White Paper – Foreign Exchange Pairs Portfolio Optimisation

Dr Kingsley Jones, Tristan Cook, Duncan Fletcher

How will Quantum Computing affect your business? QxBranch presents the first in a series of White Papers showing current, commercially available quantum computers solving real world problems. We take a common investment problem – portfolio optimisation, demonstrate how to run the problem on a D-Wave Two quantum computer and show how the D-Wave Two successfully handles a complicating feature that is challenging for classical computers to solve. We link the quantum computing capabilities of today to a transformative, quantum-powered future.

Introduction

This paper details a QxBranch-developed prototype method for using the D-Wave Two® Adiabatic Quantum Computer to construct and maintain a simple portfolio of foreign exchange currency pairs (forex pairs) using discrete optimisation.

The results demonstrate:

- that the **current** D-Wave Two can be used to solve practical, real-world problems,
- **how** the D-Wave Two is used in practice; and
- how the D-Wave Two can **improve** the robustness of a solution, compared to classical methods.

We select forex pairs rather than stocks or other instruments for demonstration purposes to take advantage of more consistent historical data (currencies are created and replaced more rarely than equities). Otherwise, this example applies to any tradable financial instrument. The methods used in this paper can scale into an operational system.

This paper:

- provides some background to the portfolio optimisation problem and the D-Wave Two architecture,
- outlines the parameters of the demonstration problem we run,
- explains how the problem is programmed onto the Quantum Computer,
- shows the results of executing the optimisation strategy over a year of trades and
- discusses how this simple example is the basis of larger, more sophisticated strategies.

Background – Portfolio Optimisation

Classical portfolio theory, as originated by Markowitz in 1952, is an important underpinning to modern quantitative finance. This approach trades off risk, as given by the covariance matrix between instrument values, against return, as given by the forecast future value or return of each instrument. It aims to maximise portfolio return and diversity simultaneously. This theory employs continuous portfolio weights and the portfolio is optimised by computationally-easy, convex quadratic programming. However, the success of such classical methods is limited by the uncertainty of key inputs – particularly the risk and returns forecasts. Typically, the uncertainty in these inputs is comparable to their forecast magnitudes due to the inherent volatility of the market. In practical terms, the optimal portfolio moves continuously with small shifts in values and forecasts, attracting significant trading costs while chasing noise.

A more robust technique for portfolio construction is to optimise **discrete** holdings of the desired securities. With discrete holdings, a more significant shift in value or forecast is required to trigger a trade. In most applications it has the further advantage of aligning better with reality – where instruments are usually held in round lots.

Unfortunately, optimisation of a discrete portfolio (a portfolio of discrete trades or round lots) is an *NP-Hard* problem. Meaning that finding the optimal solution is a hard, non-deterministic problem. A priori, you are unable to determine how long it will take to find an optimal solution on classical computer architectures.

Background – D-Wave Two Architecture

In the context of optimisation problems, the D-Wave Two quantum computer employs *quantum annealing* to find high quality solutions to hard problems in fixed time. This is its key advantage over classical computer architectures.

The D-Wave Two chip is part digital, part analogue, part quantum-mechanics. It is designed to find the **minimum** value H of an equation called an *Ising spin glass*:

$$H = \sum_{i \in V} h_i s_i + \sum_{(i,j) \in E} s_i J_{i,j} s_j$$

Where s_i are the qubits or "quantum bits" of the computer and whose value of ±1 is the **output** of a quantum program. The **inputs** of the quantum program, h_i and $J_{i,j}$ are the weightings on the qubits and the coupling strengths between them respectively. When values are assigned to the h and J terms, we have what's called the **problem Hamiltonian** or energy function that is to be minimised.

The D-Wave Two has 512 qubits, arranged in a graph which defines the available couplings. Due to hardware limitations, each qubit has at most six couplings to other qubits. In other words, for most combinations of i and j, there is no $J_{i,j}$ input available. The set E in the above equation is the available couplings.

In practice, this sparse-connectivity limitation can be addressed by a process called *embedding*, where multiple physical qubits are used to represent one logical qubit. The D-Wave Two graph of 512 sparsely connected physical qubits supports approximately 30 fully-connected logical qubits.[1]

Demonstration Problem

For demonstration purposes we pick five common forex pairs to optimise the portfolio over in a long-only strategy. We used historical data sourced from stooq.com and simulate **daily trades** based on the optimiser's recommendations.

