

Received 12 June 2024; revised 11 July 2024; accepted 6 August 2024; date of publication 9 August 2024;
date of current version 11 September 2024.

Digital Object Identifier 10.1109/TQE.2024.3441229

Multidisk Clutch Optimization Using Quantum Annealing

JOHN D. MALCOLM¹, ALEXANDER ROTH², MLADJAN RADIC²,
PABLO MARTÍN-RAMIRO³ , JON OILLARBURU³ , BORJA AIZPURUA³ ,
ROMÁN ORÚS³, AND SAMUEL MUGEL¹

¹Multiverse Computing, Centre for Social Innovation, Toronto, ON M5T 2C2, Canada

²ZF Friedrichshafen AG, Research and Development, 88046 Friedrichshafen, Germany

³Multiverse Computing, Parque Científico y Tecnológico de Gipuzkoa, 20014 Donostia/San Sebastián, Spain

Corresponding author: Borja Aizpurua (e-mail: borja.aizpurua@multiversecomputing.com).

ABSTRACT In this article, we apply a quantum optimization algorithm to solve a combinatorial problem with significant practical relevance occurring in clutch manufacturing. It is demonstrated how quantum optimization can play a role in real industrial applications in the manufacturing sector. Using the quantum annealer provided by D-Wave Systems, we analyze the performance of the quantum and quantum–classical hybrid solvers and compare them to deterministic- and random-algorithm classical benchmark solvers. The continued evolution of the quantum technology, indicating an expectation for even greater relevance in the future, is discussed, and the revolutionary potential it could have in the manufacturing sector is highlighted.

INDEX TERMS Clutch manufacturing, combinatorial optimization, quadratic unconstrained binary optimization (QUBO) formulations, quantum annealing, quantum computing applications, quantum–classical hybrid algorithms.

I. INTRODUCTION

Optimization problems are omnipresent in all areas of science and industry. In particular, combinatorial optimization problems, which consist of searching for the global minimum of an objective function over discrete variables within a very large space of possible solutions, appear in many practical applications in every industry, including sectors, such as finance, logistics, and manufacturing. For these problems, the size of solution space usually grows exponentially with the number of variables. Although specialized algorithms can be used to find (often-approximate) solutions for specific use cases, most optimization problems are intractable for sufficiently large systems.

In recent years, the state of quantum computing technology has advanced to become practically relevant [1], [2], [3]. These are machines that exploit quantum behavior (special phenomena that can occur at small scales and low temperature) to perform certain calculations, which cannot be simulated efficiently on traditional, or *classical*, computers. For example, D-Wave Systems, an early player in the field, provides access to a particular type of analog quantum computer, referred to as a quantum annealer [4]. Quantum annealers have already been used with great success to solve hard optimization problems in multiple industrial applications,

such as large-scale production [3], [5], [6] and research and development [7], [8] in manufacturing, financial investment strategy [9], shipping logistics [10], [11], [12], and mobility services [13], [14]. There is an abundance of excitement around the advances that quantum computing is expected to yield in many areas. In this article, we focus on a challenging application to a particular manufacturing quality-control problem, which is difficult to solve classically and is among the first studies applying quantum computing to a manufacturing problem (along with, for example, [5]). The problem is used as a platform to demonstrate how quantum computing can provide an edge in current industrial applications and to show its potential in the future as the technology continues to advance to evermore-powerful generations. In clutch manufacturing, one goal of shape or design optimization is to improve the friction of clutches, and it plays a key role in large-scale production. Previous works have considered such improvements by analyzing different designs and material combinations [15], [16]. The complementary approach in this article addresses the same goal, which we refer to as disk optimization, but with a different view and formulation of it. More specifically, our approach to improve the friction in a multidisk clutch consists of finding the optimal orientation for each of the disks. For this purpose, we formulate the




FIGURE 1. Disk sets in a multidisk clutch contained in the ZF 9HP transmissions. Disk set 2 (lined clutch disk) consists of the friction disks (f), which are padded with brown friction elements (see magnification). Disk set 3 (outer multidisk) comprises smooth metal disks (s).

problem as a quadratic unconstrained binary optimization (QUBO) problem [17] and introduce a quantum optimization algorithm to solve it using the D-Wave quantum annealer. First, we show that the quantum solver is able to match the performance of current classical algorithms for small problems by making direct use of the quantum processing unit (QPU). Second, we demonstrate that a hybrid quantum–classical solver shows an exceptional performance for large-scale problems, greatly outperforming classical benchmarks. Therefore, our application of a quantum optimization algorithm provides an edge over classical solutions for large problem sizes when deployed on a hybrid quantum–classical solver.

The rest of this article is organized as follows. Section II provides details on the optimization problem and its mathematical formulation. Sections III and IV describe the techniques used in this study (quantum solvers and classical benchmarks) and the testing approach, respectively. Results are presented and elaborated on, respectively, in Sections V and VI. Finally, Section VII concludes this article.

II. PROBLEM OVERVIEW

In the ZF Group, one optimization problem arises in the manufacturing of a multidisk clutch (see Fig. 1). The multidisk clutch contains the following two alternating sets of disks.

- f) Friction disks (metal disks padded with many friction elements); see brown paddings on *lined clutch disk* (label 2) in Fig. 1.
- s) Smooth metal disks; see *outer multidisk* (label 3) in Fig. 1.

Disks of the same kind are fixed to each other, while kinds (f) and (s) rotate freely relative to each other if the clutch is




FIGURE 2. Illustration of the problem setup where the clutch is made up of a stack of rotatable disks (specifically the friction disks (f) in Fig. 1), with this example showing a three-disk, five-segment problem. Disks are made up of discrete elements, which line up with fixed segments. The disks are labeled along the left and segments along the bottom, with disk shift numbers indicated on the right. Element heights $A_{k,i}$ are labeled, where the indices correspond to the k th disk and i th segment in the stack's original configuration (prior to any shifts applied). The horizontal dashed line gives the average segment height (number of disks, three times the average element height) and the red arrows indicate each segment's deviation from the mean δh_j .

open. When the clutch closes, friction disks (f) start pressing against the smooth disks (s) and assert a friction force reducing the relative rotation speed. If the clutch is closed, the friction between disks in (f) and (s) is high enough such that all disks have the same rotational speed.

The combinatorial optimization problem we are interested in arises due to manufacturing thickness variations of the disks (f) and the attached friction paddings. During manufacturing, the friction disks (f) can each be stacked in 42 different rotational positions, which are then locked in. Thickness variations of each of the friction disks in a stack (e.g., seven friction disks) add up and can lead to decreased performance of the clutch under load. Such a decreased performance would be rejected in quality control tests after manufacturing.

The optimization question is therefore: in which of the 42 possible rotations should each of the disks with friction paddings be placed to optimize the clutch performance and therefore maximize quality according to quality control?

Two of the relevant metrics to increase the clutch performance are as follows.

- 1) M1—The *standard deviation* of the thickness deviations of the friction disks stack from the mean.
- 2) M2—The *range* is defined as the difference between the highest and the lowest total thickness of the friction disks stack.

