

Portfolio Optimisation Using the D-Wave Quantum Annealer

Frank Phillipson¹ and Harshil Singh Bhatia²

¹ TNO, the Netherlands Organisation for Applied Scientific Research, The Hague, The Netherlands

² Department of Computer Science and Engineering, Indian Institute of Technology, Jodhpur, India

Abstract. The first quantum computers are expected to perform well on quadratic optimisation problems. In this paper a quadratic problem in finance is taken, the Portfolio Optimisation problem. Here, a set of assets is chosen for investment, such that the total risk is minimised, a minimum return is realised and a budget constraint is met. This problem is solved for several instances in two main indices, the Nikkei225 and the S&P500 index, using the state-of-the-art implementation of D-Wave's quantum annealer and its hybrid solvers. The results are benchmarked against conventional, state-of-the-art, commercially available tooling. Results show that for problems of the size of the used instances, the D-Wave solution, in its current, still limited size, comes already close to the performance of commercial solvers.

Keywords: Quantum Portfolio Optimisation, Quadratic unconstrained binary Optimisation, Quantum annealing, Genetic algorithm

1 Introduction

Portfolio management is the problem of selecting assets (bonds, stocks, commodities) or projects in an optimal way. Classical portfolio management, as introduced by Markowitz, focuses on efficient (expected) mean-variance combinations [20], and has led to a broad spectrum of optimisation problems: single and multi-objective [29,35], single and multi-period [33,18], without and with [10,18] transaction costs, deterministic or stochastic [26,27] in all possible combinations. One of the basic problems is the single objective, maximising expected return, under budget and risk constraints. Risk is expressed by the covariance matrix of all assets. In this paper we consider a variant minimising the risk, under budget and return constraints. This is applied in family trust and pension funds, where a specific return is needed for future liabilities and a low risk is desirable. This leads to a quadratic optimisation problem with continuous or binary variables. The variables are continuous if a fraction of the budget is allocated to an asset. Binary variables can be found when the choice is whether or not to invest in a specific asset or project.

Solving this kind of problems is not trivial. Integer quadratic programming (IQP) problems are NP-hard, the decision version of IQP is NP-complete [6]. Binary quadratic programming problems are also NP-hard in general [12], however, specific cases are polynomially solvable [17]. Classical solvers, such as IBM-CPLEX, Gurobi and Localsolver, are still getting better in solving these problems for bigger instances. Next to these classical solvers, heuristic approaches exist, based on meta-heuristics like Particle Swarms, Genetic Algorithms, Ant Colony and Simulated Annealing. A recent overview of these approaches for Portfolio Optimisation can be found in [36].

Quadratic optimisation problems with binary decision variables are expected to be a sweet-spot for near future quantum computing [32], using Quantum Annealing [22] or the Quantum Approximate Optimisation Algorithm (QAOA) on a gate model quantum computer [9]. Quantum computing is the technique of using quantum mechanical phenomena such as superposition, entanglement and interference for doing computational operations. The type of devices that are capable of doing such quantum operations are still actively being developed and are called quantum computers. We distinguish between two paradigms of quantum computing devices: gate model based (or gate based) and quantum annealers. A practically usable quantum computer is expected to be developed in the next few years. In less than ten years quantum computers are expected to outperform everyday computers, leading to breakthroughs in artificial intelligence [24], the discovery of new pharmaceuticals and beyond [23,13]. Currently, various parties, such as Google, IBM, Intel, Rigetti, QuTech, D-Wave and IonQ, are developing quantum chips, which are the basis of the quantum computer [31]. The size of these computers is limited, with the state-of-the-art being around 70 qubits for gate-based quantum computers and 5000 qubits for quantum annealers. In the meantime, progress is being made on algorithms that can be executed on those quantum computers and on the software (stack) to enable the execution of quantum algorithms on quantum hardware [28].

This also means that Portfolio Optimisation is one of the promising applications of quantum computing in finance [25]. The work of [8] presents an implementation of Markowitz's portfolio selection on D-Wave, where the expected return is maximised whilst minimising the covariance (risk) of the portfolio under a budget constraint. They formulate this as an Ising problem and solve it on the D-Wave One, having 128 qubits, for 63 potential investments within 20 μ s on the quantum processor. They indicate that the solution depends on the weights added to each of the objectives and constraints. The same is done in [34], where reverse quantum annealing is used to optimise the risk-adjusted returns by the use of the metrics of Sharpe ratio. In [21], the stock returns, variances and covariances are modelled in the graph-theoretic maximum independent set (MIS) and weighted maximum independent set (WMIS) structures under combinatorial optimisation. These structures are mapped into the Ising physics model representation of the underlying D-Wave One system. This is benchmarked against the MATLAB standard function quadprog.