Forex pairs:	USD/EUR, USD/CAD, USD/GBP, USD/CHF, USD/AUD
Timeframe:	Starting 1 October 2013 for 252 trading days – one year of trading;

The D-Wave Two architecture is limited to binary outputs, therefore each optimisation of our formulation results in either a "long" or "no-hold" recommendation for each security.

We can use the D-Wave Two to directly optimise a portfolio of approximately 30 forex pairs, but we limited it to five for simplicity in this demonstration.

We select a starting capital of US$10,000 and equally weight investment in long positions given the signals for each time-step in the period.

Rather than employ proprietary forecast algorithms in this example – which would make it difficult to evaluate the effectiveness of the portfolio construction strategy – we use look-ahead on the historical data and a tuneable forecast quality. In other words the returns forecast is calculated from a moving window on the historical data, centred on the

[1] The exact number depends on the configuration of the particular machine. The machines QxBranch uses support up to 33 fully-connected logical qubits.

day of trade. With the forecast quality dial set at 0%, the future return is estimated by using the previous day's closing returns. With the dial set at 100%, the actual (look-ahead) return of the coming day is used. When the dial is set in between, the future return is estimated by a Bayesian blend of the past and future data. This approach allows analysts who have a good feel for the quality of their own forecast algorithms to assess how well a portfolio construction strategy works, without disclosing those algorithms.

In this demonstration, the covariances between instruments are re-calculated each day from the previous 180 days' returns forecast of each instrument.

Programming the Problem

We have multiple approaches available to program the D-Wave Two quantum computer. For this problem, we chose direct expression of the problem into the computer's native form:

- one logical qubit s_i per instrument,
- returns forecast of each instrument gets negated and assigned to the h_i weighting of its corresponding qubit, and
- covariance between each instrument gets assigned to the $J_{i,j}$ coupling strength between their corresponding qubits.

Expressing the problem this way will:

- push the qubits of instruments with high returns forecasts to a +1 output state, indicating a "long" recommendation for the instrument, and
- push qubits of instruments that are correlated with positive covariance to take **opposite** output states, diversifying the portfolio and minimising risk.

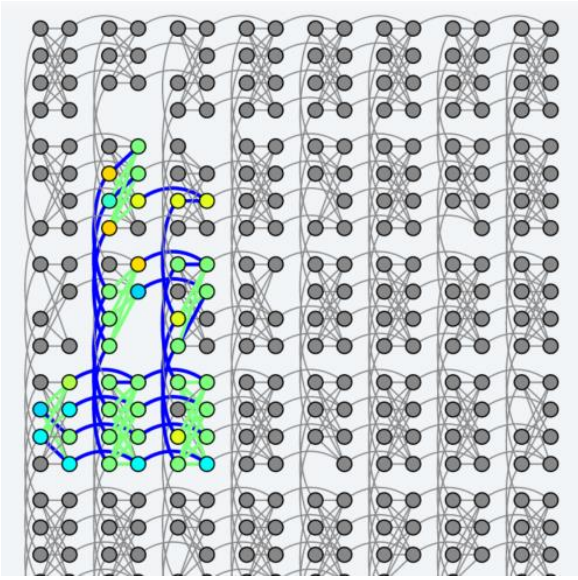

Figure 1 Example trading day input shown on the D-Wave Two chip. Dots represent physical qubits, lines represent couplings between them. Colours of dots and lines represent values currently assigned to hi & Ji,j respectively this run. Grey dots are unused regions of the chip. Missing dots are broken qubits.

As the forex pairs have non-zero covariance values with respect to each other this forms a fully-connected graph, which the D-Wave Two does not directly support. Therefore embedding is employed as discussed earlier: multiple physical qubits are used to represent each logical qubit, where the $J_{i,j}$ coupling strengths *within* the logical chain are set to -1 forcing the physical qubits to *agree*. These appear as blue lines in Figure 1.

For each day's trade optimisation, a classical computer calculates the returns forecasts and covariance matrix. It transforms those through the embedding map to calculate the derived h_i and $J_{i,j}$ values to send to the quantum computer. These numbers are submitted using a Python API[2]. No labels or other problem meta-data is sent to the quantum computer; the submitted problem is just vectors of numbers. A single day's trade recommendations are optimised by a single submission. The actual quantum computer execution takes a fixed, known time for every submission.

