

Working Paper Series — 1/2023

Finding the Optimal Currency Composition of Foreign Exchange Reserves with a Quantum Computer

Martin Veselý



Czech National Bank — Working Paper Series — 1/2023

CNB CZECH
NATIONAL
BANK

www.cnb.cz

The Working Paper Series of the Czech National Bank (CNB) is intended to disseminate the results of the CNB's research projects as well as the other research activities of both the staff of the CNB and collaborating outside contributors, including invited speakers. The Series aims to present original research contributions relevant to central banks. It is refereed internationally. The referee process is managed by the CNB Economic Research Division. The working papers are circulated to stimulate discussion. The views expressed are those of the authors and do not necessarily reflect the official views of the CNB.

Distributed by the Czech National Bank, available at www.cnb.cz

Reviewed by: Jakub Mareček (Czech Technical University in Prague)

Ivan Sutóris (Czech National Bank)

Project Coordinator: Martin Hodula

Issued by: © Czech National Bank, February 2023

Finding the Optimal Currency Composition of Foreign Exchange Reserves with a Quantum Computer

Martin Veselý *

Abstract

Portfolio optimization is an inseparable part of strategic asset allocation at the Czech National Bank. Quantum computing is a new technology offering algorithms for that problem. The capabilities and limitations of quantum computers with regard to portfolio optimization should therefore be investigated. In this paper, we focus on applications of quantum algorithms to dynamic portfolio optimization based on the Markowitz model. In particular, we compare algorithms for universal gate-based quantum computers (the QAOA, the VQE and Grover adaptive search), single-purpose quantum annealers, the classical exact branch and bound solver and classical heuristic algorithms (simulated annealing and genetic optimization). To run the quantum algorithms we use the IBM QuantumTM gate-based quantum computer. We also employ the quantum annealer offered by D-Wave. We demonstrate portfolio optimization on finding the optimal currency composition of the CNB's FX reserves. A secondary goal of the paper is to provide staff of central banks and other financial market regulators with literature on quantum optimization algorithms, because financial firms are active in finding possible applications of quantum computing.

Abstrakt

Optimalizace portfolia je v ČNB nedílnou součástí strategické alokace aktiv. Kvantové počítače představují novou technologii, která nabízí nástroje pro tuto úlohu. Z tohoto důvodu, je vhodné prozkoumat schopnosti a omezení kvantových počítačů právě v oblasti optimalizace portfolia. V tomto článku se zaměřujeme na aplikace kvantových algoritmů pro dynamickou optimalizaci portfolia vycházející z Markowitzova modelu. Konkrétně se zabýváme porovnáním algoritmů pro univerzální kvantové počítače postavené na modelu kvantových hradel (QAOA, VQE a Groverovo adaptivní vyhledávání), jednoúčelových kvantových annealerů, klasické exaktní metody větví a mezí a klasických optimalizačních heuristických algoritmů (simulované žíhání a genetická optimalizace). Kvantové algoritmy provozujeme v prostředí univerzálního kvantového počítače IBM QuantumTM. Dále využíváme kvantový annealer od společnosti D-Wave. Optimalizaci portfolia s pomocí kvantových algoritmů demonstrujeme na hledání optimální měnové struktury devizových rezerv ČNB. Sekundárním cílem tohoto článku je vytvoření literatury ke kvantové optimalizaci pro pracovníky centrálních bank a dalších regulátorů finančních trhů, jelikož finanční instituce aktivně hledají možné aplikace kvantových počítačů.

JEL Codes: C61, C63, G11.

Keywords: Foreign exchange reserves, portfolio optimization, quadratic unconstrained binary optimization, quantum computing.

*Martin Veselý, Czech National Bank, Risk Management Department, martin.vesely@cnb.cz

This work was supported by Czech National Bank Research Project No. P2/22.

The author would like to thank Simona Malovaná and Martin Hodula for invaluable advice and project supervision.

The views expressed in this paper are those of the author and not necessarily those of the Czech National Bank.

1. Introduction and Motivation

Over the last few years, knowledge of quantum computing has been penetrating the finance industry (Orús et al., 2019; Egger et al., 2020). This development has naturally attracted the attention of central banks and regulators (Hull et al., 2020). As Hull et al. (2020) is rather theoretical, we tried to assess the practical capabilities of quantum computers ourselves (Veselý, 2022b). Although the results concerning portfolio optimization with a quantum linear equation solver (the HHL algorithm) and risk evaluation with quantum Monte Carlo were discouraging, a quantum heuristic for quadratic unconstrained binary optimization (QUBO) – employed again in portfolio optimization – worked as expected. We therefore decided to further investigate quantum QUBO techniques. We will demonstrate them and assess their capabilities using dynamic portfolio optimization, in particular optimization of the currency composition of the CNB’s FX reserves, a highly important part of its strategic asset allocation process. Because of the latter, the paper contributes to the discussion on optimal reserve currency composition, although our purely technical approach to optimization is not intended as a replacement for expert judgment. On top of that, the article serves as a practical guide to quantum QUBO algorithms for staff of central banks and other financial market regulators.

Before we start our investigation of quantum QUBO algorithms in connection with portfolio optimization, it is worth listing other possible applications of QUBO in finance and general business. There are many problems which can be converted to optimization in a graph and finally to quantum QUBO. Lucas (2014) provides a (non-exhaustive) list. An outstanding example of graph optimization is the travelling salesperson problem, which can be employed to find arbitrage opportunities, as shown by Orús et al. (2019). Besides graph problems, systems of linear equations often arise in real world problems. Since the HHL algorithm designed by Harrow et al. (2009) seems to be far from practical deployment, as noted by Childs (2009) and Aaronson (2015), a linear system solver leveraging quantum QUBO was proposed by Wen et al. (2019). Linear regression, another technique often used in finance, is a particular application of linear systems. The evaluation of regression coefficients with quantum QUBO is discussed by Potok and Date (2021). Linear systems can also be employed in the numerical solution of partial differential equations in finance, for example in option pricing. For this purpose, a modification of a quantum QUBO heuristic was proposed by Fontanela et al. (2021). An example of the application of quantum computers in macroeconomics was presented by Fernández-Villaverde and Hull (2022), who designed a QUBO-based algorithm for finding the optimal choice between consumption and the capital level in a real business cycle model. Interestingly, quantum QUBO for factoring integers to prime constituents was presented by Anschuetz et al. (2019) as a temporary replacement for the algorithm designed by Shor (1994), which currently suffers from quantum hardware deficiencies. Although integer factorization is not relevant for FX reserves management, it is important for the security departments of banks and other institutions, because quantum factoring algorithms could potentially break the RSA ciphering mechanism. This shortlist of possible applications shows the importance of QUBO in finance and other fields. Interested readers can find a vast overview of real-world QUBO applications outside finance (such as the pharmaceuticals and logistics industries) employing quantum approaches in Yarkoni et al. (2022).

The focal point of this article, however, is portfolio optimization. There are several studies on portfolio optimization employing quantum QUBO. In Elsokkary et al. (2017), it is simply stated that portfolio optimization in Markowitz-like fashion can be carried out on a quantum computer. The authors of the study used binary flags only, i.e. they found out whether to invest in a particular asset or not. Rosenberg et al. (2016) worked with integer asset weights and showed how to impose constraints via a penalty function. Palmer et al. (2021) demonstrated a technique for reducing

a higher-order objective function, arising as a result of portfolio volatility targeting, back to a quadratic one solvable with QUBO algorithms. In contrast to previous studies, the authors used real-world data on the S&P 100 and S&P 500 equity indices. Mugel et al. (2022) added the time dimension to make the portfolio optimization dynamic. Additionally, they provided a performance comparison of some quantum techniques with classical ones and demonstrated the superiority of the quantum approach for problems with more than 20 binary variables. Finally, Palmer et al. (2022) presented a non-convex QUBO model for replication of a benchmark with a specified number of securities from the benchmark (the “cardinality constraint”).

To perform our QUBO problem, we will first use methods for the simulation of spin-glasses (the Ising model) on the IBM QuantumTM universal gate-based quantum computer, specifically the Quantum Approximate Optimization Algorithm (QAOA) designed by Farhi et al. (2014) and the Variational Quantum Eigensolver (VQE) proposed by Peruzzo et al. (2014). We will also employ a single-purpose quantum computer intended only for solving QUBO problems – the *quantum annealer* provided by D-Wave. These three techniques are known as the *adiabatic approach*.

As pointed out by Preskill (2018), adiabatic approach-based algorithms seem to be the first real-world applications of quantum computers. However, their main disadvantage is unproven speed-up. Despite this fact, several studies (Ding et al., 2023; Harwood et al., 2021; Egger et al., 2021; Mugel et al., 2022; Fernández-Villaverde and Hull, 2022) indicated that in particular instances the algorithms perform better than, or at least equally well as, their classical counterparts. On the other hand, Mandra et al. (2017) questioned the claimed advantages of the adiabatic approach. The study also emphasized that often only the simulated annealing developed by Metropolis et al. (1953) is used as a benchmark in the performance assessment of adiabatic approach-based algorithms. Schuld and Killoran (2022) took a “Goldilocks” stance that further research is needed, especially once quantum computers with a higher number of qubits become available. Tasseff et al. (2022) also pointed out the need for further research and provided empirical evidence of the better performance of the D-Wave quantum annealer in hardware-tailored benchmark problems. The discussion above indicates that there is no consensus on the capabilities of the algorithms. To shed more light on this issue, we will assess the performance of adiabatic approach-based algorithms ourselves. We will expand the set of classical algorithms used in the comparison to include the branch and bound exact solver designed by Land and Doig (1960) and the genetic-based heuristic algorithm proposed by Rechenberg (1973) and Schwefel (1977). To compare the QUBO approach to portfolio optimization with classical continuous quadratic programming, we will also employ the classical continuous gradient method.

Besides the algorithms described above, we will also test the algorithm proposed by Gilliam et al. (2021) based on the quantum database searching developed by Grover (1996). In contrast to adiabatic approach-based algorithms, the Grover algorithm offers proven quadratic speed-up, as does the derived QUBO algorithm. To the best of our knowledge, there is no study concerning the practical capabilities of the Grover QUBO optimizer. We will therefore address this gap in our paper.

Interestingly, all the studies on portfolio optimization mentioned above work with equity market data. We will use currency market data instead to find out whether there is any impact on the performance of quantum algorithms, particularly those without rigorously proven speed-up.

We tested all the algorithms discussed above on the currency composition optimization of the CNB’s FX reserves. We found that the best performance from the perspective of both run time and ability to find the global optimum is offered by the *hybrid heuristic* (exploiting a combination of the classical

and quantum approaches) provided by D-Wave. Moreover, the hybrid heuristic is mature enough to allow us to test a problem of practical size. In the cases of the VQE and the QAOA, we tested only toy models on real quantum hardware. Grover adaptive search was tested only on a simulator due to the limited number of qubits on the available quantum processors. The exact branch and bound algorithm is able to find the solution to the QUBO version of the problem provided that we have enough memory (more than 8 GB RAM), but the run time is still several hours. Interestingly, we found that the gradient solver implemented in MS Excel outperforms all of the algorithms tested (both the classical and the quantum ones) once the problem is formulated as a continuous problem. Therefore, we concluded that carrying out portfolio optimization with the QUBO approach is not the best option, at least for the time being. However, as portfolio optimization in binary form seems to be a hard-to-solve problem, it can serve as a testing problem for new algorithms or for the measurement of progress in quantum hardware development.

Our results do not lessen the role of quantum computing in finance. We have to bear in mind that quantum computers are not a fully mature technology yet. We need to wait several years, track the development of technology, build up our “quantum knowledge” and then pass a final verdict on the usefulness of quantum computers for finance applications. We should also critically assess which financial problems quantum computers are a suitable option for and whether classical computing may offer better algorithms.

The rest of this paper is organized as follows. The second section provides the basics on the application of QUBO in portfolio optimization, the third section introduces the quantum algorithms employed and briefly discusses classical ones, the fourth section details the currency composition optimization problem and discusses the data used and the results, and finally the fifth section concludes. Note that the basics of quantum computing and the necessary mathematical background and notation are provided in our previous article (Veselý, 2022b).

2. Quadratic Binary Portfolio Optimization

In our research, we employ portfolio optimization based on quadratic programming as introduced by Markowitz (1952). The model assumes real (continuous) weights on the assets. However, quantum computers are not able to work with real variables at their current stage of development. Durr and Hoyer (1996) presented a basic idea of how the quantum database search designed by Grover (1996) could be used for minimization of a general function but on a discrete domain. Hence, this algorithm would need too many qubits to approximate a continuous function sufficiently. Moreover, the algorithm assumes the existence of quantum memory (qRAM), which is currently rather an experimental device – see the discussion in Veselý (2022b) for details. Ayub et al. (2020) showed that a quantum continuous optimization algorithm based on Grover search would reduce the time T needed to find the global extreme of a function in the classical case to $\sqrt{T} \ln(T)$. However, the authors did not provide any technical detail on the construction of the algorithm. Verdon et al. (2019) introduced a quantum version of the gradient descent method based on the Quantum Approximate Optimization Algorithm of Farhi et al. (2014), but the authors stated that implementation of the algorithm is beyond the capabilities of current quantum computers. This short investigation therefore shows that the original continuous quadratic program used in the Markowitz model has to be modified to one with binary variables (in particular QUBO) in order to be carried out on current quantum computers.

In Veselý (2022b) we provided a toy model of portfolio optimization with QUBO and found that such optimization runs successfully on IBM QuantumTM processors. In this part, we will expand

the toy model to “real world” *dynamic portfolio optimization with discrete time steps*. As discussed by Mugel et al. (2022), there are also static and continuous-time versions of portfolio optimization. However, the discrete-time version is of greater practical value because the static version assumes immutability of portfolio composition, a condition that is hardly fulfilled in practice, and the continuous version works with the unrealistic assumptions that the portfolio changes in each infinitesimal time moment and that assets are perfectly divisible.

The QUBO Problem in General

Before turning to the particular application of QUBO in portfolio optimization, it is worth reminding ourselves of the general form of the QUBO problem. Let us denote $\mathbf{A} \in \mathbb{R}^{n,n}$, $\mathbf{b} \in \mathbb{R}^n$ (a column vector), $\mathbf{x} \in \{0; 1\}^n$ (a column vector composed of n binary variables) and $c \in \mathbb{R}$. With this notation, the QUBO problem is defined as

$$\frac{1}{2}\mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c \rightarrow \min. \quad (1)$$

The problem (1) can be rewritten in “sum form”

$$\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n a_{ij} x_i x_j + \sum_{i=1}^n b_i x_i + c \rightarrow \min. \quad (2)$$

Naturally, the QUBO problem does not involve any constraint. However, in real-world applications, constraints are often required. To incorporate them, a *penalty function* is added to function (2) and increases its value if any constraint is not fulfilled. Techniques for the construction of the penalty function are discussed by Yarkoni et al. (2022).