Both metrics are derived in Section II-A and are given in (3) and (4), respectively.

A. PROBLEM FORMULATION

The optimization problem focuses specifically on the clutch's set of friction disks (f) (see Fig. 1) and a stack of N_D rotatable disks (see Fig. 2). Each disk is made up of N_S discrete elements (the friction pads), which line up within fixed segments

used for reference. As such, there are $(N_S)^{N_D}$ distinct configurations that the stack can take. We denote the height of the i th element of the k th disk as $A_{k,i}$, with indexing starting at zero (hence, $i \in \{0, 1, \dots, N_S - 1\}$ and $k \in \{0, 1, \dots, N_D - 1\}$). The height variation of an element is $B_{k,i} = A_{k,i} - \bar{A}$, where

$$\bar{A} = \frac{1}{N_D N_S} \sum_{k=0}^{N_D-1} \sum_{i=0}^{N_S-1} A_{k,i}. \quad (1)$$

The rotation of the k th disk in the stack is captured by its shift number, $s_k \in \{0, 1, \dots, N_S - 1\}$, or the number of discrete shifts to the left relative to its initial position. Thus, following a rotation of the k th disk, the height of the element shifted into the i th fixed segment is $A_{k,i+s_k}$, where the summation in the index is understood to be modulo N_S . Finally, this takes us to the i th segment's height variation away from the mean segment height

$$\delta h_i = \sum_{k=0}^{N_D-1} B_{k,i+s_k}. \quad (2)$$

The goal of this optimization problem is to find the combination of shift numbers, which achieves as-close-as-possible-to uniformity in the full-stack height in line with the performance metrics standard deviation 1 and range 2. From the vector of segment height variations $\delta \mathbf{h} = (\delta h_0, \dots, \delta h_{N_S-1})$, these metrics 1 and 2, respectively, are computed as

$$\sigma(\delta \mathbf{h}) = \sqrt{\frac{1}{N_S} \sum_{i=0}^{N_S-1} (\delta h_i)^2} \quad (3)$$

$$\text{range}(\delta \mathbf{h}) = \max(\delta \mathbf{h}) - \min(\delta \mathbf{h}). \quad (4)$$

Note that the mean of $\delta \mathbf{h}$ is zero and thus is not included in the definition of the standard deviation.

To form this as a programmable optimization problem, we can consider all segments in tandem, which yields

$$\min_{s_0, \dots, s_{N_D-1}} \|\delta \mathbf{h}\|_n. \quad (5)$$

That is, the optimization involves minimizing the L_n -norm of $\delta \mathbf{h}$ over the shift numbers. For example, the L_∞ -norm minimizes the most extreme value in $\delta \mathbf{h}$ and is thus related to minimization of the range in segment height variations, although they are not equivalent in the sense that they do not lead to the same optimum. To see the distinction, consider a vector f with positive and negative elements f_+ and f_- , respectively. Then

$$\|f\|_\infty = \max(\|f_+\|_\infty, \|f_-\|_\infty) \quad (6)$$

$$\text{range}(f) = \|f_+\|_\infty + \|f_-\|_\infty. \quad (7)$$

While these are indeed different measures, the range gives an upper bound on the L_∞ -norm: $\text{range}(f) \geq \|f\|_\infty$. Alternatively, the L_2 -norm minimizes the vector elements' standard deviation, due to $\sigma(\delta \mathbf{h}) = \sqrt{1/N_S} \|\delta \mathbf{h}\|_2$, and naturally lends itself to the formulation of the problem as a QUBO, as discussed in Section II-B.

B. QUBO FORMULATION

To make use of the D-Wave solvers (discussed in Section III), it is necessary to formulate the optimization problem as a QUBO [17]. That is, as an objective function with only linear or quadratic dependence on binary problem variables

$$\mathcal{H} = x^T Q x \quad (8)$$

where $x \in \{0, 1\}^n$ is a vector of n binary variables and Q is an upper triangular real-valued $n \times n$ matrix. Note here that squaring a binary variable gives the same value of that variable ($x^2 = x$), and hence any self-square terms in the expansion are considered to be linear. From the general problem (5), we minimize the square of the L_2 -norm in order to fulfill the "quadratic" requirement

$$\|\delta \mathbf{h}\|_2^2 = \sum_{i=0}^{N_S-1} (\delta h_i)^2. \quad (9)$$

To map the shift numbers (s_0, \dots, s_{N_D-1}) onto binary variables, we use a unary encoding [18] (a.k.a., a "one-hot" encoding). That is, for each of the N_D disks, we introduce a row vector \mathbf{x}_k made up of N_S binary variables $x_{k,j} \in \{0, 1\}$ (or an $N_D \times N_S$ matrix of variables) such that

$$x_{k,j} = \begin{cases} 1 & \text{if } j = s_k \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

For example, $\mathbf{x}_k = (1, 0, 0, \dots, 0)$ corresponds to $s_k = 0$, $\mathbf{x}_k = (0, 1, 0, \dots, 0)$ corresponds to $s_k = 1$, and so on. This encoding allows us to take the problem variable s_k from its role as an index (not applicable in a QUBO) and into usable function variables. This is done by using the new $\{x_{k,j}\}$ variables as multiplicative assignment variables to indicate whether a particular disk element resides in a fixed segment or not. In this way, the height variation of the i th fixed segment can be written as

$$\delta h_i = \sum_{k=0}^{N_D-1} \sum_{j=0}^{N_S-1} B_{k,i+j} x_{k,j}. \quad (11)$$

Using the unary encoding requires the introduction of N_D constraints whereby one and only one of the variables in each vector \mathbf{x}_k may be equal to one. This is equivalent to requiring that the sum of the binary elements in \mathbf{x}_k be equal to one. This results in a linear equality constraint that can be added to the QUBO as a penalty term

$$\rho \left(\sum_{j=0}^{N_S-1} x_{k,j} - 1 \right)^2 \quad \forall k \in \{0, \dots, N_D - 1\} \quad (12)$$

where ρ is a Lagrange multiplier, often referred to as the penalty strength [19]. Each of the above penalty terms (one for each disk) is equal to zero when the unary-encoding constraint is satisfied and is otherwise greater than zero. Hence, adding these to the cost function with the appropriate penalty strength will force all infeasible solutions to have higher cost than the optimal feasible solution. Ideally, one could set ρ

to be arbitrarily high, which works well, for example, when performing an exhaustive search over all possible solutions and ranking them by cost. However, some of the solvers considered in this work (as discussed in Section III), have a maximum achievable resolution in the cost function. That is, relative to the full range of all possible values achievable in the cost function, the solver can only distinguish between neighboring solutions if they are sufficiently separated from each other in cost. As such, choosing the penalty strength too large can extend the range of the cost function to the point where it becomes difficult for the solver to resolve the optimal solution away from other feasible solutions. In such applications, tuning the penalty strength is an important and nontrivial endeavor.

