

Assessing Quantum Annealing to solve the Minimum vertex multicut

Ali Abbassi
Orange Research, LIST3N
University of Technology of Troyes
Chatillon-Troyes, France
ali.abbassi@utt.fr/orange.com

Yann Dujardin
Orange Research
Orange
Chatillon, France
yann.dujardin@orange.com

Eric Gourdin
Orange Research
Orange
Chatillon, France
eric.gourdin@orange.com

Philippe Lacomme
LIMOS - UMR CNRS 6158 ISIMA - Institut d'Informatique d'Auvergne
Université Clermont Auvergne
Campus Universitaire des Cézeaux 1, rue de la Chebarde 63177 Aubière, France
philippe.lacomme@isima.fr

Caroline Prodhon
LIST3N
University of Technology of Troyes
Troyes, France
caroline.prodhon@utt.fr

Abstract—Cybersecurity in telecommunication networks often leads to hard combinatorial optimization problems that are challenging to solve with classical methods. This work investigates the practical feasibility of using quantum annealing to address the Restricted Vertex Minimum Multicut Problem. The problem is formulated as a Quadratic Unconstrained Binary Optimization model and implemented on D-Wave’s quantum annealer. Rather than focusing on solution quality alone, we analyze key aspects of the quantum workflow including minor embedding techniques, chain length, topology constraints, chain strength selection, unembedding procedures, and postprocessing. Our results show that quantum annealing faces substantial hardware-level constraints limitations in embedding and scalability, especially for large instances, while hybrid quantum-classical solvers provide improved feasibility. This study offers a realistic assessment of the D-Wave system’s current capabilities and identifies crucial parameters that govern the success of quantum optimization in cybersecurity-related network problems.

Index Terms—Combinatorial Optimization; Graphs and Networks; Telecommunication Applications

I. INTRODUCTION

Modern telecommunication networks must continuously contend with the challenges of ensuring secure, reliable, and efficient communication, especially in the face of adversarial attacks or system failures. Many of the algorithmic challenges that arise in this context can be framed as hard combinatorial optimization problems

A fundamental example is the *Vertex Minimum Multicut* (VMMC) problem. Given an undirected graph and a set of terminal node pairs, the objective is to identify the smallest set of non-terminal vertices whose removal disconnects all specified terminal pairs. This problem arises in numerous real-world scenarios including intrusion containment, resilient routing, and the design of survivable communication infrastructure [7], [20]. VMMC also serves as a natural abstraction for tasks involving attack path disruption or selective isolation in security-focused network configurations.

The VMMC problem is known to be **NP-complete** even on restricted graph classes such as trees [9], [15], and its computational hardness persists under constraints that commonly occur in practical applications, such as prohibiting terminal node removals. Although polynomial-time solutions exist for special classes such as interval [15] and co-bipartite graphs [19], the general case resists tractable resolution.

Over the past decades, several approaches have been developed to cope with the problem’s inherent complexity. These include integer linear programming (ILP) models, constant-factor and logarithmic approximation algorithms [6], [12], and more recently, fixed-parameter tractable (FPT) strategies that exploit structural parameters of the input [5], [18]. However, classical algorithms often face scaling challenges with the size and connectivity of the network, especially when the number of terminal pairs grows or the graph exhibits irregular or high-degree topology.

In this work, we explore an alternative computational paradigm: quantum annealing, implemented on the D-Wave quantum processing unit (QPU). Quantum annealing is a meta-heuristic that aims to find ground state-energy (i.e., optimal or near-optimal) configurations of binary optimization problems encoded as *Quadratic Unconstrained Binary Optimization* (QUBO) models or their Ising equivalents [11], [13]. D-Wave’s architecture, designed specifically for such formulations, offers a unique opportunity to experimentally assess the performance of quantum hardware on NP-hard graph problems.

We focus specifically on tree-structured instances of the Vertex Minimum Multicut problem, for which QUBO formulations are constructed using path-based constraints. Our goal is not only to encode and solve these instances using quantum annealing, but also to evaluate the practical bottlenecks in the process. These include limitations in embedding due to hardware topology [8], chain stability [4], and the trade-offs involved in hybrid quantum-classical solvers such as those used by D-Wave’s Advantage platform [10].

