Check for updates

Transfer of knowledge through reverse annealing: a preliminary analysis of the benefits and what to share

Eneko Osaba* and Esther Villar-Rodriguez

TECNALIA, Basque Research and Technology Alliance (BRTA), Derio, Spain

Being immersed in the noisy intermediate-scale quantum (NISQ) era, current quantum annealers present limitations for solving optimization problems efficiently. To mitigate these limitations, D-Wave Systems developed a mechanism called reverse annealing, a specific type of quantum annealing designed to perform local refinement of good states found elsewhere. Despite the research activity around reverse annealing, no study has theorized about the possible benefits related to the transfer of knowledge under this paradigm. This work moves in that direction and is driven by experimentation focused on answering two key research questions: *i) is reverse annealing a paradigm that can benefit from knowledge transfer between similar problems?* and *ii) can we infer the characteristics that an input solution should meet to help increase the probability of success?* To properly guide the tests in this paper, the well-known knapsack problem has been chosen for benchmarking purposes, using a total of 34 instances composed of 14 and 16 items.

1 Introduction

A quantum annealer (QA, [1]) is a specific kind of quantum device designed to deal with optimization problems by means of a process inspired by classical simulated annealing [2]. These computers leverage quantum mechanics to efficiently explore solution spaces in an attempt to find the optimum value of a given objective function. Today, advances in quantum technologies have contributed to the building of intermediate-scale QAs that implement quantum annealing for programmable use. There are different platforms for building quantum annealers, such as optical tweezers [3] or superconducting integrated circuits [4], with the latter being the most recognized to date. Additionally, several companies are working on this

technology and building their own devices, such as NEC[1], Qilimanjaro[2], and D-Wave Systems[3].

However, all the progress made in recent years was born in the noisy intermediate-scale quantum (NISQ) era [5], when quantum devices presented great limitations in solving optimization problems efficiently, even when these problems are small- or medium-sized. As a consequence, the community as a whole is striving to devise schemes and mechanisms to address the current limitations and take advantage of the potential that quantum computing currently offers. Among the most common strategies are the design and implementation of advanced hybrid resolution schemes, which leverage the best of both computing paradigms [6]. Other alternatives include the implementation of error mitigation strategies [7] or making use of reliable quantum simulators [8].

As part of these activities, an interesting strategy studied in works such as [9–11] involves initiating the quantum evolution with a state that is an educated guess of the problem's solution. It is interesting to note that similar strategies have also been explored in gate-based quantum computing, with the warm-start initialization of quantum approximate optimization algorithms (QAOAs) being particularly well known [12–14].

In this article, we focus on a specific strategy called reverse annealing (RA, [15]), which was introduced by D-Wave Systems in 2017. The motivation behind the implementation of RA lies in the limitations of D-Wave's annealers to conduct local refinement of states previously found either through a forward anneal or by means of a classical technique. Thus, RA allows quantum local search by annealing backward from a previously specified state and then forward to a new one.

Today, many practical studies employ RA to improve a solution previously found by other solving schemes [16, 17]. A significant amount of research has also been published studying theoretical aspects of RA, such as its sensitivity to proper parameterization, mainly regarding the RA Schedule [18, 19]. Despite the existence of this buoyant research activity, the full potential of RA is yet to be revealed, as some crucial aspects of its performance still require a deeper analysis.

On this basis and bringing to the fore the proven potential of RA, this paper is committed to extending the frontiers of this promising mechanism, exploring its application in situations not yet studied by the community. To achieve this, the analysis provided in this work is driven by experimentation focused on answering two key research questions (RQ).

As mentioned before, RA was conceived as a local refinement mechanism. For this reason, in order to solve a problem using this paradigm, the QA must be initialized with a known (classical) solution. RA explores the local space around that solution for bitstrings with even lower energy. Despite the existence of intense research activity around RA, no one has theorized about the possible benefits related to the transfer of knowledge (ToK, [20]) under the umbrella of this paradigm. So far, all studies feed the annealer with an existing solution to the same problem, with this existing solution usually obtained either by a forward annealing process executed immediately before the RA [21] or by a classical optimization method [22]. However, it would

be reasonable to think that, in a real application context where the problems faced by the industry share analogies, a previous solution may be a good starting point for the resolution of a new (or perhaps not so new, as we argue) problem. At the time of writing, the use of solutions from different problems has remained unexplored. In other words, the influence of solving a specific problem through the RA paradigm by using as input a solution from another similar but not exactly the same problem has not been studied. In this context, the first of the research questions is:

RQ1: *Is RA a paradigm that can benefit from knowledge transfer between problem instances with similar characteristics?*

In relation to the analysis of RA behavior, several studies have tried to provide insights about the requirements that input solutions must meet, determining in many cases that RA is efficient when it is fed by a solution *close* enough to the optimum of the problem to be solved. However, the term *close* is usually ambiguously used [23], with studies translating this closeness into energy [24], while others into the composition of the solution itself, that is, the difference in terms of the Hamming distance between the input solution and the optimum one [25]. Thus, the second RQ posed in this paper aims to shed light on the characteristics that an input solution should have to increase the probability of success:

RQ2: *Can we infer the characteristics that an input solution should meet to help increase the probability of succeeding in an RA process?*

Moving in this direction is especially interesting for ToK because it helps to identify which solution or solutions from the sampling pool are the most promising to feed RA and to set the analysis framework for future work in this area. For clarification purposes and following the nomenclature of ToK, the term *source* refers to the problem or task whose knowledge will be leveraged to solve a *target* task [26].

To properly guide the experimentation of this paper, the well-known knapsack problem (KP, [27]) has been used as a benchmarking problem. We have chosen the KP because:

- It has been extensively used for benchmarking purposes in QC-oriented studies [28–30].
- It is an appropriate problem to be formulated as a quadratic unconstrained binary optimization (QUBO) problem [31].
- It is a complex problem to be solved by quantum algorithms, as has been previously demonstrated in the literature [32, 33].

