Post-Quantum Error-Correction for Quantum Annealers

Ramin Ayanzadeh[1,*], **John Dorband**[1], **Milton Halem**[1], **and Tim Finin**[1]

[1]Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County, Baltimore, MD 21250, United States
[*]ayanzadeh@umbc.edu

ABSTRACT

We present a general post-quantum error-correcting technique for quantum annealing, called multi-qubit correction (MQC), that views the evolution in an open-system as a Gibs sampler and reduces a set of (first) excited states to a new synthetic state with lower energy value. After sampling from the ground state of a given (Ising) Hamiltonian, MQC compares pairs of excited states to recognize virtual tunnels—i.e., a group of qubits that changing their states simultaneously can result in a new state with lower energy—and successively converges to the ground state. Experimental results using D-Wave 2000Q quantum annealers demonstrate that MQC finds samples with notably lower energy values and also improves the reproducibility of results, compared to recent hardware/software advances in the realm of quantum annealing such as reverse quantum annealing, increased inter-sample delay, and classical pre/post-processing methods.

1 Introduction

Quantum annealing is a meta-heuristic for addressing discrete (or combinatorial) optimization problems that are intractable in the realm of classical computing. While simulated annealing (a.k.a. thermal or classical annealing) uses adjustable thermal fluctuations to jump over the energy mountains, quantum annealing applies adjustable quantum fluctuations for tunneling through the (narrow-enough) energy barriers[1–6]. Quantum annealers are a special case of the adiabatic quantum computers (i.e., stoquastic open-system) that provide a hardware implementation for finding the ground state (or minimum energy configuration) of (Ising) Hamiltonians[7,8]. To solve a problem using the quantum annealers, therefore, we must define a Hamiltonian whose ground state represents the (optimum) solution of the original problem of interest[9–11].

We can form an Ising Hamiltonian whose ground state represents the optimum solution of any given problem of interest—which can be nontrivial in many real-world applications[11–16]. In practice, nevertheless, executing the corresponding quantum machine instruction (QMI) on a physical quantum annealer does not guarantee to achieve the global optimum[7,8,17]. As an example, besides thermal noise and diabatic transitions[7], various control error sources—including, but not limited to, sparse connectivity between qubits[18,19], confined annealing schedule[6], coefficients' range and precision limitations[20,21], noise and decoherence[22–26]—lower the quality of results (i.e., the energy value of the drawn samples is higher than the energy value of the ground state)[27].

Modifying some aspects of the Hamiltonian—namely adapting (and better selecting) initial and final Hamiltonians; optimizing the schedule/path function; adding a catalyst Hamiltonian (i.e., a Hamiltonian that is present only in intermediate time); and adding non-stoquastic term to the Hamiltonian—can circumvent certain drawbacks of the adiabatic quantum computers[7]. Nevertheless, the majority of these techniques are mainly for closed-systems or current generations of the physical quantum annealers cannot (fully) accommodate them. Acknowledging that adiabatic quantum computing has some inherent resistance to noise and decoherence, we need error correction (and mitigation) mechanisms for ensuring the scalability of adiabatic quantum computers (like other quantum information processing models)[28–31]. In spite of several error correction proposals for adiabatic quantum computing and quantum annealing[19,20,22,24,28–38], an accuracy-threshold theorem for adiabatic quantum computing (unlike its gate model counterpart[39]) remains elusive[24,31]. Besides, most error correction schemes (e.g., nested quantum annealing correction method[31,38]) utilize multiple physical qubits for coding every qubit that notably reduces the capacity of current quantum annealers. We view quantum annealers as a Gibs sampler—that allows diabatic transitions in an open system—and present a novel post-quantum error correction/mitigation technique for quantum annealers.

2 Method

This paper presents postprocessing methods for mitigating errors in quantum annealers, so-called post-quantum error correction, that notably improves the performance of quantum annealers in terms of reproducibility of results and finding solutions with

lower energy value. We used a D-Wave 2000Q quantum annealer by D-Wave Systems Inc. (located at Burnaby, British Colombia, Canada) to evaluate the performance of the proposed techniques. Generating random Hamiltonians is a common practice for benchmarking quantum annealers[4,17,20,40]. Hence, we generated three different types of Ising Hamiltonians as follows:

- Binary—random Ising Hamiltonians whose linear and quadratic coefficients (denoted by $\mathbf{h}$ and J, respectively) were randomly drawn from $\{-1,+1\}$, based on a Bernoulli distribution with equal probabilities for -1 and +1;
- Uniform—random Ising Hamiltonians whose linear and quadratic coefficients (denoted by $\mathbf{h}$ and J, respectively) are (double-precision) uniform random numbers in $[-1,+1]$;
- Normal—random Ising Hamiltonians whose linear and quadratic coefficients (denoted by $\mathbf{h}$ and J, respectively) are (double-precision) Normal random numbers that follow the standard Gaussian distribution, i.e., average and standard deviation are 0 and 1, respectively.

For each problem type, we generated 50 instances (random problems). Besides, we adopted the finite-range Ising model, a.k.a. EA model (Edward—Anderson)[4], to generate benchmark problems. More specifically, all randomly generated benchmark problems were compatible with the current working graph of the D-Wave quantum processor—in Chimera topology, every qubit is connected to at most six other qubits. Therefore, until the next maintenance that can change the working graph of the quantum annealer, one can directly execute them without embedding problems to a target graph. The annealing time for all experiments was 20 microseconds.

Since executing a quantum machine instruction (QMI) on a physical quantum annealer is not guaranteed to achieve the ground state of the corresponding Ising Hamiltonian, even if we request for many samples/reads, several studies have proposed software and hardware advancements to improve the performance of the quantum annealers. As an example, recent studies have revealed that using spin-reversal transforms (a.k.a. gauge transforms)—i.e., flipping qubits randomly without altering the ground state of the original Ising Hamiltonian—can reduce analog errors of the quantum annealers[41]. Similarly, applying classical postprocessing heuristics on raw samples (attained by the quantum annealers) can result in samples with lower energy values[42]. Furthermore, when we submit a problem to a D-Wave quantum annealer, it is a common practice to request for several samples/reads (e.g., up to 10,000 per QMI on the current D-Wave quantum processors). For every read (or measurement), the D-Wave QPU initializes all qubits and repeats the annealing process. When we repeat the annealing process, therefore, successive measurements are correlated to each other due to limited preparation time. From another perspective, successive measurements generally form clusters of samples (i.e., groups of identical states). Consequently, increasing the preparation time can reduce the inter-sample correlations. Lastly, reverse annealing is a recent variant of quantum annealing where we start from a known (classical) state—the annealing process exploits the neighborhood of an initial (eigen) state rather than exploring the entire landscape of the problem Hamiltonian—through increasing and then decreasing the amplitude of the transverse field[43–45] In this study, we used the following arrangements for evaluating the performance of the proposed post-quantum error correction/mitigation methods:

- QA^1—raw samples attained by a D-Wave quantum annealer;
- QA^2—applies five spin-reversal transforms on QA^1;
- QA^3—puts a longer delay between successive reads/samplings to reduce the sample-to-sample correlation, albeit longer run-time;
- QA^4—performs the optimization postprocessing, available from the D-Wave's Ocean SDK, to all raw samples (QA^1);
- QA^5— performs the sampling postprocessing, available from the D-Wave's Ocean SDK, to all raw samples (QA^1);
- QA^6—applies five spin-reversal transforms, puts a longer delay between successive reads, and performs the optimization postprocessing to raw samples (attained by QA^1);
- RA—executes reverse annealing, starting from the best sample of a standard (or forward) quantum annealing (i.e., the result of QA^1).

