scientific reports

OPEN

Multi-qubit correction for quantum annealers

Ramin Ayanzadeh[1]✉, John Dorband[2], Milton Halem[2] & Tim Finin[2]

We present *multi-qubit correction* (MQC) as a novel postprocessing method for quantum annealers that views the evolution in an open-system as a Gibbs sampler and reduces a set of 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 value— 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 improves the reproducibility of results when compared to recent hardware/software advances in the realm of quantum annealing, such as spin-reversal transforms, classical postprocessing techniques, and increased inter-sample delay between successive measurements.

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[11,12]. In practice, however, executing the corresponding quantum machine instruction (QMI) on a physical quantum annealer does not guarantee achieving the global optimum[7,8,13]. In addition to thermal noise and diabatic transitions[7], examples of control error sources include sparse connectivity between qubits[14,15], confined annealing schedule[6], coefficients' range and precision limitations[16,17], and noise and decoherence[18–22]. These error sources lower the quality of results, i.e., the energy value of the drawn samples is higher than the energy value of the ground state[23].

Modifying some aspects of the Hamiltonian by 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), or 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 suitable 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 just as we do with other quantum information processing models[24–27]. In spite of several error correction proposals for adiabatic quantum computing and quantum annealing[15,16,18,20,24–34], an accuracy-threshold theorem for adiabatic quantum computing, unlike its gate model counterpart[35], remains elusive[20,27]. Besides, most error correction schemes (e.g., nested quantum annealing correction method[27,34]) utilize multiple physical qubits for coding every qubit that notably reduces the capacity of current quantum annealers.

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) the results/samples attained by the physical quantum annealers are not reproducible over time. Applying classical postprocessing techniques to (raw) samples attained by the physical quantum annealers can mitigate these drawbacks to some extent. For example, one may apply evolutionary algorithms (e.g., genetic algorithms) or swarm intelligence techniques (such as particle swarm optimization) as a postprocessing method for quantum annealers[36]. Nevertheless, these heuristics and meta-heuristics are stochastic techniques that can result in a different solution in each try and lessen the robustness of results

[1]College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, USA. [2]Computer Science and Electrical Engineering, University of Maryland, Baltimore County, Baltimore, MD 21250, USA. ✉email: ayanzadeh@gatech.edu

1

in the application domain. Moreover, most of these techniques require hyper-parameter optimization—e.g., the number of iterations, the probability of crossover and mutation, and selection method in evolutionary algorithms—that can significantly impact the quality of ultimate results[37]. To this end, we devise a lightweight and deterministic postprocessing method that can improve the quality and robustness of results attained by physical quantum annealers. It is worth highlighting that, in this study, we did not aim to demonstrate quantum speed-up.

We view quantum annealers as a Gibbs sampler that allows diabatic transitions in an open system and present *multi-qubit correction* (MQC) as a novel postprocessing technique for quantum annealers. Unlike most studies that try to identify different types of errors and mitigate/correct them, we try to recognize the pattern(s) among the incorrect observations (or lower quality samples compared to the ground state of the given Hamiltonian) and leverage it to achieve a better solution. Measurement is the most error-prone operation on superconducting noisy intermediate-scale quantum (NISQ)[38] machines[39]—here, we look at the quantum annealers as the NISQ model of adiabatic quantum computers. Therefore, in this study, we mainly focus on bit-flips that mostly occur due to measurement errors on superconducting quantum devices. MQC compares pairs of samples and recognizes groups of qubits that we can flip their values simultaneously, and tunnels through these groups to converge to a (notably) better solution. Our experiments using the D-Wave 2000Q quantum annealers show that MQC utilizes fewer samples and finds better solutions compared to recent software/hardware advances in the realm of quantum annealing.

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 local fields and couplers, respectively[10,11,40]. 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 annealers generally yield an excited state(s) that are not necessarily a local optimum[8,13,41]. 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[41].

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 given problem Hamiltonian[4]. Owing to hardware limitations such as limited precision of coefficients, the Hamiltonian that is minimized on a physical quantum annealer can be (slightly) different from the problem Hamiltonian of interest (in application domain); therefore, the evolution is not guaranteed to result in an eigenstate of the problem Hamiltonian[7]. We start by adopting a hill-climbing[36,42] algorithm to present a postprocessing 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 a lower energy value. In each iteration of SQC, we individually toggle the value of every qubit and keep all of the 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 z

