

QUANTUM-ASSISTED GREEDY ALGORITHMS

Ramin Ayanzadeh^{1,2*}, John Dorband², Milton Halem² and Tim Finin²

¹Georgia Institute of Technology, Atlanta, GA 30332, US

²University of Maryland, Baltimore County, Baltimore, MD 21250, US

ABSTRACT

We show how to leverage quantum annealers (QAs) to better select candidates in greedy algorithms. Unlike conventional greedy algorithms that employ problem-specific heuristics for making locally optimal choices at each stage, we use QAs that sample from the ground state of a problem-dependent Hamiltonians at cryogenic temperatures and use retrieved samples to estimate the probability distribution of problem variables. More specifically, we look at each spin of the Ising model as a random variable and contract all problem variables whose corresponding uncertainties are negligible. Our empirical results on a D-Wave 2000Q quantum processor demonstrate that the proposed *quantum-assisted greedy algorithm* (QAGA) scheme can find notably better solutions compared to the state-of-the-art techniques in the realm of quantum annealing.

Index Terms— Greedy Algorithms, Optimization, Quantum Annealing, Quantum Computing

1. INTRODUCTION

Artificial intelligence (AI) is a disruptive technology that has transformed many industries. Quantum computing (QC) is a transformative computing phenomenon that promises to better address challenging problems that are intractable in the realm of classical computing. Quantum artificial intelligence (QAI) and quantum machine learning (QML) appear at the intersection of machine intelligence and QC, and aim to revolutionize application domains that are challenging for classical AI. Although promising, near-term quantum machines are susceptible to various sources of errors (such as decoherence and gate errors), and they do not include enough quantum bits (qubits) to accommodate quantum error correction policies. We are in the noisy intermediate-scale quantum (NISQ) era where we must perform computations in the presence of noise and infer the desirable output from erroneous outcomes [1, 2].

There are several models for the physical realization of quantum computers such as gate/circuit model, adiabatic, and

measurement-based QCs. For the case of gate model QC, variational quantum approaches (a.k.a. classical-quantum hybrid schemes) are the leading candidate for demonstrating the supremacy of QAI/QML models. Conversely, the limited programmability of adiabatic quantum processors—i.e., adiabatic quantum computers and quantum annealers are single-instruction (quantum) computing machines that can only sample from the ground state of a given Hamiltonian—casts doubt on their ability to outperform classical AI/ML models [3]. In this paper, we introduce a novel variational scheme for quantum annealers (QAs) and demonstrate that it can outperform QAs in finding the ground state of Ising Hamiltonians.

Greedy algorithms are a problem-solving paradigm that makes locally optimal choices at each stage and expects them to yield a globally optimum solution. Although most greedy algorithms fail to achieve the global optimum, they are often the best choice due to their efficiency [4]. As an example, greedy algorithms are widely used in sparse recovery applications at the cost of lower recovery accuracy (compared to convex optimization methods in compressive sensing [5]). Note that greedy algorithms can find globally optimum solutions if the problem exhibits optimal substructure.

We introduce a novel hybrid approach, called *quantum-assisted greedy algorithms* (QAGA), that leverages QAs to better select candidates in each stage of a greedy algorithm. At each stage, QAGA employs a QA to provide samples from the ground state of the problem and use these retrieved samples to estimate the probability distribution of problem variables. After fixing variables with negligible uncertainties, QAGA proceeds to the next stage, where the QA will solve a smaller problem with sparser couplings. Our experimental results using a D-Wave 2000Q quantum processor show that QAGA can find samples with remarkably lower energy values compared to the best-known enhancements in the realm of quantum annealing.

2. QAGA: A HYBRID APPROACH

Quantum annealing is a meta-heuristic for addressing combinatorial (or discrete) optimization problems that are intractable in the realm of classical computing. *Quantum annealers* (QAs) are a physical realization of the quantum annealing process that can draw samples from the ground

This research was supported by NASA grant (#NNH16ZDA001N-AIST 16-0091), NIH-NIGMS Initiative for Maximizing Student Development Grant (2 R25-GM55036), and a Google Lime scholarship. We thank the D-Wave Systems management team, namely Rene Copeland, for granting us access to the D-Wave 2000Q quantum annealer.

state of the given Hamiltonians at cryogenic temperatures. The quantum processing unit (QPU) by D-Wave Systems is a single-instruction quantum computing machinery that can only sample from the ground state of the following Ising Hamiltonian:

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

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 [6, 7, 8].

