



Is Quantum Computing worth the effort?



George Woodman
Quantum Lead

Quantum Computing (QC) is arguably one of the most promising research fields currently being investigated. But the very nature of this fascinating piece of technology raises questions. The various mechanisms underpinning quantum computing are quite abstract, making it very difficult to imagine how it could solve concrete problems. Is it therefore possible to leverage this technology to drive a competitive advantage – especially when we factor in the required amount of research? Simply put, is quantum worth the effort?

Granted, quantum theory is fairly complex. However, even though it can seem almost impossible for non-scientists to wrap their heads around quantum computing, one doesn't need a background in Quantum Information Theory or Computational Science to appreciate the improvement quantum computing can bring. For a specific set of problems, a quantum computer has the potential to provide a polynomial or in some cases exponential speedup to current hard computational problems. In laymen's terms, this means that as the problem size (e.g. number of inputs) increases, the speedup compared to a classical machine increases even more. Meaning that a problem which could be only solved on a very small scale today, could in the future be solved in a real scenario thanks to QC.

Optimization and Monte Carlo simulations

But what kind of problem could possibly be fit for this kind of solution? For instance, finding prime factors for large numbers. Prime factorization is very useful in internet security protocols, making QC a potential threat for cybersecurity. In general, most concrete use cases fall into a few buckets: chemistry simulation problems, optimization problems, factoring problems, and linear algebra applications for machine learning.

Optimization aims at determining the maximum or minimum value of a mathematical function. This function can be built in an endless number of ways which allows many problems to be mapped to an optimization problem. The most famous example is the traveling salesman problem – where a merchant must travel to several cities to sell their product and want to determine the shortest possible route to visit each city once. Mapped specifically for finance, this could be a way to create a portfolio trajectory whose collection of assets minimize risk and maximize expected return whilst balancing out covariance between them.

Another insurance relevant focus area is Monte Carlo simulations, which involve a function with a random variable, from which one needs to iterate a simulation million or even billions of times to form a distribution. They typically require enormous classical computing power to provide an accurate solution.

Solving “real life” problems

Numerous challenges could theoretically benefit from the calculating power offered by Quantum Computing, many of them in the financial field. For instance, due to Basel III regulations, financial institutions are required to run stress tests using Value at Risk (VaR) to determine the potential for loss. This can be

calculated using Monte Carlo simulations, and it has been proven that this calculation could be quadratically¹ sped up with a Fault Tolerant Quantum Computing.

In terms of asset allocation, quantum could help maximize returns and minimize risks by optimizing a portfolio. In the field of natural catastrophes, quantum could drastically speed up modelling of dangerous events to better predict them, allowing insurers to use this information to price bonds efficiently. Operations such as Value at Risk or Portfolio Risk could be improved by fine-tuning the potential loss predictions. Hypothetically, quantum can also have concrete applications in Nested Modelling and Pricing, with the potential to improve the accuracy of binary classifiers in credit scoring.

The limits of Quantum Computing

Although promising, the field of Quantum computing still faces challenges. One of the main issues Quantum computers face is related to the amount of stable computing power required to execute some algorithms. This is something that is still a few years away due to the immaturity of the technology. These limits can be split into five sections.

Processing Power: The most powerful Universal Quantum Computer only has physical 65 qubits today. An estimated 1.6 million is required for financial advantage.

Noise: Quantum states are inherently unstable; any outside interaction, such as

heat, can destabilize them. Fully correcting errors might be one of the most important improvements for quantum computing in the years to come.

Coherence Time: Due to a number of factors, including noise, there is a minute time that quantum states can be in a state useful for computing. This time is a strict limit to the amount of available computing power.

Entangled States: Some of the true power of quantum comes from the fact states can be entangled. However, entangling multiple states coherently is very difficult to achieve and maintain.

Memory: Currently, quantum states can't be saved and loaded up later. There is an area of research focusing on QRAM (Quantum Random Access Memory) to derive a way to hold quantum states for use later in quantum programs.

