The different forms of quantum computing skepticsm

Boaz Barak

(see also pdf version)

Quantum computing is one of the most exciting developments of computer science in the last decades. But this concept is not without its critics, often known as "quantum computing skeptics" or "skeptics" for short. The debate on quantum computing can sometimes confuse the *physical* and *mathematical* aspects of this question, and so in this essay I try to clarify those. Following Impagliazzo's classic essay, I will give names to scenarios or "potential worlds" in which certain physical or mathematical conditions apply.

Potential worlds

Superiorita is the world where it is feasible to build scalable quantum computers, and these computers have exponential advantage over classical computers. That is, in superiorita there is no fundamental physical roadblock to building large quantum computers, and hence the class BQP is a good model of computation that is physically realizable. More precisely, in superioriata the amount of resources (think dollars) that is required in order to simulate a T-gate quantum circuit grows at most polynomially or maybe even linearly (with not-too-terrible constants) in T.

The other aspect of superiorita is the *mathematical conjecture* that quantum computers offer exponential advantage over classical ones. That is, that there are functions computable by the mathematical model of (uniform) quantum circuits that require exponential time to compute by Turing machines. (In complexity jargon, this is the conjecture that $BQP \not\subseteq SUBEXP$ where the latter stands for the class $TIME(2^{n^{o(1)}})$.) *Integer factoring* is one problem that is conjectured to lie in $BQP \setminus SUBEXP$ (i.e., where quantum computers have an exponential advantage). One can also consider analogous conjectures for *sampling problems*, and some particular sampling tasks that can be achieved in quantum polynomial time have been conjectured as requiring exponential time for probabilistic Turing machines.

Superiorita is the world in which most quantum computing researchers think we live in, and, judging by the hundreds of millions of dollars of investments, many commercial companies and funding agencies as well. Note that this is a mix of both a *physical* assumption (that the model of BQP can be physically realized) and a *mathematical* assumption (that this model offers exponential speedup over classical machines). Without assuming *both* the physical and mathematical aspects of superiorita there would be no justification for investing huge efforts in building quantum computers.

In superiorita quantum computers are not a panacea and in particular they can't solve NP complete problems. Let me not wage into the (hugely important!) question of whether in superiorita the *Lattice Shortest Vector Problem* is in BQP or not, see my essay on the complexity of public key cryptography for more on this topic. Whether or not the particular problems on which quantum computing offer exponential speedup are *interesting* is a matter of taste. As far as I know, factoring large integers is not inherently interesting in its own right, and once the world moves to different encryption standards, the applications to breaking encryption will eventually disappear. However, there are other tasks where quantum computers seem to provide exponential speedups and that can be interesting in their own right in areas such as chemistry and machine learning (though one should read the fine print).

Popscitopia is the "hyper superiorita" world where quantum computers can solve NP complete problems. That is, in popscitopia quantum computers can be built, and $NP \subseteq BQP$. This is the world that is described by some popular accounts of quantum computers as being able to "run exponentially many parallel computations at once", a belief that is prevalent enough that Scott Aaronson devotes the tagline of his blog to refuting it. Most researchers in the area believe that, regardless of whether quantum computers can be physically be built, they cannot solve NP-complete problem (a belief which is essential to so called "post quantum cryptography"), and indeed so far we have no reason to think quantum computers off exponential (or even better than quadratic) speedup for such problems. But, we have no *proof* that this is the case, and indeed, some TCS researchers, as Richard Lipton, have suggested that even $NP = P$ (which in particular implies $NP \subseteq BQP$) might be true.

Skepticland is the world where it is not possible to build scalable quantum computers, though mathematically they do offer an exponential advantage. That is, in skepticland, for every function F (and more generally a promise problem or a sampling problem) that can be computed using T amount of physical resources, there is a probabilistic Boolean circuit of size polynomial in T that computes F as well. However, mathematically, like in superiorita, it is still the case in skepticland that BQP contains functions (such as integer factoring) that require exponential time to be computed classically.

Skepticland is the world that "quantum computing skeptics" such as Gil Kalai, Leonid Levin and Oded Goldreich think we live in. In this world the extended Church-Turing hypothesis hold sway and there exists some (yet unaccounted for) cost that blows up exponentially in T when trying to physically realize size T quantum circuits.
These skeptics still accept the mathematical conjecture underlying superiorita that BQP contains functions that require exponential time for deterministic or probabilistic Turing machines. Indeed, as far as I can tell, their belief in the inhrent difficulty of problems such as factoring is a large part of the intuition for why quantum computers would not be physically realizable.