To solve a problem using QAs, we must define coefficients of the Ising Hamiltonian—i.e., \mathbf{h} and J —such that the ground state of the corresponding \mathcal{H} represents a solution of the problem of interest [9]. However, physical QAs are susceptible to various error sources, and therefore executing a quantum machine instruction (QMI) on a physical QA is not guaranteed to achieve the ground state of the corresponding \mathcal{H} [10]. Recent studies have shown that applying preprocessing and postprocessing policies can improve the fidelity of QAs. Nevertheless, such policies cannot bypass the technological barriers of physical QAs, and addressing hardware drawbacks can require device-level enhancements that may span generations of QAs [10].

We present the notion of *quantum-assisted greedy algorithms* (QAGA) to improve the fidelity of the near-term QAs. Algorithm 1 illustrates the QAGA process. Let \mathbf{z}^* be the ground state of \mathcal{H} —i.e., $\mathbf{z}^* = \arg \min_{\mathbf{z}} \mathcal{H}$. Our objective is to find a sample whose corresponding Ising energy value, shown in Eq. (1), approaches the energy value of the ground state of the given Ising Hamiltonian. More specifically, we aim to find $\tilde{\mathbf{z}}$ such that:

$$|\mathcal{H}_{\mathbf{z}^*} - \mathcal{H}_{\tilde{\mathbf{z}}}| \rightarrow 0.$$

QAGA starts with $\tilde{\mathbf{z}} = \{\}$ and $\mathcal{H}^{t=0} = \mathcal{H}$. At each iteration, QAGA uses a QA to draw n samples from the ground state of \mathcal{H}^t . Let Z denotes the set of all samples, drawn by a QA from the ground state of \mathcal{H}^t , as $Z = \{\mathbf{z}^1, \mathbf{z}^2, \dots, \mathbf{z}^n\}$ where $\mathbf{z}^j \in \{-1, +1\}^N$. Every sample \mathbf{z}^j contains a measurement for all qubits. Hence, we can look at each problem variable \mathbf{z}_i as a random variable with Bernoulli distribution that takes its value from $\{-1, +1\}$. After retrieving the sample set Z , we estimate the uncertainty of every problem variable \mathbf{z}_i as follows:

$$u(\mathbf{z}_i) = 1 - \frac{|\sum_{j=1}^n \mathbf{z}_i^j|}{n}. \quad (2)$$

For every variable \mathbf{z}_i of \mathcal{H}^t that $u(\mathbf{z}_i) \leq \theta$, where $\theta \in [0, 0.5)$ specifies the threshold parameter, we fix the value of optimum solution as:

$$\tilde{\mathbf{z}}_i = \text{sgn} \left(\sum_{j=1}^n \mathbf{z}_i^j \right). \quad (3)$$

since $\theta < 0.5$, $\tilde{\mathbf{z}}_i$ is guaranteed to take its value from $\{-1, +1\}$.

After fixing the value of \mathbf{z}_i , we: (a) remove \mathbf{h}_i from \mathcal{H}^{t+1} ; (b) add the value of $\tilde{\mathbf{z}}_i J_{ij}$ (or $\tilde{\mathbf{z}}_i J_{ji}$) to \mathbf{h}_j , for $j = 1, 2, \dots, N$ and $i \neq j$; and (c) remove the coupler J_{ij} (or J_{ji}) from \mathcal{H}^{t+1} . QAGA terminates when $\mathcal{H}^t \equiv \mathcal{H}^{t+1}$. If QAGA ends without fixing all problem variables, we apply the multi-qubit correction (MQC) method [10] on the samples from the last QAGA stage and assign values for the remaining problem variables.

