Quantum technology: Denmark's next business adventure?

White paper
March 2021

https://home.kpmg/dk/da/home.html

Table of contents

Executive summary

Quantum technology has tremendous potential, and early adopters are piloting quantum-inspired algorithms to gain experience and position themselves for a competitive advantage as the technology evolves. In this project, the KPMG NewTech quantum team and the Technical University of Denmark (DTU) explored the potential of quantum computing by implementing a Markowitz model for stock portfolio optimisation on the 2000Q D-Wave quantum annealer. We successfully embedded the model in the D-Wave machine and performed a portfolio optimisation simulation with 65 stocks, benchmarking the results against two classical algorithms: brute force and simulated annealing.

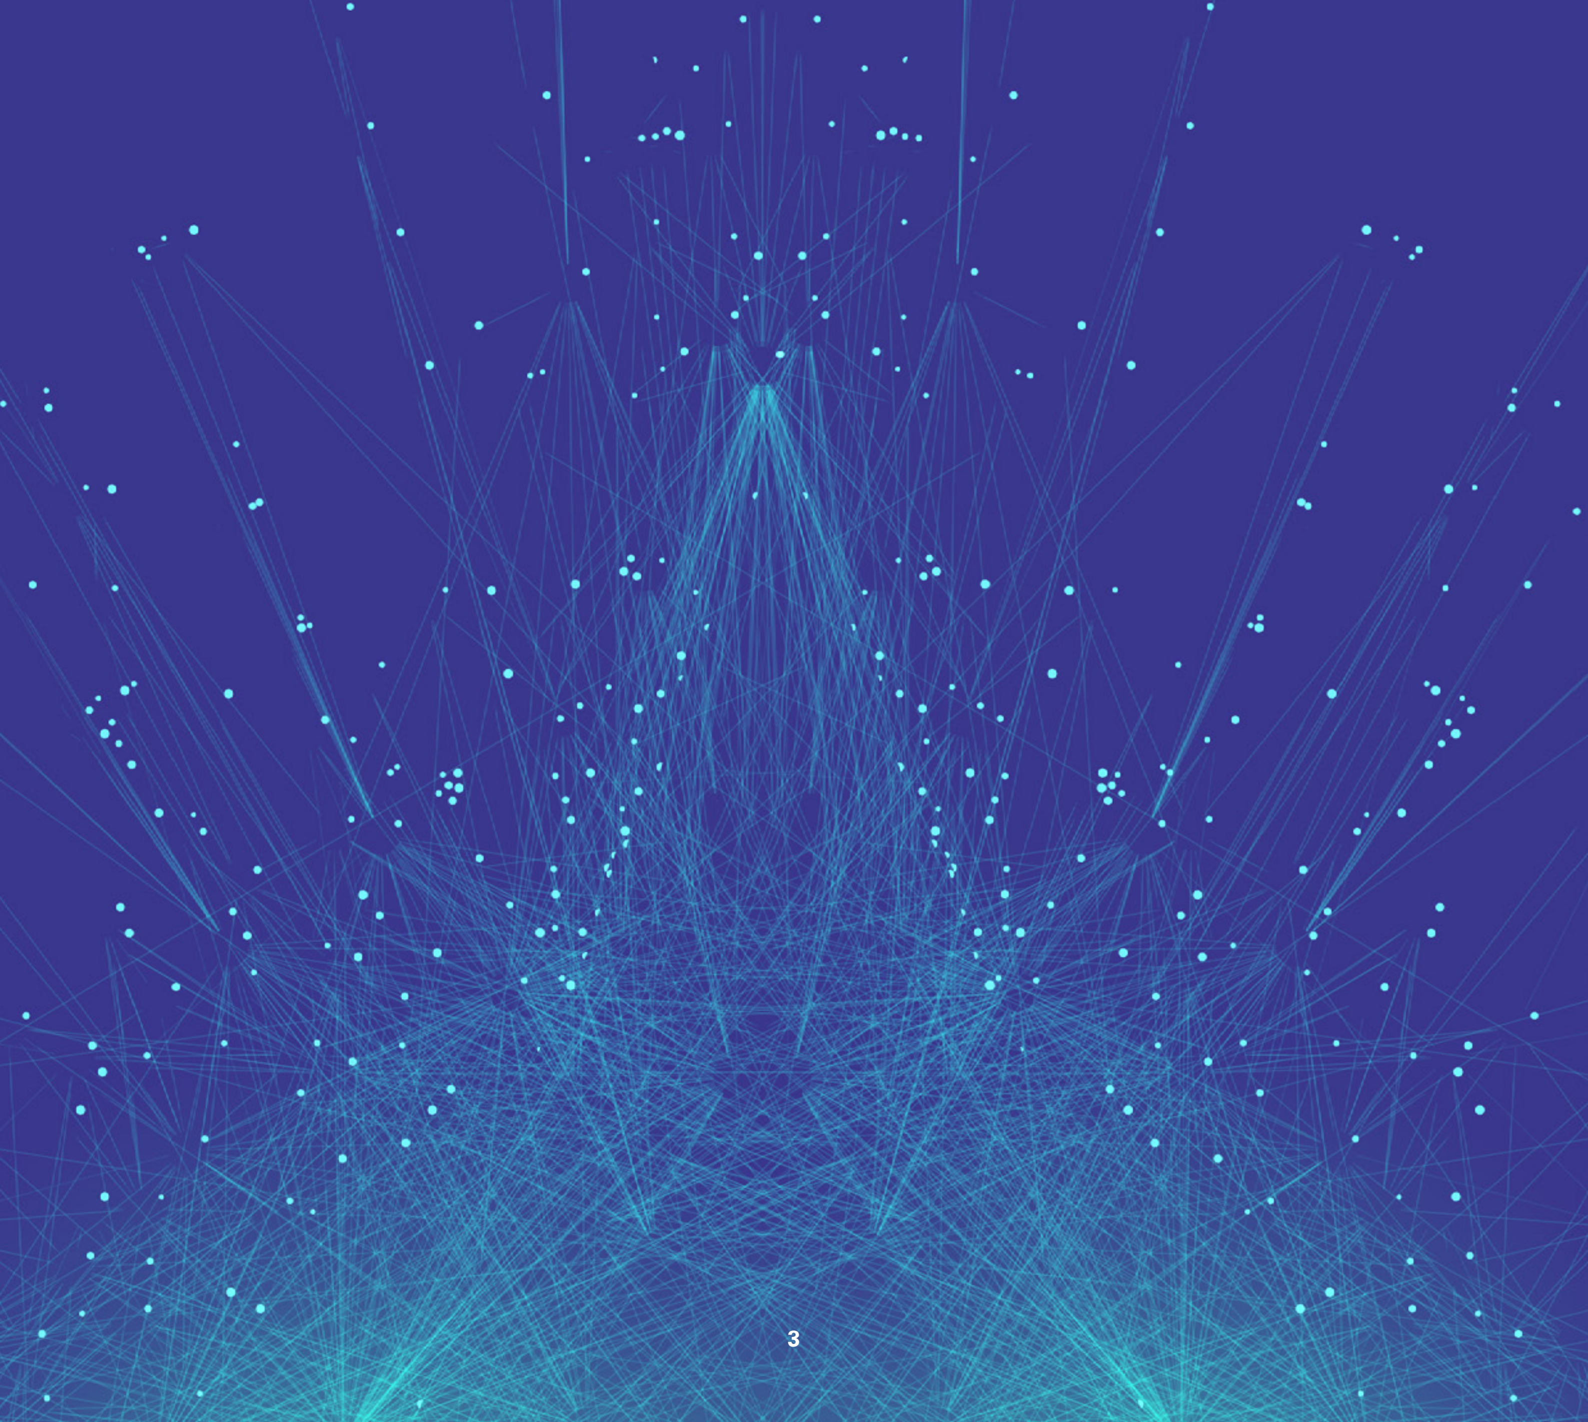

Quantum technology

Benefits and threats to the financial industry

A significant concern is quantum computing's potential to break the cybersecurity standards that we rely on every day when we connect to the internet. If the security of a company's infrastructure is jeopardised with the new technological wave, then industrial secrets and privacy will no longer be protected. It all depends on how prepared your infrastructure is for this threat, how long it will take you to make it secure and if your infrastructure is already now endangered, for example to 'download now, decrypt later' attacks. These are all concerns that should be on executives' minds, as many companies are already mobilising to understand the threat or even implementing prototype countermeasures to secure their assets.

