(article starts on next page)

Benchmarking Advantage and D-Wave 2000Q quantum annealers with exact cover problems

Dennis Willsch[1] · **Madita Willsch**[1,2] · **Carlos D. Gonzalez Calaza**[1] · **Fengping Jin**[1] · **Hans De Raedt**[1,3] · **Marika Svensson**[4,5] · **Kristel Michielsen**[1,2,6]

Abstract

We benchmark the quantum processing units of the largest quantum annealers to date, the 5000+ qubit quantum annealer Advantage and its 2000+ qubit predecessor D-Wave 2000Q, using tail assignment and exact cover problems from aircraft scheduling scenarios. The benchmark set contains small, intermediate, and large problems with both sparsely connected and almost fully connected instances. We find that Advantage outperforms D-Wave 2000Q for almost all problems, with a notable increase in success rate and problem size. In particular, Advantage is also able to solve the largest problems with 120 logical qubits that D-Wave 2000Q cannot solve anymore. Furthermore, problems that can still be solved by D-Wave 2000Q are solved faster by Advantage. We find, however, that D-Wave 2000Q can achieve better success rates for sparsely connected problems that do not require the many new couplers present on Advantage, so improving the connectivity of a quantum annealer does not per se improve its performance.

✉ Dennis Willsch
  d.willsch@fz-juelich.de

1   Institute for Advanced Simulation, Jülich Supercomputing Centre Forschungszentrum Jülich, 52425 Jülich, Germany

2   AIDAS, 52425 Jülich, Germany

3   Zernike Institute for Advanced Materials, University of Groningen, Nijenborgh 4, 9747 AG Groningen, The Netherlands

4   Jeppesen, 411 03 Gothenburg, Sweden

5   Department of Computer Science and Engineering, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

6   RWTH Aachen University, 52056 Aachen, Germany

🖄 Springer

1 Introduction

Quantum annealing is a quantum computing paradigm that relies on quantum fluctuations to solve optimization problems [1–10]. In September 2020, D-Wave Systems has released a quantum annealer with a 5000+ qubit quantum processing unit (QPU) called Advantage [11]. This system has more than twice as many qubits as its predecessor D-Wave 2000Q and an increase in qubit connectivity from 6 to 15 by using the Pegasus topology [12–15]. High expectations have been placed on its computational power, and first independent studies have become available [16–23]. For such a rapidly developing technology, it is an important task for independent researchers to study progress and test new developments.

Conceptually, there are three classes of benchmarks for quantum annealers:

(1) Comparison with detailed real-time simulations of quantum annealing systems based on solving the time-dependent Schrödinger equation [24, 25] or the time-dependent master equation [5, 6, 26–30].
(2) Direct QPU benchmarks (including comparison with other quantum annealing systems and optimization problem solvers) for problems of intermediate size that may or may not need embeddings and solve either real-world or artificial problems [11, 31–41].
(3) Benchmarks of hybrid solvers that use a combination of QPUs and CPUs or GPUs to solve large-scale application problems [12, 18, 42, 43].

The experiments reported in this paper focus on benchmarking the bare QPU performance, thus belonging to benchmarking class (2).

We assess the progress in quantum annealing technology by benchmarking both Advantage and D-Wave 2000Q with exact cover problems. The exact cover problem is an NP-complete problem [44] that has become a prominent application to study quantum annealing [45–49] and gate-based quantum computing [50–54]. In our case, the exact cover problems represent simplified aircraft scheduling scenarios. They are derived from the tail assignment problem [55] (see [53] for more information), which belongs to the class of set covering and partitioning problems that is covered in a tremendous amount of literature on both applications and solution techniques in operations research (see e.g., [56, 57]). Related aircraft assignment problems on quantum annealers have been studied in [58–60].

The essence of the present exact cover problems is shown in Fig. 1a. In this case, we are given 40 flight routes. Each route contains several out of 472 flights. The task is to find a selection of flight routes such that all 472 flights are covered *exactly* once. On the quantum annealer, each route is represented by a qubit. If a route is to be selected, the corresponding qubit ends up in the state $|1\rangle$ after the measurement. Eventually, each selected route shall be assigned to one airplane. We remark that solving the problem with a given set of routes is a simplification of the general case.

The difficulty of the problem for a quantum computer can be seen by the following counting argument: For 40 routes (i.e., 40 qubits), the number of possible selections is $2^{40} \approx 10^{12}$. For 120 routes (i.e., 120 qubits), which are the largest problems that are solved in the present benchmark set, the number of selections already grows to
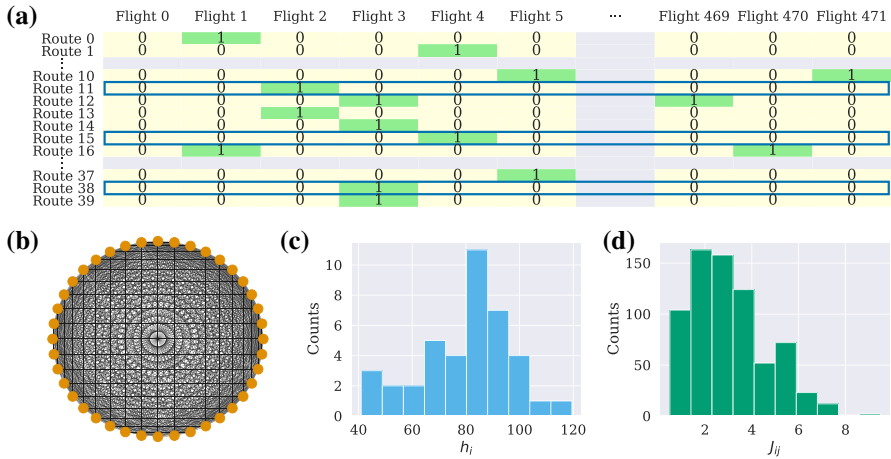
(a)

(b) **(c)** **(d)**

Fig. 1 Visualization of the exact cover problem with 40 logical qubits (instance 0). **a** Boolean matrix A defining the exact cover problem instance (see Eq. (13)). For each of the 40 routes, the matrix A indicates all flights that are covered by this route. The exact cover problem is to find a selection of routes (i.e., a subset of rows of the matrix A) such that all 472 flights are covered exactly once (meaning that the sum of the selected rows contains only ones). For this problem, rows belonging to this solution are indicated with blue boxes. The full solution is given by the ground state $|01000010100101010000001010001000000000000\rangle$ (the rightmost qubit corresponds to route 0). **(b)** Coupler graph of the Ising formulation of this problem (cf. Eq. (2)), where each nonzero J_{ij} corresponds to a black line between the 40 qubits. With 711 out of all $\binom{40}{2} = 780$ couplers being nonzero, this problem is almost fully connected (see also Appendix C). A distribution of the values of the Ising parameters of this problem is shown in **(c)** for the qubit biases h_i given by Eq. (10) and **(d)** for the couplers J_{ij} given by Eq. (11))

10^{36}. Hence, exact cover problems are well suited for benchmarking the Advantage and D-Wave 2000Q QPUs.

We find that Advantage outperforms D-Wave 2000Q on almost all problems in the present benchmark set up to 120 logical qubits. Advantage can embed and solve larger problems. Furthermore, the time-to-solution on Advantage is at least roughly a factor of two shorter. Advantage scores better success rates for all problems with almost all-to-all connectivity. Only some problems with a very sparse qubit connectivity have sometimes higher success rates on D-Wave 2000Q.

The remainder of this paper is structured as follows. In Sect. 2, we describe the mathematical details associated with the exact cover problems under investigation. In Sect. 3, we present and discuss the results that Advantage and D-Wave 2000Q have produced for both small-scale and large-scale exact cover problems. Section 4 contains our conclusions.

2 Methods

This section presents the mathematical details behind the problems under investigation and how they are solved on the quantum annealers. We first outline the type of optimization problems that Advantage and D-Wave 2000Q can solve, including the

important distinction between *physical* and *logical* qubits. Then, we describe the tail assignment and exact cover problems under investigation, along with their formulation on the D-Wave 2000Q and Advantage QPUs.