Two main instances were used in the planned experimentation, composed of 14 and 16 items, respectively. We call these seed instances *parent*-instances. From each of them, 16 new cases with similar characteristics have been created, hereinafter named *descendant*-instances. With all this, the main objective of the experimentation is to analyze, given a target-instance, the impact of feeding RA with solutions obtained from similar yet different tasks.

It is worth noting that the size of the instances was not chosen randomly. This selection was made after thoroughly analyzing the current state of the art in the conjunction between QA and KP. Additionally, our motivation has been to construct instances that are, on the one hand, large enough so that D-Wave's QA would not always solve them optimally through the forward annealing process and, on the other hand, small enough to be efficiently embedded in the quantum computer.

Lastly, we have used the Hamming distance as a reference measure to evaluate the difference between the two solutions. This measure stands out particularly for its simplicity and efficiency in calculation. Additionally, the Hamming distance is especially suited for measuring the similarity among binary strings, which are often used in quantum computing. It has also been frequently used to compare DNA sequences and error-correction codes.

2 Methods

As mentioned in the introduction, a QA is a specialized quantum device engineered to tackle optimization problems through a process inspired by classical simulated annealing. More specifically, a QA operates on the principle of adiabatic computation, where an initial, easily prepared Hamiltonian is gradually evolved from its ground state to the ground state of a final, problem-specific Hamiltonian. If this evolution is sufficiently slow, the adiabatic theorem ensures that the system stays in the ground state throughout the computation. In QA, the adiabatic theorem is deliberately relaxed, permitting the process to evolve more quickly than the adiabatic limit would normally permit. Consequently, transitions to higher energy states often occur during the evolution. Despite this, the computational model can also be considered universal [34]. The D-Wave annealer, which is the one used in this study, is based on an Ising Hamiltonian, which limits the kind of problems that can be executed on the computer. Nevertheless, this type of QA is as well-suited for solving combinatorial optimization problems as the one addressed in this article [35].

The RA method is a specific type of quantum annealing process designed to perform local refinement of good states found elsewhere [15]. To achieve this, RA resorts to a time-dependent Hamiltonian:

$$H(t) = A(s(t))H_0 + B(s(t))H_1 \qquad (1)$$

in which $s(t) \in [0,1]$, $A(s)$, and $B(s)$ are time-dependent amplitudes that must satisfy $A(0) \gg B(0)$ and $A(1) \ll B(1)$. Additionally, H_0 defines the initial Hamiltonian, while the final Hamiltonian H_1 corresponds to the unconstrained optimization problem, with a ground state that represents the computational solution. In contrast to forward annealing, RA reverses the time-evolution procedure by starting in an eigenstate of H_1 and evolving under Equation 1 in the opposite direction. Thus, $H(t) = H(1)$ at $t = 0$, and the process is divided into three different steps in which the Hamiltonian.

1. First, it evolves backward to a previously specified point s_p in the control schedule in a time $t_r = t_1$ coined *ramp-time*.
2. Then, it is paused for a time $t_p = t_2 - t_1$ (this step is optional).
3. It finally evolves forward from s_p at t_2 to the final Hamiltonian at time T for a $t_q = T - t_2$ *quench time*.

Therefore, the RA Schedule can be defined as shown in the following Equation 2 [18]

$$s(t) = \begin{cases} 1 + \dfrac{s_p - 1}{t_1} t, & 0 \le t \le t_1, \\ s_p, & t_1 \le t \le t_2, \\ s_p + \dfrac{1 - s_p}{T - t_2}(t - t_2), & t_2 \le t \le T. \end{cases} \qquad (2)$$

3 Experimentation and discussion

Before starting with the description of the experimentation, it is important to note that, at present, no metric accurately indicates how similar any two problems are to each other (in order to make them good candidates for ToK). The existence of such a metric would be very advantageous for this purpose. More specifically, similarity should be measured in terms of overlap in the energy landscape of a unified search space [36]. However, in the real world, it is not possible to determine this overlap beforehand without having executed and, therefore, solved the problem. Given this circumstance, similarity becomes a matter of intuition and is subject to the domain knowledge acquired by the expert. This situation does not detract an iota of value because, in real industrial contexts, most problems have a repetitive nature: a pool of recurring customers in routing problems, production of similar materials with stable task typologies and machines, etc. In such situations, ToK is a promising strategy to resort to past solutions in order to improve the results or speed up the computation of a new yet similar task.

With this in mind, the experiment carried out in this paper uses the well-known KP as a benchmarking problem. In a nutshell, KP consists of a set P composed of n items, describing each item p_i by profit (v_i) and weight (w_i), which must be packed into a container with a maximum capacity W. The objective, therefore, is to select a subset of items to be stored that, without exceeding W, maximizes the profit obtained. It is worth noting that the metric used in this study for measuring the quality of a solution corresponds to the energy provided by the quantum annealer, which should be minimized. That is, the less energy, the better the solution.

Aiming to provide an answer to the **RQ**s, we have designed two separate experiments focused on two *parent*-instances of the KP, which are composed of 14 (S14) and 16 (S16) items, respectively. Both instances have been generated *ad hoc* for this study with the values v and w randomly selected from {1,2,3,4}. Finally, $W = (\sum_{i=1}^{n} w_i)/2$.

Taking these cases, each *descendant*-instance has been generated by applying the following strategy.

- First, the newly generated *descendant*-instance starts out as an exact copy of the corresponding *parent*-instance.
- Then, a set P' is created, consisting of all the unique items in P along with their energy impact. Provided that each item p_i in P is described as a tuple v_i, w_i, a unique item in P' means that it differs from the others in at least one element of the tuple.
- Next, a certain number of items of the *descendant*-instance are modified following this criterion: when the purpose is to increase the energy impact of an item p_i, it is replaced by the next highest energetic item in P'. Analogously, p_i is replaced by the next lowest energy item in P' when the objective is to reduce the energy. It should be noted that the items with the highest and lowest energy are never modified in the creation of *descendant*-instances.