[2] Application Programming Interface; the software interface used by programmers to make a particular library or piece of specialised hardware do work.

When submitting the job to run on the quantum computer, we also select the number of iterations. This is the number of times the same input problem will be evaluated to try and find the minimum[3]. The probabilistic nature of quantum mechanics and inherent control errors in the computer mean that you need a large population of 1,000 or more candidate answers to have confidence you've found the optimal. The D-Wave Two quantum computer can deliver up to 35,000 iterations per second.

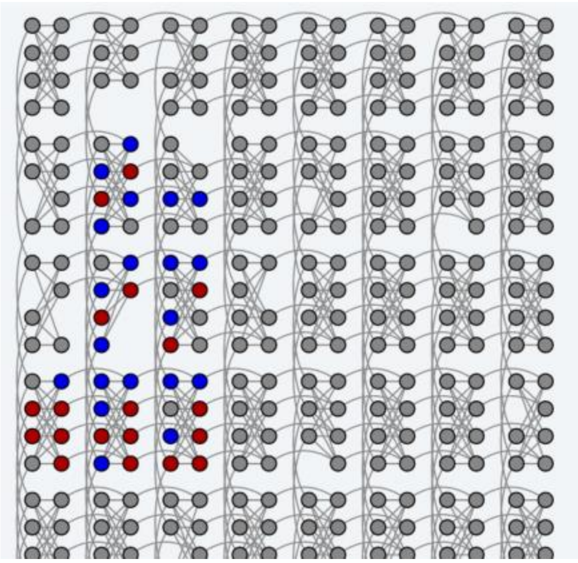

Figure 2 The result of a single iteration for a single day's trade optimisation as calculated by the D-Wave Two chip. Red dots are qubits with output value +1, indicating a "hold" recommendation for their corresponding instrument. Blue dots are qubits with output value -1, indicating "no hold".

The result returned by the quantum computer is the vector of s_i values of ±1 that resulted from each attempt to find the minimum energy H of the problem Hamiltonian. In other words, the optimal solution. Figure 2 illustrates the result of a single iteration. Figure 3 illustrates the combined results of 1,000 iterations. Note that while both days shown in Figure 3 have a clear recommendation, Day B has two reasonable secondary recommendations[4] that may be preferable due to client instructions or a desire to avoid influencing the market in particular ways. This feature of the quantum computer – the generation of multiple, alternative solutions – effectively comes for free and can be either used or ignored.

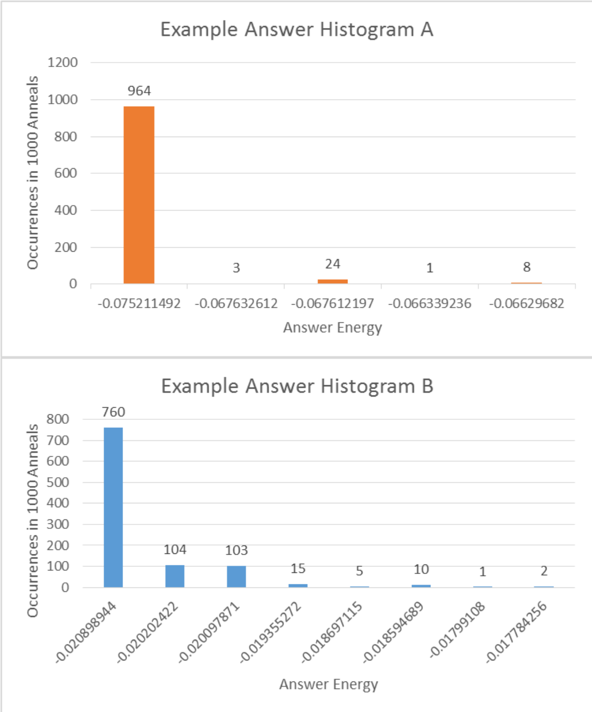

Figure 3 Histograms of occurrence of unique trade recommendations for two different days' trade. Answer Energy is the value H of the problem Hamiltonian. Note that in both cases the optimal (minimum energy) solution is the most common by a significant margin.

The simple portfolio optimiser developed for this paper selects the lowest energy combination of qubit values and, again using the classical computer, uses those as signals to simulate the recommended trades for that day. The resulting portfolio is carried over to the next day and the optimisation repeated.

