Quantum Computing

A preliminary research analysis report

by Jaime Sevilla

In this article, I discuss the relevance of Quantum Computing (QC) from the point of view of a long-termist philanthropist who wants to understand what lines of research related to QC are worth funding.

This report is split into five parts

- The first part of my analysis will be focused on the relevance of QC for Artificial Intelligence (AI), where I have spent more time thinking about the topic. This section is based on a lit review of Quantum Machine Learning and my own intuitions of the field. Counting my previous research, I have spent about a month on this topic. For this article I have spent ~2 days synthesizing my thoughts on the topic.

- The second part of my analysis will discuss cryptanalysis and cryptography. This part is mainly informed by a NIST report. I have spent about ~3 days researching this topic.

- In the third part I will cover applications of QC to efficient simulation of quantum systems, and discuss the relevance of this application for hardware design, biosecurity and Atomically Precise Manufacturing. For this section I have informally interviewed some experts on the last two topics, Jassi Pannu and Gregory Lewis on biosecurity and Eric Drexler on APM. In total I have spent about a day worth of research on each of the applications, for a total of ~3 days of research.

- In the fourth part I cover applications to industry and basic research. The information on industry applications come from publicly available information on IBM's, Google's and Microsoft's quantum research programs and educated guesses on their (lack of) downside risks. The information on applications to basic research come from some reports by academics. I have spent about two days researching this topic.

- In the fifth and final part I discuss open questions on QC Strategy that may help clarify the scope and relevance of this technology for the applications discussed before. This is guided by an educated guess on my part and no in-depth research.

In total, the time I dedicated to the elaboration of this report is about ~80 hours. The results in this report are not meant to be final, but rather educated guesses that may help inform researchers and philanthropists considering whether to further investigate some of the areas I explored.

My background is in mathematics and computer science. I have done previous research on applications of AI, and I have taken some courses on QC. I have checked my claims with some experts on quantum computing.

Executive summary

- QC could unlock a significant amount of computing power for AI. We will most likely a quadratic speed-up for general optimization problems - which would allow us to solve some problems we already know how to solve significantly faster. There is a chance that we could achieve an exponential speed-up for a narrow class of problems -which could enable us to solve problems that are unsolvable with current computers. This could have implications on AI timelines and the kind of AI designs we will first see. This area of research seems well funded by big companies (Google, IBM, Microsoft). More

- Due to its potential applications to advance transformative AI, I would be medium excited to see further research on the effects of QC compute overhang on AI timelines and possible policy implications for regulation of AI development. This line of research seems however bottlenecked by our understanding of how extra compute in general affects AI timelines, research on QC timelines and research in QC policy. More

- I expect QC developments neither hinder nor help with bottlenecks in current research agendas in AI alignment, and I see no need to currently invest in expertise in this intersection. More

- QC compromises secure communication and online transactions, which could have profound economic effects. The scale of this problem is unclear. This issue seems to be receiving a commensurate amount of attention from government orgs (NIST) and industry (ETSI), and I expect that post-quantum crypto will indeed be found to be as secure as classical public-key crypto, with only a modest additional computational overhead for participants. More

- QC could have an effect on computer hardware design. I am unsure of how promising this research line is compared to other research lines in hardware design. I'd be keen on seeing a 10 hour research project reflecting on this. More

- Very tentatively, I'd guess that the development of QC would not significantly increase bio risk nor it is a specially promising tool for mitigation. More

- QC would likely not have a major effect on Atomically Precise Manufacturing. Current research in APM is able to abstract away from the quantum mechanics of precise protein interactions and thus able to use efficient classical simulations for design. More

- QC has some promising applications with little apparent downside in medicine, agriculture (fertilizer design) and operations research. Big companies like IBM, Google and Microsoft are aware of these applications, so I would expect little

room-for-funding. It is also unclear whether there is lower hanging fruit for these applications via other avenues. More

- Some researchers are preemptively exploring QC simulations as a basic research tool in physics and chemistry, and research of limits of QC can illuminate some questions about fundamental physics. However it is not clear to me that basic research on physics is urgent, and I am also uncertain about how useful QC tools would be in chemistry research. An interview with an expert in quantum chemistry may clarify the issue. More

- Some strategic and tractable research on general Quantum Computing that would be helpful to better understand its relevance includes 1) forecasting QC timelines and 2) researching whether and how QC can be regulated. I think that a 100h research project on those questions is worth funding insofar it will allow us to better understand the VOI of strategic research on concrete applications and regulation of QC. More

- I would recommend funding some more exploratory research in QC grantmaking as a high-risk high-reward research project. I would, however, recommend against funding object-level research or pursuing careers in QC for philanthropic reasons. More

In this summary, I have given guesses on what is worth funding. Researchers may want to look into the sections they are considering investigating themselves for open questions.

