

Multi-car paint shop optimization with quantum annealing

Sheir Yarkoni^{*†}, Alex Alekseyenko[‡], Michael Streif^{*§}, David Von Dollen^{††}, Florian Neukart^{*†}, Thomas Bäck[†]

^{*}Volkswagen Data:Lab
Munich, Germany

[†]LIACS, Leiden University
Leiden, The Netherlands

[‡]Volkswagen Group of America
San Francisco, CA, USA

[§]Institute of Theoretical Physics, University Erlangen-Nürnberg
Nürnberg, Germany

Abstract—We present a generalization of the binary paint shop problem (BPSP) to tackle an automotive industry application, the multi-car paint shop (MCPS) problem. The objective of the optimization is to minimize the number of color switches between cars in a paint shop queue during manufacturing, a known NP-hard problem. We distinguish between different subclasses of paint shop problems, and show how to formulate the basic MCPS problem as an Ising model. The problem instances used in this study are generated using real-world data from a factory in Wolfsburg, Germany. We compare the performance of the D-Wave 2000Q and Advantage quantum processors to other classical solvers and a hybrid quantum-classical algorithm offered by D-Wave Systems. We observe that the quantum processors are well-suited for smaller problems, and the hybrid algorithm for intermediate sizes. However, we find that the performance of these algorithms quickly approaches that of a simple greedy algorithm in the large size limit.

Index Terms—quantum annealing, quantum computing, optimization, sequencing

I. INTRODUCTION

Quantum computing has the potential to significantly impact both research and industry in a variety of disciplines. Quantum algorithms developed expressly for fully error-corrected quantum computers are among the most promising, specifically Shor’s algorithm for factoring [1] and Grover’s algorithm for unstructured search [2], which are both asymptotically faster than their classical counterparts. Without error correcting schemes, the advantages of quantum computing are more difficult to discern. Recently, a significant milestone of quantum supremacy was reached by Google’s quantum computer, when it was shown to sample from random quantum circuits faster than any classical algorithm [3]. In the near term, specific focus has been dedicated to developing useful algorithms for so-called Noisy Intermediate-Scale Quantum (NISQ [4]) processors. Various companies such as Google,

IBM, and others have built gate-based NISQ computers with up to dozens of qubits [3], [5]. Companies such as D-Wave Systems have adopted an alternative approach, pursuing instead the physical implementation of quantum annealing algorithms, an adaptation of Adiabatic Quantum Optimization. Originally proposed as a heuristic optimization algorithm [6], the field of quantum annealing has grown significantly over time, with many potential applications being showcased in fields such as quantum simulation [7], [8], [9], quantum machine learning [10], [11], [12], optimization [13], [14], [15], and more [16], [17]. The majority of both gate-model and quantum annealing applications currently involve solving an optimization problem, typically formulated as either an Ising Model using spin variables $s \in \{-1, 1\}$, or a Quadratic Unconstrained Binary Optimization (QUBO) problem with binary variables $x \in \{0, 1\}$. These two models are equivalent under a simple change of basis, and are known to be NP-hard to minimize in the worst case [18]. As such, many interesting and complex optimization problems can be posed in one of these models [19]. Ising Models are formulated as a system of interacting spins:

$$H(s) = \sum_i h_i s_i + \sum_{i < j} J_{ij} s_i s_j, \quad (1)$$

where h_i and J_{ij} are real-valued numbers defining the strength of interactions between neighboring spins s_i, s_j . For QUBOs, we define:

$$\text{Obj}(x, Q) = x^T \cdot Q \cdot x, \quad (2)$$

where x is a vector of binary variables, and here Q is an upper-triangular matrix defining the relationships between them. The initial work of constructing an optimization problem suitable for quantum computers is finding a valid representation in

either Ising or QUBO form. In quantum annealers, these optimization problems are encoded by programming sets of qubits and couplers. The quantum processor then initializes the system of qubits in an easy-to-prepare ground state, and slowly transitions the Hamiltonian to represent the optimization problem as specified by the user. The physical and mathematical background motivating the algorithm is beyond the scope of this work, and we refer the reader to [20].