On the positive side, this technology can bring advantages in many processes, from machine learning to new optimisation methods, improvements in drug and chemical discovery, better medical imaging, and many other applications.

Of course, this is a new and complicated technology, and these applications are not easy to develop. Therefore, many companies are already developing algorithms and software that can run on quantum computers and are securing Intellectual Property Rights in anticipation of quantum computers soon being powerful enough to bring a significant competitive advantage.

The first step that every company should take at this stage of the technology is to create a solid understanding of the technology and its potential. Building a fact-based overview of the technology together with potential applications and uses within your firm, including within your own infrastructure environment, is a smart way to position yourself and harvest the full potential of the technology. KPMG can assist in this journey, from providing training to working side by side with your professionals to identify use cases and develop a Proof of Concept, helping you translate these new scientific advantages into business opportunities.

Quantum computing:

a new paradigm of computation

Many companies and organisations are working towards a functioning fault-tolerant quantum computer. IBM and Google are competing to create a superconducting quantum computer, Honeywell and IonQ are betting on trapped ion quantum computers, Xanadu is exploiting the power of light and photons, and Microsoft is working to create a topological quantum machine. These are only some of the players racing to achieve a quantum advantage through a universal quantum computer. The players are many, the technology and architectures are different and based on different materials, but in this project, we focused specifically on a quantum solver brought to market by D-Wave: the quantum annealer. One of the main reasons for our interest in this special purpose quantum device is the maturity of the technology and the early use cases implemented by several companies like Volkswagen or TIM.

The machine's quantum nature has been under discussion for quite some time, mainly because the D-Wave works differently from universal quantum computers, and for this reason, we refer to it as a quantum simulator or quantum annealer.

The first item that is essential to understand is that the quantum bits ("qubits") from D-Wave are different from the qubits of gate-based computers. A qubit is the fundamental unit of information for both systems, but in the gate-based computer, you can manipulate these units using a set of universal gates, whereas in the quantum annealer you do not have this option. Instead, you modify the environment and let the system find the ground state.
The advantage that the annealer brings is the possibility of exploiting a phenomenon called quantum tunnelling. Imagine a landscape with hills and holes, and every time you perform a computation, it is like dropping a ball into the landscape and seeing where the ball stops. Naturally, we expect to find the ball in the deepest hole (the ground state), but the advantage over classical annealing lies in the fact that our ball can "dig tunnels" between the hills while searching for the deepest hole, thus increasing the probability of hitting the ground state The D-Wave is a probabilistic machine – experiments are executed a set number of times (this is a configurable parameter), and the outcome with the highest probability is (probably) the solution to the problem.

