



Accuracy and minor embedding in subqubo decomposition with fully connected large problems: a case study about the number partitioning problem

Luca Asproni¹ · Davide Caputo^{1,2} · Blanca Silva¹ · Giovanni Fazzi³ · Marco Magagnini³

Received: 27 August 2019 / Accepted: 26 February 2020 / Published online: 19 March 2020
© The Author(s) 2020

Abstract

In this work, we investigate the capabilities of a hybrid quantum-classical procedure to explore the solution space using the D-Wave 2000QTM quantum annealer device. Here, we study the ability of the quantum hardware to solve the number partitioning problem, a well-known NP-hard optimization model that poses some challenges typical of those encountered in real-world applications. This represents one of the most complex scenario in terms of qubits connectivity and, by increasing the input problem size, we analyze the scaling properties of the quantum-classical workflow. We find remarkable results in most instances of the model; for the most complex ones, we investigate further the D-Wave Hybrid suite. Specifically, we were able to find the optimal solutions even in the worst cases by fine-tuning the parameters that schedule the annealing time and allowing a pause in the annealing cycle.

Keywords Quantum annealing · Quantum optimization · QUBO · Partitioning model

1 Introduction

Recently, the availability for the first time of quantum annealing devices from D-Wave Systems has captured the attention of both researchers and technology companies (Booth et al. 2018; Venturelli et al. 2018; Khoshaman et al. 2018; Stollenwerk et al. 2019). Besides, a growing interest is in the experimental determination of whether or not a quantum speedup can be achieved with this new class of quantum devices and what kind of working applications can be developed on such platforms (Venturelli et al. 2015; Neukart et al. 2017; Vahdat 2017; Hamze et al. 2018).

The participation of major technology powers such as Google, Lockheed Martin, and Los Alamos Laboratories continues rising, together with the scientific literature and application reports. Nevertheless, there is still a strong limitation in the usage of this model of computation for

solving real-world problems due to the limited number of qubits and couplers inside the quantum processor unit (QPU). Actually, it is well known that quantum annealers need to have a dramatically larger number of qubits and couplers in order to model the complexity of real-life problems. Especially, the limited connectivity between qubits inside the current Chimera graph architecture represents an additional obstacle in mapping large real problems in the QPU (Chancellor et al. 2016a; Chancellor et al. 2016b; Pudenz et al. 2014; Ojas et al. 2016).

Furthermore, with the release of an open-source suite spanning from the decomposing solver Qbsolv to the new Hybrid framework, D-Wave took a significant step forward towards gathering the attention from technology companies. As a matter of fact, with these technologies, it is possible to close the gap between logical qubits representation encoded in the QUBO (Quadratic Unconstrained Binary Optimization) matrix and the physical embedding of the problem into the Chimera graph (Kumar et al. 2018).

Besides, it is possible to decompose large problems into smaller subsets in such a way that they can be integrated immediately into the QPU, by providing both the combinatorial implementation required for the physical embedding and the decomposition procedure for the creation of the smaller instances. Also, the backend to be used during the computation can be specified in order

✉ Davide Caputo
da.caputo@reply.it

¹ Data Reply s.r.l., Via Nizza, 262, 10126, Turin, Italy

² Department of Mathematics and Physics, University of Salento, Via Arnesano, 73100, Lecce, Italy

³ Quantum Computing Community Of Practice, Reply S.p.A., Corso Francia, 110, 10143, Turin, Italy

to solve the model by means of either a classical or a quantum-based platform.

However, despite all the attention drawn to this crucial tool, a systematic investigation of the efficiency related to the optimization and decomposition performance has not been exhaustively conducted yet. Some studies have been developed using special techniques such as the time-to-target metric (James et al. 2015) or applying methods based on the matrix factorization (O'Malley et al. 2018) but without taking into account the capabilities of scaling up when the input size grows.

A relatively novel approach as alternative to exploiting quantum annealers to solve optimization problems is one that focuses on coherent Ising machines (CIM) (Yamamoto et al. 2017). It has been empirically shown that there exists a difference in the efficiency of such techniques, which mainly depends on the graph density that describes the problem under study, leading to the suggestion of using an approach based on CIMs to solve general Ising models (Hamerly et al. 2019). However, neither is it established that CIMs outperform quantum annealers (McGeoch et al. 2018) nor a thorough study has been pursued yet on the performances of the latter machines taking into account the possibility of scheduling the annealing process in complex scenarios.