Despite the proof-of-concept projects presented above which have been tested on quantum processors, a large-scale application has yet to be shown. Partly, this can be attributed to the mismatch between the core optimization problem and the limited scope of quantum optimization algorithms. Many industry problems involve a mix of integer, discrete, and continuous variables and constraints, while NISQ processors solve small-scale binary optimization problems. Therefore, it is important to find candidate problems that not only benefit from quantum computing theoretically, but also are amenable to the technology practically. We therefore present an automotive industry problem, the multi-car paint shop (MCPS) problem. This optimization problem is an extension of the binary paint shop problem (BPSP) which has been studied in the context of other quantum algorithms [21]. The MCPS problem represents the real-world version of the BPSP, and is a crucial step in the process of car manufacturing at Volkswagen. In this study we only solve problem instances obtained from a factory in Wolfsburg, Germany. We employ a variety of classical solvers as competition benchmarks, and compare the results to D-Wave quantum processing units (QPUs) and a hybrid quantum-classical algorithm.

The rest of the paper is structured as follows: in Section II we review the relevant existing literature for the paint shop problem. Section III motivates and defines the multi-car paint shop problem, and derives the Ising model for the problem. In Section IV we outline the experimental setup and data preparation, and Section V presents the results of our analysis. Lastly, Section VI discusses the conclusions of the scientific analysis and the lessons learned from this study.

II. PREVIOUS WORKS

The paint shop problem refers to a set of combinatorial optimization problems in the automotive industry. The objective is to color a given sequence of cars with a fixed number of colors such that the total number of color switches is minimized. This simple problem poses interesting scientific complexity-theoretic questions, which in turn have real impact for solving such problems in practice. The paint shop problem was originally posed by Epping et. al [22] as a form of coloring problem. Some clarification on nomenclature: the given car sequence can also be referred to as a word, where each car is denoted by a character. In [22], it was shown that the paint shop problem is NP-complete in both the number of colors and cars in the sequence. Furthermore, results show that, for bounded numbers of colors and unique cars, there exists a polynomial-time dynamic programming solution to these instances. Subsequent work [23] extended these results,

proving that even the simplest coloring version, with only two colors, is both NP-complete and APX-hard. Additional results show that a subset of problems meeting specific conditions can be solved in polynomial time.

Streif et. al [21] investigated solving the restricted problem, where each ensemble of cars is exactly one pair and only two colors are used (binary paint shop problem, BPSP), using a quantum algorithm. Specifically it was shown that a Quantum Approximate Optimization Algorithm (QAOA) of fixed depth can outperform classical heuristics for solving the BPSP in the infinite size limit. The authors investigate multiple classical heuristics for solving the BPSP and give bounds on the asymptotic performance of each to show this result.

III. MULTI-CAR PAINT SHOP OPTIMIZATION

One of the steps in car production at Volkswagen is painting the car body before assembly. In general, this can be viewed as a queue of car bodies that enter the paint shop, undergo the painting procedure, and exit the paint shop. It is important to note that the area of the factory immediately proceeding the paint shop is typically assembly, where car components are assembled into the car bodies. Because of the many different models and configurations being produced, designing a sequence of cars to be assembled that is optimal (also known as the car sequencing problem) is a known NP-hard problem in itself [24]. In practice it is imperative to solve the sequencing problem because of worker safety and regulatory issues, and we therefore treat the sequence of cars entering the paint shop as a fixed queue. However, the *colors* assigned to the cars within a given sequence are still randomly distributed, and thus we can focus on optimizing them.

Each car body entering the paint shop is painted independently in two steps: the first layer is called the filler, which covers the car body with an initial coat of paint, and the final color layer is the base coat, which is painted on top of the filler. The base coat is the color that matches the final customer order: blue, green, etc. However, the filler has only two possible colors: white for the lighter base coats colors, and black for darker colors. We define a customer order as the number of cars of each configuration to be painted one of the color choices. We associate each coating step (filler or base coat) with a unique class of paint shop problems, each of which are NP-complete [22], [23]. The filler optimization is referred to as the *multi-car paint shop* problem, where the cardinality of each set of unique cars in the sequence is unconstrained, but the cardinality of the color set is restricted to two (e.g., black or white). The base coat optimization is therefore an extension of the MCPS problem, where the cardinality of the color set is also unconstrained— we call this version of the problem the *multi-car multi-color paint shop* problem. In our work we focus on the filler optimization, the MCPS problem, which can be formulated natively as a binary optimization problem. We formally define the problem as follows:

Given:
a word w defining the fixed sequence of N cars (w_i denotes the i th character in w),
set of $C = \{C_1, \dots, C_M\}$ unique car ensembles,
binary choice of colors $\{W, B\}$,
function $k(C_i)$ which defines the number of C_i to be painted B in w ,
function $f(w)$ to count the number of color switches in w ,
such that: $\#C_i|_B = k(C_i), \forall C_i \in C$,
minimize: $f(w)$.

