



# Boosting Quantum Annealer Performance via Quantum Persistence

Hamed Karimi, Gili Rosenberg

## Abstract

We propose a novel method for reducing the number of variables in quadratic unconstrained binary optimization problems, using a quantum annealer to fix the value of a large portion of the variables to values that have a high probability of being optimal. The resulting problems are usually much easier for the quantum annealer to solve, due to their being smaller and consisting of disconnected components. This approach significantly increases the success rate and number of observations of the best known energy value in samples obtained from the quantum annealer, when compared with calling the quantum annealer without using it, even when using fewer annealing cycles. Use of the method results in a considerable improvement in success metrics even for problems with high-precision couplers and biases, which are more challenging for the quantum annealer to solve. The results are further enhanced by applying the method iteratively and combining it with classical pre-processing. We present results for both Chimera graph-structured problems and embedded problems from a real-world application.

**Keywords:** [adiabatic quantum computation](#), [quantum annealing](#), [variable reduction](#), [quantum persistency](#), [binary optimization](#)

## 1 Previous work

### 1.1 Quantum annealing

Quantum annealers are now commercially available [20, 6]. Manufactured by D-Wave Systems, they are designed to heuristically find low-energy states of an Ising model of the type

$$H = \sum_i h_i s_i + \sum_{\langle i,j \rangle} J_{ij} s_i s_j \quad (1)$$

on a Chimera graph [6]. The quantum annealer used in this study was the D-Wave 2X-SYS4, which has 1152 qubits, with a qubit yield of over 95% (the remainder of the qubits are inoperable).<sup>1</sup> Since numerous well known NP-hard problems have known mappings to an Ising model [28], it is hoped that many intractable problems can be solved more efficiently using quantum annealing.

There is evidence that quantum effects, namely entanglement and multi-qubit tunnelling, do indeed occur on these quantum annealing chips [26, 3, 2, 12]. It has been suggested that quantum annealers have an advantage over classical optimizers due to quantum tunnelling, which an optimizer to search the solution space of an optimization problem by passing through energy barriers. For certain problem classes, this might provide a quantum speedup [38, 13, 21, 42, 1]. For this reason, there has been much recent interest in benchmarking quantum annealers against classical solvers, although conclusive evidence of quantum speedup has not been shown to date [2, 12, 30, 22, 24, 25, 29, 40, 18, 23].

In order to solve a problem with the quantum annealer, the user must provide the biases  $h$  and the couplings  $J$  in Eq. 1. These parameters are affected by perturbations of different types, which are generally referred to as intrinsic control errors (ICE). Due to these errors, which are roughly 2%–4% of the range of the parameters, high-precision problems are challenging for the D-Wave 2X, without post-processing, motivating post-processing and error correction efforts [31, 36, 37, 46, 35, 33, 27], as well as our own efforts in this research.

The hardware graph of the D-Wave 2X quantum annealer is a Chimera graph [6] consisting of a  $12 \times 12$  lattice of complete bipartite units of 8 qubits with sparse connectivity between the unit cells. Most real-world optimization problems are unlikely to have a corresponding Ising model with this structure. To allow the solving of problems with different structure, we use an *embedding*, which is a mapping from each logical variable to one or more physical qubits, referred to as a *chain*. We connect a chain of qubits with a strong ferromagnetic (negative  $J_{ij}$ ) coupling, thereby penalizing solutions whose values are different.

When obtaining a sample from the quantum annealer for an embedded problem, it is possible that not all of the qubits in a chain will have the same state. In this case the values of the corresponding logical variables are commonly determined by either majority voting or a process referred to as energy minimization [46]. In a version of the latter that we employed in this study, we assigned an effective field to each broken chain, chose the chain with the strongest effective

<sup>1</sup>The chip at our disposal has 1100 active qubits, a working temperature of  $26 \pm 5$  mK, and a minimum annealing time of 20  $\mu$ s.

field, set its state to be opposite to that of the direction of the field (to minimize the energy), updated the effective fields, and repeated the process until no broken chains remained. This is a quick method of local error correction. We remark that finding an embedding for a graph is formally known as graph minor embedding, an NP-hard problem [8, 9] commonly solved by heuristic methods, although for some classes of problems deterministic embeddings can be found [49].

## 1.2 Classical pre-processing

Within the combinatorial optimization community, there is a significant body of work on methods for solving quadratic unconstrained binary optimization (QUBO) problems, of the form

$$\begin{aligned} \min x^T Q x \\ \text{s.t. } x \in \{0, 1\}^N, \end{aligned} \tag{2}$$

where  $Q$  is a real  $N \times N$  matrix. QUBO problems can be easily mapped to the above Ising model (with the addition of a constant term). Of interest are classical methods for pre-processing QUBO problems (see Tavares [45] for a review), in particular *persistence* [16, 4], whereby some of the variables are fixed to their values in the optimum. When a variable is fixed, its value is set and the entries in  $Q$  are updated accordingly. In the Ising-based approach, fixing a qubit adds a constant term to the energy and shifts the biases ( $h$  values) of the neighbouring qubits.

An example of classical pre-processing is fixing variables within a heuristic algorithm by maintaining a reference set of elite solutions (typically obtained by local search), and finding the variables that are often set to the same value, the idea being that they are likely to be set to that same value in the optimum [47, 48, 50]. The *backbone* of a problem is defined as the variables and corresponding values which are the same in all optima [32]. This notion has also been extended to problems with a non-degenerate optimum [39]. These techniques, as well as our method, can be described as finding an *approximate backbone* [19].

The closest existing work to our method is Chardaire et al. [7], which uses the term *thermostistical persistence* to describe variables whose values can be fixed in simulated annealing, based on their values remaining constant as the temperature is decreased. Our method relies on fixing variables whose value is the same in the solution obtained in repeated (and independent) annealing cycles, which we refer to as *quantum persistence*. We alternate between using classical pre-processing and applying our method using a quantum annealer. We find that using both methods together leads to a significant boost in the number of variables that are fixed, with a high level of confidence that they are fixed to their value in an optimal solution. Below, we use the Ising representation of this problem exclusively. We remark that due to the one-to-one mapping between Ising and QUBO problems, our method could be described equivalently using the QUBO representation.

The layout of this paper is as follows: our proposed method is presented in Section 2, numerical results are presented in Section 3 and discussed in Section 4, and concluding remarks are given in Section 5.

## 2 Method

In this section, we explain the details of our proposed method and refinements for various cases, as well as the physical intuition behind it. In Section 2.1 we introduce the basic variable-reduction method, and in Section 2.2 we present the iterative method, which builds upon the basic method. In Section 2.3 we explain our physical intuition on why and when this method is expected to succeed. In Section 2.4 we discuss challenges of the zero bias case ( $h = 0$ ) and how we mitigated them within the framework of our method, relying on the physical intuition developed in the previous section. Finally, in Section 2.5 we discuss how we applied our method to embedded problems.

### 2.1 Quantum persistence variable reduction (QPVAR)

The basic idea behind our method is that, for a given problem, we obtain a sample from the quantum annealer, find the variables that have exactly the same value across the entire sample, and fix them to that value. A quantum annealer is a heuristic solver, and a typical call to the machine involves at least hundreds, if not thousands, of quantum annealing cycles, obtaining a sample of low-energy solutions. For a general random Ising problem, with  $J$  and  $h$  chosen randomly, we observed that many spins (variables) maintain their state (value) in the sample obtained from the quantum annealer.

The sample being obtained from low-energy solutions motivated our conjecture that those variables might also maintain their values in the ground state(s) of the system.