Portfolio Optimization with QUBO

If matrix \mathbf{A} is replaced with the asset return covariance matrix $\mathbf{C} \in \mathbb{R}^{n,n}$ and multiplied by coefficient $\lambda \geq 0$ expressing risk aversion, vector \mathbf{b} is substituted with average asset returns $\mathbf{r} \in \mathbb{R}^n$ and multiplied by -1 to maximize the returns and $c = 0$, we get the optimization problem $-\mathbf{r}^T \mathbf{x} + \lambda \mathbf{x}^T \mathbf{C} \mathbf{x} \rightarrow \min$, which is nearly the portfolio optimization model proposed by Markowitz (1952). To have the model complete, we have to add a budget constraint $\mathbf{p}^T \mathbf{x} \leq B$, where $\mathbf{p} \in \mathbb{R}^n$ is a vector of the money amounts invested in the assets and $B \in \mathbb{R}$ is the total budget. We can consider that the whole budget is allocated, i.e. $\mathbf{x}^T \mathbf{p} - B = 0$, which leads to the optimization problem

$$-\mathbf{r}^T \mathbf{x} + \lambda \mathbf{x}^T \mathbf{C} \mathbf{x} + F(\mathbf{p}^T \mathbf{x} - B)^2 \rightarrow \min, \quad (3)$$

where F is the importance of the budget constraint. The value of F should be sufficiently large so that the objective function value is prohibitively high for infeasible solutions.

Still, this binary form of Markowitz optimization only tells us whether to invest in a particular asset ($x_i = 1$) or not ($x_i = 0$). However, in practice we want to know the asset weights. To introduce real weights instead of binary flags and at the same time preserve the binary nature of the optimization, we follow the suggestion of Mugel et al. (2022) and define the weight of the i^{th} asset as a *binary fraction* $w_i = \sum_{k=1}^{\ell} 2^{-k} x_i^{(k)}$. Parameter ℓ is the number of bits dedicated to expressing the weight and $x_i^{(k)}$ is a binary variable containing the k^{th} decimal place of the weight.

Replacing the binary flags x_i in (3) with weights w_i leads to the Markowitz portfolio optimization with real asset weights¹

$$-\sum_{i=1}^n r_i \sum_{k=1}^{\ell} 2^{-k} x_i^{(k)} + \lambda \sum_{i=1}^n \sum_{j=1}^n c_{ij} \left(\sum_{k=1}^{\ell} 2^{-k} x_i^{(k)} \right) \left(\sum_{k=1}^{\ell} 2^{-k} x_j^{(k)} \right) + F \left(\sum_{i=1}^n p_i \sum_{k=1}^{\ell} 2^{-k} x_i^{(k)} - B \right)^2 \rightarrow \min. \quad (4)$$

Note that we ignored the sign of weight w_i . Introducing the sign is possible,² but we will assume that short positions are forbidden, hence we can assume that the weights are non-negative. Thanks to the approach used to binarize the weights, it holds that $w_i \in \langle 0; 1 \rangle \forall i$, which is the set of constraints used in the Markowitz model.

To make the problem independent of the budget volume B , we set $B = 1$ and $p_i = 1 \forall i$, which changes the budget constraint to $\sum_i w_i = 1$. This is another constraint generally used in Markowitz portfolio optimization. Problem (4) also ignores the time dimension of the portfolio optimization. It is common in practice to rebalance the portfolio periodically because of changes in a benchmark, or to radically change the portfolio composition during or after major crises. This can be reflected in the optimization by introducing a time-dependent covariance matrix and returns and adding other binary variables containing asset weights for time $t = 1, 2, \dots, T$. What is more, any change in the portfolio structure leads to transaction costs, which we naturally want to minimize. The costs are proportional to the changes in asset weights between time periods $t - 1$ and t . We therefore introduce variables $x_i^{(k,t)}$ expressing the k^{th} decimal place of the i^{th} asset weight at time t . Note that $x_i^{(k,0)} = 0 \forall i, k$.³ The costs incurred in connection with changing the position in the i^{th} asset are $v_i^{(t)} \left| \sum_{k=1}^{\ell} x_i^{(k,t)} - x_i^{(k,t-1)} \right|$, where $v_i^{(t)}$ are unit transaction costs. To preserve the QUBO nature of our optimization problem, we replace the absolute value with the square. After the inclusion of independence of the absolute budget volume, the time dimension and transaction costs, problem (4) is transformed to

$$\begin{aligned} & \sum_{t=1}^T \left[\begin{array}{ll} & \text{time dimension} \\ - \sum_{i=1}^n r_i^{(t)} \sum_{k=1}^{\ell} 2^{-k} x_i^{(k,t)} & \text{return} \\ + \lambda \sum_{i=1}^n \sum_{j=1}^n c_{ij}^{(t)} \left(\sum_{k=1}^{\ell} 2^{-k} x_i^{(k,t)} \right) \left(\sum_{k=1}^{\ell} 2^{-k} x_j^{(k,t)} \right) & \text{risk} \\ + \mu \sum_{i=1}^n v_i^{(t)} \left[\sum_{k=1}^{\ell} 2^{-k} (x_i^{(k,t)} - x_i^{(k,t-1)}) \right]^2 & \text{transaction costs} \\ + F \left(\sum_{i=1}^n \sum_{k=1}^{\ell} 2^{-k} x_i^{(k,t)} - 1 \right)^2 & \text{sum of weights equal 1} \end{array} \right] \rightarrow \min. \quad (5) \end{aligned}$$

Note that parameter $\mu \geq 0$ is the sensitivity of a portfolio manager to the transaction costs. Problem (5) is the final version of the *discrete-time dynamic unconstrained binary portfolio optimization*

¹ In fact, the weights are not real but only rational because of the finite number of binary variables expressing them. However, this is also the case with “real” numbers in classical computing, as memory is always finite.

² Potok and Date (2021) advise rewriting the weight as $w_i = \sum_{k=1}^{\ell} 2^{-k} x_i^{(k),\text{pos}} - \sum_{k=1}^{\ell} 2^{-k} x_i^{(k),\text{neg}}$.

³ This means that the portfolio is composed of cash only at time $t = 0$.

based on the Markowitz model. Note that the more convenient matrix form of (5) is provided in Appendix A.

Computational Complexity

We now turn our attention to the computational complexity of the problem introduced above. QUBO is an exponentially complex problem on a classical computer, as adding one binary variable doubles the possible inputs of the objective function.⁴ It may seem that by moving from the continuous to the QUBO version, we have made the problem harder, but any general problem of quadratic programming with real variables is exponentially complex on a classical computer, as proved by Vavasis (1990) and Vavasis and Pardalos (1991). On the other hand, if matrix \mathbf{A} in the objective function (1) is positive definite and the constraints are convex, the quadratic program can be solved in polynomial time, as shown by Kozlov et al. (1980) and Ye and Tse (1989).

A covariance matrix is positive definite. However, adding a general penalty function can lead to an objective function with an indefinite matrix. Especially in this most difficult version of the problem, quantum algorithms could offer speed-up. What is more, if we were able to show speed-up even for cases that a classical computer can solve in polynomial time (such as the one we present in this paper), we would have stronger empirical proof of better performance of quantum algorithms in some cases of quadratic programming and we would therefore show that they are useful in portfolio optimization in general.

3. Algorithms for QUBO

This part is devoted to describing algorithms for solving QUBO problems. First, we will discuss the classical algorithms which we compare the quantum ones with. After that, we will describe the QAOA and VQE quantum variational algorithms intended for a universal gate-based quantum computer. We will then describe D-Wave's quantum annealer, the physical realization of the adiabatic approach. Finally, a QUBO algorithm based on Grover database quantum searching will be discussed.

3.1 Classical Algorithms

In this subsection we will present classical methods for quadratic optimization. The first one, the gradient method, is intended for solving continuous problems, because the original Markowitz model works with continuous variables. The method will serve as a benchmark allowing us to decide whether the transformation of Markowitz continuous portfolio optimization to QUBO and the application of quantum computers confers any advantage.

As noted in the previous part, QUBO is an exponentially complex problem on a classical computer. This renders an exhaustive search for the solution (the brute force method) impossible. Therefore, more sophisticated algorithms have been devised. First, we will describe the branch and bound algorithm, which is an exact solver for integer programming. Then, we will discuss two optimization

⁴ We use the term *exponentially complex* rather loosely. The precise statement should be that the problem belongs to the \mathbf{NP} complexity class. By stating that the problem is exponentially complex, we mean that currently there is no known algorithm capable of solving the problem in polynomial time. But this does not exclude the possibility of such an algorithm being discovered in the future, as whether $\mathbf{P} = \mathbf{NP}$ or $\mathbf{P} \neq \mathbf{NP}$ remains unresolved. Note that we will use this simplification in the rest of the paper to reflect the current capabilities of computational tools. Readers interested in learning more about complexity theory can consult Nielsen and Chuang (2010), chapter 3.

heuristics, namely simulated annealing and genetic optimization. Note that unlike the branch and bound method, these heuristics are only able to find sub-optimal solutions, as they are based on random searching in the problem domain.

Besides the above-mentioned algorithms, there is a plethora of other QUBO heuristics and exact solvers. Interested readers can consult the overview by Kochenberger et al. (2014). There are also several software packages intended for solving optimization problems including QUBO, for example CPLEXTM provided by IBM (2022) and XPressTM offered by FICO (2022). An open source alternative is BiqCrunch designed by Krislock et al. (2017).

Gradient Method

As noted above, we include gradient optimization among the classical methods for comparing the QUBO approach to portfolio optimization with continuous quadratic programming. For our purposes we will use the implementation of the method in the MS Excel solver. We will not discuss the method further, because it is widely known and long established.⁵ Interested readers can find details in Lemaréchal (2012).

Branch and Bound Algorithm

This algorithm was proposed by Land and Doig (1960) and was originally intended for integer optimization. As binary optimization is a special case of integer optimization, the algorithm can be used for QUBO as well.⁶

The main idea behind the algorithm is to cut off parts of a problem domain that offer no improvement of the objective function value found so far. First, the problem is solved with no integer constraints imposed on the variables. Once the solution is reached, the problem domain is divided into two subsets and a special constraint for some variable x_i is added for each subset. Assume that a is an integer part of the value of variable x_i . Then constraint $x_i \leq a$ is added for the first subset and $x_i \geq a + 1$ for the second one. A similar approach can be employed in the case of binary optimization. However, the constraints are simpler: constraint $x_i = 0$ is added for one subset and constraint $x_i = 1$ for the other. The lower bound of the objective function is calculated for each subset, and the subset with the higher bound is no longer taken into consideration. This effectively reduces the size of the domain which has to be searched for the optimal solution. The process is repeated until the optimal solution is found.

Note that the branch and bound algorithm can degenerate into the brute force method. This means that in the worst case it exhibits exponential complexity, as discussed by Thakoor et al. (2009). The algorithm is exact, that is, it guarantees that the optimal solution will be always found, but, as already mentioned, sometimes at the cost of a long run time.

We will use the implementation of the algorithm in the CPLEXTM software package provided by IBM (2022). The product has a free version with the number of variables and constraints limited to 1,000. Nevertheless, this is sufficient for our purposes.

⁵ The method was proposed by the famous French mathematician Louis Augustin Cauchy back in 1847.

⁶ See the application of the branch and bound algorithm to the travelling salesman problem by Little et al. (1963).

Simulated Annealing

The first heuristic intended for QUBO, simulated annealing, is inspired by the cooling of a hot material. The algorithm was first proposed by Metropolis et al. (1953) for finding thermal equilibrium in chemical processes. Later, it was improved by Corana et al. (1987) for the optimization of general continuous functions. Many special-purpose adaptations of the algorithm have been designed, for example for simulations in statistical physics (Scharf and Izrailev, 1990) and for portfolio optimization (Vesely, 2020). For some problems with an objective function defined on discrete domains, including binary optimization problems, the investigation carried out by Bertsimas and Tsitsiklis (1993) revealed superior behaviour of the algorithm in comparison with other heuristics.

The principle of simulated annealing can be briefly described as follows. As a material cools down, the particles the material is composed of find new positions so that the total energy stored in the material is the lowest possible. The configuration can be likened to a bit string composed of binary variables and the energy to the respective value of the objective function. The new configuration is found by means of random changes in the configuration of the particles. To avoid getting stuck in local minima of the objective function, a configuration with higher energy (i.e. a worse solution) is sometimes allowed.⁷ The particular implementation of the random changes (known as the *noise operator*), the *cooling schedule* (defining the speed of cooling) and the rule for temporary acceptance of a higher energy configuration can be various. We will use our own implementation of the simulated annealing algorithm based on the studies listed above and adapted to binary optimization. Technical details of the implementation are available in Appendix B, and the source code is provided by Vesely (2022a).

Genetic Optimization

The second heuristic, genetic optimization, is based on the simulation of an evolutionary process. Interestingly, the idea of solving computational problems with a biologically inspired system appeared in Turing (1950). He was trying to contradict the statement that a calculation machine cannot think, as expressed by Lovelace (1843), among others, and proposed a computational model based on the mind of a child trying to find a solution to a problem by trial and error, similarly to the evolutionary process. Evolution-based algorithms were investigated more seriously by Rechenberg (1973) and Schwefel (1977). Since then, many genetic algorithms have been proposed. A vast overview of these techniques is provided by Katoch et al. (2021).

In a nutshell, the genetic approach works as follows. The bit string representing the argument of the objective function is considered to be the genome of an organism. The objective function value assigned to the argument is a measure of the *fitness* of the organism. The goal of genetic optimization is to modify the genome so as to achieve the best possible fitness. To do so, *mutation and crossing operators* are introduced. The mutation operator randomly changes the genome; for example, it flips a randomly chosen binary variable(s). The crossing operator picks up two or more organisms from the *population* (i.e. the set of bit strings with “good fitness”) according to a predefined scheme and returns a combination thereof. If the fitness of the newly produced organism is better than that of the organisms “created” so far, it becomes a member of the population. Organisms with worse fitness are removed from the population. The exact definitions of the mutation and crossing operators and the rules for adding and removing organisms from the

⁷ This is a natural process occurring in any matter caused by fluctuations at the atomic level.

population depend on the particular genetic algorithm. We will use the implementation of the algorithm provided by Mirjalili (2018).

3.2 Quantum Adiabatic Approach

The quantum computer was originally proposed by Feynman (1985) as a tool for simulating quantum systems, a problem that is exponentially complex on a classical computer. Later, Feynman's proposal was investigated by Lloyd (1996), who conjectured that quantum systems could in some cases be simulated in polynomial time on a quantum computer. Although Lloyd's conjecture remains unproven,⁸ if we can figure out how to convert an optimization problem (QUBO in our case) into the simulation of a quantum system, we can exploit the potential power of quantum computers (see the list of articles in the Introduction providing empirical evidence that quantum computers can solve some optimization problems faster than their classical counterparts).