In practice (i.e., in the paint shop), the information required to formulate this optimization problem is always available, as it is a necessary part of fulfilling customer orders. Therefore, this MCPS problem representation above corresponds exactly to the industrial use-case of paint shop optimization on the filler line. In Fig. 1 we show a simple example of the MCPS problem with three car ensembles.

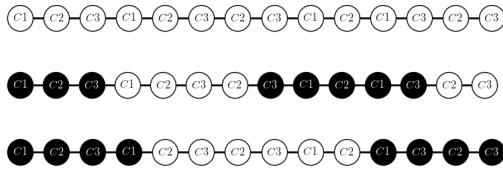


Fig. 1. Simple example of a multi-car paint shop problem with three car ensembles (C_1, C_2, C_3). The three corresponding orders are $k(C_1) = 3$, $k(C_2) = 2$, and $k(C_3) = 3$. **Top:** The fixed sequence of cars in the paint shop queue. **Middle:** Sub-optimal solution to the problem with 3 color switches. **Bottom:** Optimal solution to the problem with 2 color switches.

Despite the focus of our work on the filler, we briefly address the optimization of the base coat. Although the filler and base coat are painted independently in separate locations, computationally the two problems are not separable. Abstractly, optimizing the base coat line (i.e., solving the multi-car multi-color paint shop problem) is a straightforward generalization of the MCPS problem: we can extend the binary color variables to discrete color variables, where $k(C_i)$ denotes the number of times each color appears in a car ensemble. In practice, it is useful to consider solving the multi-color problem *after* solving the MCPS problem. Due to the aforementioned one-to-one mapping between base coat and filler color, it is still possible to permute the order of base coat colors within a contiguous sequence of filler colors to reduce the base coat color switches. Although discrete optimization problems can be represented as binary optimization problems, we do not solve the multi-color version of the problem and leave this work for future studies.

A. Ising model representation

In principle, the formulation of the MCPS as an Ising model is straightforward: we start by representing every car in the sequence w (w_i) with a single spin variable (s_i). The spin up state denotes if the car is painted black, and the spin down

state denotes if it is white. The Ising model which represents our problem can be divided into a hard constraint component and an optimization component. The optimization component is a simple Ising ferromagnet with $J = -1$ couplings between adjacent cars in the sequence:

$$H_A = - \sum_{i=0}^{N-2} s_i s_{i+1}. \quad (3)$$

This incentivizes adjacent cars to have the same color. The second component of our Ising model is the hard constraint, ensuring that the correct number of cars are colored white/black per customer orders. This is encoded in a second energy function, as a sum over independent k -hot constraint for each ensemble of cars C_i :

$$H_B(C_i) = (\#C_i - 2k(C_i)) \sum_i s_i + \sum_{i < j} s_i s_j. \quad (4)$$

Therefore, the final Ising model is the sum of the two components, $H_{\text{MCPS}} = H_A + \lambda H_B$:

$$H_{\text{MCPS}} = - \sum_{i=0}^{N-2} s_i s_{i+1} + \lambda \sum_{C_i \in C} \left[(\#C_i - 2k(C_i)) \sum_i s_i + \sum_{i < j} s_i s_j \right], \quad (5)$$

with terms as previously defined in the MCPS problem statement. In order to ensure only valid configurations of spins are encoded in the ground state, it is necessary to scale H_B by the factor λ [19]. This enforces that it is never energetically favorable to violate a constraint to reduce the energy of the system. In our study we set $\lambda = N$, the total number of cars in our sequence, which guarantees this condition.

IV. EXPERIMENTAL METHODS

A. Data sources

The MCPS problem instances we used were generated from real data taken from a Volkswagen paint shop in Wolfsburg, Germany. The reason for this is two-fold: firstly, the main goal of this work is to test the viability of quantum annealing methods in solving industrial optimization problems. It is our goal to accurately capture the complexity of the industrial use-case without relying on simplifications or randomly-generated problem instances. Secondly, the paint shop currently operates on a first-come-first-served basis, where customer orders are entered into the queue as soon as they arrive. This guarantees a certain amount of randomness (although not uniformity) in the problem instances we solve: different car models, configurations, and base coat colors appear throughout the sequences in ways we do not control. Therefore, these conditions are suitable for our analysis.