2.1 QUBO and Ising problems

The QPUs produced by D-Wave Systems are designed to solve binary quadratic models (BQMs), i.e., quadratic optimization problems over discrete variables that can each take two different values. BQMs are typically formulated as quadratic unconstrained binary optimization (QUBO) models or Ising models:

$$\text{QUBO}: \quad \min_{x_i=0,1} \left(\sum_{i \leq j} x_i Q_{ij} x_j + C_1 \right), \tag{1}$$

$$\text{Ising}: \quad \min_{s_i=\pm 1} \left(\sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j + C_2 \right), \tag{2}$$

where the indices i and j range over all qubits. In the QUBO model in Eq. (1), the problem is defined by the QUBO matrix Q with values $Q_{ij} \in \mathbb{R}$, and the binary problem variables are $x_i = 0, 1$. In the Ising model in Eq. (2), the problem is defined by the biases $h_i \in \mathbb{R}$ and the couplers $J_{ij} \in \mathbb{R}$, and the problem variables are $s_i = \pm 1$. An example distribution of h_i and J_{ij} of the problems under investigation is shown in Fig. 1c and d, respectively.

It is worth mentioning that on the D-Wave QPUs, all problems are internally converted into Ising models and (if the autoscaling feature is on) rescaled by a constant factor such that $h_i \in [-2, 2]$ and $J_{ij} \in [-1, 1]$ [61]. (Note that these ranges might be different for future QPUs.) To convert between QUBO and Ising formulation, we use the quantum annealing convention (see also [62])

$$x_i = \frac{1 + s_i}{2}, \tag{3}$$

so that $x_i = 0$ ($x_i = 1$) maps to $s_i = -1$ ($s_i = 1$). Note that in the literature on gate-based quantum computing, also the alternative $x_i = (1 - s_i)/2$ is often used [63] (which would result in a change of sign for the qubit biases, $h_i \mapsto -h_i$).

The constants C_1 and C_2 in Eqs. (1) and (2) do not affect the solution of the problem. However, they can be used to shift the energy (i.e., the value of the objective function at the solution). We use it to shift the energy of the ground state to zero so that we can conveniently determine the success rate by counting all solutions with energy zero.

2.2 Physical and logical qubits

Many problems may require nonzero couplers Q_{ij} or J_{ij} between different qubits i and j that do not physically exist on the QPUs. For instance, the 40-qubit problem

sketched in Fig. 1b has almost all-to-all connectivity. In this case, solving the problem on a QPU requires the concept of *embedding* the problem on a QPU.

Conventionally, the qubits that physically exist on a QPU are called *physical* qubits. On a D-Wave 2000Q QPU, the 2000+ physical qubits are connected in a *Chimera* topology [7]. This means that each physical qubit is connected to 6 other physical qubits on average. On an Advantage QPU, the Chimera topology has been upgraded to the *Pegasus* topology [11]. This means that nearly all of the 5000+ physical qubits are connected to 15 other physical qubits, increasing the connectivity by a factor of 2.5.

When problems require a larger connectivity than provided by the QPU, the effective connectivity between the qubits can be increased by combining several physical qubits into a *logical* qubit. The QUBO and Ising models in Eqs. (1) and (2) are typically formulated in terms of such logical qubits, and not the underlying physical qubits. The physical qubits that form a logical qubit are called a *chain*. To ensure that physical qubits within a chain function as a single logical qubit, the couplers J_{ij} between them are set to a reasonably large, negative value called `chain_strength`. If a chain between two qubits *breaks* (i.e., if the product of the Ising variables is $s_i s_j = -1$), a penalty of $2 \times$ `chain_strength` is added to the energy.

We define the `chain_strength` in terms of the *relative chain strength* RCS $\in [0, 1]$ according to

$$\texttt{chain\_strength} = \text{RCS} \times \texttt{max\_strength}, \qquad (4)$$

where `max_strength` $= \max\left(\{|h_i|\} \cup \left\{\left|J_{ij}\right|\right\}\right)$ is the maximum absolute value of all h_i and J_{ij}. For instance, for the problem sketched in Fig. 1, `max_strength` would be 119.5.

An *embedding* is a mapping from each logical qubit to a chain of physical qubits. An important property of embeddings are the chain lengths. In our experience, embeddings with too large chains may result in a poor quality of the solutions produced by a QPU. In Appendix C, we provide more details on the specific chains encountered in the embeddings (see Fig. 7).

The D-Wave Ocean SDK [64] provides algorithms to automatically generate qubit embeddings with a given value for `chain_strength`. Still, when using a QPU, finding and characterizing embeddings is an important step and may considerably affect the quality of the solution. For this reason, as part of the present benchmark, we systematically investigate different embeddings and relative chain strengths in Sect. 3.1.

Note that the Pegasus topology is an extension of the Chimera topology, so that a Chimera graph can be natively embedded in a Pegasus graph (see Appendix B and also [18]). We made use of this relation for the experiments presented in Appendix B. For all other experiments, the embeddings onto the Chimera and Pegasus topologies were generated independently using the D-Wave Ocean SDK.

2.3 Tail assignment problem

The problem instances considered in this work are derived from the tail assignment problem [55]. The tail assignment problem is a fundamental component of aircraft assignment problems, i.e., the problem of assigning flights to individual airplanes, identified by their tail number. The objective is to minimize the overall cost subject to certain constraints such as minimum connection times, airport curfews, maintenance, and preassigned activities. The general role of tail assignment problems in aircraft scheduling and their relation to the column generation technique [65] and the branch-and-price algorithm [66] is described in detail in [53].

We consider a simple form of the tail assignment problem given by

$$\text{minimize} \qquad \sum_{r=0}^{R-1} c_r x_r, \tag{5}$$

$$\text{subject to} \qquad \sum_{r=0}^{R-1} A_{rf} x_r = 1 \quad \forall f = 0, \ldots, F-1, \tag{6}$$

where $r = 0, \ldots, R-1$ enumerates all flight routes, $f = 0, \ldots, F-1$ enumerates the flights, $x_r \in \{0, 1\}$ are the Boolean problem variables with $x_r = 1$ if route r is to be selected, c_r is the cost of selecting route r, and $A \in \{0, 1\}^{R \times F}$ is the Boolean problem matrix with $A_{rf} = 1$ if flight f is contained in route r (see Fig. 1a for an example of the matrix A). Further models for the tail assignment problem can be found in [51, 53, 55, 60].

A BQM version of the tail assignment problem given by Eqs. (5) and (6) is

$$\min_{\vec{x} \in \{0,1\}^N} \left(\lambda \vec{c}^T \vec{x} + \left(A^T \vec{x} - \vec{b} \right)^2 \right), \tag{7}$$

where the number of qubits $N = R$ is given by the number of flight routes, $\vec{c} = (c_0, \ldots, c_{R-1})^T$ contains the costs, and $\vec{b} = (1, \ldots, 1)^T$ is an F-dimensional vector of ones. Note that the scaling factor λ in Eq. (7) is the inverse of the penalty multiplier that would determine the scale of the constraint $A\vec{x} = \vec{b}$ in Eq. (6). It was put in front of the objective function $\vec{c}^T \vec{x}$ so that the exact cover version of the problem corresponds to $\lambda = 0$ (see Sect. 2.4).

We obtain the QUBO formulation of the tail assignment problem by multiplying out the square in Eq. (7) and collecting all terms into the general QUBO model Eq. (1). An outline of the calculation is given in Appendix A. After doing this, we can read off the entries of the QUBO matrix as

$$Q_{ij} = \begin{cases} (2AA^T)_{ij} & (i < j) \\ (AA^T)_{ii} - (2A\vec{b})_i + \lambda c_i & (i = j) \end{cases}, \tag{8}$$

$$C_1 = \vec{b}^T \vec{b}. \tag{9}$$

Note that the qubit indices $i, j \in \{0, \ldots, N-1\}$ correspond to the previous route index $r \in \{0, \ldots, R-1\}$.

Finally, the Ising formulation of the tail assignment problem is found by using Eq. (3) to replace the qubit variables x_i by spin variables s_i in the QUBO model Eq. (1). After collecting linear, quadratic, and constant terms, we find the expressions for the coefficients of the Ising model Eq. (2):

$$h_i = \sum_j \frac{1}{2}(AA^T)_{ij} - (A\vec{b})_i + \frac{1}{2}\lambda c_i, \tag{10}$$

$$J_{ij} = \frac{1}{2}(AA^T)_{ij}, \tag{11}$$

$$C_2 = C_1 + \sum_{i<j} \frac{1}{2}(AA^T)_{ij}$$

$$+ \sum_i \frac{1}{2}((AA^T)_{ii} - (2A\vec{b})_i + \lambda c_i). \tag{12}$$

These expressions hold for general values of A and $\vec{b}$. (Details on the calculation can also be found in Appendix A.) A characteristic distribution of the biases h_i and the couplers J_{ij} is shown in Fig. 1c and d, respectively.