We improved on the basic idea in multiple ways:

- Instead of obtaining a sample in one call to the quantum annealer, we obtained a better sample by performing multiple calls, and applied a random gauge to each. It has been shown that quantum annealers perform better when random gauges are applied to problems, since this mitigates errors in the quantum hardware [34].
- Trimming the sample. We gave more weight to lower-energy solutions by ignoring higher-energy solutions of the sample, based on a given *elite\_threshold* percentile. An alternative, which we leave for future study, would be to give higher weights to lower-energy solutions with some given weight function.
- Relaxing the fixing criterion. In our basic method, we fix a variable if it maintains its state in exactly all of the solutions of the sample. We relaxed this criterion by introducing a *fixing\_threshold* parameter, such that variables are only fixed if their mean absolute value (taken across the trimmed sample) is larger than the value of this parameter.

---

**Algorithm 1** QPVAR

---

**Require:** Ising problem  $(J, h)$

Obtain sample from the quantum annealer

Narrow down the solutions to *elite\_threshold* percentile

Find the mean value of each variable in all solutions

Fix variables for which the mean absolute value is larger than *fixing\_threshold*

Update  $J$  and  $h$

**return**  $J, h$  and a mapping from fixed variables to the values to which they were fixed

---

Algorithm QPVAR summarizes the steps of our improved basic method. Fixing variables in this way can drastically reduce the size of the effective problem to be solved. However, if even one variable is fixed to a value contrary to the value of that variable in all of the optima, it becomes impossible to find an optimum by solving the reduced problem. To check whether this occurs, we solved each problem before and after fixing the variables using a heuristic solver with a long time-out, such that it has a high likelihood of finding an optimum. If all of the variables were fixed correctly, the energy of the optimum before fixing the variables is guaranteed to be equal to the energy of the optimum after fixing. We tuned the parameters such that we only fixed variables such that they took their value in one of the optima with high confidence.

We remark that setting *fixing\_threshold* to 1.0 guarantees that if an optimum was found by the quantum annealer, then no variables will be fixed contrary to their value in at least one of the optima. This can be shown by contradiction. Assume one variable is fixed contrary to its value in all of the optima. This would only occur if that variable's value was set to the contrary value in all of the solutions in the sample. However, an optimum was among the solutions in the sample, so at least one solution in the sample has that variable set to a different value (i.e., to the value in the optimum that was found), leading to a contradiction.

Decreasing *fixing\_threshold* allows one to fix more variables, since the fixing criterion is relaxed, and decreasing *elite\_threshold* focuses on fewer solutions, again fixing more variables. In both cases, the increase in the number of fixed variables comes at an increased risk of fixing variables contrary to their optimum value. Changing these parameters allows the user to optimize QPVAR for the class of problems being optimized, as well as based on the user's preference. For example, if finding the optimum is more important than reducing the problem size, the above parameters should be set to high values. Table 1 and Table 2 show the dependence of the number of fixed variables and QPVAR fixing success rate on these two parameters. The fixing success rate was defined as the percentage of problems for which the method fixed variables only to their optimal value.

Detailed results for two particular choices of parameters, and for two problem sets, are presented in Figure 1. In that experiment, QPVAR was applied on a quantum annealer sample with 2500 annealing cycles. We defined the quantum annealer's residual as the energy difference between the quantum annealer's lowest-energy solution and the best known solution (found by a heuristic solver; see Section 3 for more details), and the method's residual as the difference in energy between the best known solution before fixing variables and after fixing variables.

We see that the quantum annealer was able to find the best known solution for most of the low-precision problems (top row), but unable to find the best known solution for almost all of the higher-precision problems (bottom row). The method’s residual was almost always zero, indicating that the method almost never fixed variables incorrectly. In the few cases in which it did fix variables incorrectly, the effective problem (after fixing variables) still conserved the first or second excited states (the energy spacing was exactly 2, due to the integer biases and couplers). We also note that for those problems, the method’s residual was always lower than the quantum annealer’s residual. The mean fraction of variables fixed was 58%–70%, showing that QPVAR was able to fix most of the variables in these problems, and, as shown in Section 3.2, the iterative method presented below was able to fix even more variables.