In this work, we investigate the accuracy and the capability of the D-Wave 2000Q quantum annealer to solve problems with a significantly large input. To perform this study, we use one well-known NP-hard model: the number partitioning problem (NPP) (Mertens 2006). Thanks to the simplicity of this problem, it is easy to generate artificial instances of any size for which the optimal solution is known. Consequently, the measurement of the quality of the solution provided by the quantum annealer, along with the classical implementation of the tabu-search algorithm for the problem decomposition, will be possible even for large datasets.

2 Number partitioning problem and quantum annealing

The number partitioning problem (NPP) is defined as the task of discriminating if a given set S of positive integer numbers can be divided (partitioned) into two subsets S_1 and S_2 where the total sum of the elements in S_1 equals the total sum of the elements in S_2 . Although the NPP is an NP-complete problem, the optimization version is considered NP-hard and can be formulated in the following way: given a list of N positive integers $\{a_1, a_2, \dots, a_N\}$, the solution consists in finding a subset $A \subset \{a_1, a_2, \dots, a_N\}$ such that the difference

$$D(A) = \left| \sum_{i \in A} a_i - \sum_{i \notin A} a_i \right|, \quad (1)$$

is minimized. Throughout this work, we will refer to this difference as the *delta* between the two subsets A and $S \setminus A$. This problem is of both practical and theoretical importance: possible real applications span from multiprocessor pipeline scheduling (Kan and van Vliet 1993), where balancing and partitioning different resources can be crucial, to cryptography (Harpe et al. 1995) and all those problems requiring load balancing for I/O capacities, e.g., during databases processing (Lewis et al. 2008).

The D-Wave device implements a quantum annealing heuristics to solve sampling, optimization, and machine learning problems. Specifically, given a physical system composed of qubits, it is possible to define its Hamiltonian and initialize it in such a way that the lowest energy state corresponds to all qubits being in a superposition state of 0 and 1. Then, as the annealing proceeds, a new Hamiltonian deriving from the problem's specifications, called the problem Hamiltonian, is introduced and gradually takes over the initial energy landscape, up to a point where it contains all the energy contributions. The Hamiltonian of the system can be defined in the following way:

$$H(t) = H_I(t) + H_P(t), \quad (2)$$

where H_I is the initial Hamiltonian, H_P is the problem Hamiltonian, and their temporal evolution through the annealing is such that $H_I(0) \gg H_P(0)$ and $H_I(t_f) \ll H_P(t_f)$, being t_f the final time of the annealing.

As the problem Hamiltonian is introduced, the energy levels of the excited states are originated, increasing the probability of the system to jump from the ground state to some other excited state. In particular, there exists a critical point, the point of minimum gap, where the ground-state energy level is closest to the lowest energy level of one of the excited states. In such point, the probability of escaping the ground state is the highest, in which case the system is driven away from the global minimum.

In practice, in order to manipulate the Hamiltonian of the system, an external magnetic field is applied to the qubits. In this way, the probability of qubits falling into the 0 or 1 state is changed. The quantity that is controlled by the magnetic field, called bias or weight, is modeled by the function of the problem at hand, that is the one that has to be minimized. Moreover, it is possible to correlate qubits by entangling them. This is obtained by setting the value of a coupler, which represents the strength of the correlation between qubits that are linked together.

Hence, by letting the initial system undergo the quantum annealing process, it is possible to raise energy barriers in such a way that the energy of the system reflects the function to be minimized or maximized. If the quantum annealing is slow enough, the system is able to naturally end up in the lowest energy state, i.e., the low energy states needed

in a sampling problem or the solution of a minimization problem.

In its current implementation, the D-Wave’s quantum annealer is able to solve problems expressed in the form of an Ising glass, with a Hamiltonian written in the following form:

$$H = \sum_{i=1}^N h_i S_i + \sum_{i \neq j} c_{ij} S_i S_j \tag{3}$$

where H is the Hamiltonian encoding the problem, $S_i \in \{-1, 1\}$ are the spin values, and h_i and c_{ij} are respectively the qubits weights and the couplers coefficients of the model.

