

**LA-UR-22-33077**

**Approved for public release; distribution is unlimited.**

**Title:** Quantum Annealing vs. QAOA: 127 Qubit Higher-Order Ising Problems on NISQ Computers

**Author(s):** Pelofske, Elijah  
Baertschi, Andreas  
Eidenbenz, Stephan Johannes

**Intended for:** ISC HPC 2023, 2023-05-21/2022-12-25 (Hamburg, Germany)

**Issued:** 2024-05-21 (rev.2)



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA00001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# Quantum Annealing vs. QAOA: 127 Qubit Higher-Order Ising Problems on NISQ Computers

Elijah Pelofske<sup>\*1</sup>, Andreas Bärtschi<sup>†1</sup>, and Stephan Eidenbenz<sup>1</sup>

<sup>1</sup>CCS-3 Information Sciences, Los Alamos National Laboratory

## Abstract

Quantum annealing (QA) and Quantum Alternating Operator Ansatz (QAOA) are both heuristic quantum algorithms intended for sampling optimal solutions of combinatorial optimization problems. In this article we implement a rigorous direct comparison between QA on D-Wave hardware and QAOA on IBMQ hardware. These two quantum algorithms are also compared against classical simulated annealing. The studied problems are instances of a class of Ising models, with variable assignments of  $+1$  or  $-1$ , that contain cubic  $ZZZ$  interactions (higher order terms) and match both the native connectivity of the Pegasus topology D-Wave chips and the heavy hexagonal lattice of the IBMQ chips. The novel QAOA implementation on the heavy hexagonal lattice has a CNOT depth of 6 per round and allows for usage of an entire heavy hexagonal lattice. Experimentally, QAOA is executed on an ensemble of randomly generated Ising instances with a grid search over 1 and 2 round angles using all 127 programmable superconducting transmon qubits of `ibm_washington`. The error suppression technique digital dynamical decoupling is also tested on all QAOA circuits. QA is executed on the same Ising instances with the programmable superconducting flux qubit devices D-Wave `Advantage_system4.1` and `Advantage_system6.1` using modified annealing schedules with pauses. We find that QA outperforms QAOA on all problem instances. We also find that dynamical decoupling enables 2-round QAOA to marginally outperform 1-round QAOA, which is not the case without dynamical decoupling.

## 1 Introduction

Quantum annealing (QA) in the transverse field Ising model is an analog computation technology which utilizes quantum fluctuations in order to search for ground state solutions of a problem Hamiltonian [1–5]. D-Wave quantum annealers are programmable hardware implementations of quantum annealing which use superconducting flux qubits [6, 7].

Quantum Alternating Operator Ansatz (QAOA) is a hybrid quantum classical algorithm for sampling combinatorial optimization problems [8–10], the quantum component of which can be instantiated with a programmable gate-based universal quantum computer. The Quantum Approximate Optimization Algorithm [11] was the first variational algorithm of this type, which was then generalized to the Quantum Alternating Operator Ansatz algorithm [8].

QAOA is effectively a Trotterization of the Quantum Adiabatic Algorithm, and is overall similar to Quantum Annealing. In particular both algorithms address combinatorial optimization problems. The exact characteristics of how both QA and QAOA will scale to large system sizes is currently not fully understood, in particular because quantum hardware is still in the NISQ era [12–14]. For example, there is evidence that QAOA may be more difficult for classical computers to simulate than quantum annealing, which could make it a viable candidate for quantum advantage [15]. Quantum annealing in particular has been experimentally evaluated against classical algorithms in order to determine for what problem types and under what settings quantum annealing could provide a scaling advantage over the next best state-of-the-art classical approaches [13, 14, 16–18]. Generally these results are encouraging and show that quantum annealing can indeed sample certain problem types better than classical methods such as simulated annealing. There have been a number of studies that directly compare Quantum Annealing and QAOA for a number of different sampling tasks [19–23], however this paper presents, to the best of our knowledge, the largest direct comparison between Quantum Annealing and QAOA to date. There have been experimental QAOA implementations which used up to 40 qubits [24], 27 qubits [25], and 23 qubits [26]. There have also been QAOA experiments which had circuit depth up to 159 [27] and 148 [28].

The contributions of this article are as follows:

---

<sup>\*</sup>Email: [epelofske@lanl.gov](mailto:epelofske@lanl.gov)

<sup>†</sup>Email: [baertschi@lanl.gov](mailto:baertschi@lanl.gov)

Device name	Topology/chip name	Available qubits	Available couplers/ CNOTs	Computation type
Advantage_system4.1	Pegasus $P_{16}$	5627	40279	QA
Advantage_system6.1	Pegasus $P_{16}$	5616	40135	QA
ibm_washington	Eagle r1 heavy-hexagonal	127	142	Universal gate-model

Table 1: NISQ hardware summary at the time the experiments were executed. The hardware yield (e.g., the number of available qubits or two qubit interactions) for all of these devices can be less than the logical lattice because of hardware defects, and can also change over time if device calibration changes.

1. We provide a direct comparison between QAOA and Quantum Annealing in terms of experiments on D-Wave and IBMQ hardware. This comparison uses a comparable parameter search space for QA and QAOA, uses no minor embedding for quantum annealing, and uses short depth QAOA circuits, thus providing a fair comparison of the two algorithms. A comparison of this problem size has not been performed before to the best of our knowledge. We show that QAOA is better than random sampling, and quantum annealing clearly outperforms QAOA. A comparison against the classical heuristic algorithm simulated annealing is also presented.
2. The QAOA algorithm we present is tailored for short depth circuit construction on the heavy hexagonal lattice (CNOT depth of 6 per round), therefore allowing full usage of any heavy hexagonal topology quantum processor in the future. We use all 127 qubits of the `ibm_washington` chip in order to execute the largest QAOA circuit, in terms of qubits, to date. Each QAOA circuit uses thousands of gate operations, making these results one of the largest quantum computing experiments performed to date.
3. The Ising models that are used to compare quantum annealing and QAOA are specifically constructed to include higher order terms, specifically three variable (cubic) terms. QAOA can directly implement higher order terms, and quantum annealing requires order reduction using auxiliary variables to implement these higher order terms. This is the largest experimental demonstration of QAOA with higher order terms to date.
4. In order to mitigate errors when executing the QAOA circuits, we utilize digital dynamical decoupling. This is the largest usage of dynamical decoupling in terms of qubit system size to date, and the results show that digital dynamical decoupling improves performance for two round QAOA, suggesting that it will be useful for computations with large numbers of qubits in the noisy regime.

In Section 2 the QAOA and QA hardware implementations, and the simulated annealing implementation are detailed. Section 3 details the experimental results and how the two quantum algorithms compare, including how simulated annealing compares. Section 4 concludes with what the results indicate and future research directions. The figures in this article are generated using matplotlib [29, 30], and Qiskit [31] in Python 3. Code, data, and additional figures are available in a public Github repository <sup>1</sup>.

## 2 Methods

The Ising models are defined in Section 2.1. In Section 2.2 the QAOA circuit algorithm and hardware parameters are defined. In Section 2.3 the quantum annealing implementation is defined. Section 2.4 defines the simulated annealing implementation.

### 2.1 Ising Model Problem Instances

The NISQ computers which are used in this comparison are detailed in Table 1; the clear difference between the D-Wave quantum annealers and `ibm_washington` is the number of qubits that are available. The additional qubits available on the quantum annealers will allow us to embed multiple problem instances onto the chips. The current IBMQ devices have a graph topology referred to as the heavy-hexagonal lattice [32]. Therefore, for a direct QAOA

<sup>1</sup>[https://github.com/lanl/QAOA\\_vs\\_QA](https://github.com/lanl/QAOA_vs_QA)

and QA comparison we would want to be able to create QAOA circuits which match the logical heavy-hexagonal lattice and the quantum annealer graph topology of Pegasus. For this direct comparison we target D-Wave quantum annealers with Pegasus graph hardware [33, 34] connectivities. The two current D-Wave quantum annealers with Pegasus hardware graphs have chip id names `Advantage_system6.1` and `Advantage_system4.1`. The goal for this direct comparison is that ideally we want problems which can be instantiated on *all* three of the devices in Table 1. In particular, we want these implementations to not be unfairly costly in terms of implementation overhead. For example we do not want to introduce unnecessary qubit swapping in the QAOA circuit because that would introduce larger circuit depths which would introduce more decoherence in the computation. We also do not want to introduce unnecessary minor-embedding in the problems for quantum annealers.