A newer version of the D-Wave hardware is used in [4]. They also perform a stock selection out of a universe of U.S. listed, liquid equities based on the Markowitz formulation and the Sharpe ratio. They approach this first classically, then by an approach also using the D-Wave 2000Q. The results show that practitioners can use a D-Wave system to select attractive portfolios out of 40 U.S. liquid equities. The research has been extended to 60 U.S. liquid equities [5].

Algorithms have also been created for the gate model quantum computer. In [30] an algorithm is given for Portfolio Optimisation that runs in $\text{poly}(\log(N))$, where N is the size of the historical return data set. The number of required qubits here is also N . An alternative method for a gate model quantum computer is given by [16]. In [11] a quantum-enhanced simulation algorithm is used to approximate a computationally expensive objective function. Quantum Amplitude Estimation (QAE) provides a quadratic speed-up over classical Monte Carlo simulation. Combining QAE with quantum optimisation can be used for discrete optimisation problems like Portfolio Optimisation.

With the introduction of the 5000 qubit Advantage quantum system by D-Wave and the new hybrid solver tooling a new heuristic approach is operational for in-production quantum computing applications³. In this paper we use the new functionality of D-Wave for the portfolio management problem and compare its performance with other state-of-the-art, commercially available tooling.

There exists other work that compares D-Wave's quantum annealer to other optimisation tooling on other application areas. We name a few examples. A QUBO formulation for the railway dispatching problem has been proposed in [7]. They further benchmark their solutions for the polish railway network on the state of the art D-Wave hardware. In [2], the authors have benchmarked the 5000 Advantage system for the garden optimisation problem class. Their results demonstrate that the Advantage system and the new hybrid solvers can solve large instances in less time than its predecessor.

In the remainder of this paper first the Portfolio Optimisation problem is formulated. This is a quadratic, constrained, binary optimisation problem that can in principle be solved by commercial solvers. The solvers we use in this paper as benchmark for the quantum approach are presented in Section 3. In Section 4 the implementation of the problem on the D-Wave quantum annealer is shown. The results of solving instances of two main stock indices, the Nikkei225 and the S&P500, on the quantum annealer and the comparison with the benchmark approaches, will be presented in Section 5. We end with some conclusions and recommendations.

2 Problem description

We look at a Portfolio Optimisation problem, where we have N assets to invest in, P_1, \dots, P_N . The expected return of assets i equals μ_i and the risk of the

³ <https://www.D-Wavesys.com/press-releases/d-wave-announces-general-availability-first-quantum-computer-built-business>

asset, denoted by the standard deviation, equals σ_i . The returns of the assets are correlated, expressed by the correlation ρ_{ij} for the correlation between assets i and j . We have now the return vector $\mu = \{\mu_i\}$ and the risk matrix $\Sigma = \{\sigma_{ij}\}$ where $\sigma_{ij} = \sigma_i^2$ if $i = j$ and $\sigma_{ij} = \rho_{ij}\sigma_i\sigma_j$ if $i \neq j$.

Assume that we have a budget to select n assets out of N . We want the return to be higher than a certain value R^* and are searching for that set of n assets that realise the target return against minimal risk. We therefore define $x_i = 1$ if asset i is selected and $x_i = 0$ otherwise. This gives the following optimisation problem:

$$\min x^T \Sigma x, \quad (1)$$

$$\text{s.t. } \sum_{i=1}^N x_i = n, \quad (2)$$

$$\mu^T x \geq R^*, \quad (3)$$

which is a quadratic, constrained, binary optimisation problem.

3 Benchmarks

We want to compare the performance of the D-Wave hardware with other state of the art, commercially available tooling. For this we selected two solvers and two meta-heuristics.

The first solver is LocalSolver [1], which is a black-box local-search solver of 0-1 programming with non-linear constraints and objectives. A local-search heuristic is designed according to the following three layers: Search Strategy, Moves and Evolution Machinery. LocalSolver performs structured Moves tending to maintain the feasibility of solutions at each iteration whose evaluation is accelerated by exploiting invariants induced by the structure of the model. Unlike other math optimisation software, LocalSolver hybridises different optimisation techniques dynamically. We used an academic licensed local version.