Following this strategy, 16 *descendant*-instances have been created for each *parent*-instance, equally divided into four categories.

- X_L2L: *Lowest-Energy-to-Lower-Energy*: X% of the less energetic variables (i.e., items) of the *parent*-instance are

selected and modified so that their energetic impact is even lower.

- X_L2H: *Lowest-Energy-to-Higher-Energy*: X% of the less energetic variables of the *parent*-instance are selected and modified so that their energetic impact is greater.
- X_H2L: *Highest-Energy-to-Lower-Energy*: X% of the highest energetic items of the *parent*-instance are selected and modified so that their energetic impact is lower.
- X_H2H: *Highest-Energy-to-Higher-Energy*: X% of the highest energetic variables are selected and modified so that their energetic impact is even greater.

It is noteworthy that solutions to these problems are represented by a $(n + s)$-long vector, where n is the number of items that compose the instance, and s are slack terms introduced by the penalties. For both s14 and s16, $s = 5$, so that the decision variables for each problem are 19 and 21, respectively. In our study, the variables eventually modified correspond only to the first n items.

Systematically, knowledge transfer materializes in our experimentation by taking a *descendant*-instance as the source-task and a *parent*-instance as the target-task to be solved. That is, the goal is to solve a *parent*-instance leveraging the knowledge of a previously solved *descendant*-instance.

First, all instances have been independently executed 10 times using D-Wave's Advantage_system6.4 by means of the common forward annealing process. Regarding the parameters, the number of runs was set to 1,000. Furthermore, the Knapsack class belonging to the Qiskit v0.6.0 *Optimization Applications* open library was employed for generating the KP QUBOs[4]. Thus, Table 1 includes.

- The best energy found for each instance across the 10 independent runs.
- The energy of the *parent*-instance objective function if the best solution of the *descendant*-instance were applied, and the nominal value of the difference with respect to the energy of the best solution of the *parent*-instance. For example, let us take the row related to s16_0.2_L2L; the best solution found in the s16_0.2_L2L resolution corresponds to an energy of −10,133. This solution, when applied to the objective function of the *parent*-instance, that is, s16, produces an energy outcome of −10,712. In turn, the energy of the best solution found for the s16 instance was −10,716, which causes a difference of 4.
- The Hamming distance between the best solutions of *parent*- and *descendant*-instances when executed independently. In addition, the decomposition of the Hamming distance is added. We specify, among the total number of variables that differ (h), which ones represent an item (n_h) and which ones are slack variables (s_h). Taking the same example, the Hamming distance between the best solutions found for s16_0.2_L2L and s16 is equal to 9, meaning that the bitstrings of these solutions differ in nine bits, seven of which correspond to variables representing items and two to slack variables.

---

4   https://qiskit-community.github.io/qiskit-optimization/stubs/qiskit_optimization.applications.Knapsack.html

With all this in mind, we now proceed to answer the two questions posed in the introduction.

3.1 RQ1: is RA a paradigm that can benefit from knowledge transfer between problem instances with similar characteristics?

A second set of tests has been conducted to properly answer **RQ1**, consisting of solving each target-instance s14 and s16 by means of the RA procedure and using as the input solution the best solutions found for each *descendant*-instance through the first tests summarized in Table 1.

A fixed RA schedule has been used for all the tests, [(0.0, 1.0), (2.5, 0.5), (102.5, 0.5), (102.75, 1.0)], which has been obtained through an empirical procedure performed in the lab. More precisely, an experiment was carried out in which 25 different RA schedules were tested. Minor adjustments were made to the s_p, t_r, t_p, and t_q parameters. To assess the quality of the designed RA schedule, each one was used to solve the two target-instances 10 times each, using one random solution for the problem as input. The aforementioned schedule was selected due to its superior performance compared to the others.

For the sake of fair comparison, the annealing_time of all the forward annealing processes has been adjusted to the RA schedule duration so that all runs carried out in this research access the quantum computer for the same amount of time. We recommend papers such as [18] to readers interested in the analysis of the reverse annealing schedule. Lastly, reinitialize_state=True.

For each (source instance–target instance) combination, 10 independent runs have been executed, with Table 2 depicting the best result found among these executions, along with the average and the standard deviation. As an example, 10 independent runs have been executed to solve s14, using the best solution found for s14_0.2_L2L as the input bitstring in the RA process (i.e., source of knowledge). For this combination, the energy of the best solution found is −12,421, while the average and standard deviation are −12,420.2 and 0.74, respectively. For comparison purposes, we also depict in Table 2 the baseline results obtained by s14 and s16 using forward annealing. In these cases, the source of knowledge has been represented as "–." It is important to highlight that the optimal results for each instance have been achieved using the industry-oriented Quantagonia Hybrid Solver (QHS, [6]). In essence, this algorithm is divided into two distinguishable steps: Initially, it executes a set of *primal heuristics* to solve the entire problem. These heuristics can be classical or quantum, utilizing external services. Subsequently, QHS enhances the best solution identified by the primal heuristics through a classical Branch-and-Bound algorithm. One of the key strengths of QHS is the inclusion of an *optimality proof* mechanism, which provides the *optimality gap* along with the results. The *optimality gap* indicates the remaining potential in the optimization process. This feature is the primary reason we selected QHS as the solver to achieve optimal results for the considered instances.

The results shown in Table 2 provide answers to **RQ1**. If we pay attention to s14, we could preliminarily conclude that the sharing of knowledge does not imply an improvement with respect to the best

TABLE 1 Instances used in the experimentation. *Energy wrt. sXX* represents the energy value of the best-found solution for the descendant-instances, using as a base the objective function of sXX. Hamming distances are calculated using as reference the best solution found for the *parent*-instances and the *descendant*-instances along 10 independent runs. Hamming distances are represented by the total value (h) and divided into variables that represent items (n_h) and slack variables (s_h).