By focusing on a well-established yet computationally challenging problem, we aim to evaluate the practical viability of quantum annealing for a variant that, to the best of our knowledge, has not been previously studied in this context. In addition to proposing a QUBO formulation, we analyze the impact of hardware constraints, embedding strategies, and solver behavior on both solution quality and runtime performance. The remainder of this paper is structured as follows. Section 2 formalizes the Vertex Minimum Multicut problem, discusses its complexity landscape, and outlines the path-based formulation. Section 3 details the quantum annealing workflow and the role of D-Wave’s QPU. Section 4 presents empirical results and discusses practical observations. Section 5 concludes the paper.

II. PROBLEM STATEMENT

Multicut problems are central in communication network design, reliability, and security, where one seeks to disconnect specific terminal pairs by removing a minimum set of graph elements (edges or vertices). Several variants of the problem exist depending on the type of cut (edge vs. vertex) and whether cuts can include terminal nodes.

In this paper, we focus on the Restricted Vertex Minimum Multicut (RVMMC) problem, in which the goal is to remove a minimal number of *non-terminal vertices* such that all given terminal pairs are disconnected. For brevity, we refer to this problem as the Vertex Minimum Multicut (VMMC) throughout the remainder of the paper.

Let $G = (V, E)$ be an undirected graph and $H = \{(s_i, t_i) \mid i = 1, \dots, k\}$, where $k \in \mathbb{N}$, and $H \subseteq V \times V$ a set of terminal pairs. The objective is to find a subset $C \subseteq V$ of non terminal vertices such that for every $(s_i, t_i) \in H$, all paths connecting s_i and t_i are disrupted by at least one vertex in C .

This problem is NP-complete even on trees of bounded degree [8], and intractable on graph classes like cographs and split graphs. However, it admits polynomial-time solutions on interval graphs [15], permutation graphs, and co-bipartite graphs [19]. Moreover, the parameterized version is fixed-parameter tractable with respect to the size of the cutset [5], [18]. Related problems have been studied under different formulations. The multi-terminal vertex separator was addressed using Branch-and-Price methods in [17]. QUBO models have been proposed for the edge multicut on trees using D-Wave [8] and for the multiway cut [16] with optimized formulations. Table I summarizes key distinctions across multicut variants.

TABLE I
COMPARISON OF CUT-BASED GRAPH SEPARATION PROBLEMS.

Problem	Cut Type	Terminal Removal	Complexity (General Graph)
Min-Cut max flow	Edge	Allowed	Polynomial-time
Multiway Cut	Edge/Vertex	Depends on model	NP-Complete for $ T \geq 3$ [9]
Edge Multicut	Edge	Allowed	NP-Complete for $ H \geq 3$ [8]
Restricted Vertex Multicut	Vertex	Not Allowed	NP-Complete, even on trees [15]

We now introduce a path-based Binary Linear Program (BLP) to model VMMC. Each vertex $v \in V$ is assigned a binary variable $x_v \in \{0, 1\}$, where $x_v = 1$ if the vertex is selected for removal. Let P_{s_i, t_i} denotes the set of all $s_i - t_i$ paths, $i = 1, \dots, k \in \mathbb{N}$. The objective is to minimize the total number of removed vertices while satisfying two constraints:

$$\begin{aligned}
 \min \quad & \sum_{v \in V} x_v \\
 \text{s.t.} \quad & x_v = 0 \quad \forall v \in V_H \quad (\text{C1}) \\
 & \sum_{v \in \pi} x_v \geq 1 \quad \forall (s_i, t_i) \in H, \forall \pi \in P_{s_i, t_i} \quad (\text{C2}) \\
 & x_v \in \{0, 1\}
 \end{aligned} \tag{1}$$

The constraint (C1) forbids the removal of terminal vertices. Constraint (C2) ensures that there exists no path between each terminal pair source target vertices. This formulation does not account for redundancy among paths, but provides a suitable base for QUBO conversion and quantum optimization experiments.