Combining (9), (11), and (12) yields the cost function

$$\mathcal{H} = \sum_{i=0}^{N_S-1} \left(\sum_{k=0}^{N_D-1} \sum_{j=0}^{N_S-1} B_{k,i+j} x_{k,j} \right)^2 + \rho \sum_{k=0}^{N_D-1} \left(\sum_{j=0}^{N_S-1} x_{k,j} - 1 \right)^2 \quad (13)$$

which has to be minimized. This QUBO formulation requires $N_Q = N_D \cdot N_S$ binary variables and is fully connected. That is, by expanding the cost function into a summation of linear and quadratic terms, at least one quadratic term exists for every possible pairing of problem variables. As such, the problem amounts to the minimization of a fully connected graph problem where linear and quadratic coefficients correspond to vertex and edge weights, respectively (and any constant term corresponds to an overall offset, which is irrelevant to the minimization).

One final adjustment can be made by exploiting a global rotational symmetry in the problem whereby rotating all disks by the same shift number leaves the cost function (13) invariant. This operation is equivalent to changing the arbitrary reference point in the indexing (i.e., the choice of which segment is labeled as zero has no actual bearing on the physical clutch). As such, we are free to leave the first disk frozen in its initial configuration (a shift number of zero), reducing the problem to finding the optimal rotation of all disks relative to the first. Then, by setting $s_0 = 0$ or equivalently $\mathbf{x}_0 = (1, 0, \dots, 0)$, (13) becomes

$$\mathcal{H} = \sum_{i=0}^{N_S-1} \left(B_{0,i} + \sum_{k=1}^{N_D-1} \sum_{j=0}^{N_S-1} B_{k,i+j} x_{k,j} \right)^2 + \rho \sum_{k=1}^{N_D-1} \left(\sum_{j=0}^{N_S-1} x_{k,j} - 1 \right)^2. \quad (14)$$

Exploiting the global rotational symmetry reduces the problem size to $N_Q = (N_D - 1) \cdot N_S$ binary variables.

III. METHODS OVERVIEW

We used a variety of methods and hardware implementations to address the disk optimization problem, with the goal of ascertaining the quantum solvers' performance against selected benchmarks on a classical computer. These are listed below and each solver is detailed in the following.

- 1) ZF exact classical solver.
- 2) ZF approximate classical solver.
- 3) D-Wave simulated thermal annealer.
- 4) D-Wave quantum annealer.
- 5) D-Wave Leap hybrid.

A. ZF EXACT CLASSICAL SOLVER

This solver optimizes the range 2 and is *guaranteed* to find one of the optimal configurations (there could be multiple) of the discrete shift numbers (s_0, \dots, s_{N_D-1}) . It uses the "branch and bound" method [20], [21], consisting of two steps: the branch step followed by the bound step. The branch step divides the current problem under consideration in two or more subproblems. The bound step is responsible for estimating a lower (or upper) bound of the objective function with respect to an already divided subproblem. If the lower bound guarantees no optimal solution, then no further investigation of the whole branch is necessary and the whole branch is cut. This method is thus also sometimes called "branch and cut." Therefore, the performance of this method stands and falls with the feasibility and "sharpness" of the bound. Note that, in the worst case, the complexity of the branch and bound method is $\mathcal{O}(N_S^{N_D-1})$, identical to an exhaustive search, but an optimal solution is guaranteed.

B. ZF APPROXIMATE CLASSICAL SOLVER

As a result of the scaling mentioned above, large-scale problems are hard or even impossible to solve with the ZF exact classical solver in a reasonable time. To tackle this challenge, an approximate classical solver is proposed, subdividing the large-scale case into smaller problems, which may be optimized efficiently. In short, the N_D disks are split into K subsets, where $K - 1$ subsets contain $\lfloor N_D/K \rfloor$ disks and one subset contains the remaining $N_D - (K - 1)\lfloor N_D/K \rfloor$ disks. K is chosen such that $\lfloor N_D/K \rfloor \leq 8$ and $K \geq 3$. The entire subsets are optimized separately and then rotated relative to each other to find the optimum of the range 2. Note, that this approach does not guarantee to find the global optimum because it only searches through a fraction of the whole solution space.

C. D-WAVE SIMULATED ANNEALER

Simulated thermal annealing (typically referred to as "simulated annealing") is a random algorithm, which seeks to find the global minimum of an optimization problem while avoiding getting trapped in local minima. By making use of a parameter analogous to temperature, the system is "hot" in the beginning to allow it jump to other states in successive

iterations, so that it can escape local minima by hill climbing, and progressively “cooled” to converge to a local minimum (think of a popcorn kernel jumping around a hot kettle with many divots, which is gradually cooled). An advantage of this approach is that it allows the system to escape local minima when the temperature is sufficiently high. However, it can struggle when there are many thin but deep minima, meaning that this random algorithm is not guaranteed to converge to the optimal solution. Further details on this method can be found, for example, in [22].

In this work, we made use of the D-Wave simulated annealer available in the `dimod` python package [23]. This implementation allows for the user to input a QUBO, such as the one in (14). The algorithm scales quadratically in the number of variables, and linearly in the number of *sweeps* (steps of the algorithm), which corresponds to the number of different temperatures that are traversed. The initial state is selected at random, so that many runs of the algorithms can allow for a better exploration of the possible solutions.

D. D-WAVE QUANTUM ANNEALER

Quantum annealing is named in analogy to thermal annealing in that the physical process uses random quantum fluctuations (as opposed to thermal fluctuations). The algorithm seeks out the optimal solution as the system is slowly transitioned such that its energy landscape (determined by the *Hamiltonian*) reflects the cost function to be optimized. By leveraging quantum behavior, the system begins in the lowest energy configuration of a “mixing Hamiltonian,” [24], which places the quantum bits (*qubits*) in a complete *superposition* of all possible bitstring combinations at once. By slowly transitioning to the “problem Hamiltonian,” the quantum adiabatic theorem demonstrates that the quantum system will remain in the lowest energy configuration, ultimately arriving at that which corresponds to the optimal solution of the cost function of interest [25]. Unwanted thermal fluctuations in the hardware can introduce noise to the system so that the optimal solution is not always obtained. Hence, many samples are typically made in order to increase the likelihood of finding that optimum.

In this work, we made use of the D-Wave Advantage QPU, introduced in 2020. Here, over 5000 physical qubits (superconducting loops held at the extremely low temperature of 15 mK) are laid out in a Pegasus graph architecture [26] (where each qubit is generally connected to 15 other qubits in the system), so that it is possible to embed a fully connected graph with 180 vertices.¹ This is a major improvement compared to the previous generation D-Wave 2000Q processor, based on more than 2000 qubits on a Chimera architecture [27] (connectivity per qubit of 6) where the largest fully connected graph embedding is of size 65. In addition, the Advantage QPU is said to introduce less noise

¹Minor embedding is the problem of replicating one graph onto another of different structure. In this context, several physical qubits can be chained together with strong coupling to represent a single *logical qubit*.

than earlier generation QPU’s [28]. In the context of the QPU with reference to the minor graph embedding, we refer to a *logical qubit* (not to be confused with an error-corrected qubit [29]) as a single binary variable appearing in the objective function. A *physical qubit* is a single superconducting loop in the QPU hardware. Following the embedding, many physical qubits may be chained together to represent a single logical qubit.