To carry out a quantum system simulation, the quantum computer has to get a description of the system. In quantum physics, such description is the *Hamiltonian* (or energy operator). It characterizes the allowed energy levels of the system and the energy the system takes on at a particular level. Among energy levels, a special place is dedicated to the *ground state*, the state with the lowest possible energy and associated with the optimal solution of the optimization problem. In mathematical terms, the Hamiltonian is generally a Hermitian operator. However, in the following text, we will assume the Hamiltonian to be a Hermitian matrix \mathcal{H} . The energy of the quantum system described by \mathcal{H} being in state $|\psi\rangle$ is given by the expression $E = \langle \psi | \mathcal{H} | \psi \rangle$.⁹

A special kind of Hamiltonian is the *Ising Hamiltonian*, defined as

$$\mathcal{H}_{\text{Ising}} = \sum_{i=1}^n \sum_{j=1}^n Q_{ij} \mathbf{Z}_i \otimes \mathbf{Z}_j + \sum_{i=1}^n c_i \mathbf{Z}_i, \quad (6)$$

where \mathbf{Z}_i denotes a quantum gate composed of a \mathbf{Z} gate applied to the i^{th} qubit and identity operators applied to the other qubits, and term $\mathbf{Z}_i \otimes \mathbf{Z}_j$ means a quantum gate composed of two \mathbf{Z} gates applied to the i^{th} and j^{th} qubits while identity operators are applied to the other qubits and $Q_{ij}, c_i \in \mathbb{R}$.¹⁰ There is a clear resemblance between the Ising Hamiltonian (6) and the general QUBO problem (2). The ground state of the Ising Hamiltonian (6) can be identified with the optimal solution of problem (2). This means that we have translated the QUBO problem to the simulation of a quantum system. Technical details of the translation, mainly the relations between the coefficients and variables in the QUBO problem (2) and the Ising Hamiltonian (6), are provided by Lucas (2014).

The next step is to simulate the Ising Hamiltonian and search for its ground state.¹¹ To do so, we leverage the *adiabatic theorem*. Assume that the quantum system is described by the

⁸ Aharonov et al. (2022) showed that random circuit sampling, a problem for which Google has claimed a quantum advantage and which until recently was considered to be exponentially complex, can be solved under specific circumstances in polynomial time, although the polynomial is of high degree. This may increase the doubts about Lloyd's conjecture, but further research is still necessary.

⁹ For more information on Hamiltonians, see Nielsen and Chuang (2010), chapter 2.2.2. For details concerning the simulation of quantum systems, see de Wolf (2022), chapter 9.

¹⁰ A more precise term than gate would be operator, because Hamiltonian simulation can be carried out on a general quantum computer, not only on a gate-based machine. However, despite the name, a \mathbf{Z} operator and a \mathbf{Z} gate are mathematically the same object.

¹¹ Lloyd (1996) included a system described with Ising Hamiltonians among examples of quantum systems that could be simulated on a quantum computer with exponential speed-up if his conjecture is confirmed.

initial Hamiltonian \mathcal{H}_0 at time $t = 0$. At the same time, the system is in the ground state of \mathcal{H}_0 . Subsequently, the Hamiltonian of the system is changed. At time $t \in (0; 1)$, the system is described by the Hamiltonian

$$\mathcal{H}(t) = (1 - t)\mathcal{H}_0 + t\mathcal{H}_1. \quad (7)$$

If the change is slow enough, the system remains in the ground state of $\mathcal{H}(t)$ for any t . This means that finally (i.e. at time $t = 1$) the system is in the ground state of \mathcal{H}_1 . In our case, we set $\mathcal{H}_1 = \mathcal{H}_{\text{Ising}}$.

Let's turn our attention to Hamiltonian \mathcal{H}_0 . If \mathcal{H}_0 is simple enough, we are able to determine its ground state analytically. Such simple Hamiltonian is

$$\mathcal{H}_0 = \sum_{i=1}^n \mathbf{X}_i, \quad (8)$$

where \mathbf{X}_i denotes a quantum gate composed of an \mathbf{X} gate applied to the i^{th} qubit and identity operators applied to the other qubits. The ground state of Hamiltonian \mathcal{H}_0 is the uniformly distributed superposition $\frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle$, where $|i\rangle$ is the i^{th} basis state of a quantum system composed of n qubits. Knowing the ground state of \mathcal{H}_0 and employing the adiabatic theorem enables us to find out the ground state of $\mathcal{H}_{\text{Ising}}$ associated with the optimal solution of the QUBO problem. For a better understanding of how simulation based on the adiabatic theorem works, see the analogy in Box 1.

Box 1: Adiabatic Theorem Analogy

The adiabatic theorem can be explained with an analogy. Assume we have a sheet of paper scattered with iron filings. This is our system described by the initial Hamiltonian \mathcal{H}_0 . If no external force is applied, the system remains unchanged. Imagine that we slowly move a magnet under the sheet. The filings begin to move and follow the magnetic lines of force. In the end, we will see a typical pattern on the sheet – the iron filings oriented in the direction of the external magnetic field. This is the system described by the final Hamiltonian \mathcal{H}_1 . The systems that exist when the magnet is only partially under the sheet are described by Hamiltonians $\mathcal{H}(t)$. We now turn our attention to the requirement of carrying out the changes slowly. If we move the magnet quickly, some of the filings remain stuck because of friction between them and the sheet. Clearly, these filings resist the magnetic field and they have to have enough energy to do so. As additional energy is needed, the system is clearly not in its ground state. However, slow changes allow the filings to adapt to the increasing magnetic force without getting stuck, hence they need less energy. In other words, the system remains in the ground state all the time.

3.3 Variational Algorithms and Quantum Annealers

The previous section provided the theoretical framework for the simulation of Ising Hamiltonians and equivalent QUBO problems. However, this approach has to be implemented on a real quantum computer, which we will discuss in this part. First, we will present algorithms intended for a gate-based universal quantum computer (*variational algorithms*) and then we will turn our attention to a single-purpose *quantum annealer*.

Variational Algorithms


Aharonov et al. (2007) showed that the adiabatic and universal gate-based quantum computing models are equivalent.¹² This means that algorithms can be designed for a gate-based quantum

¹² To be precise, Aharonov et al. (2007) stated that *the model of adiabatic computation is polynomially equivalent to the standard model of quantum computation*. On the other hand, Preskill (2018) pointed out that this conclusion is valid only for noiseless qubits, and at the cost of adding a high number of ancilla qubits despite the polynomial equivalence.

computer implementing the above-described simulation of Ising Hamiltonians. What is more, the equivalence of the models allows us to expect that algorithms based on the adiabatic model will provide exponential speed-up if the conjecture by Lloyd (1996) is confirmed. At present, we have to assume that those algorithms offer unproven speed-up, but, as discussed in the Introduction, they show promising results in terms of improved computational time for some specific problems. In what follows, we describe two algorithms intended for solving QUBO problems and developed for gate-based quantum computers.

The first is the Quantum Approximate Optimization Algorithm (QAOA) introduced by Farhi et al. (2014) and improved for some kinds of QUBO problems (including portfolio optimization) by Egger et al. (2021). The second is the Variational Quantum Eigensolver (VQE) developed by Peruzzo et al. (2014) and improved by Barkoutsos et al. (2020). Both algorithms are *hybrid*, which means that they combine the classical and quantum approaches. While the evolution of the quantum system state (the hard part) runs on a quantum computer, the calculation of the energy of the quantum state (the easy part) runs on a classical computer. In the first step, a circuit describing the Hamiltonian $\mathcal{H}(t)$ is constructed. Because of parameter t , the circuit has to be parametric as well. Moreover, we have to ensure that the simulation runs slowly to preserve the assumption of the adiabatic theorem. This all means that we have to run the simulation several times with different parameters and find the circuit describing the simulation of $\mathcal{H}(t)$ in the best possible way. In other words, we solve the classical optimization problem $E = \langle \psi(\theta) | \mathcal{H}_{\text{Ising}} | \psi(\theta) \rangle \rightarrow \min_{\theta}$, where θ is a vector of parameters figuring in the quantum circuit. Since an Ising Hamiltonian comprises only two-body interactions, there is a technique described by Barkoutsos et al. (2020) allowing us to calculate E efficiently on a classical computer. The result of the optimization is the vector θ_{optim} allowing us to construct a circuit for finding the ground state of $\mathcal{H}_{\text{Ising}}$.

Figure 1: QAOA Single-Layer Circuit for Two-Variable QUBO Problem



Note: Qubits q_0 and q_1 represent binary variables x_0 and x_1 respectively. The Hadamard gates **H** prepare a uniformly distributed quantum state. Then the Ising Hamiltonian is applied: the **Rz** gate on qubit q_0 together with the two CNOT gates represent a quadratic term x_0x_1 , and the other **Rz** gates are connected with linear terms x_0 and x_1 . The **Rx** gates are related to the initial Hamiltonian \mathcal{H}_0 . Note that the **Rx** and **Rz** gates are the exponentials of the **X** and **Z** gates figuring in the initial and Ising Hamiltonians respectively. Coefficients $Q_{i,j}$ and c_i in the Ising Hamiltonian and parameters β_j and γ_j are hidden in the rotational angles of the **Rx** and **Rz** gates, i.e. $\theta_1 \dots \theta_5$.

Source: Adapted from the result obtained in the IBM QTM environment

Both the QAOA and the VQE are based on the approach described above, but they differ in the circuit employed for finding the ground state. The QAOA uses a circuit described with the quantum gate


$$\mathbf{U}(\beta, \gamma) = \prod_{j=1}^p e^{-i\beta_j \mathcal{H}_0} e^{-i\gamma_j \mathcal{H}_{\text{Ising}}}, \quad (9)$$

where p is the number of circuit layers and β_j and γ_j are parameters connected with the j^{th} layer of the circuit. Number p is a user-defined value. Increasing p leads to a higher probability of finding

the optimal solution, but also increases the complexity of the circuit. For $p \rightarrow \infty$, finding the optimal solution is guaranteed, but for finite p , the algorithm is only a heuristic. An example of a QAOA circuit is provided in Figure 1.

As can be seen from equation (9) and Figure 1, the structure of the Hamiltonian $\mathcal{H}(t)$ is reflected in the quantum circuit. In contrast, the composition of the VQE circuit is very different. The circuit is built from several layers of controlled **Z** gates and **R_Y** gates. The rotational angles of **R_Y** serve as the optimized parameters θ . The structure of the Ising Hamiltonian is not reflected in the circuit. Because of this, the VQE is more general, and the QAOA can be considered a special case of it. An example of a VQE circuit is provided in Figure 2.¹³

Figure 2: VQE Double-Layer Circuit for Three-Variable QUBO Problem



Note: The circuit has two parts – an initial layer composed only of **R_Y** gates and at least one layer composed of **R_Y** gates and controlled **Z** gates. The latter interconnects the qubits and enables the VQE to take the quadratic terms in the QUBO problem into account. The rotational angles of the **R_Y** gates are optimized so that the circuit simulates the quantum system described by the Ising Hamiltonian and finds its ground state.

Source: Adapted from Barkoutsos et al. (2020)

Quantum Annealers

Besides variational algorithms for universal gate-based quantum computers, single-purpose *quantum annealers* exploiting the adiabatic approach have been devised. In this part we will focus on the annealers provided by D-Wave.

D-Wave annealers are in fact a quantum processor based on superconducting qubits.¹⁴ For each qubit, it is possible to set the probability amplitudes of states $|0\rangle$ and $|1\rangle$ – the *bias*. Pairs of qubits can be connected together with *couplers* to prepare an entangled state.¹⁵ The probability amplitudes of the basis states that the entangled state is composed of can be set by manipulating the coupler *strength*. Setting the biases and strengths enables one to programme a Hamiltonian into the annealer. The annealer function is again based on the adiabatic theorem. First, the initial Hamiltonian \mathcal{H}_0 is programmed into the annealer, which is then set to be in state $\frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle$. In the next step, the biases and strengths are slowly changed so that the annealer behaves according to the Hamiltonian $\mathcal{H}(t)$. At the same time, it remains in the ground state of $\mathcal{H}(t)$. In the end, only the Ising Hamiltonian describes the annealer behaviour. Finally, a measurement is carried out to get the ground state of the Ising Hamiltonian (i.e. the solution of the QUBO problem).

¹³ The VQE circuit is often called *ansatz* (an educated guess), because after proper parametrization it can be used for simulating a quantum system and finding its ground state.

¹⁴ See Kjaergaard et al. (2020) for more on superconducting processors.

¹⁵ Note that because of the special topology of the processor, there is no all-to-all connectivity among the qubits. As a result, more than two qubits can be involved in the quadratic terms of an objective function.

As with any other quantum computation, superposition and entanglement are involved in the described process. On top of that, the phenomenon of *quantum tunnelling* plays a role as well. Quantum tunnelling allows the annealer to get out of local minima – “valleys” – by tunnelling through the “walls” of the valley (see Box 2 for more about quantum tunnelling).

Box 2: Quantum Tunnelling

In the classical world, it is impossible for a particle to go through an energy barrier. To overcome the barrier, the particle has to gain a sufficient amount of energy from an outside source. However, in the quantum world, the particle can “borrow” the energy ΔE for time Δt . It holds that the higher the energy, the shorter the time for which the energy can be “borrowed”. With the “borrowed” energy, the particle overcomes the barrier and then the energy is “redeemed”. This process is a consequence of Heisenberg’s relations of uncertainty. In this particular case, the second relation $\Delta E \Delta t \geq h/(4\pi)$ plays a role. Note that h is the Planck constant ($h = 6.6261 \times 10^{-34} \text{m}^2 \text{kg s}^{-1}$).

It is worth noting that due to quantum fluctuations the annealer does not have to remain in the ground state of the Hamiltonian $\mathcal{H}(t)$ all the time. As a result, the solution is only a sub-optimal one and quantum annealers have to be considered a heuristic tool similar to the QAOA and the VQE.

D-Wave also offers *hybrid heuristics*. This approach combines classical heuristic algorithms with quantum annealing. For example, the output of classical heuristics is used as the initial solution for the quantum annealer. The problem can also be broken down into sub-problems and each part solved with a different algorithm – classical or quantum. Finally, the sub-problem results are combined to get the solution of the complete problem. D-Wave offers hybrid heuristics for both unconstrained and constrained optimization.¹⁶

D-Wave (2018) provides a detailed description of quantum annealers. Further information on D-Wave hybrid heuristics is presented in D-Wave (2022a). Complete documentation for the D-Wave LeapTM and OceanTM development environments is provided in D-Wave (2022b). The source codes used in our tests are available at GitHub (Veselý, 2023).