Should we start now?

We are in the NISQ era (Noisy Intermediate-Scale Quantum), but there are still potential applications that could prove useful in this nascent form of the technology. These use cases are expected to be simple problems that require a high degree of accuracy. Mostly pointing towards asset allocation within investment management as their potential quantum algorithms were designed to provide a speedup in the near-term. Be that on a Quantum Annealer or a Universal Quantum machine. The process we go through to run these use cases no matter how small will be useful in the future.

The limit of the problem size will only loosen as the hardware improves.

Quantum computing is a complex topic and one that requires a lot of work to implement into a production setting. But, the reality with quantum computing boils down to the fact that the technology isn't an 'if' question but a 'when' so preparing for it should be a priority. Quantum does require significant work over a long period of time – but a financial company that wants to stay relevant in the long run cannot afford not to put in this effort. When to focus these efforts, where to implement them and how much effort to apply is still up for debate.

Contacts



George Woodman

Quantum Lead

george.woodman@axa.com



Marcin Detyniecki

Head of Research and Development & Group Chief Data Scientist

marcin.detyniecki@axa.com

And tomorrow?

The biggest question usually asked relating to quantum computing is: when will the technology be ready? The short answer in our current timeline is that there should be a quantum application in production bringing real business value in the next 5-10 years. The long answer as with all things is more complex.

It completely depends on the problem one wants to use the quantum computer for. For instance, solving an optimization problem with binary variables and linear constraints could benefit from the technology as soon as three years from now. Some problems could take much more time.

At the same time, classical algorithms are constantly getting improved so that the bar for QC to reach so that it can run faster is growing ever higher. But there is no doubt that once a usable application of quantum computing is proven to run faster than a classical computer, it is only matter of time, following the exponential growth, before we see occurrences dubbed 'Quantum Advantage'.

In a few years, the idea of hybrid machines leveraging both classical and quantum computers could also shape the future of computing. Since quantum computing will only improve certain parts of problems, all quantum computing will likely be hybrid.

3 questions to...



Jordanis Kerenidis

Head of Quantum Algorithms
at QC Ware Corp and Research
Director at the CNRS

Why does it make sense to get involved in quantum computing today?

“Quantum computing has the potential to revolutionize many sectors through various applications - especially in finance. This will take some time, because the hardware is not yet powerful enough to be applied at an industrial level. But one cannot just wait for the hardware to arrive. To be ready when the hardware gets there, we must start understanding how to use it now. It takes significant time and effort to be able to understand how quantum algorithms work and learn how to develop applications for these quantum machines. This is why, as an insurance policy, it is important to start working on it early on,. Starting to engage now is the best way to be ready and to integrate it in the workflow quickly and easily in the future.”

What are the main challenges for implementing quantum computing in a company such as AXA?

“The important thing to realize is that quantum computing technology is an extremely different way of computing. A quantum computer is not just a faster processor: to get the most out of it, one needs to completely change paradigm

and the type of algorithms one is running. Therefore, one of the main challenges for companies today is getting the know-how and creating intellectual property: producing new tailor-made algorithms and models and patenting them. In a few years, the quantum hardware will be accessible to everyone, like GPU clusters today – then, the companies with the competitive edge will be the ones who have developed and own these novel algorithms and models.”

What use cases do you see as being the first to bring value?

“Monte Carlo simulations, optimization and Machine Learning-related applications come to mind. For each of them, there are very good chances to see real use cases in the years to come. Optimization use cases, which can be extremely useful and impactful, might be the first ones to arrive, also because there is a lot of expertise on how to solve this kind of problems with quantum computers. Machine Learning and Monte Carlo simulations used to require much more computing power, but we recently managed to reduce the needed resources, which made it possible to see concrete use cases in the shorter term. It might still be a few years because the datasets are quite large, but we are getting there faster than initially thought.”