Algorithm 1: Single-qubit correction (SQC) heuristic for exploiting the neighborhood of excited states to find a sample with lower energy value.
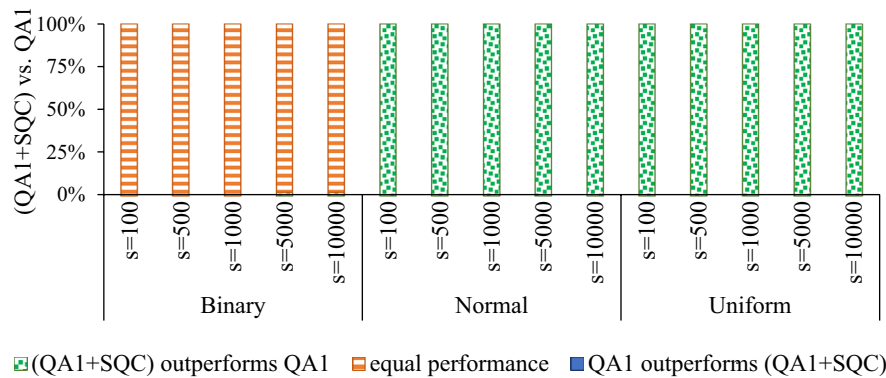
Figure 1. Performance comparison between quantum annealing (QA^1) and applying SQC to raw samples attained by quantum annealers ($QA^1 + SQC$) where s denotes the number of samples.

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. Therefore, in practice, SQC explores a broader area, i.e., the neighborhood of all excited states attained by the quantum annealers. Figure 1 illustrates the impact of applying SQC to raw samples drawn by the D-Wave quantum annealers. For uniform and normal benchmark problems, SQC always finds a better sample, and the corresponding p value is 8.8818×10^{-16} which indicates that results are statistically significant (i.e., $p < 0.05$). However, for Binary problems, both SQC and the baseline method demonstrated a similar performance (i.e., results are not statistically significant). For more information about the benchmarking, see the section "Methods".

It is worth highlighting that we do not propose SQC as a postprocessing method for quantum annealers since more efficient techniques, such as simulated annealing, can outperform SQC in terms of finding samples with lower energy values. Indeed, we will extend SQC to introduce a novel postprocessing approach that can notably improve the quality and reproducibility of results attained by the quantum annealers.

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 to find these tunnels. In this study, the term "tunnel" is analogous to the concept of quantum tunnels[43] and refers to a group of qubits that a quantum annealer simultaneously flips their values to change an excited state to a sample with a 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 a 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 of 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

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

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 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 $\mathbf{z}^2$.

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** $\mathbf{z}_i^1 = \mathbf{z}_i^2$ **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**
> > $I_{\mathbf{z}^1}^k = \sum_{i \in T^k} \mathbf{h}_i \mathbf{z}_i^1 + \sum_{i \in T^k} \sum_{j \in S} J_{ij} \mathbf{z}_i^1 \mathbf{z}_j^1$ `// Influence of sub-tunnels on` E_{Ising}
> > **if** $I_{\mathbf{z}^1}^k > 0$ **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.

The Reduce procedure presented in this paper is analogous to the crossover operation in evolutionary algorithms that acts on two potential solutions, known as parent chromosomes, and yields new solution(s), known as offspring[36,37]. 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 when compared to SQC that depends entirely on a single excited state.

When we employ physical quantum annealers, we generally request 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 a 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.
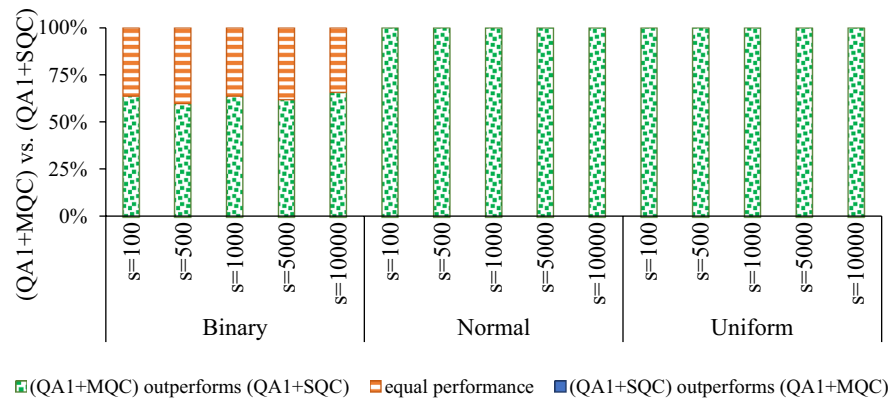
Figure 2. Performance comparison between applying MQC and SQC to raw samples attained by quantum annealers, denoted by $QA^1 + MQC$ and $QA^1 + SQC$, respectively (s denotes the number of samples).

