

On the Horizon #1 – Quantum Computing


James Sanders – Portfolio Manager
August 2021

A Brief History of Classical Computing


The last few decades have seen transformational technologies emerge that provide a bedrock for innovation across multiple industries. Transistors, their miniaturisation, and their arrangement into complex architectures at micron scales, paved the way for the internet and the digitalisation trends that form the backbone of modern technological civilisation. Behind the many conveniences and efficiencies that have arisen from this progress lie oceans of data and computation that continue to grow exponentially. Moore’s Law has largely remained intact since Gordon Moore’s updated forecast in 1975, and has expanded beyond transistors to many other areas of electronics, such as sensors, RAM, flash memory, and the pixel count in digital cameras.

Software has expanded in synchronicity with hardware, a process referred to glibly in the 90s as Andy and Bill’s Law – or, ‘what Andy giveth, Bill taketh away’ (Andy Grove was CEO of Intel during much of Bill Gates’ tenure as CEO of Microsoft). This endless hardware/software upgrade cycle has generated huge economic value over the last three decades (and some fantastic stock market returns), but also enabled incredible growth in computational capacity leading to the commercialisation of, and broad practical applications for, machine learning and artificial intelligence. Nvidia CEO Jensen Huang has since posited ‘Huang’s Law’, which adds GPU architecture and algorithmic efficiency to the mix to keep the march higher in computational capability on an exponential path, even as silicon transistors approach the limits of physics (Taiwan Semiconductor Manufacturing Co. is investing in the development of 2 nanometre transistors, just 10 times the size of a silicon atom).

Source Lines of Code (000s). Google’s codebase across all services comprises 2 billion SLOCs!



Moore’s Law (Transistors per Microprocessor, MMs) Log base 10 scale




Source: Our World in Data


What is a Quantum Computer?

The possibility of quantum computers was famously raised by physicist Richard Feynman in May 1981, when he spoke at MIT on the topic of “Simulating Physics with Computers”. In his talk, he argued no classical computer would be up to the task, concluding, “Nature isn’t classical, dammit, and if you want to make a simulation of Nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem because it doesn’t look so easy.” He has since been proven right – while techniques have been developed that allow classical supercomputers to approximate the behaviours of large numbers of atoms and complex molecules, the quantum-mechanical interactions between electrons cannot be modelled accurately, and as a result certain ‘grand challenge’ problems in chemistry, biology and materials science still elude researchers.





The white netting around this buckyball (C_{60}) shows the molecule's ground-state electron density, calculated using density function theory (DFT), a classical computing method developed in the 1960s




Caffeine has just 24 atoms ($C_8H_{10}N_4O_2$), and yet cannot be fully simulated by a classical computer. This is because there are around a 100 electrons, and each electron roughly doubles the computational difficulty

Where classical computers use bits, the fundamental unit of a quantum computer is a qubit. The key difference is that while the former represents either a 0 or a 1, a qubit can store the values 0, 1, and anything in between (e.g. 80% 0/20% 1). Qubits can also make use of quantum entanglement to be linked to one another. These difficult to intuit and strange qualities make quantum computers extremely well suited to certain classes of problem. A (somewhat) accessible explanation using the analogy of a drunken walk through a bar can be found [HERE](#).

Different Approaches

The first qubit's and quantum logic gates were created in the 1990s, and in the same decade physicist Peter Shor developed 'Shor's Algorithm', showing a quantum computer could in theory factor large numbers many orders of magnitude faster than classical computers. Theoretical progress continued in the academic world, but it has only been in the last 10-15 years that serious efforts have been made to commercialise the technology, and only in the last two years that headline-grabbing results have been achieved relating to real-world applications.





There are a number of different approaches to creating scalable quantum computers, as companies vie to gain the first mover advantage. Many tout more and more numerous qubits in their systems, though it is important to note that error rates can be as high as 99.5%, so the number of useful qubits in any of today's systems is still very small. In quantum computing, adding a single error-corrected qubit to a system doubles computational capacity. The industry estimates around 72 error-corrected or algorithmic qubits could outperform the largest classical supercomputers at certain tasks, as well as tackle problems that are beyond them. As with classical computers, software will also play an important role in improving the error rate and reducing the number of qubits required to achieve a given solution.