The other property of these problem instances that is of interest is an introduction of *higher order terms*, specifically cubic  $ZZZ$  interactions [35] also referred to as multi-body interactions [36], in addition to random linear and quadratic terms. These higher order terms require both QAOA and QA to be handle these higher order variable interactions, which is an additional test on the capability of both algorithms. QAOA can naturally handle higher order terms [37]. Implementing high order terms with QA requires introducing auxiliary variables in order to perform order reduction to get a problem structure that is comprised of only linear and quadratic terms, so that it can be implemented on the hardware, but whose optimal solutions match the optimal solutions of the original high order polynomial (for the non-auxiliary variables) [4, 38–41].

Taking each of these characteristics into account, we create a class of random problems which follow the native device connectivities in Table 1. The problem instances we will be considering are Ising models defined on the hardware connectivity graph of the heavy hexagonal lattice of the device, which for these experiments will be `ibm_washington`. For a variable assignment vector  $z = (z_0, \dots, z_{n-1}) \in \{+1, -1\}^n$ , the random Ising model is defined as

$$C(z) = \sum_{v \in V} d_v \cdot z_v + \sum_{(i,j) \in E} d_{i,j} \cdot z_i \cdot z_j + \sum_{l \in W} d_{l,n_1(l),n_2(l)} \cdot z_l \cdot z_{n_1(l)} \cdot z_{n_2(l)} \quad (1)$$

Eq. (1) defines the class of random minimization Ising models with cubic terms as follows. Any heavy hexagonal lattice is a bipartite graph with vertices  $V = \{0, \dots, n-1\}$  partitioned as  $V = V_2 \cup V_3$ , where  $V_3$  consists of vertices with a maximum degree of 3, and  $V_2$  consists of vertices with a maximum degree of 2.  $E \subset V_2 \times V_3$  is the edge set representing available two qubit gates (in this case CNOTs where we choose targets  $i \in V_2$  and controls  $j \in V_3$ ).  $W$  is the set of vertices in  $V_2$  that all have degree exactly equal to 2.  $n_1$  is a function that gives the qubit (variable) index of the first of the two neighbors of a degree-2 node and  $n_2$  provides the qubit (variable) index of the second of the two neighbors of any degree-2 node. Thus  $d_v$ ,  $d_{i,j}$ , and  $d_{l,n_1(l),n_2(l)}$  are all coefficients representing the random selection of the linear, quadratic, and cubic coefficients, respectively. These coefficients could be drawn from any distribution - in this paper we draw the coefficients from  $\{+1, -1\}$  with probability 0.5. Eq. (1) therefore defines how to compute the objective function for a given variable assignment vector  $z$ .

The heavy hexagonal topology of `ibm_washington`, along with an overlay showing one of the random problem instances with cubic terms defined on `ibm_washington`, is shown in Figure 1. Each term coefficient was chosen to be either  $+1$  or  $-1$ , which in part helps to mitigate the potential problem of limited precision for the programming control on all of the NISQ devices. 10 random instances of this class of problems are generated and sampled using QAOA and QA, the implementations of each will be discussed next.

## 2.2 Quantum Alternating Operator Ansatz

Given a combinatorial optimization problem over inputs  $z \in \{+1, -1\}^n$ , let  $C(z): \{+1, -1\}^n \rightarrow \mathbb{R}$  be the objective function which evaluates the cost of the solution vector  $z$ . For a maximization (or minimization) problem, the goal is to find a variable assignment vector  $z$  for which  $f(z)$  is maximized (or minimized). The QAOA algorithm consists of the following components:

- an initial state  $|\psi\rangle$ ,
- a **phase separating** Cost Hamiltonian  $H_C$ ,  
which is derived from  $C(z)$  by replacing all spin variables  $z_i$  by Pauli-Z operators  $\sigma_i^z$
- a **mixing** Hamiltonian  $H_M$ ; in our case, we use the standard transverse field mixer, which is the sum of the Pauli-X operators  $\sigma_i^x$
- an integer  $p \geq 1$ , the number of rounds to run the algorithm,

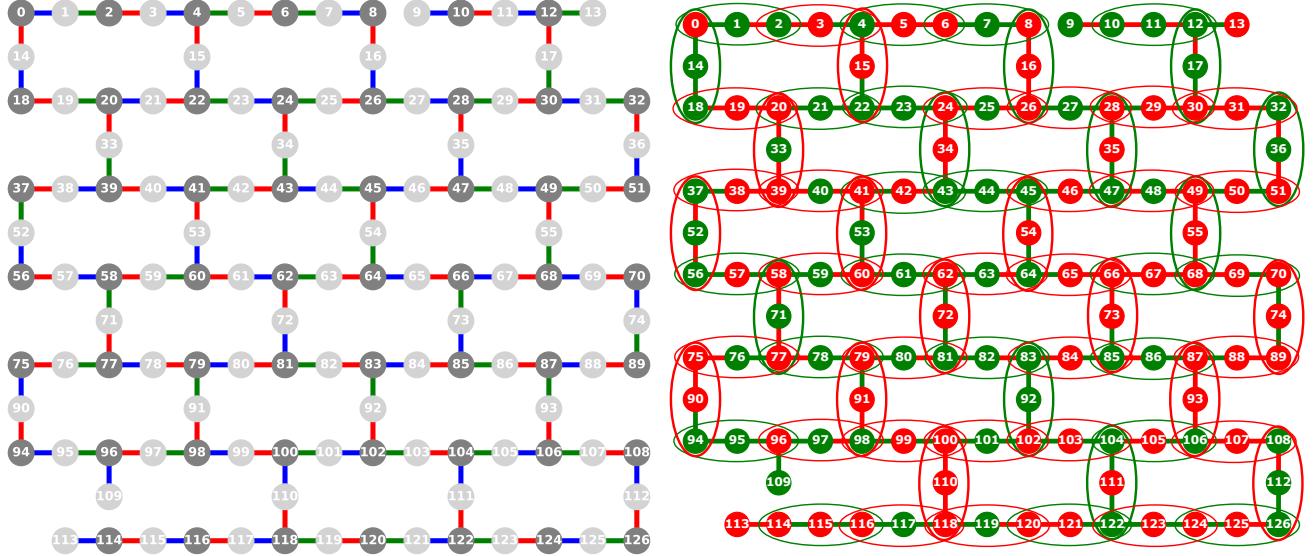


Figure 1: **Left:** `ibm_washington` graph connectivity, where qubits are connected by CNOT (also referred to as `cx`) gates. The ideal lattice is called the heavy-hexagonal lattice. Note that there are two missing graph edges from the lattice between qubits 8-9 and 109-114. The total number of qubits (nodes) is 127. The edges of the graph are three colored (red, blue, and green) such that no node shares two or more edges with the same color. The node colorings of light and dark gray show that the heavy hexagonal lattice is bipartite (meaning it can be partitioned into two disjoint sets). The three edge coloring is consistent with the QAOA circuit construction in Figure 2. **Right:** Example of a single random problem instance with cubic terms (see Eq. (1)) on the `ibm_washington` graph. The linear and quadratic terms are shown using two distinct colors (red and green). The nodes and edges colored red denote a weight of  $-1$  and the nodes and edges colored green denote a weight of  $+1$ . The cubic terms are represented by ovals around the three qubits which define the cubic variable interactions. Like the linear and quadratic terms, the color of the oval representing the cubic terms represents the sign of the weight on the terms, where green is  $+1$  and red is  $-1$ .

- two real vectors  $\vec{\gamma} = (\gamma_1, \dots, \gamma_p)$  and  $\vec{\beta} = (\beta_1, \dots, \beta_p)$ , each with length  $p$ .

The algorithm consists of preparing the initial state  $|\psi\rangle$ , then applying  $p$  rounds of the alternating simulation of the phase separating Hamiltonian and the mixing Hamiltonian:

$$|\vec{\gamma}, \vec{\beta}\rangle = \underbrace{e^{-i\beta_p H_M} e^{-i\gamma_p H_P} \cdots e^{-i\beta_1 H_M} e^{-i\gamma_1 H_P}}_{\text{round } p} |\psi\rangle \quad (2)$$

Within each round,  $H_P$  is applied first, which separates the basis states of the state vector by phases  $e^{-i\gamma f(x)}$ .  $H_M$  then provides parameterized interference between solutions of different cost values. After  $p$  rounds, the state  $|\vec{\gamma}, \vec{\beta}\rangle$  is measured in the computational basis and returns a sample solution  $y$  of cost value  $f(y)$  with probability  $|\langle y | \vec{\gamma}, \vec{\beta} \rangle|^2$ .