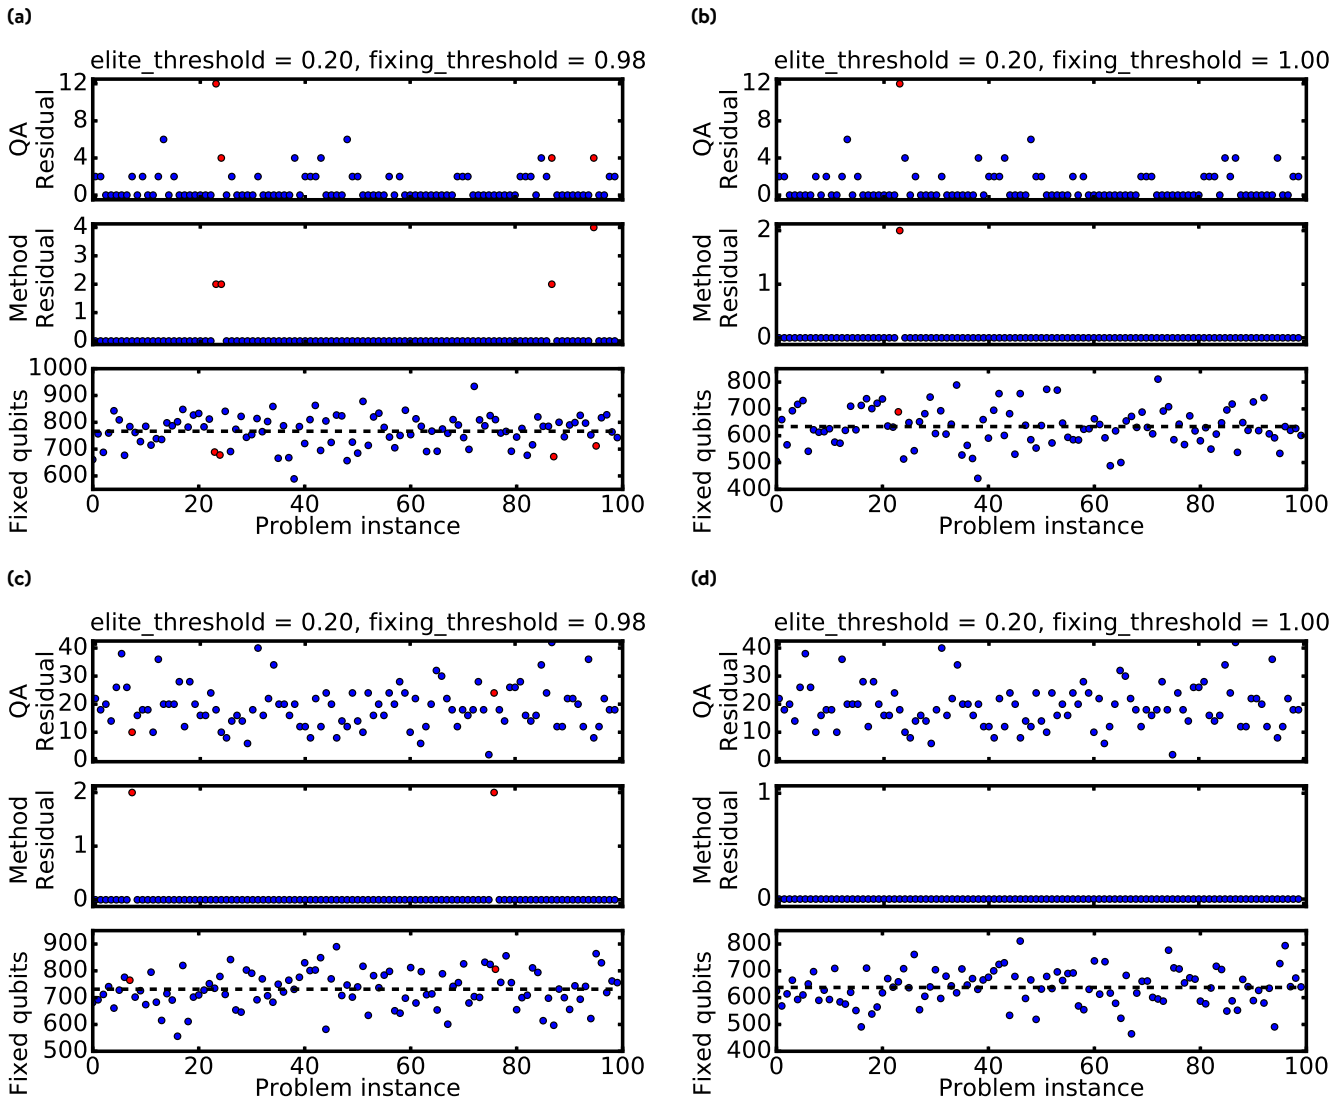


Figure 1: Results for QPVAR on Chimera graph-structured problems. **(a)** and **(b)** present results for problems in  $U_2$  ( $J, h \in \{-2, -1, 0, 1, 2\}$ , excluding 0 for  $J$ ), and **(c)** and **(d)** present results for problems in  $U_{10}$ . The parameters are indicated in the titles above each subfigure. The top panel of each subfigure shows the difference in energy between the quantum annealer’s best solution and the best known solution, for each problem, with 2500 annealing cycles using five random gauges. The middle panel shows the difference in energy between the best known solution before fixing variables and the best known solution after fixing variables (using the same 2500 cycles)—it should be zero if all variables fixed by QPVAR were fixed correctly. The bottom panel shows the number of fixed variables for each problem. A horizontal dashed line indicates the mean number of fixed variables. In all panels, the points that correspond to problems for which QPVAR fixed variables incorrectly are shown in red, and the rest are shown in blue [colour online].

**(a)** Mean number of fixed variables for  $U_2$

E / F	0.95	0.97	0.98	0.99	1.00
0.10	892.17	873.00	858.75	841.43	835.15
0.20	828.85	794.69	766.68	718.34	634.17
0.30	794.09	754.34	723.36	671.98	525.34
0.40	769.61	729.06	697.70	639.81	455.46
0.60	753.80	712.67	680.21	618.56	407.36
0.60	746.37	704.14	670.03	606.01	384.87
0.70	742.21	698.90	663.30	600.43	372.26
0.80	740.25	696.36	659.73	596.22	365.78
0.90	739.40	695.20	658.22	594.90	363.03
1.00	739.04	694.89	661.42	594.29	362.31

**(b)** Mean number of fixed variables for  $U_{10}$

E / F	0.95	0.97	0.98	0.99	1.00
0.10	842.27	819.26	815.84	815.84	815.84
0.20	797.59	757.79	731.56	686.08	638.03
0.30	760.49	718.22	686.74	638.08	515.35
0.40	732.62	688.69	656.41	602.85	427.50
0.60	715.15	669.78	634.54	579.32	368.74
0.60	704.80	656.04	620.44	563.56	332.97
0.70	696.79	646.86	612.21	551.21	312.13
0.80	692.89	642.54	607.15	547.86	300.44
0.90	692.09	641.10	603.88	543.39	294.18
1.00	690.79	639.72	605.27	541.48	289.83

Table 1: Mean number of fixed variables for different *fixing\_threshold* (columns: 'F') and *elite\_threshold* (rows: 'E') using QPVAR on Chimera graph-structured problems (with 1100 variables). Problems were chosen randomly from: **(a)**  $U_2$  ( $J, h \in \{-2, -1, 0, 1, 2\}$ , excluding 0 for  $J$ ) and **(b)**  $U_{10}$  (defined similarly).

(a) QPVAR fixing success rate for  $U_2$

E / F	0.95	0.97	0.98	0.99	1.00
0.10	0.91	0.93	0.94	0.96	0.96
0.20	0.90	0.94	0.96	0.98	0.99
0.30	0.85	0.93	0.96	0.97	0.99
0.40	0.84	0.91	0.94	0.97	0.99
0.60	0.83	0.90	0.93	0.95	0.99
0.60	0.83	0.89	0.93	0.95	0.99
0.70	0.83	0.89	0.93	0.95	0.99
0.80	0.83	0.89	0.93	0.95	0.99
0.90	0.83	0.89	0.93	0.95	0.99
1.00	0.83	0.89	0.93	0.95	0.99

(b) QPVAR fixing success rate for  $U_{10}$

E / F	0.95	0.97	0.98	0.99	1.00
0.10	0.94	0.95	0.95	0.95	0.95
0.20	0.88	0.95	0.98	0.99	1.00
0.30	0.85	0.95	0.98	0.99	1.00
0.40	0.83	0.93	0.99	0.99	1.00
0.60	0.82	0.91	0.96	0.99	1.00
0.60	0.81	0.92	0.97	0.99	1.00
0.70	0.81	0.92	0.98	0.99	1.00
0.80	0.81	0.91	0.98	0.99	1.00
0.90	0.81	0.90	0.98	0.99	1.00
1.00	0.81	0.90	0.98	0.99	1.00

Table 2: QPVAR fixing success rate for different *fixing\_threshold* (columns: 'F') and *elite\_threshold* (rows: 'E') on Chimera graph-structured problems. Problems were chosen randomly from: (a)  $U_2$  and (b)  $U_{10}$  (see the caption of Table 1 for a definition of  $U_2$  and  $U_{10}$ ). The QPVAR fixing success rate is defined as the percentage of problem instances for which the method fixed variables to values which appeared in at least one best known solution, which was found by solving the problems with a heuristic before and after fixing the variables.

We observed that after applying QPVAR, the graph structure of the reduced problems becomes simpler than that of the original problem. For example, classical techniques can benefit from having a smaller treewidth and many nodes with degree one. We return to this observation, and present related results, in Section 3.4.

## 2.2 Iterative quantum persistence variable reduction (IQPVAR)

For maximum effect, QPVAR can be applied repeatedly  $num\_steps$  times, fixing more and more variables. In order to increase the effectiveness of QPVAR, we also used a classical pre-processing method that uses roof duality and weak persistency [5] to fix more variables,<sup>2</sup> after applying QPVAR.

In addition to those variables which maintain their value in the sample, there are other variables which do not have this property but are perfectly correlated with each other and change their value coherently. We fixed additional variables by exploiting these correlations between variables. After each step of QPVAR, we calculated the pair-wise correlations of all of the variables, based on the  $correlation\_elite\_threshold$  lowest-energy solutions in the sample. We then created a graph in which the nodes were the variables and the edge weights were the correlation coefficients, disregarding all correlation coefficients with an absolute value of less than  $correlation\_threshold$ , and found the connected components in this graph. If the call to the classical pre-processor fixed at least one variable in any of the connected components, we fixed the entire component (unless it contained a frustrated loop).

The resulting problems were significantly smaller (see Section 3.4) and easier to solve for the quantum annealer (see Sections 3.2–3.3). Algorithm IQPVAR summarizes the iterative algorithm.

---

### Algorithm 2 IQPVAR

---

**Require:** Ising problem  $(J, h)$   
**for** each step of  $num\_steps$  **do**  
     $J, h \leftarrow$  QPVAR  
     $J, h \leftarrow$  Classical pre-processing  
     $J, h \leftarrow$  Fix variables via correlations  
**end for**  
**return**  $J, h$  and a mapping from fixed variables to the values to which they were fixed

---

## 2.3 Physical intuition

In the limit of very large local fields (biases), much larger than the couplings, it is obvious that to find the ground state of the system we need only consider the local fields and can ignore the couplings. Similarly, for finite local fields, Zintchenko et al. [51] showed that it is possible to find the ground state configuration of a given spin by solving only the problem restricted to the local neighbourhood of that spin. They considered spin glass problems given by Eq. 1, where  $J_{ij}$  are chosen from a normal distribution with a mean of zero and standard deviation of one, and  $h_i$ 's are chosen with standard deviation  $\sigma_h$ .