3 Results

Quantum annealers, like the quantum processing units (QPU) by D-Wave Systems, are single-instruction (quantum) computing machines that can only sample from the ground state of the following problem Hamiltonian (denoted by H_p):

$$H_p := E_{\text{Ising}}(\mathbf{z}) = \sum_{i=1}^{N} \mathbf{h}_i \mathbf{z}_i + \sum_{i=1}^{N} \sum_{j=i+1}^{N} J_{ij} \mathbf{z}_i \mathbf{z}_j, \tag{1}$$

where N denotes the number of quantum bits (qubits), spin variables $\mathbf{z} \in \{-1,+1\}^N$, and $\mathbf{h}$ and J represent linear and quadratic coefficients, respectively[10,11,46]. Quantum annealers can efficiently recognize the region of the ground state(s) of the given Hamiltonians; however, they generally fail to get to the global minimum, regardless of how close they are to the ground state. In other words, unlike classical annealing that always converges to a local (or sometimes the global) optimum, quantum annealing generally yields (first) excited state(s) that are not necessarily a local optimum[8,17,47]. Hence, we can expect that applying optimization heuristics and meta-heuristics on samples, attained by a quantum annealer, to result in new (or synthetic) sample(s) with lower energy value, specifically on systems with glassy landscapes[4]. In this section, we start with a local optimization heuristic, called single-qubit correction (SQC), and then extend it to introduce multi-qubit correction (MQC) scheme for mitigating errors in quantum annealers[47].

3.1 Single-Qubit Correction

Measuring qubits after the annealing process results in an eigenstate that is not necessarily the ground state (or even a local optimum) of the problem Hamiltonian[4]. Also, evolution in an open-system is not guaranteed to result in an eigenstate of the problem Hamiltonian[7]. We start by adopting the hill-climbing[48,49] to present a post-quantum error mitigation approach for quantum annealers, called single-qubit correction (SQC). SQC has a zero-temperature simulated annealing scheme that can relax a (raw) sample to a new/synthetic sample with lower energy value. In each iteration of SQC, we toggle the value of every qubit (individually) and keep all changes that result in a state with a lower energy value. Algorithm 1 shows how SQC exploits the neighborhood of an excited state, denoted by $\mathbf{z}$.

Input: $\mathbf{z}, \mathbf{h}$ and J
Output: $\mathbf{z}$
$N \leftarrow |\mathbf{z}|$
$lowestEnergy \leftarrow E_{\text{Ising}}(\mathbf{z}, \mathbf{h}, J)$
$terminate \leftarrow$ True
do
  for $i \leftarrow 1$ **to** N **do**
    $\mathbf{z}_i \leftarrow -\mathbf{z}_i$ `// Flip bit and check its influence`
    if $lowestEnergy > E_{Ising}(\mathbf{z}, \mathbf{h}, J)$ **then**
      $lowestEnergy \leftarrow E_{\text{Ising}}(\mathbf{z}, \mathbf{h}, J)$
      $terminate \leftarrow$ False
    else
      $\mathbf{z}_i \leftarrow -\mathbf{z}_i$ `// Flip bit back since it did not result in a sample with lower energy`
    end
  end
while not *terminate*
return $\mathbf{z}$

Algorithm 1: Single-qubit correction (SQC) heuristic for exploiting the neighborhood of an excited states to find a sample with lower energy value.

From an optimization point of view, SQC is very likely to result in a meta-stable state (i.e., SQC is sensitive to the input sample)—when the input sample is not a local optimum—which is trivial for systems with the glassy landscape. To solve a problem on a quantum annealer, we generally draw many samples (e.g., up to 10,000 samples/reads per QMI on a D-Wave quantum annealer). In practice, therefore, SQC explores a broader area, i.e., the neighborhood of all (first) excited states. It is worth highlighting that we do not propose SQC as a post-quantum error correction method for quantum annealers since more efficient techniques (namely, the simulated annealing) can outperform SQC in terms of finding samples with lower energy value. Indeed we will extend SQC to introduce a novel post-quantum error correction (or mitigation) approach that can notably improve the quality and reproducibility of results attained by the quantum annealers. Figure 1 illustrates that applying SQC to (raw) samples (drawn by the D-Wave quantum annealers) improves the quality of results for all benchmark problems.

3.2 Multi-Qubit Correction

SQC is an optimization heuristic that: (1) neglects the interactions between spins of the given problem Hamiltonian; and (2) entirely depends on one excited state as its initial state. Instead of processing one qubit at a time on samples individually, we introduce multi-qubit correction (MQC) method that treats groups of qubits as tunnels (or units) and compares pairs of samples

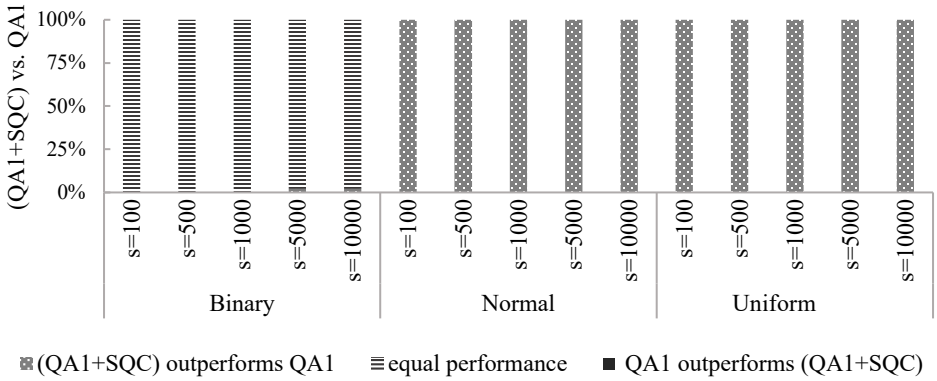

Figure 1. Performance comparison between quantum annealing (QA^1) and applying SQC to raw samples attained by quantum annealers ($QA^1 + SQC$).

to find these tunnels. In this study, the term "tunnel" is analogous to the concept of quantum tunnels[50] and refers to a group of qubits that a quantum annealer simultaneously flips their values to change an excited state to a sample with lower energy value.