E. D-WAVE LEAP HYBRID

The final solver we explored from D-Wave is their Leap hybrid. This uses a hybrid quantum–classical approach, whereby the submitted problem (14) gets strategically decomposed into subproblems, which are then passed in parallel to the D-Wave Advantage QPU described in Section III-D, as well as classical solvers [24]. Following several iterations, Leap hybrid then outputs the best solution obtained within a set run time. As with the simulated annealer and the QPU, the Leap hybrid is a probabilistic solver and is thus not guaranteed to produce the optimal solution, but has a high level of performance (as presented in Sections V and VI). For the moment, it can handle problem sizes of up to 10^6 binary variables.

As a proprietary solver, the Leap hybrid is understandably not entirely transparent. This can make it difficult to compare against other solvers in some respects. For example, a single run of this hybrid solver may make limited use of the QPU depending on the length of the queue for that quantum machine. As such, it can be somewhat unclear exactly how much the quantum component contributed to finding a solution. Nevertheless, we show below that the Leap hybrid has exceptional performance for large problem sizes. Alternative hybrid solvers are available, which provide more transparency, such as the Kerberos solver available in the `dwave-hybrid` python package [30] (also developed by D-Wave, as the name suggests). However, exploring alternative hybrid solvers was outside the scope of this project.

IV. TEST APPROACH

A. PROBLEM INSTANCES

To test the various methods, quantum and classical, we generate random matrices corresponding to N_D disks with each N_S elements. We draw the element thickness $A_{k,i}$ out of a uniform distribution over the interval $[A_0 - \frac{\Delta}{2}, A_0 + \frac{\Delta}{2}]$, where A_0 is the target element thickness and Δ is the maximal thickness variation of each element.

Relevant for the current manufacturing process in ZF are problem sizes with $N_D \leq 7$ and $N_S = 42$ corresponding to $(N_D - 1) \cdot N_S \leq 252$ binary variables (or logical qubits) in the QUBO formulation. Despite the practical limit, we explored problem sizes larger than this to analyze the performance and scaling of the different methods, giving us valuable insights in how to solve other fully connected QUBO problems in the future.

B. BENCHMARK TESTING

Each of the random-algorithm solvers (simulated annealing, quantum annealing, and hybrid) was run repeatedly a certain number of times in order to increase the probability of sampling the best solution. These solvers, in contrast to the ZF solvers, aim to minimize the QUBO in (14), which corresponds to optimizing the standard deviation in (3). The simulated annealing method acts as a random-algorithm benchmark for comparison against the quantum solvers, and the ZF solvers provide deterministic classical benchmarks. In contrast to the QUBO formulation of the optimization problem, the ZF solvers focus on minimizing the range in (4). We find, however, that an optimal solution with respect to standard deviation corresponds to a low (but not necessarily optimal) range, and vice versa. In the following, we detail some of the specifications used in the three random-algorithm solvers.

- 1) *Simulated Annealer*—35 samples were taken for each problem, using 1500 as the number of sweeps.
- 2) *Quantum Annealer*—From 20 up to 2000 samples were taken for each problem, depending on the problem size, and annealing pauses were used to enhance the success probability as in [31].²
- 3) *Hybrid Solver*—Three samples were taken for each problem. Note, however, that in a single run of the hybrid solver, the backend runs iteratively over several subproblems and takes many samples directly from the quantum annealer in the process.

V. RESULTS

In this section, we show the different results obtained by the considered solvers described in Section III under the approach described in Section IV. For varying problem sizes, the main quantities of interest are the problem metrics standard deviation, (3), and range, (4), associated with the solution that each solver outputs, as well as the runtimes. All of these results are provided in Fig. 3(a)–(c), respectively. Note that all of the solutions presented in the figure were confirmed to satisfy the constraints in (12). For the quantum annealing solver, in particular, obtaining feasible solutions without sacrificing solution quality (i.e., low values for the metrics 1 and 2) required careful tuning of the penalty strength ρ , which was carried out empirically.

A. PROBLEM METRIC COMPARISONS

To begin with, it is important to recall which quantities have been optimized by the different solvers. For the simulated annealing (pink triangles in Fig. 3), quantum annealing (blue circles), and the hybrid (orange crosses) solvers, the L_2 -norm of $\delta\mathbf{h}$ is minimized [corresponding to minimization of the standard deviation, (3)]. Data points for each of these solvers

are connected by solid lines for added clarity. For the ZF Exact (green squares) and Approximate (gray inverted triangles) solvers, it is instead the range that is minimized, see (4), with data points in the figure connected by dashed lines. Both of the quantities 1 and 2 are equally relevant and a solution is considered good or high quality if it is low in both (preferably optimal). The optimal solution with regard to the two different quantities can sometimes be different, but typically a configuration with minimal range yields a relatively low standard deviation, and vice versa. As a result, when comparing two solutions for the same problem instance, if one solver yields a lower standard deviation but a higher range than the other, we consider both solutions to be of a similar quality.

It is also important to note that the optimal metric value is instance dependent. Any perceived trends with varying problem sizes (such as oscillations) are mostly coincidental. However, the magnitude of the metrics M1 and M2 tend to increase with the problem size.

In Fig. 3(a), we see that for all problems considered, the hybrid solver yields solutions with the lowest standard deviation compared with all other solvers. While there is no guarantee from this solver that it finds the optimal solution, its consistently strong results across all investigated problems indicate that these solutions are at least close to optimal (especially for the smaller problem sizes where sampling errors are lower). Direct use of the quantum annealing solver yields good solutions, which are close to the hybrid's solutions in standard deviation (although not necessarily optimal) up to problem sizes of 40 logical qubits. Beyond that, sampling errors lead to results that are far above optimal, even for 50 qubits (well below the embedding limit of 180 qubits). We further discuss these limitations of the quantum annealing solver in Section VI. As for the benchmarks, the ZF Exact solver yields solutions that are close to optimal (if not optimal) in the standard deviation, despite its optimization of the range, see (4). The high quality of solutions is consistent for all problem sizes up to its limit around 340-qubit equivalent [recall that the ZF solvers do not use the QUBO formulation of the problem and hence do not use the binary encoding of the problem variables in (10)]. Beyond this limit, the ZF Approximate solver is applied to large problems, which yield solutions well separated from those obtained from the hybrid solver. Finally, the simulated annealing solver consistently yields suboptimal results (indicating the presence of deep local minima in the cost function). In a practical setting, the results from the simulated annealing could be sufficient to pass quality-control requirements depending on the allowed tolerance, with the advantage that this solver is run locally. However, it is also limited in that it cannot handle large problem sizes at about 400 qubits or greater, similar to the ZF Exact solver. In such instances, the hybrid solver is clearly preferred.

In Fig. 3(b), the observations in the comparison of solutions' ranges are quite similar to the standard deviation results. The main difference here is that the ZF Exact solver

²Imposing an annealing pause involves adjusting the annealing schedule as the QPU transitions from the mixing Hamiltonian to the problem Hamiltonian. As indicated in the cited work, a carefully positioned pause during the anneal can enhance the probability of finding the ground state of the problem Hamiltonian.