Instance	Best energy	Energy wrt. s14	Hamming distance	Instance	Best energy	Energy wrt. s16	Hamming distance
s14	−12422	—	—	s16	−10716	—	—
s14_0.2_L2L	−12,019	−12,419 (3)	7 (7, 0)	s16_0.2_L2L	−10,133	−10,712 (4)	9 (7, 2)
s14_0.2_L2H	−13,222	−12,421 (1)	3 (3, 0)	s16_0.2_L2H	−11,583	−10,713 (3)	5 (4, 1)
s14_0.2_H2L	−12,021	−12,391 (31)	5 (5, 0)	s16_0.2_H2L	−10,137	−10,715 (1)	4 (3, 1)
s14_0.2_H2H	−12,820	−12,388 (34)	4 (4, 0)	s16_0.2_H2H	−10,995	−10,668 (48)	12 (11, 1)
s14_0.4_L2L	−12,820	−12,388 (34)	8 (8, 0)	s16_0.4_L2L	−8,974	−10,714 (2)	5 (4, 1)
s14_0.4_L2H	−13,620	−12,387 (35)	4 (4, 0)	s16_0.4_L2H	−11,294	−10,712 (4)	5 (5, 0)
s14L_0.4_H2	−11,218	−12,294 (128)	8 (7, 1)	s16_0.4_H2L	−9,845	−10,676 (40)	8 (8, 0)
s14_0.4_H2H	−13,621	−12,418 (4)	5 (5, 0)	s16_0.4_H2H	−11,579	−10,674 (42)	8 (7, 1)
s14_0.6_L2L	−11,620	−12,297 (125)	7 (7, 0)	s16_0.6_L2L	−8,463	−10,563 (153)	11 (10, 1)
s14_0.6_L2H	−13,620	−12,416 (6)	8 (7, 1)	s16_0.6_L2H	−12,161	−10,709 (7)	10 (9, 1)
s14_0.6_H2L	−10,418	−12,297 (125)	5 (4, 1)	s16_0.6_H2L	−10,136	−10,677 (39)	8 (7, 1)
s14_0.6_H2H	−14,020	−12,294 (128)	6 (6, 0)	s16_0.6_H2H	−11,581	−10,676 (40)	4 (3, 1)
s14_0.8_L2L	−11,619	−12,139 (287)	8 (7, 1)	s16_0.8_L2L	−8,465	−10,563 (153)	6 (5, 1)
s14_0.8_L2H	−14,418	−12,413 (9)	9 (8, 1)	s16_0.8_L2H	−12,108	−10,644 (52)	14 (11, 3)
s14_0.8_H2L	−10,819	−11,925 (497)	8 (7, 1)	s16_0.8_H2L	−10,138	−10,569 (147)	5 (5, 0)
s14_0.8_H2H	−14,413	−12,287 (135)	7 (5, 2)	s16_0.8_H2H	−11,582	−10,710 (6)	7 (7, 0)

solution found by the annealer. This reflection cannot be replicated if we look at the values relative to the mean and standard deviation, where the use of solutions from *descendant*-instances clearly implies better performance. In summary, the sharing of knowledge through RA leads to more reliable and robust behavior of the annealer in the case of s14.

The results are even more interesting in the case of s16. In this case, the results offered by the forward annealing process are considerably worse as a consequence of facing a larger and more complex problem. In this specific case, the ToK between problems stands as a mechanism that directly implies superior performance for the annealer. More specifically, the use of RA under the criteria established in this research leads, in the vast majority of cases, to a more reliable and robust performance and even to an increase in the probability of obtaining the optimal solution to the *target* problem.

In summary, in view of the results shown in Table 2, we can conclude, at least preliminarily, that knowledge sharing through RA is a promising strategy for solving optimization problems. However, all that glitters is not gold, as the good performance shown in these tests reveals a clear line of investigation. This is because, despite better results, there is no clear pattern to shed light on the

requirements that two problems must meet to ensure productive knowledge sharing.

These varying results reveal that **RQ1** requires inspecting further the reasons why some *descendant*-instances produce better results than others. This inspection requires a more detailed examination of the relationship between the characteristics of the *descendant*-instances and the performance shown in Table 2.

3.2 RQ2: can we infer the characteristics that an input solution should meet to help increase the probability of succeeding in an RA process?

Although there is a widespread belief in the literature that RA is more effective when the initial state is "*close*" to the ground state [24], there is no exact definition for the term "*close*" in this context. Some papers base this concept on the energetic proximity between the input solution and the ground state [23], while others suggest that using metrics such as the Hamming distance could be beneficial, contrasting with only choosing states that are low in energy, which

TABLE 2 Results of the reverse annealing experimentation. Optimal results obtained using the industry-oriented Hybrid Solver from Quantagonia: s14 = 12,422; s16 = 10,718.