Let $\mathbf{h}$ and J denote linear and quadratic coefficients of a problem Hamiltonian that we aim to find its ground state, respectively. Also, let $\mathbf{z}^1$ and $\mathbf{z}^2$ be two (excited) states, with N spin variables, attained by a quantum annealer. We define two sets of qubits as

$$S = \{\mathbf{z_i} | \mathbf{z}_i^1 = \mathbf{z}_i^2\} \tag{2}$$

and

$$D = \{\mathbf{z_i} | \mathbf{z}_i^1 \neq \mathbf{z}_i^2\} \tag{3}$$

where $i = 1, 2, \ldots, N$. The set D is a tunnel that represents the transformation of sample $\mathbf{z}^1$ to sample $\mathbf{z}^2$, and vice versa. From a problem-solving viewpoint, finding D is rather useless. However, we can use D to find sub-tunnels that may reduce $\mathbf{z}^1$ (or $\mathbf{z}^2$) into a new sample with lower energy value. A sub-tunnel of D, denoted by T, is a subset of D where

$$J_{ij} = \begin{cases} \mathbb{R}_{\neq 0} & \mathbf{z}_i, \mathbf{z}_j \in T \\ 0 & \mathbf{z}_i \text{ or } \mathbf{z}_j \notin T. \end{cases} \tag{4}$$

In this sense, T is the closure of a set qubits connected transitively to each other, but not connected to other qubits in D. Hence, we can represent D as a partition of sub-tunnels as

$$D = \bigcup_k T^k$$

where

$$T^i \cap T^j = \{\}, \qquad \forall i, j.$$

We define the influence (or energy contribution) of T^k to $\mathbf{z}^1$ as follows:

$$I_{\mathbf{z}^1}^k = \sum_{i \in T^k} \mathbf{h}_i \mathbf{z}_i^1 + \sum_{i \in T^k} \sum_{j \in S} J_{ij} \mathbf{z}_i^1 \mathbf{z}_j^1. \tag{5}$$

Note that we have omitted the term

$$\sum_{i,j \in T^k} J_{ij} \mathbf{z}_i^1 \mathbf{z}_j^1$$

from Eq. (5) since flipping values of all qubits in T^k does not effect $I_{\mathbf{z}^1}^k$. Finally, we reduce $\mathbf{z}^1$ to a new sample, denoted by $\mathbf{z}^*$, via flipping the qubit values of all sub-tunnels that have a positive influence (or energy contribution) value. Note that $I_{\mathbf{z}^1}^k = -I_{\mathbf{z}^2}^k$; thus, applying the abovementioned process on $\mathbf{z}^2$ will result in a new sample that is identical to $\mathbf{z}^*$. Algorithm 2 shows how we

Function `Reduce` ($\mathbf{z}^1$, $\mathbf{z}^2$, $\mathbf{h}$, J) **:**

$\mathbf{z}^* \leftarrow \mathbf{z}^1$
$N \leftarrow |\mathbf{z}^*|$
$S \leftarrow \{\}$
$D \leftarrow \{\}$
for $i \leftarrow 1$ **to** N **do**
  if $\mathbf{z}^1_i = \mathbf{z}^2_i$ **then**
    $S \leftarrow S \cup \{i\}$ // Identical bits between $\mathbf{Z}^1$ and $\mathbf{Z}^2$
  else
    $D \leftarrow D \cup \{i\}$ // Different bits between $\mathbf{Z}^1$ and $\mathbf{Z}^2$
  end
end
$Adj \leftarrow []$ // Adjacency of bits in D based on J
for $i \in D$ **do**
  $Adj_i = \{\}$
end
for $i, j \in J$ **do**
  if $J_{ij} \neq 0$ & $i \in D$ & $j \in D$ **then**
    $Adj_i \leftarrow Adj_i \cup \{j\}$
    $Adj_j \leftarrow Adj_j \cup \{i\}$
  end
end
$T \leftarrow$ `ConnectedComponents` (Adj) // Sub-tunnels (groups of isolated bits in D)
for $k \leftarrow 1$ **to** $|T|$ **do**
  $I^k_{\mathbf{z}^1} = \sum_{i \in T^k} \mathbf{h}_i \mathbf{z}^1_i + \sum_{i \in T^k} \sum_{j \in S} J_{ij} \mathbf{z}^1_i \mathbf{z}^1_j$ // Influence of sub-tunnels on E_{Ising}
  if $I^k_{\mathbf{z}^1} > 0$ **then**
    for $l \in T^k$ **do**
      $\mathbf{z}^*_l \leftarrow -\mathbf{z}^*_l$ // Flip all bits of sub-tunnels with positive influence
    end
  end
end
return $\mathbf{z}^*$

Algorithm 2: Reducing two input samples $\mathbf{z}^1$ and $\mathbf{z}2$ to a new sample ($\mathbf{z}^*$) with lower energy value based on virtual tunnels.

reduce two input samples, denoted by $\mathbf{z}^1$ and $\mathbf{z}^2$, to a new sample (denoted by $\mathbf{z}^*$) whose energy value is guaranteed to be less than or equal to energy values of $\mathbf{z}^1$ or $\mathbf{z}^2$.

The presented reduce procedure in this paper is analogous to the crossover operation in evolutionary algorithms that acts on two potential solutions, known as (parent) chromosomes, and yields (two) new solution(s), known as offspring(s)[48,51]. In fact, the Reduce procedure, shown in Algorithm 2, is the extended version of SQC method, shown in Algorithm 1, that acts on a group of qubits simultaneously rather than flipping the values of qubits individually. Besides, the Reduce procedure acts on two (excited) states and is less sensitive to a single initial point, compared to SQC that entirely depends on one (excited) state. When we employ physical quantum annealers, we generally request for many samples/reads, i.e., repeating the annealing process (with different initial eigenstates), to improve the probability of achieving the ground state of the given Hamiltonian. Algorithm 3 illustrates the multi-qubit correction (MQC) method that receives a sample set as input, and tries to iteratively reduce it to a (new) sample with lower energy value. For a sample set with n samples/reads, MQC takes $\log_2(n)$ steps and in each iteration, the size of the sample set is divided by two. Figure 2 shows that MQC outperforms SQC (i.e., finds samples with lower energy value) when we apply them on raw samples attained by the D-Wave quantum annealers.