2.4 Exact cover problem

The exact cover problem is an NP-complete set partitioning problem [44]. In matrix form, it can be written as

$$\min_{x_r=0,1} \sum_{f=0}^{F-1} \left(\sum_{r=0}^{R-1} A_{rf}x_r - 1 \right)^2, \tag{13}$$

and its purpose can be directly understood from Fig. 1a: The selection of routes (i.e., rows) with $x_r = 1$ has to be such that each flight f in the problem matrix A is covered exactly once.

The exact cover problem corresponds to the feasibility version of the tail assignment problem given by the sole constraints in Eq. (6), without the objective function in Eq. (5). Formally, the exact cover problems are obtained by setting $\lambda = 0$ in Eq. (7), which yields Eq. (13).

Hence, the QUBO (Ising) coefficients of the exact cover problem are given by Eq. (8) (Eqs. (10) and (11)) for $\lambda = 0$ (see also [47]). Note that the explicit expressions for the constants in Eqs. (9) and (12) are not relevant for the solution of the problem on the quantum annealer. However, it is convenient to add them to the resulting energies to ensure that the energy minimum is zero, because then we can determine the success rate by counting the occurrences of samples with energy zero.

The problems have been generated using the column generation method as described in detail in [53]. Additional properties of the exact cover problems, including the

Fig. 2 Success rates for exact cover problems with 30–40 qubits (instance 0) as a function of the relative chain strength (bottom axis) on **a** D-Wave 2000Q and **b** Advantage with default annealing time $20\,\mu$s. The scan of the relative chain strength is repeated for 10 different, randomly generated embeddings (represented by different colors) and with 10 repetitions each to gather statistics. Markers indicate the corresponding standard deviation above and below the mean. Filled areas between the markers are guides to the eye. The curves for the success rates as a function of the RCS are representative of the other problem instances 1, 2, and 3 characterized in Appendix C
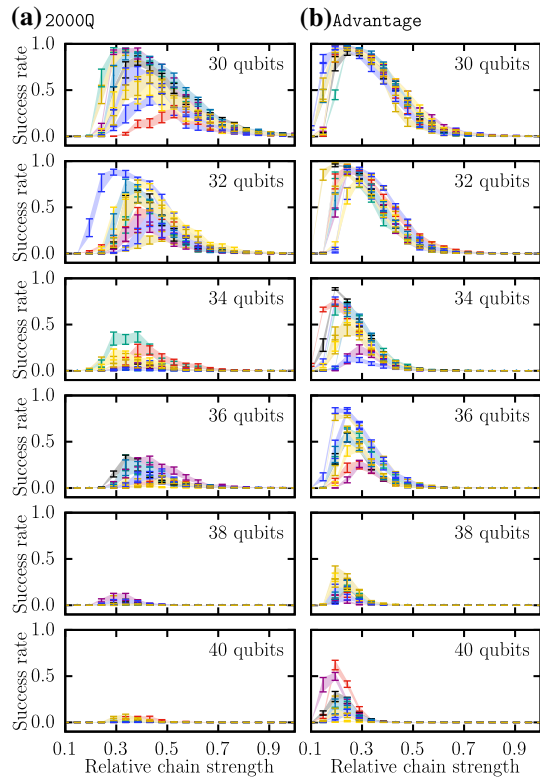

number of logical couplers and physical qubits used in the generated embeddings, can be found in Appendix C.

3 Results

In this section, we present the benchmark results for Advantage and D-Wave 2000Q. We first consider small and intermediate exact cover problems with 30–40 logical qubits and almost full connectivity, with a focus on comparing different embeddings and annealing times. Then, we proceed to large exact cover problems with up to 120 logical qubits, with a focus on the success rate and the time that it takes the QPUs to solve the problems. We note that all considered problem sizes, although large for current quantum annealers, are still small for classical solvers and can be solved within a few minutes using, for instance, a linear programming solver. For the sake of completeness, we also present results for tail assignment problems with $\lambda \neq 0$ (cf. Eq. (7)) in Appendix D.

The experiments reported in this section were performed in August and September 2020 using the solver `DW_2000Q_VFYC_6` for D-Wave 2000Q and the solvers `Advantage_beta` and `Advantage_system1.1` for Advantage. We stress that the goal of the present study was to compare the different QPUs under equivalent

Fig. 3 Success rates as a function of the annealing time on **a** D-Wave 2000Q and **b** Advantage, averaged over 10 repetitions. Different markers indicate different problem instances: 0 (markers with error bars), 1 (pluses), 2 (crosses), 3 (circles). For each instance, the runs in each panel are performed with the same embeddings and relative chain strengths, characterized in Fig. 7 in Appendix C. For instance 0 (highlighted in color), these parameters correspond to the best configuration of the corresponding panel in Fig. 2. Additionally, for instance 0, the standard deviation from the 10 repetitions is indicated by the filled areas between the error bars
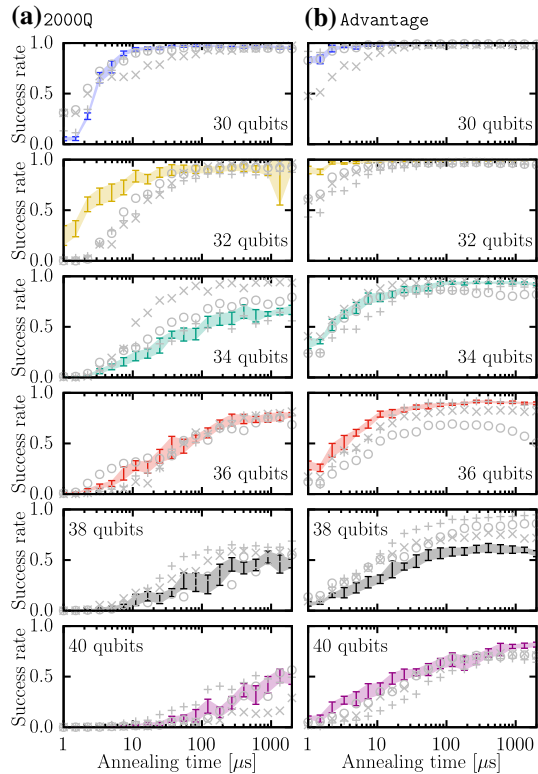
conditions. For this reason, we did not make use of any post-processing features to solely improve the quality of the results. Unless otherwise stated, all QPU settings were left at their default values.

3.1 Densely connected problems with 30–40 qubits

The problems with $N = 30, 32, 34, 36, 38, 40$ qubits and almost full connectivity (cf. Appendix C) are exact cover problems from aircraft scheduling scenarios with 472 flights. Each qubit represents a flight route that contains some of the 472 flights (cf. Fig. 1). We remark that by construction, the ground state of each problem instance is unique and contains 9 qubits in state $|1\rangle$. We obtain the success rate by counting the number of samples with energy zero. (Note that the value of the constant Eq. (12) can be used to shift the energy accordingly.)

3.1.1 Characterizing embeddings and chain strengths

For each problem with $N = 30, 32, 34, 36, 38, 40$ qubits, we generate 10 different embeddings on both D-Wave 2000Q and Advantage. For each embedding, we scan the RCS (cf. Eq. (4)). The results are shown in Fig. 2.

We see that the results on Advantage (Fig. 2b) are generally much better than on D-Wave 2000Q, especially for larger problems. The reason for this is that the 30–40 qubit problems are almost fully connected (cf. Figs. 1b and 6). Hence, in such cases, the user can clearly profit from the much larger connectivity between qubits on the Pegasus topology.

This observation is in line with the fact that the chains on Advantage for the same problem are much shorter (see Fig. 7b). We also see that for increasing problem size, generating multiple embeddings and tuning the chain strengths is crucial to obtain good results when using the bare QPUs.

3.1.2 Varying the annealing time

Next, we study the influence of the annealing time on the success rate. The reason is that for a quantum annealer, the annealing time is a central parameter with an expectable influence on the success rate: From the adiabatic theorem [67], comparably long annealing times are expected to lead to reasonable success rates. Given that the system is not completely isolated from an environment, however, long annealing times may lead to an increasing influence of noise on the system and thus in fact degrade the performance. Therefore, experiments beyond the default annealing time are vital for reliable benchmarks of quantum annealers.

