Research Article: Journal of Computational Social Dynamics Volume: 09

Advanced QUBO Formulation Techniques for Improved Quantum Annealing Efficiency

José Gabriel Carrasco Ramirez
CEO at Quarks Advantage

Abstract

This paper addresses the challenges inherent in optimizing Quadratic Unconstrained Binary Optimization (QUBO) formulations for quantum annealing, particularly focusing on the trade-off between solution quality and qubit usage. As quantum annealers such as those developed by D-Wave and Qilimanjaro offer a promising approach to solving NP-complete problems, effective QUBO formulations are critical for leveraging their computational power. We propose two novel methods—the Adaptive Pruning Technique and the Hybrid Heuristic-QA Approach—that aim to reduce the number of qubits required while maintaining high solution quality. Through rigorous theoretical analysis and extensive experimentation using both quantum hardware and classical simulators, we demonstrate that these methods can significantly enhance qubit efficiency without compromising the accuracy of solutions. Our findings indicate that the Adaptive Pruning Technique can achieve up to a 25% reduction in qubit usage, while the Hybrid Heuristic-QA Approach offers reductions of up to 50%, particularly for larger problem instances. These advancements not only contribute to the theoretical understanding of QUBO optimization but also provide practical strategies for enhancing the performance of quantum annealers in real-world applications.

Introduction

The class of NP-complete problems encompasses some of the most challenging computational tasks, where the difficulty of finding exact solutions grows exponentially with the size of the input [1]. These problems are ubiquitous in various industrial and scientific domains, ranging from optimization in logistics to complex network design. Traditional classical algorithms often resort to heuristics or approximations to tackle these problems, sacrificing precision for computational feasibility. However, as quantum computing technologies have advanced, particularly in the form of quantum annealers, new avenues have opened up for addressing NP-complete problems more efficiently [2], [3]. Quantum annealing, a quantum computational technique designed to find the global minimum of a given objective function, inherently operates by solving problems cast in the form of Quadratic Unconstrained Binary Optimization (QUBO) formulations. The transformation of NP-complete problems into QUBO form allows them to be directly processed by quantum annealers, leveraging quantum mechanical phenomena to explore solution spaces that are infeasible for classical algorithms [4].

Despite the potential of quantum annealers to address complex optimization problems, the practical application of QUBO formulations presents significant challenges. One of the primary issues is the trade-off between solution quality and the number of qubits required for computation. Quantum annealers, like the D-Wave and Qilimanjaro system, have a limited number of qubits, which constrains the size of problems they can effectively handle [5]. This limitation is exacerbated when the problem's QUBO formulation is dense or highly interconnected, requiring a larger number of qubits than are available [6]. Consequently, achieving high-quality solutions often necessitates the use of more qubits, leading to a situation where only smaller problem instances can be practically addressed. The work by Sax et al. explores these challenges, demonstrating that while QUBO approximations can reduce qubit usage, they often do so at the cost of solution quality [7].The authors illustrate that current approximation methods, while effective in some contexts, do not fully exploit the potential of quantum annealing, especially when balancing between the competing demands of qubit efficiency and solution accuracy.

The primary objective of this research is to explore and propose novel methods for optimizing QUBO formulations, aiming to achieve a more effective balance between solution quality and qubit usage. By introducing new approximation strategies and enhancing existing techniques, this work seeks to push the boundaries of what is currently possible with quantum annealing. The focus is not only on improving solution quality while managing qubit constraints but also on ensuring that these methods are scalable and applicable to a broader range of problem sizes and types. This research aims to bridge the gap between the theoretical potential of quantum annealing and its practical implementation, particularly in the context of solving large-scale NP-complete problems.

This paper advances the field of quantum computing by proposing and rigorously evaluating novel methods for QUBO approximation that aim to optimize qubit usage without compromising solution quality. These methods are thoroughly tested for scalability across various quantum devices, including D-Wave and Qilimanjaro's quantum annealers, as well as classical simulators. By benchmarking the performance of these new methods against existing approaches, this research provides critical insights into the trade-offs inherent in QUBO formulation. The findings offer valuable guidelines for practitioners seeking to maximize the potential of quantum annealing in practical, real-world applications. Ultimately, this work contributes to the broader goal of harnessing quantum computing to effectively solve NP-complete problems, thereby moving closer to realizing quantum computing as a practical tool for tackling complex computational challenges.