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**
        $\mathbf{z}^A \leftarrow \hat{Z}.\texttt{Pop()}$
        $\mathbf{z}^B \leftarrow \hat{Z}.\texttt{Pop()}$
        $\mathbf{z}^{AB} \leftarrow \texttt{Reduce}(\mathbf{z}^A, \mathbf{z}^B, \mathbf{h}, J)$
        $Z.\texttt{Append}(\mathbf{z}^{AB})$
    end
    if $|\hat{Z}| > 0$ **then**
        $Z.\texttt{Append}(\hat{Z}.\texttt{Pop()})$
    end
end
$\mathbf{z}^* \leftarrow Z.\texttt{Pop()}$

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 the worst case) is equal to the lowest energy value of input samples.

Figure 2 shows that for normal and uniform benchmark problems MQC always outperforms SQC (i.e., finds a better sample with lower energy value), and the corresponding p value is 8.8818×10^{-16}. For Binary benchmark problems, the success rate of MQC (i.e., MQC outperforms SQC) ranges from 60 to 66%, and the corresponding p values ranges from 9.3132×10^{-10} to 1.1641×10^{-10} which indicate that results are statistically significant. We remark that SQC was not able to outperform MQC in any arrangement.

Figure 3 demonstrates 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 field of quantum annealing, such as spin-reversal transforms, optimization, and sampling postprocessing methods, and increased inter-sample delay between successive measurements. For normal benchmark problems, MQC always outperforms all baselines (i.e., finds a better sample with lower energy value), and the corresponding p value is 8.8818×10^{-16}. For uniform benchmark problems, the success rate of MQC (i.e., MQC outperforms the baseline method) ranges from 98 to 100%, and the corresponding p values ranges from 1.7764×10^{-15} to 8.8818×10^{-16}. For Binary benchmark problems, the success rate of MQC (i.e., MQC outperforms the baseline method) ranges from 46 to 100%, and the corresponding p values ranges from 1.1921×10^{-7} to 8.8818×10^{-16}. We remark that none of the baseline methods was not able to outperform MQC in any arrangement, and results are statistically significant (i.e., $p < 0.05$) in all arrangements. These results explain that MQC requires notably fewer samples to visit the ground state of the given Hamiltonian with a high enough probability compared to recent software and hardware advances in the realm of quantum annealing. For more information about the benchmarking, see the section "Methods".