Introduction

Quantum Computing (QC) is a disruptive technology that may not be too far ahead on the horizon. Small proof-of-concept quantum computers have already been built [IBMQ] and major obstacles to large-scale quantum computing are being heavily researched [IEEE].

At its essence, QC allows a computer to hold a state that is a complex linear combination of classical states, shortcutting some computations. This allows quantum computers to solve certain problems faster or using less memory — but not all, and adapting classical algorithms to achieve even modest speed-ups using QC is an arduous task.

For a gentle introduction to QC, I recommend “Quantum computing for the very curious”, an interactive essay by Andy Matuschak and Michael Nielsen [QCVC].

Among its potential uses, QC will allow breaking classical cryptographic codes, simulate large quantum systems and faster search and optimization [WOLF].

Some big companies (Google, Microsoft, IBM) and startups (Rigetti) have been racing towards creating the first quantum hardware capable of showcasing a quantum speedup by solving a problem using quantum processors that we do not know how to efficiently solve with classical computers (this is known as the quantum supremacy milestone). Google recently claimed this milestone [QSPSP], showcasing a quantum processor capable of

sampling in 0.02 seconds a probability distribution that would take the current faster supercomputer (IBM's Summit as of this writing) 2.5 days to sample [IBMQS].

We are likely to first see special-purpose machines, but there are some theoretical proposals for general-purpose computation schemes [QCT].

Ultimately, timelines for large scale, general purpose QC are unclear. Some quantum computing researchers claim that it is 30 years away[1], but it could be substantially sooner or later.

The investment in quantum computing technologies is quite substantial. In 2015 about 7,000 people worldwide, with a combined budget of about $1.5bn, were working on quantum-technology research [ECON]. By way of comparison, in 2018 there were about 22,000 experts worldwide working on Artificial Intelligence [AITALENT], with a combined budget of $37.5bn [AISPEND].

Quantum computing and AI

In the next two sections, we discuss the relation between QC and AI from three points of view: whether QC is important for AI capabilities, the implications for AI strategy and whether it is relevant for technical AI alignment.

We conclude that a profound understanding of Quantum Computing is unlikely to help with technical AI alignment, albeit better forecasting of QC timelines and research analysis of the interaction between QC and advanced AI research may be of strategic interest.

Big companies like Google are heavily investing in quantum machine learning. It seems like QC will not lead to generalization insights directly, but the computational speedup may enable new AI technologies.

Quantum computing for AI capabilities

Some early research suggests that QC could be used to speed up some core optimization problems with wide applicability. Some of these problems include:

- Quadratically faster unstructured search using Grover's algorithm [GROVER],
- Finding extrema over finite domains [CHAN],
- Sampling complex distributions [QSP], and
- Solving large, sparse, well-conditioned systems of linear equations exponentially faster [HHL].

Early research in the newly found field of Quantum Machine Learning (QML) seems promising, and sufficiently powerful quantum computers can in principle provide computational speedups for key machine learning algorithms and subroutines such as data

[1] "For general applications, 30 years is "not an unrealistic timescale", says physicist John Preskill at the California Institute of Technology in Pasadena." [BQS:HUQC]

fitting, principal component analysis, Bayesian inference, Monte Carlo methods, support vector machines, Boltzmann machines, and recommendation systems [CVQNN].

On the other hand, enthusiasm seems to have waned after some recent dequantization results, where purported quantum speed ups have been matched by classical advances [TANG].

Overall, it seems like the effects of quantum computing availability in AI research is comparable to that of compute overhang. This is because research in quantum algorithms mostly aims to find faster versions of already known classical algorithms and quantum routines naturally want to be isolated subroutines. Thus we expect to live in a world of QC as a transparent service that researchers and programmers can abstract away from, rather than of quantum computing as a paradigm shift in how we approach algorithm design.

At best, QC speeds up some algorithms exponentially. At worst, Grover's algorithm can provide a quadratic speedup of almost any parallelizable search procedure. In practice, there are some caveats towards achieving exponential speedups [QML:RFP].

It is unclear whether the overhead and cost of design will make this worth the effort, but it remains a possibility - arguably the deep learning revolution was enabled in a similar way by advances in parallel computing.

On the other hand, the brain does not seem to use quantum speedups[2], which is evidence of quantum speedups being hard to exploit or not that useful for general-purpose AI.

In any case, Google is publicly justifying their investment in QC due to its applications to Quantum Machine Learning, and created a Quantum AI Lab group. IBM and Microsoft seem also interested in these applications and have funded some research in this area. On the startup side Rigetti is also an important player in general QC, and Xanadu seems particularly focussed on ML [XANADU].