A total of 104,334 cars were used in this study. To reproduce real-world conditions as faithfully as possible, the car sequences used are multiple independent sets of car sequences, each representing one week of continuous production, which

are stitched together as one continuous block. The data is collected over a period of one year, roughly once every six weeks, to avoid seasonal biases in customer order preferences. There are 121 unique car configurations in the data set. Of those, 13 of the configurations appeared only once in the data set, and therefore do not need optimization at all. Typically this indicates that either a custom configuration was built that cannot normally be ordered, or a prototype assembled for testing purposes. Rather than exclude these from the study, we include them in the optimization, because fixing a color at one location in the sequence influences the adjacent cars (and consequently the total number of color switches) in a non-trivial way. We do, however, eliminate the spin variable from the Ising model by conditioning on the color.

B. MCPS problem sizes

To generate a variety of input sizes from the full data set, we partition the data into different sized sequences without permuting the car order. This is motivated by the amount of cars that need to be optimized for different purposes. For example, the paint shop used as a basis for this analysis has a queue capacity of roughly 300 cars. This is not the *total* capacity of the paint shop, but rather the maximum number of cars that can be physically inside the paint shop queue before they are painted. We consider this a rough lower bound on the problem size for industrially-relevant problem instances. An upper bound is more difficult to establish. From a theoretical point of view, there is merit in investigating the behavior of quantum systems in the infinite size limit, as in [21]. From an industrial perspective, car orders can be placed weeks to months in advance, which would yield problem sizes of $10^3 - 10^5$ variables. In reality, real-time and last-minute adjustments (due to manufacturing problems, supply chain issues, or imperfections in painting) can happen on a daily basis. We limit the analysis to problems of up to 3000 variables, which roughly corresponds to a few days worth of production. The data partitioning is performed by dividing the entire data set into equal chunks for each problem size N . Each partition is considered a candidate instance, yielding a total of $\lfloor \frac{104,344}{N_{\text{cars}}} \rfloor$ partitions per problem size. We then mine these partitions and select suitable partitions to generate problem instances from. This is due to the fact that there could be little frustration within some partitions. For example, it is possible that all cars of any one configuration (C_i) *all* need to be painted either black or white. While in general this is an accurate reflection of production, for experimental analyses this scenario is not useful. Therefore we deem a data partition to be a usable MCPS instance if the total number of non-fixed cars is at least 70% of the cars in the partition. Meaning that, in a 10-car data partition, at least 7 of the cars must have the freedom of being painted either color. We show the total number of partitions for the various problem sizes and how many partitions were valid problem instances in Table I. For our experiments we randomly selected 50 valid instances to test at each N , except for the largest problem size of which we use all 34 valid instances.

TABLE I
PROBLEM SIZES AND NUMBER OF PROBLEM INSTANCES GENERATED FROM THE DATA SET PARTITIONS.

Problem size (cars)	Num. partitions	Num. instances (% of partitions)
10	10,433	172 (1.6%)
30	3,477	418 (12.0%)
100	1,043	756 (72.5%)
300	347	341 (98.3%)
1000	104	102 (98.1%)
3000	34	34 (100%)

C. Classical, quantum, and hybrid solvers

We now review the solvers used in our experiments:

Random. Without optimization, we consider any assignment of colors to cars in a sequence where the orders are fulfilled to be a valid, but not necessarily optimal, solution to the problem. Thus, we can trivially generate random sets of valid solutions by uniformly assigning the color black to $k(C_i)$ cars for each car ensemble C_i . While far from optimal, this solution to the MCPS problem may be preferable if other steps in the car manufacturing process are valued over the painting step. This was indeed the case for the data obtained from the paint shop in Wolfsburg, and thus random valid solutions serve as the baseline the competition algorithms are tasked with beating. For our study, we generate $2N_{\text{cars}}$ random valid solutions at each problem size to estimate the number of color switches that would occur naturally in the paint shop.

Black-first. This is the simplest algorithm that is used to solve the MCPS problem. Starting at the beginning of the sequence, we greedily assign the color black to every car until a white color *must* be assigned to the next car, meaning all orders $k(C_i)$ have been fulfilled. In greedily obtained solutions the number of color switches grows linearly with the number of cars and is sub-optimal except for a minority of cases [23], [21]. Nonetheless, it serves as a good benchmark for more sophisticated optimization algorithms, as the improvement over random grows with the problem size.

Simulated annealing. The simulated annealing (SA) metaheuristic is a well-known optimization routine inspired by the physical annealing of metals [25]. Starting from an initial random state, potential solutions' individual variables are flipped randomly according to an acceptance criterion at every step (called a sweep) based on a “temperature” parameter. As this temperature is “cooled”, it becomes increasingly unlikely that an energetically unfavorable move is accepted, at which point the search is terminated. This algorithm has been used in previous quantum computing benchmarking studies [26], [27], [28]. We use the open-source D-Wave Python library Dimod implementation of SA [29].