3.4 Grover Adaptive Search

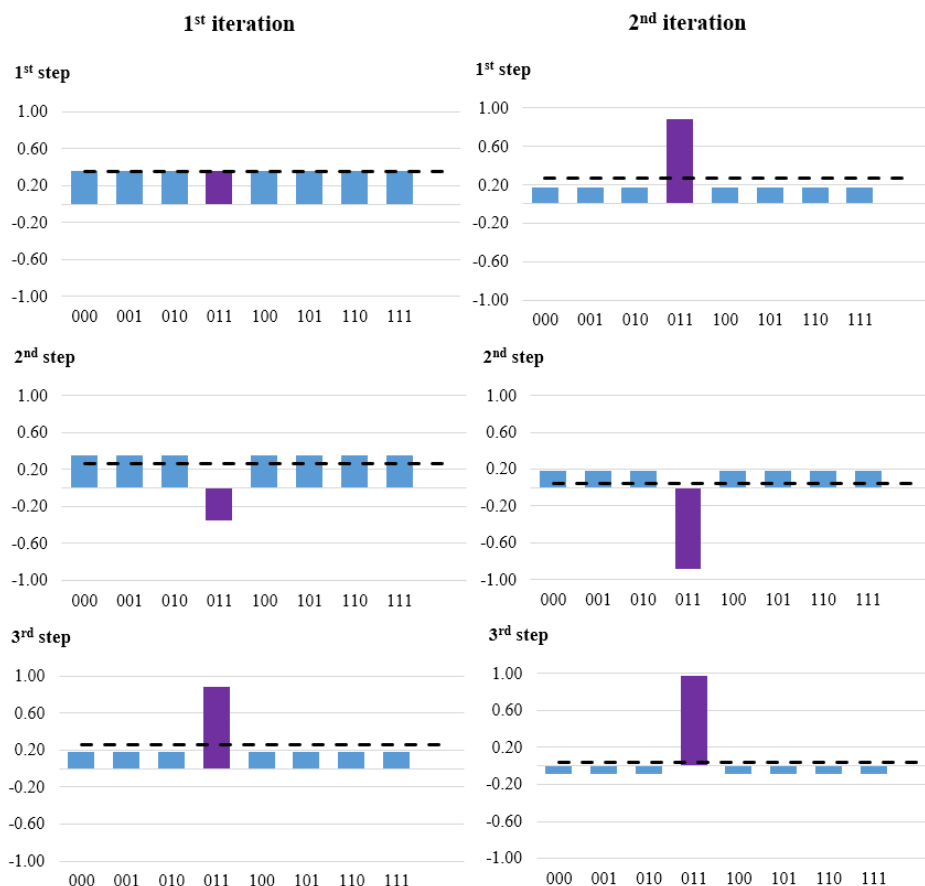
In this section, we will discuss the QUBO method based on the algorithm designed by Grover (1996). Originally, the algorithm was intended for fast search in an unordered database, i.e. a database without an index or a special structure. As we have no index, with a classical computer we have to go through the database records one by one and check whether the record currently picked out is the one we are searching for. In the worst case, the searched record is the last one. Clearly, the problem has linear complexity $O(N)$, where N is the number of records. With Grover quantum search, the complexity drops to $O(\sqrt{N})$, i.e. the algorithm provides quadratic speed-up. Unfortunately, the practicality of the Grover algorithm for database searching was questioned by Viamontes et al. (2005).

Notwithstanding the practical limitation connected with the original purpose of the algorithm, based on ideas presented by Durr and Hoyer (1996), the Grover algorithm was modified to serve as a QUBO algorithm by Gilliam et al. (2021). In contrast to adiabatic approaches, the Grover algorithm offers proven quadratic speed-up. The complexity of QUBO is exponential ($O(2^n)$, n being the number of variables) on a classical computer. The Grover algorithm thus reduces the complexity to $O(2^{n/2})$, which is still exponential in n , but any improvement helps. However, it is important to note that in the current era of noisy quantum computers, the Grover algorithm is not able to provide the

¹⁶ In our case, the constrained heuristic allows constraint $\sum_i w_i^{(t)} = 1 \forall t$ to be expressed outside the objective function.

declared quadratic speed-up. As shown by Chen et al. (2022), its complexity remains the same as in the case of a classical unstructured search algorithm.¹⁷

Figure 3: Illustration of Grover Algorithm Action



Note: Assume that we want to find a record identified with state $|011\rangle$ (highlighted in violet). In the first iteration, a uniformly distributed superposition of all the records is prepared (i.e. the probability of measuring any basis state is 12.5%, or the probability amplitude is 0.3536). Next, the probability amplitude of state $|011\rangle$ is inverted. This leads to a decrease in the average probability amplitude (depicted with a dashed black line). After that, the probability amplitudes are “mirrored” around the average. The quantum state gained is inputted into the second iteration. Then marking and “mirroring” is carried out again. As a result, the probability amplitude of state $|011\rangle$ is close to 1. We found the desired record after two iterations. On a classical computer, the search could take up to eight iterations.

Source: Author’s own creation

First, we will outline how the original Grover algorithm works. Assume that a database contains $N = 2^n$ records, where n is the number of qubits used for representation of the database records. First of all, superposition $\frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle$, containing all the database records with the same probability, is prepared. The next step is to mark the basis state representing the record we are searching for by inverting its probability amplitude. As a result, the average probability amplitude of the quantum state decreases. After that, all the probability amplitudes are “mirrored” around the average amplitude. This leads to an increase in the probability amplitude of the marked state and a decrease in the other amplitudes. The marking and “mirroring” steps are repeated until the probability

¹⁷ Interestingly, this holds for any algorithm offering quadratic speed-up, including quantum Monte Carlo.

amplitude of the marked state is not close to 1. Finally, the qubits are measured to reveal the marked state, or in other words to return the database record we were searching for. Note that for a database with N records, we have to repeat the marking and “mirroring” $R = \lceil \pi\sqrt{N}/4 \rceil$ times. The described process is graphically illustrated in Figure 3 on a search for state $|011\rangle$ in a database containing eight records.


We now turn our attention to the technical realization of the Grover algorithm. The circuit implementing the algorithm is depicted in Figure 4. Without loss of generality, we will assume that we are searching for record $|011\rangle$ in a database stored within three qubits. We refer to those qubits as *working qubits*. At the beginning, the working qubits are set to be in equally distributed superposition with Hadamard gates. Regardless of the number of working qubits, one *ancilla qubit* is added. The ancilla is initialized to state $|-\rangle$. This means that at the beginning, the quantum computer is in state

$$|\psi_{\text{init}}\rangle = \left(\frac{1}{\sqrt{8}} \sum_{i=0}^7 |i\rangle \right) \otimes |-\rangle. \quad (10)$$

After that, a so-called *oracle* is applied to mark the record we are searching for. In fact, the oracle is a multi-qubit controlled \mathbf{X} gate negating the target qubit if a logical expression based on the control qubits values is true. In our example, the oracle negates the ancilla if the working qubits are in state $|011\rangle$, otherwise the ancilla remains unchanged. Negation of $|-\rangle$ leads to state $-|-\rangle$, hence state $|\psi_{\text{init}}\rangle$ is changed to¹⁸

$$|\psi_{\text{oracle}}\rangle = \left(\frac{1}{\sqrt{8}} \sum_{i \in \{0,1,2,4,5,6,7\}} |i\rangle - \frac{1}{\sqrt{8}} |011\rangle \right) \otimes |-\rangle. \quad (11)$$

Figure 4: Grover Algorithm Circuit



Note: Panel A depicts a general Grover algorithm circuit for three working qubits. The part of the circuit between the dashed lines (i.e. the oracle and the diffusion operator) is repeated R times. Note that w stands for working qubits and anc for ancilla qubit. Panel B shows the oracle for marking state $|011\rangle$. For a higher number of qubits, the circuit can be scaled up following the patterns in the initial phase and the diffusion operator. Interested readers can learn more about the circuit construction using the interactive applet created by Prokopenya (2011).

Source: Adapted from Prokopenya (2011)

¹⁸ It holds that $\mathbf{X}|-\rangle = \frac{1}{\sqrt{2}}(\mathbf{X}|0\rangle - \mathbf{X}|1\rangle) = \frac{1}{\sqrt{2}}(|1\rangle - |0\rangle) = -\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) = -|-\rangle$. It may seem that the minus sign before the state can be ignored, as it is a global quantum phase. However, state $|-\rangle$ is incorporated into multi-qubit state $|\psi_{\text{init}}\rangle$ and the sign is changed only for state $|011\rangle$, hence -1 is in fact the relative phase.

Next, a so-called *diffusion operator* implementing the “mirroring” is applied. The logic behind the construction of the diffusion operator is beyond the scope of this article. However, details are provided by Grover (2001).¹⁹ The oracle and the diffusion operator of the circuit are applied R times. Finally, the working qubits are measured to get the desired record from the database.

Note that multi-controlled \mathbf{X} and \mathbf{Z} gates appear in the circuit. Techniques for decomposing such gates into more elementary ones are demonstrated in Barenco et al. (1995). However, in practice we can use the implementation of the Grover algorithm available in the Qiskit libraries.

Having introduced the general Grover algorithm, we will now outline its adaptation for solving QUBO problems. Assume we are solving a problem with n variables and we have m qubits for storing the respective objective function values. Those $n + m$ working qubits are initialized to the quantum state

$$|\psi_{\text{init}}\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle_n |f(i) - y\rangle_m, \quad (12)$$

where $|i\rangle$ is the i^{th} basis state of an n -qubit system, $f(i)$ is the objective function value for the bit string encoded in state $|i\rangle$ and y is the lowest value of the objective function found so far. Unless such y exists, it is initialized to $y := f(s)$, where s is a random n bits long string.

In fact, state (12) is a uniformly distributed superposition consisting of all the possible arguments of the objective function and the respective function values. Note that qubits $|f(i) - y\rangle_m$ are entangled with qubits $|i\rangle_n$ to preserve the relation between function arguments i and function values $f(i)$. It is worth noting that the parameters of the quantum gates in the circuit preparing state $|\psi_{\text{init}}\rangle$ are based on the coefficients of the objective function rather than on the actual function values. Otherwise, we would need to list all the function values ex-ante and the algorithm would be reduced to the brute force method without any speed-up.

If, for a certain argument k , it holds that $f(k) < y$, then state $|k\rangle_n$ should be marked, because we have found a better solution than y . This means that the oracle is quantum gate \mathbf{O} defined as

$$\mathbf{O}|i\rangle_n |f(i) - y\rangle_m = \text{sgn}[f(i) - y] |i\rangle_n |f(i) - y\rangle_m \quad \forall i \in \{0; 1 \dots 2^n - 1\}. \quad (13)$$

After state $|k\rangle_n$ has been marked, its probability amplitude is amplified with the diffusion operator (the same one as shown in Figure 4) and the state stored in the first n working qubits is measured to get k . The function value $f(k)$ is determined and a new quantum state $|\psi_{\text{init}}\rangle$ is prepared, but with $y := f(k)$. The algorithms run until the value of the objective function decreases. As can be seen, this adaptation of the Grover algorithm follows the hybrid approach, as the quantum computer has to be reprogrammed in each iteration. Technical details on the construction of the circuit preparing state $|\psi_{\text{init}}\rangle$ and the design of the oracle \mathbf{O} are presented in Gilliam et al. (2021).

¹⁹ The article contains ideas leading to the construction of the Grover algorithm. Interestingly, the algorithm is a by-product of Schrödinger equation discretization.

4. Optimization of the FX Reserves Currency Composition

In this part, we will discuss the actual application of quantum computers in portfolio optimization. In particular, we will show the capabilities of the approaches presented in Section 3 as regards finding the optimal currency composition of the CNB's FX reserves. First, we will provide a detailed description of the optimization problem. After that, we will show how the underlying data for the optimization are derived, and finally we will discuss the results from the technical and briefly from the economic perspectives.

4.1 Formulation of the Optimization Problem

The goal of the optimization problem is to find the optimal currency composition of the CNB's FX reserves after three major crises of the 21st century: the Great Recession of 2007–2009, the European debt crisis of 2011 and the Covid crisis of 2020. Since three periods are involved, the optimization is dynamic. The optimal compositions should generate the maximum possible return in CZK and at the same time the return volatility and the transaction costs incurred in connection with reserve currency rebalancing should be minimized. We require the sum of the weights of all the currencies involved to equal 1 for each time period. The weights have to be non-negative (i.e. short positions are forbidden) and less than or equal to one. Apart from these requirements, we impose no further constraint on the currency composition. We include all the currencies we currently have in the reserves in the optimization regardless of they were part of the reserves in the periods mentioned above. This means that the following currencies are included: AUD, CAD, CNY, EUR, GBP, JPY, SEK, USD and gold.

Note that for the purposes of this paper, we abstract from the particular asset classes included in the reserves and optimize the currency composition only from the perspective of the FX returns in CZK. We could include individual asset classes (such as government bonds, agency bonds and equities) for each currency in the model, but at the cost of additional variables. This would make the model too complex for solving on a quantum computer. Another possible approach would be two-level optimization, i.e. first find the optimal currency composition and then find the asset composition for each currency. However, we leave this exercise for further research.

As discussed above, we carry out the optimization for three time periods. To anchor the periods, we looked for troughs (i.e. local minima) of the MSCI World Developed Markets index. The index comprises more than 1,000 shares listed at exchanges in developed countries.²⁰ Thanks to its world-wide coverage, the index is suitable for identifying global crises. The path of the index from 2000 to 2021 is shown in Figure 5. A detailed analysis revealed that the lowest points for the Great Recession, the European debt crisis and the Covid crisis were reached on 9 March 2009, 22 September 2011 and 18 March 2020 respectively. Note that we did not include the year 2022 in our analysis, because the crisis connected with the Russian–Ukrainian conflict is still ongoing.

The above-described version of the optimization problem is the *practical* one. Because of the different capabilities of the individual quantum algorithms and computers we used, we also formulated two simpler versions of the practical problem. The first is a *testing* version comprising only the time period connected with the Great Recession and all nine currencies. The second is a *toy* version comprising the time period connected with the European debt crisis and only three currencies (AUD, CAD and gold). To reduce the number of weights in the toy

²⁰ Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Hong Kong, Ireland, Israel, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, the United Kingdom and the USA.

version further, we removed the constraint $\sum_i w_i = 1$ from the objective function and substituted $w_{\text{gold}} = 1 - w_{\text{AUD}} - w_{\text{CAD}}$ into the function.²¹

Figure 5: MSCI World Developed Markets Index 2000–2021



Note: The troughs of the Great Recession, the European debt crisis and the Covid crisis are highlighted with red circles. Note that we do not include the dot-com crisis of 2002 (the blue circle) in our analysis.

Source: Author’s own creation based on Bloomberg data

Since the optimization is carried out on a quantum computer, the objective function is formulated in QUBO form (5). This means that we binarized the asset weights. We used 14, 10 and 3 binary variables per weight for the practical, testing and toy versions respectively. The accuracy of the weights is therefore 0.006%, 0.1% and 12.5% respectively. The total number of binary variables involved is 378, 90 and 6 respectively. The numbers of binary variables used for binarizing the weights come from the maximum number of variables the given algorithm can work with (the toy and testing versions) and from the requirement to ensure sufficient accuracy (the practical version). We set risk aversion $\lambda = 10$, the sensitivity to transaction costs $\mu = 20$ and the importance of the constraint that the sum of the weights is equal to one $F = 100$ (for D-Wave constrained hybrid heuristic $F = 0$). These numbers are chosen arbitrarily.

