Ambainis and I were able to give an $O\left(\sqrt{N}\right)$-qubit quantum communication protocol for the so-called *disjointness problem*, thereby matching a lower bound of Razborov [45] and solving an annoying open problem. For me, though, the real motivation for this work was that it connected quantum computing to spacetime geometry—and in particular to the "holographic principle," which is a universal upper bound on the number of bits that can be stored in a given region of spacetime. As an example of this connection, Ambainis and I were able to answer the following question: in a universe with positive cosmological constant (like the one we seem to inhabit), and assuming only quantum mechanics and the holographic principle, how large a database could a "quantum robot" ever search for a specific entry, before most of the database receded past the robot's cosmological horizon?

1.6 Beyond Quantum Computing

Even though all the ingredients were in place by the 1960's, the field of quantum computing didn't take off until the mid-1990's. This raises an obvious question: *what else* about physics might computer scientists have overlooked in studying efficient computation? As an example, could quantum field theory or quantum gravity lead to a yet-more-powerful model of computation? Two years ago I wrote a survey paper about these issues entitled "NP-complete Problems and Physical Reality" [11]. My basic thesis was that the intractability of NP-complete problems might eventually come to be seen as a principle of *physics*, analogous to the Second Law of Thermodynamics or the impossibility of superluminal communication.

Motivated by this perspective, I've long been interested in the computational effects of various changes to the known laws of physics. I already mentioned one example in Section 1.1, my work on the computational power of hidden-variable theories [12]. Let me now mention two other examples.

In a paper on "Quantum Computing, Postselection, and Probabilistic Polynomial-Time" [13], I studied the power of quantum computers with *postselection*: that is, the ability to measure a qubit and *assume* the outcome will be $|1\rangle$ (or equivalently, discard all runs of the computation where the outcome is $|0\rangle$). Postselection arises frequently in discussions about the many-worlds interpretation, the anthropic principle, and hypothetical nonlinearities in the Schrödinger equation. Surprisingly, I showed that quantum computers with postselection have exactly the power of the well-studied classical complexity class PP: that is, PostBQP = PP.

The second example concerns some recent work with John Watrous [21]. Watrous and I studied the power of quantum computers in the presence of *closed timelike curves* (CTC's). Basically, a CTC is a region of spacetime where we impose the condition that the initial and final states are the same, or $\rho_{\text{initial}} = S\left(\rho_{\text{initial}}\right) = \rho_{\text{final}}$ where S is the quantum operator acting along the CTC. Watrous and I showed that, in the presence of CTC's, *quantum computers are no more powerful than classical computers*: both of them have exactly the power of the class PSPACE.

1.7 Classical Results from Quantum Arguments

Just as mathematicians use complex numbers to prove theorems about the real numbers, so we know today that one can use ideas from quantum computing to prove new results about *classical* computing (as well as to provide simpler proofs of known results). Two of my best-known results played a role in bringing this phenomenon to wider attention.

The first concerns the problem of *local search*. Here we're given an undirected graph $G = (V, E)$, as well as black-box access to a function $f : V \to \mathbb{Z}$. Our goal is to find a *local minimum* of f: that is, a vertex

$v \in V$ such that $f(v) \leq f(w)$ for all edges $(v,w) \in E$. Aldous [22] showed in 1983 that, if G is the Boolean hypercube $\{0,1\}^n$, then any randomized algorithm needs to query the black box $2^{n/2-o(n)}$ times to find a local minimum of f. Using Ambainis's quantum adversary method [23], I showed in [15] that any quantum algorithm needs to query the black box $\Omega\left(2^{n/4}/n\right)$ times. The surprising part was that a small tweak to the quantum argument yielded an improvement of Aldous's *classical* lower bound: from $2^{n/2-o(n)}$ to the nearly-optimal $\Omega\left(2^{n/2}/n^2\right)$. Using a quantum argument, I also obtained the first classical lower bounds for local search on cubes of constant dimension. These results were subsequently improved and extended by Santha and Szegedy [48], Zhang [55], and Sun and Yao [52].