III. QUANTUM OPTIMIZATION

Quantum computing introduces a novel paradigm for addressing combinatorial optimization problems by leveraging quantum mechanical phenomena such as superposition and tunneling, providing an efficient scheme to efficiently solve large-scale instances to explore large solution spaces [1]. Although general NP-complete problems, including SAT, are widely believed to be not solvable in polynomial time by quantum computers (that is, they are not in BQP) [1], quantum approaches can still offer substantial advantages in practice. In particular, Grover’s algorithm provides at most a quadratic speedup in the unstructured search setting [14], reinforcing the theoretical limits of black-box optimization. However, these limitations do not rule out the possibility of significant quantum speedups on structured problem instances commonly encountered in real-world settings [2].

Motivated by this nuance, quantum heuristic methods such as quantum annealing are receiving increasing attention for solving NP-hard combinatorial problems where structure can be exploited. In this work, we focus on the application of quantum annealing to instances of the VMMC. By formulating the latter as a Quadratic Unconstrained Binary Optimization (QUBO) model, we explore the effectiveness of D-Wave’s quantum processing unit.

A. Quantum Computing

Optimization problems can be addressed using gate-based quantum algorithms, such as Grover’s adaptive search or the Quantum Approximate Optimization Algorithm (QAOA), or using analog approaches such as Quantum Annealing (QA) [1], [3]. For a quantum device to process an optimization problem, it must be formulated in a compatible mathematical model. This is typically a degree-two Boolean polynomial encoded as either a QUBO model or an Ising Hamiltonian

for most of the analogue machines in contrast to higher degree Hamiltonian for digital machines. In gate-based quantum computers, the QUBO is mapped to quantum circuits consisting of unitary operations acting on qubits in a Hilbert space. The evolution of the quantum state leads to a probabilistic measurement, collapsing to a classical solution.

In contrast, quantum annealers like D-Wave's QPU physically realize optimization by embedding the QUBO or Ising model directly into hardware. The goal is to find the ground state—i.e., the lowest energy configuration—of the encoded cost function. The system is initialized in the ground state of an easy-to-prepare Hamiltonian and gradually evolved toward the problem Hamiltonian under a time-dependent Schrödinger equation [11]. If this evolution is sufficiently slow, the adiabatic theorem ensures that the system remains in its ground state throughout, yielding the optimal or near-optimal solution.

In both models, the quality of the final solution critically depends on the formulation of the problem, hardware constraints, and noise resilience, making the encoding strategy a central design component of quantum optimization workflows [21].

B. Quantum Annealing

Quantum annealing (QA) is a metaheuristic to solve combinatorial optimization problems by exploiting quantum fluctuations, especially quantum tunneling, to escape local minima and explore low-energy configurations. QA relies on a physical process governed by the *quantum adiabatic theorem* [11].

In this framework, the optimization problem is encoded into the ground state of a *problem Hamiltonian* H_p , acting on a Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ associated with n qubits. The system is initialized in the ground state of a *driver Hamiltonian* $H_0 \in \mathcal{L}(\mathcal{H})$, where $\mathcal{L}(\mathcal{H})$ denotes the space of bounded linear operators on \mathcal{H} . The full system evolves according to a time-dependent Hermitian operator:

$$H(t) = A(t)H_0 + B(t)H_p, \quad t \in [0, T], \quad (2)$$

where $A(t), B(t) \in C^\infty([0, T], \mathbb{R})$ are smooth, real-valued annealing schedules satisfying $A(0) > 0, B(0) = 0$ and $A(T) = 0, B(T) > 0$. Smoothness ensures differentiability at all orders, which is essential for the adiabatic condition.

If the system evolves slowly enough and the instantaneous spectral gap $\Delta(t) = E_1(t) - E_0(t)$ between the ground state and the first excited state remains strictly positive for all $t \in [0, T]$, the adiabatic theorem guarantees that the system remains in its instantaneous ground state throughout the evolution, ending in a state close to the ground state of H_p . On D-Wave quantum annealers, H_p is implemented as an Ising Hamiltonian:

$$H_p = \sum_{i=1}^n h_i Z_i + \sum_{1 \leq i < j \leq n} J_{ij} Z_i Z_j, \quad (3)$$

where Z_i is the Pauli- Z operator acting on qubit i , with eigenvalues ± 1 , and $h_i, J_{ij} \in \mathbb{R}$ are tunable scalar parameters representing local fields and pairwise couplings, respectively.