Finally, **Classicatopia** is the world where $BQP \subseteq BPP$ and more generally

any function, promise problem, or sampling problem that can be solved by (uniform) quantum circuits can be solved by probabilistic Turing machines with a polynomial overhead. In this world quantum computers *can* be physically realized, but only because they are no more powerful than classical computers. Hence the Extended Church-Turing holds but for a completely different reason than in Skepticland. In Classicatopia we can simulate the entire physical world using a classical computer. One advocate of this world is Ed Fredkin (who interestingly was the person who motivated Richard Feynmann to propose the possiblity of quantum computers in the first place). Also, several researchers (such as Peter Sarnak) have suggested that the marquee problem of integer factoring can be solved by polynomial-time Turing machines.

Truth and beauty

At this point I should probably talk about the evidence for the probability of *truth* of each of these scenarios, and discuss the latest advances in experimental works building quantum computers. But frankly I'd be just parroting stuff I Googled, since I don't really know much about these works beyond second or third hand reports.

Rather, I'd like to talk about which of these worlds is more *beautiful*. *Beauty* is in some ways as important for science as truth. Science is not just a collection of random facts but rather a coherent framework where these facts fit together. If a conjecture is "ugly" in the sense that it does not fit with our framework then this can be evidence that it is false. When such "ugly ducklings" turn out to be true then this means we need to change our standards of beauty and come up with a new framework in which they fit. This is often how progress in science is made.

While I am not a physicst, I believe that *quantum mechanics* itself followed exactly such a trajectory. (I am probably making some historical, physical, and maybe even mathematical mistakes below, but I hope thebigger picture description is still accurate; however please do correct me in the comments!)

The ancient greek philospher Democritus is often quoted as saying *"Nothing exists except atoms and empty space, everything else is opinion."* This saying is usually interpreted as an emprical *hypothesis* about the world, or to use mathematical jargon, a *conjecture*. But I think this is really more of a *definition*. That is, one can interpret Democritus as not really making a concrete physical theory but defining the allowed space for all physical theories: any theory of the world should involve particles that mechnically and deterministically evolve following some specific and local rules.

Over the coming years, scientists such as Newton, Leibniz and Einstein, took this prescription to heart and viewed the role of physics as coming up with every more general and predictive theories within the democritus model of deterministic particulars with no randomness, intent, or magic such as "action at a

distance". In the late 1910's, Emmy Noether proved some remarkable theorems that derived conservation laws from physical theories based only on the fact that they satisfy certain *symmetries* (see also my recent post). While the mechanical clockwork theories satisfied such symmetries, one could think of more general classes theories that satisfy them as well. Noether's theorems showed that even *non-clockwork theories* could still satisfy a more general notion of "mathematical beauty".

At the time Noether's Theorems were just a very useful mathemtical tool, but soon nature gave some indications that she prefers Noether's notion of beauty to Democritus'. That is, a series of experiments led to the introduction of the distinctly "non clockwork" theory of quantum mechanics. Giving up on the classical notion of beauty was not easy for physicsts, and many (most famously Einstein) initially thought of quantum mechanics as a temporary explanation that eventually will be replaced by a more beautiful "Democritus-approved" theory. But Noether's results allowed to make quantum mechanics not just *predictive* but *beautiful*. As Nima Harkani-Hamed says:

> Newton's laws, even though they were the first way we leaned how to think about classical physics, were not the right way to make the jump to quantum mechanics. … [Rather] because the underlying ideas of the action– and everything just really ports beautifully through, from classical to quantum physics, only the interpretation changes in a fundamental way– all of Noether's arguments, all of Emmy Noether's arguments about conservation laws go through completely unscathed. It's absolutely amazing. All these arguments about conservation laws, many other things change, tons of other things changed when we went from classical to quantum. But our understanding of the conservation laws, even though they're come up with by this classical physicist a hundred years ago, are equally true in quantum mechanics today.