The second example of a classical harvest from quantum techniques concerns my result (mentioned in Section 1.6) that postselected quantum computation has exactly the power of PP. A year after proving this result, I realized that it yields a half-page proof of the celebrated Beigel-Reingold-Spielman Theorem [29] from classical complexity theory: that PP is closed under intersection. Proving PP closed under intersection had been a well-known open problem since 1972, until Beigel et al. settled it using fairly heavy machinery in 1991. My proof finally lets us pinpoint what went wrong with previous attempts to give a simple proof of the theorem: they were missing quantum mechanics!

1.8 Circuit Lower Bounds

It's a *bit* premature to tackle the P versus NP question. But like many complexity theorists, I've long been interested both in the logical status of the P versus NP question [5], and in developing the sorts of circuit lower bound techniques that would eventually be needed to address it. Recently my interest in circuit lower bounds merged with my interest in quantum computing in at least two unexpected ways.

First, I was able to give one of the first examples of a provably-nonrelativizing circuit lower bound [16]. Vinodchandran [54] had shown earlier that for every positive integer k, there exists a language in the complexity class PP that does not have circuits of size n^k. I generalized Vinodchandran's result from classical to quantum circuits. More importantly, using the same "polynomial method" that I had previously used to prove quantum lower bounds, I gave an oracle relative to which PP has *linear*-size circuits. This implies that Vinodchandran's lower bound was nonrelativizing: in other words, it was necessarily "sensitive" to the fact that no oracle was present. Since the work of Baker, Gill, and Solovay in the mid-1970's [27], complexity theorists have known that nonrelativizing lower bounds will be needed to attack the P versus NP question, but have had almost no examples of such bounds.

The second example grew out of my work on Sure/Shor separators [8], discussed in Section 1.2. Razborov and Rudich [46] famously argued that any proof of $\mathsf{P} \neq \mathsf{NP}$ (or more precisely, of $\mathsf{NP} \not\subset \mathsf{P/poly}$) would need to be "non-naturalizing." Loosely speaking, this means that such a proof would need to zero in on a specific property of (say) the SAT function, and *not* a property that the SAT function has in common with random Boolean functions. Unfortunately, we know of almost no non-naturalizing lower bound techniques for any reasonable model of computation. As part of my investigation into the tree size of quantum states [8], I showed that the so-called "manifestly orthogonal tree size" (MOTS) of a random coset state over $\mathbb{Z}_2^n$ is exponential in n. To do so, I used an argument that relies essentially on the coset structure of such states, and that therefore appears not to naturalize in the Razborov-Rudich sense.

2 Future Work

Obviously, it's vastly easier to summarize what I've already done than to say what I *will* do—since if I knew, I presumably would've done it already! And indeed, rereading my graduate-school applications gives me very little confidence in my ability to predict which problems I'm going to tackle next.

In the short term, I've been thinking about copy-protection of quantum software. In the classical world, copy-protecting software is trivially impossible—not that that's stopped people from trying! But what if your program is a quantum state? The key question is this: given a Boolean function f, is there a state $|\psi_f\rangle$ that lets you efficiently evaluate f, but that *doesn't* let you efficiently prepare more states with which to evaluate f? So far, I've been able to prove three results about this question. First, I can give a nontrivial class of programs that *can* be quantumly copy-protected, under a new computational assumption that's related to (but stronger than) the hardness of the Nonabelian Hidden Subgroup Problem. Second, I can give a nontrivial class of programs that *can't* be quantumly copy-protected. Finally, I can give a "quantum oracle" relative to which *all* programs can be quantumly copy-protected, except the ones that are like my second example. The last result uses two components that might be of independent interest: an explicit construction of "pseudorandom quantum states"; and a common generalization of the No-Cloning Theorem and the quantum search lower bound.