In the computational basis, the Hamiltonian is diagonal and defines a classical energy function:

$$E(\mathbf{s}) = \sum_i h_i s_i + \sum_{i < j} J_{ij} s_i s_j, \quad \mathbf{s} \in \{-1, +1\}^n,$$

whose minimization corresponds to identifying the ground state of H_p .

C. QUBO Formulation

Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ be a vector of binary decision variables. A general *pseudo-Boolean function* is a multivariate polynomial [8] $f : \{0, 1\}^n \rightarrow \mathbb{R}$ defined as:

$$f(\mathbf{x}) = \sum_{S \subseteq \{1, \dots, n\}} c_S \prod_{j \in S} x_j, \quad (4)$$

where $c_S \in \mathbb{R}$. We define $\deg(f) = \max\{|S| : c_S \neq 0\}$, when $\deg(f) = 2$, the function is called a *quadratic pseudo-Boolean function*, which can be written as:

$$f(\mathbf{x}) = \sum_{i=1}^n u_i x_i + \sum_{1 \leq i < j \leq n} w_{ij} x_i x_j. \quad (5)$$

This yields the canonical QUBO form:

$$\min_{\mathbf{x} \in \{0, 1\}^n} \mathbf{x}^\top Q \mathbf{x}, \quad Q \in \mathbb{R}^{n \times n}, \quad (6)$$

where Q is a symmetric cost matrix. In this model, the diagonal entries Q_{ii} represent linear terms and the off-diagonal entries Q_{ij} capture pairwise interactions. The QUBO formulation is closely related to the Ising model in physics. Via a variable substitution $x_i = \frac{1+s_i}{2}$ where $s_i \in \{-1, +1\}$, one can map the QUBO into an equivalent Ising Hamiltonian. This allows direct hardware implementation on quantum annealers such as D-Wave, which evolve the system toward the ground state of the encoded energy landscape. In practice, the BLP is transformed into QUBO form by incorporating constraint violations as penalty terms. These penalty terms are weighted by hyperparameters that must be tuned to enforce the constraints while avoiding numerical instabilities or misrepresentations of the energy landscape [13].

Let $M_1, M_2 \in \mathbb{R}$ denote the penalty coefficients for the terminal, and path constraints, respectively. Given binary variables $x_v \in \{0, 1\}$, the QUBO formulation becomes:

$$\begin{aligned} \mathbf{x}^\top H_p \mathbf{x} &= \sum_{v \in V} x_v \\ &+ M_1 \left(\sum_{v \in V_H} x_v \right)^2 \quad (C1) \\ &+ M_2 \sum_{(s_i, t_i) \in H} \sum_{\pi \in P_{s_i, t_i}} \left(\sum_{v \in \pi} (1 - x_v) \right)^2 \quad (C2) \end{aligned} \quad (7)$$

This expression encodes the objective and constraints as penalty terms into a single quadratic function that the quantum annealer minimizes. While this formulation avoids slack

variables, alternative versions may introduce auxiliary binary slack variables to transform inequality constraints into equalities [13]. For example, a path constraint of the form $\sum_{v \in \pi} x_v \geq 1$ can be reformulated as

$$\left(\sum_{v \in \pi} x_v + \sum_{j=0}^{\lceil \log_2 M_\pi \rceil - 1} 2^j \cdot y_j^\pi - 1 \right)^2, \forall \pi \in P_{s_i, t_i}$$

where $y_j^\pi \in \{0, 1\}$ are binary slack variables and M_π is the number of nodes in the path. This formulation increases expressiveness but also introduces additional variables and quadratic terms, which may impact embedding feasibility and resource usage on quantum hardware.

D. D-Wave’s Quantum Annealing Workflow

Once the VMMC problem is formulated as a QUBO, solving it on D-Wave’s quantum annealer involves a nontrivial compilation pipeline. In particular, the path-based encoding of terminal separations introduces a formulation overhead, as the number of constraints scales with the set of enumerated paths. The resulting QUBO graph must then be embedded into D-Wave’s hardware topology—such as Pegasus or Zephyr—which imposes strict qubit connectivity constraints. Embedding logical variables onto chains of physical qubits is itself NP-hard and requires additional qubits and couplers [4].