Contracting variables with negligible uncertainties results in a new Ising Hamiltonian (\mathcal{H}^{t+1}), which is smaller and sparser compared to \mathcal{H}^t . Hence, at each iteration of QAGA, the remaining Ising Hamiltonian becomes easier to solve with physical QAs. Note that contracting variables reduces the number of qubits (N) correspondingly. Finally, we can apply a classical local optimization heuristic, such as the single-qubit correction (SQC) method [10], to increase the probability of finding the ground state.

```

Input:  $\mathcal{H}, \theta$ 
Output:  $\tilde{\mathbf{z}}$ 
 $\tilde{\mathbf{z}} \leftarrow \{\}$ 
 $\mathcal{H}^t \leftarrow \mathcal{H}$ 
 $\mathcal{H}^{t+1} \leftarrow \{\}$ 
while  $\mathcal{H}^{t+1} \neq \mathcal{H}^t$  do
     $\mathcal{H}^{t+1} \leftarrow \mathcal{H}^t$ 
     $Z \leftarrow \{\mathbf{z}^1, \mathbf{z}^2, \dots, \mathbf{z}^n\} = \arg \min_{\mathbf{z}} \mathcal{H}^t$ 
    for  $i \leftarrow 0$  to  $N$  do
        if  $u(\mathbf{z}_i) \leq \theta$  then
             $\tilde{\mathbf{z}}_i \leftarrow \text{sgn} \left( \sum_{j=1}^n \mathbf{z}_i^j \right)$ 
             $\mathcal{H}_{\mathbf{h}_i}^{t+1} \leftarrow \{\}$ 
            for  $j \leftarrow 0$  to  $N$  do
                if  $J_{ij} \in \mathcal{H}^{t+1}$  then
                     $\mathcal{H}_{\mathbf{h}_j}^{t+1} \leftarrow \mathcal{H}_{\mathbf{h}_j}^{t+1} + J_{ij} \tilde{\mathbf{z}}_i$ 
                     $\mathcal{H}_{J_{ij}}^{t+1} \leftarrow \{\}$ 
                end
                if  $J_{ji} \in \mathcal{H}^{t+1}$  then
                     $\mathcal{H}_{\mathbf{h}_j}^{t+1} \leftarrow \mathcal{H}_{\mathbf{h}_j}^{t+1} + J_{ji} \tilde{\mathbf{z}}_i$ 
                     $\mathcal{H}_{J_{ji}}^{t+1} \leftarrow \{\}$ 
                end
            end
        end
    end
end
if  $|\tilde{\mathbf{z}}| < N$  then
     $\hat{\mathbf{z}} \leftarrow \text{MQC}(Z)$ 
     $\tilde{\mathbf{z}} \leftarrow \hat{\mathbf{z}} \cup \tilde{\mathbf{z}}$ 
end
return  $\tilde{\mathbf{z}}$ 

```

Algorithm 1: Quantum-assisted greedy algorithm (QAGA) for minimizing an Ising Hamiltonian

3. RESULTS

We use randomly generated Ising Hamiltonians for evaluating the performance of the proposed QAGA and compare it with MQC, which is the state-of-the-art technique in the realm of QA [10]. We employ three different types of benchmark problems: (a) coefficients drawn uniformly from $\{-1, +1\}$ (binary coefficients); (b) coefficients drawn uniformly from $[-1, +1]$ (uniform coefficients); and (c) coefficients drawn from the standard normal distribution (normal coefficients). For our evaluations, we use a D-Wave 2000Q QA. Since randomly generated problems are not compatible with the working graph of QAs, we use the minor-embedding heuristic [11] for embedding the arbitrary random graphs to the Chimera topology of the D-Wave 2000Q quantum processors. QAGA iteratively contracts variables whose uncertainties are negligible; thus, the Ising Hamiltonian in QAGA has a dynamic structure, and we need to apply the embedding in all iterations of QAGA. Thus, QAGA needs to embed the remaining Ising Hamiltonian to an executable QMI on the target QA.

We compare the performance of QAGA to QA with ten spin-reversal transforms and longer inter-sample delay (denoted by QA), which is a notably stronger baseline. As our second baseline, we apply MQC to the result of the first baseline (denoted by MQC). For each QMI, we request 1,000 samples for all methods.