Performance analysis on s14		Performance analysis on s16	
Source of knowledge	(best, avg., std.)	Source of knowledge	(best, avg., std.)
—	(−12,422, −12,418.9, 2.02)	—	(−10,716, −10,712.0, 3.00)
s14_0.2_L2L	(−12,421, **−12,420.2, 0.74**)	s16_0.2_L2L	(**−10,716, −10,714.7, 1.27**)
s14_0.2_L2H	(**−12,422, −12,421.6, 0.48**)	s16_0.2_L2H	(**−10,718, −10,716.2, 0.97**)
s14_0.2_H2L	(**−12,422, −12,421.0, 0.63**)	s16_0.2_H2L	(**−10,718, −10,716.6, 0.80**)
s14_0.2_H2H	(**−12,422, −12,420.8, 0.60**)	s16_0.2_H2H	(−10,714, **−10,711.5, 1.43**)
s14_0.4_L2L	(−12,421, **−12,420.6, 0.48**)	s16_0.4_L2L	(−10,717, **−10,715.7, 0.90**)
s14_0.4_L2H	(**−12422, −12,420.8, 0.60**)	s16_0.4_L2H	(−10,717, **−10,716.0, 1.00**)
s14_0.4_H2L	(−12,421, **−12,419.4, 1.11**)	s16_0.4_H2L	(−10,716, **−10,714.8, 0.87**)
s14_0.4_H2H	(**−12,422, −12,420.6, 0.80**)	s16_0.4_H2H	(−10,717, **−10,714.9, 1.22**)
s14_0.6_L2L	(−12,421, **−12,420.7, 0.45**)	s16_0.6_L2L	(−10,715, **−10,713.8, 0.74**)
s14_0.6_L2H	(−12,421, **−12,418.9, 1.04**)	s16_0.6_L2H	(−10,715, **−10,713.9, 0.94**)
s14_0.6_H2L	(**−12,422, −12,420.6, 0.66**)	s16_0.6_H2L	(−10,716, **−10,715.1, 0.94**)
s14_0.6_H2H	(−12,421, **−12,420.4, 0.66**)	s16_0.6_H2H	(**−10,718, −10,715.8, 0.87**)
s14_0.8_L2L	(−12,421, **−12,419.7, 1.10**)	s16_0.8_L2L	(**−10,718, -10715.0, 1.34**)
s14_0.8_L2H	(−12,420, **−12,419.1, 1.30**)	s16_0.8_L2H	(−10,714, −10,710.7, 1.55)
s14_0.8_H2L	(−12,421, **−12,419.5, 1.02**)	s16_0.8_H2L	(**−10,717, −10,716.5, 0.50**)
s14_0.8_H2H	(−12,421, **−12,420.2, 0.60**)	s16_0.8_H2H	(**−10,717, −10,715.6, 0.79**)

Bold results equalize or improve the performance of the forward annealing baseline (which is −12422 for s14 and −10716 for s16).

may be far from the correct ground state in terms of the quantum annealing procedure [25].

Thus, experimentally answering this second question can be done by sorting the results depicted in Table 2, using as references both energy differences and Hamming distances (represented in Table 1). This is the purpose of the plots shown in Figures 1, 2. To facilitate the visualization of the outcomes, we show the results obtained by solving each target-instance through forward annealing in the leftmost part of these images. It should also be noted that the figures focused on the energy difference are based on the results shown in column *Energy wrt. sXX* of Table 1. That is, on the energy value that each best-found solution returns using as a basis the objective function of its related target-instance. In addition, the total Hamming distance h has been used as a reference in Figures 1b, 2b.

On the one hand, by carefully analyzing the results of Figures 1a, 2a, it is difficult to detect a clear correlation between the energy difference among the solutions and their good performance as an input instance in an RA-based optimization process. Although the closest solutions offer competitive results, there are many cases that make it impossible to draw a generalizable conclusion.

In the case of s14, s14_0.8_L2H has emerged as the worst instance for feeding the RA mechanisms despite its energy value being close to the ground-state solution (−12,413 vs. −12,422, +9). In contrast, s14_0.6_H2L has shown adequate performance despite its remoteness in terms of energy (−12,297 vs. −12,422, +125). Turning our attention to s16, it is also possible to detect that closer solutions produce good results. Even so, there are contradictory cases such as s16_0.6_L2H, s16_0.8_H2L, or s16_0.8_L2L. Regarding s16_0.6_L2H, its performance has proven to be poor, even being near the reference solution (−10,709 vs. −10,716, +7). Finally, the other two *descendant*-instances mentioned have emerged as good alternatives to feed the RA despite being among the most distant in terms of energy.

On the other hand, a clear trend is detected by analyzing Figures 1b, 2b, where the Hamming distance emerges as a key factor in finding good solutions for the target-instances. In other words, using as a source of knowledge a *descendant*-instance whose best solution minimizes the Hamming distance regarding *parent*-instances has proven to be efficient. Even a quick glance at the results obtained leads to the conclusion that as the
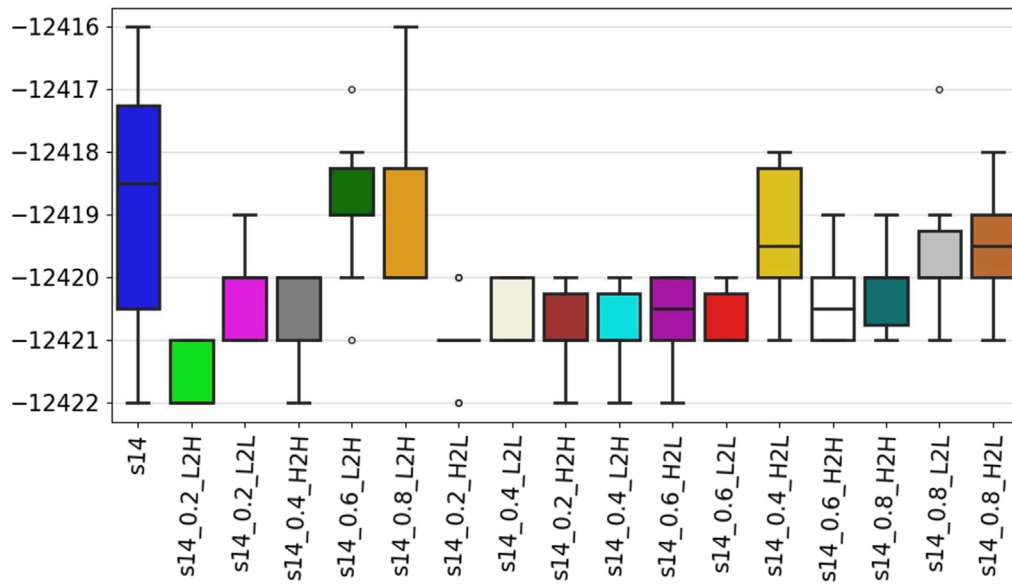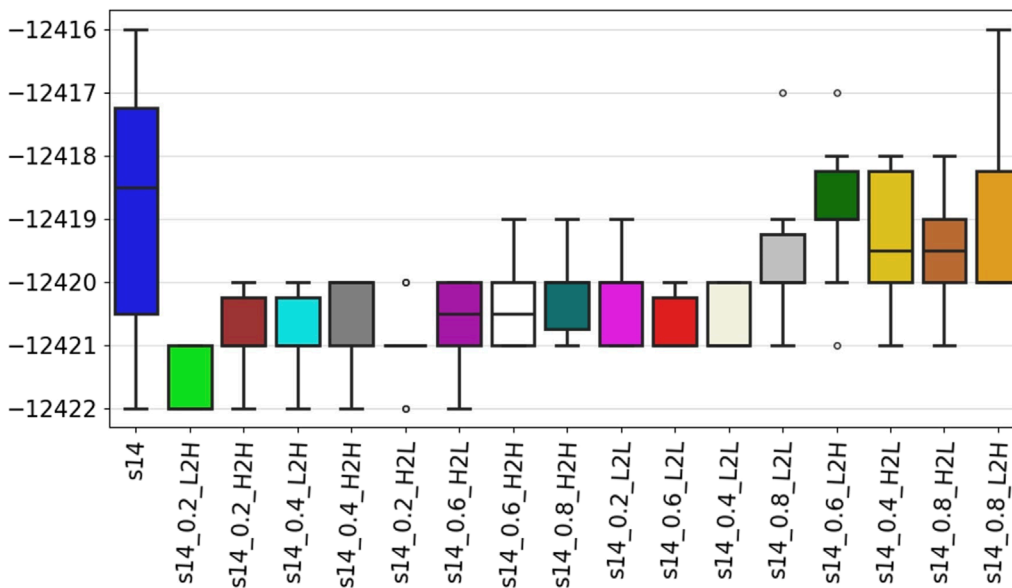
FIGURE 1
Results related to the resolution of s14 using as input the best solution found for each *descendant*-instance over 10 independent runs. The blue boxplots represent the baseline results obtained through forward annealing. The lower the energy value, the better the solution. **(a)** Results for s14 sorted by energy difference. **(b)** Results for s14 sorted by Hamming distance.