4.2 Data and Methodology

The underlying raw data we used are the FX rates between CZK and all the other currencies involved. The rates were expressed in the notation $ccy\text{CZK}$, i.e. the number of CZK units that one unit of the currency ccy is worth. The data come from the CNB’s FX rates list and Bloomberg. Based on the bid and ask rates we calculated the mid rate.

The daily CZK return on the i^{th} currency is given by the expression $100(FX_t^{(i)}/FX_{t-1}^{(i)} - 1)$, where $FX_t^{(i)}$ is the mid rate for day t . The daily returns were averaged over a time period starting on the day the trough of the crisis was reached and spanning five years to get an estimation of the expected

²¹ The objective function for the three assets is provided by Markowitz (1952), page 83.

returns. Note that for the Covid crisis the series ends on 31 December and is therefore only 1.75 years long. The averaged returns were annualized. Based on the daily returns, we calculated the variances in the returns and the covariances among the returns of all currencies. The variances and covariances were also annualized. Transaction costs are based on bid/ask spreads. On each day, for the i^{th} currency we calculated the relative spread $100(FX_{\text{ASK}}^{(i)} - FX_{\text{BID}}^{(i)})/FX_{\text{MID}}^{(i)}$. The relative spreads were averaged over the same time spans as the returns. The expected returns, average transaction costs and covariance matrices for all three time periods are provided in Appendix C.

Based on the expected returns, the average transaction costs and the covariance matrices, we carried out the optimization using a continuous gradient solver implemented in MS Excel. The results for the practical, testing and toy versions of the optimization are provided in Table 1. The run time of the gradient method was 2,109, 375 and 94 milliseconds for the practical, testing and toy model respectively.²² The optimal value of the objective function was 0.0921, 0.0177 and 0.088 for the practical, testing and toy model respectively. We find that in all instances, the eigenvalues of the matrix defining the objective function are positive, i.e. the matrix is positive definite. This means that we are dealing with the “easy version” of the quadratic optimization problem, which makes it harder for quantum computers to outperform classical ones. Based on these results, we will assess the capabilities of quantum algorithms in practical applications.

Table 1: Results of Optimization with MS Excel Gradient Solver

Version	Time period	USD	EUR	AUD	CAD	GBP	SEK	JPY	CNY	Gold
Practical	Great Recession	0.0%	31.7%	19.8%	1.2%	5.9%	35.6%	0.0%	0.0%	5.7%
	Debt Crisis	6.2%	42.6%	15.2%	0.0%	6.4%	23.5%	0.0%	4.8%	1.4%
	Covid	0.0%	20.1%	32.7%	0.6%	12.7%	21.2%	0.0%	3.7%	8.9%
Testing	Great Recession	0.0%	26.2%	17.1%	0.0%	0.0%	52.9%	0.0%	0.0%	3.8%
Toy	Debt Crisis			31.1%	54.7%					14.2%
Toy – bin	Debt Crisis			37.5%	50.0%					12.5%

Note: In the toy version, only AUD, CAD and gold are included. Since only three binary variables are used for the representation of each weight, the optimal solution that can be found in this setting is different from the one obtained using the gradient method. This solution is depicted in row *Toy – bin*.

Source: Author’s own calculations

It is important to note that these benchmark values concern portfolio optimization with real asset weights, whereas the other algorithms tested are intended for the QUBO version of the problem. As the purpose of this article is to assess the practical capabilities of quantum algorithms, we compared them with the option most available to bank analysts, the MS Excel solver, despite the different nature (i.e. continuous vs binary) of the algorithms.

4.3 Results

In this part, we will discuss the results we obtained. First, we will present the results for all three versions of the problem from the perspective of the ability of the algorithms employed to find the optimal solution and the run time. Second, we will briefly discuss the economic interpretation of the optimal results.

²² Note that the tests of all the classical algorithms were carried out on a PC equipped with 8 GB RAM and a four-core processor with clock speed 3.3 GHz.

Note that in the case of the toy version, we ran the optimization only once, as we immediately obtained the optimal result. In the case of the testing and practical versions, we repeated the optimization 100 times for the D-Wave hybrid constrained heuristic and quantum annealer (in particular, we used the *D-Wave Advantage System 4.1* quantum processor) and 10 times for the classical heuristics, as the run times were much longer in the latter case. After that, we picked the solution nearest to the optimal one in terms of the objective function value. For the branch and bound algorithm, we ran the calculation only once, as the algorithm is exact. Note that for the testing and practical versions, we show the currencies' weights in condensed graphical form; the numerical results are presented in Appendix D.

Toy Version

The toy version was devised to demonstrate portfolio optimization on IBM QuantumTM, due to the limited number of qubits on the available quantum processors. In particular, we used the Nairobi and Oslo processors, both of which have seven qubits. Such a low number of qubits allowed us to run the portfolio optimization with at most three currencies and to use only three binary variables to represent each currency weight. This would imply nine binary variables, but one weight was expressed with the other two, hence we actually needed six qubits to run the VQE and the QAOA.

Both the VQE and the QAOA were run on the Nairobi and Oslo real quantum processors with success. As both algorithms are iterative, we compared the number of iterations necessary for real processors and the noiseless simulator. The comparison is depicted in Table 2.

Table 2: Results of Optimization with QAOA and VQE for Toy Version

Processor	VQE	QAOA
Simulator	26	119
Nairobi	61	118
Oslo	61	154

Note: The table presents the number of iterations needed to find the optimal solution of the toy version. The number of repetitions (shots in IBM terminology) in each iteration was set automatically to 4,000.

Source: Author's own calculations

As can be seen from the table, the number of iterations is higher in the case of real processors. The only exception is the QAOA run on the Nairobi processor, where the number of iterations is lower. We found that the VQE necessitated less iterations than the QAOA. We attribute this better performance to the fact that the VQE is a more general algorithm than the QAOA (the VQE assumes no specific structure of the Hamiltonian describing the optimization problem) and thus the quantum circuit is less complicated. The run time of one iteration of both the VQE and the QAOA is in the order of seconds. Moreover, the batch-processing approach used in the IBM QuantumTM environment increases the total run time to hours. Together with the low number of qubits, this fact disqualifies optimization algorithms designed for gate-based quantum computers from real-world application for the time being.

The situation was even more complicated in the case of Grover adaptive search. On top of the six qubits used for the binary variables, we also needed one ancilla qubit and qubits for storing the objective function values. After experimenting on a simulator, we found that another six qubits are

necessary for this purpose. In total, we would need 13 qubits to run Grover adaptive search. Hence, we tested the algorithm on the simulator only. Note that we achieved the desired results.

We also solved the toy version with D-Wave heuristics and classical methods. The run times are compared in Table 3. The table also contains the classification of the algorithm (continuous/QUBO, exact/heuristic and classical/quantum/hybrid) to show the differences between the approaches. The best algorithm is classical simulated annealing, followed by the D-Wave annealer. Interestingly, the slowest approach is the D-Wave hybrid constrained heuristic. However, we should not draw strong conclusions from the results based on the toy version with only six binary variables. In such cases, the majority of the time is consumed by setting up the problem, with the actual solution taking up only a fraction of the run time. In examples with more variables, the relation is rather the opposite.

Table 3: D-Wave and Classical Method Run Times for Toy Version

Method	Type	Run time (ms)
Simulated annealing	Classical heuristic QUBO	8.98
D-Wave annealer ^a	Quantum heuristic QUBO	26.01
Brute force method	Classical exact QUBO	69.93
Genetic optimization	Classical heuristic QUBO	80.56
Gradient method ^b	Classical exact continuous	94.00
Branch and bound	Classical exact QUBO	375.00
D-Wave hybrid	Hybrid heuristic QUBO	2,993.00 ^c

Note:

^a The number of calculation repetitions (the *number of reads* in D-Wave terminology) was 1,000.

^b The gradient method serves as the benchmark.

^c The run time on the quantum annealer was 63.17 ms; the rest of the time was consumed by classical heuristics.

Source: Author's own calculations

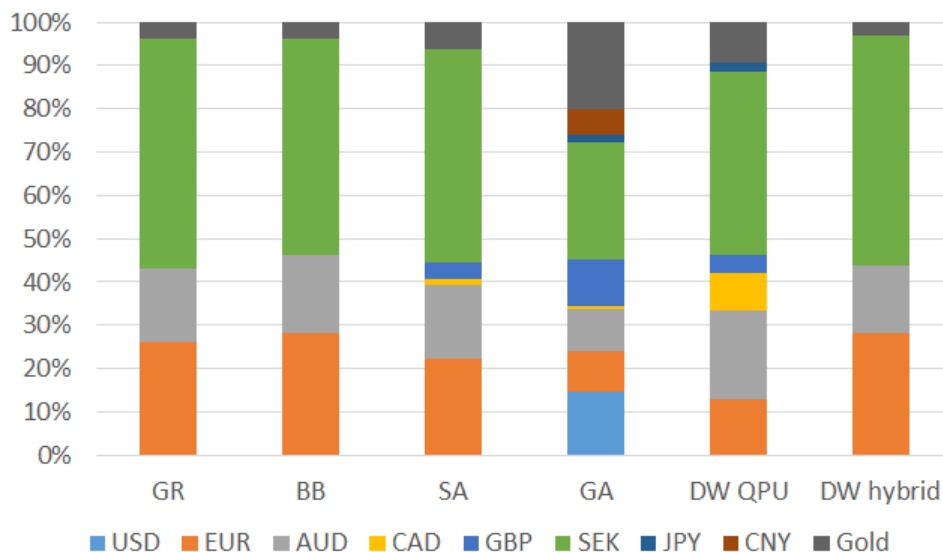
Testing Version

The testing version comprises 90 binary variables and thus cannot be optimized on IBM QuantumTM. The results obtained with the D-Wave and classical algorithms are compared in Figure 6.

As can be seen in Figure 6, no algorithm was able to find the global optimum. The nearest sub-optimal solution was provided by the D-Wave hybrid heuristic (objective function value 0.01791), followed by the branch and bound algorithm (0.01810), simulated annealing (0.01905), the D-Wave annealer (0.02516) and genetic optimization (0.05578).

The D-Wave hybrid approach and simulated annealing outperformed the purely quantum D-Wave annealer. We attribute this result mainly to the limited connectivity of the D-Wave quantum processor. The problem has to be adapted to the processor topology, which leads to an increase in the number of ancilla qubits serving as intermediaries in the links between the working qubits. The intermediaries add additional noise and therefore increase the probability of quantum state decoherence. In contrast, the hybrid approach uses the quantum processor for solving less complex sub-problems, which eliminates the need for a large number of ancillas and hence reduces the probability of decoherence. This means that the hybrid approach leverages the capabilities of the quantum annealer as much as possible and at the same time the combination with the classical approach reduces the imperfections of the current quantum hardware.

Figure 6: Best Solutions for Testing Version



Note: *GR* stands for classical gradient algorithm, *BB* for branch and bound algorithm, *SA* for simulated annealing, *GA* for genetic-based algorithm, *DW QPU* for D-Wave quantum annealer and *DW hybrid* for D-Wave constrained hybrid heuristic. In the case of the D-Wave annealer, the number of repetitions (*reads* in D-Wave terminology) was set to 5,000.

Source: Author's own calculations

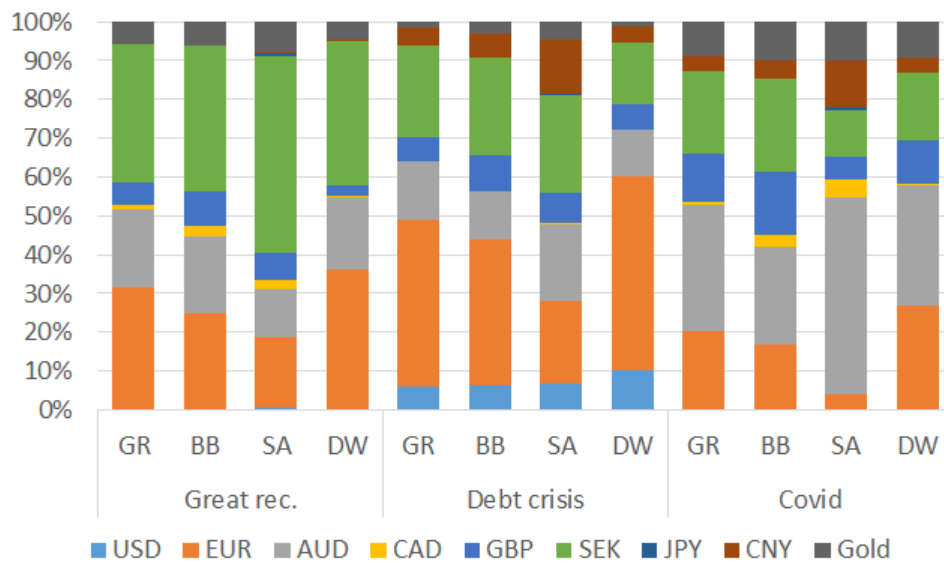
Despite the fact that the branch and bound algorithm is an exact solver, it was not able to find the optimum. However, the reason is purely technical. The CPLEX™ environment ended with a message that the optimization could not continue and we found that the RAM was full. Increasing the RAM size would allow us to find the global optimum of the problem.

Note that classical genetic optimization was far from the optimal solution. This is probably caused by the fact that, in contrast to simulated annealing, solutions with a higher value of the objective function are always rejected. As a result, the algorithm can get stuck in a local minimum.

The run times of the algorithm were the following: 1,069 ms for the D-Wave annealer, 5,011 ms for the D-Wave hybrid heuristic (of that, 15.232 ms was spent on the quantum annealer), 485 seconds for simulated annealing, 1,382 seconds for genetic optimization and 5,880 seconds for the branch and bound algorithm. This clearly demonstrates the superiority of the D-Wave hybrid approach in the realm of heuristics. In the case of the branch and bound algorithm, despite the very long run time, we would be able to find the global optimum if we had enough memory.

Practical Version

As the practical version contains 378 binary variables, the D-Wave software was not able to embed the problem in the quantum processor in reasonable time. We also realized that the genetic algorithm is not able to find any feasible solution (i.e. having a sum of weights equal to 1). Therefore, we only employed simulated annealing, the branch and bound algorithm and the D-Wave hybrid constrained heuristic for the practical version. The solutions obtained are shown in Figure 7.

Figure 7: Best Solutions for Practical Version

Note: *GR* stands for classical gradient algorithm, *BB* for branch and bound algorithm, *SA* for simulated annealing and *DW* for D-Wave constrained hybrid heuristic.