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)
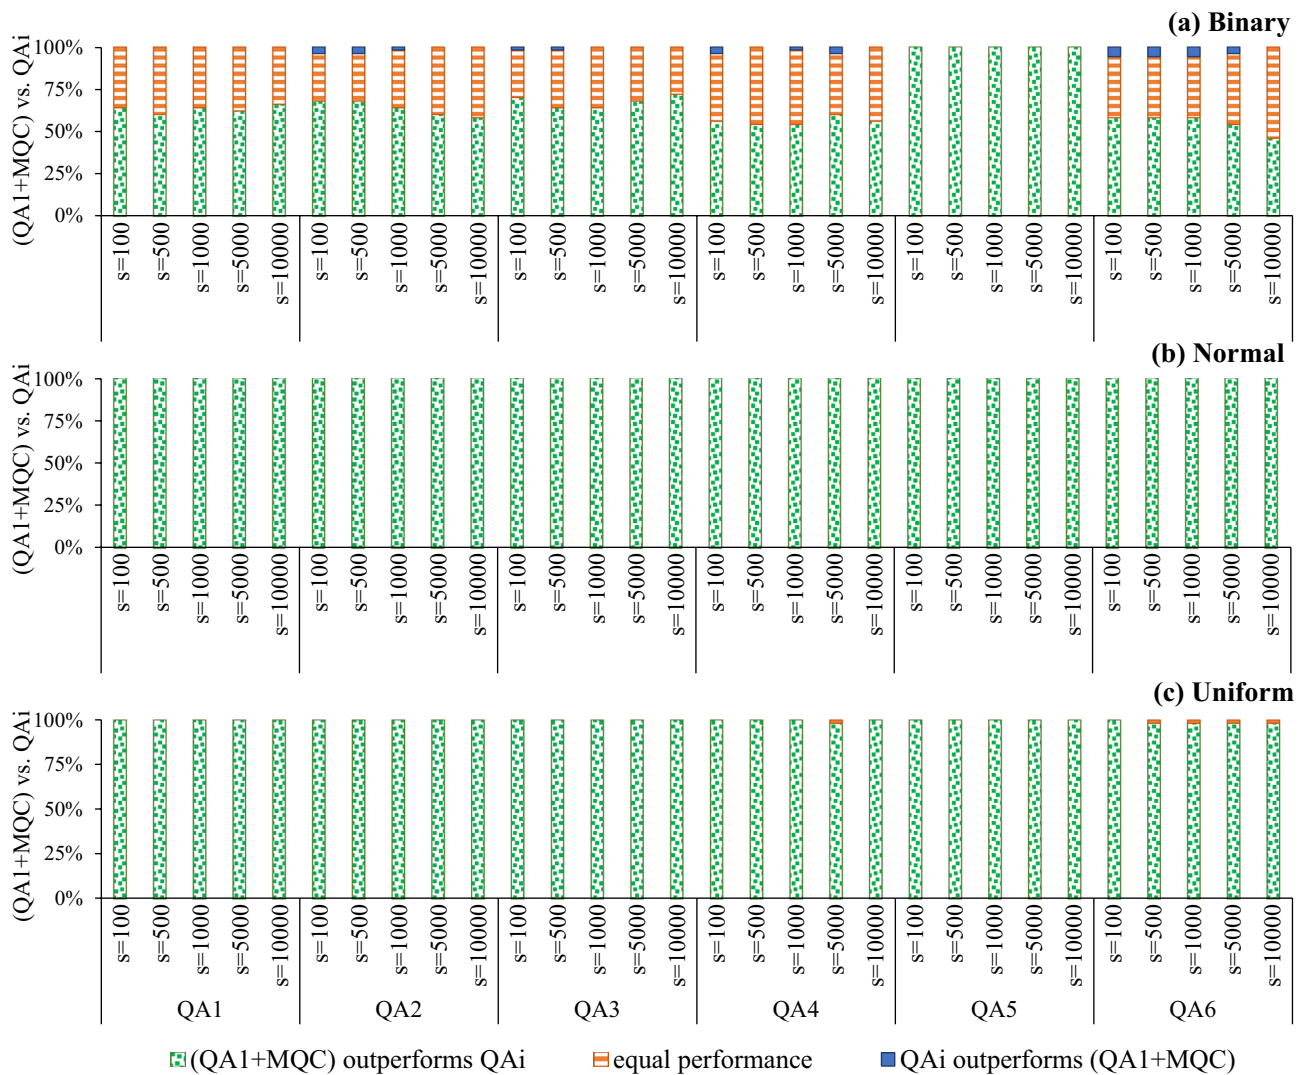
Figure 3. Comparing the performance of applying MQC to raw samples attained by a quantum annealer ($QA^1 + MQC$) with recent software/hardware advances—namely quantum annealing with five spin-reversal transforms (QA^2), quantum annealing with longer inter-sample delay (QA^3), quantum annealing with optimization 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) where s denotes the number of samples.

the results/samples attained by the physical quantum annealers are not reproducible over time. According to Anderson localization[44], 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 is 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 repeated experiments. Figure 3 explains that applying the sampling postprocessing on raw samples attained by the D-Wave quantum annealers (QA^5) can significantly lower the quality of results. In other words, applying the sampling postprocessing method can result in a sample with a higher energy value. On this basis, we omitted QA^5 from Table 1. These results reveal that applying MQC on (raw) samples attained by the D-Wave quantum processors can notably improve the reproducibility of results.

Randomized MQC. MQC is a postprocessing heuristic that iteratively reduces a set of samples to a smaller sample set whose energy values are lower than the previous iteration(s). Therefore, the performance of MQC mostly depends on samples that we obtain from the physical quantum annealers. When we repeat the annealing process on a D-Wave quantum annealer, successive measurements are correlated to each other due to limited preparation time. Hence, 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 become ineffective. Figure 4

Coefficients	Samples	QA^1	QA^2	QA^3	QA^4	QA^6	QA^1 + 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

Table 1. Comparing the robustness (or reproducibility of results) 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), 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.
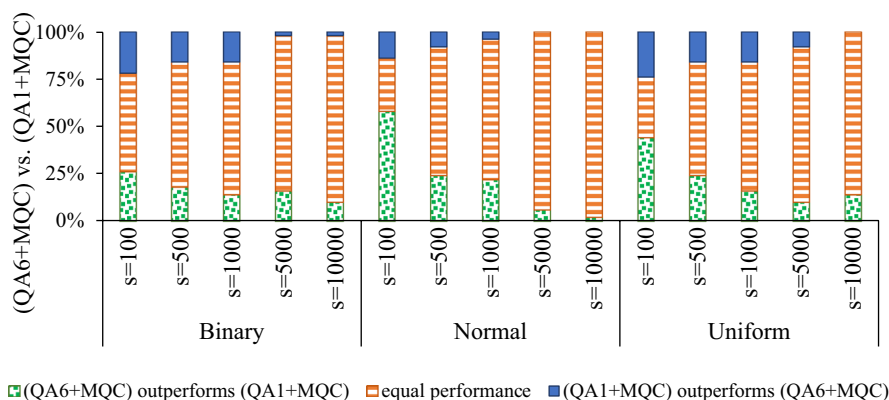