The size of the relevant neighbourhood  $l$  depends on the strength of the local fields  $h$ . In particular, Zintchenko et al. showed that for spin glass problems on a square lattice  $l \propto \sigma_h^{-0.8}$ . The length scale  $l$  corresponds to the correlation length of the system, and correlations between spins with a distance larger than this length decay exponentially. As a result, by considering a cluster of spins around a given spin  $s_0$  of size larger than  $l$ ,  $s_0$  becomes uncorrelated to the boundary spins of the cluster. By solving the problem on the spins in the cluster with an arbitrary configuration on the boundary, we can find the ground state configuration of  $s_0$ .

It is expected that at finite fields and low energies, due to finite correlations, these clusters define a larger effective spin which experiences an effective field and possibly some small coupling with other clusters. As a result, each cluster will align with its effective local field with high probability. Some of these clusters will experience larger effective fields compared to the others and will therefore be more robust under perturbations. A good low-energy sampler can observe and detect these clusters. This could explain the existence of spins which maintain their value in the low-energy samples obtained from the quantum annealer. The effective low-energy excitations of the system can be described by the effective interaction between these clusters. Flipping the clusters, which results in the spins in each cluster flipping

---

<sup>2</sup>We used the function `fix_variables` in D-Wave Systems' SAPI 2.3.1, which is the solver API used to access the quantum annealer [10].



coherently with other spins in the same cluster, justifies the additional fixing of variables via correlation, as explained in the previous section.

Following this argument, we expect that detecting these clusters and fixing spins becomes more difficult as the correlation length increases, in the limit of smaller local fields. Therefore, we conjectured that QPVAR would be challenged in the limit of  $h \rightarrow 0$ . In order to test this, we chose the couplers uniformly at random over the set  $\{-5, -4, \dots, 4, 5\}$ , but chose  $h_i$  from the set  $\{-n, \dots, n\}$ , increasing  $n$  from 1 to 10. This has the effect of changing the standard deviation of the bias  $\sigma_h$ . The expected result was that increasing  $\sigma_h$  shortens the correlation length, and hence more variables should have been fixed by QPVAR. Figure 2 presents the results of this experiment, which validate our intuition.




Figure 2: Number of fixed qubits versus the range of  $h$ , for  $J$  in  $U_5$ , for different pairs of parameters, *elite\_threshold* and *fixing\_threshold*. QPVAR was applied using 2500 annealing cycles. The couplers were chosen uniformly at random over the set  $\{-5, -4, \dots, 4, 5\}$ , and the biases were chosen from the set  $\{-n, \dots, n\}$ , increasing  $n$  from 1 to 10, which has the effect of changing the standard deviation of the bias  $\sigma_h$ .

For the case of  $h = 0$ , the system has  $\mathcal{Z}_2$  symmetry,  $\vec{s} \rightarrow -\vec{s}$ . In this case, the sample might contain many solutions, along with their degenerate pairs and QPVAR in this form will not work. We present some modifications to deal with this situation in the following section. A simple approach to take is to simply break the  $\mathcal{Z}_2$  symmetry by choosing a random spin and force its value to  $+1$  or  $-1$ . With this approach, we were able to observe the correlation between the randomly chosen spin and the other spins.

For this experiment we created twenty Chimera graph-structured problems, based on the hardware graph of the D-Wave 2X (with 1100 active qubits), with  $h = 0$  and  $J$  chosen uniformly at random over the set  $\{-5, -4, \dots, 4, 5\}$ . For each problem, we first obtained a sample from the quantum annealer. Then we iterated over the qubits, choosing one as our reference qubit. For each reference qubit, we chose all of the solutions such that this qubit's state was up, summed the states of all of the qubits, and saved the results. Finally, we averaged over all of the qubits that had the same length of shortest path for any reference qubit. See Figure 3 for the average of the sum over the solutions (averaged over the problems and the reference qubits) as a function of the length of the shortest path.




Figure 3: Average correlation between a fixed reference qubit and other qubits as a function of the shortest path. The couplers were chosen uniformly at random over the set  $\{-5, -4, \dots, 4, 5\}$  and the biases were set to zero. For each problem, we obtained a sample from the quantum annealer, chose all of the solutions such that a reference qubit's state was 'up', summed the states of all of the qubits, and averaged over all of the qubits that had the same shortest path length.

#### 2.4 Special considerations for $h = 0$ problems

If all of the biases are zero, the system is time reversal invariant, leading to a two-fold degeneracy of all states. Therefore, we expect that the sample will contain many degenerate solution pairs. Since all variables would have an expected value of zero in the sample, this would lead to QPVAR fixing zero variables. For this reason, a modification was necessary for this case.

Without loss of generality, we fixed one variable of our choice to be 'up', thereby breaking the  $\mathbb{Z}_2$  symmetry. Then we iterated over the solutions in the sample, and for any solutions in which that variable's value was set to 'down', we multiplied the entire solution by minus one, thereby recovering the degenerate pair for that solution. This had the result that the expected state of each variable was no longer zero, and QPVAR was able to fix a non-zero number of variables.

There is freedom to choose any variable in the system to be fixed in the initial step. This motivated the idea that we should fix the variable that is the most correlated with other variables, which would allow us to fix those variables as well. As in Section 2.1, we found the connected components in the correlation graph. In the initial step, we chose any variable from the largest connected component, and fixed all of the variables that were connected to it to their respective values. Using this method, in the initial pre-fixing step, we were able to fix a number of variables equal to the size of the largest connected component, instead of just one.

In our experiments, we noted that the variance in the number of fixed variables was larger for the zero bias problems than for the non-zero bias problems. In particular, in the zero bias case, in some problems a large number of variables was fixed and for other problems a small number of variables was fixed. This motivated the question of how to decide, in advance, if a particular problem could benefit from the modified IQPVAR. Katzgraber et al. [23] found that problems for which the distribution of the spin glass order parameter is bimodal are easier in principle. We also observed a correlation between bimodality of the distribution of overlaps in the low-energy states and the number of fixed qubits. For this

reason, we suggest using the overlap of low-energy states in a small sample obtained from a quantum annealer as a heuristic for predicting whether IQPVAR will be useful for particular zero bias problems.

## 2.5 Methods for embedded problems

Most Ising (or QUBO) problems that result from real applications are not Chimera graph-structured, so they require a graph minor embedding problem to be solved first, yielding a non-trivial mapping from variables (logical variables) to qubits (physical qubits). In order to use IQPVAR to solve such embedded problems, we devised two different algorithms. The first, which we refer to as the *logical method*, involved obtaining a sample from the quantum annealer at each step and converting it to a sample in the logical space by fixing the broken chains using an energy minimization method (see Section 1.1). IQPVAR was then applied in the logical space, fixing logical variables.

An alternative method, which we refer to as the *physical method*, relied instead on embedding the problem and then fixing the states of the physical qubits at each step. In addition to fixing variables as in IQPVAR, we added two other fixing criteria. The first stipulated that if more than a fraction *majority\_length\_threshold* of the qubits in a chain were fixed to the same state, then the remaining qubits in that chain were also fixed to that state (since it was assumed that that is the state of the corresponding logical variable in the optimum). For the second criterion, a relaxed fixing threshold *chain\_fixing\_threshold* (typically 0.80–0.95) was imposed, and the qubits to be fixed were identified as in Section 2.1. If fixing these qubits in a chain of greater length than *absolute\_min\_length* (typically 3) resulted in all of the qubits being fixed, then they were actually fixed; otherwise, they were ignored.