I also hope to go further with my work on the learnability of quantum states [14], discussed in Section 1.4. In principle, my pretty-good tomography approach could extend the experimental frontier in any area of physics that uses quantum state tomography. However, there are at least two things that will need to be done to convince experimentalists to try the new approach. First, my quantum-state reconstruction algorithm will need to be implemented, and its performance characterized empirically. Second, special classes of quantum states will need to be identified that can be learned not only with a linear number of measurements, but also with a reasonable amount of computation.

2.1 The Back Burner

So much for the immediate future. Over the longer term, many fundamental open questions in quantum complexity theory remain on my back burner. Let me now describe four examples of such questions, together with the progress on them that I've been able to make so far.

(1) Is BQP contained in the polynomial-time hierarchy? Alternatively, can we at least give an oracle relative to which $\mathsf{BQP} \not\subset \mathsf{PH}$? This has remained perhaps the most embarrassing open problem in all of quantum complexity theory, ever since Bernstein and Vazirani [32] posed it in 1993.

In [7], I undertook a detailed study of Recursive Fourier Sampling (RFS), which is almost the only candidate problem we have for proving an oracle separation between BQP and PH. I showed that the known quantum algorithm for RFS was optimal, which ruled out the possibility of using RFS to place BQP outside of PH with a non-constant number of alternations. More recently, Greg Kuperberg and I [20] proposed a new approach to the problem: one that wouldn't be based on RFS at all, but rather on "dequantizing" quantum oracles. At this point, however, we don't even have an oracle relative to which $\mathsf{BQP} \not\subset \mathsf{AM}$.

(2) Can quantum computers ever obtain more than the quadratic speedup of Grover's algorithm for "unstructured" problems? (Note that Shor's algorithm solves a highly structured problem—namely period-finding—and therefore doesn't count.) There are actually at least three ways to make this question precise, and all of them are open. First, does $\mathsf{BPP} = \mathsf{BQP}$ relative to a random oracle with probability 1? Second, is there a total Boolean function f for which $\mathrm{D}(f) = \omega\left(\mathrm{Q}(f)^2\right)$, where $\mathrm{D}(f)$ is the deterministic query complexity of f, and $\mathrm{Q}(f)$ is the quantum query complexity? Beals et al. [28] showed that $\mathrm{D}(f) = O\left(\mathrm{Q}(f)^6\right)$ for all total f, but improving their exponent of 6 has remained open for a decade—see my paper [6] for a possible approach to the question. Third, is there any problem invariant under permutation symmetry (like the collision and element distinctness problems), for which quantum computing gives more than a quadratic advantage? Currently, we can't even rule out an *exponential* advantage for such functions; to do so would require a far-reaching generalization of my collision lower bound.

(3) What is the power of quantum one-way communication? Here we consider two players Alice and Bob, who have inputs x and y respectively. We then ask how many qubits Alice needs to transmit to Bob, for Bob to evaluate a known Boolean function $f(x, y)$. The central open question is this: is there a total Boolean function f for which $\mathrm{Q}^1(f) = o\left(\mathrm{R}^1(f)\right)$, where $\mathrm{Q}^1(f)$ is the quantum one-way communication complexity of f and $\mathrm{R}^1(f)$ is the randomized one-way communication complexity? For partial Boolean functions, Gavinsky et al. [34] recently achieved an *exponential* separation between

$\mathrm{R}^1(f)$ and $\mathrm{Q}^1(f)$, and I showed [10, 14] that their result is in some sense optimal. For total functions, by contrast, I'd conjecture that $\mathrm{R}^1(f)$ and $\mathrm{Q}^1(f)$ are asymptotically equivalent. In [10], I managed to prove this conjecture in some interesting special cases: for example, the case where Alice is given $\langle x, y\rangle \in \mathbb{F}_p^2$ and Bob is given $\langle a, b\rangle \in \mathbb{F}_p^2$, and the problem is to decide whether $y \equiv ax + b \pmod p$. But the general case remains wide open.