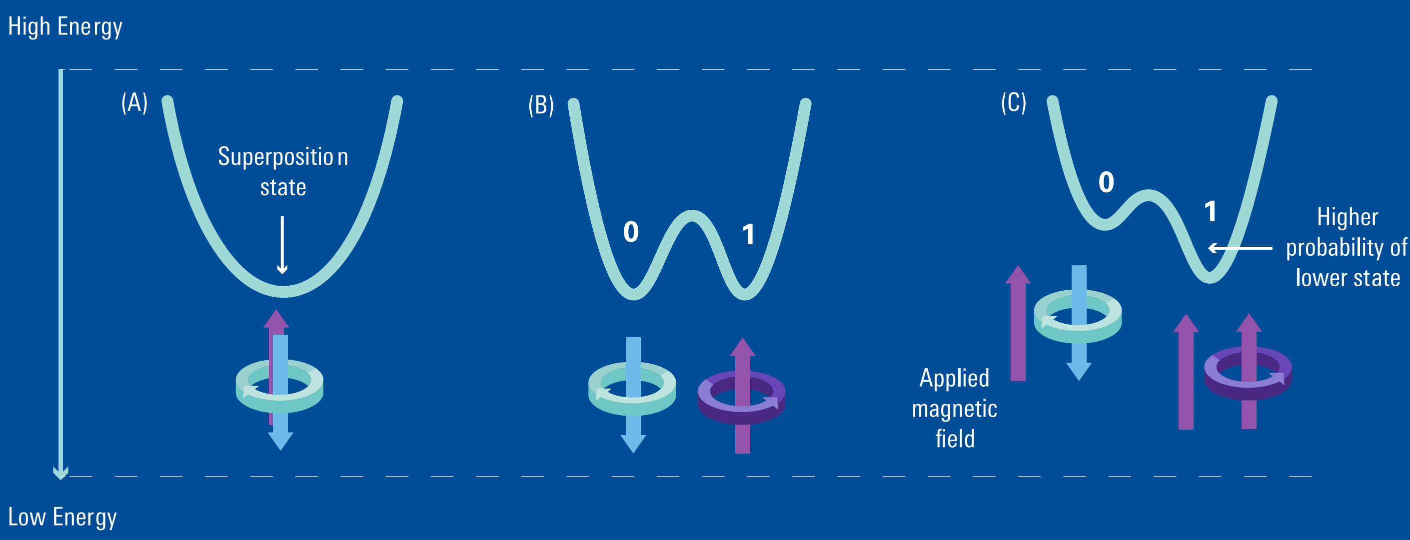

Here we can see the D-Wave process. We start from a very simple landscape (a) and we modify it through the application of magnetic fields. In (b), if we imagine dropping a ball from the middle, we end up with a 50% probability in both holes. In fact, both holes are a minimum energy state for the problem. On the other hand, in (c), we expect the ball to fall in the deepest hole, representing “1” and this will indeed happen with higher probability.

The quantum annealer is designed to handle specific optimisation problems, including Ising and Quadratic Unconstrained Binary Optimisation (QUBO) models. Therefore, whenever you want to embed a problem into D-Wave, you need to translate it into a QUBO model. This is not an easy task, and many studies have been conducted on these reductions.

The quantum annealer has a promising range of applications, from molecule design and protein folding to traffic and logistic optimisation and radio cell optimisation for 5G networks. In financial services, potential use cases include X-Value Adjustment (XVA) and risk modelling. We chose to analyse the potential of D-Wave for portfolio optimisation.

A promising option for near-term solutions

Quantum technology exploits effects like superposition, entanglement and tunnelling that are properties of the quantum world. This new technological wave promises to bring advantages in many fields, with potential applications usually divided into three categories: computation, communication and sensing.

- Quantum computing seeks to take advantage of properties from the quantum world to speed up computations and tackle problems currently intractable with classical computers.
- Quantum communication plays a fundamental role in both secure communication and quantum information transmission.
- Quantum sensors exploit quantum effects to accurately detect tiny changes in speed, gravity and electric or magnetic fields, bringing more accurate medical imaging, radars or accelerometers.

While significant improvements have been made regarding ease of access to the quantum annealer due to the new Ocean Software Development Kit, identifying problems that are a good fit for the D-Wave machine is critical, and implementing the models requires a new way of thinking.

In addition to formulating a problem as a QUBO model, the optimisation of technical parameters, like chain strength or annealing time, is very important to obtain good performance. We anticipate that many tricks and techniques are yet to be developed in this field. Imagine the classical computer science community, where thousands of experts constantly work on optimisations. In a new field such as quantum computing, resources and time will be needed to develop a similar knowledge level.

While we do not believe that the D-Wave will have the same impact as fault-tolerant gate-based quantum computers, it can be a promising option for near-term solutions, and further improvements in the hardware will improve the quality of the solution and address current limits to scalability. We continue to explore this simulator’s potential in other fields, and with the newly released ‘Pegasus’ architecture, we hope to see many more algorithms running across industries.

Portfolio optimization using quantum technology

KPMG NewTech's quantum team has, in corporation with the Technical University of Denmark (DTU) and a European bank, explored the potential of quantum computing for determining which stocks to buy and sell for maximum return and minimum risk. The quantum annealer proved to perform better than other benchmark methods; however, identifying the right set of parameters is a challenge and the current size of the annealer puts a natural limit on the number of assets that can be embedded.