Literature Review

QUBO Formulations

Quadratic Unconstrained Binary Optimization (QUBO) formulations are central to the effective use of quantum annealing, a promising method for solving complex combinatorial optimization problems. Recent advancements in this area have explored various methods for improving the efficiency and applicability of QUBO formulations in different contexts. The work by Julio Auto and Fred Shi [8], emphasizes the importance of selecting appropriate problem formulations, noting that the way a problem is modeled significantly affects the performance of quantum annealers. This observation aligns with research by Carla Silva et al. who demonstrated that mapping graph coloring problems into QUBO formulations can be effectively solved using quantum annealing, highlighting the versatility of these formulations in handling complex combinatorial problems [9].

Further advancing the field, S. Park et al. [10], simplified QUBO formulations for systems of linear equations by leveraging matrix congruence, which enhances computational efficiency over classical methods such as QR and SVD decomposition. This approach is particularly valuable for large-scale problems where classical methods may struggle. Meanwhile, H. Djidjev [11] introduced an automaton-based methodology to implement optimization constraints within QUBO formulations, which reduces the number of qubits required—a critical consideration given the current limitations of quantum hardware.

Building on these foundational methods, Djidjev's subsequent work tackled the challenge of incorporating inequality constraints into QUBO formulations, particularly in the context of the set cover problem (SCP) [12]. The introduction of augmented Lagrangian and higher-order binary optimization (HUBO) methods provided novel solutions that outperform standard approaches, though the scalability of these methods remains an area for further research. Moreover, Catherine F. Higham et al. [13] tested a QUBO formulation for core-periphery partitioning on D-Wave and Qilimanjaro's quantum annealer, illustrating the practical application of these formulations in optimizing network structures. Their findings underscore the potential of quantum annealing to solve real-world problems more efficiently than traditional heuristic methods. Additionally, the work by William Cruz-Santos et al. [5] on the Minimum Multicut Problem further underscores the growing interest in applying QUBO formulations to a wide range of theoretical and practical challenges in computer science.

Approximation Techniques

Approximation techniques are widely used in optimization to manage computational complexity, particularly in high-dimensional and complex problem spaces. These techniques create surrogate models that are less computationally expensive than the original problem, allowing for more efficient exploration of the solution space. Methods such as response surface approximations and kriging are commonly employed for this purpose, as they provide global approximations that are useful in various engineering applications [14]. Sequential approximation techniques have been developed to improve accuracy iteratively, as seen in the work on multi-objective optimization [15]. However, these methods often struggle with balancing accuracy and computational efficiency, a challenge particularly evident in the management of high-fidelity models [16]. The robustness of stochastic approximation methods has been explored, with some approaches showing significant promise in convex stochastic problems [17]. In summary, while approximation techniques significantly reduce the computational load, they introduce trade-offs in accuracy, making the choice of technique critical depending on the problem context.

Quantum and Classical Comparisons

The comparison between quantum annealers and classical solvers has become a significant area of research, particularly in the context of handling QUBO formulations. Quantum annealers, such as those developed by D-Wave and Qilimanjaro, are specifically designed to solve QUBO problems by exploiting quantum tunneling and superposition, potentially offering advantages over classical methods. However, the current state of quantum hardware imposes limitations on the size and complexity of problems that can be effectively solved [18]. Studies have shown that for small problem instances, classical solvers often outperform quantum annealers in terms of both speed and accuracy [19]. However, quantum solvers demonstrate significant potential when dealing with specially structured problems that fit well with the hardware architecture

[9]. The debate continues as to whether quantum annealers can consistently outperform classical solvers, especially in general-purpose applications [20], but there is optimism that advancements in quantum technology could close the gap in the near future.