We first select the best embedding and relative chain strength for each problem, based on the results from the previous section. For problem instance 0, for example, this configuration corresponds to the position of the peaks in each panel in Fig. 2. For the selected configuration, we then replace the default annealing time of 20 μs by 20 different, logarithmically spaced annealing times in the QPU annealing time range [1 μs, 2000 μs]. The results for the 30–40 qubit problems and all four problem instances are shown in Fig. 3.

Comparing D-Wave 2000Q and Advantage, we can make the following observations: First, Advantage reaches higher maximum success rates, typically for the longest annealing time. Second, the success rates on Advantage are already reasonable for short annealing times. This means that Advantage is typically faster than D-Wave 2000Q (see also the comparison of QPU access times in Fig. 4b in the following section). Finally, the fluctuations over 10 repetitions are smaller on Advantage (see the filled areas in Fig. 3). We note that this observation also holds for the three instances whose statistics are not indicated in Fig. 3 (gray markers). Thus, we conclude that Advantage shows a demonstrable advantage over D-Wave 2000Q.

3.2 Large problems with 50–120 qubits

In this section, we consider exact cover problems of a larger size with 50 to 120 logical qubits. The goal is to assess the performance of the QPUs for larger but more sparsely connected problem instances. For each problem size, we consider six problem instances. Each corresponds to an aircraft scheduling problem of the type sketched in Fig. 1a with 535 flights. As before, the ground state of these exact cover problems is unique and known. It has 40 qubits in state $|1\rangle$.
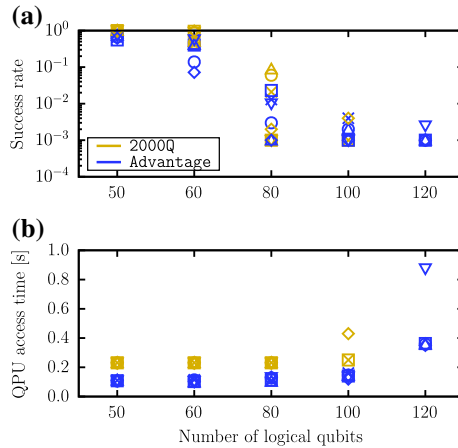
Fig. 4 Solution of exact cover problems with 50–120 logical qubits on both D-Wave 2000Q (yellow) and Advantage (blue). Shown are **(a)** the success rate as a function of the number of logical qubits and **(b)** the QPU access time (which mainly consists of programming, annealing, and readout times [68]) . Different markers indicate different problem instances: 0 (crosses), 1 (squares), 2 (circles), 3 (up-pointing triangles), 4 (down-pointing triangles), 5 (diamonds). The embedding and the relative chain strength for these runs was selected using the same procedure as in Sect. 3.1

These large problems require many more physical qubits so that their embedded versions can occupy a large part of the QPUs. For this reason, the success rates may be smaller, and especially D-Wave 2000Q may not be able to solve the largest problems anymore.

Indeed, as Fig. 4a shows, only Advantage can solve five out of the six 120 qubit instances. The success rates for the 80 and 100 qubit instances are comparable for D-Wave 2000Q and Advantage. Only for the 50 and 60 qubit instances, we see that D-Wave 2000Q yields sometimes better success rates than Advantage.

The reason for this can be understood by studying the number of couplers required for these problems: As shown in Fig. 6 in Appendix C, the 50 and 60 qubit problems need comparably little couplers, so they have a sparser connectivity. Therefore, many of the physical couplers present in the Pegasus topology on Advantage are not required. In other words, the problems can already be embedded well enough on the Chimera topology. This observation is in line with a similar observation for Advantage reported in [18].

We conjecture that in this case, the additional unused connections between qubits on Advantage may disturb the annealing process (additional experimental evidence for this conjecture is given in Appendix B), because even if they are not used, they still exist physically. The reason for our conjecture lies in a well-known physical mechanism: If two subsystems of a larger system are coupled, the system's Hamiltonian contains a corresponding interaction term. For flux qubits manufactured by D-Wave Systems, this interaction is often mediated by tunable interqubit inductive couplers [5]. Such an interaction is always on, due to the existence of a physical link between the subsystems, even if its scale is very weak or if the coupler is "programmed to zero." During the time evolution (which may take comparably long for an annealing process), the

wave function of the system typically develops an error over time, which may have a detrimental impact on the function of a quantum computer. The effect can be easily observed for a system of coupled spins, for example the NMR quantum computing system [69]. Furthermore, it can also be found in superconducting qubit systems; see, e.g., the cross talk experiment in [70], the performance of two-qubit gates mediated by a coupler such as the resonator in [71] or the flux-tunable transmon in [72], or the theoretical study on ZZ couplings in [73]. We leave a detailed simulation of the present system for future research.

Finally, the large problems may require more time on the QPUs, so that they are well suited to compare QPU run times between D-Wave 2000Q and Advantage. Unless mentioned otherwise, the same number of reads and annealing times have been used for D-Wave 2000Q and Advantage to ensure a fair comparison. The timing results are shown in Fig. 4b.

For 50–80 qubits, the annealing time is less than or equal to $20\,\mu s$, so for 1000 reads it does not make up a significant fraction of the QPU access time. Under these conditions, we see that D-Wave 2000Q still needs at least twice as long to solve the problem. Thus, we infer that it is a speedup in programming and readout times [68] that make Advantage faster than D-Wave 2000Q.

For problems with 100 and 120, the annealing time needs to be significantly increased to find the ground state, which is visible in the QPU access times. For instance, problem instance 5 for 100 qubits on D-Wave 2000Q (the single yellow diamond at 100 qubits in Fig. 4b) corresponds to an annealing time of $200\,\mu s$. The same annealing time is required for problem instances 1, 2, 3, and 5 for 120 qubits on Advantage (the blue cluster at 120 qubits in Fig. 4b). Still, Advantage solves these problems faster than D-Wave 2000Q solves the corresponding 100 qubit problem. Only instance 4 for 120 qubits stands out (the blue down-pointing triangle in Fig. 4b): Here, an annealing time of $2000\,\mu s$ was required to find a solution. In order not to exceed the maximum run time on Advantage, the number of reads was reduced to 400. Thus, $400 \times 2000\,\mu s = 0.8\,s$ makes up the largest part of the QPU access time in this case.

4 Conclusion

In this paper, we have benchmarked the performance of the 2000+ qubit quantum annealer D-Wave 2000Q and the 5000+ qubit quantum annealer Advantage. The benchmark suite consists of intermediate and large exact cover problems from aircraft scheduling scenarios with both sparse and dense logical qubit connectivity.

We observed a considerable increase in performance on Advantage. First, Advantage was able to solve exact cover problems with up to 120 logical qubits that were unsolvable on D-Wave 2000Q. Second, the success rates produced by Advantage were almost always higher. Third, Advantage is approximately twice as fast as D-Wave 2000Q in terms of both programming and readout times. Additionally, the required annealing times to solve a problem on Advantage are often shorter. Finally, the fluctuations in success rates over several repetitions on Advantage were smaller.

A large part of the increase in performance can be attributed not only to the larger number of physical qubits, but rather to the increase in qubit connectivity: Every qubit in the Pegasus topology is connected to 15 other qubits, as compared to 6 other qubits in the Chimera topology used in D-Wave 2000Q. We observed chain lengths in the embeddings that were roughly smaller by a factor of two. We could only observe better performance on D-Wave 2000Q for problems with very sparse qubit connectivity or when using the same embedding on D-Wave 2000Q and Advantage. The sparse problems could already be well embedded on the Chimera topology. We conclude that increasing the number of couplers does not necessarily improve the performance of a quantum annealer. Instead, whether an improvement or a degradation is observed depends on the connectivity of each problem instance. On the one hand, increasing the number of couplers allows to embed bigger problem instances, and it may also reduce the required chain length and thus the chain strength. On the other hand, we observed that if only a few couplers are needed, it may be better to use a processor with lower connectivity (see also [18]).

We also observed that, when the same embedding is used on the D-Wave 2000Q processor with connectivity 6 and the Advantage processor with connectivity 15, the results on the D-Wave 2000Q processor are significantly better (see Appendix B), suggesting that the additional couplers introduce additional noise if they are not needed. In particular, this supports the conclusion that for our experiments, the better embeddings (and not a reduced noise level) are the reason for the improved performance of Advantage over D-Wave 2000Q. Thus, a higher connectivity only yields an advantage if the embedding of the problem can actually be improved. An interesting project for future work would be a systematic study of the question whether, for a given embedding, there is some "minimal improvement" required to overcome the influence of additional unused couplers to also yield an improvement in the success probability.