A complete formulation of the NPP as Ising spin glass has been provided in Ref. (Lucas 2014). The Hamiltonian for this type of problem can be defined by assuming an increase in the energy when the total of amplitudes associated with positive spin states is different from that of amplitudes with negative spins. According to this formulation, it is possible to use the following relation:

$$H = \left(\sum_{i=1}^N a_i S_i \right)^2 \tag{4}$$

with $S_i = \pm 1$ the spin values indicating the subset to which the i th element belongs and a_i the element of the set A . It follows that if the ground state has $H > 0$ there is no exact solution of the specific problem and the ground state is the one minimizing the mismatch between the two subsets.

Expression 4 is fundamental to write the NPP in QUBO form as it constitutes the link between the Ising Hamiltonian 3 supported by quantum annealers and the formulation of quadratic binary optimization models. The amplitude of integer numbers $a_i, \forall i = 1, \dots, N$, in the form of biases and couplers influence the value of spin variables, while elevating the sum to second power guarantees that, like the Hamiltonian of the Ising glass, this relation allows interactions up to second order. With 4, it is straightforward to obtain a QUBO model for the NPP.

In order to formulate the problem as a Quadratic Unconstrained Binary Optimization model, we first have to convert our $S_i = \pm 1$ into binary variables of the form $q_i \in \{0, 1\}$. This can be done by using the following simple relation:

$$q_i = \frac{S_i + 1}{2} \tag{5}$$

where q_i is the i th variable and S_i is the spin value. Now, the original Ising problem can be mapped into the QUBO form:

$$\min \sum_{i,j} Q_{ij} x_i x_j \tag{6}$$

where x represents a binary variable and Q is the so-called QUBO matrix containing the weights of qubits (h_i in Eq. 3)

in the diagonal and the couplers coefficients (c_{ij} in Eq. 3) in the (i, j) elements. This matrix will be symmetric ($c_i c_j = c_j c_i$), allowing a reduction in the number of variables by selecting only $i \leq j$ and setting the remaining terms to zero, leading to an upper-triangular matrix.

Having the QUBO matrix, it is possible to submit it to the QPU and retrieve a solution of the optimization problem. However, the connectivity between qubits required by the NPP is that of a complete graph, which is yet to be supported by any modern quantum annealer that provides a fairly high number of qubits. To overcome this and similar problems, the D-Wave device operates a minor embedding of the problem onto its Chimera architecture. Specifically, through the D-Wave Hybrid tool, it is possible to decompose the problem either by identifying which set of variables, if their value is changed, maximally contributes to changing the energy of the system, as we have done in this work, or by implementing a custom decomposing heuristics. Once the smaller problem is found, since the connectivity it requires may not match that of the D-Wave’s Chimera graph, one can either run the built-in tabu-search heuristics to optimally map such problem onto the annealer’s architecture, or choose a self-constructed minor embedding strategy. The QPU is able to run the annealing for the subproblems and the classical device operates a merging phase, in which the global solution of the original problem is retrieved from subproblems. Since we use the Hybrid tool default tabu-search strategy and the minor embedding process is NP-hard itself, we should stress that more sophisticated alternatives might yield better performances in terms of solution quality.

In Fig. 1, the time required to solve the NPP on classical hardware by using the D-Wave Qbsolv is reported as a function of the input set size. The elaboration time increases exponentially while a structured procedure is applied in




Fig. 1 Execution time of tabu-search for increasing input size. Classical partitioning of a set with the classical embedded tabu-search as backend. The red line is the exponential fitting $t = Ae^{x/B}$ where $B=340$ a.u. and x is the size of the input. Blue points represent measured data. The blue line in the bottom part is the deviation of the experimental point from the value of the fitting

order to find the minimum: a number of subproblems are generated, handled, and finally merged into a global solution of the NPP. The exponential increase in the execution time confirms the NP-hardness of the problem when approached with classical hardware and formulations. When the problem is submitted to the QPU, the execution time changes and paves the way for a wide range of investigations of the D-Wave Hybrid tool. Moreover, this peculiar model allows us to study what happens in one of the worst case scenarios from the perspective of the qubits connectivity: a fully connected graph, where the number of couplers and weights precision play a central role (Denchev et al. 2016; Venturelli et al. 2015).

3 Results

In order to investigate the capabilities of the D-Wave Hybrid tool, we solve multiple NPP examples of increasing size. For each fixed problem size, we use 10 different datasets and collect statistics of the results. For experimental purposes, we choose the data in such a way that the ground state of the corresponding Ising models is $H = 0$, i.e., there is a single partition of the set of numbers.