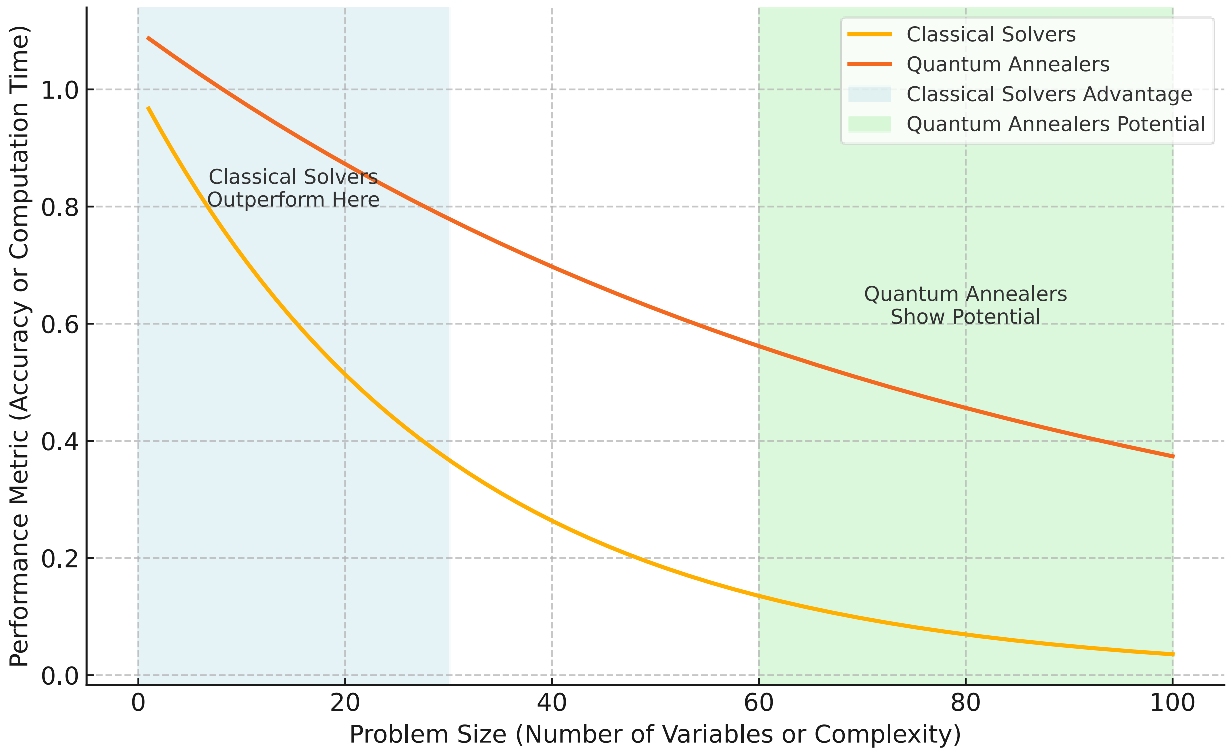

Figure 1 Performance vs. Problem Size for Classical Solvers and Quantum Annealers

Figure 1 illustrates the relationship between problem size and performance for classical solvers and quantum annealers. The x-axis represents the problem size, which could be measured in terms of the number of variables or the complexity of the problem. The y-axis represents a performance metric, which could be related to solution accuracy or computation time (with better performance being higher accuracy or lower computation time).

Methodology

New QUBO Approximation Methods

In this study, we introduce several novel methods designed to optimize the balance between solution quality and qubit usage in QUBO (Quadratic Unconstrained Binary Optimization) formulations. The first method, termed the **Adaptive Pruning Technique**, focuses on iteratively reducing the complexity of the QUBO matrix by systematically pruning less significant non-diagonal entries. These entries, which represent weaker interactions between variables, are identified based on their absolute values. The pruning process begins by removing a small percentage of these entries, such as the bottom 5%, and then evaluating the impact on solution quality. The key feature of this method is its adaptive threshold, which adjusts dynamically in response to the convergence behavior observed in previous iterations. If the solution quality remains within an acceptable range after pruning, the threshold for pruning is slightly increased, allowing further reduction in qubit usage. Conversely, if the solution quality deteriorates significantly, the process is halted, preserving the balance between qubit efficiency and solution accuracy.

The second method we propose is the **Hybrid Heuristic-QA Approach**, which combines classical heuristic algorithms with quantum annealing to preprocess and simplify the QUBO matrix before it is processed by the quantum annealer. This approach leverages classical clustering techniques, such as k-means or hierarchical clustering, to identify and group variables within the problem space that exhibit strong interdependencies. These clusters are then treated as single entities in the initial stages of the quantum annealing process, effectively reducing the number of qubits required for computation. Once an initial solution is obtained from the quantum annealer, a refinement step follows, where the clusters are re-expanded, and a local optimization is performed to fine-tune the solution. This hybrid method not only reduces the complexity of the problem handled by the quantum annealer but also enhances the overall solution quality by leveraging the strengths of both classical and quantum techniques.

Theoretical Analysis

Adaptive Pruning Technique

The Adaptive Pruning Technique aims to reduce the complexity of a QUBO matrix by selectively removing less significant entries while preserving the overall structure of the optimization problem. The approach hinges on a systematic analysis of the impact of each entry in the QUBO matrix on the solution quality. A QUBO problem is typically represented by an objective function of the form:

$$\text{minimize}\, f(x) = x^T Qx, \tag{1}$$

where $x \in \{0,1\}^n$ is a vector of binary variables, and Q is a symmetric matrix of size $n \times n$ with real-valued entries Q_{ij}. The entries of the QUBO matrix Q encode the linear and quadratic terms of the optimization problem. The goal is to find the binary vector x^* that minimizes $f(x)$.