The aim of QAOA is to prepare the state  $|\vec{\gamma}, \vec{\beta}\rangle$  from which we can sample a solution  $y$  with high cost value  $f(y)$ . Therefore, in order to use QAOA the task is to find angles  $\vec{\gamma}$  and  $\vec{\beta}$  such that the expectation value  $\langle \vec{\gamma}, \vec{\beta} | H_P | \vec{\gamma}, \vec{\beta} \rangle$  is large ( $-H_P$  for minimization problems). In the limit  $p \rightarrow \infty$ , QAOA is effectively a Trotterization of the Quantum Adiabatic Algorithm, and in general as we increase  $p$  we expect to see a corresponding increase in the probability of sampling the optimal solution [42]. The challenge is the classical outer loop component of finding the good angles  $\vec{\gamma}$  and  $\vec{\beta}$  for all rounds  $p$ , which has a high computational cost as  $p$  increases.

Variational quantum algorithms, such as QAOA, have been a subject of large amount of attention, in large part because of the problem domains that variational algorithms can address (such as combinatorial optimization) [43]. One of the challenges however with variational quantum algorithms is that the classical component of parameter selection, in the case of QAOA this is the angle finding problem, is not solved and is even more difficult when noise is present in the computation [44]. Typically the optimal angles for QAOA are computed exactly for small problem instances [20, 45]. However, in this case the angle finding approach we will use is a reasonably high resolution

gridsearch over the possible angles. Note however that a fine gridsearch scales exponentially with the number of QAOA rounds  $p$ , and therefore is not advisable for practical high round QAOA [9, 11]. Exactly computing what the optimal angles are for problems of this size would be quite computationally intensive, especially with the introduction of higher order terms. We leave the problem of exactly computing the optimal QAOA angles to future work.

Figure 2 describes the short depth QAOA circuit construction for sampling the higher order Ising test instance. This algorithm can be applied to any heavy hexagonal lattice topology, which allows for executing the QAOA circuits on the 127 variable instances on the IBMQ `ibm_washington` backend. For the class of Ising models with higher order terms defined in Section 2.1, the QAOA angle ranges which are used are  $\gamma_1, \dots, \gamma_p \in [0, \pi)$  and  $\beta_1, \dots, \beta_{p-1} \in [0, \pi), \beta_p \in [0, \frac{\pi}{2}]$  where  $p$  is the number of QAOA rounds. Note that the halving of the angle search space for  $\beta$  applies when  $p = 1$ . For optimizing the angles using the naive grid search for  $p = 1$ ,  $\beta_0$  is varied over 60 linearly spaced angles  $\in [0, \frac{\pi}{2}]$  and  $\gamma_0$  is varied over 120 linearly spaced angles  $\in [0, \pi]$ . For the high resolution gridsearch for  $p = 2$ ,  $\beta_1$  is varied over 5 linearly spaced angles  $\in [0, \frac{\pi}{2}]$  and  $\gamma_0, \gamma_1$ , and  $\beta_0$  are varied over 11 linearly spaced angles  $\in [0, \pi]$ . Therefore, for  $p = 2$  the angle gridsearch uses 6655 separate circuit executions (for each of the 10 problem instances), and for  $p = 1$  the angle gridsearch uses 7200 separate circuit executions. Each circuit execution used 10,000 samples in order to compute a robust distribution for each angle combination.

In order to mitigate decoherence on idle qubits, digital dynamical decoupling (DDD) is also tested for all QAOA circuits. Dynamical Decoupling is an open loop quantum control technique error suppression technique for mitigating decoherence on idle qubits [46–51]. Dynamical decoupling can be implemented with pulse level quantum control, and digital dynamical decoupling can be implemented simply with circuit level instructions of sequences of gates which are identities [50]. Note that digital dynamical decoupling is an approximation of pulse level dynamical decoupling. Dynamical decoupling has been experimentally demonstrated for superconducting qubit quantum processors including IBMQ devices [46, 52, 53]. Dynamical decoupling in particular is applicable for QAOA circuits because they can be relatively sparse and therefore have idle qubits [46]. DDD does not always effective at consistently reducing errors during computation (for example because of other control errors present on the device [46, 49]), and therefore the raw QAOA circuits are compared against the QAOA circuits with DDD in the experiments section. In order to apply the DDD sequences to the OpenQASM [54] QAOA circuits, the `PadDynamicalDecoupling`<sup>2</sup> method from Qiskit [31] is used, with the `pulse_alignment` parameter set based on the `ibm_washington` backend properties. The circuit scheduling algorithm that is used for inserting the digital dynamical decoupling sequences is ALAP, which schedules the stop time of instructions as late as possible<sup>3</sup>. There are other scheduling algorithms that could be applied which may increase the efficacy of dynamical decoupling. There are different DDD gate sequences that can be applied, including Y-Y or X-X sequences. Because the X Pauli gate is already a native gate of the IBMQ device, the X-X DDD sequence is used for simplicity.

Note that the variable states for the optimization problems are either  $-1$  or  $+1$ , but the circuit measurement states are either  $0$  or  $1$ . Therefore once the measurements are made on the QAOA circuits, for each variable in each sample the variable state mapping of  $0 \rightarrow 1, 1 \rightarrow -1$  is performed. For circuit execution on the superconducting transom qubit `ibm_washington`, circuits are batched into *jobs* where each job is composed of a group of at most 250 circuits - the maximum number of circuits for a job on `ibm_washington` is currently 300, but we use 250 in order to reduce job errors related to the size of jobs. Grouping circuits into jobs is helpful for reducing the total amount of compute time required to prepare and measure each circuit. When submitting the circuits to the backend, they are all first locally transpiled via Qiskit [31] with `optimization_level=3`. This transpilation converts the gateset to the `ibm_washington` native gateset, and the transpiler optimization attempts to simplify the circuit where possible. The QAOA circuit execution on `ibm_washington` spanned a large amount of time, and therefore the backend versions were not consistent. The exact backend software versions were 1.3.7, 1.3.8, 1.3.13, 1.3.15, 1.3.17.

### 2.3 Quantum Annealing

Quantum annealing is a proposed type of quantum computation which uses quantum fluctuations, such as quantum tunneling, in order to search for the ground state of a user programmed Hamiltonian. Quantum annealing, in the case of the transverse field Ising model implemented on D-Wave hardware, is explicitly described by the system given in Eq. (3). The state begins at time zero purely in the transverse Hamiltonian state  $\sum_i \sigma_i^x$ , and then over the course of the anneal (parameterized by the *annealing time*) the user programmed Ising is applied according the function  $B(s)$ . Together,  $A(s)$  and  $B(s)$  define the anneal schedules of the annealing process, and  $s$  is referred to

<sup>2</sup>[https://qiskit.org/documentation/locale/bn\\_BN/stubs/qiskit.transpiler.passes.PadDynamicalDecoupling.html](https://qiskit.org/documentation/locale/bn_BN/stubs/qiskit.transpiler.passes.PadDynamicalDecoupling.html)

<sup>3</sup>[https://qiskit.org/documentation/apidoc/transpiler\\_passes.html](https://qiskit.org/documentation/apidoc/transpiler_passes.html)

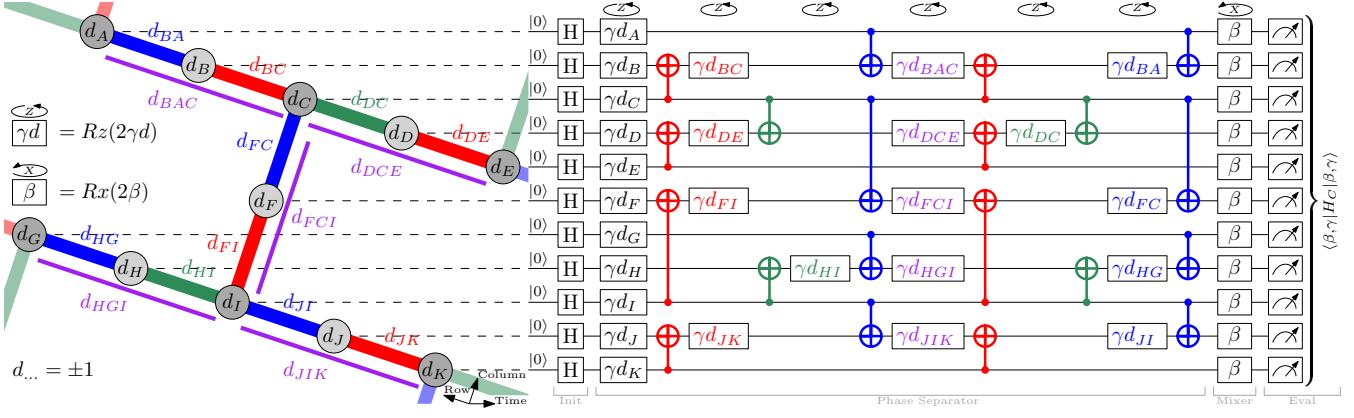


Figure 2: A 1-round QAOA circuit: (left) The problem instance is a hardware-native bipartite graph with an arbitrary 3-edge-coloring given by König’s line coloring theorem. (right) Any quadratic term (colored edge) gives rise to a combination of two CNOTs and a Rz-rotation in the phase separator, giving a CNOT depth of 6 due to the degree-3 nodes. When targeting the degree-2 nodes with the CNOT gates, these constructions can be nested, leading to no overhead when implementing the three-qubit terms: these always have a degree-2 node in the middle (see Eq. (1)).