Given the amount of interest by big players, it seems unlikely that a marginal dollar would make much of a difference in this line of research, and we might not want to speed up the development of AI in any case [OPENPHIL:AI].

The key questions one may ask are:

- What AI designs would be particularly favored by QC technology versus raw improvements in compute?
- Are there any "dormant" AI algorithms that are currently too computationally intensive to be useful but are sped up by quantum? For example, would Bayes net architectures be significantly augmented by improved Boltzman sampling[3]?

[2] Quote from Scott Aaronson: "The brain is a hot, wet environment, and it's hard to understand how long-range coherence could be maintained there. (With today's understanding of quantum error-correction, this is no longer a knock-down argument, but it's still an extremely strong one.)" [SAP]

[3] From the edx course on QML (U of Toronto): "The roots of probabilistic graphical models go back to the 1980s, with a strong connection to Bayesian statistics. The story resembles that of neural networks: they have been around for over three decades and they need massive computational power.

- How much does it matter to get an exponential speed up versus a "merely" polynomial sped up?

Quantum computing for AI Strategy

To the extent that quantum computing will unlock compute, and to the extent that further compute will speed up development in AI, I expect QC to be relevant to AI timelines.

Furthermore, to the extent that we expect AI development driven by increases in compute to lead to more opaque AI, advances QC might increase the risk of unaligned and opaque advanced AI systems.

It is unclear whether we will find problems where an exponential speedup can be effectively achieved. And while it seems likely that asymptotic polynomial speedups are possible it is unclear how constrained will they be by the inconvenience of setting up reliable QC architecture and bottlenecks in preparation time.

Nevertheless, if quantum computing ends up being a core component of advanced AI systems then this fact would have profound political implications - barring a major surprise, I expect quantum computers to be used as specialized computer hardware by major companies, rather than as a commodity. This is because we currently believe that QC requires very extreme physical conditions to work [IEEE].

This means for example that while QC services might be made available to a wider public through cloud platforms, the hardware itself should be relatively easy to track and regulate.

Some open questions one may look at at QC and AI strategy are:

- How will advances in QC affect AI timelines?
- How would expected developments in how QC is developed and deployed affect AI race dynamics?

In general research in this intersection would need to be preceded by a cross domain strategic analysis of QC, forecasting its development and researching political levers that could be used to regulate this technology.

As far as I know, nobody is working on this particular intersection of domains.

Quantum computing for AI alignment

Long story short, I do not believe that QC is a critical area of knowledge for advancing current research agendas of technical AI alignment, and I would weakly recommend against pursuing a career in it for this purpose or funding research in this intersection.

However, unlike in the case of deep learning, the requirements for computational resources remain out of reach. These models require sampling a distribution, and very often it is the Boltzmann distribution. Since quantum computers can give samples from this distribution, we can hope that quantum hardware can enable these models the same way graphics processing units enabled deep learning." [UEDX:QML]

For the full discussion of my reasoning and a list of open questions, I refer the reader to my latest paper on the topic [QC+AIA].

The biggest caveat to this conclusion is that if we expect QC to be a core component of advanced AI systems this may allow us to make some assumptions about advanced AI that may simplify research in alignment.

If one was set to research this intersection, some promising questions are:

- We could imagine a verifier that leverages access to QC to oversee a potentially unsafe agent without access to QC. In this vein, how can asymmetry of QC resources be exploited for safety purposes?
- In the case where QC is relevant for the design of advanced AIs, can we expect to have an actual quantum agent in the future, or will it just be a classical agent with access to quantum subroutines, in a Comprehensive AI Services fashion [CAIS]?

As far as I know, nobody is working on this intersection of domains.

Quantum computing and cryptanalysis

In this section, we expose one key risk associated with quantum computing: the failure of modern cryptography.

While analyzing the magnitude of this risk falls beyond the scope of this report, I give a brief overview of the problem, talk about who is working on this and briefly talk about possible solutions.

In essence, we conclude that this is a potentially important and tractable problem but that seems already well attended by the government and industry.

Post-quantum cryptography

Quantum computing will allow us to efficiently solve certain problems that are conjectured to be really hard to solve - and we rely on those problems to be hard to solve for modern-day cryptography to be secure.

Particularly, the security of all commonly used public-key schemes (RSA, elliptic curve and finite field-based cryptography) relies on the factorization of large quasi-prime numbers being hard to compute. We, however, know an efficient algorithm for integer factorization using quantum computing - Shor's algorithm [NIST:PQC].