Quantum technology is entering the business market today, and financial sectors are exploring early applications. Banks are rapidly recognising this potential, creating in-house teams, securing Intellectual Property Rights, and, even in some cases, testing models on simulators and commercial quantum annealers.

The global KPMG quantum hub has conducted a project using a modelling approach from the paper "Financial Portfolio Management using D-Wave Quantum Optimizer: The Case of Abu Dhabi Securities Exchange," where Nada Elsokkary et al. formulated a Markowitz model based on expected return, covariance, and budget. This model allows one to generate high return and diversified portfolio respecting a given budget. Our focus was on analysing the potential of D-Wave through first-hand experience.

We found that accessing the quantum annealer is reasonably easy thanks to the new Ocean Software Development Kit (SDK). Nevertheless, creating and adapting a model to meet the requirements of the annealer is not straightforward. As mentioned, the D-Wave handles a specific class of problems well, and translating the problem into a form that can be solved by the quantum annealer is the first step.

Once the model is created, it must be tested. Here, we validated whether our program was working by comparing the performance with a classical brute force method for a low number of assets (up to 25). Once we assured that the results provided by D-Wave matched the optimal solution provided by the brute force method, we scaled the problem up. [In the figure below, we compared the time it takes for bf with the time for DW machine, the time using bf grows exponentially, while the time for the annealer is essentially constant].

We successfully embedded up to 65 stocks and ran our algorithms over Nasdaq data, benchmarking the results against simulated annealing and brute force. The brute force was naturally worst in terms of time and was mainly used to check the solution's quality up to 25 stocks (intractable for higher number of assets). For more than 25 assets, we used simulated annealing to estimate the quality of the solution obtained from quantum annealing, and we generally found quantum annealing to have better solution than simulated annealing. However, we did encounter scalability issues with the quality of the results obtained by the 2000Q D-Wave, finding that as the portfolio size increased, the model continued to find a solution but started to collapse to a normal distribution if we maintained the parameters unchanged.