as the *anneal fraction*. The standard anneal schedule that is used is a linear interpolation between  $s = 0$  and  $s = 1$ .

$$H = -\frac{A(s)}{2} \left( \sum_i^n \sigma_i^x \right) + \frac{B(s)}{2} \left( H_{Ising} \right) \quad (3)$$

The adiabatic theorem states that if changes to the Hamiltonian of the system are sufficiently slow, the system will remain in the ground state of problem Hamiltonian, thereby providing a computational mechanism for computing the ground state of optimization problems. The user programmed Ising  $H_{Ising}$ , acting on  $n$  qubits, is defined in Eq. (4). The quadratic terms and the linear terms combined define the optimization problem instance that the annealing procedure will ideally find the ground state of. As with QAOA, the objective of quantum annealing is to find the variable assignment vector  $z$  that minimizes the cost function which has the form of Eq. (4).

$$H_{Ising} = \sum_i^n h_i \sigma_i^z + \sum_{i < j}^n J_{ij} \sigma_i^z \sigma_j^z \quad (4)$$

The goal is to be able to implement the Ising models defined in Section 2.1 on D-Wave quantum annealers. In order to implement the higher order terms, we will need to use order reduction in order to transform the cubic terms into linear and quadratic terms [4, 38–41]. This order reduction will result in using additional variables, usually called *auxiliary* or *slack* variables. Figure 3 shows the embeddings of the problem instances onto the logical Pegasus  $P_{16}$  graph, including the order reduction procedure which is used. The order reduction procedure outlined in Figure 3 allows for direct embedding of the order reduced polynomials onto the hardware graph, regardless of whether the cubic term coefficient is  $+1$  or  $-1$ . This order reduction ensures that the ground state(s) of the cubic term are also the ground states of the order reduced Ising. Additionally, this order reduction ensures that for every excited state of the cubic term, there are no slack variable assignments which result in the original variables having an energy less than or equal to the ground state of the original cubic term. This order reduction procedure allows any problem in the form of Eq. (1) to be mapped natively to quantum annealing hardware which accepts problems with the form of Eq. (4). Importantly, this procedure does not require minor-embedding, even including the auxiliary variables.

In order to get more samples for the same QPU time, the other strategy that is employed is to embed multiple independent Ising model instances onto the hardware graph and thus be able to execute several instances in the same annealing cycle(s). This technique is referred to as *parallel quantum annealing* [40, 55] or *tiling*<sup>4</sup>. Figure 3 (right) shows the parallel embeddings on a logical Pegasus graph. Because some of the logical embeddings may use a qubit or coupler which is missing on the actual hardware, less than 6 parallel instances can be tiled onto the chips to be executed at the same time. For **Advantage\_system4.1**, 2 independent embeddings of the problem instances

<sup>4</sup><https://dwave-systemdocs.readthedocs.io/en/samplers/reference/composites/tiling.html>

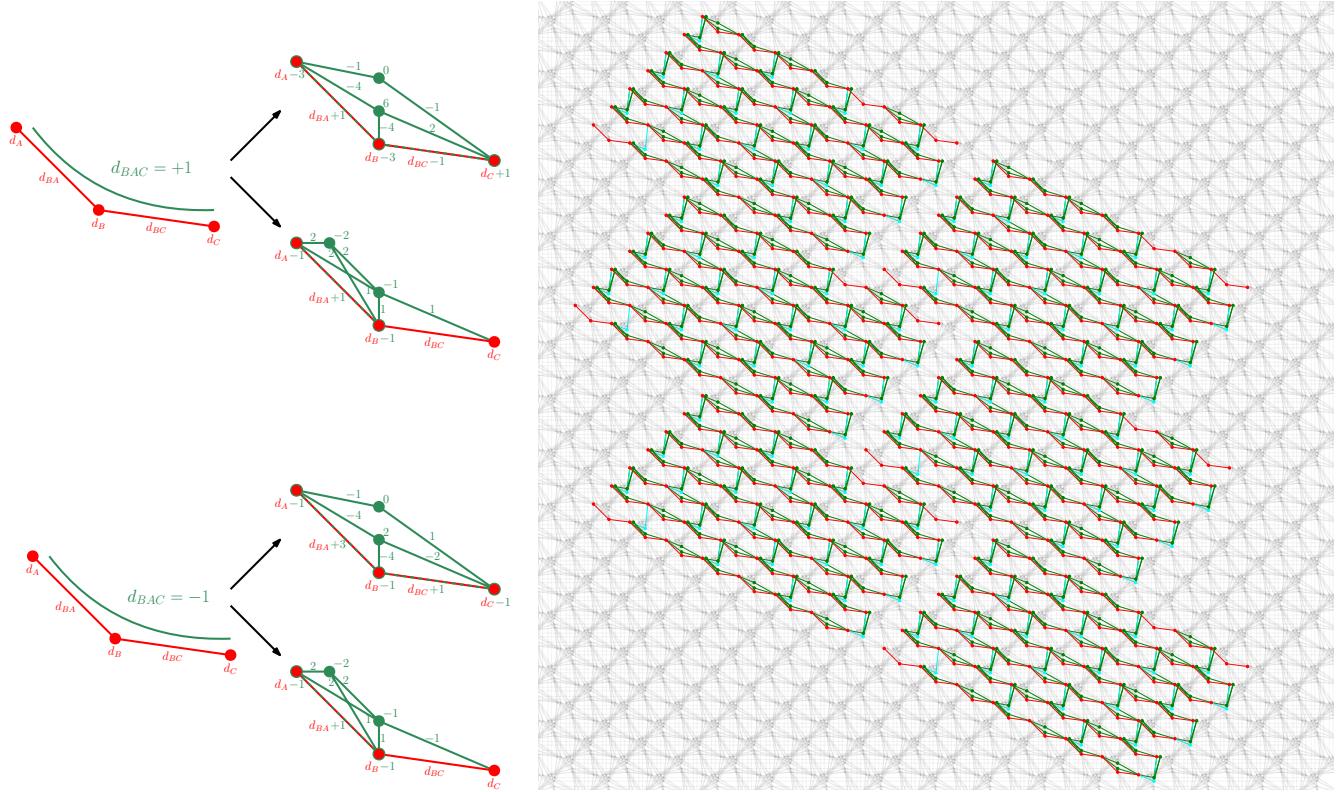


Figure 3: **(left)** Two different embeddings for cubic  $+1/-1$  terms. Each embedding needs two slack variable qubits. Our overall embedding alternates between these two cubic term embeddings. Any embedding with only one slack variable needs a 4-clique between the slack and the three original variables, which is not possible to embed for consecutive cubic terms. **(right)** Embedding structures of the problem instances with higher order terms embedded in parallel (independently) 6 times onto the logical Pegasus  $P_{16}$  graph. The view of this graph has been slightly partitioned so that not all of the outer parts of the Pegasus chip are drawn. The light grey qubits and couplers indicate unused hardware regions. The cyan coloring on nodes and edges denote the vertical qubits and CNOTs on the `ibm_washington` hardware graph (see Figure 1). The red coloring on nodes and edges denote the horizontal lines of qubits and CNOTs on `ibm_washington`. The green nodes and edges denote the order reduction auxiliary variables. Note that the top right hand and lower left hand qubits are not present on the `ibm_washington` lattice - but for the purposes of generating the embeddings, these extra qubits are filled in to complete the lattice.

could be created without encountering missing hardware. For `Advantage_system6.1`, 3 independent embeddings of the problem instances could be created. The structure of the heavy-hexagonal lattice onto Pegasus can be visually seen in Figure 3; the horizontal heavy-hex lines (Figure 1) are mapped to diagonal Pegasus qubit lines that run from top left to bottom right of the square Pegasus graph rendering. Then the vertical heavy-hexagonal qubits are mapped to QA qubits in between the diagonal qubit lines.