Moreover, my outsider impression is that with time physicsts have learned to accept and even *grow to love* quantum mechanics, to the degree that today many would not *want* to live in a purely classical world. If you wonder how anyone could ever love such a monstrosity, note that, as Scott Aaronson likes to say, there is a sense in which the relation between quantum and classical physics is analogous to the relation between the ℓ_2 and ℓ_1 norms. I think most mathematicians would agree that the former norm is "more beautiful" than the latter.

My personal opinion

So, which is the most beautiful world, superiorita or skepticland?

If you've asked me that question a decade ago, I would have answered "skepticland" without hesitation. Part of the reason I got into computer science is

that I was never good at physics and didn't particularly like it. I also thought I could avoid caring about it. I believed that ultimately the world is a Turing machine or cellular automata and whether it has 5 or 12 particles is about as interesting as whether the computer I'm typing this on uses big endian or little endian representation for integers. When I first heard about quantum computing I was hoping very much that there is some inherent reason it can never work so I can avoid dealing with the uglyness of quantum mechanics and its bracket notation.

But as I've learned more about quantum mechanics, I've grown not just to accept it as a *true* theory but also *beautiful*, and with this to also accept quantum information and computation theory as a beautiful generalization of information and computation in its own right. At the moment I don't see any beautiful alternative theory (to use Aaronson's terms, a "Sure/Shor separator") from the skeptics. The closest we have to such a theory comes from Gil Kalai, but as far as I can tell it posits *noise* as a new fundamental property of nature (the Ka-la-ee constant?). Noise here is not the usual interpretation of quantum probabilities or the uncertainty principle. It seems to be more similar to the engineering form of noise as inaccuracies in measurements or errors in transmissions. While these can be serious issues (for example, I believe that friction is a large part why actually building Babbage's Analytical Engine was so difficult). But as far as I can tell, these engineering difficulties are not fundamental barriers and with sufficient hard work and resources the noise can be driven down to as close to zero as needed.

Moreover some of the predictions involve positing noise that scales with number of qubits in the computer. It seems to require nature to "know" that some physical system in fact corresponds to a logical qubit, and moreover that two distant physical systems are part of the same quantum computer. (I should say that Gil Kalai disagrees with this interpretation of such scaling.) While one could argue that this is not more counterintuitive than other notions of quantum mechanics such as destructive interference, entanglement, and collapse under measurements, each one of those notions was only accepted following unequivocal experimental results, and moreover they all follow from our modelling of quantum mechanics via unitary evolutions.

The bottom line is that, as far as I can tell, superiorita is the most beautiful and evidence-supported world that is currently on offer.

Will we see a mega-qubit quantum computer?

The current experimental efforts are aimed at building a 50 qubit quantum computer. This sounds impressive until I remember that the VIC 20 I played with as a third-grader more than thirty years ago already had 5K (i.e., about 40,000 bits) of memory. So, will we ever see a quantum computer big enough to run Frogger? (not to mention Ultima IV)

The answer to this question depends not just on the science but also on economics and policy as well. Suppose that (with no real justification) that eventually we will able to produce a quantum computer at a cost of 1000 dollars per qubit. Then a million qubit machine will cost a billion dollar to build. The current applications of quantum computers do not seem to justify this cost. As I mentioned, once we transition to different cryptosystems, the motivation for factoring integers will be significantly lessenned, and while simulating quantum systems can be important, it's hard to see it as forming the basis for a billion dollar business. Of course, this can all change with a single theory paper, just as Peter Shor revolutionized the field of quantum computing with a single algorithm.

Moreover I hope that at some point, policy makers and the public at large will stop viewing computer science just through the lens of applications, and start seeing it also as a fundamental science in its own right. The large Hardron Collider apparantly cost about 13 billion dollars to build and operate, and yet the same analysis calls it a "bargain" in terms of the benefit from both technologies invented and scientific discovery. The case can be made that building a large scale quantum computer would be no less important to science, and would offer no less benefit to society. Indeed, a quantum computer offers literally an exponential number of potential experiments one can run on it. Moreover, there is absolutely no reason to think that Shor gave the final word on breakthrough algorithms that could use such a computer for tasks that a priori seem to have nothing to do with physics. In that vein, I hope that whatever bodies that fund experimental quantum computing research realize that at least part of their investment should go into theoretical work in quantum (and also classical, as the two are intertwined) algorithm design.

Acknowledgements: Thanks to Gil Kalai and Scott Aaronson for comments on earlier versions of this post. Needless to say that are not responsible for anything that I said here.