## 3 Results

---

This section is organized as follows. In Section 3.1 we describe the procedure we used for benchmarking IQPVAR. In Section 3.2 we present results for Chimera graph-structured problems. In Section 3.3 we present results for a case study on embedded problems, namely optimal trading trajectory problems. In Section 3.4 we show the connected component sizes in the effective problems after applying our method, and in Section 3.5 we study the effect of applying a post-processing step to the quantum annealer’s sample prior to applying our method.

### 3.1 Benchmarking procedure

We defined problem sets in which the biases and couplers were chosen uniformly from a given set of integers, for example,  $U_2 \in \{-2, -1, 0, 1, 2\}$  (0 was excluded for the couplers but included for the biases) and a Sidon set  $S_{28} \in \{-28, -19, -13, -8, 8, 13, 19, 28\}$ . For each of the problem sets, we generated 100 random instances. To get a baseline result, we used the quantum annealer to solve each problem with 50,000 annealing cycles (using five random gauges with 10,000 cycles each). We then applied IQPVAR on each problem, with *num\_steps* = 4 and 2500 annealing cycles per step, finally using the quantum annealer to solve each reduced problem with 10,000 cycles (for a total of 20,000 cycles). This allowed us to compare solving a problem set with the quantum annealer alone versus solving the problem set by first applying IQPVAR.

As solution metrics, we used the number of occurrences of the best known energy value, the energy residual, and the success rate. The number of occurrences is the mean number of occurrences of the best known energy value for each sample for which at least one such solution was found. The energy residual is the difference in energy between the best solution found and the best known solution. The success rate is the percentage of problem instances for which the best known energy was found. An often used benchmarking practice is to measure the number of annealing cycles required for a 99% confidence of seeing the best known energy at least once. However, in practice, this measurement can be impractical for hard problems or high-precision problems, since it requires a very large number of annealing cycles, and measuring the number of occurrences provides an easier to measure proxy. We also calculated the number of problem instances for which the lowest-energy solution found by our method combined with the quantum annealer was at least as good as the lowest-energy solution found by the quantum annealer alone.

In order to determine the success rate of the quantum annealer, we solved each problem instance with a heuristic with a long time-out, which gave a high confidence of finding the optimum (but not a guarantee). For the Chimera graph-structured problems we used an implementation of the Hamze-de Freitas-Selby (HFS) algorithm [43, 44, 17], and for the embedded problems we used an implementation of multi-start tabu 1-opt search [14, 15] with 1000 repetitions

per problem instance. The IQPVAR fixing success rate was defined as the percentage of problem instances for which IQPVAR fixed variables to values which appeared in at least one best known solution. This was checked by finding the best known solution before and after fixing variables. If our method succeeded, the global optimum should be preserved; if it failed, then the energy value of the optimum would be higher after fixing the variables than before them.

### 3.2 Chimera graph-structured problems

Results for the number of fixed variables in each step for the Chimera graph-structured problem sets, for  $h = 0$  and  $h \neq 0$ , are presented in Table 3. For each step, we show a breakdown of the number of variables fixed, so that one can identify how many variables were fixed by QPVAR combined with variables fixed based on correlation and due to classical pre-processing (see Section 1.2), as well as the total number of variables fixed. Table 4 compares the success rate of the quantum annealer alone with the success rate of IQPVAR and the quantum annealer used together, and shows the IQPVAR fixing success rate.

	Set	0		1		2		3		4		Total
		m	c	m	c	m	c	m	c	m	c	
$h \neq 0$	$U_2$	–	10.6	517.4	168.4	69.4	48.1	52.3	32.9	45.1	19.5	963.7
	$U_5$	–	3.3	519.9	137.7	94.8	48.6	63.8	32.5	54.7	20.9	976.3
	$U_{10}$	–	2.5	502.9	124.9	101.2	49.6	68.5	29.3	64.8	22.9	966.7
	$S_{28}$	–	0.0	549.9	96.6	92.2	38.4	61.7	27.6	62.1	25.2	953.7
	$U_{100}$	–	1.6	507.6	122.1	84.9	43.8	54.8	30.8	52.4	21.8	919.8
$h = 0$	$U_2$	46.8	9.0	9.1	2.9	28.0	5.9	61.9	10.3	199.3	24.1	397.3
	$U_5$	43.4	6.0	13.0	3.4	24.2	4.8	50.8	7.0	131.1	12.8	296.6
	$U_{10}$	39.3	4.5	11.1	2.4	32.3	4.9	45.8	7.4	88.8	9.5	245.9
	$S_{28}$	47.1	4.7	16.7	3.1	24.8	3.5	35.9	4.9	76.3	7.8	224.7
	$U_{100}$	33.5	5.2	8.1	1.9	10.9	2.2	13.8	2.5	22.5	3.5	104.2

Table 3: Mean number of fixed variables for Chimera graph-structured problems, by problem set and by step. Step 0 is a pre-step in which only the classical pre-processing was run. For the zero bias ( $h = 0$ ) case, the pre-step also included fixing the largest connected component, as described in Section 2.4. For each step of the four steps of IQPVAR, the mean number of variables fixed by QPVAR and correlations (‘m’) and the classical pre-processing (‘c’) are shown. The final column shows the mean total number of fixed variables.

For the first step of our method, the parameters *fixing\_threshold* and *elite\_threshold* were tuned to 1.0 and 0.3 (respectively) by generating a problem set with 100 problem instances in  $U_5$ , and scanning different parameter values. The parameters for the remaining three steps were set based on trial and error, to 0.2, 0.15, and 0.1 for *elite\_threshold* and to 1.0 for all steps for *fixing\_threshold* (to preserve the optimum, as described above). The reason we decreased *fixing\_threshold* for the later steps is that as the method progresses, increasingly more qubits are fixed and the problems become simpler, so we become increasingly confident that the quantum annealer will solve the problem. The correlation graph parameters *correlation\_threshold* and *correlation\_elite\_threshold* were set to 1 and 0.4, respectively, for all steps. These might not be the optimal parameters—selecting optimal values is left for future work.

### 3.3 Embedded problems – case study

A challenge in studying the performance of our method on embedded problems is the large range of problems available, as well as the necessity to control the parameters, given the objective of studying the dependence of the method’s performance on the precision of the biases and couplers. In order to perform a case study, we obtained an adjacency matrix for a particular problem on which we generated random problem instances.

The problem we studied was an optimal trading trajectory (OTT) problem [41], which is a multi-period portfolio optimization problem. In the OTT problem, the objective is to maximize the future returns of discrete financial assets, given forecast returns and risk, and taking into account transaction costs. We chose the largest problem instance that could be embedded on the chip available to us, which had eight assets, five time steps, and a bit depth of two, giving a total of 80 logical variables, which required 872 qubits with a longest chain of 14.