Figure 3 shows that applying MQC on raw samples, attained by the D-Wave quantum processors in sampling from the ground state of benchmark Ising models, outperforms recent software/hardware advances in the realm of quantum annealing (i.e., MQC finds samples with lower energy value)—including, but not limited to, spin-reversal transforms and increased inter-sample delay between successive measurements. In the same way, Fig. 4 demonstrates that MQC outperforms reverse quantum annealing[43–45] where we initialize a quantum annealing process with the result of a standard (or forward) quantum annealing.

Input: $Z = \left[\mathbf{z}^1, \mathbf{z}^2, \ldots, \mathbf{z}^n\right]$, $\mathbf{h}$ and J
Output: $\mathbf{z}^*$
while $|Z| > 1$ **do**
  $\hat{Z} \leftarrow Z$
  $Z \leftarrow \{\}$
  while $|\hat{Z}| > 1$ **do**
    $\mathbf{z}^A \leftarrow \hat{Z}$.`Pop()`
    $\mathbf{z}^B \leftarrow \hat{Z}$.`Pop()`
    $\mathbf{z}^{AB} \leftarrow$ `Reduce` $(\mathbf{z}^A, \mathbf{z}^B, \mathbf{h}, J)$
    Z.`Append` $(\mathbf{z}^{AB})$
  end
  if $|\hat{Z}| > 0$ **then**
    Z.`Append` $(\hat{Z}$.`Pop()`$)$
  end
end
$\mathbf{z}^* \leftarrow Z$.`Pop()`

Algorithm 3: Multi-qubit correction (MQC) method for reducing a set of samples to a new sample whose energy value is less than all input samples or (in worst case) is equal to the lowest energy value of input samples.

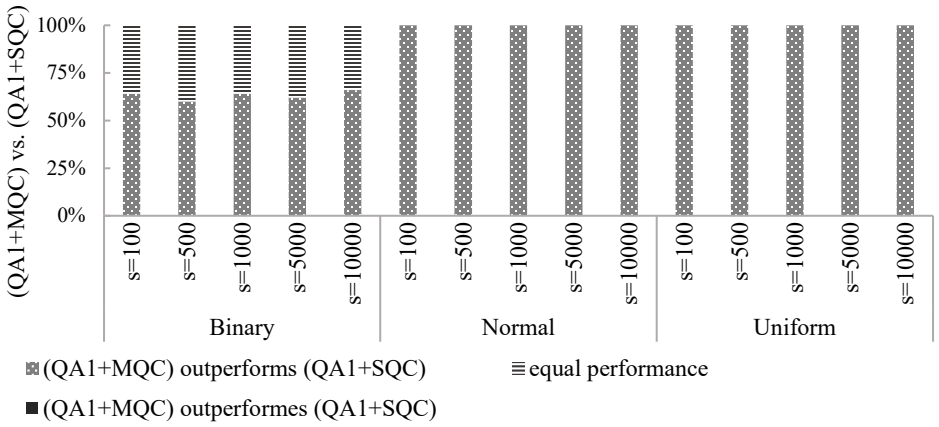

Figure 2. Performance comparison between applying MQC and SQC to raw samples attained by quantum annealers, denoted by $\mathrm{QA}^1 + MQC$ and $\mathrm{QA}^1 + SQC$, respectively.

From an application perspective, problem-solving with a physical quantum annealer has two drawbacks: (1) quantum annealers can yield excited states (rather than the ground state of the given Hamiltonian); and (2) results/samples attained by the physical quantum annealers are not reproducible over time. According to Anderson localization[52], as an example, the energy gap between the ground and first excited states is shrunk close to the end of the annealing. The landscape of glassy Hamiltonians generally includes many excited states and a physical quantum annealer is likely to relax to one of these excited states. In conclusion, problem-solving with quantum annealers results in a distribution of (potential) ground states, and the variance of the corresponding energy values are large enough to lessen the reproducibility of results. We repeated the aforementioned methods 50 times and Table 1 includes the variance of the energy values of these repeats. These results reveal that applying post-quantum error correction/mitigation techniques on (raw) samples, attained by the D-Wave quantum processors, can notably improve the reproducibility of results.

3.3 Randomized MQC

MQC is a postprocessing meta-heuristic that iteratively reduce a set of samples to a smaller sample set whose energy values are lower than the previous iteration(s). The performance of MQC, therefore, depends on samples that we obtain from the physical quantum annealers. For example, when we repeat the annealing process on a D-Wave quantum annealer, successive measurements are correlated to each other due to limited preparation time. Moreover, successive measurements generally form clusters of samples (i.e., groups of identical states). Since applying the Reduce procedure on identical input samples yields the same sample, early iterations of applying MQC on raw samples (attained by the physical quantum annealers) can