Hamming distance increases, the performance of the RA degrades significantly.

This finding, despite being almost unexplored in studies related to RA, is in line with the premises of the classical computing field known as transfer optimization [20]. There are several studies related to this classical research stream that conclude that for efficient knowledge sharing, there must be at least a *partial overlapping* in the optimal solutions to the problems at hand. More specifically, two problems present a *partial overlapping* if the optima of both tasks are exactly the same in the unified search space regarding a subset of variables and different from the remaining variables [37].

Furthermore, conclusions related to **RQ2** are also coherent with the work proposed in [10]. That paper is focused on a specific kind of quantum annealing called *biased annealing*, which involves complementing the driver Hamiltonian with a longitudinal bias
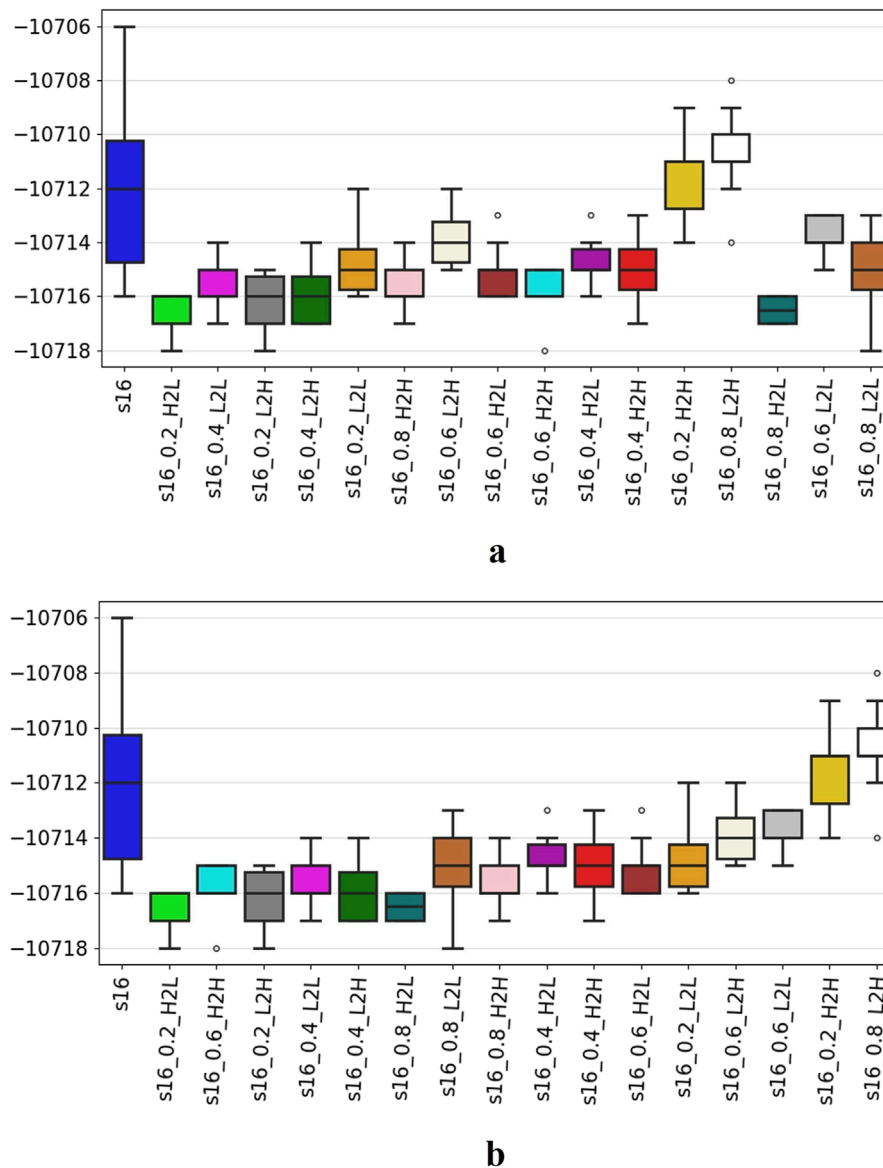
FIGURE 2
Results related to the resolution of s16 using as input the best solution found for each *descendant*-instance over 10 independent runs. The blue boxplots represent the baseline results obtained through forward annealing. The lower the energy value, the better the solution. **(a)** Results for s16 sorted by energy difference. **(b)** Results for s16 sorted by Hamming distance.

field. The importance of the Hamming distance between the driver Hamiltonian and the ground state is highlighted in a specific process called *iterative biased quantum annealing*, despite the same author minimizing its impact in a subsequent study when solving problems of large size [38].