The second solver is Gurobi [15], from Gurobi Optimisation, Inc., which is a powerful optimiser designed to run in multi-core with capability of running in parallel mode. Gurobi uses the Branch and Bound Algorithm to solve Mixed-Integer Programming (MIP) models. It is based on four basic principles: pre-solve, cutting planes, heuristics, and parallelism. Each node in the Branch and Bound search tree is a new MIP. For our current MIQP (Mixed Integer Quadratic Programming), a Simplex Algorithm is used to solve the root node. We used an academic licensed local version of Gurobi.

The third benchmark is a standard MATLAB implementation of Genetic Algorithms, which is a method for solving both constrained and unconstrained optimisation problems based on a natural selection process that mimics biological evolution. The algorithm repeatedly modifies a population of individual solutions. It is a stochastic, population-based algorithm that searches randomly by mutation and crossover among population members.

The last benchmark is the Simulated Annealing approach as implemented by D-Wave in their Ocean environment. Their sampler implements the simulated annealing algorithm, based on the technique of cooling metal from a high temperature to improve its structure (annealing). This algorithm often finds good solutions to hard optimisation problems.

4 Implementation

In this paper we use the newly introduced (2020) functionality of D-Wave for the portfolio management problem and compare its performance with the commercial solvers described in Section 3. In this section we describe how the problem can be implemented on the D-Wave hardware. For this, the QUBO representation will be derived, a method to find the parameters is given and D-Wave's quantum and hybrid algorithms are explained.

4.1 QUBO representation

The D-Wave hardware solves Ising or QUBO problems. The QUBO [14] is expressed by the optimisation problem:

$$\text{QUBO: } \min/\max y = x^t Q x, \quad (4)$$

where $x \in \{0, 1\}^n$ are the decision variables and Q is a $n \times n$ coefficient matrix. Another formulation of the problem, often used, equals

$$\text{QUBO: } \min/\max H = x^t q + x^t Q x, \quad (5)$$

or a combination of multiple of these terms

$$\text{QUBO: } \min/\max H = \lambda_1 \cdot H_1 + \lambda_2 \cdot H_2 + \dots, \quad (6)$$

where $\lambda_1, \lambda_2, \dots$ are weights that can be used to tune the problem and include constraints into the QUBO. For already a large number of combinatorial optimisation problems the QUBO representation is known [14,19]. Many constrained binary programming problems can be transformed easily to a QUBO representation. Assume that we have the problem

$$\min y = c^t x, \text{ subject to } Ax = b, \quad (7)$$

then we can bring the constraints to the objective value, using a penalty factor λ for a quadratic penalty:

$$\min y = c^t x + \lambda(Ax - b)^t(Ax - b). \quad (8)$$

Using $P = Ic$, the matrix with the values of vector c on its diagonal, we get

$$\min y = x^t P x + \lambda(Ax - b)^t(Ax - b) = x^t P x + x^t R x + d = x^t Q x, \quad (9)$$

where matrix R and the constant d follow from the matrix multiplication and the constant d can be neglected, as it does not influence the optimisation problem.

A QUBO problem can be easily translated into a corresponding Ising problem of N variables s_i ($i = 1, \dots, N$) with $s_i \in \{-1, 1\}$ given by :

$$\min y = \sum_{i=1}^N h_i s_i + \sum_{i=1}^N \sum_{j=i+1}^N J_{ij} s_i s_j. \quad (10)$$

The Ising model and QUBO model are related by $s_i = 2x_i - 1$.

4.2 QUBO formulation

We now create a QUBO formulation for the problem given in Equation (1)-(3). In case Equation (3) is an equality, this problem can be translated easily to a QUBO formulation:

$$\min (\lambda_0 x^T \Sigma x + \lambda_1 (\sum_{i=1}^N x_i - n)^2 + \lambda_2 (\mu^T x - R^*)^2). \quad (11)$$

In the case of inequalities in Equation (3) we have to add additional K slack variables y_k ($k = 1, \dots, K$), where $K = \lceil \log_2(\sum_{i=1}^N (\mu_i)) \rceil$. Scaling the μ values (in thousands, millions, etc.) will help reduce the number of variables. This leads to

$$\min (\lambda_0 x^T \Sigma x + \lambda_1 (\sum_{i=1}^N x_i - n)^2 + \lambda_2 (\mu^T x - R^* - \sum_{k=1}^K 2^k y_k)^2). \quad (12)$$