3.1. Experiment A

For every problem in this experiment, we generate a random graph of size 50 with a specified sparsity rate ($s \in \{0.05, 0.25, 0.5, 0.75, 1.0\}$). More specifically, we randomly select edges from a complete graph with $N = 50$ nodes where the sparsity rate s denotes the probability of selecting edges. We then set values of the corresponding biases and couplers randomly. The uncertainty threshold in QAGA is $\theta = 0.0$ —i.e., all spins must have the same value so QAGA can fix them for the next stage. We obtained this threshold empirically by evaluating the performance of QAGA on a small problem set. We study the sensitivity of QAGA to θ in the following experiment.

Figure 1 illustrates the performance comparisons between QAGA and QA in minimizing Ising Hamiltonians with binary, uniform and normal coefficients. For each case, we generate 100 random problems. In this experiment, all randomly generated Ising Hamiltonians include 50 spin variables ($N = 50$). For each sparsity rate, the corresponding column illustrates the number of times that QA has found a better sample (compared to QAGA), the number of times that QA and QAGA demonstrated the same performance (i.e., best samples from both methods had identical Ising energies), and the number of times that QAGA has outperformed QA.

In the same manner, Fig. 2 illustrates the performance comparisons between QAGA and MQC in minimizing the same 100 randomly generated Ising Hamiltonians with bi-

nary, uniform and normal coefficients. Note that all arrangements in this experiment (e.g., number of variables, sparsity rate, number of spin-reversal-transforms, etc.) were identical to the previous experiment. Similar to Fig. 1, each column in Fig. 2 represents the number of times that MQC has found a sample with lower energy, MQC and QAGA had similar performance, and QAGA resulted in a sample with a lower Ising energy value.

3.2. Experiment B

In this experiment, we aim to study the impact of the threshold parameter on the performance of QAGA. To this end, we measure the average number of stages (iterations) that QAGA takes to converge. Table 1 illustrates the average number of iterations that QAGA takes to solve 100 random benchmark problems with $N = 50$ variables and normal coefficients.

Table 1. Average number of iterations for QAGA in solving 100 random benchmark problems with different thresholds

θ	sparsity rate (s)				
	0.05	0.25	0.50	0.75	1.00
0.25	3.25	2.80	3.15	3.40	3.30
0.15	3.60	3.40	3.35	3.45	3.65
0.05	4.10	4.45	4.30	3.20	4.05
0.00	5.50	6.20	2.10	2.05	2.30

Table 1 reveals that for Chimera-like problems (i.e., sparse problems where $s = 0, 05$ or 0.25), when $\theta \rightarrow 0$, QAGA takes more iterations (i.e., slower convergence). On the other hand, for dense Ising Hamiltonians, where $s \rightarrow 1$ (i.e., clique like problems), when $\theta \rightarrow 0$, QAGA converges quickly—the maximum number of iterations appears on $\theta \sim 0.1$.

4. CONCLUSION

Quantum annealers are a type of adiabatic quantum computer that can sample from the ground state of Hamiltonians. Unfortunately, several technological barriers preclude physical QAs from attaining the ground state of the given problem Hamiltonians. We introduce the notion of *quantum-assisted greedy algorithms* (QAGA) that employs QAs for making globally optimum choices at each stage of a greedy algorithm. QAGA views QAs as a physical process that naturally draws samples from the ground state of Ising Hamiltonians (i.e., a problem-dependent Boltzmann distribution) at cryogenic temperatures. Combining QAs and greedy algorithms addresses the limitations of both and results in remarkably better solutions, albeit executing multiple QMIs for one problem.