Source: Author's own calculations

Similarly to the testing version, no algorithm was able to find the optimal solution. The branch and bound algorithm, the D-Wave hybrid heuristic and simulated annealing found sub-optimal solutions with objective function values of 0.09325, 0.09395 and 0.10710 respectively in 39,429, 5.048 and 6,333 seconds respectively. These results make the D-Wave hybrid heuristic the winner in terms of run time (only 5 seconds). Concerning the closeness to the optimal solution, the branch and bound algorithm achieved the best result, despite its very long run time (almost 11 hours).

The reason why none of the QUBO heuristics tested was able to find the global optimum lies in their nature – they search in a small subset of all the possible solutions. On the one hand, this allows us to find at least some solution, but on the other hand, we cannot be sure if the solution is optimal. Moreover, binarization of the real weights leads to an objective function with considerably different coefficients, as terms containing 2^{-k} have significantly different values for different k . In the case of quantum methods, terms with low coefficients can be obscured by noise. In the case of classical heuristics, the algorithm can get stuck in a local minimum despite the measures taken, as in the case of simulated annealing.

Concerning the exact solvers, the branch and bound algorithm again faced the issue of a lack of memory. If we had large enough RAM, we would be able to find the exact solution, but at the cost of a long run time.

Economic Interpretation of the Optimal Results

The actual currency composition of the CNB FX reserves differs significantly from the optimal one that we found for all three time periods. The reserves are composed of all nine mentioned currencies, but the presented optimal solutions contain only some of them.

With the exception of the “Debt crisis” time period, the weight of USD is zero. This is caused by its relatively high volatility and low returns in comparison with correlated currencies such as AUD and CAD. However, the American dollar is the primary reserve currency and the prominent currency in international trade. Moreover, the markets for AUD and CAD government bonds are much smaller and less liquid than the US one. Therefore, USD has its place in the FX reserves despite the mentioned disadvantages.

Because of the lower volatility and higher returns of CNY and the high correlation between CNY and JPY, the optimization algorithms prefer the Chinese yuan to the Japanese yen (the weight of the yen is zero). However, until recently, investments in CNY were constrained by the Chinese authorities. For this reason, JPY has been more attractive as a reserve currency.

The Swedish krona and the euro are the currencies with the lowest correlation with the others, so the optimization algorithms increase their weight significantly in order to diversify the reserves. A high share of EUR makes sense, as the Czech Republic is part of the EU and the EUR market is deep and highly liquid. In contrast, the Swedish market is the smallest of those included in our analysis. The share of SEK should be therefore reduced and the share of EUR even increased.

These three examples clearly illustrate the need to provide optimization algorithms with additional information, for example maximum (or conversely minimum) shares for currencies that suffer from some deficiencies (or conversely are too important to be omitted). Such information has to be based on expert judgment.²³ However, we leave more practical optimization for further research once quantum computers are mature enough to easily incorporate constraints into optimization problem.

Another issue is connected with the estimation of the covariance matrices, expected returns and average transaction costs. To get a meaningful estimation, relatively long time series are needed. For example, we tried to estimate the covariance matrix for the Great Recession period with a one-year-long series. However, the optimization ended with 100% allocation to the Australian dollar. Clearly, such portfolio is poorly diversified, and we had to prolong the time series to five years. This led to another problem, as the five-year time series covers not only the recovery after the Great Recession, but also part of the European debt crisis. On top of that, we have to bear in mind that our analysis is retrospective. In practice, we would need to carry out a forward-looking estimation of the covariances and returns, because we would be interested in the future currency composition of the reserves. In the end, we put these deficiencies of our model aside, because the main purpose of this article is to examine the capabilities of quantum optimization algorithms. We leave a discussion of the estimations of inputs to portfolio optimization to further research.

5. Conclusion

The main aim of this paper was to assess the capabilities of quantum algorithms for portfolio optimization and compare them with their classical counterparts. In particular, we tested algorithms for universal gate-based quantum computers (the VQE, the QAOA and Grover adaptive search) on IBM QuantumTM, the quantum annealers and hybrid heuristics provided by D-Wave, the classical branch and bound algorithm (implemented in IBM CPLEXTM) and two classical heuristics

²³ As Markowitz (1952) puts it: *To use the E-V [expected return-variance] rule in the selection of securities we must have procedures for finding reasonable μ and σ . These procedures, I believe, should combine statistical techniques and the judgment of practical men. My feeling is that the statistical computations should be used to arrive at a tentative set of μ and σ . Judgment should then be used in increasing or decreasing some of these μ and σ on the basis of factors or nuances not taken into account by the formal computations.*

(simulated annealing and a genetic-based algorithm). We tested the algorithms from the perspective of run time and ability to find the optimal solution. To carry out the analysis, we constructed a binary version of Markowitz dynamic portfolio optimization in order to find the optimal currency composition of the CNB's FX reserves. The continuous gradient method implemented in MS Excel served as a benchmark for the purposes of comparison, as the original version of Markowitz optimization is continuous.

We found that none of the QUBO algorithms tested was able to find the optimal solution discovered by the classical continuous gradient method implemented in MS Excel, although the D-Wave hybrid constrained heuristic was very close in the case of the testing version of the problem, as was the branch and bound algorithm in the case of the practical version. However, this all means that the QUBO approach seems not to be a suitable option for portfolio optimization. Even with a usual office PC, we are able to find the global optimum of the problem in seconds using the classical continuous gradient method implemented in MS Excel. For the purposes of central banks, this is plainly sufficient. The QUBO version of portfolio optimization is useful only when the problem is purely binary, i.e. when we only want to decide whether or not to incorporate certain assets into the portfolio and we do not need to know the weights of those assets.

However, this conclusion does not mean that quantum QUBO algorithms are useless. As mentioned earlier, they can be used to solve problems which are binary by definition (such as the travelling salesperson problem or the assignment problem), but we leave this to further research. Moreover, the inability of the algorithms to find the solution to the problem could mean that the problem is hard to solve (we can see this from the excessive run time of the exact branch and bound algorithm). Therefore, it can be used as an interesting testing problem for new algorithms or for measuring progress in the development of quantum computers.

Concerning the capabilities of quantum annealers, we found that the current noisy quantum computers are not able to solve problems of practical size themselves, but they can be used to enhance classical algorithms. The D-Wave hybrid heuristic seems to be just one step from real-world application. Although this algorithm was not always the fastest, it was able to find the solution closest to the optimal one in reasonable time, especially when compared with classical algorithms, both heuristic and exact.

Regarding the QAOA, the VQE and Grover adaptive search, we will have to wait until the number of qubits offered by universal gate-based quantum computers is high enough before we can assess their ability to solve problems of practical size. Grover adaptive search has an even longer road ahead, as the consumption of qubits is higher than in case of the QAOA and the VQE. Moreover, as the Grover algorithm has guaranteed speed-up, we suspect that qubit quality will also be an important factor in its performance.

For the time being, this article should serve as a complement to the introductory literature on quantum computing for economists in general and central bankers and regulators in particular. We expect that it will take several years for quantum computers to become a fully mature technology. Moreover, it is possible that Lloyd's conjecture on exponential speed-up for quantum system simulation will be proved. This would mean that quantum computers will be helpful in solving QUBO problems. In the meantime, we will "keep an eye" on the development of quantum computers and certainly return to this research once they are on the verge of industry-scale deployment.

Disclaimer

This article is intended to be a purely academic research and educational material. The article is neither an offer to buy or sell any security nor a solicitation to do so. The article is not investment advice. Neither the author nor the CNB is responsible for any losses (including but not limited to financial and reputational ones) incurred in connection with the implementation of the methods described in the article. The author declares that no result presented in the article has been implemented as a part of the CNB's investment strategy or used as an input to its strategic asset allocation decision-making. The use of the software products discussed in the article is non-profit-oriented (research and educational purposes).

References

- AARONSON, S. (2015): “Quantum Machine Learning Algorithms: Read the Fine Print.” *Nature Phys*, 11(4):291–293.
- AHARONOV, D., W. VAN DAM, J. KEMPE, Z. LANDAU, S. LLOYD, AND O. REGEV (2007): “Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation.” *SIAM Journal of Computing*, 37(1):166–194.
- AHARONOV, D., X. GAO, Z. LANDAU, Y. LIU, AND U. VAZIRANI (2022): “A Polynomial-time Classical Algorithm for Noisy Random Circuit Sampling.” *arXiv:2211.03999v1 [quant-ph]*.
- ANSCHUETZ, E. R., J. OLSON, A. ASPURU-GUZI, AND Y. CAO (2019): *Variational Quantum Factoring*. In *Quantum Technology and Optimization Problems*, pages 74–85. Springer International Publishing.
- AYUB, C., M. CEBERIO, AND V. KREINOVICH (2020): *How Quantum Computing Can Help with (Continuous) Optimization*. In Ceberio, M. and V. Kreinovich, editors, *Decision Making under Constraints*, pages 7–14. Springer International Publishing.
- BARENCO, A., C. H. BENNETT, R. CLEVE, D. P. DIVINCENZO, N. MARGOLUS, P. SHOR, T. SLEATOR, J. A. SMOLIN, AND H. WEINFURTER (1995): “Elementary Gates for Quantum Computation.” *Physical Review A*, 52(5):3457.
- BARKOUTSOS, P. K., G. NANNICINI, A. ROBERT, I. TAVERNELLI, AND S. WOERNER (2020): “Improving Variational Quantum Optimization using CVaR.” *Quantum*, 4:256.
- BERTSIMAS, D. AND J. TSITSIKLIS (1993): “Simulated Annealing.” *Statistical Science*, 8(1): 10–15.
- CHEN, S., J. COTLER, H.-Y. HUANG, AND J. LI (2022): “The Complexity of NISQ.” *arXiv:2210.07234v1 [quant-ph]*.
- CHILDS, A. M. (2009): “Equation Solving by Simulation.” *Nature Physics*, 5(12):861.
- CORANA, A., M. MARCHESI, C. MARTINI, AND S. RIDELLA (1987): “Minimizing Multimodal Functions of Continuous Variables with the Simulated Annealing Algorithm.” *ACM Transactions on Mathematical Software*, 13(3):262–280.
- D-WAVE (2022): “Hybrid Solvers for Quadratic Optimization.” (<https://www.dwavesys.com/media/soxph512/hybrid-solvers-for-quadratic-optimization.pdf>)
- D-WAVE (2018): “Getting Started with the D-Wave System.” (https://docs.dwavesys.com/docs/latest/doc_getting_started.html)
- D-WAVE (2022): “D-Wave Ocean Software Documentation.” (<https://docs.ocean.dwavesys.com/en/stable/>)
- DE WOLF, R. (2022): “Quantum Computing: Lecture Notes.” *arXiv:1907.09415v3 [quant-ph]*.
- DING, Y., J. GONZALEZ-CONDE, L. LAMATA, J. D. MARTÍN-GUERRERO, E. LIZASO, S. MUGEL, X. CHEN, R. ORÚS, E. SOLANO, AND M. SANZ (2023): “Towards Prediction of Financial Crashes with a D-Wave Quantum Computer.” *Entropy*, 25(2):323.
- DURR, C. AND P. HOYER (1996): “A Quantum Algorithm for Finding the Minimum.” *arXiv:quant-ph/9607014*.

- EGGER, D. J., C. GAMBELLA, J. MAREČEK, S. MCFADDIN, M. MEVISSSEN, R. RAYMOND, A. SIMONETTO, S. WOERNER, AND E. YNDURAIN (2020): “Quantum Computing for Finance: State of the Art and Future Prospects.” *IEEE Transactions on Quantum Engineering*, 1:1–24.
- EGGER, D. J., J. MAREČEK, AND S. WOERNER (2021): “Warm-starting Quantum Optimization.” *Quantum*, 5:479.
- ELSOKKARY, N., F. S. KHAN, D. L. TORRE, T. S. HUMBLE, AND J. GOTTLIEB (2017): “Financial Portfolio Management using D-Wave Quantum Optimizer: The Case of Abu Dhabi Securities Exchange.” *IEEE High-performance Extreme Computing 2017 Conference Proceedings*.
- FARHI, E., J. GOLDSTONE, AND S. GUTMANN (2014): “A Quantum Approximate Optimization Algorithm.” *arXiv:1411.4028v1 [quant-ph]*.
- FERNÁNDEZ-VILLAYERDE, J. AND I. HULL (2022): “Dynamic Programming on a Quantum Annealer: Solving the RBC Model.” (https://www.sas.upenn.edu/~jesusfv/Solving_Quantum.pdf)
- FEYNMAN, R. P. (1985): “Quantum Mechanical Computers.” *Optics News*, 11(2):11–20.
- FICO (2022): “Xpress Optimization Help.” (<https://www.fico.com/fico-xpress-optimization/docs/latest/solver/optimizer/HTML/GUID-3BEAAE64-B07F-302C-B880-A11C2C4AF4F6.html>)
- FONTANELA, F., A. JACQUIER, AND M. OUMGARI (2021): “Short Communication: A Quantum Algorithm for Linear PDEs Arising in Finance.” *SIAM Journal on Financial Mathematics*, 12(4):98–114.
- GILLIAM, A., S. WOERNER, AND C. GONCIULEA (2021): “Grover Adaptive Search for Constrained Polynomial Binary Optimization.” *Quantum*, 5:428.
- GROVER, L. K. (2001): “From Schrödinger’s Equation to the Quantum Search Algorithm.” *Pramana: Journal of Physics*, 56(2):333–348.
- GROVER, L. K. (1996): “A Fast Quantum Mechanical Algorithm for Database Search.” *Proceedings, 28th Annual ACM Symposium on the Theory of Computing (STOC)*, 212–219.
- HARROW, A. W., A. HASSIDIM, AND S. LLOYD (2009): “Quantum Algorithm for Linear Systems of Equations.” *Physical Review Letters*, 103(15).
- HARWOOD, S., C. GAMBELLA, D. TRENÉV, A. SIMONETTO, D. BERNAL, AND D. GREENBERG (2021): “Formulating and Solving Routing Problems on Quantum Computers.” *IEEE Transactions on Quantum Engineering*, 2:1–17.
- HULL, I., O. SATTATH, E. DIAMANTI, AND G. WENDIN (2020): “Quantum Technology for Economists.” *Sveriges Riksbank Working Paper Series*, 398.
- IBM (2022): “User’s Manual for CPLEX.” (<https://www.ibm.com/docs/en/icos/12.8.0.0?topic=cplex-users-manual>)
- KATOCH, S., S. S. CHAUHAN, AND V. KUMAR (2021): “A Review on Genetic Algorithm: Past, Present, and Future.” *Springer: Multimedia Tools and Applications*, 50(5):8091–8126.
- KJAERGAARD, M., M. E. SCHWARTZ, J. BRAUMULLER, P. KRANTZ, J. I. WANG, S. GUSTAVSSON, AND W. D. OLIVER (2020): “Superconducting Qubits: Current State of Play.” *Annual Review of Condensed Matter Physics*, 11:369–395.