4.3 Finding QUBO parameters

It is known that the performance of the D-Wave hardware is depending strongly on the choice of the penalty coefficient λ . When determining the penalty coefficients, we can set $\lambda_0 = 1$ and look for good values for λ_1 and λ_2 . In Figure 1 it is shown that the choice of λ_1 greatly influences the best found solution. Values for λ_1 lower than 500 gives an invalid solution, i.e., the solutions do not meet the constraints given in the original problem. Values for λ_1 higher than 500 give allowed solutions. A value of 500 or just above would be optimal. Rule of thumb is that the gain of violating a constraint must be lower than the costs. For λ_1 this means that violating the associated constraint in the optimal solution x^* for stock i gives a benefit of $\sum_{j=1}^N \sigma_{ij} x_j^*$. Around the optimal value, the biggest benefit would be for stock i for which the sum of the smallest n values of σ_{ij} is the largest, meaning

$$\widehat{\lambda}_1 = \max_i \sum_{j=1}^n \sigma_{i\{j\}},$$




Fig. 1. Relation between Risk and setting of λ . Open dots are violating the original constraints, while closed dots are valid solutions.

where $\sigma_{i\{j\}}$ represents the j -th smallest covariance value for asset i .

This is more complicated for λ_2 . Again, we have to look at the optimal solution x^* and exchange a zero and one in this solution, such that n stocks are chosen, leading to x' . This is done such that $(x^* - x')^T \Sigma (x^* - x') / \mu^T (x^* - x')$ is maximised. The procedure we used is as follows:

1. $A1$ = average difference between smallest n sums $S_i = \sum_{j=1}^n \sigma_{i\{j\}}$,
2. $A2$ = average positive difference in μ_i between these n stocks,
3. $\widehat{\lambda}_2 = A1/A2$.

Note that these values are a first estimation and can be the starting point for an eventual grid search of the optimal parameters. In Table 1 we show the results of the grid search and the comparison with the estimated values. It is shown that the estimation is accurate (enough), such that the grid search can be skipped in practice.

4.4 Solving QUBO

The most recent hardware, D-Wave Advantage, features a qubit connectivity based on Pegasus topology. The QUBO problem has to be transformed to this structure. Due to the current limitation of the chip size, a compact formulation of the QUBO and an efficient mapping to the graph is required. This problem is known as Minor Embedding. Minor Embedding is NP-Hard and can be handled by D-Wave's System automatically, hence we do not attempt to optimise it. While some qubits in the chip are connected using external couplers, the D-Wave QPU (Quantum Processing Unit) is not fully connected. Hence a problem

variable has to be duplicated to multiple connected qubits. Those qubits should have the same value, meaning the weight of their connection should be such that it holds in the optimisation process. All these qubits representing the same variables are part of a so-called chain, and their edge weights is called the chain strength (γ), which is an important value in the optimisation process. In [3], it has been indicated that if γ is large enough, the optimal solutions will match $\gamma \geq \sum_{ij} |Q_{ij}|$. However the goal is to find the smallest value to avoid rescaling the problem. The problem of finding the smallest γ is NP-Hard. On a higher level, D-Wave offers hybrid solvers. These solvers implement state of the art classical algorithms with intelligent allocation of the QPU to parts of the problem where it benefits most. These solvers are designed to accommodate even very large problems. This means that also most parameters of the embedding are set automatically. By default, samples are iterated over four parallel solvers. The top branch implements a classical Tabu Search that runs on the entire problem until interrupted by another branch completing. The other three branches use different decomposers to sample out a part of the current sample set to different samplers. The D-Wave's System sampler uses the energy impact as the criteria for selection of variables.

5 Results

We solved the Portfolio Optimisation problem for two indices, the Nikkei225 and the S&P500. We solved several instances for both indices on the D-Wave annealer, using the implementation described in Section 4, and all four benchmark approaches described in Section 3. The instances were run on a Intel(R) Core(TM) i7-8565U CPU@1.80Ghz 8GB RAM personal computer. We used the newest version of the D-Wave Leap environment solvers, Hybrid Binary Quadratic Model Version 2, for binary problems. The solver is, due to the underlying QPU, a stochastic solver. For this, we ran the hybrid solver five times for each problem instance and performed a parameter grid search for λ_1 and λ_2 for each problem instance.

We used the implementation first to select a number of stocks of the Nikkei225 index. The Nikkei225 is a stock market index for the Tokyo Stock Exchange. It has been calculated daily by the Nihon Keizai Shimbun (The Nikkei) newspaper since 1950. It is a price-weighted index, operating in the Japanese Yen, and its components are reviewed once a year. The Nikkei measures the performance of 225 large, publicly owned companies in Japan from a wide array of industry sectors. We took quarterly data of the Nikkei225 index of the last five year. The used return (μ) of each stock is the five year return, and the covariance σ_{ij} the calculated covariance over all quarters of the last 5 years. From this index, we took a subset N of which n are selected to minimise the risk under a budget and return (R^*) constraint. Note that (R^*) and n are related here: $R^* = 3000$ and $n = 20$ means an average 5-year return of the portfolio of $3000/20 = 150$ percent.