Figure 4. Comparing the performance of applying MQC to raw samples attained by a quantum annealer (QA^1 + MQC) with applying MQC to samples attained by the quantum annealing with five spin-reversal transforms, longer preparation time and optimization postprocessing (QA^6 + MQC) where s denotes the number of samples.

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 optimization postprocessing). For normal benchmark problems, results are statistically significant when we draw fewer than 5000 samples. p values of uniform and Binary problems are less than 0.05 only when we draw 10,000 and 5000 samples, respectively.

Figure 4 suggests that the quality of the input sample set can impact the performance of MQC. We propose the *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. Note that when $r = 1$ MQC and RMQC are identical. For $r > 1$, RMQC is guaranteed to outperform MQC—albeit r times more (classical) computation time/overhead.
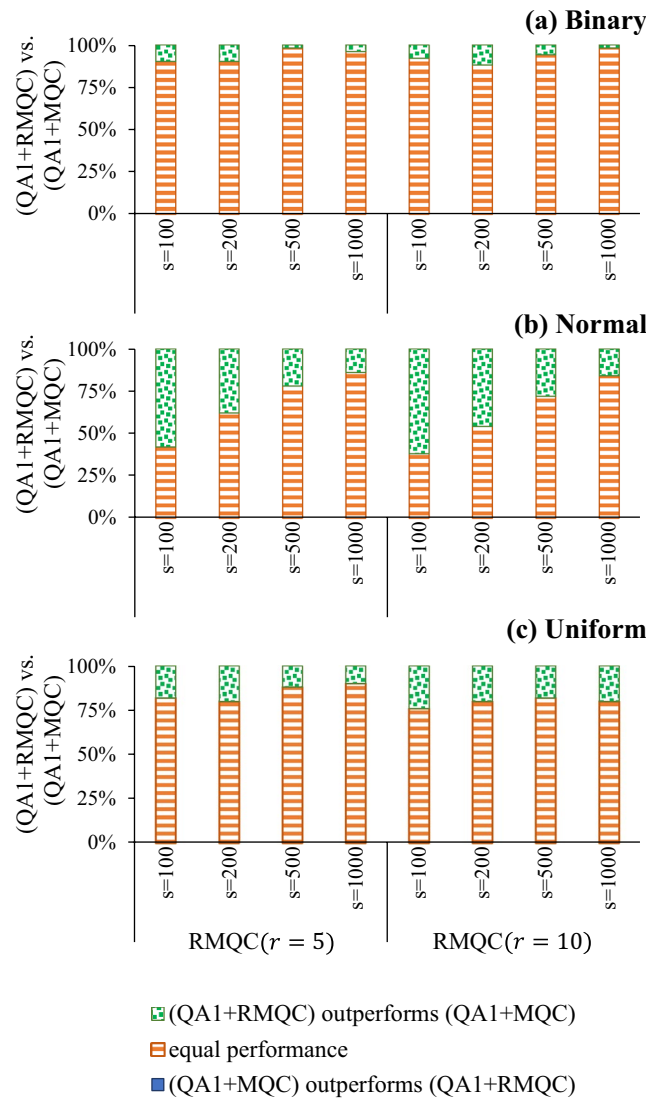
Figure 5. Performance comparison between MQC (denoted by $QA^1 + MQC$) and RMQC with $r = 5$ and 10, denoted by $QA^1 + RMQC$ (s represents the number of samples).

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
$\quad \mid \quad r \leftarrow r - 1$
$\quad \mid \quad Z^*.\texttt{Append}(\texttt{MQC}(Z))$
$\quad \mid \quad Z \leftarrow \texttt{Shuffle}(Z)$
while $r > 0$
$\mathbf{z}^* \leftarrow \texttt{MQC}(Z^*)$
return $\mathbf{z}^*$

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

Figure 5 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 normal and uniform benchmark problems, results are always statistically significant (i.e., $p < 0.05$). More specifically, when $r = 5$, $p \in \left[0.0078, 1.8626 \times 10^{-9}\right]$ and $p \in [0.0313, 0.001]$ for normal and uniform benchmark problems, respectively. When $r - 10$,

$p \in [0.004, 4.6566 \times 10^{-10}]$ and $p \in [0.002, 0.0003]$ for normal and uniform benchmark problems, respectively. For Binary problems, however, results of MQC and RMQC are not statistically significant—for sufficient number of samples MQC and RMQC demonstrate similar performance on Binary benchmark problems.