IBM – IBM has been working on building quantum computers for two decades and announced a 65 qubit system last year. They aim to release a >1,000 qubit system by 2023. Their approach requires cooling close to absolute zero using dilution refrigerators, and creating artificial atoms called superconducting transmon qubits. Any interaction from the outside world through electromagnetism or vibration (even cosmic rays) can cause their qubits to decohere, and calculations can only be run for a few milliseconds. Consequently they need to built a system with a large number of qubits to end up with sufficient useful qubits to apply to practical problems (and find a way to cool them – they are currently custom-designing the world's largest dilution fridge). IBM has built about 20 quantum computers, and the smaller ones can be used for free in the Cloud, as IBM tries to encourage the next generation of quantum programmers to come forward. IBM's IQX platform has had more than 120,000 users running more than 12 million experiments, and has produced more than 180 scientific papers. Commercialisation will likely take the form of Quantum as a Service (QaaS) over the Cloud – the nature of the system will make it unsuitable for on premises enterprise use. This may preclude use cases requiring low latency (such as time-sensitive financial services applications for example).


Development Roadmap

IBM Quantum





GOOGL – Alphabet announced they had achieved quantum supremacy in October 2019, claiming they had solved a problem on their 54 qubit Sycamore quantum computer in 3 minutes 20 seconds that would take a classical computer with 1 million processing units (equivalent to 100,000 desktop computers) 10,000 years to solve. This was disputed by IBM, who said it could be completed by the hypothesised classical computer in 2.5 days (not peer reviewed), and many others took issue with the highly esoteric problem chosen (involving random numbers). Alphabet's quantum computer is also based on an extremely cold, superconducting system, with all the associated scaling and coherence problems. Last month they announced a breakthrough in error correction, though they need further breakthroughs to achieve the fault tolerance threshold required - their paper suggests a practical quantum computer of this kind might need 1,000 to 10,000 error-corrected qubits for each logical qubit. Google are targeting a 'useful, error-corrected quantum computer' by the end of the decade, aiming for a 1,000,000 qubit system by 2029. Like IBM, commercialisation will be in the form of Quantum as a Service (QaaS) over the Cloud – the nature of the system may make it unsuitable for on premises enterprise use.



Google's Sycamore quantum computer (54 qubits). It is hard to imagine what one that has been scaled 18,500x to 1,000,000 qubits might look like!


PsiQuantum (private) – PsiQuantum aims to build a >1,000,000 quantum computer by 2025 that will have sufficient error correction and fault tolerance metrics to offer 'hundreds' of logical qubits. Their system is based on photonics which has a number of advantages over the superconducting route: 1) most of it can operate at room temperature (just need to cool the superconducting single photon detector), 2) it can interface with optical semiconductors and fibreoptics quite easily which are mature technologies and therefore easy to scale, 3) photons occur naturally and are perfectly identical, so no need to 'create' artificial qubits. The photonic approach does have issues that are still looking for technical solutions – for example, photons travel at the speed of light which presents problems with storing quantum memory. PsiQuantum has a plan for overcoming this, and has partnered with Global Foundries to scale their system. The company recently raised \$450MM in a Series D round at a \$3.15BN valuation.

Xanadu Quantum (private) – This Canadian company is also taking the photonic approach, and has made their latest 8 qubit chip available on the Cloud. Made of silicon nitride, this also uses mature semiconductor knowhow, and can operate at room temperature which the company says will help them scale. They are also building an



open-source, hardware agnostic software platform for quantum machine learning. Xanadu recently raised \$100MM in a Series B, at a \$400MM valuation.

HON – Honeywell announced their first quantum computer in June 2020, using trapped Ytterbium-171 (Yb) ions to serve as qubits. Like the photonic approach, these have the benefit of being naturally occurring and identical, with the additional benefit of being much easier to ‘trap’ compared to a photon. The trapped ion approach has higher fidelity than either of the previous two approaches, and the physics has been more or less proven over the last 25 years, leaving scaling to engineers to a certain extent. Their latest 10 qubit system is already being tested by paying subscription customers (which include DHL, Merck, and JPMorgan Chase). Their approach involves physically shuttling qubits around the chip (a quantum charged coupled device), which poses some engineering challenges as they scale to dozens or hundreds of ions. Whether or not this can be overcome will become apparent when they attempt junction transport in their 3rd generation chip (referencing Honeywell’s generational roadmap below).