Lastly, all the instances used in this paper and the complete set of results shown are openly available [39].

4 Conclusions and further work

In this work, we have preliminarily studied the influence of transferring knowledge between similar tasks through the reverse

annealing mechanism implemented by D-Wave. To do that, two research questions have been posed:

> **RQ1**: *Is RA a paradigm that can benefit from knowledge transfer between problem instances with similar characteristics?*
>
> **RQ2**: *Can we infer the characteristics that an input solution should meet to help increase the probability of succeeding in an RA process?*

First, answering **RQ1** and in view of the positive results obtained, the transfer of knowledge in quantum computing appears to be a promising research avenue that undoubtedly deserves further investigation. As for RA, it is commonly applied as a local refinement procedure right after an optimization process carried out by forward annealing or a classical optimization technique. In this work, RA

serves as a mechanism for knowledge transfer, allowing the reuse of a solution obtained in an independent optimization process.

It is important to highlight that, aside from the results regarding performance, knowledge sharing pursues savings in the use of computational resources. These savings are mainly materialized by resorting to previous outcomes when solving a new task under the assumption that industries are likely to face tasks that have much in common with each other. This is especially useful in real-world environments where companies accumulate results from past planning exercises and continually receive new tasks that are sometimes not very distinct from those already executed.

Second, answering **RQ2** and according to the results depicted, the closeness in terms of energy is not related to the performance of RA, and only the closeness with respect to the Hamming distance is. This means that the neighborhood must be based on solution codification and not energy, making it a priority to organize the coding of the target problem so that it fits as much as possible with the source problem through a unified search space [36].In light of these positive results, several open research questions should be further studied in subsequent investigations, which can be summarized as

- ■ Future work of considerable importance will likely focus on advancing toward the generalization of the conclusions drawn from this study. To reach this ambitious goal, the following steps are essential:

  - Inspired by the findings described in [38], further experiments should be carried out with larger instances, with the intention of determining whether the conclusions are replicable in such cases.
  - Further study the impact of the `RA Schedules` on knowledge transfer. To do this, the preliminary study already described in Section 3 will be revisited, and the knowledge transfer paradigm will be examined under the mentioned 25 `RA Schedules`.
  - Carry out a thorough analysis of computational resources and time.
  - Examine whether the findings from the experimentation can be replicated in other optimization problems, such as the traveling salesman problem or the bin packing problem.
  - Study the potential limitations when applying RA to more real-world scenarios.

- ■ Fairly compare RA as a local refinement method (using the same instance for source- and target-tasks) with RA as a transfer knowledge procedure.
- ■ Define a similarity measure for transfer of knowledge purposes. This means, given a potential source-task and its results, designing a metric that calculates how good a candidate source-task is for knowledge transfer purposes. The findings provided by the experimentation of this paper, particularly that focused on Hamming distance, are a good starting point for formulating this metric. Thus, a validated metric accepted

by the community would be a milestone for the transfer of knowledge research line.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found in the article/supplementary material.

Author contributions

Funding

Acknowledgments

Conflict of interest

Authors EO and EV-R were employed by TECNALIA.

Publisher's note

References

1. Rajak A, Suzuki S, Dutta A, Chakrabarti BK. Quantum annealing: an overview. *Philosophical Trans R Soc A* (2023) 381:20210417. doi:10.1098/rsta.2021.0417

2. Bertsimas D, Tsitsiklis J. Simulated annealing. *Stat Sci* (1993) 8:10–5. doi:10.1214/ss/1177011077

3. Ebadi S, Keesling A, Cain M, Wang TT, Levine H, Bluvstein D, et al. Quantum optimization of maximum independent set using rydberg atom arrays. *Science* (2022) 376:1209–15. doi:10.1126/science.abo6587

4. Bunyk PI, Hoskinson EM, Johnson MW, Tolkacheva E, Altomare F, Berkley AJ, et al. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Trans Appl Superconductivity* (2014) 24:1–10. doi:10.1109/tasc.2014.2318294

5. Preskill J. Quantum computing in the nisq era and beyond. *Quantum* (2018) 2:79. doi:10.22331/q-2018-08-06-79

6. Osaba E, Villar-Rodriguez E, Gomez-Tejedor A, Oregi I. Hybrid quantum solvers in production: how to succeed in the nisq era? *arXiv preprint arXiv:2401.10302* (2024) 423–34. doi:10.1007/978-3-031-77738-7_35

7. Bhalgamiya B, Perera D, Novotny MA. Quantum annealing error mitigation using mirror symmetries on different generations of quantum annealers. In: *2022 IEEE international conference on quantum computing and engineering (QCE)*. IEEE (2022). p. 844–6.

8. Céleri LC, Huerga D, Albarrán-Arriagada F, Solano E, de Andoin MG, Sanz M. Digital-analog quantum simulation of fermionic models. *Phys Rev Appl* (2023) 19:064086. doi:10.1103/physrevapplied.19.064086

9. Perdomo-Ortiz A, Venegas-Andraca SE, Aspuru-Guzik A. A study of heuristic guesses for adiabatic quantum computation. *Quan Inf Process* (2011) 10:33–52. doi:10.1007/s11128-010-0168-z

10. Graß T. Quantum annealing with longitudinal bias fields. *Phys Rev Lett* (2019) 123:120501. doi:10.1103/physrevlett.123.120501

11. Ohkuwa M, Nishimori H, Lidar DA. Reverse annealing for the fully connected p-spin model. *Phys Rev A* (2018) 98:022314. doi:10.1103/physreva.98.022314

12. Tate R, Farhadi M, Herold C, Mohler G, Gupta S. Bridging classical and quantum with sdp initialized warm-starts for qaoa. *ACM Trans Quan Comput* (2023) 4:1–39. doi:10.1145/3549554

13. Beaulieu D, Pham A. Max-cut clustering utilizing warm-start qaoa and ibm runtime. *arXiv preprint arXiv:2108.13464* (2021).