For our studies, we first construct the QUBO matrix for each problem, and then we define the tabu-search heuristics as the algorithm that splits the original problem into the subproblems, preparing them to be embedded on the Chimera graph.

Figure 2 a shows the QUBO matrix defining the connectivity of qubits required by the specific NPP instance and with regular patterns related to the number amplitudes in the dataset. With the problem being formulated as an Ising model, all variables are coupled in pairs, resulting in

a dense (upper-triangular) QUBO matrix. Such connectivity is the most complex to handle and can thus be an issue for current quantum hardwares, making it interesting to investigate the quantum annealer performance.

The distribution of partition deltas for each different problem size is summarized in Fig. 2b. We produced 10 different datasets to be partitioned for every problem size, and we computed the value of delta for all these instances. For each problem size, we have built a boxplot of deltas coming from the solution of the NPP.

The combination of quantum annealing with the classical minor embedding heuristics is able to find the optimal solution in most cases. This is achieved especially when the problem is very small (and, as a consequence, computationally easy) or when its size is significantly higher. In fact, for our smallest problem and for those with input size greater than 450 binary variables, we are able to optimally solve the 10 different NPP instances. On the other hand, for middle-sized problems, not all distributions of data allow qubits to reach the ground state. As a result, we obtain the optimal solutions only for a subset of the given problems.

From the point of view of the computational time, we do not record an outperformance of classical approaches. The reason lies in the hybridity of the algorithm: while on one hand quantum annealing takes place in only a few dozens of microseconds, on the other hand the complexity and size of our problems require decomposition and minor embedding phases to be solved by the D-Wave 2000Q, which are computationally expensive. The latter, in particular, is itself a NP-hard problem that is handled by a classical device and thus leads to a significant increase in computational time.

Figure 2 c and d report the density distribution of each of the 10 datasets used for 2 different problem sizes (200 and




Fig. 2 QUBO matrix and delta distributions over multiple datasets. **a** QUBO matrix of one instance of data with problem size equal to 100. The entries are scaled and the intensity of colors is used as a means to summarize the main characteristics of the plot: the diagonal is made up of negative values, the lower triangular part is 0 and the upper one has no null entries. **b** Boxplots of deltas for different input problem sizes

computed over 10 datasets for each size with dots representing the values of the delta in each instance. These values were saturated to 50; therefore, such numerical value is to be interpreted as the result of a bad solution. **c–d** Kernel density estimation of the distribution of input data for problems with, respectively, 200 and 500 variables, showing the data from all 10 instances in each plot

500 variables). As explained above, the quality of the results on the bigger model exceeds the one on intermediate sizes. Comparing both density distributions, we can conclude that this behavior is fundamentally related to the fact that a shift of the distribution curve to lower values leads to a dataset containing more solution degeneracy for lower energy states and consequently simpler to solve even when the problem is bigger.

An effective method to enhance the exploration of the solution space is the direct manipulation of the annealing schedule (Ottaviani and Amendola 2018; Marshall et al. 2019). This distinctive technique can be used to improve the quality of the solution in the cases described before in which we could not reach the ground state. Indeed, in contrast to what we did with the first approach, where the annealing has been used without interfering with the spontaneous process, we exploit now the capability of the D-Wave solver API to manipulate directly the scheduling of the cycle. To accomplish this, we define the time instant at which the cycle has to be stopped and resumed, as well as the value of the persistent current powering the adiabatic relaxation. This entire procedure is referred to as annealing pause.

The top panel of Fig. 3 is a sketch of the evolution over time of the initial and problem Hamiltonians, as the time-scheduled moves forward, compared with their theoretical state if no pause is scheduled. In both cases, the problem Hamiltonian grows while the initial one decreases, but in the time-schedule case we find an instant (determined by the user) when the annealing is paused and, as a consequence, the two terms of the system Hamiltonian remain constant. Once the pause is finished, the normal scheduling is resumed and continues its cycle. At the end of the process, the initial Hamiltonian vanishes and the energy of the system is determined by the problem alone.