IonQ (DMYI) – IonQ traces back to a similar lineage as Honeywell’s quantum computer, as IonQ co-founder Chris Monroe demonstrated the first quantum gate in 1995 and co-authored the 2002 paper on a quantum charged coupled device (QCCD) with David Wineland (who went on to win a Nobel Prize in 2012). IonQ also uses trapped Ytterbium ions, held in 3D space by 100 tiny electrodes. A key difference between IonQ’s approach and Honeywell’s is that in IonQ’s system all qubits are linked to each other, so no shuttling or physically wired connections are required (this is achieved through electromagnetic repulsion mediated by laser pulses). IonQ has an 11 qubit chip available for software experimentation and testing on Microsoft Azure Quantum, Amazon Braket, and Google Cloud. Their 5th generation chip has 32 qubits (22 algorithmic qubits) and will be available on three cloud platforms in the coming months. IonQ counts Google





An IonQ ion trap and vacuum chamber



Ventures, Amazon, Samsung, Lockheed Martin, Hewlett Packard, and Bosch amongst their early investors (some of whom are also clients). More recently, Softbank took a stake in the company to actively expand partnerships in Asia and Japan. Earlier this year, IonQ announced a plan to list on the NYSE through dMY Technology Group III (DMYI), a publicly traded SPAC. The rationale was to raise all of the anticipated funding needs for their commercialisation roadmap in one go (\$650MM including the PIPE), to raise their public profile, and to have listed shares to use as currency to attract the best talent going forwards. The funds were raised at a \$2BN equity valuation (\$1.4BN ex cash).

Rising Prominence


The race to develop and deploy useful quantum computers has taken off in the last couple of years, with superpowers such as the US and China politicising ‘supremacy’ in the technology and some of the largest companies in the world ploughing billions of dollars into the effort. When I first wrote on the topic in September 2017, only Google and IBM really featured in the press; since then there has been an explosion in quantum start-ups both on the hardware and software side:






Who Could Create Value with Quantum Computing?

Distribution of quantum-computing use cases, 2019, %



Estimated value at stake¹



McKinsey estimate the quantum computing market could reach \$1 trillion by 2035. This could prove conservative - I am reminded of the President of IBM's forecasts in 1943: "I think there is a world market for maybe five computers."

¹ Approximate timing for medium term is by the year 2025; for long term, by the year 2035. Experts consider these values at stake to be a snapshot in time. Fully developed quantum computing will lead to additional value within and shifts between industry verticals. Source: Expert interviews; McKinsey analysis

Practical Applications

In May this year, US climate envoy John Kerry said, "Fifty per cent of the reductions we have to make to get to net-zero by 2050 or 2045 are going to come from technologies that we don't yet have. That's part of the challenge. But look at what we did to push the creation of vaccines, look at what we did to go to the moon, look at what we did to invent the internet." Achieving net-zero carbon emissions by 2050, while maintaining our standard of living, will be one of the greatest challenges faced by modern civilisation. Quantum computing can make a significant contribution to this effort on a number of fronts:

Renewable Energy – Today's commercial photovoltaic cells have 15-20% efficiency, which would require 0.6% of the landmass of the US to meet total electricity demand² (21,000 square miles, equivalent to the whole of West Virginia). PV efficiency of >50% has been achieved in the lab, but insights from a quantum computer could drive efficiency to the theoretical maximum (dendrimer molecules could potentially be designed to harvest light at near 100% efficiency). The result of this would be only the states of Delaware (2,500sqm) and Rhode Island (1,500sqm) need to be covered in solar panels!

Artificial Photosynthesis – The ability to model and understand this process fully would conquer a grand challenge of science and would open the door to harnessing photosynthesis for solar fuel (storing solar energy in chemical bonds), natural carbon fixation, and photocatalytic water splitting to generate biohydrogen.

Nitrogen Fixing – The Haber-Bosch process is widely used as a an artificial nitrogen fixation method in the production of fertilizers. It requires a blend of nitrogen and hydrogen (usually derived from natural gas) to be compressed to 200-400 atmospheres and heated to 450-650⁰ C. This process is extremely important in agriculture, however it is also very energy intensive and has a huge carbon footprint (up 3 tonnes of carbon dioxide is released for every tonne of ammonia produced). As a result, it is estimated to account for around 1.2%

² <https://css.umich.edu/factsheets/photovoltaic-energy-factsheet>



of anthropogenic CO₂ emissions (about equivalent to the CO₂ output of the UK). Plants achieve nitrogen fixing through symbiotic relationships with microorganisms, and can do it sitting in a plant pot on your windowsill with just sunlight to power the process. Quantum computing may help us replicate the efficiency found in nature, and advance us beyond the Haber-Bosch process which was conceived over a century ago during World War I.