- KOCHENBERGER, G., J.-K. HAO, F. GLOVER, M. LEWIS, Z. LÜ, H. WANG, AND Y. WANG (2014): “The Unconstrained Binary Quadratic Programming Problem: A Survey.” *Journal of Combinatorial Optimization*, 28(1):58–81.
- KOZLOV, M. K., S. P. TARASOV, AND L. G. KHACHIYAN (1980): “The Polynomial Solvability of Convex Quadratic Programming.” *USSR Computational Mathematics and Mathematical Physics*, 20(5):223–228.
- KRISLOCK, N., J. MALICK, AND F. ROUPIN (2017): “BiqCrunch: A Semidefinite Branch-and-bound Method for Solving Binary Quadratic Problems.” *ACM Transactions on Mathematical Software*, 43(4):1–23.
- LAND, A. H. AND A. G. DOIG (1960): “An Automatic Method of Solving Discrete Programming Problems.” *Econometrica*, 28(3):497–520.
- LEMARÉCHAL, C. (2012): “Cauchy and the Gradient Method.” *Documenta Mathematica*, Extra Vol. Optimization Stories:251–254.
- LITTLE, J., K. G. MURTY, D. W. SWEENEY, AND C. KAREL (1963): “An Algorithm for the Traveling Salesman Problem.” *Operations Research*, 11(6):972–989.
- LLOYD, S. (1996): “Universal Quantum Simulators.” *Science*, 273(5278):1073–1078.
- LOVELACE, A. A. (1843): “Notes to Sketch of the Analytical Engine Invented by Charles Babbage, by L. F. Menabrea.” *Scientific Memoirs, London*, 3:433–460.
- LUCAS, A. (2014): “Ising Formulations of Many NP Problems.” *Frontiers in Physics*, 2.
- MANDRA, S., H. G. KATZGRABER, AND C. THOMAS (2017): “The Pitfalls of Planar Spin-Glass Benchmarks: Raising the Bar for Quantum Annealers (Again).” *Quantum Science and Technology*, 2(3).
- MARKOWITZ, H. (1952): “Portfolio Selection.” *The Journal of Finance*, 7(1):77–91.
- METROPOLIS, N., A. W. ROSENBLUTH, M. N. ROSENBLUTH, A. H. TELLER, AND E. TELLER (1953): “Equation of State Calculation by Fast Computing Machines.” *The Journal of Chemical Physics*, 21(6):1087–1092.
- MIRJALILI, S. (2018): “The Genetic Algorithm (GA) : Selection + Crossover + Mutation + Elitism.” (<https://www.mathworks.com/matlabcentral/fileexchange/67435-the-genetic-algorithm-ga-selection-crossover-mutation-elitism>)
- MUGEL, S., C. KUCHKOVSKY, E. SÁNCHEZ, S. FERNÁNDEZ-LORENZO, J. LUIS-HITA, E. LIZASO, AND R. ORÚS (2022): “Dynamic Portfolio Optimization with Real Datasets using Quantum Processors and Quantum-Inspired Tensor Networks.” *Physical Review Research*, 4(1):013006.
- NIELSEN, M. A. AND I. L. CHUANG (2010): *Quantum Computation and Quantum Information*. Cambridge University Press, tenth edition. ISBN: 978-1-107-00217-3
- ORÚS, R., S. MUGEL, AND E. LIZASO (2019): “Quantum Computing for Finance: Overview and Prospects.” *Reviews in Physics*, 4.
- PALMER, S., S. SAHIN, R. HERNÁNDEZ, S. MUGEL, AND R. ORÚS (2021): “Quantum Portfolio Optimization with Investment Bands and Target Volatility.” *arXiv:2106.06735v3 [q-fin.PM]*.

- PALMER, S., K. KARAGIANNIS, A. FLORENCE, A. RODRIGUEZ, R. ORÚS, H. NAIK, AND S. MUGEL (2022): “Financial Index Tracking via Quantum Computing with Cardinality Constraints.” *arXiv:2208.11380v1 [quant-ph]*.
- PERUZZO, A., J. MCCLEAN, P. SHADBOLT, M.-H. YUNG, X.-Q. ZHOU, P. J. LOVE, A. ASPURU-GUZYK, AND J. L. O’BRIEN (2014): “A Variational Eigenvalue Solver on a Photonic Quantum Processor.” *Nature Communications*, 5(5):4213.
- POTOK, T. AND P. DATE (2021): “Adiabatic Quantum Linear Regression.” *Nature: Scientific Reports*, 11(1):21905.
- PRESKILL, J. (2018): “Quantum Computing in the NISQ Era and Beyond.” *Quantum*, 2:79.
- PROKOPENYA, A. (2011): “Quantum Circuit Implementing Grover’s Search Algorithm.” (<http://demonstrations.wolfram.com/QuantumCircuitImplementingGroversSearchAlgorithm/>)
- RECHENBERG, I. (1973): *Evolutionsstrategie. Optimierung Technischer Systeme nach Prinzipien der Biologischen Evolution*. Frommann-Holzboog, first edition. ISBN: 978-3-7728-0374-1
- ROSENBERG, G., P. HAGHNEGAHDAR, P. GODDARD, P. CARR, K. WU, AND M. L. DE PRADO (2016): “Solving the Optimal Trading Trajectory Problem Using a Quantum Annealer.” *IEEE Journal of Selected Topics in Signal Processing*, 10(6):1053–1060.
- SCHARF, R. AND F. M. IZRAILEV (1990): “Dyson’s Coulomb Gas on a Circle and Intermediate Eigenvalue Statistics.” *Journal of Physics A: Mathematical and General*, 23(6):963–977.
- SCHULD, M. AND N. KILLORAN (2022): “Is Quantum Advantage the Right Goal for Quantum Machine Learning?” *PRX Quantum*, 3(3).
- SCHWEFEL, H.-P. (1977): *Numerische Optimierung von Computer-Modellen Mittels der Evolutionsstrategie*. Birkhäuser Basel, first edition. ISBN: 978-3-7643-0876-6
- SHOR, P. W. (1994): “Algorithms for Quantum Computation: Discrete Logarithms and Factoring.” *Proceedings 35th Annual Symposium on Foundations of Computer Science, IEEE*, 124–134.
- TASSEFF, B., T. ALBASH, Z. MORRELL, M. VUFFRAY, A. Y. LOKHOV, S. MISRA, AND C. COFFRIN (2022): “On the Emerging Potential of Quantum Annealing Hardware for Combinatorial Optimization.” *arXiv:2210.04291v1 [math.OC]*.
- THAKOOR, N., V. DEVARAJAN, AND J. GAO (2009): “Computation Complexity of Branch-and-bound Model Selection.” *IEEE 12th International Conference on Computer Vision*, 1895–1900.
- TURING, A. M. (1950): “Computing Machinery and Intelligence.” *Mind*, 59(236):433–460.
- VAVASIS, S. A. (1990): “Quadratic Programming is in NP.” *Information Processing Letters*, 36(2):73–77.
- VAVASIS, S. A. AND P. M. PARDALOS (1991): “Quadratic Programming with One Negative Eigenvalue is NP-hard.” *Journal of Global Optimization*, 1:15–22.
- VERDON, G., J. M. ARRAZOLA, K. BRÁDLER, AND N. KILLORAN (2019): “A Quantum Approximate Optimization Algorithm for Continuous Problems.” *arXiv:1902.00409v1 [quant-ph]*.
- VESELÝ, M. (2020): “Portfolio Optimisation Problems with Hard-to-Optimise Objective Functions.” *Evolving Practices in Public Investment Management: Proceedings of the Seventh Public Investors Conference*, 143–166.

- VESELÝ, M. (2022): “Simulated Annealing for Binary Optimization.” (<https://www.mathworks.com/matlabcentral/fileexchange/116150-simulated-annealing-for-binary-optimization>)
- VESELÝ, M. (2023): “Portfolio Optimization with D-Wave Quantum Annealer.” (<https://github.com/QuantumQu/PtfOptim>)
- VESELÝ, M. (2022): “Application of Quantum Computers in Foreign Exchange Reserves Management.” *CNB Working Paper No. 2/2022*.
- VIAMONTES, G. F., I. L. MARKOV, AND J. P. HAYES (2005): “Is Quantum Search Practical?” *IEEE Computing in Science and Engineering*, 7(3):62–70.
- WEN, J., X. KONG, S. WEI, B. WANG, T. XIN, AND G. LONG (2019): “Experimental Realization of Quantum Algorithms for a Linear System Inspired by Adiabatic Quantum Computing.” *Physical Review A*, 99(1):012320.
- WIEGELE, A. (2007): “Biq Mac Library - Binary Quadratic and Max-Cut Library.” (<https://biqmac.aau.at/biqmaclib.html>)
- YARKONI, S., E. RAPONI, T. BÄCK, AND S. SCHMITT (2022): “Quantum Annealing for Industry Applications: Introduction and Review.” *Reports on Progress in Physics*, 85(10):104001.
- YE, Y. AND E. TSE (1989): “An Extension of Karmarkar’s Projective Algorithm for Convex Quadratic Programming.” *Mathematical Programming*, 44(1):157–179.

Appendix A: Matrix Formulation of the Portfolio Optimization Problem

In this appendix, we present the matrix form of the objective function (5). We will use symbol $\mathbf{1}$ for a matrix or vector (depending on the context) in which all the elements are equal to 1. Symbol \mathbf{O} is used for a zero matrix. We assume risk aversion λ and sensitivity to transaction cost μ to be time-independent. Note that in contrast to part 2 we use Dirac notation for linear algebra objects to make the expressions easily readable.

First, we rewrite the objective function in terms of weights $w_i \in \mathbb{R}$ and after that we binarize the function with binary variables $x_i \in \{0; 1\}$ to get the QUBO form of the problem.

Condition $\sum_i w_i^{(t)} = 1 \forall t \in \{1, 2, \dots, T\}$ is expressed as $F(\sum_i w_i^{(t)} - 1)^2 = F\left[(\sum_i w_i^{(t)})^2 - 2\sum_i w_i^{(t)} + 1\right]$. In matrix form $F(\langle w^{(t)} | \mathbf{1} | w^{(t)} \rangle - 2\langle w^{(t)} | \mathbf{1} \rangle + 1)$.

Transaction costs at time t are given by the expression $\sum_i v_i^{(t)} (w_i^{(t)} - w_i^{(t-1)})^2$. Considering \mathbf{V}_t to be a diagonal matrix containing unit transaction cost (the i^{th} diagonal element is $v_i^{(t)}$), the matrix form of the transaction cost term is $\langle w^{(t)} - w^{(t-1)} | \mathbf{V}_t | w^{(t)} - w^{(t-1)} \rangle = \langle w^{(t)} | \mathbf{V}_t | w^{(t)} \rangle + \langle w^{(t-1)} | \mathbf{V}_t | w^{(t-1)} \rangle - \langle w^{(t-1)} | \mathbf{V}_t | w^{(t)} \rangle - \langle w^{(t)} | \mathbf{V}_t | w^{(t-1)} \rangle$. Note that $\mathbf{V}_{T+1} = \mathbf{O}$, because we assume no further changes in the portfolio after time T .

Together with covariance matrix \mathbf{C}_t , the quadratic part of the objective function referring to time period t becomes $\mathbf{A}_t = \lambda \mathbf{C}_t + \mu (\mathbf{V}_t + \mathbf{V}_{t+1}) + F \mathbf{1}$. Besides this, we have mixed terms $\mu \langle w^{(t-1)} | \mathbf{V}_t | w^{(t)} \rangle$ and $\mu \langle w^{(t)} | \mathbf{V}_t | w^{(t-1)} \rangle$ connecting consecutive time periods. Putting all that together, we arrive at the final matrix describing the quadratic part of function (5)

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_1 & -\mu \mathbf{V}_2 & \mathbf{O} & \mathbf{O} & \mathbf{O} & \dots & \mathbf{O} \\ -\mu \mathbf{V}_2 & \mathbf{A}_2 & -\mu \mathbf{V}_3 & \mathbf{O} & \mathbf{O} & \dots & \mathbf{O} \\ \mathbf{O} & -\mu \mathbf{V}_3 & \mathbf{A}_3 & -\mu \mathbf{V}_4 & \mathbf{O} & \dots & \mathbf{O} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \mathbf{O} & \mathbf{O} & \dots & \mathbf{O} & -\mu \mathbf{V}_{T-1} & \mathbf{A}_{T-1} & -\mu \mathbf{V}_T \\ \mathbf{O} & \mathbf{O} & \dots & \mathbf{O} & \mathbf{O} & -\mu \mathbf{V}_T & \mathbf{A}_T \end{pmatrix} \quad (\text{A1})$$

The linear part of function (5) consists of asset returns and the linear term of $(\sum_i w_i^{(t)} - 1)^2$, therefore $\langle b_t | = -\langle r_t | - 2F \langle \mathbf{1} |$, hence the whole vector $\langle b |$ is $(\langle b_1 | \langle b_2 | \dots \langle b_T |)$. The constant part $c_t = F$ is generated by the absolute term of $F(\sum_i w_i^{(t)} - 1)^2$. Since we assume T time periods, in total $c = FT$. All these steps allow us to express function (5) in the form $\langle w | \mathbf{A} | w \rangle + \langle b | w \rangle + c$.

The next step involves replacing weights w_i with binary variables x_i . It holds that $w_i^{(t)} = \sum_{k=1}^{\ell} 2^{-k} x_i^{(k,t)}$. Introducing vector $\langle s | = (2^{-1} \ 2^{-2} \ \dots \ 2^{-\ell})$, we can write $x_i^{(t)} = \langle s | x_i^{(t)} \rangle$, where $|x_i^{(t)} \rangle$ is a vector of binary variables describing weight $w_i^{(t)}$. Quadratic terms such as $w_i^{(t)} w_j^{(t)}$ can clearly be rewritten as $\langle x_i^{(t)} | s \rangle \langle s | x_j^{(t)} \rangle$. Expression $|s \rangle \langle s| := \mathbf{Q}$ is a matrix of type $\ell \times \ell$. With these new symbols, we get the final shape of the ‘‘binarized’’ objective function. In matrix \mathbf{A} , we replace each element a_{ij} with matrix $a_{ij} \mathbf{Q}$ and we substitute $b_i \langle s |$ for each element in vector $\langle b |$. Constant c is unaffected by the binarization.

Appendix B: Simulated Annealing for Binary Optimization

In this appendix, we provide details on the simulated annealing implementation we used in the practical part of this paper (the MatLab source code is provided in (Veselý, 2022a)). As noted in the general description of the algorithm in Section 3.1, the particular implementation of simulated annealing is characterized by a specific noise operator and a cooling schedule. Our noise operator consists of two steps:

1. Generate integers i and j from a uniform distribution, and swap the values of x_i and x_j .
2. Generate random numbers $p_i, p_j \sim U(0, 1)$. If $p_i > 0.5$, then $x_i := \bar{x}_i$, otherwise x_i remains unchanged. Do the same for j .