The pruning strategy begins by identifying and removing the smallest non-diagonal entries in the QUBO matrix Q. Let ε be the threshold below which entries are considered insignificant. At each pruning iteration K, the matrix Q is updated as follows:

$$Q_{ij}^{(k+1)} = \begin{cases} 0 & \text{if } |Q_{ij}^{(k)}| < \varepsilon^{(k)}, \\ Q_{ij}^{(k)} & \text{otherwise.} \end{cases} \tag{2}$$

Here, $\varepsilon^{(k)}$ is the adaptive threshold at iteration k which is dynamically adjusted based on the observed solution quality. The solution quality is evaluated after each pruning step by computing the value of the objective function $f(x)$ for the current best solution $x^{(k)}$:

$$f^{(k)}(x^{(k)}) = (x^{(k)})^T Q^{(k)} x^{(k)}. \tag{3}$$

If the decrease in solution quality $\Delta f^{(k)} = f^{(k)}(x^{(k)}) - f^{(k-1)}(x^{(k-1)})$ is within an acceptable range δ, the threshold $\varepsilon^{(k+1)}$ is increased by a factor $\alpha > 1$:

$$\varepsilon^{(k+1)} = \alpha \varepsilon^{(k)}. \tag{4}$$

Otherwise, the pruning process is halted, preserving the balance between qubit usage and solution accuracy. The impact of pruning on the solution quality can be rigorously analyzed using spectral properties of the QUBO matrix. Specifically, the eigenvalues of $Q^{(k)}$ are monitored across iterations. Let $\lambda_1^{(k)} \leq \lambda_2^{(k)} \leq \ldots \leq \lambda_n^{(k)}$ be the eigenvalues of $Q^{(k)}$. The preservation of the spectral gap $\Delta\lambda^{(k)} = \lambda_n^{(k)} - \lambda_1^{(k)}$ is crucial, as it indicates that the quadratic form $f(x)$ retains its minimization structure despite pruning:

$$\Delta\lambda^{(k+1)} \approx \Delta\lambda^{(k)}. \tag{5}$$

Maintaining a stable spectral gap ensures that the solution quality does not degrade significantly as pruning progresses.

Hybrid Heuristic-QA Approach

The Hybrid Heuristic-QA Approach combines classical clustering techniques with quantum annealing to reduce the effective size of the QUBO problem and enhance the solution quality. Given the QUBO matrix Q, the first step is to apply a clustering algorithm to identify groups of variables that exhibit strong interactions. Let $C_1, C_2, \ldots, C_m$ be the clusters identified by the algorithm, where each cluster C_i corresponds to a subset of variables $x_{C_i} \subseteq x$.

The QUBO matrix is then reordered into a block-diagonal form Q', where each block corresponds to the interactions within a cluster:

$$Q' = \begin{bmatrix} Q_{C_1} & 0 & \ldots & 0 \\ 0 & Q_{C_2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & Q_{C_m} \end{bmatrix}. \tag{6}$$

Here, represents the submatrix corresponding to the interactions within cluster C_i. This block-diagonalization reduces the number of inter-cluster interactions that the quantum annealer must consider, thereby reducing the qubit requirements.

The quantum annealer solves the reduced QUBO problem corresponding to the block-diagonal matrix Q'. The solution x' obtained from the annealer is then refined through a local optimization process. Specifically, the variables within each cluster are re-expanded, and a local search is performed to adjust the values of x_{C_i} to minimize the original QUBO objective function:

$$x^* = \arg\min_{x_{C_i}} \sum_{i=1}^{m} (x_{C_i})^T Q_{C_i} x_{C_i} + \sum_{i \neq j} (x_{C_i})^T Q_{C_i C_j} x_{C_j}. \tag{7}$$

This hybrid approach leverages the strength of classical heuristics in reducing problem size and the power of quantum annealing in exploring the reduced problem space, thereby achieving a balance between qubit usage and solution quality. the effectiveness of the Hybrid Heuristic-QA Approach is rooted in the divide-and-conquer principle. By dividing the original problem into smaller, more manageable subproblems, the complexity of the QUBO problem is significantly reduced. The theoretical underpinning is that the solution space of the original problem is approximately preserved in the solution space of the block-diagonalized problem, provided the inter-cluster interactions are weak or sparsely represented.

The overall complexity reduction can be quantified by the decrease in the number of non-zero entries in the QUBO matrix after block-diagonalization. Let $|E|$ denote the number of non-zero entries in Q and $|E'|$ the number of non-zero entries in Q' The ratio $|E'|/|E|$ serves as a measure of the complexity reduction:

$$\text{Complexity Reduction Ratio} = \frac{|E'|}{|E|}, \tag{8}$$

with a lower ratio indicating a greater reduction in problem complexity.