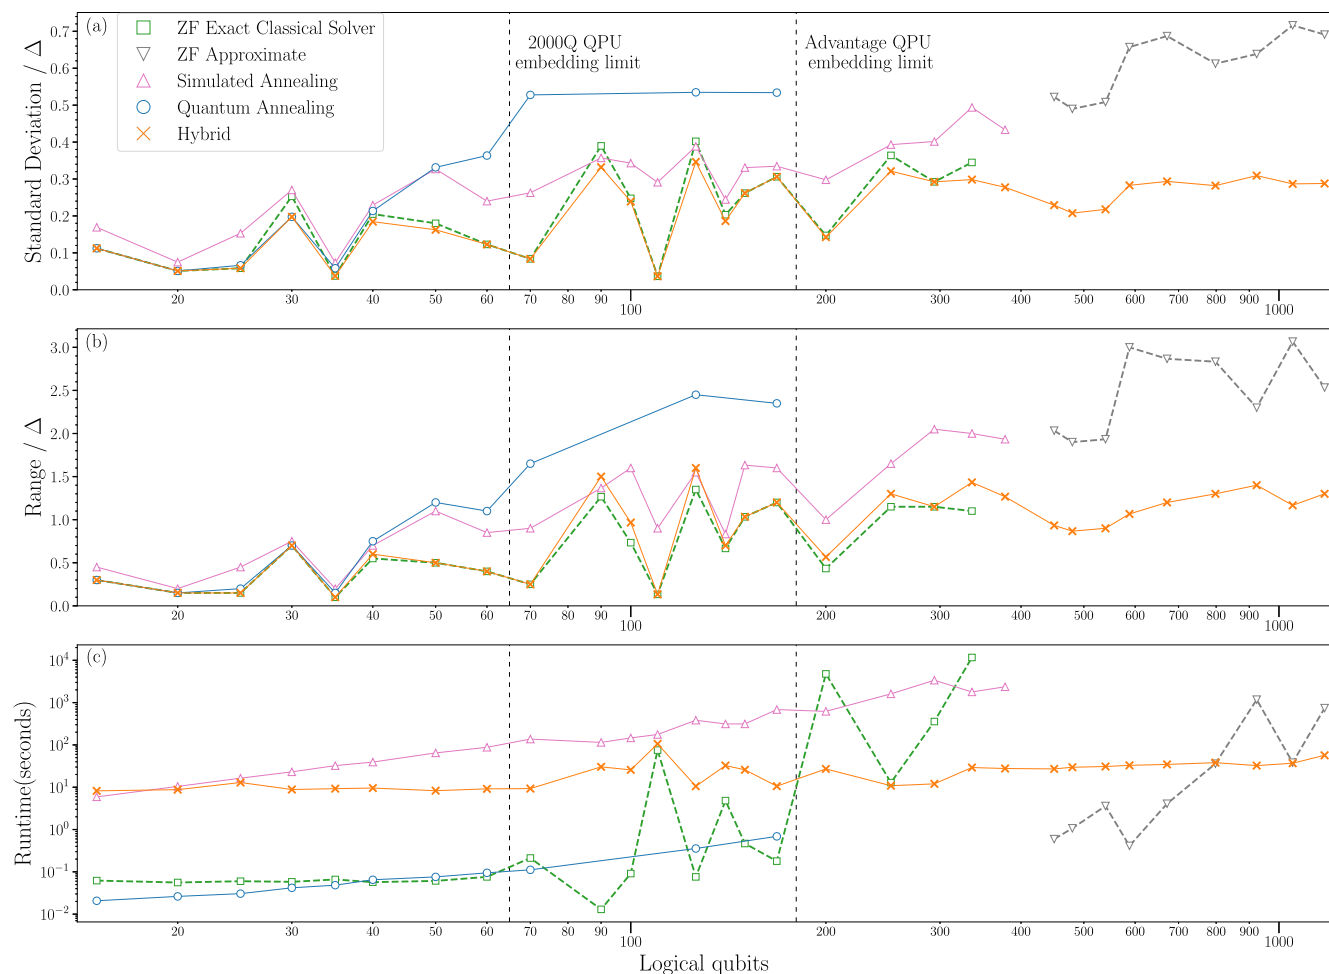


FIGURE 3. Comparison of (a) standard deviation [see (3)], (b) range [see (4)], and (c) runtimes for the different solvers across independent random disk optimization problem instances. For added clarity, solvers which optimized the standard deviation (range) are connected by solid (dashed) lines. Each problem corresponds to a fully connected graph with a certain number of logical qubits (e.g., the problem with $N_D = 7$ disks and $N_S = 42$ segments corresponds to $(N_D - 1)N_S = 252$ qubits). Vertical dashed lines indicate the embedding limit of a fully connected graph in the Pegasus architecture of the D-Wave Advantage QPU on the right (used for both the quantum annealer and hybrid solvers) and in the Chimera architecture of the D-Wave 2000Q QPU (indicated only for reference). Note that the vertical scale in (c) and the horizontal scale in all three panels are logarithmic.

yields solutions with the guaranteed optimal range compared with the other solvers. Interestingly, in many instances, the hybrid solver manages to find a solution with the same minimal value of the range, despite that solver’s optimization of the standard deviation. There are even instances, such as the one with 50 qubits, where the hybrid solver matches the minimal range and still manages to find a lower standard deviation as compared to the ZF Exact solver.

B. RUNTIME COMPARISONS

Fig. 3(c) provides the total runtimes of each solver for the various problem sizes. Note that comparisons of the absolute runtimes (i.e., the levels) are not fundamentally relevant given the different platforms each solver was run on, but are still useful for practical reasons. For example, the simulated annealing solver was run locally on a MacBook Pro laptop, the ZF solvers on a ThinkPad, and the classical components of the Hybrid solver are run on supercomputing clusters.

More relevant is the relative scaling of each solver’s runtime with the problem size that does not depend on the underlying hardware (beyond the classical versus quantum qualities). Furthermore, for the random-algorithm solvers (simulated annealing, quantum annealing, and hybrid), a different, often more-relevant, measure is time to solution (which considers the probability of finding the optimal solution relative to a given level of confidence), rather than total runtime. However, since the simulated and quantum annealing solvers did not find the optimal solution in most instances, we chose to instead present total runtime for a fixed number of samples.

For both the ZF Exact and ZF Approximate solvers, the runtime appears to be super exponential (increasing slope in the log–log plot), although the runtime of these approaches is very instance dependent. For problem sizes corresponding to about 200 qubits or more (or even earlier, depending on usage requirements), the runtimes in the ZF Exact solver become prohibitive at the order of an hour or more.

At all problem sizes (small, medium, and large), the simulated annealing presents among the longest runtimes and scales quadratically (according to theory). From a runtime perspective in practical applications, this approach becomes infeasible for medium-size problems and beyond.

Making direct use of the quantum annealing solver has very low runtimes, even for a problem size approaching its embedding limit. For problem sizes below 40 qubits, where this solver was able to find close-to-optimal solutions (if not optimal), it had the fastest runtime compared to all other solvers (even when including internet latency). This is particularly promising in that the quantum computing technology is still rather nascent and the solution quality is only expected to improve. This topic is discussed further in Section VI.