(4) What is the power of quantum proofs with multiple provers? Formally, Kobayashi et al. [40] defined the complexity class $\mathsf{QMA}(2)$, which is the same as QMA (Quantum Merlin-Arthur) except that now the verifier receives *two* quantum proofs that are guaranteed to be unentangled with each other. In the classical case, receiving two proofs instead of one would make no difference, but in the quantum case, we know essentially nothing about how much difference it makes. We don't have an oracle separation (even a quantum oracle separation) between $\mathsf{QMA}(2)$ and QMA. We don't know if $\mathsf{QMA}(2)$ is contained in EXP. We don't even know if $\mathsf{QMA}(2)$ proofs can be amplified to exponentially small error probability. A recent observation of mine begins to explain why these problems are so difficult. I showed that if $\mathsf{NP} \subseteq \mathsf{SPACE}(\mathrm{polylog}(n))$ in the unrelativized world, then $\mathsf{QMA}(2)^U \subseteq \mathsf{BQPSPACE}^U$ for all quantum oracles U—and hence, we can't hope for (say) a quantum oracle relative to which $\mathsf{QMA}(2) \not\subset \mathsf{BQPSPACE}$.

These are all problems that I've worked on in the past, that have repaid me with partial results, that I still consider to be of fundamental importance, and that I intend to return to as soon as I see a promising new line of attack.

2.2 New Directions

At no point in my research career have I been exclusively interested in quantum computing. As an example, two years ago I undertook a detailed study of the complexity of Bayesian agreement protocols [9]. My main result was the following: suppose two Bayesian agents Alice and Bob share a common prior but have different information, and let X be any $[0, 1]$ random variable that they both want to estimate (for example, the probability of rain tomorrow). Then there exists a communication protocol that exchanges only $O\left(\frac{1}{\delta\varepsilon^2}\right)$ bits of information, and that causes Alice and Bob's expectations of X to agree within ε with probability at least $1 - \delta$ over their shared prior. Furthermore, the computations needed for this protocol can in a certain sense be performed efficiently. This result is basically the complexity-theoretic version of Aumann's celebrated Agreement Theorem [26]: that Bayesian agents with common priors can never "agree to disagree." My result shows that agreeing to disagree is problematic not merely "in the limit" of common knowledge, but even for agents subject to realistic constraints on computation and communication.

Not only has my research strayed often from my core interest of quantum complexity theory, I have not even been averse to writing code (!) when the need arose. Prior to graduate school, I wrote three papers with large implementation components: one on algorithms for Boolean function query complexity measures [4], one on clustering documents by author [2], and one on optimizing link layout in hypertext systems [1]. I began my graduate studies focusing on AI rather than theory, and I still maintain a strong interest in complexity-theoretic questions arising from more applied areas such as machine learning, automated reasoning, and bounded rationality.

Within quantum computing, I've worked on many problems of a less foundational nature than the ones described in Section 1. For example, in joint work with Daniel Gottesman [19], we gave a new, more efficient classical algorithm for simulating *stabilizer circuits*: quantum circuits that consist entirely of CNOT, Hadamard, and $\frac{\pi}{4}$-phase gates. We also showed that the problem of simulating stabilizer circuits is in a subclass of P called $\oplus\mathsf{L}$, and hence is presumably not even universal for *classical* computation. More to the point, we actually implemented our simulation algorithm, and showed that it makes practical the simulation of stabilizer circuits on up to about 3000 qubits (the main limitation being available memory). I'm pleased that our simulation software is now being used to help design fault-tolerant quantum architectures (see Metodi et al. [42] for example).

I'll now try to give a sense of some new directions that my research *might* take. This list is by no means exhaustive.