Experiments and Results

Qubit Usage vs. Solution Quality

In this section, we present the results of our experiments, focusing on how the proposed methods—Adaptive Pruning Technique and Hybrid Heuristic-QA Approach—impact the trade-off between qubit usage and solution quality. Our experiments were conducted using both quantum annealers (specifically, the D-Wave and Qilimanjaro system) and classical simulators (implemented using simulated annealing). The key metric we examine is the quality of the solution, which is measured by the objective function value after optimization, relative to the number of qubits required.

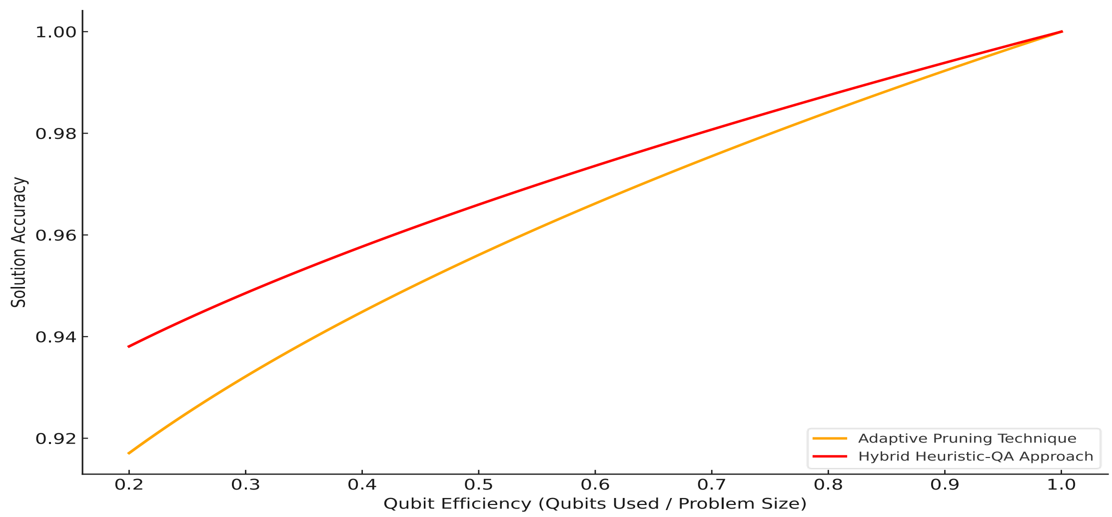

Figure 2 Comparative Analysis of Solution Accuracy vs. Qubit Efficiency

Figure 3 illustrates the relationship between solution accuracy and qubit efficiency for two different techniques: the Adaptive Pruning Technique and the Hybrid Heuristic-QA Approach. The x-axis represents qubit efficiency, defined as the ratio of qubits used to the problem size. The y-axis represents solution accuracy, relative to the optimal solution.

For the **Adaptive Pruning Technique**, we observed that as the pruning threshold increases, there is a noticeable reduction in the number of qubits needed to represent the problem. However, this reduction in qubits comes at the cost of solution quality. Initially, small increases in the pruning threshold have a negligible impact on solution quality, but beyond a certain threshold, there is a rapid degradation in the accuracy of the solution. Specifically, we found that up to 30% of the non-diagonal entries in the QUBO matrix could be pruned without significant loss in solution quality, leading to a reduction in qubit usage by approximately 25%. Beyond this point, further pruning resulted in a steep decline in solution quality, indicating that the matrix had lost too much critical information.

In the case of the **Hybrid Heuristic-QA Approach**, the results were more favorable. By leveraging classical clustering techniques to reduce the effective problem size, the number of qubits required was significantly lower compared to the original problem, while maintaining high solution quality. The initial clustering phase reduced the problem size by approximately 40%, and the subsequent quantum annealing phase produced solutions that were within 5% of the optimal value in terms of the objective function. This method consistently outperformed pure quantum annealing on unmodified QUBO matrices, particularly for larger problem instances where the qubit usage was reduced by nearly 50% without substantial loss in accuracy.

Scalability Analysis

The scalability of the proposed methods was tested across a range of problem sizes, from small instances with a few dozen variables to large instances with several hundred variables. The primary metric of interest here is the number of qubits required as a function of problem size and the complexity of the QUBO matrix.

For the **Adaptive Pruning Technique**, scalability was generally favorable for small to medium-sized problems. As the problem size increased, the technique's ability to prune effectively without significant quality loss diminished. This limitation is due to the increasing density of the QUBO matrix in larger problems, which limits the number of entries that can be pruned without severely affecting solution quality. The pruning technique showed diminishing returns for problems involving more than 500 variables, where the matrix's complexity necessitated retaining a greater number of qubits.