14. Tate R, Moondra J, Gard B, Mohler G, Gupta S. Warm-started QAOA with custom mixers provably converges and computationally beats goemans-williamson's max-cut at low circuit depths. *Quantum* (2023) 7:1121. doi:10.22331/q-2023-09-26-1121

15. Developers D-W. Reverse quantum annealing for local refinement of solutions. *Tech Rep* (2021) 14-1018A.

16. Haba R, Ohzeki M, Tanaka K. Travel time optimization on multi-agv routing by reverse annealing. *Scientific Rep* (2022) 12:17753. doi:10.1038/s41598-022-22704-0

17. Carugno C, Ferrari Dacrema M, Cremonesi P. Evaluating the job shop scheduling problem on a d-wave quantum annealer. *Scientific Rep* (2022) 12:6539. doi:10.1038/s41598-022-10169-0

18. Grant E, Humble TS, Stump B. Benchmarking quantum annealing controls with portfolio optimization. *Phys Rev Appl* (2021) 15:014012. doi:10.1103/physrevapplied.15.014012

19. Golden J, O'Malley D. Reverse annealing for nonnegative/binary matrix factorization. *Plos one* (2021) 16:e0244026. doi:10.1371/journal.pone.0244026

20. Gupta A, Ong YS, Feng L. Insights on transfer optimization: because experience is the best teacher. *IEEE Trans Emerging Top Comput Intelligence* (2017) 2:51–64. doi:10.1109/tetci.2017.2769104

21. Pelofske E, Hahn G, Djidjev H. Initial state encoding via reverse quantum annealing and h-gain features. *IEEE Trans Quan Eng* (2023) 4:3102221–1. doi:10.1109/tqe.2023.3319586

22. Venturelli D, Kondratyev A. Reverse quantum annealing approach to portfolio optimization problems. *Quan Machine Intelligence* (2019) 1:17–30. doi:10.1007/s42484-019-00001-w

23. Ding J, Spallitta G, Sebastiani R. Effective prime factorization via quantum annealing by modular locally-structured embedding. *Scientific Rep* (2024) 14:3518. doi:10.1038/s41598-024-53708-7

24. Passarelli G, Yip KW, Lidar DA, Nishimori H, Lucignano P. Reverse quantum annealing of the p-spin model with relaxation. *Phys Rev A* (2020) 101:022331. doi:10.1103/physreva.101.022331

25. Pelofske E. Mapping state transition susceptibility in quantum annealing. *Phys Rev Res* (2023) 5:013224. doi:10.1103/physrevresearch.5.013224

26. Villar-Rodriguez E, Osaba E, Oregi I, Romero SV, Ferreiro-Vélez J. On the transfer of knowledge in quantum algorithms *arXiv preprint arXiv:2501.14120*.

27. Salkin HM, De Kluyver CA. The knapsack problem: a survey. *Naval Res Logistics Q* (1975) 22:127–44. doi:10.1002/nav.3800220110

28. Bożejko W, Burduk A, Pempera J, Uchroński M, Wodecki M. Optimal solving of a binary knapsack problem on a d-wave quantum machine and its implementation in production systems. *Ann Operations Res* (2024) 1–16. doi:10.1007/s10479-024-06025-1

29. Wilkening S, Lefterovici AI, Binkowski L, Perk M, Fekete S, Osborne TJ. A quantum algorithm for the solution of the 0-1 knapsack problem. *arXiv preprint arXiv:2310.06623* (2023).

30. Awasthi A, Bär F, Doetsch J, Ehm H, Erdmann M, Hess M, et al. Quantum computing techniques for multi-knapsack problems. In: *Science and information conference*. Springer (2023). p. 264–84.

31. Glover F, Kochenberger G, Du Y. A tutorial on formulating and using qubo models. *arXiv preprint arXiv:1811* (2018) 11538.

32. Pusey-Nazzaro L. Adiabatic quantum optimization fails to solve the knapsack problem. *arXiv preprint arXiv:2008.07456* (2020).

33. Van Dam W, Eldefrawy K, Genise N, Parham N. Quantum optimization heuristics with an application to knapsack problems. In: *2021 IEEE international Conference on quantum Computing and engineering (QCE) (IEEE)* (2021). p. 160–70.

34. Mizel A, Lidar DA, Mitchell M. Simple proof of equivalence between adiabatic quantum computation and the circuit model. *Phys Rev Lett* (2007) 99:070502. doi:10.1103/physrevlett.99.070502

35. Yang Z, Zolanvari M, Jain R. A survey of important issues in quantum computing and communications. *IEEE Commun Surv and Tutorials* (2023) 25:1059–94. doi:10.1109/comst.2023.3254481

36. Gupta A, Ong YS, Feng L. Multifactorial evolution: toward evolutionary multitasking. *IEEE Trans Evol Comput* (2015) 20:343–57. doi:10.1109/tevc.2015.2458037

37. Da B, Ong YS, Feng L, Qin AK, Gupta A, Zhu Z, et al. Evolutionary multitasking for single-objective continuous optimization: Benchmark problems, performance metric, and baseline results. *arXiv preprint arXiv:1706.03470* (2017).

38. Grass T. Quantum annealing sampling with a bias field. *Phys Rev Appl* (2022) 18:044036. doi:10.1103/physrevapplied.18.044036

39. Osaba E, Villar E. Benchmark dataset and results for transfer of knowledge through reverse annealing. *experimentation* (2024). doi:10.17632/yr8dg923wg.1