The runtimes for the hybrid solver are fairly consistent on the order of 10 s, based on the minimum fixed runtime (3 s) for each of the three samples made per problem. Fluctuations seen in the runtime are due to internet latency and other time-costing overhead processes. There is a shallow, almost imperceptible, increase in runtime (beyond 1024 variables, the minimum runtime of the solver per sample gradually grows from the baseline 3 s). For small- and medium-size problems, the ZF Exact solver is faster and yields similar-quality results as the hybrid. However, for large problem sizes, the hybrid solver is incredibly fast while still yielding high-quality results. For a problem corresponding to 336 qubits, the ZF Exact solver took nearly 3.5 h (11 659 s) while the hybrid solved it in only 29 s. The ZF Solver was unable to handle problems larger than this, while the hybrid continued to yield high-quality results with a runtime of only 56 s on the largest problem considered (1176 qubits). The specifications indicate that it can handle problem sizes of up to 10^6 variables. In this work, the calculations were not pushed to these boundaries due to a lack of readily available benchmarks.

VI. DISCUSSION

Based on the results presented in Section V, the different scales of problems show differences in the performance of the solvers. For small problems, below 50 qubits, all solvers find high-quality solutions except maybe the simulated annealing. However, in terms of runtime, the quantum annealing solver is fastest—by several orders of magnitude than most solvers, and slightly better than the ZF Exact solver. For small problem sizes, where sampling errors are reduced, taking 20 to 100 samples typically suffices to find (near-)optimal results, and so runtimes could be reduced without sacrificing solution quality as compared to larger problem sizes.

For medium-size problems, 50 to 350 qubits, the quantum annealing solver no longer performs well. This is due to noise in the machine arising from thermal fluctuations, which lead to sampling errors. In particular, as the number of problem variables increases linearly, embedding of the fully connected graph problem onto the Pegasus architecture (16-qubit connectivity) requires that evermore physical qubits be chained together to represent a single logical qubit.

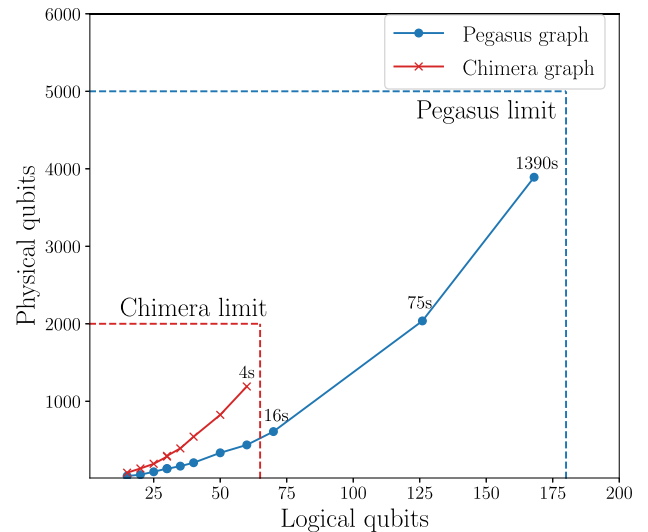


FIGURE 4. Number of physical qubits required for the embedding of fully connected graph problems on the Pegasus architecture (relevant to the advantage QPU used in this study) and the Chimera architecture (in the previous-generation 2000Q QPU, included for reference). The data points included here correspond to the problem sizes tested with the quantum annealing solver in Fig. 3. Numbers assigned to a subset of the data points indicate the runtime in seconds for calculating the minor embedding.

This relationship is demonstrated in Fig. 4, which shows the number of physical qubits used for each of the problem sizes considered with this solver in blue (as well as the embedding limit assuming an availability of 5000 physical qubits). The physical qubit chain is achieved by imposing a strong ferromagnetic coupling across all qubits in the chain in an attempt to fix them all at the same value (either 0 or 1), which then corresponds to the value of the corresponding logical qubit. However, for large problem sizes, chain breaks are more likely to occur (where the strong coupling fails to ensure all physical qubits in the chain are equal valued, potentially leading to an incorrect assignment in the corresponding logical qubit) because the chains are longer and are thus more susceptible to sampling errors. In addition, there is more competition between the chains' couplings and the couplings that encode the optimization problem itself, obfuscating the details of the problem and making it more difficult for the solver to discern the optimal solution.

While there are clear limitations to the performance when making direct use of the quantum annealing solver, we point out again that this technology is still in its early stages and that there are major improvements with each new generation of quantum annealer from D-Wave. To demonstrate this, Figs. 3 and 4 indicate the limits of embedding fully connected graphs onto the previous generation QPU, the 2000Q, compared with the Advantage QPU used in this study. The latter figure shows that the increased connectivity in the Advantage QPU architecture means that fewer physical qubits are required to represent each logical qubit in the same size problem. Had we made use of the 2000Q QPU, we would have

only been able to run problem sizes up to about 65 qubits and the quality of the solutions would have been reduced relative to the Advantage QPU results.

As an aside, we also include the runtimes (in seconds) in Fig. 4 for determining the embedding of a subset of the problems considered (using the `minorminer` python package from D-Wave [32]). Minor graph embedding is itself an NP-Hard problem and if done for each run of the quantum annealer solver would completely wash out the fast runtimes seen in Fig. 3(c). However, since all problems are known to be on fully connected graphs, we are able to precompute the embedding for each problem size and store this in a library for subsequent and repeated use. As such, we did not include the embedding times in Fig. 3(c).

It is in the latter half of this medium-size region (at about 200 qubits and greater) that the hybrid solver becomes the preferred approach. For all problem sizes, it yields consistently high-quality results, and at this scale of problems, it has a faster runtime than the ZF Exact solver. Also, in this region, the runtimes of the ZF Exact solver increase rapidly, to a point where its use is prohibitively expensive in both runtime and memory constraints. For large problem sizes, 350 qubits or more, the ZF Approximate solver has reasonable runtimes (up until about 800 qubits), but the solution quality is much poorer than in the hybrid solver.

To rank the different solvers considered in this work, direct use of the quantum annealing solver is preferred for small problem sizes due to its capacity to find high-quality solutions in very short times. However, hardware limitations in that solver mean that both the hybrid and ZF Exact solvers outperform it for medium-sized problems. In this region, these two algorithms yield similar results in terms of solution quality. In particular, we observe that the hybrid solver finds solutions with slightly better standard deviation, whereas the ZF Exact solver finds solutions with a smaller range. For large-size problems, the hybrid solver shows an impressive performance, clearly outperforming the ZF Approximate solver in terms of solution quality and runtime. In this region, the ZF Exact solver has no data points as it becomes prohibitively expensive in both runtime and memory, highlighting the benefits of the hybrid solver.