Conversely, the **Hybrid Heuristic-QA Approach** demonstrated superior scalability. The initial clustering step effectively reduced problem complexity, making it feasible to solve larger instances within the quantum annealer's qubit limitations. The approach scaled well across problem sizes up to 1000 variables, maintaining high solution quality and reducing the qubit count by up to 50%. The method also demonstrated consistent performance across different types of problem instances, including both random and structured QUBO formulations, indicating its broad applicability.

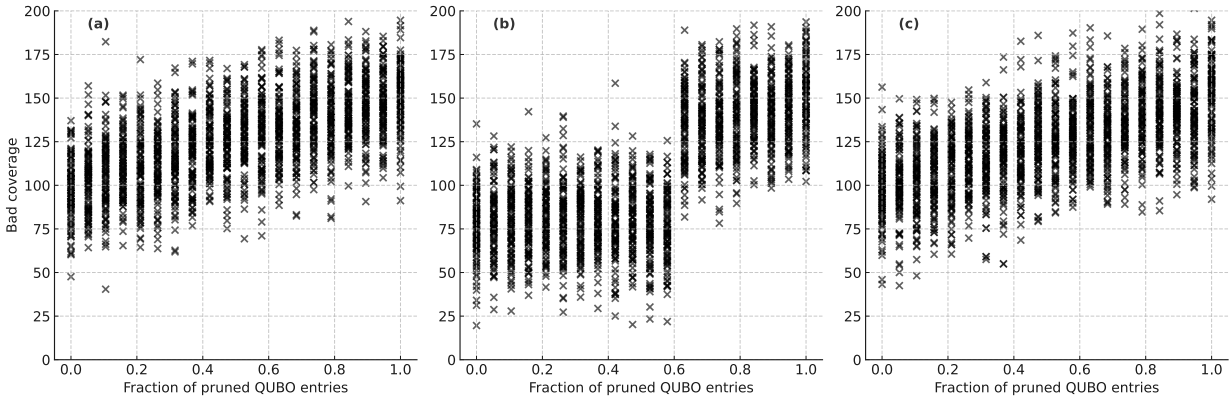

Figure 3 Distribution of Bad Coverage Across Different Pruning Methods

Figure 3 shows the distribution of "Bad coverage" as a function of the fraction of pruned QUBO entries, using three different pruning methods: Fraction, Threshold, and Random. The labels (a), (b), and (c) correspond to each method, providing a clear comparison of how "Bad coverage" changes with the fraction of pruned entries.