The idea behind a possible improvement of solution quality with an annealing pause is that, by letting the system evolve under constant Hamiltonian terms for a while, the probability of following the minimum energy adiabatic pathway, and thus finding the ground state, increases (Marshall et al. 2019). For this reason, the choice of pause duration and time instant at which it should start and end play a fundamental role. Bad choices for these parameters may leave the solution quality unaffected or even worsened.

The middle and bottom panels of Fig. 3 show the results of the analysis on two problems, one of size 200 and the other of size 300, for which the uncontrolled annealing performed worst. For each problem, we have paused the annealing after 10 μs , let the system rest for 10, 40, 60, 100, and 120 μs respectively halfway through the flow of current and finally let the annealing end. This whole process was repeated 5 times for each problem.

The best energy configuration in terms of distance from the ground state for the two problems analyzed here were

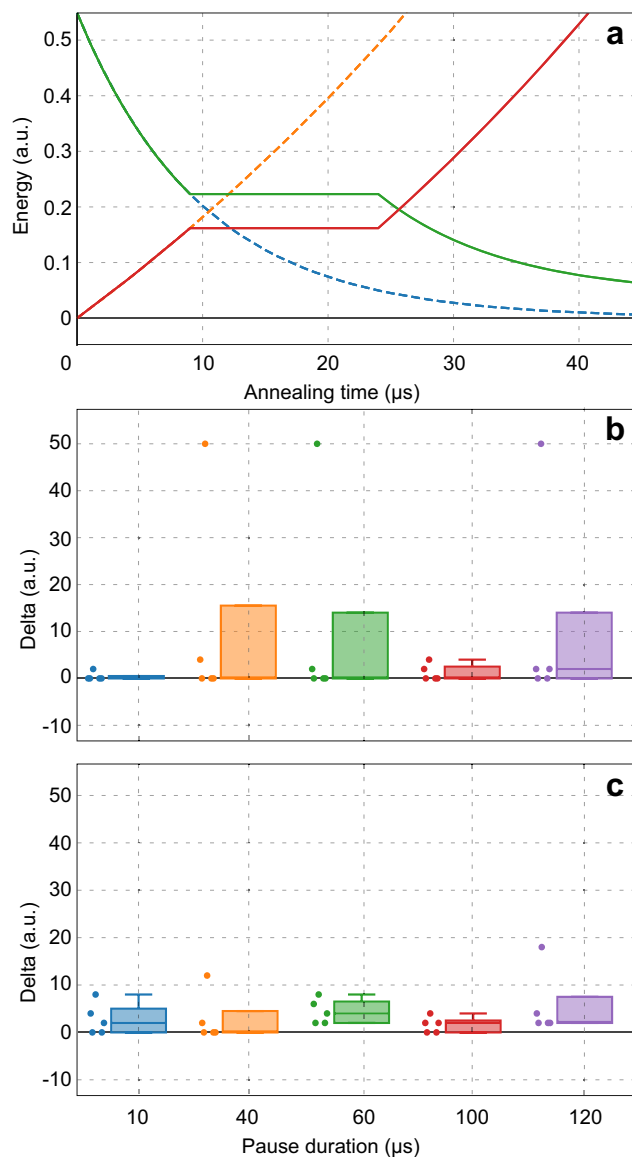


Fig. 3 Annealing cycle and boxplots of deltas with pause. **a** Sketch representation of the terms contributing to the system Hamiltonian as a function of the time. The solid red, dashed orange, solid green, and dashed blue lines represent respectively the problem Hamiltonian with and without pause and the initial Hamiltonian with and without pause. **b–c** Boxplots of deltas found over multiple runs of the annealing with the same pause starting point, same value of persistent current but different pause duration times. The dots represent the values of (eventually saturated) deltas for every run. **b** shows boxplots for one of the problems with 300 variables yielding a very bad solution, while **c** shows the same for one of the problem with 200 variables

not achieved with the same parameters settings. In fact, every instance requires different values of the pause starting point, duration, and persistent current. Nevertheless, all of our choices greatly improved the results previously obtained with the uncontrolled annealing, even though not all of them led to the optimal solution. We were able to record considerable results multiple times, proving that the

introduction of the pause can increase the accuracy of the annealing.

This improvement in the quality of the results is due to the effect of the pause on the search region of the solution space: by pausing the flow of the persistent current, and hence the annealing, we are able to widen the exploitation of the energy landscape and, as a consequence, the probability of finding the global minimum.