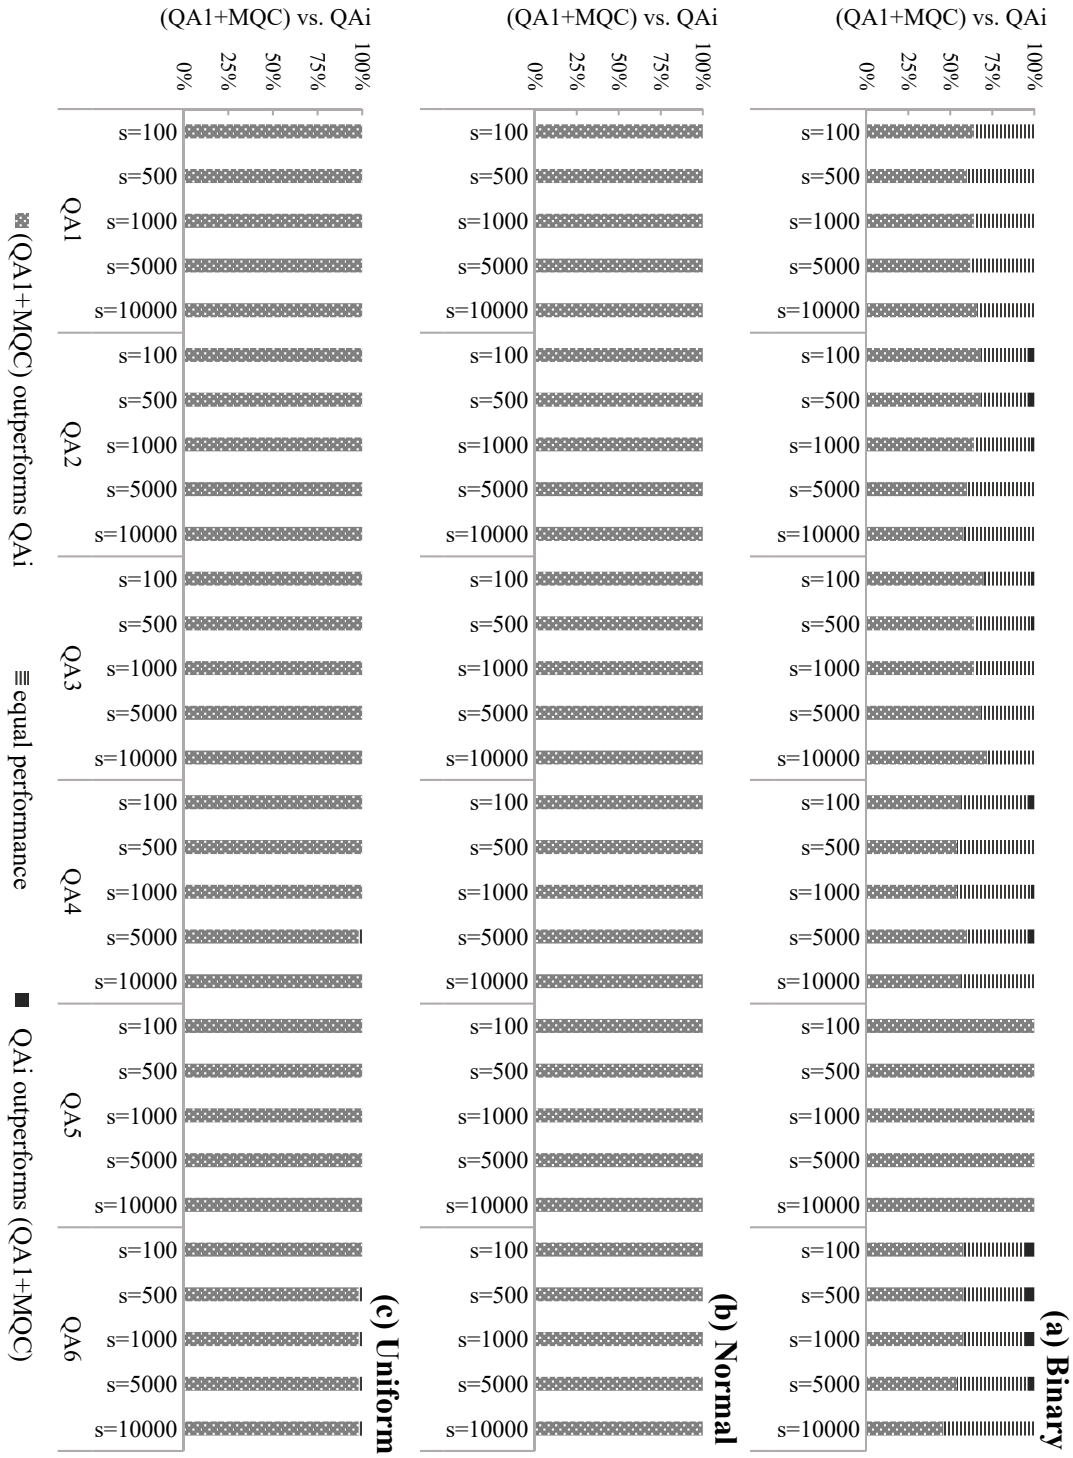

Figure 3. Comparing the performance of applying MQC to raw samples attained by a quantum annealer ($QA^1 + MQC$) with recent software/hardware advances—namely quantum annealing with five spin-reversal transforms (QA^2), quantum annealing with longer inter-sample delay (QA^3), quantum annealing with optimization postprocessing (QA^4), quantum annealing with sampling postprocessing (QA^5), and quantum annealing with five spin-reversal transforms, longer preparation time and optimization postprocessing (QA^6).

become ineffective. Figure 5 compares the performance of applying MQC to samples attained by standard quantum annealing (i.e., raw samples attained by the D-Wave quantum annealers) and enhanced quantum annealing (i.e., quantum annealing with spin-reversal transforms, longer inter-sample delay and classical postprocessing).

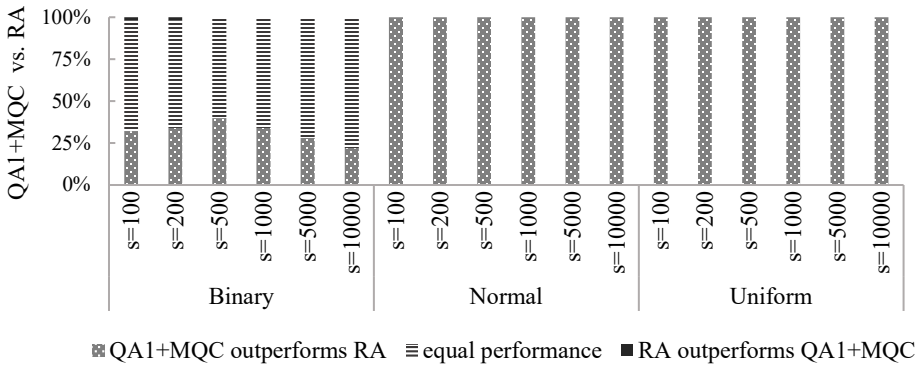

Figure 4. Comparing the performance of reverse quantum annealing (RA) with applying MQC to raw samples attained by a quantum annealer $(\mathrm{QA}^1 + MQC)$.

Table 1. Comparing the robustness (or reproducibility of results) of applying MQC to raw samples attained by a quantum annealer $(\mathrm{QA}^1 + MQC)$ with recent software/hardware advances—namely quantum annealing with five spin-reversal transforms (QA^2), quantum annealing with longer inter-sample delay (QA^3), quantum annealing with optimization postprocessing (QA^4), and quantum annealing with five spin-reversal transforms, longer preparation time and optimization postprocessing (QA^6). Each element represents the variance of energy values from repeating the corresponding method 50 times.

Coefficients	Samples	QA^1	QA^2	QA^3	QA^4	QA^6	$\mathrm{QA}^1 + \mathrm{MQC}$
Binary	100	1.3696	2.8304	1.4544	2.7264	5.2416	0.9984
	200	0.9936	3.4576	0.9664	1.4656	8.1936	0.8704
	500	0.7696	2.3936	0.9104	1.3584	3.9184	0.6400
	1000	0.4816	1.9584	0.9984	0.8464	3.3936	0.2944
Uniform	100	1.6558	1.1323	0.6401	0.6606	1.1116	0.0704
	200	0.8674	0.9988	0.7189	0.4792	1.2079	0.0171
	500	0.5471	0.9898	0.5596	0.4471	0.7513	0.0092
	1000	0.6006	0.5860	0.5792	0.2757	0.7606	0.0049
Normal	100	5.9365	4.0989	2.5668	0.4009	1.0033	0.0563
	200	3.8940	3.5666	2.4667	0.3343	0.5284	0.0071
	500	2.8559	2.0883	2.0424	0.4017	0.4056	0.0000
	1000	2.7226	2.3799	1.5881	0.2467	0.2899	0.0000