(1) As alluded to in Section 1.5, supernova observations indicate that we live in a universe with a positive cosmological constant $\Lambda > 0$ [47]. Combined with the holographic principle, this suggests that no computation carried out in the physical world can ever involve more than $1/\Lambda \approx 10^{122}$ bits. I want to think more about the computational implications of a positive cosmological constant. To this end, I've defined a formal model of "inflationary Turing machines," which models computation in a d-dimensional spacetime that is expanding exponentially. I've formulated several open questions in this model: for example, for a fixed Λ, are d-dimensional inflationary Turing machines more powerful than $(d-1)$-dimensional machines? I haven't yet been able to answer these questions.

(2) Now that quantum computing has left the Extended Church-Turing Thesis hanging by a fingernail, it's natural to ask whether the *original* Church-Turing Thesis—the one about computability theory—is also open to attack. While I don't believe that any of the existing attacks on the Church-Turing Thesis have been successful, years ago I formulated a program for studying this question in a new way. My program begins with the following question: *is there any sensible model of computation that can solve its own halting problem?* To study this question, I look abstractly at the set of possible program behaviors—for example, $\{\text{Accept}, \text{Reject}, \text{LoopForever}\}$—and more specifically, the way in which those behaviors can depend on the behaviors of another program that is given as input. The problem of characterizing the possible "theories of computation" then boils down to one of characterizing the *monoids* and *clones* that do not contain any fixed-point-free elements. At this point I set the problem aside.

(3) Another one of my goals is to develop a theory of *procrastinating algorithms*: algorithms that, whenever possible, try to put off all the actual work until later in the computation. I can show that, in a certain precise sense, the n-bit XOR function can be evaluated by a procrastinating algorithm whereas the n-bit OR function cannot. However, it remains an open problem to characterize the set of all Boolean functions that can be evaluated by procrastinating algorithms. I intend to work on this problem as soon as I have the time.

(4) I've long been interested in applying tools from computational complexity theory to understand Darwinism. One way of doing this would start from Valiant's beautiful notion of 'evolvability' [53]. A different way would start from a speculation of Richard Dawkins: that natural selection is basically the *only* mechanism for producing adaptive complexity; or in other words, that if we ever encounter extraterrestrial life, then we should assume that it ultimately arose by some trial-and-error process like natural selection. Does Dawkins's conjecture have any complexity-theoretic counterpart? The following is one attempt to provide one:

    Given any randomized algorithm for solving a PLS*-complete local optimization problem, it is possible to modify that algorithm, adding only a small overhead in complexity, so that as it proceeds the algorithm outputs an "evolutionary path" that starts with an arbitrary solution and ends with a locally optimal one.*

    Here PLS, or Polynomial Local Search, is the class of local optimization problems defined by Johnson, Papadimitriou, and Yannakakis [38]. Note that there is no hope of proving such a conjecture unconditionally, since any reasonable formalization of it would imply $\mathsf{P} \neq \mathsf{NP}$. However, for a certain reasonable formalization, I believe I can prove that the above conjecture *does* hold relative to an oracle, by extending my query complexity lower bound for local search [15].

Whether I continue attacking the core questions of quantum complexity theory, branch out into other areas, or (most likely) do some of both, there are two properties that I think will characterize my research for the foreseeable future. First, any research project I undertake will use ideas from theoretical computer science in a nontrivial way. Second, whether the inspiration comes from physics, cosmology, or even economics or biology, any problem that captures my interest will have a philosophical *point*: some conceptual issue at stake beyond the difference between $O\left(n^{3/2}\right)$ and $O\left(n^{4/3}\right)$. I will feel, in other words, like the universe would be a different place if the conjectured result were true than if it weren't. I can't say with certainty where the problems that produce this feeling will come from in the future, but seeking them out, and then trying to solve them, has been a viable strategy for me in the past.

References

[1] S. Aaronson. Optimal demand-oriented topology for hypertext systems. In *Proceedings of ACM SIGIR*, pages 168–177, 1997.