When using the bare QPUs, it is essential to scan several embeddings and chain strengths to find optimal results. Furthermore, it is important to tune the annealing time. We note that besides the bare QPUs, we have also submitted all exact cover problems of our benchmark set to D-Wave's hybrid solver services, which use a combination of QPUs and classical solvers to solve much larger problems [12]. All exact cover problems could be solved by the hybrid solvers `hybrid_v1` (using D-Wave 2000Q) and `hybrid_binary_quadratic_model_version2` (using Advantage) on September 14, 2020. See [18] for more detailed benchmarks of the hybrid solvers with problems of up to 12000 variables.

Our benchmark study confirms the consistent increase in both size and performance of quantum annealers over the past years. For the future, it is an interesting question whether D-Wave Systems will prove capable of keeping up the steep progress of doubling qubit numbers and increasing performance and qubit connectivity at the same time.

Appendix A: Derivation of QUBO and Ising models

In this appendix, we outline the derivation of the QUBO and Ising coefficients from the BQM in Eq. (7) for the tail assignment and exact cover problems under investigation.

To obtain the QUBO formulation, we first multiply out the square in Eq. (7),

$$\left(A^T \vec{x} - \vec{b}\right)^2 = \left(A^T \vec{x} - \vec{b}\right)^T \left(A^T \vec{x} - \vec{b}\right) \tag{A1}$$

$$= \vec{x}^T A A^T \vec{x} - \vec{b}^T A^T \vec{x} - \vec{x}^T A \vec{b} + \vec{b}^T \vec{b} \tag{A2}$$

$$= \sum_{ij} x_i (A A^T)_{ij} x_j - \sum_i (2A\vec{b})_i x_i + \vec{b}^T \vec{b}. \tag{A3}$$

After splitting the first sum into three parts, $\sum_{ij} = \sum_{i<j} + \sum_{i>j} + \sum_i$, and exchanging $i \leftrightarrow j$ in the second part (since the matrix AA^T is symmetric), we obtain the upper triangular coefficients of the QUBO matrix, $Q_{ij} = (2AA^T)_{ij}$ for $i < j$, and the first part of the diagonal coefficients $(AA^T)_{ii}$. The second sum in Eq. (A3) yields, after using $x_i = x_i x_i$, the remaining part of the diagonal coefficients. The last term yields the constant contribution to the QUBO model. Combining this with the linear term $\lambda \vec{c}^T \vec{x}$ in Eq. (7), we obtain all coefficients of the QUBO model $\sum_{i \leq j} x_i Q_{ij} x_j + C_1$,

$$Q_{ij} = \begin{cases} (2AA^T)_{ij} & (i < j) \\ (AA^T)_{ii} - (2A\vec{b})_i + \lambda c_i & (i = j) \end{cases}, \tag{A4}$$

$$C_1 = \vec{b}^T \vec{b} = F. \tag{A5}$$

To obtain the coefficients h_i and J_{ij} of the corresponding Ising model, we replace the qubit variables by spin variables, $x_i = (1 + s_i)/2$ (cf. Eq. (3)),

$$\sum_{i \le j} x_i Q_{ij} x_j = \sum_{i \le j} \frac{(1 + s_i)}{2} Q_{ij} \frac{(1 + s_j)}{2} \tag{A6}$$

$$= \sum_{i < j} \frac{(1 + s_i)}{2} (2AA^T)_{ij} \frac{(1 + s_j)}{2} + \sum_i \frac{(1 + s_i)}{2} ((AA^T)_{ii} - (2A\vec{b})_i)$$

$$+ \lambda c_i) \frac{(1 + s_i)}{2} \tag{A7}$$

$$= \sum_{i < j} \frac{1}{2} (AA^T)_{ij} s_i + \sum_{i > j} \frac{1}{2} (AA^T)_{ij} s_i + \sum_i \frac{1}{2} (AA^T)_{ii} s_i$$

$$- \sum_i (A\vec{b})_i s_i + \sum_i \frac{1}{2} \lambda c_i s_i \tag{A8}$$

$$+ \sum_{i < j} \frac{1}{2} (AA^T)_{ij} s_i s_j \tag{A9}$$

$$+ \sum_{i < j} \frac{1}{2} (AA^T)_{ij} + \sum_i \frac{1}{2} ((AA^T)_{ii} - (2A\vec{b})_i + \lambda c_i), \tag{A10}$$

where we used that the matrix AA^T is symmetric and that $s_i^2 = 1$. Note that this calculation does not make use of prior knowledge about the values of A and $\vec{b}$.

We can then identify the coefficients of the Ising model $\sum_i h_i s_i + \sum_{i < j} J_{ij} s_i s_j + C_2$ as

$$h_i = \sum_j \frac{1}{2} (AA^T)_{ij} - (A\vec{b})_i + \frac{1}{2} \lambda c_i, \tag{A11}$$

$$J_{ij} = \frac{1}{2} (AA^T)_{ij}, \tag{A12}$$

$$C_2 = C_1 + \sum_{i < j} \frac{1}{2} (AA^T)_{ij} + \sum_i \frac{1}{2} ((AA^T)_{ii} - (2A\vec{b})_i + \lambda c_i). \tag{A13}$$

Appendix B: Additional experiments for the conjecture about unused couplers

Based on the results for the sparse 50–120 qubit problems presented in Fig. 4, we formulated the conjecture that the additional couplers physically present on Advantage may disturb the results if they are not needed. One might wonder whether the effect could be mitigated by using logical chains with redundant physical qubits in the embeddings.

To provide additional support for our conjecture from the dense 30–40 qubit problems, we perform the following experiments: We take the best Chimera embedding
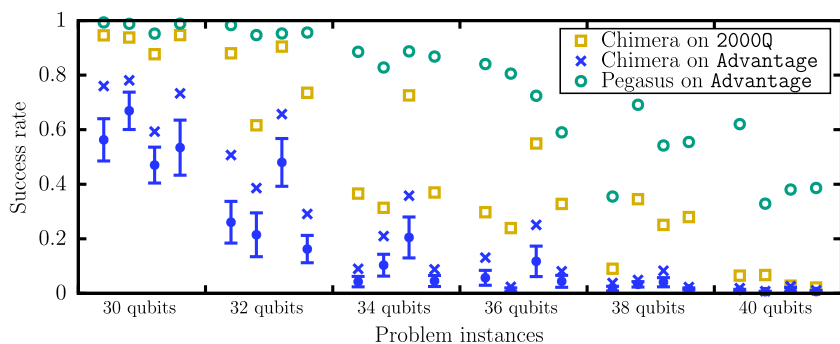
Fig. 5 Additional experiments to study the conjecture about unused couplers. Shown are the maximum success rates found for each embedding and relative chain strength, using the best Chimera embedding on D-Wave 2000Q (yellow squares), the same Chimera embedding but used on Advantage (blue crosses), and the best Pegasus embedding on Advantage (green circles). Additionally, for "Chimera on Advantage," solid blue circles and error bars indicate the mean and standard deviation, respectively, over all different positions where the Chimera embedding can be placed on the Pegasus graph of Advantage (i.e., over all possible values of Δx, Δy, and z in Eq. (B1))

found for D-Wave 2000Q (for problem instance 0, these correspond to the maxima in Fig. 2a) and use this same embedding to solve the problem on Advantage.

Since the Pegasus graph on Advantage is quite large, it can embed the Chimera subgraphs in several different places (cf. Fig. 2 in [18]). We run experiments for all possible relative displacements $\Delta x \in \{-15, \ldots, 15\}$, $\Delta y \in \{-15, \ldots, 15\}$, and $z \in \{0, 1, 2\}$ that generate valid embeddings. Physical qubits $q_C \in \{0, \ldots, 2048 - 1\}$ on the Chimera graph are mapped to physical qubits $q_P \in \{0, \ldots, 5760 - 1\}$ on the Pegasus graph using the formula

$$q_P = 60 + 180x + y + 60z$$
$$+ 15(q_C \bmod 4) + \begin{cases} 0 & (\text{if } q_C \bmod 8 = 0, 1, 2, 3) \\ 3060 - 179(x - y) - 120z & (\text{if } q_C \bmod 8 = 4, 5, 6, 7) \end{cases},$$
$$(B1)$$

where $x = \lfloor (q_C \bmod 128)/8 \rfloor + \Delta x$, $y = \lfloor q_C/128 \rfloor + \Delta y$, and $z = 0, 1, 2$ denote the corresponding unit cell within the graphs. (This formula has been obtained by comparing and relating the different physical qubit labels on D-Wave 2000Q and Advantage; the same result can also be achieved using tools from D-Wave NetworkX [74].)