Regardless, quantum-resistant public key schemes - Post-Quantum Cryptography (PQC) - are being heavily researched [NIST:WEB]. Modern-day cryptography has two key applications - signing messages in a way nobody else can forge (signature schemes) and encrypting messages so only the intended recipient can read them (key establishment schemes). The security of current hash-based signatures against quantum attacks is well understood - doubling the size of the hashes should suffice to make the hashes quantum-resistant. We have several candidates for key establishment schemes, like those

based on lattices and codes, but we do not have quantum security guarantees as strong as against classical attacks.

The security of the most widely used symmetric key (AES) and hash function (SHA3) algorithms could also be compromised by QC due to the quadratic improvement provided by Grover's algorithm, but this particular via of attack can be neutralized by doubling the key size and hash size respectively.

Overall it seems like to avoid the cryptographic problems that would come with the advent of quantum computing we would need to:

1) Produce official standards for PQC. For this to happen, we would need to research whether we can get some security guarantees on the lattice-based or code-based schemes currently designed. Should that endeavour fail, we would need to research new quantum-resistant encryption schemes.
2) Promote the transition to post-quantum cryptography before quantum attacks can compromise key infrastructure

Given the amount of information I was able to find online, I weakly believe that current research trying to solve the first challenge is well-aimed, well-funded and well-supplied in terms of talent.

The second challenge is, in my opinion, a bigger bottleneck. Some related research questions:

- Who would need to act for PQC to be implemented? How quickly have the relevant actors reacted to technical security issues before? What policy and industry compliance tools do standards organizations have?
- All considered PQC schemes for signatures, encryption, hashes and symmetric key rely on bigger key sizes, which would break some other internet protocol standards. How big of a problem is that? What can we do to mitigate this problem?
- How much damage could be done if large-scale QC arrives before post-quantum cryptography is widely implemented?
- Which PQC schemes are more promising? What is their room-for-funding?
- What timelines do the organization and people working on this problem have in mind? How likely is a sudden burst of development that catches people working in PCQ unaware?

Some organizations like ETSI (nonprofit, mainly supported by industry members) and NIST (US public organization) seem invested in pushing for a timely change towards PQC. This makes me guess that PQC is receiving an adequate amount of resources.

Quantum cryptography

With quantum cryptography we refer to techniques the exploit quantum mechanics to ensure secure information exchange between two parties.

Quantum cryptography is with respect to QC a separate but closely related field of study. Their relations are that 1) quantum cryptography is expected to be quantum-attack resistant and 2) both fields rely on the same principles of quantum information theory.

The most promising quantum cryptography scheme is quantum key distribution (QKD), which allows the exchange of private keys, with an information-theoretic guarantee that any attempts to eavesdrop on the key exchange will be detected.

This guarantee requires some assumptions to work. One of them is that the communication channel used must be authenticated, ie the sender must be aware of the true identity of the receiver. This means that QKD works as an encryption scheme, but would need to be complemented by an authentication scheme to allow practical secure communication. Some quantum-secure Message Authentication Codes have been proposed [QSMAC].

There are already some commercial implementations of QKD (using the BB84 protocol) [IDQ] [MAGICQ] [QLAB] [SQN]. These implementations require specialized, expensive hardware, which makes it impractical outside of a handful of applications like bank-to-bank communication.

Infamously, implementations of BB84 have been hacked due to imperfect implementation that invalidates the assumptions needed for theoretical security [XQL]. Furthermore, the BB84 has a known, possible deal-breaking vulnerability via Photon Number Splitting.

Overall, despite QKD having some theoretical guarantees that other PQC proposed schemes currently lack, the need for specialized hardware and the repeated failures at a secure implementation make me sceptical this is a practical solution to PQC.

Moreover, research in QKD implementations seems to be well attended and funded by startups and the military, so I'd expect little room-for-funding for philanthropists.

Quantum computing and simulation of quantum systems

One of the most promising applications of QC is efficient simulations of physical systems [UQS]. This opens the doors to many classes of efficient *in silico* design spaces. We reason about the applicability of efficient simulation of quantum systems to three critical risk areas: computer hardware design, biosecurity and Atomically Precise Manufacturing (APM).

Further research on the interaction effects on these areas would be interesting to conduct.

Quantum computing and computer hardware design

Compute overhang could fuel the design of advanced AI designed through raw search procedures (eg evolution-inspired design) rather than first principles, which arguably leads to more opaque AI systems.

We previously discussed how QC could lead to compute improvement on quantizable algorithms, but it could also indirectly lead to further compute overhang by enabling the design of AI-specific accelerators like improved Tensor Processing Units (TPUs).

Overall I would guess that it is unclear to which extent compute improvements lead to advances on AI capabilities. I think there is moderate evidence in favor of the conclusion that to the extent that it does, it differentially hinders interpretability and AI Alignment.