To address scalability and noise-related issues, we adopt D-Wave’s hybrid solver based on the racing strategy [10]. This approach concurrently evaluates multiple classical-quantum solver instances, dynamically prioritizing those with faster convergence. decomposition enables better utilization of both classical heuristics and quantum tunneling for refining partial solutions.

IV. NUMERICAL ANALYSIS

This section presents a performance evaluation of QUBO for the Vertex Minimum Multicut VMMC problem on tree-structured instances. The study investigates solver behavior, D-Wave hardware bottlenecks, and the effect of parameter tuning. We generate random unweighted trees where each terminal pair shares a unique path. This guarantees feasibility of the path-based QUBO formulation while preserving nontrivial embedding complexity. Each instance is built as a random spanning tree with carefully selected non-adjacent terminal pairs to avoid degenerate cuts. The study is focused on 9 types of instances varying from number of terminals pair equals to 3 until 100 and number of vertices from 20 to 400

To ensure realistic performance, we focus on instances where the size and structure reflect practical limitations in embedding and qubit resources. This is particularly relevant when scaling up the number of terminal pairs.

A. QUBO Parameter Tuning

We conduct a grid search over penalty coefficients, using simulated annealing to estimate feasibility likelihoods and ground state convergence. Monte Carlo sampling is also

applied to validate penalty robustness. Penalty weights were calibrated empirically to avoid excessive dominance (which could flatten the energy landscape) or weakness (which risks infeasible solutions). Final values are scaled based on graph size and the number of terminal pairs.

B. Solvers and Evaluation Metrics

Performance is assessed using four key metrics. First, solution quality is evaluated through the *optimality gap*, defined as $\frac{f_{\text{found}} - f^*}{f^*} \times 100\%$, where f^* is the optimal cost obtained via the MILP solver when tractable. In larger cases, where MILP is computationally infeasible, only relative comparison across solvers is reported. Sampling time, which includes annealing and postprocessing duration, is measured to reflect the computational cost. Embedding difficulty is quantified using chain statistics (length, break count), while feasibility rate captures the proportion of solutions that satisfy all constraints imposed by the original BLP formulation

TABLE II
SOLVER BACKENDS.

Type	Backend
Quantum Annealing	D-Wave Advantage 4.1 QPU
Classical Heuristic	Simulated Annealing (custom)
Exact Solver	CPLEX
Hybrid Solver	Leap Hybrid Solver (Racing)

Additional metrics include QPU sampling time and solution feasibility. Simulated Annealing is tuned using inverse temperature ranges (0.1, 10.0) and (0.1, 20.0), sweep counts {100, 200, 500}, and both geometric and linear cooling schedules. Shot counts vary from 10 to 1000. For each configuration, we log the penalty weights M_1, M_2 , the chosen annealing schedule, and sampling parameters. Chain statistics are only discussed in relation to feasibility breakdowns when relevant.

C. D-Wave Backend Configuration

Experiments were executed on D-Wave’s Advantage System 4.1 (eu-central-1), which supports the Pegasus P16 topology. Embedding was performed using the *MinorMiner* algorithm with emphasis on minimizing chain lengths and avoiding chain breaks. The annealing time was fixed at $20\mu\text{s}$, selected within the system’s available range.

The hybrid solver was run in *racing mode*, where multiple classical and quantum samplers are launched in parallel. This improves response time and solution robustness. Coupling strength J was tuned over the extended range $[-2, 1]$ to stabilize qubit chains and reduce fragmentation.