We used the embedding and adjacency matrix of the original problem, but chose random biases and couplers. This preserved the structure of the original problem, while allowing us to control the precision of the couplers and biases. To

	Set	Quantum annealer			IQPVAR			Better	Fix
		Success	Residual	Freq	Success	Residual	Freq		
$h \neq 0$	$U_2$	92.00	0.16	285.13	99.00	0.02	3330.35	100.00	99.00
	$U_5$	21.00	3.34	6.43	88.00	0.38	759.90	100.00	92.00
	$U_{10}$	2.00	10.76	2.00	79.00	0.84	559.29	100.00	99.00
	$S_{28}$	0.00	29.84	0.00	56.00	4.96	484.11	100.00	95.00
	$U_{100}$	0.00	123.46	0.00	43.00	27.58	179.40	100.00	92.00
$h = 0$	$U_2$	57.00	1.10	19.28	72.00	0.62	1480.56	93.00	99.00
	$U_5$	1.00	7.92	1.00	20.00	5.44	55.10	82.00	95.00
	$U_{10}$	1.00	18.36	1.00	7.00	13.50	65.14	83.00	93.00
	$S_{28}$	0.00	42.68	0.00	6.00	35.14	312.83	74.00	87.00
	$U_{100}$	0.00	218.60	0.00	1.00	207.20	7.00	57.00	88.00

Table 4: Success metrics by problem set for the quantum annealer with and without IQPVAR, for Chimera graph-structured problems. ‘Success’ is the percentage of problems (out of 100) for which the best known energy value was observed at least once. ‘Residual’ is the mean energy difference between the lowest energy found and the best known solution. ‘Freq’ is the mean number of occurrences of the best known solution’s energy value in all solutions from the quantum annealer, averaged over the problem instances for which the best known solution’s energy value was seen at least once. The total number of annealing cycles was 50,000 for ‘Quantum annealer’ and 20,000 for ‘IQPVAR’ (2500 at each of the four steps of IQPVAR and 10,000 afterwards) for the finite field ( $h \neq 0$ ) problems and 50,000 for the zero bias ( $h = 0$ ) problems (2500 at each of the four steps of IQPVAR, and 40,000 afterwards). ‘Better’ shows the percentage of problems for which the lowest-energy solution found by IQPVAR was equal to or better than the lowest-energy solution found by the quantum annealer. The final column, ‘Fix’, shows the percentage of problems for which IQPVAR fixed variables such that their values occur in a best known solution of the original problem.

embed the problems, we first scaled them down to the range  $[-0.5, +0.5]$  and then used the *embed\_problem* function in D-Wave Systems’ SAPI 2.3.1 (such that the intra-chain couplers were set to  $-1$ ).

Results for the number of fixed variables in each step for the OTT problem sets are presented in Table 5, for both the logical and physical methods (see Section 2.5). For each step, we show a breakdown of the number of variables fixed, so that one can identify how many variables were fixed in total due to QPVAR, the correlations, and the classical pre-processing (see Section 1.2), as well as the mean total number of fixed variables. Table 6 compares the success rate of the quantum annealer used on its own with the success rate of IQPVAR and the quantum annealer used together, and shows IQPVAR’s fixing success rate.

	Set	1	2	3	4	c	Total
Logical	$U_2$	4.33	13.87	10.93	11.48	2.95	43.56
	$U_5$	4.22	11.45	10.69	12.52	3.17	42.05
	$U_{10}$	4.86	11.93	10.65	9.57	2.34	39.35
	$U_{100}$	2.85	8.28	8.02	7.15	3.94	30.24
	Physical	$U_2$	2.93	15.70	10.32	7.54	4.14
$U_5$		2.09	13.82	10.72	7.23	3.77	37.63
$U_{10}$		2.74	14.66	8.53	5.79	3.46	35.18
$U_{100}$		1.61	11.69	8.84	5.75	4.80	32.69

Table 5: Mean number of fixed variables for embedded optimal trading trajectory (OTT) problems, by problem set, by step, and by method (logical or physical methods). For each of the four steps of IQPVAR, the mean number of fixed variables is shown due to fixing variables first via classical pre-processing and then via QPVAR and correlations. The last step consisted of fixing variables using only classical pre-processing. The final column shows the total mean number of fixed variables.

The parameter *fixing\_threshold* was set to 1.0 for all steps in both methods. The parameter *elite\_threshold* was set to 0.2, 0.2, 0.15, and 0.1 for the four steps of both methods. The correlation graph parameters *correlation\_threshold* and *correlation\_elite\_threshold* were set to 0.95 and 0.2, respectively, for all steps of the logical method and 1 and 0.2 for all steps of the physical method. The parameters *chain\_fixing\_threshold*, *chain\_elite\_threshold*, and

	Set	Quantum annealer			IQPVAR			Better	Fix
		Success	Residual	Freq	Success	Residual	Freq		
Logical	$U_2$	60.00	0.67	6.40	79.00	0.38	2059.87	93.00	91.00
	$U_5$	69.00	0.33	17.43	77.00	0.24	2085.05	86.00	93.00
	$U_{10}$	71.00	0.24	13.48	77.00	0.23	2179.08	88.00	98.00
	$U_{100}$	57.00	0.37	5.67	67.00	0.28	933.79	83.00	96.00
Physical	$U_2$	60.00	0.67	6.40	61.00	0.57	1105.31	86.00	81.00
	$U_5$	69.00	0.33	17.43	66.00	0.32	1403.32	78.00	86.00
	$U_{10}$	71.00	0.24	13.48	68.00	0.32	1246.50	81.00	89.00
	$U_{100}$	57.00	0.37	5.67	65.00	0.34	730.80	84.00	90.00

Table 6: Success metrics by problem set for the quantum annealer with and without IQPVAR, for embedded OTT problems, for the logical and physical methods. See the caption of Table 4 for a definition of the metrics.

*majority\_length\_threshold* were set to 0.95, 0.2, and 0.51, respectively, for all steps of the physical method.

### 3.4 Problem decomposition

We have already shown that IQPVAR reduces the number of variables. However, an additional benefit of the method is that the resulting (effective) problems tend to be decomposed into smaller connected components, many of them small enough to be solved quickly, even exhaustively. In addition, the larger-sized components have a simpler graph structure, such as much smaller treewidth than the original problem, which dynamic programming could benefit from. Figure 4 shows a histogram of the sizes of the connected components for  $U_5$ , after four steps.




Figure 4: Frequency of connected component sizes in the resulting 100 problems after 4 steps of IQPVAR for the  $U_5$  problem set.

### 3.5 Effect of post-processing

The performance of the quantum annealer can be improved by post-processing. We hypothesized that post-processing should improve the results of our method. In particular, it would increase the method’s fixing success rate, as well as improve the success metrics compared with the quantum annealer alone (with post-processing). The results of running the quantum annealer with and without post-processing and with and without QPVAR are presented in Table 7.

Post-processing was done via the *optimization* flag of the function *solve\_ising* in D-Wave Systems’ SAPI 2.3.1, which performs a local search on each of the solutions returned from the quantum annealer [11].

Post-process	Quantum annealer			QPVAR			Fix
	Success	Residual	Freq	Success	Residual	Freq	
Without	0.00	123.46	0.00	1.00	96.16	1.00	99.00
With	60.00	9.24	49.65	87.00	4.20	458.54	100.00

Table 7: Success metric for the quantum annealer without post-processing, with post-processing, for the quantum annealer alone (‘Quantum annealer’) and with QPVAR, for problem set  $U_{100}$ . See the caption of Table 3 for a definition of the columns.

## 4 Discussion

For Chimera graph-structured problems, for the non-zero bias case ( $h \neq 0$ ), Table 3 shows that at least 84% of the variables were fixed on average in four steps of IQPVAR. Notably, classical pre-processing on its own was only able to fix at most 1% of the variables. Table 4 shows that for all problem sets, the results with IQPVAR have significantly better success metrics, that is, higher success rates, lower mean-energy residuals, and a higher mean number of occurrences of the best known energy value. This is compared to using the quantum annealer alone, despite using less annealing cycles. In addition, the lowest-energy solution found with IQPVAR was always of an equal or lower energy value than the lowest energy solution found by the quantum annealer alone. The percentage of problems for which variables were fixed incorrectly was low, less than 8%, and could be made lower by choosing a higher *elite\_threshold*.

The zero bias problems ( $h = 0$ ) are known to be harder, but IQPVAR still fixed 10%–36% of the variables. The results with IQPVAR show improved success metrics, also fixing variables incorrectly for up to 13% of the problems. For these problems, the mean number of variables fixed increased at every step, in contrast to the non-zero bias problems, where it decreased. The explanation for this is that as variables are fixed, they induce non-zero biases on the neighbouring qubits, so that after each additional step of the method, we observe more qubits with non-zero bias, and the resulting problem becomes more like the non-zero bias problems.