In order to optimize the quantum annealing parameters, with relatively similar complexity to the angle parameter search done for QAOA, the forward anneal schedule with pausing is optimized over a gridsearch. Pausing the anneal at the appropriate spot can provide higher chances of sampling the ground state [56]. Figure 4 shows this anneal schedule search space - importantly the annealing times used in these schedule are also optimized for. The total number of QA parameters which are varied are 9 anneal fractions, 9 pause durations, and 4 annealing times (10, 100, 1000, 2000 microseconds). Therefore, the total number of parameter combinations which are considered in the grid search is 324. 2000 microseconds is the longest annealing time available on the current D-Wave quantum annealers. The number of anneals sampled for each D-Wave job was 500. The annealing times and the anneal schedules were varied in a simple grid search. Readout and programming thermalization times are both set to 0 microseconds. All other parameters are set to default, with the exception of the modified annealing schedule.

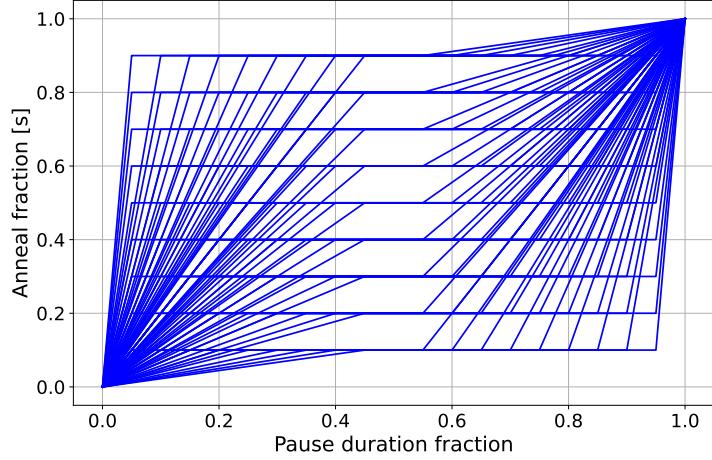


Figure 4: All modified (forward) quantum annealing schedules which are tested in order to find the best anneal schedule with a pause. The symmetric pause inserted into the normal linearly interpolated schedule defining the  $A(s)$  and  $B(s)$  functions can provide better ground state sampling probability. The anneal fraction at which this pause occurs is varied between 0.1 and 0.9 in steps of 0.1. The pause duration, as a fraction of the total annealing time, is also varied between 0.1 and 0.9 in steps of 0.1. Although not shown in this figure, the annealing times are also varied between 10, 100, 1000, and 2000 microseconds.

## 2.4 Simulated Annealing implementation

In order to provide a reasonable basis of comparison, the 10 Ising model problem instances are also sampled using simulated annealing. Simulated annealing is a standard high accuracy and general purpose classical heuristic algorithm [57], and has been used as a reasonable comparison against quantum algorithms [13]. The simulated annealing implementation that we utilize is an open source implementation<sup>5</sup>. The settings we use are all set to default and 1000 samples are drawn for each Ising model. The simulated annealing implementation does not natively handle higher order terms, and therefore order reduction must be applied to the Ising model's before being sampled by simulated annealing. Order reduction introduces additional variables into the computation. The order reduction is performed using the python package *dimod*<sup>6</sup>. The order reduction penalty **strength** is set to 2, which ensures that the optimal solution of the original higher order Ising matches the order reduced Ising model (excluding the ancillary variables introduced by the order reduction).

## 3 Results

Figures 5 and 6 combined show the detailed energy distributions for all 10 cubic Ising models sampled using the best parameter choices found for QA and QAOA. These histograms include the four variants of QAOA - 1 and 2 rounds with and without digital dynamical decoupling. The histograms include 10000 random samples (binomial distribution with  $p = 0.5$ ) on the 10 Ising models.

**QA performs better than QAOA:** The most notable observation across these histograms is that clearly quantum annealing results in better variable assignments compared to all tested variations of QAOA; this clear stratification of the algorithms capabilities is consistent across all 10 problem instances. Notice that the minimum energies achieved by QAOA (marked by the solid vertical lines) do not reach the energy distribution sampled by the quantum annealers. The characteristics of each of the 10 problem instances are slightly different, but this trend is very clear.

**QAOA performs better than random sampling:** Both QA and QAOA sampled better solutions than the 10000 random samples. Although an obvious observation from the distributions in Figures 6 and 5, it is not trivial that the QAOA samples had better objective function values compared to random sampling. The reason this is not trivial is because at sufficient circuit depth, which is not difficult to reach, the computation will entirely decohere and the computation will not be meaningful. This result is encouraging because it shows that short depth

<sup>5</sup><https://github.com/dwavesystems/dwave-neal>

<sup>6</sup><https://github.com/dwavesystems/dimod>

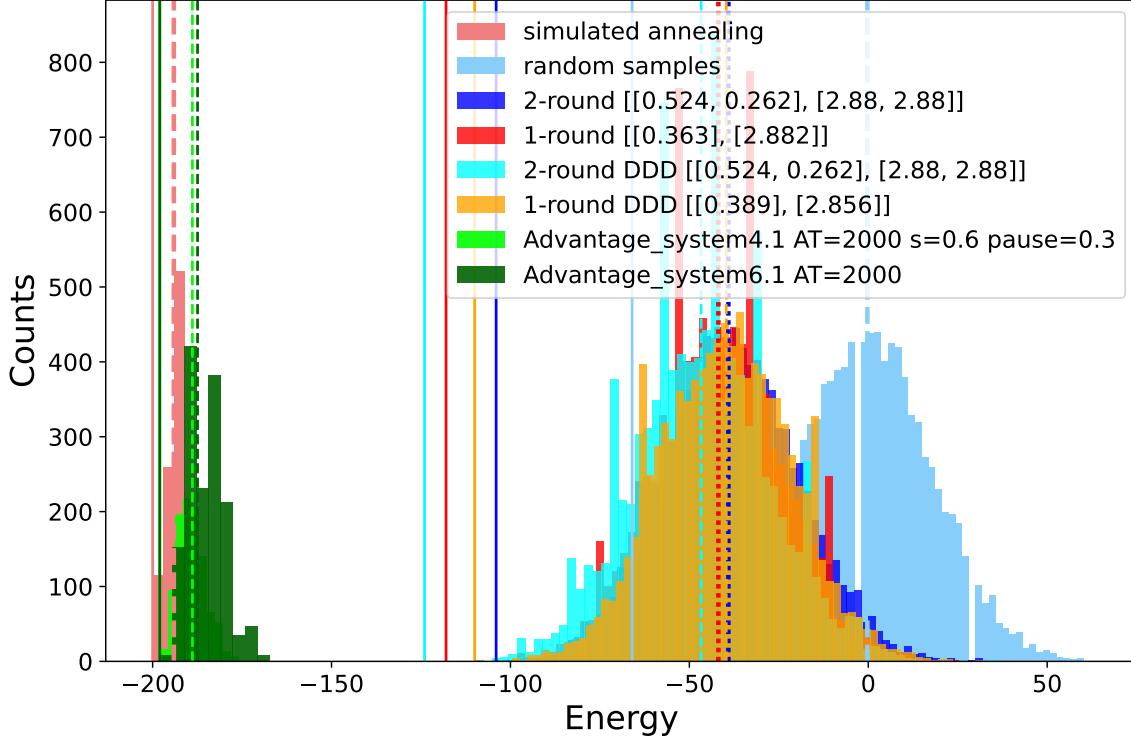


Figure 5: Direct objective function (e.g. energy) histogram comparison of QA and QAOA results for one of the 10 minimization problem instances. A distribution of simulated annealing energies are also shown to provide a comparison against a reasonable classical heuristic. Here the energies being plotted are the full energy spectrum for the parameters which gave the minimum mean energy across the parameter grid searches performed across the QA and QAOA parameters. The optimal parameter combination for each distribution is given in the figure legend. For QA parameters, the annealing time in microseconds, the forward anneal schedule (symmetric) pause fraction, and anneal fraction, are given in the legend. If the default linearly interpolated quantum annealing schedule performed the best, only the annealing time reported in the legend. For the QAOA angle parameters, the format is  $[\beta, \gamma]$ , and are rounded to 3 decimal places. The mean for each dataset is marked with vertical dashed lines and the minimum energy found in each dataset is marked with solid vertical lines. The energy histogram plots for the other 9 Ising models are shown in Figure 6.

circuit constructions, combined with increasing scale of near term quantum computers, can begin to yield relevant computations for larger system sizes (in this case, 127 variables).

**The effect of digital dynamical decoupling:** The dataset shown in Figure 6 also allows for a direct quantification of how successful the digital dynamical decoupling passes were at improving the QAOA circuit executions. Table 2 shows a comparison of the four QAOA implementations. For 2-round QAOA, DDD improved the mean sample energy for 10 out of the 10 Ising models. For 1-round QAOA, DDD improved the mean sample energy for 4 out of the 10 problem instances. This shows that digital dynamical decoupling does not uniformly improve the performance of the QAOA circuits. This suggests that the qubits in the 2-round QAOA circuits have more available idle time compared to the 1-round QAOA circuits, which would allow for DDD to improve the circuit performance. The 2-round QAOA results had better average energy compared the 1-round results in 6 out of the 10 problem instances.