Battery Chemistry – improving the energy density of batteries is mission critical for electrification and the build out of renewable energy capacity. Better cathodes in lithium-ion batteries, different/cheaper chemical compositions, novel separator materials for anode-less/solid state Li batteries could all help make cheaper, lighter, safer, and higher capacity batteries for mobility and energy storage. Similar improvements could be found for hydrogen fuel cells, perhaps through better proton-exchange membranes that allow for lower (and therefore much cheaper) purities of hydrogen as fuel.

Nuclear Fusion – one of the problems (there are others) holding back a commercial fusion reactor are the difficulties in modelling plasma fluid dynamics. Quantum computers could model these interactions which would help guide design of the optimal toroid cross-section and the shape of the magnetic containment fields. Currently this is done almost by informed trial and error, and given the decade long lead times between new experimental reactors, this goes some way to explain why fusion has remained elusive for so long.

Aside from the important matter of climate change, other commercial applications include:

Quantum Machine Learning – The increasingly capable library of machine learning algorithms and deep learning through neural networks will likely continue to advance and evolve alongside quantum computing, which is years behind in terms of software development. IonQ recently outperformed classical machine learning algorithms when it comes to human handwriting recognition, when trained on the same dataset, though they achieved only about 30% of the success rate on image recognition tasks (this was on their 11 qubit chip, will be interesting to see what the 32 qubit chip can do). In the next two years we may see a quantum advantage in machine learning tasks. To give a real world example, Google and DeepMind used ML techniques to achieve a -40% reduction in energy used cooling Google's datacentres. They still consume \$500MM in energy per annum, so further optimisations are worth millions of dollars. The same goes for optimisations in other areas, such as digital advertising.

Logistics – The travelling salesman algorithm is famously computationally intensive. It seems simple: start with a set of cities and find a route that goes through each city once, ending where it started. The problem gets much more complex when it tries to solve for the shortest among all possible routes. One solution is to map every round-trip route and find the shortest ones. The more cities, however, the more round-trips to map. Ten cities mean more than 3.6 million possibilities. For sixteen cities this rises to more than 20 trillion. If stops were cities, a delivery drover would visit at least 100 cities a day, and calculating the perfect route would be practically impossible. Logistics company's use classical computers to optimise their routes (UPS estimate theirs saves them 100MM miles/\$250MM per year on their 66,000 routes per day), but they are still just approximations. The same problems are faced by the airline industry, and future smart cities may need sophisticated routing capabilities to control millions of connected mobility vehicles, from scooters, to autonomous cars and VTOL passenger craft navigating around each other in three dimensions.

Financial Services – Quantitative hedge funds, arbitrageurs, and real-time risk management systems would all benefit from quantum computing. This is an area of active interest, with multiple initiatives ongoing in financial services to see who quantum computers might improve portfolio optimisation, reduce fraud, and better quantify risk.



Cryptography – RSA encryption protects much of the internet and the financial system. RSA-2048 bit encryption keys would take a classical computer over 300 trillion years to break (so not many have bothered to try). This type of encryption will likely fall to quantum computers some time in the 2030s (elliptical curve cryptography will fall sooner due to the smaller key size), however work is already underway to replace these types of cryptography with quantum alternatives. Bitcoin HODLers should be monitoring developments in this area closely, as their encryption keys may not be as safe as they thought.

Quantum Database Search – For the same reasons Quantum computers are suited to cryptography, they are also able to search large, unsorted datasets very quickly. Classical computers can take a period of time proportional to the number of elements in the search, potentially having to search through all the elements one by one to arrive at the answer. Using Grover’s algorithm, quantum computers can return a query in a timespan proportional to the square root of the number of elements (quadratically faster).

Drug Discovery – As highlighted in previous notes, innovation in medicine has stagnated for years, and may be on the cusp of a renaissance. Part of this will be driven by the new tools available to researchers through advancements in genomic sequencing and editing, and our recent breakthroughs in proteomics. These will continue to drive breakthroughs in the coming years, however quantum computing will provide tailwind to accelerate this even further. We are in the early stages of zeroing in on the mechanism of senescence – if the biochemistry behind this final frontier falls to quantum computing by the next decade, what does it mean for the longevity of everyone alive today?