D. Findings and Discussion

We evaluate nine instances of increasing complexity by simultaneously scaling the number of vertices and terminal pairs. Instance I1 is the smallest, with $|H| = 3$ and $|V| = 24$, while instance I9 is the largest, with $|H| = 100$ and $|V| = 450$. As shown in Fig. 1, we compare the minimum QUBO energy and the size of the associated cutset obtained from four solver types. Simulated annealing (SA) consistently achieves better

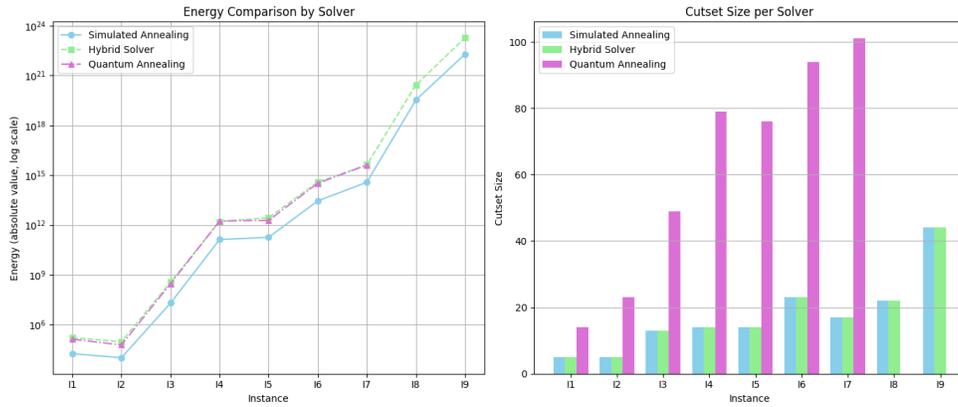


Fig. 1. Left: Solver energy comparison (absolute value log scale) across instances. Right: Associated cutset sizes indicating solution feasibility.

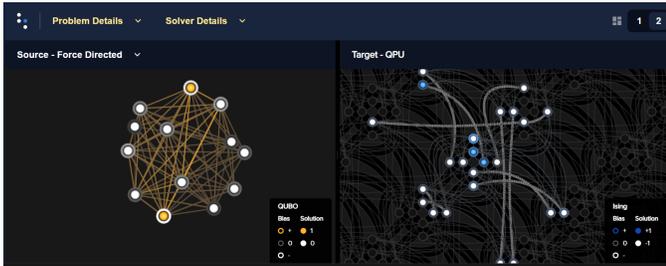


Fig. 2. Source-to-QPU mapping and chain structure using D-wave inspector tool for a small instance.

cutset sizes, indicating better quality solutions—about 10 to 15% gap—, although it does not always reach the lowest energy configurations. Hybrid solvers tend to minimize energy more effectively than classical methods, yet this advantage does not necessarily translate into better solutions due to limitations in the energy landscape encoding.

A key observation is that the quantum annealing solver often achieves low-energy solutions that are infeasible with respect to the original VMMC constraints. For example, in Instance I7, both the hybrid and quantum annealers reach comparable QUBO energies, but the quantum solver’s cutset deviates by over 300% from the exact solution. Beginning with Instance I5, quantum annealing fails to yield any feasible solution, and for instances I8 and I9, embedding the QUBO into the QPU was not possible regardless of the embedding strategy. Multiple approaches were attempted, including MinorMiner, LazyFixedEmbeddingComposite, Uniform Torque Compensation, and clique-based embeddings, all without success.

In smaller instances, successful embeddings were possible due to lower overhead. However, as the number of terminal pairs increases, so does the number of path constraints and quadratic terms in the QUBO, leading to complex connectivity requirements. These exacerbate the embedding process and introduce chain instability, often preventing valid solutions from being recovered during unembedding.

To better understand degradation in quantum annealing performance, we analyzed variable embeddings (Fig. 2). Im-

proper chain strengths often result in broken chains, which introduce infeasibility even when energy values appear low. In the visualized instance, increasing chain strength improves embedding stability, enabling the recovery of feasible solutions. This highlights the importance of careful unembedding strategies and accurate majority-vote decoding in extracting valid variable assignments.

The non-linear relationship between logical variables and physical qubits is further illustrated in Fig. 3. Even instances with the same number of logical variables exhibit drastically different qubit counts, depending on the density of interaction terms. This is intrinsic to how QUBOs encode pairwise constraints and reflects the limitations of the Pegasus topology, which supports an average chain length of about 5 under ideal connectivity.