For the embedded OTT problems, Table 5 shows that 38%–54% of the variables were fixed by the logical method and 41%–51% were fixed by the physical method. In this case, the classical pre-processing was initially unable to fix any of the variables. Table 6 shows that for all problem sets, the results with IQPVAR have significantly improved success metrics. In addition, the method’s lowest-energy solution was as good or better than the lowest-energy solution for the quantum annealer alone for the majority of the problems.

We found that the logical method was better than the physical method, for these problem sets and for the specific parameters chosen, giving improved success metrics. In addition, the logical method fixed variables incorrectly less often than the physical method. We conjecture that this is due to the physical method being more sensitive to the quality of the embedding. In particular, the optimum of the embedded (i.e., physical) problem is not guaranteed to include a solution with unbroken chains. In principle, this could be guaranteed by increasing the strength of the intra-chain couplings, but the limited range of the biases and couplers in the quantum annealer limits our ability to do so. The physical method fixes variables with the objective of reaching the optimum of the embedded problem. However, if the embedded problem does not include an unbroken solution as one of the optima, the solution to which the physical method converges might not lead to the optimum of the logical problem.

Our results clearly show better success metrics for both Chimera graph-structured problems and embedded problems, even for high-precision problems that appear to be very difficult for the quantum annealer. Our intuition is that the quantum annealer is better at finding low-energy states than it is at finding optima. For difficult problems, for which the



expected number of occurrences of the optimum would be very low, looking only at the lowest-energy solution found does not make full use of the information in the sample obtained from the quantum annealer. It appears that the sample contains additional information about the structure of the low-energy states, information that allows our method to boost the performance of the quantum annealer.

We expect that the adjustable parameters in our method could be tuned based on the class of problems being optimized to yield the best results for that class of problems. In our study, we purposefully chose parameters such that they were not over-fit. For example, we did not choose different parameters for the high-precision problems compared to the low-precision problems, even though we observed that for the low-precision problems better results can be achieved by setting *elite\_threshold* lower (since these problems are easier for the quantum annealer). The objective here was to choose parameters that are a good all-around choice for general use of our method.

In Section 2.1, we described how to find correlation-based connected components. In principle, all of the variables in a connected component could be reduced to a single variable, by identification (i.e., by adding a strong coupling between them), such that they are either equal, or opposite, to that single variable, making full use of the information in the correlations. The issue with doing this for Chimera graph-structured problems is that the resulting problem may no longer be Chimera graph-structured (since new ferromagnetic couplers must be introduced), and as such might not fit on the chip (i.e., it would be impossible to find an embedding). In the case of embedded problems, the same issue arises, but in that case there is a more realistic possibility that an embedding could be found for the new problem. This is not guaranteed, however, and it would require an additional and possibly costly intermediate embedding step. Another possibility for utilizing our method is as a pre-processing step for a classical solver, as such solvers are typically not limited to a particular structure.

Each subsequent generation of quantum annealers has exhibited lower noise (i.e., ICE), therefore future quantum annealers are likely to show further reductions. The results in Section 3.5 give an indication that an improvement in this regard should also result in improved performance of our method. Table 7 shows that the results for the quantum annealer alone, as well as the results for the method, were both drastically improved, with increased success rate, lower mean residual, and a higher mean number of occurrences.

## 5 Conclusions and future work

---

We have presented a new method (IQPVAR) for iterative quantum processing of QUBO problems, which yields a significant reduction in the number of variables in a problem. Consequently, it gives a pronounced increase in success rates and number of occurrences of the best known energy value, when compared against a quantum annealer alone, even when using less annealing cycles. For this reason, at least for the problem sets studied herein, this evidence suggests strongly that this is a better way to use the quantum annealer to greater potential, for both Chimera graph-structured problems and embedded problems. More-specific tuning of parameters for specific problem classes could yield even better results, and future research might uncover improvements to the method presented in this work.

It has been shown that the quantum annealer exhibits quantum speedup compared with simulated annealing when global collective tunnelling occurs [12]. We suggest that our method would benefit even from local tunnelling, a hypothesis that we leave for future study. The performance of the method and fine-tuning of the parameters for other problem classes is another possible avenue for future study.



## Acknowledgements

---

The authors would like to thank Dominic Marchand, Pooya Ronagh, and Brad Woods for their insightful comments, Marko Bucyk for editing the manuscript, and Alex Selby for the use of his implementation of the Hamze-de Freitas-Selby (HFS) algorithm, available for public use on GitHub [43]. This work was supported by Mitacs.

## References

---

- [1] Demian A Battaglia, Giuseppe E Santoro, and Erio Tosatti. Optimization by quantum annealing: Lessons from hard satisfiability problems. *Physical Review E*, 71(6):066707, 2005.
- [2] Sergio Boixo, Troels F Rønnow, Sergei V Isakov, Zihui Wang, David Wecker, Daniel A Lidar, John M Martinis, and Matthias Troyer. Evidence for quantum annealing with more than one hundred qubits. *Nature Physics*, 10(3):218–224, 2014.
- [3] Sergio Boixo, Vadim N Smelyanskiy, Alireza Shabani, Sergei V Isakov, Mark Dykman, Vasil S Denchev, Mohammad H Amin, Anatoly Yu Smirnov, Masoud Mohseni, and Hartmut Neven. Computational multiqubit tunnelling in programmable quantum annealers. *Nature communications*, 7, 2016.
- [4] Endre Boros, Peter L Hammer, and Gabriel Tavares. Preprocessing of unconstrained quadratic binary optimization. 2006.
- [5] Endre Boros, Peter L Hammer, and Gabriel Tavares. Preprocessing of unconstrained quadratic binary optimization preprocessing of unconstrained quadratic binary optimization. 2006.
- [6] Paul I Bunyk, Emile M Hoskinson, Mark W Johnson, Elena Tolkacheva, Fabio Altomare, Andrew J Berkley, Roy Harris, Jeremy P Hilton, Trevor Lanting, Anthony J Przybysz, et al. Architectural considerations in the design of a superconducting quantum annealing processor. *Applied Superconductivity, IEEE Transactions on*, 24(4):1–10, 2014.
- [7] Pierre Chardaire, Jean Luc Lutton, and Alain Sutter. Thermostatistical persistency: A powerful improving concept for simulated annealing algorithms. *European Journal of Operational Research*, 86(3):565–579, 1995.
- [8] Vicky Choi. Minor-embedding in adiabatic quantum computation: I. the parameter setting problem. *Quantum Information Processing*, 7(5):193–209, 2008.
- [9] Vicky Choi. Minor-embedding in adiabatic quantum computation: II. minor-universal graph design. *Quantum Information Processing*, 10(3):343–353, 2011.
- [10] D-Wave Systems. Private communication, 2016.
- [11] D-Wave Systems. SAPI 2.3.1 documentation, 2016.
- [12] Vasil S Denchev, Sergio Boixo, Sergei V Isakov, Nan Ding, Ryan Babbush, Vadim Smelyanskiy, John Martinis, and Hartmut Neven. What is the computational value of finite range tunneling? *arXiv preprint arXiv:1512.02206*, 2015.
- [13] AB Finnila, MA Gomez, C Sebenik, C Stenson, and JD Doll. Quantum annealing: a new method for minimizing multidimensional functions. *Chemical physics letters*, 219(5):343–348, 1994.

