Finance Applications of Adiabatic Quantum Optimisation

This document summarises the technical capability and applicability to finance of the current generation of Adiabatic Quantum Optimisation (AQO) computers, such as the D-Wave® quantum computers.

Introduction

Adiabatic quantum computers leverage quantum effects to rapidly solve complex optimization problems. This capability yields three classes of likely application areas in finance: natural fit optimisation, machine learning and Monte Carlo-based pricing and risk analysis.

Figure 1 shows the relative technical difficulty, estimated timeline to emergence, and estimated impact of applications within these areas.

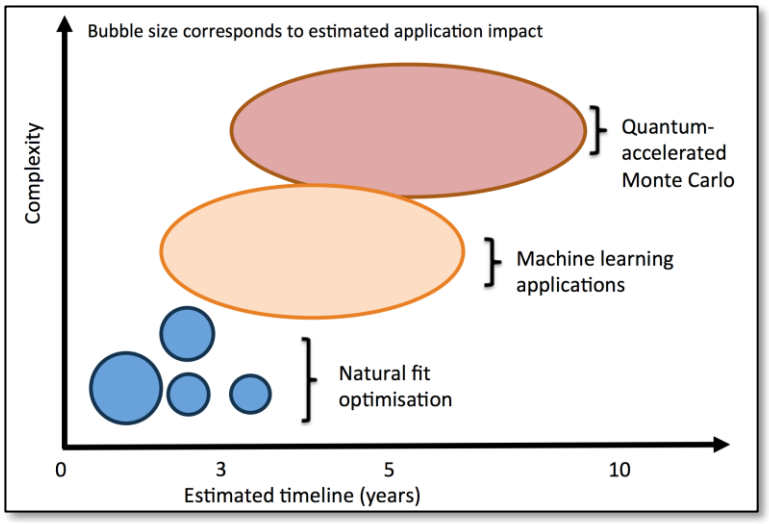

Figure 1: AQO application areas

Natural fit optimisation

Problems such as portfolio optimisation, pattern analysis within graphs of network traffic or trades and route optimisation for capital movement are natural fits to the underlying structure of AQO. AQO systems are able to solve the *hard* variant of many of these problems. For example, including trading costs and cardinality in portfolio optimisation introduces non-linear, integer programming which is known to be NP-Hard. The D-Wave quantum computer has been shown to out-perform current classical solvers of such problems when searching for near-optimal solutions.

Quantum-accelerated Monte Carlo

Monte Carlo simulation is used extensively in finance for risk analysis and market pricing. Any substantial improvement in Monte Carlo simulation performance – either in precision or speed – is likely to have significant impact. Research has shown that *Universal Quantum Computers,* the next generation due in 5–10 years, will almost certainly accelerate Markov Chain Monte Carlo simulation. In the meantime, AQO can be used to augment Monte Carlo simulation by using quantum-accelerated machine learning techniques on surrounding processes.

Machine Learning

AQO can enhance several machine learning techniques. Finance problems like price factorisation, principle component analysis and black swan analysis can be solved using the mathematics of machine learning. Key machine learning constructs such as regularised regression and boosted decision models can be optimised using the D-Wave, performing particularly well on non-convex functions. Recent research has also shown that algorithms built for the D-Wave have improved convergence characteristics over classical computers in training a deep belief network with restricted Boltzmann machines.

Example applications of machine learning	
Risk modelling	Asset pricing
Monte Carlo methods	Multi-state arbitrage
Portfolio optimisation	Stress testing
Market segmentation	Market making
Causal analysis	Collateral allocation
Sentiment analysis	Cyber security
Anomaly detection	Market simulation

D-Wave Capability

The D-Wave 2X, manufactured by Vancouver-based D-Wave Systems, is the third generation of the world's first commercially available quantum computer. It has unique capabilities for complex optimisation via *Adiabatic Quantum Optimisation*. The machine can rapidly find the optimal or near-optimal solution to high-dimensional problems expressed as a Quadratic Unconstrained Binary Optimisation (QUBO).

Research on the current generation of hardware shows that for most optimisation problems the machine can find a near-optimal solution in a time proportional to the square root of the problem dimensionality.

This probabilistic nature of the machine is the source of its benefits. Its advantages are strongest when:

- Time to solution is a factor and a near-optimal solution suffices
- A diversity of near-optimal solutions is useful as an input into further analysis
- Useful near-optimal solutions exist at the margins of a particular function (i.e., the solution space contains important 'black swans')

The current limitations of the machine are the number of qubits, the number of connections available between the qubits and the error rate.