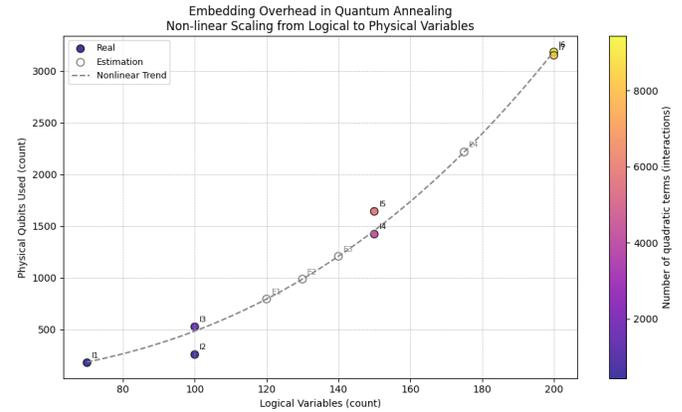


Fig. 3. Embedding overhead: non-linear scaling from logical to physical qubits in D-Wave quantum annealing.

Finally, Fig. 4 compares solver runtimes. While the QPU’s physical annealing time remains low—dominated by sub-millisecond sampling—the total runtime is significantly affected by the compilation pipeline, including QUBO instantiation, embedding, and unembedding. Simulated annealing provides the best trade-off between runtime and solution quality for small to medium instances. Quantum annealing

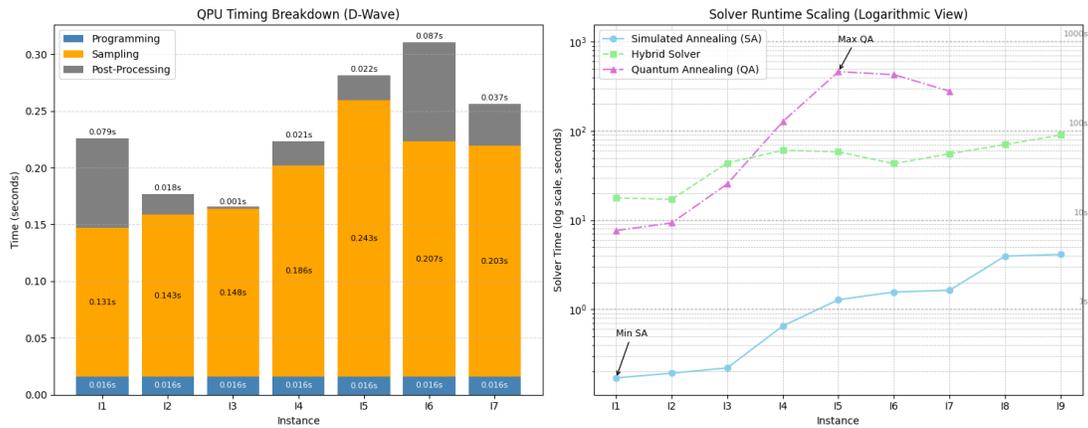


Fig. 4. Solver runtime comparison and QPU timing breakdown across instances.

achieves competitive runtimes for small instances but quickly becomes impractical due to embedding failures. The hybrid solver balances both regimes by decomposing QUBOs into tractable subproblems, making it a compelling choice for intermediate instance sizes.

V. CONCLUSION AND FUTURE WORK

This work presents an experimental evaluation of quantum annealing for the Restricted Vertex Minimum Multicut (VMMC) problem, modeled as a QUBO and solved using D-Wave’s quantum hardware. While the problem is naturally suited to combinatorial optimization, encoding it as a path-based QUBO introduces significant overhead as the number of terminal pairs increases.

Although quantum annealing often reaches low-energy configurations, these do not consistently yield feasible solutions due to broken qubit chains and embedding limitations. The embedding process, which maps logical variables to physical hardware qubits, introduces non-linear overhead stemming from hardware connectivity constraints and QUBO density. As a result, scalability suffers and solution quality degrades in larger instances. Hybrid quantum-classical solvers demonstrate greater robustness by decomposing the problem into tractable substructures, achieving feasible solutions at larger scales. Simulated annealing also remains competitive in both runtime and cutset quality.

Future work will focus on leveraging problem structure to reduce the number of quadratic interactions in the QUBO, enabling more compact embeddings. More strategic control over chain strength and adaptive embedding heuristics may improve solution feasibility in quantum-only runs.