Some lines of research which I consider tentatively promising:

- Are there any "dormant" AI algorithms whose theory is well understood but are currently considered infeasible due to compute requirements?
- How much compute could we achieve through QC-fueled hardware design?
- Can QC hardware-design lead to recursive improvements in QC hardware?
- How necessary is QC-fueled design for hardware improvements? Are there easier ways of achieving hardware improvements?

Quantum computing and biosecurity

This section is based on my conversations with Jassi Pannu and Gregory Lewis about biological security and OpenPhil's report on biosecurity [OPENPHIL:BIO].

Among the different types of biological risks we face, Quantum Computing would have the greatest effect on synthetic biology risks, where computational design might be an important part of the pipeline.

However, after talking with the above mentioned experts I came away with the impression that artificial biological design is not overly constrained by any single bottleneck. Improvements to protein simulation would be a minor, welcome improvement, but would not unlock a drastically different design space.

As an analogy, Lewis pointed me to how the genome project gave us the complete human genome, but we were not able to immediately operationalize that knowledge. Biological systems have many parts that interact with one another in multiple ways, and progress feels mostly incremental rather than blocked on a few key problems.

In any case, further research would be required to better understand how these interactions would play out. Some questions I would be keen to see research on include:

- How much of a bottleneck is computer design for dual use synthetic biology research versus laboratory design? And for vaccine development?
- What concrete quantizable biological simulation algorithms are we aware of? How widely applicable are them to synthetic biology? What speed-up do they achieve?

If I had to take a wild guess, I would probably err on the side of quantum computing being on net slightly positive for preparedness versus biorisks, but in the end not worth prioritizing over promoting better oversight, creating stockpiles of important medical countermeasures or more traditional avenues of research towards the creation of broad-spectrum virucides.

Quantum computing and atomically precise manufacturing

This section has been informed by a short chat with Eric Drexler and OpenPhil's report on Atomically Precise Manufacturing (APM) [OPP:APM].

For an introduction to APM risks I refer the reader to OpenPhil's report. Taking one quote from the article for context:

> "**What is the problem?** If created, atomically precise manufacturing would likely radically lower costs and expand capabilities in computing, materials, medicine, and other areas. However, it would likely also make it substantially easier to develop new weapons and quickly and inexpensively produce them at scale with an extremely small manufacturing base. In addition, some argue that it would help make it possible to create tiny self-replicating machines that could consume the Earth's resources in a scenario known as "grey goo," but such machines would have to be designed deliberately and we are highly uncertain of whether it would be possible to make them."

Arguably, one of the key bottlenecks in APM has been the design of proteins with specific properties. While some promising candidates have been identified (like beta solenoids), having a more ample catalogue of selection could increase the chances of the success of APM. Efficient simulation of protein folding is among the applications of QC, so we could naively suspect an interesting interaction effect.

However, Drexler argues that 1) protein fold *prediction* is not the relevant problem to solve for APM, but rather protein engineering (protein fold *design*) and 2) current techniques for classical simulation of proteins are efficient and accurate enough for designing APM machines[4].

As I understand it, the distinction between protein folding and protein engineering is that they are in a sense the inverse of one another.

- In protein folding, we are given a particular protein chain and we are tasked with the prediction of how it will fold.
- In protein engineering, we have a particular shape that we want to achieve and we want to design a protein chain that will take that particular shape.

In his seminal paper Drexler illustrates some efficient approaches to protein engineering [DRX].

Furthermore, while it is true that past design was bottlenecked on efficient computation, Drexler argues that current techniques of simulation, which abstract away the quantum calculation of configuration-dependent electronic energies in favor of classical approximations, are fairly efficient and accurate enough for APM.

[4] Quote from Eric Drexler: *"I don't see simulation of quantum systems as important to APM or other (essentially non-electronic) atomically precise technologies. These are, in effect, classical systems (mechanical, not electronic degrees of freedom), and existing methods predict energies and forces well enough that better methods would make little practical difference."*

Technically speaking, as far as we know QC offers exponential speedups for the evolution of quantum systems, but polynomial for finding the ground state (stable state). The latter is the one we are interested in for protein engineering.

Furthermore, Drexler argues that the loss of accuracy can be thought of as a small perturbation in the system. While proteins often are sensitive to small perturbations, the kind of proteins used for APM are engineered for high stability, which arguably counters this effect. This is true of soft APM, and even more true of hard APM, as more stiff materials as less affected by this "perturbation".

Overall I am not excited about research in this intersection. If one was set on researching it, one possible question to look at is:

- Which algorithms with immediate applications to protein engineering are known to be quantizable?

Quantum computing in industry and academia

In this section we briefly talk about other applications of QC that are interesting for industry and academic purposes, and reason about potential downside risks that a philanthropist may want to look into mitigating.