The results are shown in Fig. 5. For each valid embedding, we scan the relative chain strengths (analogously to Fig. 2). We only consider results for the best relative chain strengths found for each embedding. As expected from Fig. 2, the best Pegasus embeddings on Advantage (green circles) outperform the best Chimera embeddings on D-Wave 2000Q (yellow squares). However, the Chimera embeddings on Advantage (blue crosses) perform significantly worse than on D-Wave 2000Q. In particular, as the blue circles with error bars show, this worse performance is independent of $(\Delta x, \Delta y, z)$,
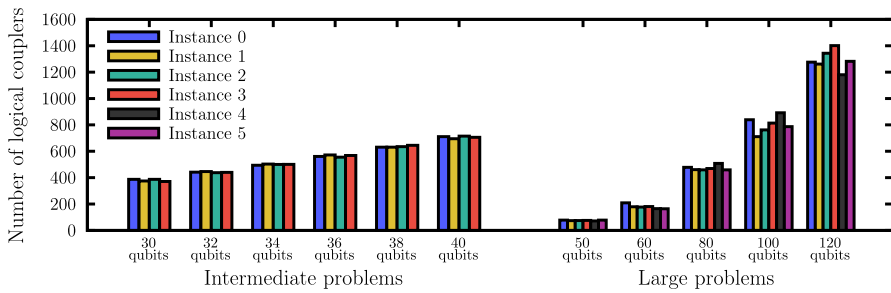
Fig. 6 Number of logical couplers required for the exact cover problems in the present benchmark set. The intermediate problems with 30–40 qubits contain 4 different problem instances for each problem size (left). The large problems with 50–120 qubits contain 6 different problem instances (right). Note that after the embedding, the numbers of physical qubits and couplers required on the QPUs may be much larger (cf. Fig. 7a)

i.e., independent of where the Chimera embeddings are placed on the larger Pegasus graph.

We draw two conclusions from these results: First, the better overall performance on Advantage can be tied most strongly to the improved connectivity of the device, which causes the best embeddings to have much shorter chain lengths (cf. Fig. 7b). Second, when the same embeddings (and thus the same chains) are used on both devices, the main difference lies in the many additional unused physical qubits coupling to the chains. Therefore, the significantly worse performance shown in Fig. 5 provides direct support for the conjecture that the extra couplers on Advantage can negatively impact the performance if they are not needed.

Appendix C: Exact cover problem details

In this appendix, we provide details on the exact cover problems used in the present benchmark set. For each problem instance, Fig. 6 shows the number of couplers required between the qubits. Note that the 30–40 qubit problems require almost all-to-all connectivity.

In Fig. 7, we provide details on the generated embeddings for the intermediate 30–40 qubit problems. In Fig. 7a, we list the number of physical qubits required in the embeddings on D-Wave 2000Q and Advantage. As these problems require almost full connectivity, the number of physical qubits is much larger than the number of logical qubits, especially on D-Wave 2000Q. We see that D-Wave 2000Q needs more than twice as many physical qubits, but also the slope as a function of the logical qubits is steeper. This is reasonable as also the number of logical couplers increases (cf. Fig. 6), and with sparser connectivity on the Chimera topology, more physical qubits need to be chained into a logical qubit.

This trend is also visible when looking at the chain lengths shown in Fig. 7b: While the chains on Advantage stay almost constant for increasing problem sizes, the chains on D-Wave 2000Q grow longer on average. Especially for 40 qubits, chains on D-Wave
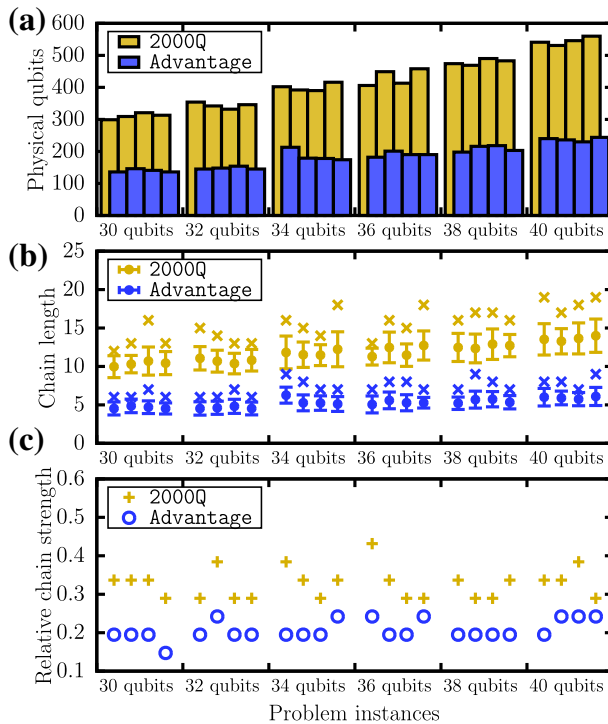
Fig. 7 Details on the optimal embeddings found for the intermediate exact cover problems with $N = 30, 32, 34, 36, 38, 40$ logical qubits on D-Wave 2000Q (yellow) and Advantage (blue). For each N, the four problem instances 0, 1, 2, and 3 are shown from left to right on the bottom axes. **(a)** Number of physical qubits needed for the best embeddings; **(b)** lengths of the corresponding physical qubit chains (cf. Sect. 2.2), where the markers with error bars indicate the mean and the standard deviation, and crosses indicate the maximum chain lengths; **(c)** optimal relative chain strengths, taken from the positions of the corresponding peaks in Fig. 2

2000Q can be up to 19 physical qubits long. Such chains are almost always broken and may lead to a wrong value for the logical qubit that they represent (cf. Sect. 2.2).

In Fig. 7c, we plot the relative chain strengths that produced the best results. For each $N = 30, 32, 34, 36, 38, 40$, the first point (instance 0) corresponds to the peak with the optimal success rate in the corresponding panel in Fig. 2.

Appendix D: Tail assignment problems with $\lambda \neq 0$

The tail assignment problem introduced in Sect. 2.3 contains an objective function that represents the cost associated with each flight route (cf. Eq. (5)). Depending on the magnitude of these cost terms, the multiplier λ in the BQM version of the problem (see Eq. (7)) has to be adjusted to put a reasonable weight on the objective function with respect to the constraints. Therefore, we test several values of λ for a 25-qubit problem. For each λ, we generate 10 embeddings (cf. Sect. 2.2) and evaluate the success rate
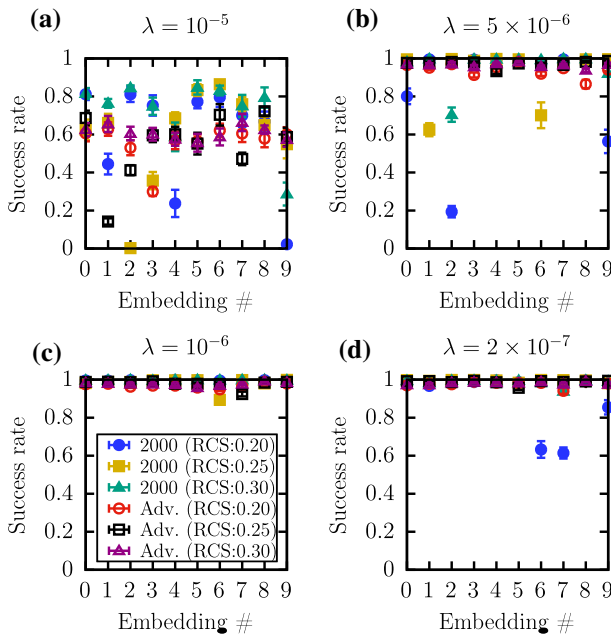
Fig. 8 Success rates for the tail assignment problem on D-Wave 2000Q (filled markers) and Advantage (unfilled markers) as a function of the relative chain strength RCS given by Eq. (4) (see legend). Each run is performed for 10 different embeddings (bottom axis). For a given embedding, the markers indicate the mean and the standard deviation for 10 repetitions using the same embedding. The formulation of the tail assignment problem given by Eq. (7) for 25 qubits is solved for decreasing values of the scaling factor **(a)** $\lambda = 10^{-5}$, **(b)** $\lambda = 5 \times 10^{-6}$, **(c)** $\lambda = 10^{-6}$, **(d)** $\lambda = 2 \times 10^{-7}$

on D-Wave 2000Q and Advantage. The unique ground state was found using both a linear program solver and exact enumeration of all 2^{25} states on a GPU. We remark that the 25-qubit tail assignment problem solved in this section is the same problem that was investigated as the largest problem instance in [51].