In Table 1 we show the 20 test instances we created on the Nikkei225, each having a different combination of (N, n, R^*) . For each instance we show the size of the coefficient matrix of the QUBO ($\text{Size}(Q)$), the calculated (λ) and optimal $\widehat{\lambda}_2$ value of the parameters and the resulting value of the objective function of the best solution found by the hybrid solver. Also the best known solution is shown.

N	n	R^*	$\text{Size}(Q)$	$\widehat{\lambda}_1$	$\widehat{\lambda}_2$	λ_1	λ_2	HQPU	Best
50	10	0	60	76	0	70	0	256	256*
50	25	0	60	320	0	365	0	3,035	3,035*
50	10	1,200	60	76	0.04	76	0.1	321	321*
50	25	1,200	60	320	0.10	365	0.1	3,058	3,058*
100	20	0	111	137	0	100	0	809	809*
100	50	0	111	607	0	700	0	10,999	10,999*
100	20	2,500	111	137	0.05	100	0.1	1,126	1,126*
100	50	2,500	111	607	0.09	700	0.1	11,065	11,065*
150	20	0	161	116	0	100	0	628	628*
150	50	0	161	502	0	500	0	7,993	7,993*
150	20	3,000	161	116	0.04	100	0.1	1,271	1,256*
150	50	3,000	161	502	0.10	500	0.1	8,315	8,270*
200	20	0	211	100	0	100	0	530	529
200	50	0	211	428	0	500	0	6,000	5,950
200	20	3,250	211	100	0.04	100	0.1	1,308	1,262
200	50	3,250	211	428	0.08	500	0.1	6,859	6,500
225	20	0	236	88	0	88	0	484	482
225	50	0	236	385	0	385	0	5,468	5,379
225	20	3,500	236	88	0.05	100	0.1	1,504	1,455
225	50	3,500	236	385	0.07	385	0.1	6,610	6,131

Table 1. 20 instances for the Nikkei225 index with their parameters and HQPU solutions. For the best solution, values marked with * have been proven optimal

In Table 2 the results of all benchmark tooling per test instance is given. In the table the results are shown of the hybrid solver (HQPU), Simulated Annealing (SA), Genetic Algorithm (GA), the Gurobi solver (GB) and LocalSolver (LS). We see that LocalSolver finds for all instances the best solution. Gurobi also finds this solution, however, it runs out of memory (locally) for instances bigger than $N = 150$. For the instances it solves, it proves optimality by closing the optimality gap to 0%. LocalSolver is not able to close this optimality gap within reasonable time for most larger instances. The HQPU gives reasonable results, optimal for the smaller instances and well within 5% of the optimal solution for the larger cases, with exception for the last instance. The Simulated Annealing stays close to the HQPU solution and the Genetic Algorithm implementation underperforms in comparison with the other solvers.

The performance is also depending on the computation time. In Table 3 the calculation times are listed. Here only the time the solver requires are mentioned,

N	n	R^*	HQPU	SA	GA	GB	LS	Best Solution
50	10	0	256	256	256	256	256	256*
50	25	0	3,035	3,035	3,236	3,035	3,035 [†]	3,035*
50	10	1,200	321	321	321	321	321	321*
50	25	1,200	3,058	3,058	3,236	3,058	3,058	3,058*
100	20	0	809	809	1,102	809	809 [†]	809*
100	50	0	10,999	10,999	12,318	10,999	10,999	10,999*
100	20	2,500	1,126	<i>1,135</i>	1,450	1,126	1,126	1,126*
100	50	2,500	11,065	11,065	12,397	11,065	11,065	11,065*
150	20	0	628	628	1,047	628	628 [†]	628*
150	50	0	7,993	<i>8,321</i>	10,561	7,993	7,993 [†]	7,993*
150	20	3,000	<i>1,271</i>	<i>1,260</i>	2,937	1,256	1,256 [†]	1,256*
150	50	3,000	<i>8,315</i>	8,270	12,659	8,270	8,270 [†]	8,270*
200	20	0	<i>530</i>	<i>534</i>	1,151	-	529 [†]	529
200	50	0	<i>6,000</i>	<i>6,030</i>	9,397	-	5,950 [†]	5,950
200	20	3,250	<i>1,308</i>	<i>1,312</i>	2,525	-	1,262 [†]	1,262
200	50	3,250	6,859	6,983	11,804	-	6,500 [†]	6,500
225	20	0	<i>484</i>	<i>489</i>	1,143	-	482 [†]	482
225	50	0	<i>5,468</i>	<i>5,472</i>	9,438	-	5,379 [†]	5,379
225	20	3,500	<i>1,504</i>	1,754	4,276	-	1,455 [†]	1,455
225	50	3,500	6,610	6,835	11,802	-	6,131 [†]	6,131