Industrial applications of quantum computing

Companies like IBM, Google and Microsoft justify their investment in QC appealing to potential industrial applications. Those include logistics, energy use optimization, financial modelling (via portfolio optimization), agriculture (via better fertilizers), weather forecasting (via efficient quantum simulation) and medicine (via *in silico* automatic drug discovery) [IBM].

As far as I am aware, these applications have no salient long-term risks associated with them. They do have a tempting upside (in this regard, research in better fertilizers seems particularly promising), but it is unclear how tractable they are and whether QC is the lowest hanging fruit research in those areas. In any case, big companies like IBM, Microsoft and Google seem well-aware of them and in a good position to research them so I expect little room-for-funding.

Some research questions broadly related to these applications include:

- What risks may stem from these applications?
- What are alternative avenues of research in these applications? How favorably do they compare to QC?
- If we expect that QC is likely to lead to enormous advances in one of these areas like fertilizer design, how can we ensure that the benefits are captured by society? What intellectual property institutions do we want to develop to ensure the technology benefits as many as possible without discouraging investment from industry? How will the centralized nature of QC (due to hardware upfront costs) affect this issue?

Quantum computing for basic scientific research

In this section we look at how QC may lead to further basic research in physics and chemistry.

Intuitively it seems like polynomial speed ups are less significant when regarding applications for basic research, as the bottleneck of research seems intuitively to be experiment design rather than execution, so reducing eg an experiment that takes a week to an hour is an improvement but not a game-changer. In contrast, an exponential speedup may signify the difference between an experiment being realizable at all or not.

For this reason we will focus mainly on the place where we might expect an exponential speedup relevant to research - quantum simulation.

The ability to efficiently simulate quantum systems makes QC an interesting playground from theoretical physics. For example, some physicists are toying with proposals for QC tools to do basic research in quantum field theory, for example to better understand quantum chromodynamics [SQFTQC].

Beyond the development of research tools that use QC, many physicists also look into the study of QC as a way to understand some basic facts about the universe.

For example, computer scientists like Scott Aaronson [NP+PHYS] suspect that QC is a universal computing scheme, in the sense that it could efficiently simulate any other physical process.

This is an interesting fact about the universe, with predictive power. For example, if this is true and we also knew that the class of problems that can be efficiently solved by a quantum computer (BQP) is different from the class of problems which can be efficiently verified in a classical computer (NP) then we could rule out closed timeline curves in relativity theory [NP+PHYS].

However, I do not see immediate applications for this knowledge, which make me by default doubt its urgency, as some others have pointed out that we have some good reasons to doubt the returns of basic research in physics nowadays [NINTIL].

Another area of knowledge that could take advantage of quantum speedups is **chemistry**.

Some examples of concrete applications in the area are simulation of nitrogen fixation processes, simulation of catalytic mechanisms of enzymes and determination of the spin of arbitrary wavefunctions [CHEM].

We suspect that classical computers cannot efficiently perform some of these arbitrary simulations due to exponential memory requirements. However, as we remarked in the section about APM we are only rarely interested in the behaviour of arbitrary systems, but rather some very concrete ones with practical applications that we might be better off studying with classical computational techniques or laboratory experiments.

Overall it is unclear to me how key of an enabler would quantum computing be for advancing basic research in chemistry. I'd recommend interviewing an expert working on this intersection and chemists to further clarify this issue.

Some open questions in this intersection are:

- What concrete basic research tools do we already know would be possible with quantum computing?
- What speedup do they achieve versus classic techniques?
- How widely used are the classical analogues?
- How useful and applicable is the research that would be enabled by these QC tools?

Crossdomain strategic research in quantum computing

Instrumental to understanding the applications above, there are some general questions about quantum computing that would be useful to understand better.

Timelines of quantum computing are unusually tractable - all the theoretical basis for quantum computing is already well researched, and what remains to be solved are engineering problems that seem overall tractable [IEEE].

Creating a map of actors in QC also seems like a tractable problem, because the specialized hardware and big upfront investment required restricts who can enter the R&D scene.

There are also some interesting research directions on regulation and forecasting usage of QC, which could also be applied to other technologies such as distributed computing. Relatedly, would be interesting to find technologies that are analogues to QC and reason about what insights can be imported from one technology to another.

In addition, it would be interesting to create a compendium of problems that can be sped up using quantum computing, develop better intuitions around what algorithms are susceptible to being quantized, and identify problems on which a significant amount of compute is being poured at and whose resolution could be sped up with QC.

In summary, these are some strategic questions around QC development that I think could be interesting research leads:

- What are QC timelines? What are major milestones ahead and when do we expect to reach them?
- Who are the major players in QC? Exactly how many resources are they investing in QC R&D? What are their motivations?
- How easy will be QC to regulate? Should we expect QC to be concentrated in the hands of a few big companies?
- Will QC be made available to the public through cloud services? How cheap will QC compute be?
- What are some useful parallels to compare QC development trajectories to (eg is parallel computing a useful analog? Tensor Processing Units? Why? Why not?)?