The parameters must be tuned wisely: too long pauses could make the system escape from energy points near the ground state, while too short pauses could not be effective at all. At the same time, if we schedule a pause after the system has overcome the minimum gap between the energy levels, i.e. it has already been in some equilibrium state, we get no benefit from this whole procedure; likewise, a pause scheduled too early has no effect on the probability of obtaining the global minimum because the chances of escaping the ground state would still be high.

It has been shown empirically that finding the appropriate time to start the pause and its duration is a technique that is likely to increase the computation performance of the quantum annealing, yielding much better solutions at the cost of only little more QPU computational time (Ottaviani and Amendola 2018).

The possibility of scheduling annealing cycles proves to be of crucial importance for complex tasks. The improvement in solution quality recorded in this work contributes to confirming quantum annealing as a valid approach to solve computationally heavy problems that are typical of a number of areas, such as artificial intelligence. Indeed, the possibility to find a very accurate global minimum is fundamental in different contexts, spanning from the reinforcement learning penalty models to the weights optimization behind the functioning of neural networks (McKiernan et al. 2019; Sun et al. 2018). In such field, the overwhelming difficulty of tasks, due both to problem size and model complexity, puts to the center of the attention the trade-off between solution quality and computational time required to solve problems. We have here empirically shown that quantum annealing, when thoroughly tuned, can be very efficient in terms of optimal solution finding and, as the hardware keeps being improved and the classical part of the algorithm becomes less expensive, is expected to reduce time issues as well.

4 Conclusions

In this work, we studied the capabilities of the D-Wave quantum annealer and the D-Wave Hybrid framework to approach problems in complex scenarios. For our analysis, we selected a fully connected model: the NPP, which poses

an enormous challenge to the currently available QPUs and the architecture they are based on.

Two different analyses were done: accuracy of the outcome when the input size scales up and the impact of the annealing pause on the solution quality.

For the first part, we conducted our analysis on a number of small-to-large size problems to investigate the behavior of the quantum annealer on a level of complexity which is potentially that of real-life problems. One interesting result was found: a discontinuous accuracy with the problem size. While high-quality results were found at small problems, there is a counter-intuitive behavior as the problem dimension increases: a dip in the accuracy for medium-sized problems and a recovery as size increases. This effect was explained by the value distribution within the dataset: lower values in the input allow higher accuracy of the result, even when the size of the problem is rising.

The medium-sized problems were studied in more detail by applying pauses during the annealing cycle, allowing the system to explore the solution space with a modified equilibrium. Our results prove that with the correct parameter tuning it is possible to improve dramatically the accuracy of the solution, obtaining optimal results in cases that had proven to be troublesome in a non-altered context.

Acknowledgments We acknowledge the support of the Universities Space Research Association, Quantum AI Lab Research Opportunity Program. Also, we thank Davide Venturelli for fruitful discussions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Booth KE, Do M, Beck JC, Rieffel E, Venturelli D, Frank J (2018) Comparing and integrating constraint programming and temporal planning for quantum circuit compilation. In: Twenty-Eighth international conference on automated planning and scheduling
- Chancellor N, Szoke S, Vinci W, Aeppli G, Warburton PA (2016a) Maximum-entropy inference with a programmable annealer. *Sci Rep* 6:22318. EP –. Article
- Chancellor N, Zohren S, Warburton PA, Benjamin SC, Roberts S (2016b) A direct mapping of max k-sat and high order parity checks to a chimera graph. *Sci Rep* 6:37107. EP –. Article