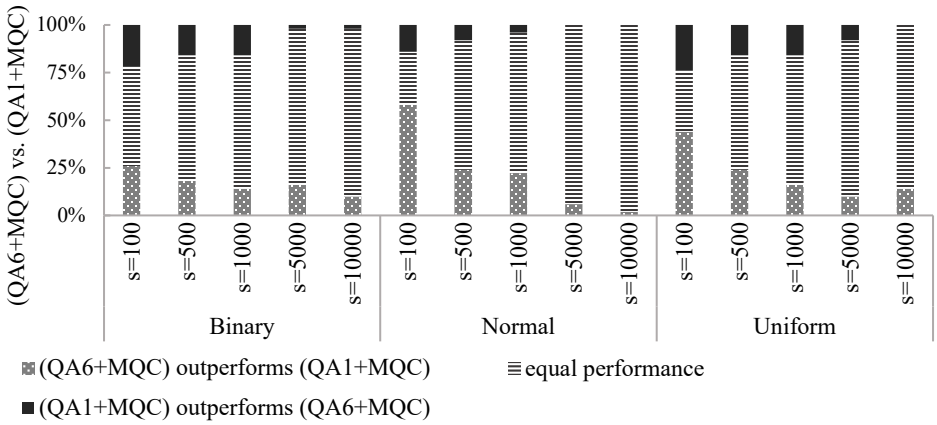

Figure 5. Comparing the performance of applying MQC to raw samples attained by a quantum annealer $(\mathrm{QA}^1 + MQC)$ with applying MQC to samples attained by the quantum annealing with five spin-reversal transforms, longer preparation time and optimization postprocessing $(\mathrm{QA}^6 + MQC)$.

Figure 5 shows that the performance of MQC depends on the quality of the input sample set. We propose a randomized multi-qubit correction (RMQC) scheme, presented in Algorithm 4, that repeats MQC on shuffled sample sets. In RMQC, we repeat the MQC method r times. We start with the raw sample set (similar to MQC) and then shuffle the sample set in each iteration. Finally, we apply MQC on r samples (results from applying MQC on shuffled sample set) to obtain the final solution.

Input: $Z = \left[\mathbf{z}^1, \mathbf{z}^2, \ldots, \mathbf{z}^n\right], r > 0, \mathbf{h}$ and J
Output: $\mathbf{z}^*$
$Z^* \leftarrow \{\}$
do
  $r \leftarrow r - 1$
  Z^*.`Append(MQC(Z))`
  $Z \leftarrow$ `Shuffle(Z)`
while $r > 0$
$\mathbf{z}^* \leftarrow$ `MQC` (Z^*)
return $\mathbf{z}^*$

Algorithm 4: Randomized multi-qubit correction (RMQC).

Figure 6 illustrates the performance of RMQC for $r = 5$ and 10, and compares it with applying MQC to raw samples, attained by the D-Wave quantum annealers. For $r > 1$, RMQC is guaranteed to outperform MQC—albeit r times more (classical) computation time/overhead.

4 Discussion

Owing to various technological barriers—namely diabatic transitions, thermal noise and a vast range of control errors—quantum annealing in a real device (i.e., open-system and stoquastic) is necessarily susceptible to errors[7,38]. While several studies have proposed various error correction approaches for adiabatic quantum computers, most error-correcting schemes for closed-systems are not applicable to quantum annealers[7]. Moreover, error correction techniques (namely nested quantum annealing correction method[31,38]) utilize multiple physical qubits for representing logical qubits that notably reduces the capacity of current quantum annealers. Quantum annealers can draw many high-quality samples in near-constant time (i.e., crystals with only a few defects). However, they generally fail to find the global minimum, specifically when the energy gap(s) between the ground and first excited state(s) is small. In this sense, we can view the open-system quantum annealing process as a Gibbs distribution sampler[31].

In this study, we first studied the impact of applying a local search heuristic, called single qubit correction (SQC), on raw samples drawn by a D-Wave 2000Q quantum annealer. Figure 1 reveals that applying SQC on raw samples for a given Ising Hamiltonian with normal and uniform coefficients always results in a sample with lower energy value. In other words, none of the drawn samples for normal and uniform problems were a local optimum because SQC was able to perform a local search and find another sample with lower energy value. On the other side, roughly all drawn samples for binary problems were a local optimum and applying SQC was not able to improve the quality of the attained samples. There are two possibilities for this observation: (1) Ising Hamiltonians with binary (or discrete) coefficients are easier problems and sampling with a D-Wave quantum annealer is very likely to result in the ground state or (at least) the first excited state; or (2) owing to the precision limitations (e.g., 8–9 bits precision on the D-Wave quantum annealers), the Ising Hamiltonian that is being minimized by a physical quantum annealer is different from the given Ising Hamiltonian (with double precision). Indeed, not only SQC exploits the neighborhood of an input sample for finding a sample with lower energy, but it can also remediate the precision limitations of physical quantum annealers.

We extended SQC to introduced a novel post-quantum error correction/mitigation method, called multi-qubit correction (MQC). The first premise behind MQC is the idea that quantum annealers can draw high-quality samples from the ground state of the given problem Hamiltonian (i.e., crystals with only a few defects). In other words, all samples attained by a quantum annealer partially represent the ground state of the given problem Hamiltonian, although they can contain bits in error. The second premise behind MQC is that we must flip subsets of multiple bits simultaneously for relaxing any excited state to a ground state. From another perspective, flipping bits individually (like how SQC tries to exploit the neighborhood of the measured samples) is (very) likely to result in a local optimum. For every reduction, MQC takes two samples and bitwise compares them to determine which bits are the same and which ones are different. Although identical bits are more likely to be correct, we are not interested in them because we do not know whether they are correct. On the other side, if a bit is different between the two samples, one of the samples has the correct bit value. The objective of MQC is to find groups of isolated bits that flipping them simultaneously can yield a sample with lower energy value. Note that if there is only one isolated group, then flipping all the bits in the group will only change one of the two samples into the other.