- What kind of algorithms can be quantized? Can we develop reliable intuitions and heuristics to tell us which tasks would significantly benefit from a quantum speedup?
- How are supercomputer resources currently used? Which of these tasks could be quantized?
- Weighting in all promising applications, is accelerating development in QC good or bad in expectation? If the latter, can key actors be persuaded to reduce their investment in QC?

Conclusions

In total I've spent about ~80 hours researching and writing this report. It is quite likely that if I was to spend another ~100 hours on the report I would be able to identify some other areas of concern or significantly change some of my conclusions.

In general, I believe that looking at the benefits and risks associated with new technologies through a philanthropic lense is a good high-risk high-reward strategy to find . QC is a good candidate for this task, and I'd be excited about funding or conducting further analysis in this topic.

I would however weakly recommend against premature grants to object level research; as seen in the previous sections, most research in most areas associated with QC development and associated risk mitigation seem to be 1) bottlenecked by other previous research questions, 2) current evidence points against QC being very relevant or 3) already being well researched by government, industry and/or academia.

I would also recommend against funding and/or pursuing careers in QC for the same reasons. This is because I expect that current talent in QC is already enough to conduct the exploratory high risk research outlined above.

This article was written by Jaime Sevilla. This work was partially supported by the Future of Humanity Institute summer fellowship program and partly by a grant made by the Effective Altruism Foundation.

I want to thank Jesse Riedel for sharing his previous research and substantive comments on the article, Pablo Moreno for discussion on Quantum Computing, Jassi Pannu and Gregory Lewis for discussion on biological risks, Eric Drexler for discussion on Atomically Precise Manufacturing, Max Daniel for mentorship and Luisa Rodriguez for general feedback.

Further thanks to Matīss Apinis, Nolan Winker and Daniel Kokotajlo for proofreading the article.

Citing this work

Please cite this work as
Jaime Sevilla, "Quantum Computing - A Preliminary research analysis report", Oxford (2019).

Bibliography

Quantum big picture

[IBMQ] IBM, "IBM Unveils World's First Integrated Quantum Computing System for Commercial Use", https://newsroom.ibm.com/2019-01-08-IBM-Unveils-Worlds-First-Integrated-Quantum-Computing-System-for-Commercial-Use, New York (2019).

[IEEE] C.G. Almodever et al, "The engineering challenges in quantum computing", IEEE, March 2017. https://ieeexplore.ieee.org/document/7927104

[BQS:HUQC] Michael Brooks, "Beyond quantum supremacy: the hunt for useful quantum computers". https://www.nature.com/articles/d41586-019-02936-3

[QSPSP] Quantum supremacy using a programmable superconducting processor https://www.nature.com/articles/s41586-019-1666-5

[IBMQS] https://www.ibm.com/blogs/research/2019/10/on-quantum-supremacy/

[WOLF] Ronald de Wolf, "The Potential Impact of Quantum Computers on Society". https://arxiv.org/pdf/1712.05380.pdf

[QCVC] Andy Matuschak and Michael A. Nielsen, "Quantum Computing for the Very Curious", https://quantum.country/qcvc, San Francisco (2019).

[QCT] Ethan Bernstein and Umesh Vazirani, "Quantum Complexity Theory", https://epubs.siam.org/doi/abs/10.1137/S0097539796300921, July 2006.

[SAP] Scott Aaronson, "PHYS771 Lecture 10.5: Penrose" https://www.scottaaronson.com/democritus/lec10.5.html

[ECON] "Quantum technology is beginning to come into its own" https://www.economist.com/news/essays/21717782-quantum-technology-beginning-come-its-own

[AITALENT] "Global AI Talent Report 2018" https://jfgagne.ai/talent/

[AISPEND] "Global Spending on AI Systems to Hit $98 Billion by 2023 – IDC" https://adtmag.com/articles/2019/09/04/ai-spending.aspx

Quantum Computing and Artificial Intelligence

[TANG] Ewin Tang, "A quantum-inspired classical algorithm for recommendation systems". https://arxiv.org/abs/1807.04271

[QML:RFP] Scott Aaronson, "Quantum Machine Learning Algorithms: Read the Fine Print". https://scottaaronson.com/papers/qml.pdf

[GROVER] Lov K Grover, “A fast quantum mechanical algorithm for database search”, https://arxiv.org/abs/quant-ph/9605043

[CHAN] Man Chan, “Quantum algorithms for finding extrema with unary predicates”. https://www.semanticscholar.org/paper/Quantum-algorithms-for-finding-extrema-with-unary-Chan/139a3623a6119d3bdb778cc0a853a465b69a3562