**Optimal parameter choices - QAOA:** The optimal 2-round QAOA angles for all 10 problems with and without dynamical decoupling is the same. The optimal 1-round QAOA angles are not consistent across all problems, and even vary between the with and without DDD circuit executions. However, even though the exact optimal angle assignments are not consistent across all problems the, they are very close to each other which is notable because it indicates that the optimal angles may be identical or nearly identical but the search space is being obscured by the noise in the computation.

**Optimal parameter choices - QA:** Figure 6 also allows examination of how stable the different parameters are,

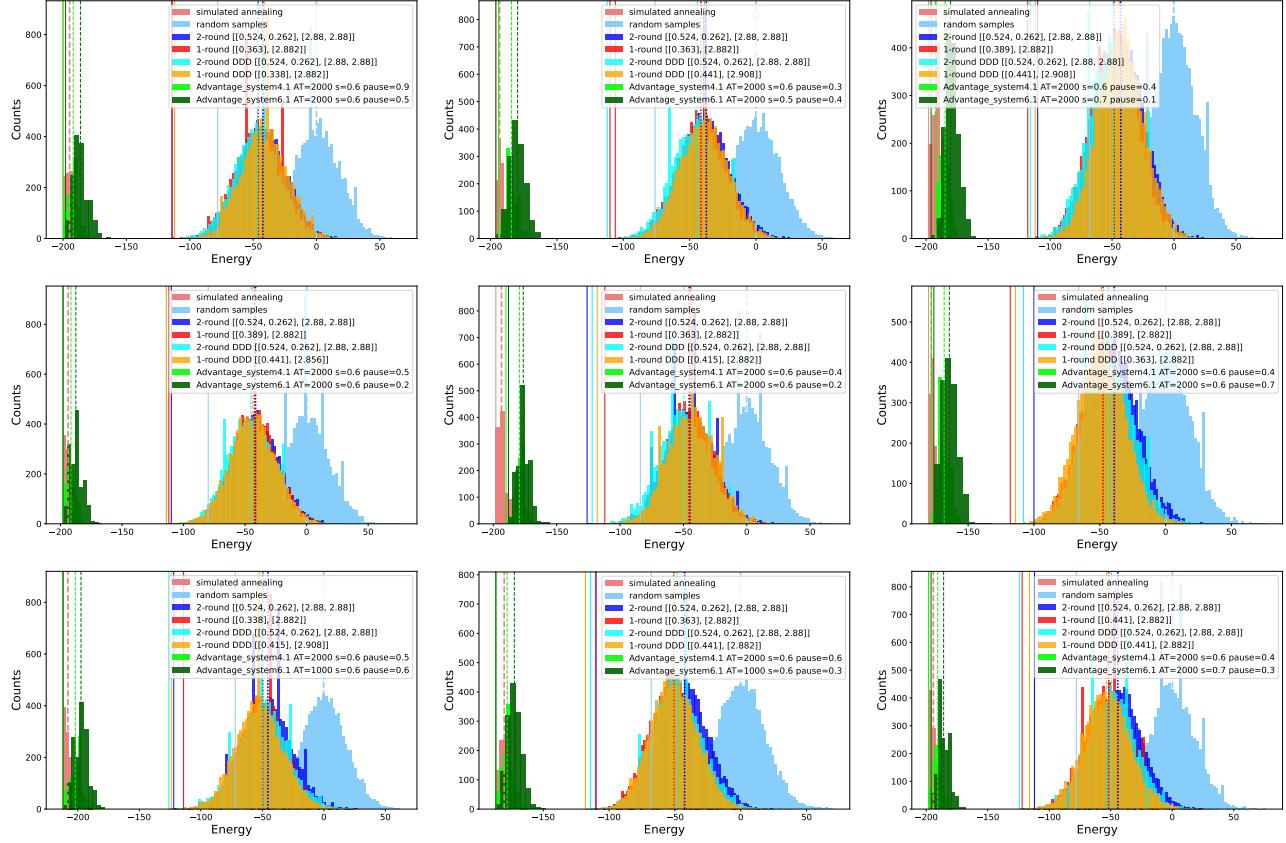


Figure 6: Direct energy histogram comparison of QA and QAOA results for the other nine problem instances, continuing from Figure 5. The mean of each energy distribution is marked with vertical dashed lines, and the minimum energy of each dataset is marked with vertical solid lines. Note that for several of the distributions there are overlapping minimum energies.

	$p = 1$	$p = 2$	$p = 1$ with DDD	$p = 2$ with DDD
$p = 1$ (no DDD) better than -	-	10/10	5/10	4/10
$p = 2$ (no DDD) better than -	0/10	-	2/10	0/10
$p = 1$ (with DDD) better than -	5/10	8/10	-	4/10
$p = 2$ (with DDD) better than -	6/10	10/10	6/10	-

Table 2: How the four different QAOA implementations, one and two rounds with and without DDD, compare against each other in terms of in how many of the 10 random instances each method was better than the other three methods in terms of mean objective function value across the 10000 samples (for the best angle combination). There is a clear finding in the order of performance of the four methods;  $p = 2$  with no DDD performed the worse,  $p = 1$  with no DDD performed the next best,  $p = 1$  with DDD performed the next best, and  $p = 2$  with digital dynamical decoupling performed the best overall.

both across the 10 Ising models but also within each problem instance. In the case of quantum annealing, but the optimal annealing times are always 2000 and the optimal pause schedule is not incredibly consistent with pause fraction durations ranging from 0.1 to 0.9 and with anneal fractions  $s$  ranging from 0.5 to 0.7.

**D-Wave devices performance differences:** One last observation from Figure 6 is that there a small but consistent performance difference between the two quantum annealers; the slightly older generation `Advantage_system4.1` yields lower mean energy than `Advantage_system6.1`. Simulated annealing is comparable to the quantum annealing distributions, with simulated annealing performing marginally better than the quantum annealing distributions.

## 4 Discussion

It is of considerable interest to determine how effective quantum annealing and QAOA are at computing the optimal solutions of combinatorial optimization problems. Combinatorial optimization problems have wide reaching applicability, and being able to solve them faster or to get better heuristic solutions is a very relevant topic in computing. In this article, we have presented experimental results for a fair direct comparison of QAOA and quantum annealing, implemented on the state-of-the-art currently accessible quantum hardware via cloud computing. We leave more detailed benchmarking against state of the art classical solvers on these Ising model instances to future work. This research has specifically found the following:

1. Quantum annealing finds higher quality solutions to the random test Ising models with higher order terms compared to the short depth QAOA  $p = 1$  and  $p = 2$  circuits, with reasonably fine grid searches over the QAOA angles and quantum annealing schedules with pauses.
2. QAOA performs noticeably better than random sampling - this is mostly due to the short depth QAOA circuit constructions which allow reasonably robust computations to be executed without the qubits decohering on current quantum computers.
3. The short depth QAOA circuit construction is notable because it allows for higher order terms in the Ising, and is scalable to a heavy-hexagonal lattice of any size, therefore this circuit construction can be used for future implementations of QAOA on devices with heavy-hexagonal lattices for heavy-hex native Ising models.
4. Dynamical decoupling can improve the computation of QAOA on NISQ computers.

## 5 Acknowledgments

This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). The research presented in this article was supported by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20220656ER and the NNSA’s Advanced Simulation and Computing Beyond Moore’s Law Program at Los Alamos National Laboratory. This research used resources provided by the Darwin testbed at Los Alamos National Laboratory (LANL) which is funded by the Computational Systems and Software Environments subprogram of LANL’s Advanced Simulation and Computing program (NNSA/DOE). This research used resources provided by the Los Alamos National Laboratory Institutional Computing Program. We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team. The authors would like to thank the anonymous reviewers for their helpful comments which helped to improve the manuscript.