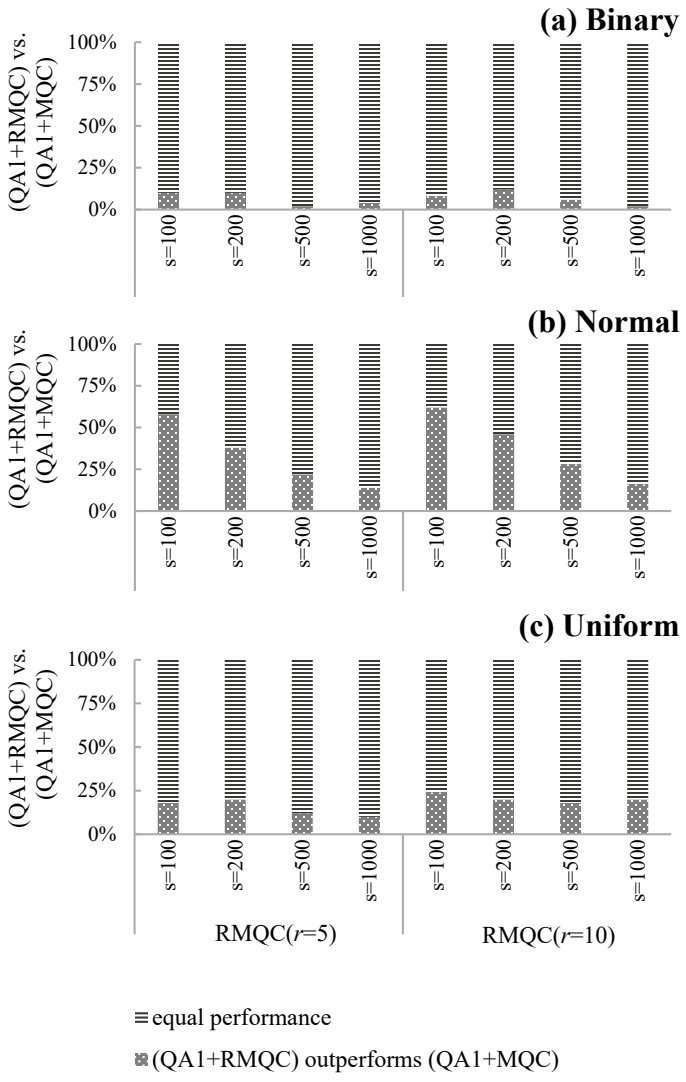

Figure 6. Performance comparison between MQC (denoted by $QA^1 + MQC$) and RMQC with $r = 5$ and 10, denoted by $QA^1 + RMQC$.

Figure 2 explains that applying MQC outperforms SQC. More specifically, for random normal and uniform Ising problems, applying MQC to raw samples always (i.e., in 100% of the employed benchmark problems) results in a sample with lower energy, compared to applying SQC to the same raw samples. On binary problems, nevertheless, MQC was able to outperform SQC in about 63% of cases and they were a tie in approximately 37% of the problems. Figures 3 and 4 demonstrate that MQC outperforms recent software/hardware advances in the realm of quantum annealing—namely the reverse quantum annealing, increasing the inter-sample delays, and applying classical pre/post-processing methods. More specifically, for normal and uniform random Ising problems, MQC always (100% of the used benchmark problems) finds a sample with lower energy value. For binary problems, MQC almost results in a sample with lower energy value. It is worth noting that, however, applying spin-reversal transforms, longer inter-sample delay, optimization postprocessing, and reverse annealing resulted in a better solution in about 2.3%, 1%, 2% and 1% of random binary problems, respectively. Besides, employing all enhancements (denoted by QA^6) was able to outperform MQC in about 5% of random binary problems. From an application viewpoint, near-term quantum processors provide a (noisy) distribution of the ground state(s). Hence, the results of the physical quantum annealers are not well reproducible, mostly due to thermal noise. Table 1 reveals that applying MQC notably improves the robustness of the D-Wave quantum annealers.

Successive measurements on the D-Wave quantum annealers are correlated to each other due to limited preparation time. From another perspective, successive measurements generally form clusters of samples (i.e., groups of identical states). Consequently, technological barriers (e.g., the limited delay between successive reads and the thermal noise) can lessen the performance of MQC.

Figure 5 shows that applying MQC to samples attained by enhanced quantum annealing (denoted by QA^6)—i.e., using five spin-reversal transforms, increasing the preparation time, and performing optimization postprocessing—mostly results in better solutions, compared to applying MQC on raw samples (denoted by QA^1). It is worth highlighting that, nevertheless, increasing the number of reads/samples shrinks the gap between the performance of applying MQC to QA^1 to and QA^6. In this sense, we extended MQC and introduced randomized MQC (RMQC) that re-applies MQC on a shuffled sample set. Figure 6 displays that RMQC is guaranteed to outperform MQC (for $r > 1$), albeit notably more (classical) computations.

Acknowledgements

This research has been supported by NASA grant (#NNH16ZDA001N-AIST 16-0091), NIH-NIGMS Initiative for Maximizing Student Development Grant (2 R25-GM55036), and the Google Lime scholarship. We would like to thank the D-Wave Systems management team for granting access to the D-Wave 2000Q quantum processor.

References

1. Amara, P., Hsu, D. & Straub, J. E. Global energy minimum searches using an approximate solution of the imaginary time schrödinger equation. *The J. Phys. Chem.* **97**, 6715–6721 (1993).
2. Finnila, A., Gomez, M., Sebenik, C., Stenson, C. & Doll, J. Quantum annealing: a new method for minimizing multidimensional functions. *Chem. physics letters* **219**, 343–348 (1994).
3. Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse ising model. *Phys. Rev. E* **58**, 5355 (1998).
4. Das, A. & Chakrabarti, B. K. Colloquium: Quantum annealing and analog quantum computation. *Rev. Mod. Phys.* **80**, 1061 (2008).
5. Ohzeki, M. & Nishimori, H. Quantum annealing: An introduction and new developments. *J. Comput. Theor. Nanosci.* **8**, 963–971 (2011).
6. Nishimori, H. & Takada, K. Exponential enhancement of the efficiency of quantum annealing by non-stoquastic hamiltonians. *Front. ICT* **4**, 2 (2017).
7. Albash, T. & Lidar, D. A. Adiabatic quantum computation. *Rev. Mod. Phys.* **90**, 015002 (2018).
8. Ayanzadeh, R., Halem, M. & Finin, T. Reinforcement quantum annealing: A hybrid quantum learning automata. *Sci. Reports* **10**, 1–11 (2020).
9. Lucas, A. Ising formulations of many np problems. *Front. Phys.* **2**, 5 (2014).
10. McGeoch, C. C. Theory versus practice in annealing-based quantum computing. *Theor. Comput. Sci.* (2020).
11. Ayanzadeh, R., Halem, M. & Finin, T. SAT++: A quantum programming paradigm. *Univ. Maryland, Baltim. Cty.* (2019).
12. Ayanzadeh, R., Mousavi, S., Halem, M. & Finin, T. Quantum annealing based binary compressive sensing with matrix uncertainty. *arXiv preprint arXiv:1901.00088* (2019).
13. Mousavi, S., Taghiabadi, M. M. R. & Ayanzadeh, R. A survey on compressive sensing: Classical results and recent advancements. *arXiv preprint arXiv:1908.01014* (2019).
14. Ayanzadeh, R., Halem, M. & Finin, T. SAT-based compressive sensing. *arXiv preprint arXiv:1903.03650* (2019).
15. Ayanzadeh, R., Halem, M. & Finin, T. An ensemble approach for compressive sensing with quantum annealers. In *IEEE International Geoscience and Remote Sensing Symposium* (2020).
16. Ayanzadeh, R. *Leveraging Artificial Intelligence to Advance Problem-Solving with Quantum Annealers*. Ph.D. thesis, University of Maryland, Baltimore County (2020).
17. Ayanzadeh, R., Halem, M., Dorband, J. & Finin, T. Quantum-assisted greedy algorithms. *arXiv preprint arXiv:1912.02362* (2019).
18. Cai, J., Macready, W. G. & Roy, A. A practical heuristic for finding graph minors. *arXiv preprint arXiv:1406.2741* (2014).
19. Vinci, W., Albash, T., Paz-Silva, G., Hen, I. & Lidar, D. A. Quantum annealing correction with minor embedding. *Phys. Rev. A* **92**, 042310 (2015).
20. Pudenz, K. L., Albash, T. & Lidar, D. A. Quantum annealing correction for random ising problems. *Phys. Rev. A* **91**, 042302 (2015).
21. Dorband, J. E. Extending the d-wave with support for higher precision coefficients. *arXiv preprint arXiv:1807.05244* (2018).