The cost overhead of making each trade was not included in this simulation.

Results - Timing

It took approximately 0.5 seconds for the D-Wave Two to find the optimal solution for each period's trades. It took a further 10-15 seconds

[3] Technically "annealed". The D-Wave Two finds the minimum of the problem Hamiltonian by a process called *quantum annealing.* Each time you let the computer attempt to find the global minima, you are letting it *anneal*. The process has some similarities to the commonly-used *simulated annealing*.

[4] Considered reasonable because they are close to the primary answer's energy.

for internet round-tripping between the classical part of the algorithm and the quantum computer. Therefore the trading frequency could be increased to 2-3 times per minute with the current physical architecture. Co-locating the classical and quantum computers would allow the trading frequency to be significantly increased.

Network latency, job queueing and problem size were the only source of variation in calculation time for this problem. Fixed time-to-solution for a given problem size is an inherent feature of the quantum program.

Results – Portfolio Performance

Using the forex pairs and timeframe described in the previous sections, we ran the same portfolio optimisation simulation with the forecast quality dial set from 0% to 100% in 10% increments. The overall performance of the portfolio construction strategy is shown in Figure 4 & Figure 5. Recall that 0% forecast quality means forecast return is drawn only from past data, while 100% forecast quality is the return of only look-ahead data. 50% forecast quality is an equal probability each day of selecting one of the two. Over the year simulated, the portfolio returns range from a 10% loss under low quality forecasts up to a 70% gain under high quality forecasts. In assessing this performance, recall that the signals on each instrument are simply binary (long or no-hold) and that the portfolio is re-balanced daily to equally weight the long positions. To illustrate the nature of the trade recommendations, Figure 6 shows the portfolio activity over the 252 trading days for a forecast quality setting of 20%.

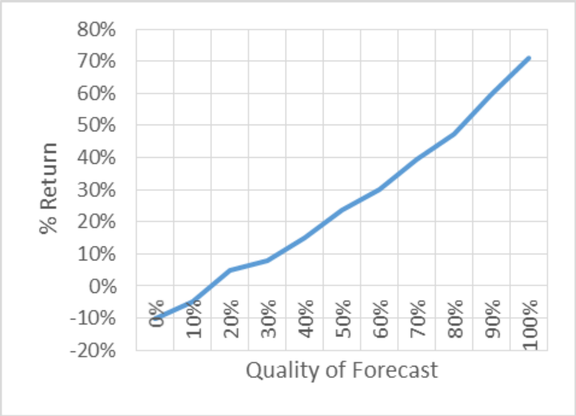

Figure 4 Total portfolio return after a year of simulated trades, as a function of forecast quality