Tabu search. Another metaheuristic used to solve QUBOs based on individual variable flips is called Tabu search [30], [31]. In each candidate solution, variables' states are flipped based on the likelihood they contribute to the optimality of the solution, and revisiting the same variables is discouraged, based on the “tabu tenure” list. Flips which worsen the

solution quality are admissible if no other move is possible, which provides a trade-off between quality of solution and globality of the search. The Dimod Python library contains an implementation of Tabu search as well [29].

D-Wave 2000Q. D-Wave QPUs have been tested extensively in literature by embedding Ising models directly onto the processors using an approach of minor-embedding. This process maps logical variables in the Ising model to chains of qubits in the target graph. The maximum size of problems which can then be solved directly by the QPU is limited by both the number of qubits in the processor and their connectivity. The D-Wave 2000Q generation of QPUs have a relatively sparse connectivity graph, named Chimera, where each qubit has at most degree 6. The processor used in this study had 2041 functional qubits.

D-Wave Advantage. The newest-generation QPU provided by D-Wave has a different topology and a significantly higher qubit count than its predecessor. The new topology, named Pegasus, has a maximum degree of 15, and the QPU used in our study contained 5436 functional qubits. For further information regarding the D-Wave QPU topologies and the differences between them, we refer the reader to [32].

D-Wave Hybrid Solver. At large instance sizes (300 cars and higher) it was no longer possible to embed problems directly onto both D-Wave QPUs. Therefore we employed a state-of-the-art hybrid quantum-classical algorithm developed and maintained by D-Wave System called the Hybrid Solver Service (HSS). This algorithm is designed to solve arbitrarily-structured QUBOs and Ising models of up to 10^4 variables, while also leveraging access to a QPU in its inner loop. It is accessible via the same API as D-Wave’s QPUs. Despite the power of such a hybrid algorithm, the QPU it uses cannot be programmed directly by the user, and thus we treat this algorithm as an optimization black-box with a single timeout parameter.

V. RESULTS AND DISCUSSION

We interpret our results relative to two different regimes: small-scale (10-100 cars) and industrial (300-3000 cars). Each solver used in these experiments was tuned in good-faith, but not necessarily optimally. Meaning, considerable effort was made to ensure solvers were being used to their strengths, but fully optimizing over all sets of hyperparameters for the solvers was deemed out of scope. We compared all solvers’ performance in terms of their “improvement” over the random solver, defined as the difference in $f(w)$ between the best solution obtained by each solver and the random solver. This metric is representative of the real-world expected improvement using each of the solvers. In Table. II we report the median for each solver. We quote the median to be less susceptible to tails of the distribution. The results therefore represent the typical MCPS case at each problem size, rather than the expected value of each solver’s overlap with the ground states. The solutions obtained from all solvers were post-processed (if needed) to ensure that the $k(C_i)$ constraint was satisfied in each problem. We show how often this occurred in Table II

as well. We consider this necessary in order to interpret the solutions relative to the application, since it is trivial to reduce the number of color switches by ignoring the customer order constraint. In Fig. 2 we highlight the empirical scaling of our results in the industrial limit.

Solving problems directly with both QPUs was only possible for problem sizes 10-100. Embeddings were generated using the standard D-Wave embedding tool Python package [33]. The chain strengths required by each QPU was different depending on the length of the chains in the embeddings. We calculated the algebraic chain strength $\text{chain}_i = |h_i| + \sum_{j \in \text{adj}(i)} |J_{ij}|$ for every spin in the Ising model, and introduce an additional scaling parameter $s = [0.1, 0.2, \dots, 1]$. Thus, the chain strength is defined as $s \cdot \max(\text{chain}_i)$, where s was optimized per problem size using a subset of the instances (10 per size), and $50 \cdot N$ samples per instance. Optimal s was defined as the value which yielded the highest frequency of valid solutions relative to the constraints of the MCPS problem in Eq. 5 (not chain breaks). We found that the D-Wave Advantage QPU had optimal $s = 0.3$, whereas the D-Wave 2000Q QPU had optimal $s = 0.45$. This is consistent with the fact that the Advantage QPU required shorter chains to embed the same Ising models as the 2000Q. Each QPU was then sampled for $500 \cdot N$ samples per problem size N , with annealing time $t_a = 1\mu\text{s}$. We used N spin-reversal transforms for each sample set, as it has been shown that there are diminishing returns between 100-1000 samples per transform [34].