[2] S. Aaronson. Stylometric clustering: a comparison of data-driven and syntactic features. www.scottaaronson.com/papers/sc.doc, 2001.

[3] S. Aaronson. Quantum lower bound for the collision problem. In *Proc. ACM STOC*, pages 635–642, 2002. quant-ph/0111102.

[4] S. Aaronson. Algorithms for Boolean function query properties. *SIAM J. Comput.*, 32(5):1140–1157, 2003.

[5] S. Aaronson. Is P versus NP formally independent? *Bulletin of the European Assocation for Theoretical Computer Science*, (81), October 2003.

[6] S. Aaronson. Quantum certificate complexity. In *Proc. IEEE Conference on Computational Complexity*, pages 171–178, 2003. ECCC TR03-005, quant-ph/0210020.

[7] S. Aaronson. Quantum lower bound for recursive Fourier sampling. *Quantum Information and Computation*, 3(2):165–174, 2003. ECCC TR02-072, quant-ph/0209060.

[8] S. Aaronson. Multilinear formulas and skepticism of quantum computing. In *Proc. ACM STOC*, pages 118–127, 2004. Journal version to appear in SICOMP. quant-ph/0311039.

[9] S. Aaronson. The complexity of agreement. In *Proc. ACM STOC*, pages 634–643, 2005. ECCC TR04-061.

[10] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005. quant-ph/0402095. Conference version in Proceedings of CCC'2004.

[11] S. Aaronson. NP-complete problems and physical reality. *SIGACT News*, March 2005. quant-ph/0502072.

[12] S. Aaronson. Quantum computing and hidden variables. *Phys. Rev. A*, 71(032325), 2005. quant-ph/0408035 and quant-ph/0408119.

[13] S. Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. Roy. Soc. London*, A461(2063):3473–3482, 2005. quant-ph/0412187.

[14] S. Aaronson. The learnability of quantum states. quant-ph/0608142, 2006.

[15] S. Aaronson. Lower bounds for local search by quantum arguments. *SIAM J. Comput.*, 35(4):804–824, 2006. ECCC TR03-057, quant-ph/0307149. Conference version in Proceedings of ACM STOC'2004.

[16] S. Aaronson. Oracles are subtle but not malicious. In *Proc. IEEE Conference on Computational Complexity*, pages 340–354, 2006. ECCC TR05-040.

[17] S. Aaronson. QMA/qpoly is contained in PSPACE/poly: de-Merlinizing quantum protocols. In *Proc. IEEE Conference on Computational Complexity*, pages 261–273, 2006. quant-ph/0510230.

[18] S. Aaronson and A. Ambainis. Quantum search of spatial regions. *Theory of Computing*, 1:47–79, 2005. quant-ph/0303041.

[19] S. Aaronson and D. Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70(052328), 2004. quant-ph/0406196.

[20] S. Aaronson and G. Kuperberg. Quantum versus classical proofs and advice. To appear in Theory of Computing. quant-ph/0604056, 2006.

[21] S. Aaronson and J. Watrous. Closed timelike curves make classical and quantum computing equivalent. In preparation, 2006.

[22] D. Aldous. Minimization algorithms and random walk on the d-cube. *Annals of Probability*, 11(2):403–413, 1983.

[23] A. Ambainis. Quantum lower bounds by quantum arguments. *J. Comput. Sys. Sci.*, 64:750–767, 2002. Earlier version in ACM STOC 2000. quant-ph/0002066.

[24] A. Ambainis. Polynomial degree and lower bounds in quantum complexity: collision and element distinctness with small range. *Theory of Computing*, 1:37–46, 2005. quant-ph/0305179.

[25] A. Ambainis, J. Kempe, and A. Rivosh. Coins make quantum walks faster. In *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2005. To appear. quant-ph/0402107.

[26] R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

[27] T. Baker, J. Gill, and R. Solovay. Relativizations of the P=?NP question. *SIAM J. Comput.*, 4:431–442, 1975.