While full quantum advantage remains elusive, our findings show that hybrid workflows and tailored formulations can effectively contribute to real-world cybersecurity optimization.

REFERENCES

- [1] A. Abbas et al “Challenges and opportunities in quantum optimization,” *Nature Reviews Physics*, 2024.
- [2] S. Aaronson, “How much structure is needed for huge quantum speedups?,” *arXiv:2209.06930*, 2022.
- [3] A. Bochkarev, R. Heese, S. Jäger, P. Schiewe, and A. Schöbel, “Quantum computing for discrete optimization: A highlight of three technologies,” *arXiv:2409.01373v1*, 2024.
- [4] A. Bertuzzi, D. Ferrari, A. Manzalini, and M. Amoretti, “Evaluation of quantum and hybrid solvers for combinatorial optimization,” *arXiv:2403.10455*, 2024.
- [5] N. Bousquet, J. Daligault, and S. Thomassé, “Multicut is FPT,” in *Proc. 43rd ACM Symp. Theory Comput. (STOC)*, 2011, pp. 459–468.
- [6] G. Calinescu, H. Karloff, and Y. Rabani, “Multicuts in unweighted graphs and digraphs with bounded degree and bounded tree-width,” *J. Algorithms*, vol. 48, no. 2, pp. 333–359, 2003.
- [7] M.-C. Costa, D. Coudert, and B. Jaumard, “Minimal multicut and maximal integer multiflow: A survey,” *Eur. J. Oper. Res.*, vol. 162, no. 1, pp. 55–69, 2005.
- [8] W. Cruz-Santos, C. Salas, and J. C. López-Sánchez, “A QUBO formulation of minimum multicut problem instances in trees for D-Wave quantum annealers,” *Sci. Rep.*, vol. 9, p. 1721, 2019.
- [9] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis, “The complexity of multiterminal cuts,” *SIAM J. Comput.*, vol. 23, no. 4, pp. 864–894, 1994.
- [10] D-Wave Systems Inc., “Hybrid workflow solvers,” 2024.
- [11] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, “Quantum computation by adiabatic evolution,” *arXiv:0001106*, 2000.
- [12] N. Garg, V. V. Vazirani, and M. Yannakakis, “Approximate max-flow min-(multi)cut theorems and their applications,” *SIAM J. Comput.*, vol. 25, no. 2, pp. 235–251, 1996.
- [13] F. Glover, G. Kochenberger, and Y. Du, “A tutorial on formulating and using QUBO models,” *arXiv:1811.11538*, 2019.
- [14] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proc. 28th Annu. ACM Symp. Theory Comput. (STOC)*, 1996, pp. 212–219.
- [15] J. Guo and R. Niedermeier, “Complexity and exact algorithms for vertex multicut in interval and bounded treewidth graphs,” *Eur. J. Oper. Res.*, vol. 186, no. 2, pp. 542–553, 2008.
- [16] S. Heidari, M. J. Dinneen, and P. Delmas, “An equivalent QUBO model to the minimum multi-way cut problem,” *CDMTCS Research Reports*, Tech. Rep. CDMTCS-565, 2022.
- [17] Y. Magnouche, A. R. Mahjoub, and S. Martin, “The multi-terminal vertex separator problem: Branch-and-cut-and-price,” *Discrete Appl. Math.*, vol. 286, pp. 168–189, 2020.
- [18] D. Marx and I. Razgon, “Fixed-parameter tractability of multicut parameterized by the size of the cutset,” in *Proc. 43rd ACM Symp. Theory Comput. (STOC)*, 2011, pp. 469–478.
- [19] C. Papadopoulos, “Restricted vertex multicut on permutation graphs,” *Discrete Appl. Math.*, vol. 160, no. 12, pp. 1791–1797, 2012.
- [20] H. S. Stone, “Multiprocessor scheduling with the aid of network flow algorithms,” *IEEE Trans. Softw. Eng.*, vol. SE-3, no. 1, pp. 85–93, 1977.
- [21] S. E. Venegas-Andraca, J. F. Barrera, M. L. Boyer, and A. E. Carrasco, “A cross-disciplinary introduction to quantum annealing-based algorithms,” *Contemp. Phys.*, vol. 59, no. 2, pp. 174–196, 2018.