We found that for the 10 car instances both QPUs matched the consensus best results between all solvers (median of $f(w) = 2$), and for the 30 car instances very near the best results ($f(w) = 5$ as opposed to SA’s and HSS’s 4). For the 100 car instances, with 50,000 samples per problem, the 2000Q QPU found valid solutions for 22/50 instances, as opposed to 37/50 for the Advantage QPU. From this we conclude that 50,000 samples is insufficient for the QPUs, but due to limited time availability we could not take more samples. We note that it was also possible to embed 47/50 of the 300 car instances onto the Advantage QPU. However, due to the poor performance on smaller sizes with limited resources, we did not evaluate those problems. Furthermore, due to the limited problem sizes that could be solved with the QPUs and the quality of results, we do not include these results in Fig. 2.

The Tabu solver was given $\lfloor N/3 \rfloor$ seconds per problem as its timeout parameter, and all other parameters were set to their default value. This solver struggled to find valid solutions past the 10 car instances. Due to the post-processing technique, which greedily corrected each sample to satisfy order $k(C_i)$, the Tabu results were essentially a worse version of the greedy algorithm at all problem sizes but the smallest.

Simulated annealing (SA) has many tunable hyperparameters: number of sweeps (N_{sweeps}), number of samples (N_{samples}), and (inverse) temperature (β) schedule. In our experiments we fixed the schedule to $\beta = [0.01, 10]$, interpolated geometrically using N_{sweeps} . We set $N_{\text{sweeps}} = 10 \cdot N$ and $N_{\text{samples}} = 20 \cdot N$,

where N is the number of cars. The solutions obtained by SA (shown in Fig. 2) were consistently better than the greedy algorithm. Using the given parameters SA was able to provide valid solutions for 50/50, 49/50, and 44/50 for the 30, 50, and 100 car instances. However, the timescales necessary to obtain results were prohibitive from extending the experiments: 300 variable problems were terminated after running for 24 hours without returning a solution. We include the SA results in Fig. 2 due to their high quality at small sizes. Using a single-threaded SA implementation, run-time of the algorithm was on the order of seconds to minutes for the 10 and 30 car instances, and between 1-3 hours for the 100 car instances. We note that SA was the only solver which was allotted quadratically scaling computing resources: both sweeps and samples scaled with N . This is necessary for SA to be competitive, and exemplifies the trade-off between results quality and algorithmic run-time when using heuristics. The D-Wave HSS was given equal time to the Tabu solver, given its only parameter is the timeout: $\lfloor N/3 \rfloor$ seconds per problem. The HSS was the only solver to consistently provide better solutions than the greedy algorithm for all problem sizes. The improvement continued to grow with increasing problem size, shown in Fig. 2. However, the gap between the HSS and the greedy algorithm shrank with increasing problem size, performing only slightly better than greedy for the 3000 car instances.

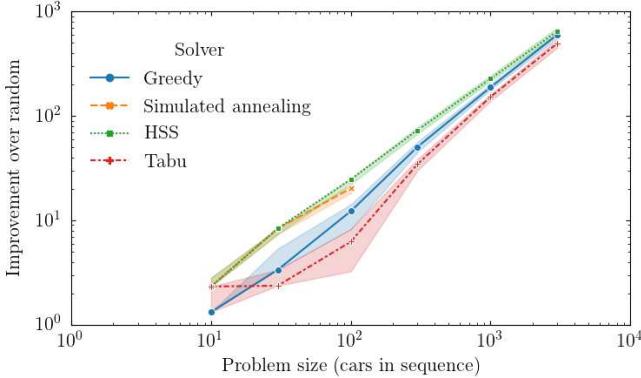


Fig. 2. Number of color switches within the sequence shown as improvement over random configurations of orders.

Despite the simplicity of the Ising model representation of the MCPS problem, almost all solvers exhibited difficulty in finding valid solutions. This is particularly evident in the performance of the two QPU models. The results degrading rapidly from 30 to 100 cars indicate that the problem became more difficult to solve disproportionately to the increase in system size. The SA and Tabu solvers exhibited similar trends, despite the increase in resources allotted to them. While the HSS was the best-performing algorithm, it also missed valid solutions for some problems of intermediate size (100 and 300 car instances). We identify two possible issues: first is the connectivity of the problem graph. Each order $k(C_i)$ requires a separate k -hot constraint which is represented using