22. Lidar, D. A. Towards fault tolerant adiabatic quantum computation. *Phys. Rev. Lett.* **100**, 160506 (2008).

23. Deng, Q., Averin, D. V., Amin, M. H. & Smith, P. Decoherence induced deformation of the ground state in adiabatic quantum computation. *Sci. reports* **3**, 1–6 (2013).

24. Pudenz, K. L., Albash, T. & Lidar, D. A. Error-corrected quantum annealing with hundreds of qubits. *Nat. communications* **5**, 1–10 (2014).

25. Gardas, B., Dziarmaga, J., Zurek, W. H. & Zwolak, M. Defects in quantum computers. *Sci. reports* **8**, 4539 (2018).

26. Gardas, B. & Deffner, S. Quantum fluctuation theorem for error diagnostics in quantum annealers. *Sci. reports* **8**, 1–8 (2018).

27. King, A. D. & McGeoch, C. C. Algorithm engineering for a quantum annealing platform. *arXiv preprint arXiv:1410.2628* (2014).

28. Jordan, S. P., Farhi, E. & Shor, P. W. Error-correcting codes for adiabatic quantum computation. *Phys. Rev. A* **74**, 052322 (2006).

29. Sarovar, M. & Young, K. C. Error suppression and error correction in adiabatic quantum computation: non-equilibrium dynamics. *New J. Phys.* **15**, 125032 (2013).

30. Mizel, A. Fault-tolerant, universal adiabatic quantum computation. *arXiv preprint arXiv:1403.7694* (2014).

31. Vinci, W., Albash, T. & Lidar, D. A. Nested quantum annealing correction. *npj Quantum Inf.* **2**, 1–6 (2016).

32. Young, K. C., Blume-Kohout, R. & Lidar, D. A. Adiabatic quantum optimization with the wrong hamiltonian. *Phys. Rev. A* **88**, 062314 (2013).

33. Young, K. C., Sarovar, M. & Blume-Kohout, R. Error suppression and error correction in adiabatic quantum computation: Techniques and challenges. *Phys. Rev. X* **3**, 041013 (2013).

34. Ganti, A., Onunkwo, U. & Young, K. Family of [[6 k, 2 k, 2]] codes for practical and scalable adiabatic quantum computation. *Phys. Rev. A* **89**, 042313 (2014).

35. Bookatz, A. D., Farhi, E. & Zhou, L. Error suppression in hamiltonian-based quantum computation using energy penalties. *Phys. Rev. A* **92**, 022317 (2015).

36. Matsuura, S., Nishimori, H., Albash, T. & Lidar, D. A. Mean field analysis of quantum annealing correction. *Phys. review letters* **116**, 220501 (2016).

37. Mishra, A., Albash, T. & Lidar, D. A. Performance of two different quantum annealing correction codes. *Quantum Inf. Process.* **15**, 609–636 (2016).

38. Matsuura, S., Nishimori, H., Vinci, W. & Lidar, D. A. Nested quantum annealing correction at finite temperature: p-spin models. *Phys. Rev. A* **99**, 062307 (2019).

39. Aliferis, P., Gottesman, D. & Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. *arXiv preprint quant-ph/0504218* (2005).

40. King, A. D., Hoskinson, E., Lanting, T., Andriyash, E. & Amin, M. H. Degeneracy, degree, and heavy tails in quantum annealing. *Phys. Rev. A* **93**, 052320 (2016).

41. Pelofske, E., Hahn, G. & Djidjev, H. Optimizing the spin reversal transform on the d-wave 2000q. *arXiv preprint arXiv:1906.10955* (2019).

42. Golden, J. K. & O'Malley, D. Pre-and post-processing in quantum-computational hydrologic inverse analysis. *arXiv preprint arXiv:1910.00626* (2019).

43. Ottaviani, D. & Amendola, A. Low rank non-negative matrix factorization with d-wave 2000q. *arXiv preprint arXiv:1808.08721* (2018).

44. Venturelli, D. & Kondratyev, A. Reverse quantum annealing approach to portfolio optimization problems. *Quantum Mach. Intell.* **1**, 17–30 (2019).

45. Yamashiro, Y., Ohkuwa, M., Nishimori, H. & Lidar, D. A. Dynamics of reverse annealing for the fully connected p-spin model. *Phys. Rev. A* **100**, 052321 (2019).

46. Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194 (2011).

47. Dorband, J. E. A method of finding a lower energy solution to a qubo/ising objective function. *arXiv preprint arXiv:1801.04849* (2018).

48. Engelbrecht, A. P. *Computational intelligence: an introduction* (John Wiley & Sons, 2007).

49. Russell, S. J. & Norvig, P. *Artificial intelligence: a modern approach* (Malaysia; Pearson Education Limited,, 2016).

50. Denchev, V. S. *et al.* What is the computational value of finite-range tunneling? *Phys. Rev. X* **6**, 031015 (2016).

51. Shahamatnia, E., Ayanzadeh, R., Ribeiro, R. A. & Setayeshi, S. Adaptive imitation scheme for memetic algorithms. In *Doctoral Conference on Computing, Electrical and Industrial Systems*, 109–116 (Springer, 2011).

52. Altshuler, B., Krovi, H. & Roland, J. Anderson localization makes adiabatic quantum optimization fail. *Proc. Natl. Acad. Sci.* **107**, 12446–12450 (2010).