LA-UR-22-33077

## References

- [1] Tadashi Kadowaki and Hidetoshi Nishimori. “Quantum Annealing in the Transverse Ising Model”. In: *Physical Review E* 58.5 (1998), pp. 5355–5363. DOI: [10.1103/physreve.58.5355](https://doi.org/10.1103/physreve.58.5355). arXiv: [arXiv:cond-mat/9804280](https://arxiv.org/abs/cond-mat/9804280). URL: <https://doi.org/10.1103%2Fphysreve.58.5355>.
- [2] Satoshi Morita and Hidetoshi Nishimori. “Mathematical Foundation of Quantum Annealing”. In: *Journal of Mathematical Physics* 49.12 (2008), p. 125210. DOI: [10.1063/1.2995837](https://doi.org/10.1063/1.2995837). arXiv: [arXiv:0806.1859](https://arxiv.org/abs/0806.1859).
- [3] Arnab Das and Bikas K Chakrabarti. “Quantum Annealing and Analog Quantum Computation”. In: *Reviews of Modern Physics* 80.3 (2008), p. 1061. DOI: [10.1103/revmodphys.80.1061](https://doi.org/10.1103/revmodphys.80.1061). arXiv: [arXiv:0801.2193](https://arxiv.org/abs/0801.2193).
- [4] Philipp Hauke et al. “Perspectives of quantum annealing: methods and implementations”. In: *Reports on Progress in Physics* 83.5 (2020), p. 054401. DOI: [10.1088/1361-6633/ab85b8](https://doi.org/10.1088/1361-6633/ab85b8). arXiv: [arXiv:1903.06559](https://arxiv.org/abs/1903.06559). URL: <https://dx.doi.org/10.1088/1361-6633/ab85b8>.
- [5] Sheir Yarkoni et al. “Quantum Annealing for Industry Applications: Introduction and Review”. In: *Reports on Progress in Physics* 85.10 (2022), p. 104001. DOI: [10.1088/1361-6633/ac8c54](https://doi.org/10.1088/1361-6633/ac8c54). arXiv: [arXiv:2112.07491](https://arxiv.org/abs/2112.07491). URL: <https://doi.org/10.1088%2F1361-6633%2Fac8c54>.

[6] T. Lanting et al. “Entanglement in a Quantum Annealing Processor”. In: *Phys. Rev. X* 4 (2 2014), p. 021041. DOI: 10.1103/PhysRevX.4.021041. arXiv: arXiv:1401.3500. URL: <https://link.aps.org/doi/10.1103/PhysRevX.4.021041>.

[7] Andrew D. King et al. “Coherent quantum annealing in a programmable 2,000-qubit Ising chain”. In: *Nature Physics* 18.11 (2022), pp. 1324–1328. DOI: 10.1038/s41567-022-01741-6. arXiv: arXiv:2202.05847. URL: <https://doi.org/10.1038/s41567-022-01741-6>.

[8] Stuart Hadfield et al. “From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz”. In: *Algorithms* 12.2 (2019), p. 34. DOI: 10.3390/a12020034. arXiv: arXiv:1709.03489. URL: <https://doi.org/10.3390%2Fa12020034>.

[9] Jeremy Cook, Stephan Eidenbenz, and Andreas Bärtschi. “The Quantum Alternating Operator Ansatz on Maximum k-Vertex Cover”. In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 2020, pp. 83–92. DOI: 10.1109/QCE49297.2020.00021. arXiv: arXiv:1910.13483.

[10] Zhihui Wang et al. “XY mixers: Analytical and numerical results for the quantum alternating operator ansatz”. In: *Physical Review A* 101.1 (2020). DOI: 10.1103/physreva.101.012320. URL: <https://doi.org/10.1103/2Fphysreva.101.012320>.

[11] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. *A Quantum Approximate Optimization Algorithm*. 2014. DOI: 10.48550/ARXIV.1411.4028. arXiv: arXiv:1411.4028. URL: <https://arxiv.org/abs/1411.4028>.

[12] Phillip C. Lotshaw et al. “Scaling Quantum Approximate Optimization on Near-term Hardware”. In: *Scientific Reports* 12.1 (2022). DOI: 10.1038/s41598-022-14767-w. arXiv: arXiv:2201.02247. URL: <https://doi.org/10.1038/s41598-022-14767-w>.

[13] Tameem Albash and Daniel A. Lidar. “Demonstration of a Scaling Advantage for a Quantum Annealer over Simulated Annealing”. In: *Phys. Rev. X* 8 (3 2018), p. 031016. DOI: 10.1103/PhysRevX.8.031016. arXiv: arXiv:1705.07452. URL: <https://link.aps.org/doi/10.1103/PhysRevX.8.031016>.

[14] Andrew D King et al. “Scaling advantage over path-integral Monte Carlo in quantum simulation of geometrically frustrated magnets”. In: *Nature communications* 12.1 (2021), pp. 1–6. DOI: 10.1038/s41467-021-20901-5. arXiv: arXiv:1911.03446.

[15] Edward Farhi and Aram W Harrow. *Quantum Supremacy through the Quantum Approximate Optimization Algorithm*. 2016. DOI: 10.48550/ARXIV.1602.07674. arXiv: arXiv:1602.07674. URL: <https://arxiv.org/abs/1602.07674>.

[16] Salvatore Mandrà et al. “Strengths and weaknesses of weak-strong cluster problems: A detailed overview of state-of-the-art classical heuristics versus quantum approaches”. In: *Physical Review A* 94.2 (2016). DOI: 10.1103/physreva.94.022337. arXiv: 1604.01746. URL: <https://doi.org/10.1103%2Fphysreva.94.022337>.

[17] Sergio Boixo et al. “Evidence for quantum annealing with more than one hundred qubits”. In: *Nature Physics* 10.3 (2014), pp. 218–224. DOI: 10.1038/nphys2900. arXiv: 1304.4595. URL: <https://doi.org/10.1038/2Fnphys2900>.

[18] Byron Tasseff et al. *On the Emerging Potential of Quantum Annealing Hardware for Combinatorial Optimization*. 2022. DOI: 10.48550/ARXIV.2210.04291. arXiv: arXiv:2210.04291. URL: <https://arxiv.org/abs/2210.04291>.

[19] Thomas Lubinski et al. *Optimization Applications as Quantum Performance Benchmarks*. 2023. DOI: 10.48550/ARXIV.2302.02278. arXiv: arXiv:2302.02278. URL: <https://arxiv.org/abs/2302.02278>.

[20] Elijah Pelofske et al. “Sampling on NISQ Devices: ”Who’s the Fairest One of All?””. In: *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2021. DOI: 10.1109/qce52317.2021.00038. arXiv: arXiv:2107.06468. URL: <https://doi.org/10.1109%2Fqce52317.2021.00038>.

[21] Hayato Ushijima-Mesigwa et al. “Multilevel Combinatorial Optimization across Quantum Architectures”. In: *ACM Transactions on Quantum Computing* 2.1 (2021). ISSN: 2643-6809. DOI: 10.1145/3425607. arXiv: arXiv:1910.09985. URL: <https://doi.org/10.1145/3425607>.

[22] Michael Streif and Martin Leib. *Comparison of QAOA with Quantum and Simulated Annealing*. 2019. DOI: 10.48550/ARXIV.1901.01903. arXiv: arXiv:1901.01903. URL: <https://arxiv.org/abs/1901.01903>.

[23] Elijah Pelofske, Andreas Bärtschi, and Stephan Eidenbenz. “Quantum Annealing vs. QAOA: 127 Qubit Higher-Order Ising Problems on NISQ Computers”. In: *Lecture Notes in Computer Science*. Springer Nature Switzerland, 2023, pp. 240–258. DOI: 10.1007/978-3-031-32041-5\_13. arXiv: arXiv:2301.00520. URL: [https://doi.org/10.1007%2F978-3-031-32041-5\\_13](https://doi.org/10.1007%2F978-3-031-32041-5_13).

[24] Guido Pagano et al. “Quantum Approximate Optimization of the Long-Range Ising Model with a Trapped-Ion Quantum Simulator”. In: *Proceedings of the National Academy of Sciences* 117.41 (2020), pp. 25396–25401. DOI: 10.1073/pnas.2006373117. arXiv: arXiv:1906.02700. URL: <https://doi.org/10.1073%2Fpnas.2006373117>.

[25] Johannes Weidenfeller et al. “Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware”. In: *Quantum* 6 (Dec. 2022), p. 870. ISSN: 2521-327X. DOI: 10.22331/q-2022-12-07-870. arXiv: arXiv:2202.03459. URL: <https://doi.org/10.22331/q-2022-12-07-870>.

[26] Matthew P. Harrigan et al. “Quantum Approximate Optimization of Non-Planar Graph Problems on a Planar Superconducting Processor”. In: *Nature Physics* 17.3 (2021), pp. 332–336. DOI: 10.1038/s41567-020-01105-y. arXiv: arXiv:2004.04197. URL: <https://doi.org/10.1038%2Fs41567-020-01105-y>.

[27] Pradeep Niroula et al. “Constrained Quantum Optimization for Extractive Summarization on a Trapped-ion Quantum Computer”. In: *Scientific Reports* 12.1 (2022), pp. 1–14. DOI: 10.1038/s41598-022-20853-w. arXiv: arXiv:2206.06290.