Our empirical results on a D-Wave 2000Q quantum processor demonstrate that QAGA finds samples with remarkably lower energy values compared to the multi-qubit correc-

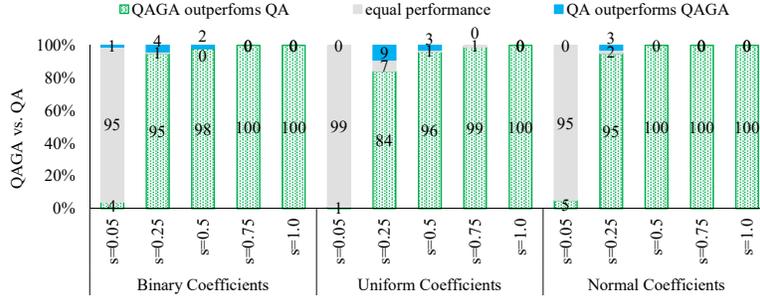


Fig. 1. Performance comparison between QAGA and QA (with spin-reversal-transforms and inter-sample delays) in solving 100 random benchmark problems with $N = 50$ spin variables and different sparsity rates s

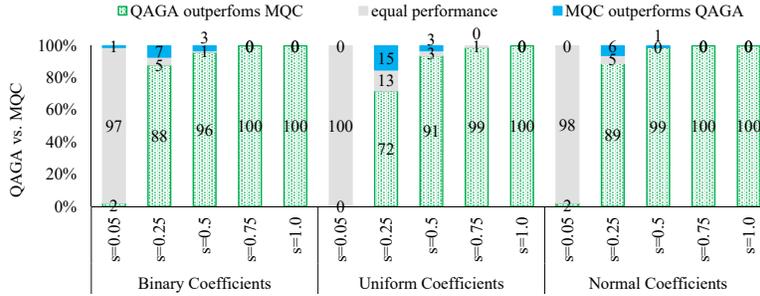


Fig. 2. Performance comparison between QAGA and MQC in solving 100 random benchmark problems with $N = 50$ spin variables and different sparsity rates s

tion (MQC) method that is the state-of-the-art technique in the realm of QA. For sparse problems (i.e., the structure of the problem is close to the Chimera architecture), the performance of the QAGA approaches to MQC. When the sparsity decreases, however, QAGA shows supremacy in terms of finding samples with lower energy values.

5. REFERENCES

- [1] John Preskill, “Quantum computing in the nisq era and beyond,” *Quantum*, vol. 2, pp. 79, 2018.
- [2] Hsin-Yuan Huang, Richard Kueng, and John Preskill, “Information-theoretic bounds on quantum advantage in machine learning,” *Physical Review Letters*, vol. 126, no. 19, pp. 190505, 2021.
- [3] Michael A Nielsen and Isaac L Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2010.
- [4] Ronald A DeVore and Vladimir N Temlyakov, “Some remarks on greedy algorithms,” *Advances in computational Mathematics*, vol. 5, no. 1, pp. 173–187, 1996.
- [5] Seyedahmad Mousavi, Mohammad Mehdi Rezaee Taghiabadi, and Ramin Ayanzadeh, “A survey on compressive sensing: Classical results and recent advancements,” *arXiv preprint arXiv:1908.01014*, 2019.
- [6] Mark W Johnson, Mohammad HS Amin, Suzanne Gildert, Trevor Lanting, Firas Hamze, Neil Dickson, R Harris, Andrew J Berkley, Jan Johansson, Paul Bunyk, et al., “Quantum annealing with manufactured spins,” *Nature*, vol. 473, no. 7346, pp. 194, 2011.
- [7] Catherine C McGeoch, “Theory versus practice in annealing-based quantum computing,” *Theoretical Computer Science*, vol. 816, pp. 169–183, 2020.
- [8] Ramin Ayanzadeh, *Leveraging Artificial Intelligence to Advance Problem-Solving with Quantum Annealers*, Ph.D. thesis, U. of Maryland, Baltimore County, 2020.
- [9] Ramin Ayanzadeh, Milton Halem, and Tim Finin, “Reinforcement quantum annealing: A hybrid quantum learning automata,” *Scientific Reports*, vol. 10, no. 1, pp. 1–11, 2020.
- [10] Ramin Ayanzadeh, John Dorband, Milton Halem, and Tim Finin, “Multi-qubit correction for quantum annealers,” *Scientific Reports*, vol. 11, no. 1, pp. 1–12, 2021.
- [11] Jun Cai, William G Macready, and Aidan Roy, “A practical heuristic for finding graph minors,” *arXiv preprint arXiv:1406.2741*, 2014.