Table 2. Results of all the methods. Values marked with * are proven optimal. Values marked with [†] for LocalSolver are not proven optimal, solver stops before GAP=0. Bold values in HQPU, SA, GA columns represent that they achieved optimal solutions while italicised values represent solutions within 5 % of the solution . All solutions in GB and LS are optimal in nature.

the time to build the problem is out of scope. Best solving times are realised by Gurobi. LocalSolver performs well regarding to finding a good solution. However, for most cases the optimality gap is not closed to 0% which means the solver keeps running until the maximum admitted time is achieved. Simulated Annealing and Genetic algorithms have increasing calculation times for the instances. The HQPU is quite fast and independent from the instance size here. The times displayed in the table correspond to one call to the hybrid solver, where we used 5 calls per instance. D-Wave allows the user to set a time limit. When the user does not specify this, a minimum time limit is calculated and used. For all instances here we did not adjust the time limit.

The second exercise we perform on the bigger S&P500. The S&P500 is a stock market index that measures the stock performance of 500 large companies listed on stock exchanges in the United States. It is one of the most commonly followed equity indices. The S&P500 index is a capitalisation-weighted index and the 10 largest companies in the index account for 26% of the market capitalisation of the index: Apple Inc., Microsoft, Amazon.com, Alphabet Inc., Facebook, Johnson & Johnson, Berkshire Hathaway, Visa Inc., Procter & Gamble and JPMorgan Chase. For this index we created five instances, which are depicted with the results of the solvers in Table 4 and Table 5.

N	n	R^*	HQPU	SA	GA	GB	LS
50	10	0	1.0	8	58	< 1	2 (210)
50	25	0	1.3	7	55	< 1	2 (>600)
50	10	1,200	1.0	2	53	< 1	2 (35)
50	25	1,200	0.9	2	61	< 1	2 (20)
100	20	0	1.1	9	89	< 1	3 (>600)
100	50	0	1.1	20	88	< 1	2 (76)
100	20	2,500	1.0	9	82	< 1	2 (556)
100	50	2,500	1.1	9	86	< 1	3 (43)
150	20	0	1.3	37	128	< 1	2 (>600)
150	50	0	1.4	51	130	< 1	4 (>600)
150	20	3,000	1.4	2	71	< 1	3 (>600)
150	50	3,000	1.2	7	74	< 1	16 (>600)
200	20	0	1.2	26	157	-	3 (>600)
200	50	0	1.1	33	161	-	4 (>600)
200	20	3,250	1.1	28	106	-	21 (>600)
200	50	3,250	1.1	32	128	-	7 (>600)
225	20	0	1.2	31	169	-	3 (>600)
225	50	0	1.1	43	168	-	12 (>600)
225	20	3,500	1.3	32	131	-	165 (>600)
225	50	3,500	1.2	37	140	-	207 (>600)

Table 3. Pure solver times in seconds of all the methods. For LocalSolver the time to find the best binary solution is given and between brackets the time to prove optimality. The HQPU time is per query, in the analysis we used 5 queries per problem and a parameter grid search was performed prior to these runs. Pure solver means time required just for solving the problem and doesn't include the initialisation time or the time needed to send the problem to the Leap network.

Again we see that Gurobi is not able to run instances bigger than $N = 150$. LocalSolver finds the best solutions but is not able to close the optimality gap. In the three largest instances the gap stayed 100% within the allowed 600s. Genetic algorithms performed poorly again. The hybrid approach and Simulated Annealing were able to stay close to the LocalSolver solutions, within 3% except for the last instance. Although we gave the HQPU more time to solve, we adjusted the time limit by hand, it stayed on 18% above the best solution found by LocalSolver, as the Simulated Annealing solver did. The solving times of the HQPU stay quite reasonable as compared to the other solvers.