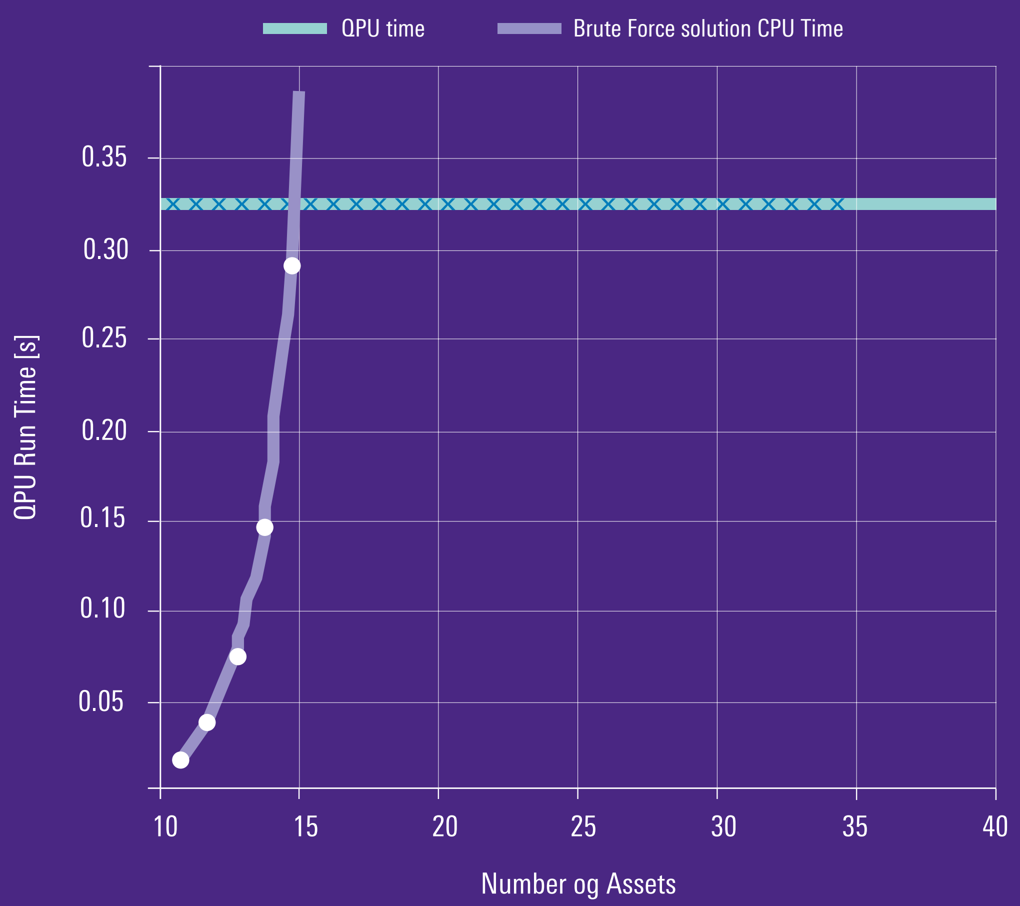

This figure compares the time used from D-Wave to solve the optimization compared to brute force. We can clearly see the exponential separation as the number of assets grows

A deeper look into the details

$$\text{Minimize}: \theta_0 \sum_i (-\alpha_i E(R_i)) + \theta_1 \sum_i \sum_j \alpha_i \alpha_j \text{Cov}(R_i, R_j) + \theta_2 \left(\sum_i \alpha_i A_i - B \right)^2$$

Let us look at the model taken from Nada Elsokkary et al. "Financial Portfolio Management using D-Wave Quantum Optimizer: The Case of Abu Dhabi Securities Exchange. Tech. rep. Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 2017":

- E(Ri) is the expected return of the stock Ri, given by the data, e.g. Nasdaq data
- Cov(Ri, Rj) is the covariance between stock Ri and Rj, representing the diversification of the portfolio
- Ai is the cost of the stock Ri
- B is the total budget for investment
- αi are binary values (0,1) that indicate whether a given stock is purchased or not

Inputs to the model are the covariance matrix, the expected return of each stock, together with the price and budget. We expect as output a string of 0's and 1's indicating whether to buy or not buy a given stock.

The theta parameters are there to tune the optimisation, and they sum to 1. In particular:

- θ0 controls how much the expected return should play a role in our model
- θ1 controls the diversification of the portfolio
- θ2 controls how tightly we stick to the budget

Note that the number of possible combinations of the portfolio grows exponentially with the number of assets (2x, where x is the number of assets), and the brute force model becomes computationally intractable very quickly.

The tasks required for implementation were adapting the model into a QUBO formulation and the tuning of the D-Wave parameters, for example, for annealing time and chain strength.