The results as a function of the embedding are shown in Fig. 8. We see that for relatively large λ, the success rates for D-Wave 2000Q fluctuate strongly as a function of the embedding. In such cases, it may become possible to violate some of the constraints to obtain a better value of the objective function (cf. Eq. (7)). However, as λ approaches zero (Fig. 8b–d) and the problem approaches its exact cover version, most embeddings yield the optimal solution with unit probability, especially on Advantage (unfilled markers).

References

1. Apolloni, B., Carvalho, C., de Falco, D.: Quantum stochastic optimization. Stoch. Process. Their Appl. **33**, 233 (1989)
2. Finnila, A., Gomez, M., Sebenik, C., Stenson, C., Doll, J.: Quantum annealing: a new method for minimizing multidimensional functions. Chem. Phys. Lett. **219**, 343 (1994)

3. Kadowaki, T., Nishimori, H.: Quantum annealing in the transverse Ising model. Phys. Rev. E **58**, 5355 (1998)

4. Brooke, J., Bitko, D., Rosenbaum, T.F., Aeppli, G.: Quantum annealing of a disordered magnet. Science **284**, 779 (1999)

5. Harris, R., Johnson, M.W., Lanting, T., Berkley, A.J., Johansson, J., Bunyk, P., Tolkacheva, E., Ladizinsky, E., Ladizinsky, N., Oh, T., Cioata, F., Perminov, I., Spear, P., Enderud, C., Rich, C., Uchaikin, S., Thom, M.C., Chapple, E.M., Wang, J., Wilson, B., Amin, M.H.S., Dickson, N., Karimi, K., Macready, B., Truncik, C.J.S., Rose, G.: Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor. Phys. Rev. B **82**, 024511 (2010)

6. Johnson, M.W., Amin, M.H.S., Gildert, S., Lanting, T., Hamze, F., Dickson, N., Harris, R., Berkley, A.J., Johansson, J., Bunyk, P., Chapple, E.M., Enderud, C., Hilton, J.P., Karimi, K., Ladizinsky, E., Ladizinsky, N., Oh, T., Perminov, I., Rich, C., Thom, M.C., Tolkacheva, E., Truncik, C.J.S., Uchaikin, S., Wang, J., Wilson, B., Rose, G.: Quantum annealing with manufactured spins. Nature **473**, 194 (2011)

7. Bunyk, P.I., Hoskinson, E.M., Johnson, M.W., Tolkacheva, E., Altomare, F., Berkley, A.J., Harris, R., Hilton, J.P., Lanting, T., Przybysz, A.J., Whittaker, J.: Architectural considerations in the design of a superconducting quantum annealing processor. IEEE Trans. Appl. Supercond. **24**, 1 (2014)

8. Job, J., Lidar, D.: Test-driving 1000 qubits. Quantum Sci. Technol. **3**, 030501 (2018)

9. Hauke, P., Katzgraber, H.G., Lechner, W., Nishimori, H., Oliver, W.D.: Perspectives of quantum annealing: methods and implementations. Rep. Prog. Phys. **83**, 054401 (2020)

10. Nath, R.K., Thapliyal, H., Humble, T.S.: A review of machine learning classification using quantum annealing for real-world applications. SN Comput. Sci. **2**, 365 (2021)

11. McGeoch, C., Farré, P.: The D-wave advantage system: an overview, Tech. Rep. (D-Wave Systems Inc, Burnaby, BC, Canada, D-Wave Technical Report Series 14-1049A-A(2020)

12. McGeoch, C., Farré, P., Bernoudy, W.: D-wave hybrid solver service + advantage: technology update, Tech. Rep. (D-Wave Systems Inc, Burnaby, BC, Canada, D-Wave User Manual 09-1109A-V(2020)

13. D-Wave Systems: Technical description of the D-wave quantum processing unit, Tech. Rep. (D-Wave Systems Inc., Burnaby, BC, Canada, 2020) D-Wave User Manual 09-1109A-V

14. Boothby, K., Bunyk,P., Raymond, J., Roy, A.: Next-generation topology of D-wave quantum processors, arXiv:2003.00133 [quant-ph] (2020)

15. King, A.D., Bernoudy, W.: Performance benefits of increased qubit connectivity in quantum annealing 3-dimensional spin glasses, arXiv:2009.12479 [quant-ph] (2020)

16. Cohen, J., Alexander, C.: Picking efficient portfolios from 3,171 US common stocks with new quantum and classical solvers, arXiv:2011.01308 [quant-ph] (2020)

17. Kuramata, M., Katsuki, R., Nakata, K.: Larger sparse quadratic assignment problem optimization using quantum annealing and a bit-flip heuristic algorithm, arXiv:2012.10135 [quant-ph] (2020)

18. Calaza, C.D.G., Willsch, D., Michielsen, K.: Garden optimization problems for benchmarking quantum annealers, arXiv:2101.10827 [quant-ph] (2021)

19. Birdal, T., Golyanik,V., Theobalt, C., Guibas, L.: Quantum permutation synchronization, arXiv:2101.07755 [quant-ph] (2021)

20. Bhatia, H.S., Phillipson, F.: Performance analysis of support vector machine implementations on the D-wave quantum annealer. In: Paszynski, M., Kranzlmüller, D., Krzhizhanovskaya, V.V., Dongarra, J.J., Sloot, P.M.A. (eds.) Computational Science - ICCS. Springer International Publishing, Cham (2021)

21. Fox, D.M., Branson, K.M.,Walker, R.C.: mRNA codon optimization on quantum computers, bioRxiv 10.1101/2021.02.19.431999 (2021)

22. Rahman, S.A.,Lewis, R.,Mendicelli, E.,Powell, S.: SU(2) lattice gauge theory on a quantum annealer, arXiv:2103.08661 [hep-lat] (2021)

23. Phillipson, F., Wezeman, R.S., Chiscop, I.: Indoor–outdoor detection in mobile networks using quantum machine learning approaches. Computers **10**, 71 (2021)

24. Willsch, M., Willsch, D., Jin, F., De Raedt, H., Michielsen, K.: Real-time simulation of flux qubits used for quantum annealing. Phys. Rev. A **101**, 012327 (2020)

25. Willsch, M.: Study of quantum annealing by simulating the time evolution of flux qubits, Ph.D. thesis, RWTH Aachen University, Aachen (2020a)

26. Boixo, S., Albash, T., Spedalieri, F.M., Chancellor, N., Lidar, D.A.: Experimental signature of programmable quantum annealing. Nat. Commun. **4**, 2067 (2013)