6 Conclusions and Further research

Quantum computing is still in its early stage. However, the moment of quantum computers outperforming classical computers is coming closer. In this phase, hybrid solvers, using classical methods combined with quantum computing, are promising. The quantum paradigm that is already maturing is quantum annealing combined with the Hybrid Solvers that D-Wave is offering. In this paper we showed the performance of this hybrid approach, applied to a quadratic optimi-

N	n	R^*	Size(Q)	HQPU	SA	GA	GB	LS	Best
100	50	3500	112	5518	5518	6959	5518	5518 [†]	5518*
200	50	3500	213	3121	3141	5896	-	3121 [†]	3121
300	50	3500	314	2492	2525	6825	-	2485 [†]	2485
400	50	3500	414	1537	1677	6846	-	1504 [†]	1504
500	50	3500	515	1287	1570	9014	-	1286 [†]	1286

Table 4. Results of all the methods. Values marked with * are proven optimal. Values marked with [†] for LocalSolver are not proven optimal, solver stops before GAP=0.

N	n	R^*	HQPU	SA	GA	GB	LS
100	50	3500	1.0	8	61	< 1	2 (>600)
200	50	3500	1.3	28	92	-	3 (>600)
300	50	3500	1.6	52	124	-	7 (>600)
400	50	3500	1.9	92	155	-	18 (>600)
500	50	3500	5.8	135	221	-	16 (>600)

Table 5. Pure solver times in seconds of all the methods. For LocalSolver the time to find the best binary solution is given and between brackets the time to prove optimality. The HQPU time is per query, in the analysis we used 5 queries per problem and a parameter grid search was performed prior to these runs.

sation problem occurring in finance, Portfolio Optimisation. Here a portfolio of assets is selected, minimised the total risk of the portfolio. This portfolio has to meet some return and budget constraints.

We suggested an implementation of this problem on the newest D-Wave quantum annealer using the hybrid solver. The performance on a problem where 50 stocks needs to be selected from an existing index, the Nikkei225 or the S&P500, was already reasonable, in comparison with the performance of other commercially available solvers. Of these solvers, LocalSolver performed the best, finding the optimal solution in all cases very quickly. However, proving optimality, and thus finishing the optimisation task, was not realised within reasonable time. The hybrid quantum solver was able to find a solution within 3% of the optimal solution for the S&P500 index up to 400 stocks. Further improvement of the hybrid solvers and the enlargement of the QPU in the coming years lead to the expectation that this computing paradigm is close to real business applications.

For further research we recommend to dive into the hybrid approach and look for improvements. The used method is a standard method, whereas D-Wave offers the opportunity to tailor the methodology in the Leap environment. Also other quantum optimisation algorithms could be used, like the Quantum Approximate Optimization Algorithm algorithm on a gate model quantum computer.

Acknowledgments

The authors want to thank Hans Groenewegen for providing the financial data underlying this research and Irina Chiscop for her valuable review and comments.

References

1. Benoist, T., Estellon, B., Gardi, F., Megel, R., Nouioua, K.: Localsolver 1.x: A black-box local-search solver for 0-1 programming. *4OR quarterly journal of the Belgian, French and Italian Operations Research Societies* pp. 299–316 (2011)
2. Calaza, C.D.G., Willsch, D., Michielsen, K.: Garden optimization problems for benchmarking quantum annealers (2021)
3. Coffrin, C.J.: Challenges with chains: Testing the limits of a d-wave quantum annealer for discrete optimization. Tech. rep., Los Alamos National Lab.(LANL), Los Alamos, NM (United States) (2019)
4. Cohen, J., Khan, A., Alexander, C.: Portfolio optimization of 40 stocks using the dwave quantum annealer. arXiv preprint arXiv:2007.01430 (2020)
5. Cohen, J., Khan, A., Alexander, C.: Portfolio optimization of 60 stocks using classical and quantum algorithms. arXiv preprint arXiv:2007.08669 (2020)
6. Del Pia, A., Dey, S.S., Molinaro, M.: Mixed-integer quadratic programming is in np. *Mathematical Programming* **162**(1-2), 225–240 (2017)
7. Domino, K., Koniorczyk, M., Krawiec, K., Jałowicki, K., Gardas, B.: Quantum computing approach to railway dispatching and conflict management optimization on single-track railway lines (2021)
8. Elsokkary, N., Khan, F.S., La Torre, D., Humble, T.S., Gottlieb, J.: 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)
9. Farhi, E., Harrow, A.W.: Quantum supremacy through the quantum approximate optimization algorithm. arXiv preprint arXiv:1602.07674 (2016)
10. Fulga, C., Stanojević, B.: Single period portfolio optimization with fuzzy transaction costs. In: 20th International Conference EURO Mini Conference” Continuous Optimization and Knowledge-Based Technologies”, EurOPT 2008. Vilnius Gediminas Technical University (2008)
11. Gacon, J., Zoufal, C., Woerner, S.: Quantum-enhanced simulation-based optimization. arXiv preprint arXiv:2005.10780 (2020)
12. Garey, M.R., Johnson, D.S.: *Computers and intractability*, vol. 174. freeman San Francisco (1979)
13. Gemeinhardt, F.G.: *Quantum Computing: A Foresight on Applications, Impacts and Opportunities of Strategic Relevance*. Ph.D. thesis, Universität Linz (2020)
14. Glover, F., Kochenberger, G., Du, Y.: A tutorial on formulating and using qubo models. arXiv preprint arXiv:1811.11538 (2018)
15. Gurobi Optimization, L.: Gurobi optimizer reference manual (2020), <http://www.gurobi.com>
16. Kerenidis, I., Prakash, A., Szilágyi, D.: Quantum algorithms for portfolio optimization. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. pp. 147–155 (2019)
17. Li, D., Sun, X., Gu, S., Gao, J., Liu, C.: Polynomially solvable cases of binary quadratic programs. In: *Optimization and Optimal Control*, pp. 199–225. Springer (2010)