Discussion

Owing to various technological barriers such as 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[7,34]. While several studies have proposed various error correction approaches for adiabatic quantum computers, most of them are not applicable to quantum annealers[7]. Moreover, error correction techniques such as the nested quantum annealing correction method[27,34] use multiple physical qubits for representing logical qubits that notably reduce the capacity of current quantum annealers. Quantum annealers can draw many high-quality samples in near-constant time. In other words, the annealing time does not depend on the number of qubits. However, they generally fail to find the global minimum, specifically when the energy gap between the ground and the first excited state(s) is small.

While most studies try to recognize specific types of errors and mitigate/correct them, we show that one can recognize the pattern(s) of incorrect observations (or lower quality samples) and leverage it to improve the fidelity of the quantum annealers. In this sense, we view the open-system quantum annealing process as a Gibbs distribution sampler[27] and exploit the neighborhood of the drawn samples to find a better solution. We first examined the impact of applying a local search heuristic, called *single qubit correction* (SQC), on raw samples drawn by a D-Wave 2000Q quantum annealers. 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 a 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 a lower energy value. On the other hand, roughly all drawn samples for binary problems were a local (or global) optimum, and applying SQC could not 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 (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). Previous studies have mentioned that random Ising problems might not be hard-enough problems for quantum annealers (and even classical optimization techniques)[45,46]. In practice, nevertheless, we see that applying SQC/MQC can notably improve the fidelity of the D-Wave quantum annealers. More specifically, while for most random Binary problems applying SQC/MQC does not result in a better solution—the D-Wave quantum annealer can itself find the best (achievable) solution—applying MQC to random Normal and Uniform Ising Hamiltonians always results in a better solution. Hence, we can conclude that (in practice) easy to solve random Ising problems can be challenging for NISQ machines—we look at quantum annealers as the NISQ model of adiabatic quantum computers. Indeed, not only does 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 introduce a novel postprocessing 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. 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 is that we must simultaneously flip subsets of multiple bits to relax 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 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 hand, 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 such that simultaneously flipping them can yield a sample with a lower energy value. 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 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 than applying SQC to the same set of 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. Figure 3 demonstrates that MQC outperforms recent software/hardware advances in the realm of quantum annealing—namely increasing the inter-sample delays and applying classical pre/post-processing methods that are available in Ocean SDK. More specifically, for normal and uniform random Ising problems, MQC always (100% of the used benchmark problems) finds a sample with a lower energy value.

For binary problems, MQC almost results in a sample with a lower energy value. However, it is worth noting that applying spin-reversal transforms, longer inter-sample delay, and optimization postprocessing resulted in a better solution in about 2.3%, 1%, and 2% 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.

MQC might be considered related to sample persistence[47] or meta-heuristics (such as evolutionary algorithms and swarm intelligence techniques)[37]. MQC is a deterministic postprocessing method—i.e., for a fixed set of input samples, it will always yield the same solution—but most postprocessing schemes are stochastic. Hence, applying MQC (notably) improves the reproducibility of results attained by the quantum annealers. Moreover, most meta-heuristics require hyperparameter optimization (e.g., number of iterations, and probabilities of crossover and mutation in evolutionary algorithms); however, run-time and performance of MQC mainly depend on the size and quality of the input sample set, respectively. More specifically, for n samples, MQC performs $\log_2 n$ iterations, and in iteration i it compares $n/2^i$ pairs of samples.

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 4 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. RMQC is guaranteed to outperform MQC (for $r > 1$), albeit notably more (classical) computations.

Methods

This paper presents a novel postprocessing method for quantum annealers, called *multi-qubit correction* (MQC), that notably improves quantum annealers' performance in terms of reproducibility of results and finding solutions with lower energy values.

Quantum hardware.

For all evaluations, we used a D-Wave 2000Q quantum annealer by D-Wave Systems Inc. (located at Burnaby, British Colombia, Canada). The annealing time for all experiments was 20 microseconds.

Study cases.