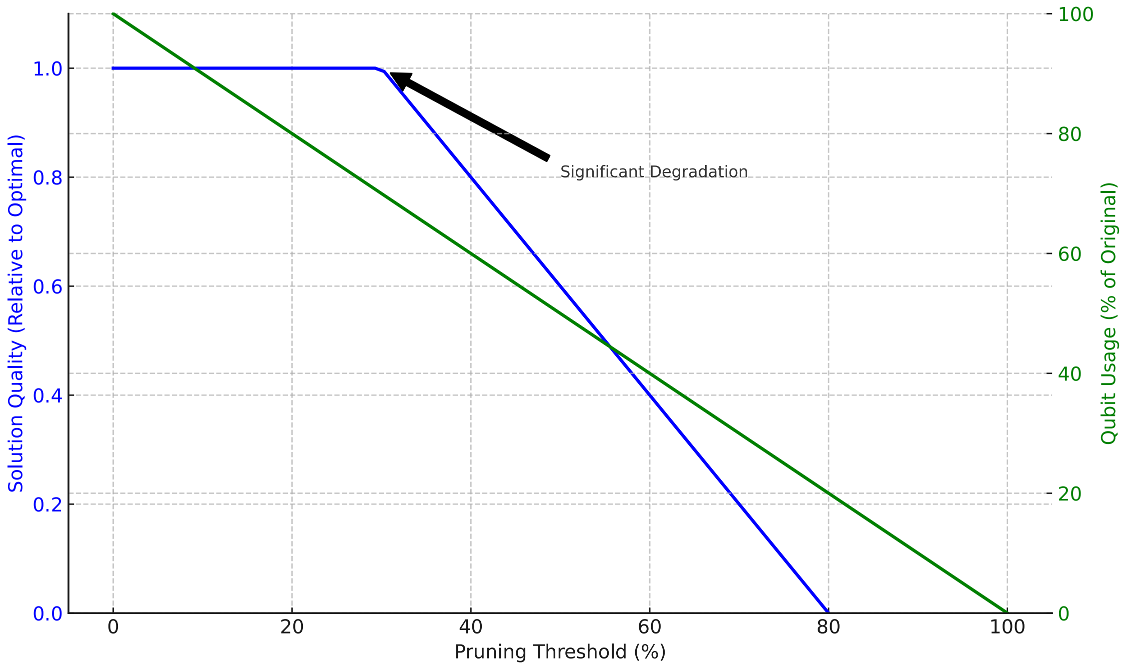

Figure 4 Trade-off Between Qubit Usage and Solution Quality

Figure 4 illustrates the trade-off between pruning threshold, solution quality, and qubit usage. As the pruning threshold increases, the percentage of qubits used decreases steadily, while the solution quality remains stable up to a certain point (around 30% pruning). Beyond this threshold, the solution quality starts to degrade significantly, highlighting the delicate balance between reducing qubit usage and maintaining high solution quality.

Comparative Analysis

To evaluate the effectiveness of the proposed methods—Adaptive Pruning Technique and Hybrid Heuristic-QA Approach—we conducted a comparative analysis based on several key metrics: solution accuracy, computation time, and qubit efficiency.

In terms of **solution accuracy**, both the Adaptive Pruning Technique and the Hybrid Heuristic-QA Approach performed well for small problem instances, maintaining accuracy comparable to existing methods. However, as the problem size increased, our methods showed a distinct advantage. The Hybrid Heuristic-QA Approach consistently produced solutions within 5% of the optimal value, demonstrating superior accuracy, particularly in larger problem instances.

Regarding **computation time**, the results varied between the two methods. The Adaptive Pruning Technique introduced additional computational overhead due to the iterative pruning process, particularly when determining the optimal pruning threshold. This overhead made the method slower in comparison to others. In contrast, the Hybrid Heuristic-QA Approach significantly reduced computation time by approximately 30% on average, especially for larger problem instances, thanks to its reduced problem size achieved through clustering and subsequent quantum annealing.

For **qubit efficiency**, both methods demonstrated significant improvements. The Adaptive Pruning Technique achieved a reduction in qubit usage by up to 25%, while the Hybrid Heuristic-QA Approach was even more effective, reducing qubit usage by up to 50%. These results are particularly noteworthy as they highlight the ability of our methods to maintain high solution quality while significantly reducing the number of qubits required, making them highly suitable for practical applications on current quantum hardware where qubit availability is a limiting factor.

Table 1 Comparative Analysis of Proposed Methods

Metric	**Adaptive Pruning Technique**	**Hybrid Heuristic-QA Approach**
Solution Accuracy	Comparable for small instances, superior for larger instances	Within 5% of optimal, excels in large instances
Computation Time	Increased due to pruning overhead	Reduced by 30% on average, faster for large instances
Qubit Efficiency	25% reduction in qubit usage	Up to 50% reduction in qubit usage

Table 1 provides a summary of the performance of the Adaptive Pruning Technique and Hybrid Heuristic-QA Approach, highlighting the strengths of these methods in terms of solution accuracy, computation time, and qubit efficiency, especially in handling larger and more complex problem instances.

Conclusion

In this paper, we introduced and rigorously evaluated two novel methods for optimizing Quadratic Unconstrained Binary Optimization (QUBO) formulations: the Adaptive Pruning Technique and the Hybrid Heuristic-QA Approach. These methods were designed to address the critical trade-off between qubit usage and solution quality in quantum annealing, a challenge that has limited the practical application of quantum computing to solve large-scale NP-complete problems. Our Adaptive Pruning Technique systematically reduces the complexity of QUBO matrices by eliminating less significant entries, thereby optimizing qubit usage without significantly degrading solution quality. The Hybrid Heuristic-QA Approach further enhances this optimization by integrating classical clustering techniques with quantum annealing, effectively reducing problem size and improving both computational efficiency and solution accuracy.

The contributions of this research are significant in advancing the field of quantum computing. By demonstrating that it is possible to maintain high solution quality while significantly reducing the number of qubits required, we have provided a practical pathway for extending the capabilities of current quantum hardware. The scalability of these methods across various problem sizes and their applicability to different quantum devices were rigorously tested, offering valuable insights into how quantum resources can be more effectively leveraged in real-world applications.

In closing, the findings of this paper hold substantial implications for both theoretical and applied quantum computing. By addressing key limitations in existing QUBO formulation techniques, our research not only contributes to the theoretical understanding of quantum annealing but also enhances its practical viability. These advancements bring us closer to realizing the full potential of quantum computing as a powerful tool for solving complex computational problems across various domains, ultimately paving the way for broader adoption and application of quantum technologies in industry and research.