27. Albash, T., Rønnow, T., Troyer, M., Lidar, D.: Reexamining classical and quantum models for the D-Wave one processor. Eur. Phys. J. Spec. Top. **224**, 111 (2015)
28. Albash, T., Vinci, W., Mishra, A., Warburton, P.A., Lidar, D.A.: Consistency tests of classical and quantum models for a quantum annealer. Phys. Rev. A **91**, 042314 (2015)
29. Boixo, S., Smelyanskiy, V.N., Shabani, A., Isakov, S.V., Dykman, M., Denchev, V.S., Amin, M.H., Smirnov, A.Y., Mohseni, M., Neven, H.: Computational multiqubit tunnelling in programmable quantum annealers. Nat. Commun. **7**, 10327 (2016)
30. Marshall, J., Venturelli, D., Hen, I., Rieffel, E.G.: Power of pausing: advancing understanding of thermalization in experimental quantum annealers. Phys. Rev. Appl. **11**, 044083 (2019)
31. Boixo, S., Rønnow, T.F., Isakov, S.V., Wang, Z., Wecker, D., Lidar, D.A., Martinis, J.M., Troyer, M.: Evidence for quantum annealing with more than one hundred qubits. Nat. Phys. **10**, 218 (2014)
32. Rønnow, T.F., Wang, Z., Job, J., Boixo, S., Isakov, S.V., Wecker, D., Martinis, J.M., Lidar, D.A., Troyer, M.: Defining and detecting quantum speedup. Science **345**, 420 (2014)
33. Hall, J., Novotny, M., Neuhaus, T., Michielsen, K.: A study of spanning trees on a D-wave quantum computer. Phys. Proc. **68**, 56 (2015)
34. Hen, I., Job, J., Albash, T., Rønnow, T.F., Troyer, M., Lidar, D.A.: Probing for quantum speedup in spin-glass problems with planted solutions. Phys. Rev. A **92**, 042325 (2015)
35. McGeoch, C.C.: Benchmarking D-wave quantum annealing systems: some challenges, in *Electro-Optical and Infrared Systems: Technology and Applications XII; and Quantum Information Science and Technology*, Vol. 9648, edited by D. A. Huckridge, R. Ebert, M. T. Gruneisen, M. Dusek, and J. G. Rarity, International Society for Optics and Photonics (SPIE, 2015) pp. 264 – 273
36. Novotny, M., Hobl, Q.L., Hall, J., Michielsen, K.: Spanning tree calculations on d-wave 2 machines. J. Phys. Conf. Ser. **681**, 012005 (2016)
37. Li, R.Y., Di Felice, R., Rohs, R., Lidar, D.A.: Quantum annealing versus classical machine learning applied to a simplified computational biology problem. npj Quantum Inf. **4**, 14 (2018)
38. Willsch, M., Willsch, D., Jin, F., De Raedt, H., Michielsen, K.: Benchmarking the quantum approximate optimization algorithm. Quantum Inf. Process. **19**, 197 (2020)
39. Willsch, D., Willsch, M., De Raedt, H., Michielsen, K.: Support vector machines on the D-Wave quantum annealer. Comput. Phys. Commun. **248**, 107006 (2020)
40. Cavallaro, G., Willsch, D., Willsch, M., Michielsen,K., Riedel, M.: Approaching remote sensing image classification with ensembles of support vector machines on the D-wave quantum annealer, in *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium* (2020) pp. 1973–1976
41. Domino, K., Koniorczyk, M., Krawiec, K., Jałowiecki, K., Gardas, B.: Quantum computing approach to railway dispatching and conflict management optimization on single-track railway lines, arXiv:2010.08227 [cs.ET] (2021)
42. Mugel, S., Kuchkovsky, C., Sanchez, E., Fernandez-Lorenzo, S., Luis-Hita, J., Lizaso, E., Orus, R.: Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks, arXiv:2007.00017 [quant-ph] (2020)
43. Grozea, C., Hans, R., Koch, M., Riehn, C., Wolf, A.: Optimising rolling stock planning including maintenance with constraint programming and quantum annealing, arXiv:2109.07212 [cs.AI] (2021)
44. Karp, R.M.: Reducibility among combinatorial problems, in *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*, edited by R. E. Miller, J. W. Thatcher, and J. D. Bohlinger Springer US, Boston, MA, 1972) pp. 85–103
45. Farhi, E., Goldstone, J., Gutmann, S., Lapan, J., Lundgren, A., Preda, D.: A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. Science **292**, 472 (2001)
46. Choi, V.: Adiabatic quantum algorithms for the NP-complete maximum-weight independent set, exact cover and 3SAT problems, arXiv:1004.2226 [quant-ph] (2010)
47. Lucas, A.: Ising formulations of many NP problems. Front. Phys. **2**, 5 (2014)
48. Cao, Y., Jiang, S., Perouli, D., Kais, S.: Solving set cover with pairs problem using quantum annealing. Sci. Rep. **6**, 33957 (2016)
49. Sax, I., Feld, S., Zielinski, S., Gabor, T., Linnhoff-Popien,C., Mauerer, W.: Approximate approximation on a quantum annealer, in *Proceedings of the 17th ACM International Conference on Computing Frontiers*, CF '20 (Association for Computing Machinery, New York, NY, USA, 2020) pp. 108–117

50. Farhi, E., Goldstone, J., Gutmann, S.: A quantum approximate optimization algorithm, arXiv:1411.4028 [quant-ph] (2014)
51. Vikstål, P., Grönkvist, M., Svensson, M., Andersson, M., Johansson, G., Ferrini, G.: Applying the quantum approximate optimization algorithm to the tail-assignment problem. Phys. Rev. Appl. **14**, 034009 (2020)
52. Bengtsson, A., Vikstål, P., Warren, C., Svensson, M., Gu, X., Kockum, A.F., Krantz, P., Križan, C., Shiri, D., Svensson, I.-M., Tancredi, G., Johansson, G., Delsing, P., Ferrini, G., Bylander, J.: Improved success probability with greater circuit depth for the quantum approximate optimization algorithm. Phys. Rev. Appl. **14**, 034010 (2020)
53. Svensson, M., Andersson, M., Grönkvist, M., Vikstål, P., Dubhashi, D., Ferrini, G., Johansson, G.: A heuristic method to solve large-scale integer linear programs by combining branch-and-price with a quantum algorithm, arXiv:2103.15433 [quant-ph] (2021)
54. Willsch, D., Willsch, M., Jin, F., Michielsen, K., De Raedt, H.: GPU-accelerated simulations of quantum annealing and the quantum approximate optimization algorithm, arXiv:2104.03293 [quant-ph] (2021)
55. Grönkvist, M.: The tail assignment problem, Ph.D. thesis, Chalmers University of Technology and Göteborg University (2005)
56. Ernst, A., Jiang, H., Krishnamoorthy, M., Sier, D.: Staff scheduling and rostering: a review of applications, methods and models. Eur. J. Oper. Res **153**, 3 (2004)
57. Tahir, A., Desaulniers, G., El Hallaoui, I.: Integral column generation for the set partitioning problem. EURO J. Transp. Logist. **8**, 713 (2019)
58. Stollenwerk, T., Lobe, E., Jung, M.: Flight gate assignment with a quantum annealer. In: *Quantum Technology and Optimization Problems*. Springer International Publishing, pp. 99–110 (2019)
59. Stollenwerk, T., O'Gorman, B., Venturelli, D., Mandrà, S., Rodionova, O., Ng, H., Sridhar, B., Rieffel, E.G., Biswas, R.: Quantum annealing applied to de-conflicting optimal trajectories for air traffic management. IEEE Trans. Intell. Transp. Syst. **21**, 285 (2020)
60. Martins, L.: Applying quantum annealing to the tail assignment problem, Ph.D. thesis, University of Porto (2020)
61. D-Wave Systems, D-wave solver properties and parameters, Tech. Rep. (D-Wave Systems Inc., Burnaby, BC, Canada, 2021) D-Wave User Manual 09-1169A-S
62. D-Wave Systems, D-Wave Problem-Solving Handbook, Tech. Rep. (D-Wave Systems Inc., Burnaby, BC, Canada, 2020) D-Wave User Manual 09-1171A-G
63. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press, New York (2010)
64. D-Wave Systems, D-Wave Ocean SDK, (2020c), https://github.com/dwavesystems/dwave-ocean-sdk, release 2.5.0
65. Desrosiers, J., Lübbecke, M.E.: A primer in column generation. In: Desaulniers, G., Desrosiers, J., Solomon, M.M. (eds.) Column Generation. Springer, Boston (2005)
66. Barnhart, C., Johnson, E.L., Nemhauser, G.L., Savelsbergh, M.W.P., Vance, P.H.: Branch-and-price: column generation for solving huge integer programs. Op. Res. **46**, 316 (1998)
67. Born, M., Fock, V.: Beweis des Adiabatensatzes. Z. Phys. **51**, 165 (1928)
68. D-Wave Systems, Solver Computation Time, Tech. Rep. (D-Wave Systems Inc., Burnaby, BC, Canada, 2021) D-Wave User Manual 09-1107B-A
69. De Raedt, H., Michielsen, K., Hams, A., Miyashita, S., Saito, K.: Quantum spin dynamics as a model for quantum computer operation. Eur. Phys. J. B **27**, 15 (2002)
70. Willsch, D.: Supercomputer simulations of transmon quantum computers, Ph.D. thesis, RWTH Aachen University, Aachen (2020b)
71. Willsch, D., Nocon, M., Jin, F., De Raedt, H., Michielsen, K.: Gate-error analysis in simulations of quantum computers with transmon qubits. Phys. Rev. A **96**, 062302 (2017)
72. Lagemann, H., Willsch, D., Willsch, M., Jin, F., De Raedt, H., Michielsen, K.: Numerical analysis of effective models for flux-tunable transmon systems, in preparation (2021)
73. Xu, X., Ansari, M.: ZZ freedom in two-qubit gates. Phys. Rev. Appl. **15**, 064074 (2021)
74. D-Wave Systems: D-Wave NetworkX, (2021c), https://github.com/dwavesystems/dwave-networkx