- Denchev VS, Boixo S, Isakov SV, Ding N, Babbush R, Smelyanskiy V, Martinis J, Neven H (2016) What is the computational value of finite-range tunneling? *Phys Rev X* 6:031015
- Hamerly R, Inagaki T, McMahan PL, Venturelli D, Marandi A, Onodera T, Ng E, Langrock C, Inaba K, Honjo T, Enbutsu K, Umeki T, Kasahara R, Utsunomiya S, Kako S, Kawarabayashi K-i, Byer RL, Fejer M, Mabuchi H, Englund D, Rieffel E, Takesue H, Yamamoto Y (2019) Experimental investigation of performance differences between coherent ising machines and a quantum annealer. *Sci Adv* 5(5)
- Hamze F, Jacob DC, Ochoa AJ, Perera D, Wang W, Katzgraber HG (2018) From near to eternity: spin-glass planting, tiling puzzles, and constraint-satisfaction problems. *Phys Rev E* 97(4):043303
- Harpes C, Kramer GG, Massey JL (1995) A generalization of linear cryptanalysis and the applicability of matsui's piling-up lemma. In: Guillou LC, Quisquater J-J (eds) *Advances in Cryptology — EUROCRYPT '95*. Springer, Berlin, pp 24–38
- James K, Sheir Y, Mayssam MN, Jeremy PH, Catherine CM (2015) Benchmarking a quantum annealing processor with the time-to-target metric. arXiv:1508.05087
- Kan AHGR, van Vliet A (1993) Probabilistic analysis of packing and partitioning algorithms (e. g. coffman, jr. and george s. lueker). *SIAM Rev* 35(1):153–154
- Khoshaman A, Vinci W, Denis B, Andriyash E, Amin MH (2018) Quantum variational autoencoder. *Quantum Sci Technol* 4(1):014001
- Kumar V, Bass G, Tomlin C, Dulny J (2018) Quantum annealing for combinatorial clustering. *Quantum Inf Process* 17(2):39
- Lewis M, Kochenberger G, Alidaee B (2008) A new modeling and solution approach for the set-partitioning problem. *Comput Oper Res* 35(3):807–813
- Lucas A (2014) Ising formulations of many np problems. *Front Phys* 2:5
- Marshall J, Venturelli D, Hen I, Rieffel E (2019) Power of pausing: advancing understanding of thermalization in experimental quantum annealers. *Phys Rev Appl* 11:044083
- McGeoch CC, Bernoudy W, King J (2018) Comment on "scaling advantages of all-to-all connectivity in physical annealers: the coherent ising machine vs d-wave 2000q. arXiv:1807.00089
- McKiernan KA, Davis E, Alam MS, Rigetti C (2019) Automated quantum programming via reinforcement learning for combinatorial optimization. arXiv:1908.08054
- Mertens S (2006) The easiest hard problem: number partitioning. *Computational Complexity and Statistical Physics* 125(2):125–139
- Neukart F, Compostella G, Seidel C, von Dollen D, Yarkoni S, Parney B (2017) Traffic flow optimization using a quantum annealer. *Frontiers in ICT* 4:29
- Ojas P, Jeremy W, Luke S, Andrew L, Jonathan M, Aidun J (2016) Benchmarking adiabatic quantum optimization for complex network analysis. arXiv:1604.00319
- O'Malley D, Vesselinov VV, Alexandrov BS, Alexandrov LB (2018) Nonnegative/binary matrix factorization with a d-wave quantum annealer. *PLoS One* 13(12):e0206653
- Ottaviani D, Amendola A (2018) Low rank non-negative matrix factorization with d-wave 2000q. arXiv preprint arXiv:1808.08721
- Pudenz KL, Albash T, Lidar DA (2014) Error-corrected quantum annealing with hundreds of qubits. *Nat Commun* 5:3243. EP – Article
- Stollenwerk T, O'Gorman B, Venturelli D, Mandrà S, Rodionova O, Ng H, Sridhar B, Rieffel EG, Biswas R (2019) Quantum annealing applied to de-conflicting optimal trajectories for air traffic management. *IEEE transactions on intelligent transportation systems*
- Sun H, Chen X, Shi Q, Hong M, Fu X, Sidiropoulos ND (2018) Learning to optimize: training deep neural networks for interference management. *IEEE Trans Signal Process* 66(20):5438–5453
- Vahdat A (2017) Toward robustness against label noise in training deep discriminative neural networks. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) *Advances in Neural Information Processing Systems 30*, Curran Associates, Inc., pp 5596–5605
- Venturelli D, Do M, Rieffel E, Frank J (2018) Compiling quantum circuits to realistic hardware architectures using temporal planners. *Quantum Sci Technol* 3(2):025004
- Venturelli D, Mandrà S, Knysh S, O'Gorman B, Biswas R, Smelyanskiy V (2015) Quantum optimization of fully connected spin glasses. *Phys Rev X* 5:031040
- Yamamoto Y, Aihara K, Leleu T, Kawarabayashi K-i, Kako S, Fejer M, Inoue K, Takesue H (2017) Coherent ising machines-optical neural networks operating at the quantum limit. *npj Quantum Inf* 3(1):49

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