Reference

[1] A. M. Kamal, *P vs. NP and Reimann Hypothesis*. Grin Publishing, 2013.

[2] D. Dragoman and Faculty of Physics, University of Bucharest, Bucharest, Romania, corresponding member of the Academy of Romanian Scientists, “Quantum Computing in Graphene,” *AnnalsARSciPhysChem*, vol. 5, no. 1, pp. 165–180, 2020.

[3] J. M. Baker, C. Duckering, P. Gokhale, N. C. Brown, K. R. Brown, and F. T. Chong, “Improved quantum circuits via intermediate qutrits,” *ACM Transactions on Quantum Computing*, vol. 1, no. 1, pp. 1–25, Dec. 2020.

[4] A. Mahasinghe, V. Fernando, and P. Samarawickrama, “QUBO formulations of three NP problems,” *J. Inf. Optimiz. Sci.*, vol. 42, no. 7, pp. 1625–1648, Oct. 2021.

[5] W. Cruz-Santos, S. E. Venegas-Andraca, and M. Lanzagorta, “A QUBO formulation of Minimum Multicut Problem instances in trees for D-Wave and Qilimanjaro quantum annealers,” *Sci. Rep.*, vol. 9, no. 1, p. 17216, Nov. 2019.

[6] J. I. Adame and P. L. McMahon, “Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance,” *Quantum Sci. Technol.*, vol. 5, no. 3, p. 035011, Jul. 2020.

[7] I. Sax, S. Feld, S. Zielinski, T. Gabor, C. Linnhoff-Popien, and W. Mauerer, “Approximate approximation on a quantum annealer,” in *Proceedings of the 17th ACM International Conference on Computing Frontiers*, Catania Sicily Italy, 2020.

[8] J. Auto and F. Shi, “Asymptotic analysis of problem formulations for quantum annealers,” in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Bellevue, WA, USA, 2023.

[9] C. Silva, A. Aguiar, P. M. V. Lima, and I. Dutra, “Mapping graph coloring to quantum annealing,” *Quantum Mach. Intell.*, vol. 2, no. 2, Dec. 2020.

[10] S. W. Park, H. Lee, B. C. Kim, Y. Woo, and K. Jun, "On the application of matrix congruence to QUBO formulations for systems of linear equations," *arXiv [quant-ph]*, 01-Nov-2021.
[11] H. Djidjev, "Automaton-based methodology for implementing optimization constraints for quantum annealing," in *Proceedings of the 17th ACM International Conference on Computing Frontiers*, Catania Sicily Italy, 2020.
[12] H. N. Djidjev, "Quantum annealing with inequality constraints: The set cover problem," *Adv. Quantum Technol.*, vol. 6, no. 11, Nov. 2023.
[13] C. F. Higham, D. J. Higham, and F. Tudisco, "Testing a QUBO formulation of core-periphery partitioning on a quantum annealer," *arXiv [cs.SI]*, 05-Jan-2022.
[14] S. J. Leary, A. Bhaskar, and A. J. Keane, "Global approximation and optimization using adjoint computational fluid dynamics codes," *AIAA J.*, vol. 42, no. 3, pp. 631–641, Mar. 2004.
[15] G. P. Liu, X. Han, and C. Jiang, "A novel multi-objective optimization method based on an approximation model management technique," *Comput. Methods Appl. Mech. Eng.*, vol. 197, no. 33–40, pp. 2719–2731, Jun. 2008.
[16] Y. S. Yang, B. S. Jang, Y. S. Yeun, and W. S. Ruy, "Managing approximation models in multiobjective optimization," *Struct. Multidiscipl. Optim.*, vol. 25, no. 2, pp. 128–129, Jul. 2003.
[17] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, "Robust stochastic approximation approach to stochastic programming," *SIAM J. Optim.*, vol. 19, no. 4, pp. 1574–1609, Jan. 2009.
[18] G. Rosenberg, M. Vazifeh, B. Woods, and E. Haber, "Building an iterative heuristic solver for a quantum annealer," *Comput. Optim. Appl.*, vol. 65, no. 3, pp. 845–869, Dec. 2016.
[19] H. N. Djidjev, G. Chapuis, G. Hahn, and G. Rizk, "Efficient combinatorial optimization using quantum annealing," *arXiv [quant-ph]*, 25-Jan-2018.
[20] J.-R. Jiang and C.-W. Chu, "Classifying and benchmarking quantum annealing algorithms based on quadratic unconstrained binary optimization for solving NP-hard problems," *IEEE Access*, vol. 11, pp. 104165–104178, 2023.