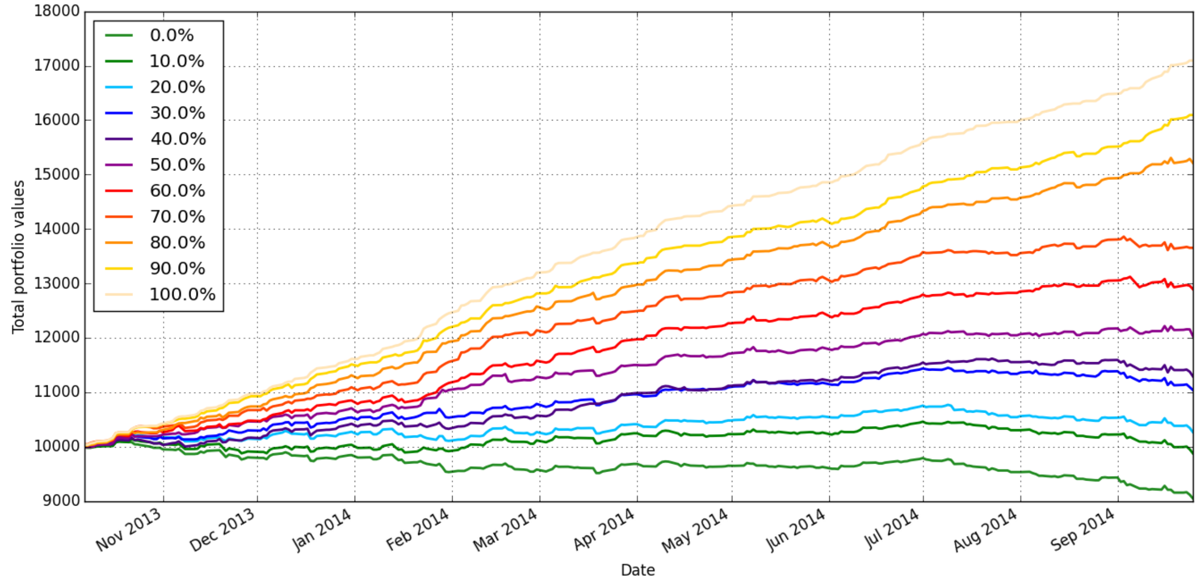

Figure 5 One year's simulated portfolio trading activity for forecast accuracies from 0% to 100% in 10% increments. The lines plot the value of the portfolios over time in USD.

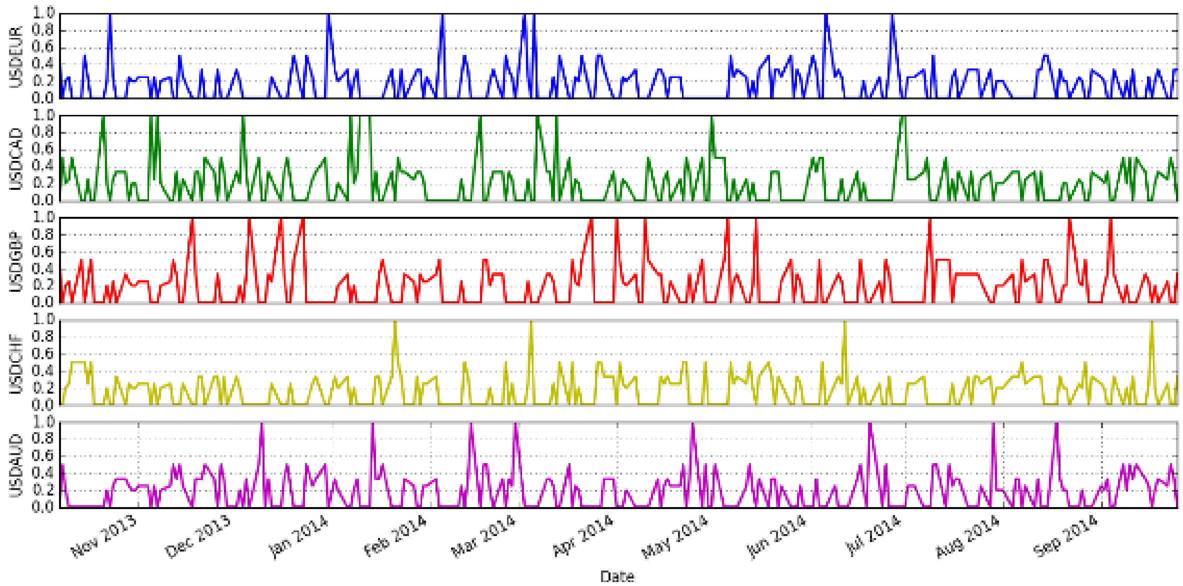

Figure 6 One year's simulated, equally-weighted long signals for 20% forecast quality.

Conclusion

Through this paper we've shown that the current D-Wave Two quantum computer can be used to solve a practical, real-world problem: the optimisation of a portfolio of foreign exchange currency pairs to maximise return on investment.

In doing this we've demonstrated one method for programming a mathematical problem onto the quantum computer – that of directly mapping the problem data onto a fully-connected Ising spin glass which is then embedded into the physical qubits of the machine.

Furthermore, this demonstration shows that the D-Wave Two quantum computer can successfully solve a hard variation on the classical Markowitz theory, by introducing discrete lots of portfolio holdings. This more robust method is in a class that classical computer-based algorithms struggle to find optimal answers in fixed time.

From this example we can sketch a roadmap to a more sophisticated, operational system:

- extension to long-short-no hold strategies,
- expansion to optimise over more instruments – noting that there are techniques called *divide and conquer* that allow for quantum-powered optimisation of more instruments than will fit on the chip,
- introduction of trade costs to penalise over-active trading,
- introduction of investment grade return and risk forecast algorithms – noting that these will be client-specific, not QxBranch provided, and
- development of the classical driver of the optimiser into enterprise grade software.

Most importantly, the technology roadmap for real-world problems such as this positions organisations to take immediate advantage of future developments in quantum computing.