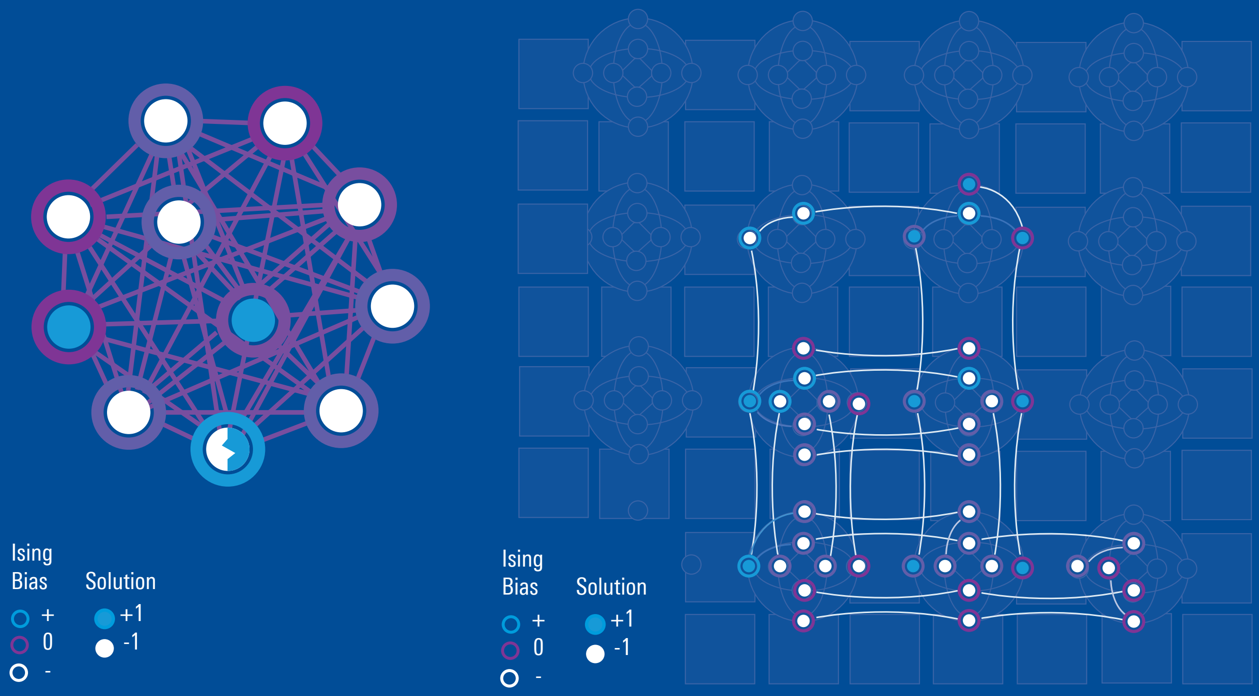

Embedding a problem into D-Wave means connecting different qubits and finding the configuration with the lowest energy state. The architecture of the chip, though, only permits a few direct connections. When direct connections are not possible, the user can create chains of physical qubit that will be considered a logical qubit. We expect the chain to have the same colour at the end of the annealing. If the chain "breaks", the solution is not reliable. Example of an embedding of 11 assets into the D-Wave machine through the Ocean SDK. On the left, you can see a chain break, which means that the problem solved by the machine is different from the one proposed. This figure aims to represent problem interconnection on the left and the actual embedding on the chip on the right.

To check the solution's validity, we compared the results obtained by D-Wave with a classical check, brute force, for solutions up to 25 stocks. The global minimum was found by D-Wave, confirming the validity of the simulator. However, as the problem grew, we needed to use bigger chains and we experienced a drop in the quality of the solution. Even though the model worked, it failed to find a good solution and started to collapse to a normal distribution.

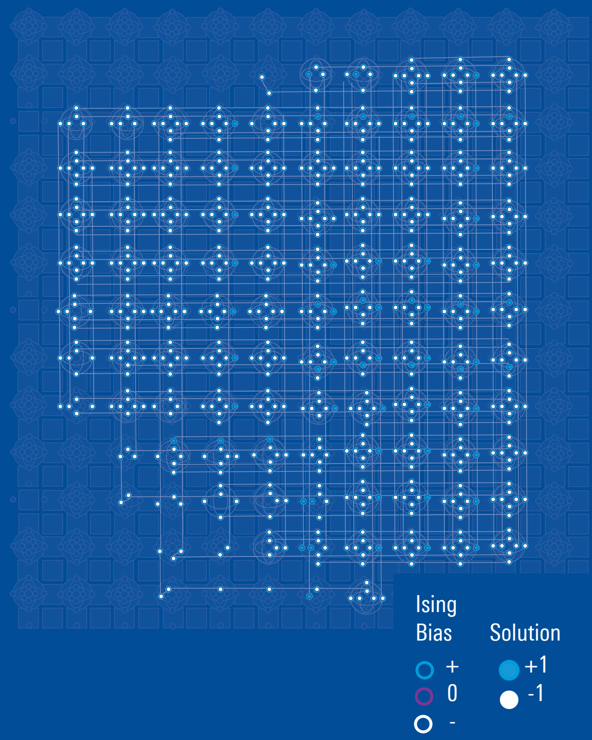

This figure represents a large embedding of stocks. We can see that to represent the optimisation of 65 assets, we use almost all the qubits available in the chip. This is due to the interconnectivity.

Acknowledgments

We would like to thank Mathias Boisen (DTU) for his technical implementation and insights during the project together with Ulrich Busk Hoff (DTU) and Ulrik Andersen (DTU) for supervising the steps of the project.

Contact

If you would like to know more about the project, the results and challenges please contact:

Troels Steenstrup Jensen

Senior Manager,
Head of Machine Learning
and Quantum Technologies

trsjensen@kpmg.com

Marco Ugo Gambetta

Consultant - NewTech

mgambetta@kpmg.com