18. Liagkouras, K., Metaxiotis, K.: Multi-period mean–variance fuzzy portfolio optimization model with transaction costs. *Engineering applications of artificial intelligence* **67**, 260–269 (2018)
19. Lucas, A.: Ising formulations of many np problems. *Frontiers in Physics* **2**, 5 (2014)
20. Markowitz, H.: *Harry Markowitz: selected works*, vol. 1. World Scientific (2009)
21. Marzec, M.: Portfolio optimization: Applications in quantum computing. *Handbook of High-Frequency Trading and Modeling in Finance* pp. 73–106 (2016)
22. McGeoch, C.C.: Adiabatic quantum computation and quantum annealing: Theory and practice. *Synthesis Lectures on Quantum Computing* **5**(2), 1–93 (2014)
23. Möller, M., Vuik, C.: On the impact of quantum computing technology on future developments in high-performance scientific computing. *Ethics and Information Technology* **19**(4), 253–269 (2017)
24. Neumann, N., Phillipson, F., Versluis, R.: Machine learning in the quantum era. *Digitale Welt* **3**(2), 24–29 (2019)
25. Orus, R., Mugel, S., Lizaso, E.: Quantum computing for finance: overview and prospects. *Reviews in Physics* **4**, 100028 (2019)
26. Pang, T., Hussain, A.: A stochastic portfolio optimization model with complete memory. *Stochastic Analysis and Applications* **35**(4), 742–766 (2017)
27. Pang, T., Varga, K.: Portfolio optimization for assets with stochastic yields and stochastic volatility. *Journal of Optimization Theory and Applications* **182**(2), 691–729 (2019)
28. Piattini, M., Peterssen, G., Pérez-Castillo, R., Hevia, J.L., Serrano, M.A., Hernández, G., de Guzmán, I.G.R., Paradela, C.A., Polo, M., Murina, E., et al.: The talavera manifesto for quantum software engineering and programming. In: *QANSWER*. pp. 1–5 (2020)
29. Radulescu, M., Radulescu, C.Z.: A multi-objective approach to multi-period: Portfolio optimization with transaction costs. In: *Financial Decision Aid Using Multiple Criteria*, pp. 93–112. Springer (2018)
30. Reberntrost, P., Lloyd, S.: Quantum computational finance: quantum algorithm for portfolio optimization. *arXiv preprint arXiv:1811.03975* (2018)
31. Resch, S., Karpuzcu, U.R.: Quantum computing: an overview across the system stack. *arXiv preprint arXiv:1905.07240* (2019)
32. Ronagh, P., Woods, B., Iranmanesh, E.: Solving constrained quadratic binary problems via quantum adiabatic evolution. *Quantum Information & Computation* **16**(11-12), 1029–1047 (2016)
33. Skaf, J., Boyd, S.: Multi-period portfolio optimization with constraints and transaction costs. In: *Working Manuscript*. Citeseer (2009)
34. Venturelli, D., Kondratyev, A.: Reverse quantum annealing approach to portfolio optimization problems. *Quantum Machine Intelligence* **1**(1-2), 17–30 (2019)
35. Xidonas, P., Mavrotas, G., Hassapis, C., Zopounidis, C.: Robust multiobjective portfolio optimization: A minimax regret approach. *European Journal of Operational Research* **262**(1), 299–305 (2017)
36. Zanjirdar, M.: Overview of portfolio optimization models. *Advances in Mathematical Finance and Applications* **5**(4), 1–16 (2020)