In Summary

The development of useful quantum computers has been a slow and slightly niche calling over the last three decades, latterly outshone by the billions of dollars of value created by real world applications for machine learning and AI, which continue to grow at an exponential rate, and extend into all sectors of the economy. These technologies will continue to evolve for years to come, and quantum computers will likely converge into a symbiotic relationship with this complex computational infrastructure. Equally, there is a place for several different types of quantum computer, as each will have strengths and weaknesses, and have their own set of useful applications. Many of these applications have not yet been envisaged, nor can we predict what might be possible once we weave these new systems into our existing supercomputers, datacentres, and the Cloud. However this takes shape, we can say with confidence that we are entering the quantum age, and this fundamental technology has the potential to accelerate civilisation-shaping change in the coming years.





LEGAL DISCLOSURES

FOR EU INVESTORS

The information contained in this document is issued by Green Ash Partners LLP (hereinafter "Green Ash"), at 11 Albemarle Street, Mayfair, London, W1S 4HH who is authorised and regulated by the Financial Conduct Authority: Firm Reference Number (FRN) – 500315. This disclaimer is not intended to exclude or restrict any liability under the rules of the Financial Conduct Authority ("the FCA Rules") or FSMA.

This presentation reflects the opinion of Green Ash as of the date of issue. This document is not intended for distribution, publication, or use in any jurisdiction where such distribution, publication, or use would be unlawful, nor is it directed to any person or entity to which it would be unlawful to direct such a document. This presentation is for information purposes only and does not constitute an offer or a recommendation to purchase or sell any security. It does not constitute investment research or a research recommendation and is not intended for distribution to the public or a large number of persons. The opinions herein do not take into account individual clients' circumstances, objectives, or needs. Before entering into any transaction, each client is urged to consider the suitability of the transaction to his particular circumstances and to independently review, with professional advisors as necessary, the specific risks incurred, in particular at the financial, regulatory, and tax levels.

All examples of financial strategies/investments set out in this document are for illustrative purposes only and do not represent future performance. The information and analysis contained herein have been based on sources believed to be reliable. However, Green Ash does not guarantee their timeliness, accuracy, or completeness, nor does it accept any liability for any loss or damage resulting from their use. All information and opinions as well as the prices indicated are subject to change without notice. Past performance is no guarantee of current or future returns and you may consequently get back less than he invested. From time to time the partners of Green Ash Partners LLP may enter into personal transactions in the securities and strategies discussed in this presentation. The firm has a personal account dealing policy that manages such conflicts and ensures the fair treatment of its clients.

This document may not be reproduced (in whole or in part), transmitted, modified, or used for any public or commercial purpose without the prior written permission of Green Ash.

FOR SWISS INVESTORS

This document is issued by Green Ash Partners LLP. This document is exclusively intended for qualified investors within the meaning of article 10 paragraph 3, 3bis, 3ter and 4 of the Swiss Collective Investment Schemes Act ("CISA"). The information and data presented in this document are not to be considered as an offer or solicitation to buy, sell or subscribe to any securities or financial instruments. Information, opinions and estimates expressed in this document reflect a judgment at its original date of publication and are subject to change without notice. Green Ash Partners LLP has not taken any steps to ensure that the securities referred to in this document are suitable for any particular investor and this document is not to be relied upon in substitution for the exercise of independent judgment. Tax treatment depends on the individual circumstances of each investor and may be subject to change in the future. You should obtain specific professional advice before making any investment decision. The value and income of any of the securities or financial instruments mentioned in this document can fall as well as rise and, as a consequence, investors may receive back less than originally invested. Risk factors are not intended to be reproduced in full in this document. Past performance is no guarantee nor a reliable indicator of future results. This material is not intended to be a substitute to the full, legal documentation and to any information which investors must obtain from their financial intermediaries acting in relation to their investment in the funds mentioned in this document. Please note that none of the management company, the registrar and transfer agent, the central administration or the custodian of the relevant fund has independently verified any information contained herein and no party makes any representation or warranty as to the accuracy, completeness, or reliability of such information. Further information about Woodman SICAV - Green Ash Global High Yield Fund, its prospectus, its KIIDs and its latest annual and semi-annual report may be obtained free of charge, in English language, from MultiConcept Fund Management S.A., Luxembourg, the CH Legal Representative, Credit Suisse Funds AG, the appointed distributors or online at www.greenash-partners.com and www.credit-suisse.com/ Multiconcept.