[28] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. Earlier version in IEEE FOCS 1998, pp. 352-361. quant-ph/9802049.

[29] R. Beigel, N. Reingold, and D. Spielman. PP is closed under intersection. *J. Comput. Sys. Sci.*, 50(2):191–202, 1995.

[30] P. Benioff. Space searches with a quantum robot. In S. J. Lomonaco and H. E. Brandt, editors, *Quantum Computation and Information*, Contemporary Mathematics Series. AMS, 2002. quant-ph/0003006.

[31] C. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. quant-ph/9701001.

[32] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997. First appeared in ACM STOC 1993.

[33] G. Brassard, P. Høyer, and A. Tapp. Quantum algorithm for the collision problem. *ACM SIGACT News*, 28:14–19, 1997. quant-ph/9705002.

[34] D. Gavinsky, J. Kempe, I. Kerenidis, R. Raz, and R. de Wolf. Exponential separations for one-way quantum communication complexity, with applications to cryptography. quant-ph/0611209, 2006.

[35] 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.

[36] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. ACM STOC*, pages 212–219, 1996. quant-ph/9605043.

[37] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt. Scalable multiparticle entanglement of trapped ions. *Nature*, 438:643–646, 2005. quant-ph/0603217.

[38] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Comput. Sys. Sci.*, 37:79–100, 1988.

[39] H. Klauck, R. Špalek, and R. de Wolf. Quantum and classical strong direct product theorems and optimal time-space tradeoffs. In *Proc. IEEE FOCS*, pages 12–21, 2004. quant-ph/0402123.

[40] H. Kobayashi, K. Matsumoto, and T. Yamakami. Quantum Merlin-Arthur proof systems: are multiple Merlins more helpful to Arthur? In *ISAAC*, pages 189–198, 2003. quant-ph/0306051.

[41] S. Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1:29–36, 2005. quant-ph/0304162.

[42] T. S. Metodi, D. D. Thaker, A. W. Cross, F. T. Chong, and I. L. Chuang. A quantum logic array microarchitecture: scalable quantum data movement and computation. In *38th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 305–318, 2005. quant-ph/0509051.

[43] 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.

[44] R. Raz. Quantum information and the PCP theorem. In *Proc. IEEE FOCS*, 2005. quant-ph/0504075.

[45] A. A. Razborov. Quantum communication complexity of symmetric predicates. *Izvestiya Math. (English version)*, 67(1):145–159, 2003. quant-ph/0204025.

[46] A. A. Razborov and S. Rudich. Natural proofs. *J. Comput. Sys. Sci.*, 55(1):24–35, 1997.

[47] S. Perlmutter and 31 others (Supernova Cosmology Project). Measurements of Ω and Λ from 42 high-redshift supernovae. *Astrophysical Journal*, 517(2):565–586, 1999. astro-ph/9812133.

[48] M. Santha and M. Szegedy. Quantum and classical query complexities of local search are polynomially related. In *Proc. ACM STOC*, pages 494–501, 2004.

[49] Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. In *Proc. IEEE FOCS*, pages 513–519, 2002. quant-ph/0112086.

[50] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997. Earlier version in IEEE FOCS 1994. quant-ph/9508027.

[51] E. Skovsen, H. Stapelfeldt, S. Juhl, and K. Mølmer. Quantum state tomography of dissociating molecules. *Phys. Rev. Lett.*, 91(9), 2003. quant-ph/0301135.

[52] X. Sun and A. C. Yao. On the quantum query complexity of local search in two and three dimensions. In *Proc. IEEE FOCS*, pages 429–438, 2006.

[53] L. G. Valiant. Evolvability. ECCC TR06-120, 2006.

[54] N. V. Vinodchandran. A note on the circuit complexity of PP. ECCC TR04-056, 2004.

[55] S. Zhang. New upper and lower bounds for randomized and quantum local search. In *Proc. ACM STOC*, pages 634–643, 2006. quant-ph/0603034.