Generating random Hamiltonians is a common practice for benchmarking quantum annealers[4,13,16,48]. Hence, we generated three different types of Ising Hamiltonians as follows:

- Binary—random Ising Hamiltonians whose linear and quadratic coefficients (denoted by **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 **h** and *J*, respectively) are (double-precision) uniform random numbers in $[-1, +1]$;
- Normal—random Ising Hamiltonians whose linear and quadratic coefficients (denoted by **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) and 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 D-Wave 2000Q quantum processor's working graph—in Chimera topology, every qubit is connected to at most six other qubits. Therefore, until the next maintenance that can change the quantum annealer's working graph, one can directly execute them without embedding problems to a target graph.

We remark that random Ising problems might not be hard-enough problems for quantum annealers (and even classical optimization techniques)[45]. To demonstrate the quantum speed-up, one need to use proper benchmarks[46]. It is worth highlighting that our objective in this study was not to show/claim quantum speed-up.

Baselines.

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 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 (also known as 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[49]. Similarly, applying classical postprocessing heuristics on raw samples (attained by the quantum annealers) can result in samples with lower energy values[50,51]. Furthermore, when we submit a problem to a D-Wave quantum annealer, it is a common practice to request several samples/reads (i.e., 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. Therefore, when we repeat the annealing process, 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). Hence, increasing the preparation time can reduce the inter-sample correlations. In this study, we used the following arrangements for evaluating the performance of the proposed postprocessing 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).

Evaluations. To compare MQC with baseline methods, we count how many times the *expected to win* method outperforms the baseline (i.e., finds a sample with lower energy value), ignoring cases where both methods can find an identical solution (in terms of energy value). Since case studies used here are all random problems, of interest is how statistically significant such results are. Hence, we performed hypothesis testing. Our null hypothesis in all evaluations was that the *expected to win* method and the baseline are not different, and the alternate hypothesis was that the *expected to win* method outperforms the baseline (i.e., can always find a sample with lower energy value). To this end, we calculate the p value as:

$$p = \frac{1}{2^{n_b+n_w}} \sum_{k=n_b}^{n_b+n_w} \binom{n_b + n_w}{k}. \tag{6}$$

where n_b represents the number of times that the *expected to win* method was able to outperform the baseline and n_w is the number of times that the baseline was able to outperform the *expected to win* method. This represents the probability of the null hypothesis is true. By convention, $p < 0.05$ disproves the null hypothesis (i.e., the *expected to win* indeed outperforms the baseline) and is considered to be statistically significant, and $0.05 \leq p \leq 0.95$ indicates that results are not statistically significant. While $p \geq 0.95$ also disprove the null hypothesis (i.e., results are statistically significant), it also rejects the alternate hypothesis and indicates that the expected winner has been chosen incorrectly[52].

Data Availability. MQC code (in Python) is publicly available at https://doi.org/10.5281/zenodo.5142230.

References

1. Amara, P., Hsu, D. & Straub, J. E. Global energy minimum searches using an approximate solution of the imaginary time Schrödinger equation. *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. Phys. Lett.* **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. Rep.* **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.* **816**, 169–183 (2020).
11. Ayanzadeh, R. *Leveraging Artificial Intelligence to Advance Problem-Solving with Quantum Annealers*. Ph.D. thesis, University of Maryland, Baltimore County (2020).
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. Ayanzadeh, R., Halem, M., Dorband, J. & Finin, T. Quantum-assisted greedy algorithms. arXiv preprint arXiv:1912.02362 (2019).
14. Cai, J., Macready, W. G. & Roy, A. A practical heuristic for finding graph minors. arXiv preprint arXiv:1406.2741 (2014).
15. Vinci, W., Albash, T., Paz-Silva, G., Hen, I. & Lidar, D. A. Quantum annealing correction with minor embedding. *Phys. Rev. A* **92**, 042310 (2015).
16. Pudenz, K. L., Albash, T. & Lidar, D. A. Quantum annealing correction for random Ising problems. *Phys. Rev. A* **91**, 042302 (2015).
17. Dorband, J. E. Extending the d-wave with support for higher precision coefficients. arXiv preprint arXiv:1807.05244 (2018).
18. Lidar, D. A. Towards fault tolerant adiabatic quantum computation. *Phys. Rev. Lett.* **100**, 160506 (2008).
19. Deng, Q., Averin, D. V., Amin, M. H. & Smith, P. Decoherence induced deformation of the ground state in adiabatic quantum computation. *Sci. Rep.* **3**, 1–6 (2013).
20. Pudenz, K. L., Albash, T. & Lidar, D. A. Error-corrected quantum annealing with hundreds of qubits. *Nat. Commun.* **5**, 1–10 (2014).
21. Gardas, B., Dziarmaga, J., Zurek, W. H. & Zwolak, M. Defects in quantum computers. *Sci. Rep.* **8**, 4539 (2018).
22. Gardas, B. & Deffner, S. Quantum fluctuation theorem for error diagnostics in quantum annealers. *Sci. Rep.* **8**, 1–8 (2018).
23. King, A. D. & McGeoch, C. C. Algorithm engineering for a quantum annealing platform. arXiv preprint arXiv:1410.2628 (2014).
24. Jordan, S. P., Farhi, E. & Shor, P. W. Error-correcting codes for adiabatic quantum computation. *Phys. Rev. A* **74**, 052322 (2006).
25. Sarovar, M. & Young, K. C. Error suppression and error correction in adiabatic quantum computation: Non-equilibrium dynamics. *New J. Phys.* **15**, 125032 (2013).
26. Mizel, A. Fault-tolerant, universal adiabatic quantum computation. arXiv preprint arXiv:1403.7694 (2014).
27. Vinci, W., Albash, T. & Lidar, D. A. Nested quantum annealing correction. *npj Quantum Inf.* **2**, 1–6 (2016).