The cooling schedule is defined as follows. First, the temperature is set to $T = 1$. For concrete temperature T , the noise operator is applied several times (the maximum number of iterations is defined by the user). The temperature is decreased by a user-defined decrement once the maximum number of iterations is reached. The process repeats until $T > 0$.

Sometimes, a solution with a worse value of the objective function is allowed to avoid getting stuck in a local minimum.²⁴ A decision whether to preserve the worse solution is made based on the *Metropolis criterion* (Corana et al., 1987). First, the following value is calculated

$$q = \exp\left(\frac{f_{\text{prev}} - f_{\text{curr}}}{k_B T}\right), \quad (\text{B1})$$

where f_{prev} and f_{curr} are the previous and current values of the objective function respectively, T is the temperature and $k_B > 0$ is the ‘‘Boltzmann constant’’. Second, a random number $p \sim U(0, 1)$ is generated and if $q > p$, the worse solution is preserved. Note that in the default setting $k_B = 1$, but the user can alter the constant.²⁵ Increasing the value of k_B leads to a higher probability of preserving the worse solution. Note that allowing the worse solution to be preserved requires storing the best solution found so far.

The implementation described above was tested on two types of binary optimization problem:

1. Linear objective functions $\langle b|x \rangle \rightarrow \min$ with $b_i \sim U(-1, 1)$. Clearly, the minimum occurs at the point satisfying $x_i = 1 \ \forall b_i \leq 0$ and $x_i = 0$ otherwise. We tested linear functions with up to 50,000 variables.
2. Quadratic objective functions $\langle x|\mathbf{A}|x \rangle \rightarrow \min$, where matrices \mathbf{A} were taken from the library of testing functions for QUBO provided by Wiegele (2007). We tested functions with up to 500 variables. In particular, we used the following problems from the library: *be100.1*, *be200.8.1*, *be250.1*, *gka1f* and *gka5f*.

²⁴ Note that we use the term local minimum loosely. Since the QUBO problem is a discrete problem, any solution is a local minimum. However, in cases where the objective function value is low for a particular solution in comparison with other ones, the algorithm can become trapped in that solution. This is the actual meaning we assign to *get stuck in a local minimum*.

²⁵ The actual value of the Boltzmann constant arising in thermodynamics is $k_B = 1.380649 \times 10^{-23} \text{JK}^{-1}$. However, due to the magnitude of the objective function values involved in portfolio optimization, the natural Boltzmann constant is too small.

In Table B1, we provide execution times for all the tested instances to document the capabilities of our implementation. For all the problems tested, our simulated annealing algorithm found the solution. It is worth emphasizing that the algorithm parameters (i.e. the number of iterations, the temperature-decreasing speed and the “Boltzmann constant”) had to be fine-tuned in order to achieve the optimum for a particular objective function. There is no scientific approach to setting the parameters; it is rather a matter of experience and intuition. However, two rules of thumb can be used:

1. With an increasing number of variables, the maximum number of iterations and the “Boltzmann constant” should be increased, while the temperature should decrease more slowly. This gives the algorithm enough time to scan a higher number of bit strings in the objective function domain.
2. With an increasing value of the objective function, the “Boltzmann constant” should be increased to properly scale the term $f_{\text{prev}} - f_{\text{curr}}$ in the Metropolis criterion.

Table B1: Results of Simulated Annealing Implementation Tests

Number of variables	Optimum found	Execution time		Additional details
		SA	Brute force (est.)	
<i>Random linear objective function</i>				
50	yes	0.08 s	13 d	
100	yes	0.25 s	$\approx 10^{13}$ y	
1000	yes	1.90 s	$\approx 10^{284}$ y	
10000	no	36 s	$\approx 10^{2993}$ y	opt: -2516.8697, found: -2516.8497
	yes	15 m 38 s		
50000	no	1 m 23 s	$\approx 10^{15034}$ y	opt: -19856.4264, found: -19854.8727
	yes	3 h 7 m		
<i>Quadratic objective function</i>				
100	yes	0.18 s	$\approx 10^{13}$ y	
200	yes	2.00 s	$\approx 10^{43}$ y	
250	yes	4.20 s	$\approx 10^{58}$ y	
500	yes	13 s	$\approx 10^{134}$ y	sparse matrix (10% non-zero elements) full matrix
	yes	24 s		

Note: The tests were carried out on a PC equipped with 8 GB RAM and a four-core processor working at 3.3 GHz. As no optimization for parallel computing was implemented, effectively only one core was used. The estimation of the brute force method execution time was based on the assumption that the calculation of the objective function value for a particular bit string and the comparison of the value obtained with the best one found so far takes one nanosecond (10^{-9} s). The assumption is based on the gigahertz (10^9 Hz) speed of the CPU clock. As more than one instruction (or processor clock tick) is needed for the described evaluation, our estimations are probably overly optimistic. However, they well document the need for more sophisticated methods for finding the optimum of exponentially scaling problems. Note that the age of the Universe is around 13.77×10^9 years.

Source: Author’s own calculations

Appendix C: Expected Returns, Transaction Costs and Covariance Matrices

Great Recession

	USD	EUR	AUD	CAD	GBP	SEK	JPY	CNY	Gold
<i>Returns</i>	-0.80%	0.05%	5.00%	1.87%	2.02%	5.90%	-1.10%	1.73%	6.28%
<i>Costs</i>	0.26%	0.27%	0.23%	0.22%	0.21%	0.27%	0.20%	0.23%	0.25%
<i>CovVar</i>									
USD	1.90%	0.67%	0.77%	1.11%	1.20%	0.44%	1.81%	1.89%	1.24%
EUR	0.67%	0.48%	0.36%	0.45%	0.54%	0.35%	0.74%	0.67%	0.53%
AUD	0.77%	0.36%	1.34%	0.94%	0.73%	0.53%	0.69%	0.76%	0.90%
CAD	1.11%	0.45%	0.94%	1.28%	0.84%	0.51%	0.96%	1.10%	0.95%
GBP	1.20%	0.54%	0.73%	0.84%	1.29%	0.46%	1.22%	1.20%	1.01%
SEK	0.44%	0.35%	0.53%	0.51%	0.46%	0.91%	0.41%	0.44%	0.50%
JPY	1.81%	0.74%	0.69%	0.96%	1.22%	0.41%	2.78%	1.80%	1.51%
CNY	1.89%	0.67%	0.76%	1.10%	1.20%	0.44%	1.80%	1.91%	1.25%
Gold	1.24%	0.53%	0.90%	0.95%	1.01%	0.50%	1.51%	1.25%	3.87%

European Debt Crisis

	USD	EUR	AUD	CAD	GBP	SEK	JPY	CNY	Gold
<i>Returns</i>	6.12%	1.75%	0.84%	0.89%	2.51%	1.38%	0.74%	5.36%	1.40%
<i>Costs</i>	0.22%	0.18%	0.31%	0.20%	0.27%	0.19%	0.13%	0.18%	0.17%
<i>CovVar</i>									
USD	1.25%	0.35%	0.69%	0.87%	0.75%	0.42%	1.02%	1.23%	0.67%
EUR	0.35%	0.26%	0.24%	0.28%	0.29%	0.24%	0.35%	0.35%	0.26%
AUD	0.69%	0.24%	1.28%	0.87%	0.63%	0.42%	0.62%	0.71%	0.70%
CAD	0.87%	0.28%	0.87%	1.12%	0.68%	0.41%	0.69%	0.87%	0.65%
GBP	0.75%	0.29%	0.63%	0.68%	1.01%	0.39%	0.59%	0.76%	0.42%
SEK	0.42%	0.24%	0.42%	0.41%	0.39%	0.69%	0.37%	0.42%	0.35%
JPY	1.02%	0.35%	0.62%	0.69%	0.59%	0.37%	1.73%	1.01%	0.95%
CNY	1.23%	0.35%	0.71%	0.87%	0.76%	0.42%	1.01%	1.29%	0.67%
Gold	0.67%	0.26%	0.70%	0.65%	0.42%	0.35%	0.95%	0.67%	3.00%

Covid Crisis

	USD	EUR	AUD	CAD	GBP	SEK	JPY	CNY	Gold
<i>Returns</i>	-5.72%	-4.50%	4.34%	-0.25%	0.00%	-1.22%	-9.35%	-0.64%	4.02%
<i>Costs</i>	0.33%	0.21%	0.41%	0.34%	0.25%	0.44%	0.30%	0.29%	0.34%
<i>CovVar</i>									
USD	0.96%	0.44%	0.30%	0.54%	0.40%	0.24%	0.72%	0.81%	0.39%
EUR	0.44%	0.35%	0.18%	0.26%	0.23%	0.22%	0.40%	0.39%	0.30%
AUD	0.30%	0.18%	0.74%	0.47%	0.40%	0.35%	0.22%	0.31%	0.41%
CAD	0.54%	0.26%	0.47%	0.70%	0.39%	0.29%	0.33%	0.49%	0.28%
GBP	0.40%	0.23%	0.40%	0.39%	0.60%	0.25%	0.36%	0.37%	0.30%
SEK	0.24%	0.22%	0.35%	0.29%	0.25%	0.52%	0.21%	0.25%	0.30%
JPY	0.72%	0.40%	0.22%	0.33%	0.36%	0.21%	0.86%	0.63%	0.54%
CNY	0.81%	0.39%	0.31%	0.49%	0.37%	0.25%	0.63%	0.82%	0.47%
Gold	0.39%	0.30%	0.41%	0.28%	0.30%	0.30%	0.54%	0.47%	2.45%

Note: Returns stands for expected returns, costs for average transaction costs and *CovVar* for covariance matrix.

Source: Author's own calculation based on the CNB's FX rates list and Bloomberg data

Appendix D: Numerical Results

Testing Version

Algorithm	USD	EUR	AUD	CAD	GBP	SEK	JPY	CNY	Gold
Grad.	0.0%	26.2%	17.1%	0.0%	0.0%	52.9%	0.0%	0.0%	3.8%
BB	0.0%	28.0%	18.2%	0.0%	0.0%	49.9%	0.0%	0.0%	3.9%
SA	0.0%	22.2%	17.1%	1.4%	3.9%	49.1%	0.2%	0.0%	6.2%
GA	14.6%	9.6%	9.7%	0.5%	11.0%	27.1%	1.7%	5.8%	20.3%
D-Wave QPU	0.2%	12.7%	20.4%	8.6%	4.4%	42.2%	2.1%	0.3%	9.2%
D-Wave hyb.	0.1%	28.0%	15.5%	0.2%	0.0%	53.0%	0.0%	0.2%	2.9%

Practical Version

Period	Algorithm	USD	EUR	AUD	CAD	GBP	SEK	JPY	CNY	Gold
Great Rec.	Grad.	0.0%	31.7%	19.8%	1.2%	5.9%	35.6%	0.0%	0.0%	5.7%
	BB	0.0%	25.0%	19.7%	2.6%	9.0%	37.5%	0.0%	0.0%	6.2%
	SA	0.5%	18.4%	12.4%	2.1%	7.2%	50.5%	0.8%	0.6%	7.5%
	D-Wave hyb.	0.1%	36.2%	18.5%	0.3%	2.7%	37.1%	0.1%	0.2%	4.8%
Debt crisis	Grad.	6.2%	42.6%	15.2%	0.0%	6.4%	23.5%	0.0%	4.8%	1.4%
	BB	6.2%	37.5%	12.5%	0.0%	9.4%	25.0%	0.0%	6.3%	3.1%
	SA	6.6%	21.4%	19.6%	0.4%	7.9%	25.0%	0.5%	14.2%	4.4%
	D-Wave hyb.	10.1%	50.0%	12.1%	0.1%	6.5%	15.9%	0.0%	4.1%	1.1%
Covid	Grad.	0.0%	20.1%	32.7%	0.6%	12.7%	21.2%	0.0%	3.7%	8.9%
	BB	0.0%	16.9%	25.0%	3.1%	16.4%	23.9%	0.0%	4.7%	10.0%
	SA	0.2%	3.9%	50.6%	4.6%	5.7%	12.4%	0.6%	12.2%	9.9%
	D-Wave hyb.	0.0%	26.7%	31.2%	0.2%	11.3%	17.3%	0.0%	4.2%	9.1%

Note: *Grad.* stands for classical gradient algorithm, *BB* for branch and bound algorithm, *SA* for simulated annealing, *GA* for genetic-based algorithm, *D-Wave QPU* for D-Wave quantum annealer and *D-Wave hyb.* for D-Wave constrained hybrid heuristic. The gradient solver is highlighted in red, as it is the benchmark.

Source: Author's own calculation

CNB Working Paper Series (since 2022)

WP 1/2023	Martin Veselý	<i>Finding the optimal currency composition of foreign exchange reserves with a quantum computer</i>
WP 10/2022	Martin Hodula Milan Szabo Josef Bajžík	<i>Retail fund flows and performance: Insights from supervisory data</i>
WP 9/2022	Jiří Gregor Jan Janků Martin Melecký	<i>From central counter to local living: Pass-through of monetary policy to mortgage lending rates in districts</i>
WP 8/2022	Simona Malovaná Martin Hodula Zuzana Gric Josef Bajžík	<i>Borrower-based macroprudential measures and credit growth: How biased is the existing literature?</i>
WP 7/2022	Martin Časta	<i>How credit improves the exchange rate forecast</i>
WP 6/2022	Milan Szabo	<i>Meeting investor outflows in Czech bond and equity funds: Horizontal or vertical?</i>
WP 5/2022	Róbert Ambriško	<i>Nowcasting macroeconomic variables using high-frequency fiscal data</i>
WP 4/2022	Jaromír Baxa Jan Žáček	<i>Monetary policy and the financial cycle: International evidence</i>
WP 3/2022	Martin Hodula Milan Szabo Lukáš Pfeifer Martin Melecký	<i>Cooling the mortgage loan market: The effect of recommended borrower-based limits on new mortgage lending</i>
WP 2/2022	Martin Veselý	<i>Application of quantum computers in foreign exchange reserves management</i>
WP 1/2022	Vojtěch Molnár	<i>Price level targeting with imperfect rationality: A heuristic approach</i>

CNB Research and Policy Notes (since 2022)

RPN 2/2022	Jan Filáček Lucie Kokešová Matějková	<i>Disclosing dissent in monetary policy committees</i>
RPN 1/2022	Oxana Babecká Kucharčuková Jan Brůha Petr Král Martin Motl Jaromír Tonner	<i>Assessment of the nature of the pandemic shock: Implications for monetary policy</i>

CZECH NATIONAL BANK
Na Příkopě 28
115 03 Praha 1
Czech Republic

ECONOMIC RESEARCH DIVISION
Tel.: +420 224 412 321
Fax: +420 224 412 329
<http://www.cnb.cz>
e-mail: research@cnb.cz

ISSN 1803-7070