[QSP] A. P. Lund, Michael J. Bremner & T. C. Ralph. “Quantum sampling problems, BosonSampling and quantum supremacy”. https://www.nature.com/articles/s41534-017-0018-2

[HHL] Aram W. Harrow, Avinatan Hassidim & Seth Lloyd, “Quantum algorithm for solving linear systems of equations”. https://arxiv.org/abs/0811.3171

[QC+AIA] Jaime Sevilla & Pablo Moreno, “Implications of Quantum Computing for Artificial Intelligence Alignment Research”, https://arxiv.org/ftp/arxiv/papers/1908/1908.07613.pdf, August 2019.

[CVQNN] Nathan Killoran et al, “Continuous-variable quantum neural networks”, https://arxiv.org/abs/1806.06871 (see introduction section).

[QML:CP] “Quantum machine learning: a classical perspective”, https://arxiv.org/abs/1707.08561v3, July 2017.

[QML:UT] University of Toronto, “Quantum Machine Learning”, edX https://www.edx.org/course/quantum-machine-learning-2

[OPENPHIL:AI] Holden Karnofski, “Potential Risks from Advanced Artificial Intelligence: The Philanthropic Opportunity”, https://www.openphilanthropy.org/blog/potential-risks-advanced-artificial-intelligence-philanthropic-opportunity

[CAIS] K. Eric Drexler, “Reframing Superintelligence: Comprehensive AI Services as General Intelligence”, https://www.fhi.ox.ac.uk/reframing/

Quantum Computing and Cryptanalysis

[NIST:WEB] National Institute of Standards, “Post-Quantum Cryptography” https://csrc.nist.gov/Projects/Post-Quantum-Cryptography

[NIST:PQC] L. Chen et al, “Report on Post-Quantum Cryptography”, National Institute of Standards and Technology Internal Report 8105. https://nvlpubs.nist.gov/nistpubs/ir/2016/nist.ir.8105.pdf

[ETSI] European Telecommunications Standards Institute. https://www.etsi.org/

[QSMAC] Dan Boneh & Mark Zhandry, “Quantum-Secure Message Authentication Codes”. https://link.springer.com/chapter/10.1007/978-3-642-38348-9_35

[XQL] Feihu Xu, Bing Qi, Hoi-Kwong Lo, “Experimental demonstration of phase-remapping attack in a practical quantum key distribution system”. https://arxiv.org/abs/1005.2376

[IDQ] IDQuantique. https://www.idquantique.com/

[MAGICQ] MagicQ. https://www.magiqtech.com/

[QLAB] Quintessence Labs. https://www.quintessencelabs.com/

[SQN] SequreNet. https://www.cbinsights.com/company/sequrenet

Quantum computing and simulation

[UQS] Seth Lloyd, “Universal Quantum Simulators”.
http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.654.7909&rep=rep1&type=pdf

Biosecurity

[OPENPHIL:BIO] Open Philanthropy Project, “Biosecurity”. January 2014.
https://www.openphilanthropy.org/research/cause-reports/biosecurity

Atomically Precise Manufacturing

[OPENPHIL:APM] Nick Beckstead, “Risks from Atomically Precise Manufacturing”.
https://www.openphilanthropy.org/research/cause-reports/atomically-precise-manufacturing

[DRX] K. Eric Drexler, “Molecular engineering: An approach to the development of general capabilities for molecular manipulation”. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC348724/

Industry quantum

[IBM] IBM, “Applications of quantum computing”.
https://www.research.ibm.com/ibm-q/learn/quantum-computing-applications/

[GOOGLE] Google, “Google AI Quantum”.
https://ai.google/research/teams/applied-science/quantum/

[MS] Microsoft, “Quantum Computing”.
https://www.microsoft.com/en-us/research/research-area/quantum/

[RIG] Rigetti. https://www.rigetti.com/

[XANADU] Xanadu. https://www.xanadu.ai/

Physics

[SQFTQC] John Preskill, “Simulating quantum field theory with a quantum computer”
https://arxiv.org/abs/1811.10085

[NINTIL] José Luis Ricón, “Is (useful) physics over?”, Nintil (2018-10-03), available at
https://nintil.com/is-useful-physics-over/

[NP+PHYS] Scott Aaronson, “NP-complete Problems and Physical Reality”
https://www.scottaaronson.com/papers/npcomplete.pdf

Chemistry

[CHEM] Kenji Sugisaki et al, “Quantum chemistry on quantum computers: quantum simulations of the time evolution of wave functions under the S2 operator and determination of the spin quantum number S”. https://pubs.rsc.org/en/content/articlelanding/2019/cp/c9cp02546d