VII. CONCLUSION

We have successfully applied a quantum optimization algorithm to tackle the optimization problem of improving the configuration of friction pads in a clutch, a challenging quality-control problem in manufacturing. As far as we are aware, this is among the first studies applying quantum computing to a manufacturing assembly problem (in league with, for example, [5]). We presented the mathematical formulation of the problem in Section II-A, a minimization of the range (for the deterministic classical benchmark algorithms) of height variations or the L_2 -norm (for the random algorithms, including the quantum approaches that were the main focus of this work). For the latter, we derived a QUBO formulation using unary encoding for the shifts

in order to meet the problem-formulation requirements for quantum annealing. We ran the algorithm on two quantum solvers (direct use of the quantum annealing solver and the quantum-classical hybrid solver) on randomly generated problem instances of varying sizes and compared against classical benchmarks. In a more general sense, this work is applicable to the optimization of fully connected QUBO problems with unary-encoding constraints.

From the results presented in Section V and elaborated on in Section VI, we found that for small problem sizes (below about 50 qubits), the quantum annealing solver yields high-quality solutions (although not always optimal) with very short runtimes. The performance of this solver is only expected to improve over time with further advances in the underlying technology. For medium-size problems up to about 350 qubits, the hybrid solver and the ZF exact solver were on par in solution quality. However, in terms of runtime, the latter rapidly scales upward as the number of variables is increased. In contrast, the hybrid solver remarkably has only limited scaling in runtime while maintaining high quality of solutions. Using the hybrid solver, we were able to solve large problems of up to 1200 fully connected variables, with relatively short runtime. Furthermore, much larger problem sizes could have been run using the hybrid solver, but without a useful benchmark for comparison we limited our investigations there.

With further advances in the quantum technology (extending, for example, the evolution of the D-Wave quantum annealer's capabilities highlighted in Fig. 4), the benefits of making direct use of the QPU in combinatoric optimization are only expected to improve. In addition to such improvements from technological advances, we believe that the current generation's performance could be made better with techniques available now that were outside the scope of this project. For one, the penalty strength in (14) could be optimized in a better way: an ideal penalty strength would limit the sampling of infeasible solutions picked by the solver, while preserving the details of the energy landscape of feasible solutions. While we tuned the penalty strength empirically in this work, there are ideas available in the literature that suggest more systematic approaches (for example, in [33] and [34]).

From this work, we point out the strong performance of quantum-classical hybrid approaches, such as in D-Wave's leap hybrid solver used here, which can lead to significantly better results than only classical or only quantum approaches. The results presented in this work demonstrate the competitive advantage that quantum computing can already offer in challenging problems encountered in the manufacturing sector. In our point of view, this study underlines the revolutionary potential of quantum computing in the manufacturing sector and the role it could play in the reshaping of industry.

In addition to our findings, future work could explore the use of domain-wall encoding to further enhance the efficiency of encoding optimization problems into physical qubits. Domain-wall encoding represents a variable with N

possible states using $N-1$ binary variables. This encoding method not only reduces the number of required qubits but also requires fewer and simpler penalty terms, which can improve the performance of quantum annealers. The potential benefits of domain-wall encoding have been demonstrated in various studies, showing improved embedding efficiency and performance in quantum optimization problems [35], [36], [37]. By incorporating domain-wall encoding, larger problem instances could be encoded more efficiently, potentially allowing quantum annealers to tackle even more complex and large-scale industrial problems.

ACKNOWLEDGMENT

This work is a collaboration between ZF Group and Multiverse Computing. The authors would like to acknowledge and thank Amy Friske, Samuel Palmer, Serkan Sahin, and Devansh Trivedi for their helpful discussions.