[28] Dylan Herman et al. *Portfolio Optimization via Quantum Zeno Dynamics on a Quantum Processor*. 2022. DOI: 10.48550/ARXIV.2209.15024. arXiv: arXiv:2209.15024. URL: <https://arxiv.org/abs/2209.15024>.

[29] Thomas A Caswell et al. *matplotlib/matplotlib*. Version v3.4.3. DOI: 10.5281/zenodo.5194481.

[30] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55.

[31] Qiskit contributors. *Qiskit: An Open-source Framework for Quantum Computing*. 2023. DOI: 10.5281/zenodo.2573505.

[32] Christopher Chamberland et al. “Topological and Subsystem Codes on Low-Degree Graphs with Flag Qubits”. In: *Phys. Rev. X* 10 (1 2020), p. 011022. DOI: 10.1103/PhysRevX.10.011022. arXiv: arXiv:1907.09528. URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.011022>.

[33] Stefanie Zbinden et al. “Embedding algorithms for quantum annealers with chimera and pegasus connection topologies”. In: *International Conference on High Performance Computing*. Springer. 2020, pp. 187–206. DOI: 10.1007/978-3-030-50743-5\_10.

[34] Nike Dattani, Szilard Szalay, and Nick Chancellor. *Pegasus: The second connectivity graph for large-scale quantum annealing hardware*. 2019. DOI: 10.48550/ARXIV.1901.07636. arXiv: arXiv:1901.07636. URL: <https://arxiv.org/abs/1901.07636>.

[35] C. H. Tseng et al. “Quantum simulation of a three-body-interaction Hamiltonian on an NMR quantum computer”. In: *Phys. Rev. A* 61 (1 1999), p. 012302. DOI: 10.1103/PhysRevA.61.012302. arXiv: arXiv: quant-ph/9908012. URL: <https://link.aps.org/doi/10.1103/PhysRevA.61.012302>.

[36] Nicholas Chancellor, Stefan Zohren, and Paul A Warburton. “Circuit design for multi-body interactions in superconducting quantum annealing systems with applications to a scalable architecture”. In: *npj Quantum Information* 3.1 (2017), pp. 1–7. DOI: 10.1038/s41534-017-0022-6. arXiv: arXiv:1603.09521.

[37] Colin Campbell and Edward Dahl. “QAOA of the Highest Order”. In: *2022 IEEE 19th International Conference on Software Architecture Companion (ICSA-C)*. 2022, pp. 141–146. DOI: 10.1109/ICSA-C54293.2022.00035. arXiv: arXiv:2111.12754.

[38] Elisabetta Valiante et al. “Computational Overhead of Locality Reduction in Binary Optimization Problems”. In: *Computer Physics Communications* 269 (2021), p. 108102. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2021.108102>. arXiv: arXiv:2012.09681. URL: <https://www.sciencedirect.com/science/article/pii/S0010465521002149>.

[39] Hiroshi Ishikawa. “Transformation of General Binary MRF Minimization to the First-Order Case”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33.6 (2011), pp. 1234–1249. DOI: 10.1109/TPAMI.2010.91.

[40] Elijah Pelofske et al. “Quantum Annealing Algorithms for Boolean Tensor Networks”. In: *Scientific Reports* 12.1 (2022). DOI: 10.1038/s41598-022-12611-9. arXiv: arXiv:2107.13659. URL: <https://doi.org/10.1038%2Fs41598-022-12611-9>.

[41] Shuxian Jiang et al. “Quantum Annealing for Prime Factorization”. In: *Scientific reports* 8.1 (2018), pp. 1–9. DOI: 10.1038/s41598-018-36058-z. arXiv: arXiv:1804.02733.

[42] John Golden et al. “Numerical Evidence for Exponential Speed-Up of QAOA over Unstructured Search for Approximate Constrained Optimization”. In: *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, Sept. 2023. DOI: 10.1109/qce57702.2023.00063. arXiv: arXiv:2202.00648. URL: <http://dx.doi.org/10.1109/QCE57702.2023.00063>.

[43] M. Cerezo et al. “Variational quantum algorithms”. In: *Nature Reviews Physics* 3.9 (2021), pp. 625–644. DOI: 10.1038/s42254-021-00348-9. arXiv: arXiv:2012.09265. URL: <https://doi.org/10.1038/s42254-021-00348-9>.

[44] Samson Wang et al. “Noise-Induced Barren Plateaus in Variational Quantum Algorithms”. In: *Nature communications* 12.1 (2021), pp. 1–11. DOI: 10.1038/s41467-021-27045-6. arXiv: arXiv:2007.14384.

[45] Yingyue Zhu et al. “Multi-round QAOA and advanced mixers on a trapped-ion quantum computer”. In: *Quantum Science and Technology* 8.1 (2022), p. 015007. DOI: 10.1088/2058-9565/ac91ef. arXiv: arXiv: 2201.12335. URL: <https://dx.doi.org/10.1088/2058-9565/ac91ef>.

[46] Siyuan Niu and Aida Todri-Sanial. “Effects of Dynamical Decoupling and Pulse-Level Optimizations on IBM Quantum Computers”. In: *IEEE Transactions on Quantum Engineering* 3 (2022), pp. 1–10. DOI: 10.1109/tqe.2022.3203153. arXiv: arXiv:2204.01471. URL: <https://doi.org/10.1109/tqe.2022.3203153>.

[47] Dieter Suter and Gonzalo A. Álvarez. “Colloquium: Protecting quantum information against environmental noise”. In: *Rev. Mod. Phys.* 88 (4 2016), p. 041001. DOI: 10.1103/RevModPhys.88.041001. URL: <https://link.aps.org/doi/10.1103/RevModPhys.88.041001>.

[48] Lorenza Viola, Emanuel Knill, and Seth Lloyd. “Dynamical Decoupling of Open Quantum Systems”. In: *Phys. Rev. Lett.* 82 (12 1999), pp. 2417–2421. DOI: 10.1103/PhysRevLett.82.2417. arXiv: arXiv:quant-ph/9809071. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.82.2417>.

[49] Mustafa Ahmed Ali Ahmed, Gonzalo A. Álvarez, and Dieter Suter. “Robustness of dynamical decoupling sequences”. In: *Physical Review A* 87.4 (2013). DOI: 10.1103/physreva.87.042309. arXiv: arXiv:1211.5001. URL: <https://doi.org/10.1103/physreva.87.042309>.

[50] Ryan LaRose et al. “Mitiq: A software package for error mitigation on noisy quantum computers”. In: *Quantum* 6 (2022), p. 774. DOI: 10.22331/q-2022-08-11-774. arXiv: arXiv:2009.04417. URL: <https://doi.org/10.22331/q-2022-08-11-774>.

[51] Youngseok Kim et al. “Scalable error mitigation for noisy quantum circuits produces competitive expectation values”. In: *Nature Physics* (2023). DOI: 10.1038/s41567-022-01914-3. arXiv: arXiv:2108.09197. URL: <https://doi.org/10.1038/s41567-022-01914-3>.

[52] Nic Ezzell et al. “Dynamical decoupling for superconducting qubits: A performance survey”. In: *Physical Review Applied* 20.6 (Dec. 2023). ISSN: 2331-7019. DOI: 10.1103/physrevapplied.20.064027. arXiv: arXiv:2207.03670. URL: <http://dx.doi.org/10.1103/PhysRevApplied.20.064027>.

[53] Bibek Pokharel et al. “Demonstration of Fidelity Improvement Using Dynamical Decoupling with Superconducting Qubits”. In: *Phys. Rev. Lett.* 121 (22 2018), p. 220502. DOI: 10.1103/PhysRevLett.121.220502. arXiv: arXiv:1807.08768. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.220502>.

[54] Andrew W. Cross et al. *Open Quantum Assembly Language*. 2017. DOI: 10.48550/ARXIV.1707.03429. arXiv: arXiv:1707.03429. URL: <https://arxiv.org/abs/1707.03429>.

[55] Elijah Pelofske, Georg Hahn, and Hristo N. Djidjev. “Parallel Quantum Annealing”. In: *Scientific Reports* 12.1 (2022). DOI: 10.1038/s41598-022-08394-8. arXiv: arXiv:2111.05995. URL: <https://doi.org/10.1038/s41598-022-08394-8>.

[56] Jeffrey Marshall et al. “Power of Pausing: Advancing Understanding of Thermalization in Experimental Quantum Annealers”. In: *Phys. Rev. Appl.* 11 (4 2019), p. 044083. DOI: 10.1103/PhysRevApplied.11.044083. arXiv: arXiv:1810.05881. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.11.044083>.

[57] Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. “Optimization by simulated annealing”. In: *science* 220.4598 (1983), pp. 671–680. DOI: 10.1126/science.220.4598.671.