- [14] Fred Glover. Tabu search-part I. *ORSA Journal on computing*, 1(3):190–206, 1989.
- [15] Fred Glover. Tabu search-part II. *ORSA Journal on computing*, 2(1):4–32, 1990.
- [16] Peter L Hammer, Pierre Hansen, and Bruno Simeone. Roof duality, complementation and persistency in quadratic 0–1 optimization. *Mathematical programming*, 28(2):121–155, 1984.
- [17] Firas Hamze and Nando de Freitas. From fields to trees. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 243–250. AUAI Press, 2004.
- [18] Itay Hen, Joshua Job, Tameem Albash, Troels F Rønnow, Matthias Troyer, and Daniel A Lidar. Probing for quantum speedup in spin-glass problems with planted solutions. *Physical Review A*, 92(4):042325, 2015.
- [19] He Jiang and Jifeng Xuan. *Backbone guided local search for the weighted maximum satisfiability problem*. INTECH Open Access Publisher, 2009.
- [20] MW Johnson, MHS Amin, S Gildert, T Lanting, F Hamze, N Dickson, R Harris, AJ Berkley, J Johansson, P Bunyk, et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, 2011.
- [21] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse Ising model. *Physical Review E*, 58(5):5355, 1998.
- [22] Helmut G Katzgraber, Firas Hamze, and Ruben S Andrist. Glassy Chimeras could be blind to quantum speedup: Designing better benchmarks for quantum annealing machines. *Physical Review X*, 4(2):021008, 2014.
- [23] Helmut G Katzgraber, Firas Hamze, Zheng Zhu, Andrew J Ochoa, and H Munoz-Bauza. Seeking quantum speedup through spin glasses: The good, the bad, and the ugly. *Physical Review X*, 5(3):031026, 2015.
- [24] Andrew D King. Performance of a quantum annealer on range-limited constraint satisfaction problems. *arXiv preprint arXiv:1502.02098*, 2015.
- [25] James King, Sheir Yarkoni, Mayssam M. Nevisi, Jeremy P. Hilton, and Catherine C. McGeoch. Benchmarking a quantum annealing processor with the time-to-target metric. *arXiv preprint arXiv:1508.05087*, 2015.
- [26] T Lanting, AJ Przybysz, A Yu Smirnov, FM Spedalieri, MH Amin, AJ Berkley, R Harris, F Altomare, S Boixo, P the, et al. Entanglement in a quantum annealing processor. *Physical Review X*, 4(2):021041, 2014.
- [27] Wolfgang Lechner, Philipp Hauke, and Peter Zoller. A quantum annealing architecture with all-to-all connectivity from local interactions. *Science advances*, 1(9):e1500838, 2015.
- [28] Andrew Lucas. Ising formulations of many NP problems. *Frontiers in Physics*, 2(5), 2014.
- [29] Salvatore Mandrà, Zheng Zhu, Wenlong Wang, Alejandro Perdomo-Ortiz, and Helmut G Katzgraber. Strengths and weaknesses of weak-strong cluster problems: A detailed overview of state-of-the-art classical heuristics vs quantum approaches. *arXiv preprint arXiv:1604.01746*, 2016.
- [30] Catherine C McGeoch and Cong Wang. Experimental evaluation of an adiabatic quantum system for combinatorial optimization. In *Proceedings of the ACM International Conference on Computing Frontiers*, page 23. ACM, 2013.
- [31] Anurag Mishra, Tameem Albash, and Daniel A Lidar. Performance of two different quantum annealing correction codes. *Quantum Information Processing*, 15(2):609–636, 2016.
- [32] Rémi Monasson, Riccardo Zecchina, Scott Kirkpatrick, Bart Selman, and Lidror Troyansky. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400(6740):133–137, 1999.
- [33] Fernando Pastawski and John Preskill. Error correction for encoded quantum annealing. *Phys. Rev. A*, 93:052325, May 2016.
- [34] Alejandro Perdomo-Ortiz, Joseph Fluegemann, Rupak Biswas, and Vadim N Smelyanskiy. A performance estimator

for quantum annealers: Gauge selection and parameter setting. *arXiv preprint arXiv:1503.01083*, 2015.

- [35] Alejandro Perdomo-Ortiz, Bryan O’Gorman, Joseph Fluegemann, Rupak Biswas, and Vadim N Smelyanskiy. Determination and correction of persistent biases in quantum annealers. *Scientific Reports*, 6:18628, 2016.
- [36] Kristen L Pudenz, Tameem Albash, and Daniel A Lidar. Error-corrected quantum annealing with hundreds of qubits. *Nature communications*, 5, 2014.
- [37] Kristen L Pudenz, Tameem Albash, and Daniel A Lidar. Quantum annealing correction for random Ising problems. *Physical Review A*, 91(4):042302, 2015.
- [38] P Ray, BK Chakrabarti, and Arunava Chakrabarti. Sherrington-Kirkpatrick model in a transverse field: Absence of replica symmetry breaking due to quantum fluctuations. *Physical Review B*, 39(16):11828, 1989.
- [39] F Romá and S Risau-Gusman. Backbone structure of the edwards-anderson spin-glass model. *Physical Review E*, 88(4):042105, 2013.
- [40] Troels F. Rønnow, Zihui Wang, Joshua Job, Sergio Boixo, Sergei V. Isakov, David Wecker, John M. Martinis, Daniel A. Lidar, and Matthias Troyer. Defining and detecting quantum speedup. *Science*, 345(6195):420–424, 2014.
- [41] Gili Rosenberg, Poya Haghnegahdar, Phil Goddard, Peter Carr, Kesheng Wu, and Marcos López de Prado. Solving the optimal trading trajectory problem using a quantum annealer. In *Proceedings of the 8th Workshop on High Performance Computational Finance*, page 7. ACM, 2015.
- [42] Giuseppe E Santoro, Roman Martoňák, Erio Tosatti, and Roberto Car. Theory of quantum annealing of an Ising spin glass. *Science*, 295(5564):2427–2430, 2002.
- [43] Alex Selby. QUBO-Chimera. <https://github.com/alex1770/QUBO-Chimera>, 2013.
- [44] Alex Selby. Efficient subgraph-based sampling of Ising-type models with frustration. *arXiv preprint arXiv:1409.3934*, 2014.
- [45] Gabriel Tavares. *New algorithms for Quadratic Unconstrained Binary Optimization (QUBO) with applications in engineering and social sciences*. PhD thesis, Rutgers University, Graduate School – New Brunswick, 2008.
- [46] Walter Vinci, Tameem Albash, Gerardo Paz-Silva, Itay Hen, and Daniel A Lidar. Quantum annealing correction with minor embedding. *Physical Review A*, 92(4):042310, 2015.
- [47] Yang Wang, Zhipeng Lü, Fred Glover, and Jin-Kao Hao. Effective variable fixing and scoring strategies for binary quadratic programming. In *Evolutionary Computation in Combinatorial Optimization*, pages 72–83. Springer, 2011.
- [48] Yang Wang, Zhipeng Lü, Fred Glover, and Jin-Kao Hao. Backbone guided tabu search for solving the UBQP problem. *Journal of Heuristics*, pages 1–17, 2013.
- [49] Arman Zaribafiyani, Dominic J J Marchand, and Seyed Saeed Changiz Rezaei. Systematic and deterministic graph-minor embedding for Cartesian products of graphs. *arXiv preprint arXiv:1602.04274*, 2016.
- [50] Weixiong Zhang. Configuration landscape analysis and backbone guided local search.: Part i: Satisfiability and maximum satisfiability. *Artificial Intelligence*, 158(1):1–26, 2004.
- [51] Iliia Zintchenko, Matthew B Hastings, and Matthias Troyer. From local to global ground states in Ising spin glasses. *Physical Review B*, 91(2):024201, 2015.

Academic Paper – <https://arxiv.org/abs/1606.07797>

Visit [1QBit.com](http://1QBit.com) for more information.

**1QBit**