REFERENCES

- [1] L. Gyongyosi and S. Imre, "A survey on quantum computing technology," *Comput. Sci. Rev.*, vol. 31, pp. 51–71, 2019, doi: [10.1016/j.cscs.2018.11.002](https://doi.org/10.1016/j.cscs.2018.11.002).
- [2] R. Orús, S. Mugel, and E. Lizaso, "Quantum computing for finance: Overview and prospects," *Rev. Phys.*, vol. 4, 2019, Art. no. 100028, doi: [10.1016/j.revip.2019.100028](https://doi.org/10.1016/j.revip.2019.100028).
- [3] A. Luckow, J. Klepsch, and J. Pichlmeier, "Quantum computing: Towards industry reference problems," *Digitale Welt*, vol. 5, pp. 38–45, 2021, doi: [10.1007/s42354-021-0335-7](https://doi.org/10.1007/s42354-021-0335-7).
- [4] K. Boothby et al., "Architectural considerations in the design of a third-generation superconducting quantum annealing processor," 2021, *arXiv:2108.02322*, doi: [10.48550/arXiv.2108.02322](https://doi.org/10.48550/arXiv.2108.02322).
- [5] S. Yarkoni et al., "Multi-car paint shop optimization with quantum annealing," in *2021 IEEE Int. Conf. Quantum Comput. Eng. (QCE)*, 2021, pp. 35–41, doi: [10.1109/QCE52317.2021.00019](https://doi.org/10.1109/QCE52317.2021.00019).
- [6] M. Streif et al., "Beating classical heuristics for the binary paint shop problem with the quantum approximate optimization algorithm," *Phys. Rev. A*, vol. 104, no. 1, 2021, Art. no. 012403, doi: [10.1103/PhysRevA.104.012403](https://doi.org/10.1103/PhysRevA.104.012403).
- [7] A. Glos, A. Kundu, and Ö. Salehi, "Optimizing the production of test vehicles using hybrid constrained quantum annealing," *SN Comput. Sci.*, vol. 4, no. 5, 2023, Art. no. 609, doi: [10.1007/s42979-023-02071-x](https://doi.org/10.1007/s42979-023-02071-x).
- [8] T. Hubregtsen, C. Segler, J. Pichlmeier, A. Sarkar, T. Gabor, and K. Bertels, "Integration and evaluation of quantum accelerators for data-driven user functions," in *2020 IEEE 21st Int. Symp. Qual. Electron. Des.*, 2020, pp. 329–334, doi: [10.1109/ISQED48828.2020.9137029](https://doi.org/10.1109/ISQED48828.2020.9137029).
- [9] S. Mugel et al., "Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks," *Phys. Rev. Res.*, vol. 4, no. 1, 2022, Art. no. 013006, doi: [10.1103/PhysRevResearch.4.013006](https://doi.org/10.1103/PhysRevResearch.4.013006).
- [10] Ö. Salehi, A. Glos, and J. A. Mischak, "Unconstrained binary models of the travelling salesman problem variants for quantum optimization," *Quantum Inf. Process.*, vol. 21, no. 2, 2022, Art. no. 67, doi: [10.1007/s11128-021-03405-5](https://doi.org/10.1007/s11128-021-03405-5).
- [11] S. Feld et al., "A hybrid solution method for the capacitated vehicle routing problem using a quantum annealer," *Front. ICT*, vol. 6, 2019, Art. no. 13, doi: [10.3389/fict.2019.00013](https://doi.org/10.3389/fict.2019.00013).
- [12] A. E. Moylett, N. Linden, and A. Montanaro, "Quantum speedup of the traveling-salesman problem for bounded-degree graphs," *Phys. Rev. A*, vol. 95, no. 3, 2017, Art. no. 032323, doi: [10.1103/PhysRevA.95.032323](https://doi.org/10.1103/PhysRevA.95.032323).
- [13] S. Yarkoni et al., "Quantum shuttle: Traffic navigation with quantum computing," in *Proc. 1st ACM SIGSOFT Int. Workshop Architectures Paradigms Eng. Quantum Softw.*, 2020, pp. 22–30, doi: [10.1145/3412451.3428500](https://doi.org/10.1145/3412451.3428500).
- [14] H. Hussain et al., "Optimal control of traffic signals using quantum annealing," *Quantum Inf. Process.*, vol. 19, pp. 1–18, 2020, doi: [10.1007/s11128-020-02815-1](https://doi.org/10.1007/s11128-020-02815-1).
- [15] V. Thacker and E. Vijayaragavan, "Design optimization of clutch cushion disc with the integration of finite element analysis and design of experiments," *IOP Conf. Ser.: Mater. Sci. Eng.*, vol. 402, no. 1, 2018, Art. no. 012053, doi: [10.1088/1757-899X/402/1/012053](https://doi.org/10.1088/1757-899X/402/1/012053).
- [16] K. Virmani, T. Madhogaria, and P. Baskar, "Design optimization of friction lining of a clutch plate," *Mater. Today: Proc.*, vol. 46, pp. 8009–8024, 2021, doi: [10.1016/j.matpr.2021.02.775](https://doi.org/10.1016/j.matpr.2021.02.775).
- [17] M. Lewis and F. Glover, "Quadratic unconstrained binary optimization problem preprocessing: Theory and empirical analysis," *Netw.*, vol. 70, no. 2, pp. 79–97, 2017, doi: [10.1002/net.21751](https://doi.org/10.1002/net.21751).
- [18] S. Kak, "Generalized unary coding," *Circuits, Syst., Signal Process.*, vol. 35, no. 4, pp. 1419–1426, 2016, doi: [10.1007/s00034-015-0120-7](https://doi.org/10.1007/s00034-015-0120-7).
- [19] F. Glover et al., "Quantum bridge analytics I: A tutorial on formulating and using QUBO models," *Ann. Operations Res.*, vol. 314, no. 1, pp. 141–183, 2022, doi: [10.1007/s10479-022-04634-2](https://doi.org/10.1007/s10479-022-04634-2).
- [20] R. J. Dakin, "A tree-search algorithm for mixed integer programming problems," *Comput. J.*, vol. 8, no. 3, pp. 250–255, 1965, doi: [10.1093/comjnl/8.3.250](https://doi.org/10.1093/comjnl/8.3.250).
- [21] A. H. Land and A. G. Doig, "An automatic method of solving discrete programming problems," *Econometrica*, vol. 28, no. 3, pp. 497–520, 1960, doi: [10.2307/1910129](https://doi.org/10.2307/1910129).
- [22] P. J. M. Van Laarhoven et al., "Simulated annealing," in *Simulated Annealing: Theory and Applications*. Dordrecht, The Netherlands: Springer, 1987, pp. 7–15, doi: [10.1007/978-94-015-7744-1](https://doi.org/10.1007/978-94-015-7744-1).
- [23] D-Wave Systems, "dimod: Python package for solving problems with discrete variables," 2022, GitHub, [Online]. Available: <https://github.com/dwavesystems/dimod>
- [24] D-Wave Systems, "Getting started with D-wave solvers," 2021. [Online]. Available: https://docs.dwavesys.com/docs/latest/doc_getting_started.html#doc-gsg
- [25] T. Albash and D. A. Lidar, "Adiabatic quantum computation," *Rev. Modern Phys.*, vol. 90, no. 1, 2018, Art. no. 015002, doi: [10.1103/RevModPhys.90.015002](https://doi.org/10.1103/RevModPhys.90.015002).
- [26] N. Dattani, S. Szalay, and N. Chancellor, "Pegasus: The second connectivity graph for large-scale quantum annealing hardware," 2019, *arXiv:1901.07636*, doi: [10.48550/arXiv.1901.07636](https://doi.org/10.48550/arXiv.1901.07636).
- [27] Y. Cao et al., "Solving set cover with pairs problem using quantum annealing," *Sci. Rep.*, vol. 6, no. 1, 2016, Art. no. 33957, doi: [10.1038/srep33957](https://doi.org/10.1038/srep33957).
- [28] D-Wave Systems, "The D-wave advantage system: An overview," 2020. [Online]. Available: https://www.dwavesys.com/media/s3qbjp3s/14-1049a-a_the_d-wave_advantage_system_an_overview.pdf
- [29] L. Egan et al., "Fault-tolerant control of an error-corrected qubit," *Nature*, vol. 598, no. 7880, pp. 281–286, 2021, doi: [10.1038/s41586-021-03928-y](https://doi.org/10.1038/s41586-021-03928-y).
- [30] D-Wave Systems, "D-wave release 0.6.7," 2020. [Online]. Available: https://docs.ocean.dwavesys.com/en/stable/docs_hybrid/reference/reference.html#module-hybrid.reference.kerberos
- [31] J. Marshall et al., "Power of pausing: Advancing understanding of thermalization in experimental quantum annealers," *Phys. Rev. Appl.*, vol. 11, no. 4, 2019, Art. no. 044083, doi: [10.1103/PhysRevApplied.11.044083](https://doi.org/10.1103/PhysRevApplied.11.044083).
- [32] D-Wave Systems, minorminer. Source code. 2022. GitHub, [Online]. Available: <https://github.com/dwavesystems/minorminer>
- [33] M. Ohzeki, "Breaking limitation of quantum annealer in solving optimization problems under constraints," *Sci. Rep.*, vol. 10, no. 1, 2020, Art. no. 3126, doi: [10.1038/s41598-020-60022-5](https://doi.org/10.1038/s41598-020-60022-5).
- [34] T. Huang et al., "QROSS: QUBO relaxation parameter optimisation via learning solver surrogates," in *2021 IEEE 41st Int. Conf. Distrib. Syst. Workshops*, 2021, pp. 35–40, doi: [10.1109/ICDCSW53096.2021.00013](https://doi.org/10.1109/ICDCSW53096.2021.00013).
- [35] N. Chancellor, "Domain wall encoding of discrete variables for quantum annealing and QAOA," *Quantum Sci. Technol.*, vol. 4, no. 4, 2019, Art. no. 045004, doi: [10.1088/2058-9565/ab33c2](https://doi.org/10.1088/2058-9565/ab33c2).
- [36] C. McGeoch and P. Farré, "Understanding domain-wall encoding theoretically and experimentally," *Phil. Trans. Roy. Soc. A*, vol. 381, no. 2241, 2023, doi: [10.1098/rsta.2021.0410](https://doi.org/10.1098/rsta.2021.0410).
- [37] M. W. Johnson et al., "Comparing QUBO models for quantum annealing: Integer encodings for permutation problems," *Int. Trans. Oper. Res.*, vol. 32, no. 1, pp. 18–37, Jan. 2025, doi: [10.1111/itor.13471](https://doi.org/10.1111/itor.13471).