a fully-connected graph. This yields sub-cliques within each problem that increase with the problem size. In QPUs, denser problems create longer chains and higher chain strengths. For classical solvers, these sub-cliques create rugged landscapes which make single-flip optimization algorithms significantly less useful. Therefore, that the maximum sub-clique of the MCPS problem graph increases as a function of the number of cars is a bottleneck for performance. Secondly, the normalization terms necessary to encode the MCPS problem as an Ising model also scale with problem size. In Eq. 5, we set $\lambda = N$ to ensure the constraints are valid in the ground state of the Ising problem. Therefore, we observe that the numerical precision necessary to formulate the MCPS problem scales as $1/N$. This effect compresses the gaps between the local (and global) minima, making it harder to differentiate between them. Using a direct embedding approach for QPUs requires even higher precision due to chains: the chain strengths scale with N , and therefore the encoding precision as $1/N^2$. This may be prohibitively low in large-scale problems and analog devices.

VI. CONCLUSIONS

In this paper we presented a potential application for quantum computing, the multi-car paint shop problem. We used real-world data from a paint shop in Wolfsburg, Germany, to generate problem instances and benchmark the performance of competition QUBO and Ising solvers. We found that for small problems of up to 30 cars the two QPUs tested, the hybrid algorithm, and simulated annealing were able to significantly improve results over production conditions (random distributions of colors). At intermediate and large problem sizes (100 cars and up), only the D-Wave HSS was able to consistently beat the greedy algorithm. Although the improvement over random grew with system size, the performance of the HSS and the greedy algorithm converged in the large size limit. We find that our methods demonstrated the potential of quantum annealing, and specifically the viability of hybrid quantum-classical approaches in solving industrially relevant optimization problems. We identified potential bottlenecks in performance and gave insights into how the formulation of the MCPS problem—both in terms of problem graph structure and precision—could affect the performance of the different solvers we used in our experiments. We note that in our experimental setup, we only tested whether the presented algorithms were capable of solving the industrial problems we generated, and make no claims in regards to the ultimate scaling of these algorithms. We defer a more thorough evaluation of algorithmic scaling and performance optimality for future work.

Overall, we found that the performance of the HSS, D-Wave 2000Q QPU, and Advantage QPU show the potential of quantum annealing to solve industrial problems. Specifically, we note that the instance sizes tested using the QPUs are approaching the industrially-relevant limit. In the future, we will work on finding methods to mitigate the bottlenecks presented in this work, and other potential applications that lend themselves naturally to quantum optimization.

TABLE II

RESULTS FOR ALL SOLVERS. WE PRESENT THE PERCENTAGE OF PROBLEMS FOR WHICH EACH SOLVER FOUND VALID SOLUTIONS ($\%_{\text{valid}}$) AND THE MEDIAN NUMBER OF COLOR SWITCHES ($f(w)$) FOR EACH PROBLEM SIZE N CARS.

	Greedy		HSS		Tabu		SA		2000Q		Advantage		Random	
	$\%_{\text{valid}}$	$f(w)$												
$N = 10$	100%	3	100%	2	86%	2	100%	2	100%	2	100%	2	100%	4.325
$N = 30$	100%	9	100%	4	26%	10	98%	4	100%	5	100%	5	100%	12.367
$N = 100$	100%	23	92%	10.5	0%	29	88%	15	44%	31	74%	28.5	100%	35.25
$N = 300$	100%	57	96%	34.5	0%	73	-	-	-	-	-	100%	107.58	
$N = 1000$	100%	160	100%	121	0%	196	-	-	-	-	-	100%	348.35	
$N = 3000$	100%	455	100%	406.5	0%	558.5	-	-	-	-	-	100%	1054.41	