28. Young, K. C., Blume-Kohout, R. & Lidar, D. A. Adiabatic quantum optimization with the wrong Hamiltonian. *Phys. Rev. A* **88**, 062314 (2013).
29. 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).
30. 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).
31. Bookatz, A. D., Farhi, E. & Zhou, L. Error suppression in Hamiltonian-based quantum computation using energy penalties. *Phys. Rev. A* **92**, 022317 (2015).
32. Matsuura, S., Nishimori, H., Albash, T. & Lidar, D. A. Mean field analysis of quantum annealing correction. *Phys. Rev. Lett.* **116**, 220501 (2016).
33. Mishra, A., Albash, T. & Lidar, D. A. Performance of two different quantum annealing correction codes. *Quantum Inf. Process.* **15**, 609–636 (2016).
34. 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).
35. Aliferis, P., Gottesman, D. & Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. arXiv preprint arXiv: quant-ph/0504218 (2005).
36. Engelbrecht, A. P. *Computational Intelligence: An Introduction* (Wiley, 2007).
37. 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).
38. Preskill, J. Quantum computing in the nisq era and beyond. arXiv preprint arXiv:1801.00862 (2018).
39. Tannu, S. S. & Qureshi, M. K. Mitigating measurement errors in quantum computers by exploiting state-dependent bias. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 279–290 (2019).
40. Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194 (2011).
41. Dorband, J. E. A method of finding a lower energy solution to a qubo/Ising objective function. arXiv preprint arXiv:1801.04849 (2018).
42. Russell, S. J. & Norvig, P. *Artificial Intelligence: A Modern Approach* (Malaysia; Pearson Education Limited, 2016).
43. Denchev, V. S. *et al.* What is the computational value of finite-range tunneling?. *Phys. Rev. X* **6**, 031015 (2016).
44. Altshuler, B., Krovi, H. & Roland, J. Anderson localization makes adiabatic quantum optimization fail. *Proc. Natl. Acad. Sci.* **107**, 12446–12450 (2010).
45. Katzgraber, H. G., Hamze, F. & Andrist, R. S. Glassy chimeras could be blind to quantum speedup: Designing better benchmarks for quantum annealing machines. *Phys. Rev. X* **4**, 021008 (2014).
46. Marshall, J., Martin-Mayor, V. & Hen, I. Practical engineering of hard spin-glass instances. *Phys. Rev. A* **94**, 012320 (2016).
47. Karimi, H. & Rosenberg, G. Boosting quantum annealer performance via sample persistence. *Quantum Inf. Process.* **16**, 166 (2017).
48. 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).
49. Pelofske, E., Hahn, G. & Djidjev, H. Optimizing the spin reversal transform on the d-wave 2000q. arXiv preprint arXiv:1906.10955 (2019).
50. Borle, A. & McCarter, J. On post-processing the results of quantum optimizers. In *International Conference on Theory and Practice of Natural Computing*, 222–233 (Springer, 2019).
51. Golden, J. K. & O'Malley, D. Pre-and post-processing in quantum-computational hydrologic inverse analysis. arXiv preprint arXiv: 1910.00626 (2019).
52. Chen, J., Stollenwerk, T. & Chancellor, N. Performance of domain-wall encoding for quantum annealing. arXiv preprint arXiv: 2102.12224 (2021).

Acknowledgements

Author contributions

R.A. and J.D. conceived algorithms, and R.A. carried out implementations. M.H. and T.F. performed analysis.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to R.A.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.