REFERENCES

- [1] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994.
- [2] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC ’96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.
- [3] Frank. Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [4] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [5] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. F. Bogorin, M. Brink, L. Capelluto, O. Günlük, T. Itoko, N. Kanazawa, et al. Demonstration of quantum volume 64 on a superconducting quantum computing system. *Quantum Science and Technology*, 6(2):025020, 2021.
- [6] T. Kadowaki and H. Nishimori. Quantum annealing in the transverse ising model. *Phys. Rev. E*, 58:5355–5363, Nov 1998.
- [7] A. D. King et al. Observation of topological phenomena in a programmable lattice of 1,800 qubits. *Nature*, 560(7719):456–460, Aug 2018.
- [8] R. Harris et al. Phase transitions in a programmable quantum spin glass simulator. *Science*, 361(6398):162–165, 2018.
- [9] M. Streif, F. Neukart, and M. Leib. Solving quantum chemistry problems with a d-wave quantum annealer. In Sebastian Feld and Claudia Linnhoff-Popien, editors, *Quantum Technology and Optimization Problems*, pages 111–122, Cham, 2019. Springer International Publishing.
- [10] F. Neukart, D. Von Dollen, and C. Seidel. Quantum-assisted cluster analysis on a quantum annealing device. *Frontiers in Physics*, 6:55, 2018.
- [11] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko. Quantum boltzmann machine. *Physical Review X*, 8(2), May 2018.
- [12] S. Yarkoni, A. Kleshchonok, Y. Dzerin, F. Neukart, and M. Hilbert. Semi-supervised time series classification method for quantum computing. *Quantum Machine Intelligence*, 3(1):12, Apr 2021.
- [13] D. Venturelli, D. JJ Marchand, and G. Rojo. Quantum annealing implementation of job-shop scheduling. *arXiv preprint arXiv:1506.08479*, 2015.
- [14] F. Neukart, G. Compostella, C. Seidel, D. von Dollen, S. Yarkoni, and B. Parney. Traffic flow optimization using a quantum annealer. *Frontiers in ICT*, 4:29, 2017.
- [15] T. Stollenwerk, B. O’Gorman, D. Venturelli, S. Mandrà, O. Rodionova, H. Ng, B. Sridhar, E. G. Rieffel, and R. Biswas. Quantum annealing applied to de-conflicting optimal trajectories for air traffic management. *IEEE Transactions on Intelligent Transportation Systems*, 21(1):285–297, 2020.
- [16] D. Venturelli and A. Kondratyev. Reverse quantum annealing approach to portfolio optimization problems. *Quantum Machine Intelligence*, 1(1):17–30, 2019.
- [17] N. Nishimura, K. Tanahashi, K. Suganuma, M. J. Miyama, and M. Ohzeki. Item listing optimization for e-commerce websites based on diversity. *Frontiers in Computer Science*, 1:2, 2019.
- [18] F. Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241, 1982.
- [19] Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2:5, 2014.
- [20] M. W. Johnson et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, 2011.
- [21] M. Streif, S. Yarkoni, A. Skolik, F. Neukart, and M. Leib. Beating classical heuristics for the binary paint shop problem with the quantum approximate optimization algorithm. *Phys. Rev. A*, 104:012403, Jul 2021.
- [22] T. Epping, W. Hochstättler, and P. Oertel. Complexity results on a paint shop problem. *Discrete Applied Mathematics*, 136(2):217–226, 2004. The 1st Cologne-Twente Workshop on Graphs and Combinatorial Optimization.
- [23] P. Bonsma, T. Epping, and W. Hochstättler. Complexity results on restricted instances of a paint shop problem for words. *Discrete Appl. Math.*, 154(9):1335–1343, June 2006.
- [24] Tamás Kis. On the complexity of the car sequencing problem. *Operations Research Letters*, 32(4):331–335, 2004.
- [25] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [26] J. King, S. Yarkoni, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch. Benchmarking a quantum annealing processor with the time-to-target metric, 2015.
- [27] H. G. Katzgraber, F. Hamze, Z. Zhu, A. J. Ochoa, and H. Munoz-Bauza. Seeking quantum speedup through spin glasses: The good, the bad, and the ugly. *Phys. Rev. X*, 5:031026, Sep 2015.
- [28] J. King, S. Yarkoni, J. Raymond, I. Ozfidan, A. D. King, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch. Quantum annealing amid local ruggedness and global frustration. *Journal of the Physical Society of Japan*, 88(6):061007, 2019.
- [29] D-Wave Systems has produced an open-source library in Python (dimod) for solvers that optimize QUBOs and Ising Hamiltonians. More information can be found here: https://docs.ocean.dwavesys.com/en/stable/docs_dimod/.
- [30] Fred Glover. Tabu search—part I. *ORSA Journal on Computing*, 1(3):190–206, August 1989.
- [31] Fred Glover. Tabu search—part II. *ORSA Journal on Computing*, 2(1):4–32, February 1990.
- [32] K. Boothby, P. Bunyk, J. Raymond, and A. Roy. Next-Generation Topology of D-Wave Quantum Processors, 2020.
- [33] J. Cai, W. G. Macready, and A. Roy. A practical heuristic for finding graph minors. Code available at: <https://github.com/dwavesystems/minorminer>.
- [34] J. Raymond, S. Yarkoni, and E. Andriyash. Global warming: Temperature estimation in annealers. *Frontiers in ICT*, 3:23, 2016.