Minor Embedding for Max-Cut: Empirical Analysis of Quantum Annealing's Combinatorial Bottleneck

Ron LAVI*, Alberto MORAGLIO†, Ofer M. SHIR‡

*Migal - The Galilee Research Institute, Israel

Email: ronl@migal.org.il

†Department of Computer Science, University of Exeter, Exeter, United Kingdom

Email: a.moraglio@exeter.ac.uk

‡Tel-Hai - University of Kiryat Shmona in the Galilee, Israel

Email: ofersh@telhai.ac.il

***Abstract*—Quantum optimization by annealing constitutes a stochastic global search over QUBO-formulated combinatorial optimization problems. Its performance depends on pre-solving the Minor Embedding (ME) challenge—the $\mathcal{NP}$-hard problem of mapping logical graphs representing QUBOs onto the quantum hardware. Despite determining solution feasibility and quality, ME remains understudied as a standalone optimization challenge. Current understanding is largely limited to synthetic graphs on isolated topologies, leaving the generalizability of structural predictors across hardware generations unclear. We conduct the first cross-topology study analyzing the MINORMINER heuristic—the *de facto* industry standard—over ≈550 benchmark graph instances across D-Wave's Chimera, Pegasus, and Zephyr architectures. Our analyses reveal a functional dissociation between embedding feasibility and cost: topological independence (via Maximum Independent Set ratio) predicts feasibility by identifying structural 'slack', while *edge density*—reflecting clique-like proximity—governs the resource 'toll' (chain length). Since ME is a purely graph-theoretic problem, agnostic to the combinatorial origin of the source graph, these structural predictors apply to any QUBO whose logical graph falls within the studied structural regime. This dissociation remains invariant across all architectures, establishing ME as a benchmark for Evolutionary Computation, where the local distribution of edges, rather than global density alone, determines algorithmic difficulty.**

***Index Terms*—Quantum Annealing, Minor Embedding, Combinatorial Optimization, D-Wave, MINORMINER, embeddability, maximum independent set, edge density.**

I. Introduction

Quantum Annealing (QA) is a specialized form of quantum computing designed for Combinatorial Optimization (CO) problems. From a computational intelligence perspective, QA operates as a global stochastic search heuristic that exploits quantum tunneling and superposition to explore solution landscapes, potentially offering advantages over classical optimization methods for certain problem classes (see, e.g., [1]). As QA hardware—primarily produced by D-Wave Systems—continues to mature and scale, understanding the practical factors influencing performance becomes increasingly critical for both theoretical advancement and real-world applications. Solving CO problems on quantum annealers requires formulating them as QUBO models, which define logical graphs of binary decision variables and their interactions. A fundamental challenge in quantum annealing lies in the mismatch between these arbitrary logical structures and the sparse, fixed connectivity of quantum hardware. This challenge manifests as the ME problem: mapping logical problem graphs onto the constrained physical topology of the quantum processor. Despite being a crucial preprocessing step that directly impacts solution quality, computational overhead, and problem-solving feasibility, ME remains surprisingly understudied in the QA literature. Most research either abstracts away embedding considerations entirely or treats them as a necessary nuisance mentioned only in passing.

This research gap has practical consequences. Embedding quality, typically measured by average chain length (ACL), directly impacts solution accuracy through chain breaks, where physical qubits disagree. Moreover, the decision version of the embedding problem—determining whether a valid embedding exists—is $\mathcal{NP}$ [2]; the optimization variant, minimizing chain length, is at least as hard. In practice, this pre-solving step often dominates total solution time, with embedding overhead exceeding QPU execution time by orders of magnitude. For densely connected problems, chain-length costs can reach ratios of 10:1 or higher, severely limiting effective problem sizes. These bottlenecks make understanding the relationship between problem structure and embedding difficulty critical, yet this remains poorly characterized. Current heuristics exhibit substantial instability, with ACL varying by 4+ qubits for identical instances [3].

The development of effective ME strategies has progressed largely within commercial settings, where proprietary constraints limit knowledge dissemination. This asymmetry between industrial practice and published research motivates the need for systematic, independent empirical studies that can inform the broader quantum optimization community.

Research Questions: We address the following questions:

(RQ1) To what extent does *edge density* — commonly conceived as the dominant predictor of ME performance — hold as a sufficient explanatory metric across diverse, established benchmark graph instances?

(RQ2) Which intrinsic graph-theoretic properties govern embedding *feasibility* versus embedding *cost*, and do these predictive relationships hold consistently across the Chimera, Pegasus, and Zephyr hardware generations?

To address these questions, we conduct a systematic empirical study using the Max-Cut problem as our test case. Max-Cut is an ideal benchmark for this investigation: it is $\mathcal{NP}$-hard, well-studied in both classical and quantum optimization communities, and possesses a broad spectrum of established benchmarks. Max-Cut admits direct QUBO formulation without auxiliary variables or penalty terms, allowing us to isolate structural embedding effects from encoding artifacts [4], [5].

Our **contribution** is threefold: (i) we provide a cross-topology empirical study evaluating ME using MINORMINER – the *de facto* industry standard – across ≈550 benchmark graph instances, providing a rigorous baseline for $\mathcal{NP}$-hard mapping on Chimera, Pegasus, and Zephyr architectures; (ii) we establish a structural predictor hierarchy, identifying the Maximum Independent Set (MIS) ratio as a superior predictor of embedding feasibility over traditional edge density; and (iii) we identify a functional dissociation where the MIS ratio governs feasibility while edge density dictates cost, a robust pattern that generalizes across all three hardware generations.

The remainder of this paper is organized as follows. Section II provides background on QA, ME, and the Max-Cut problem, including a review of related work on embedding algorithms and performance studies. Section III describes our experimental methodology, including problem generation, embedding procedures, and performance evaluation. Section IV presents our experimental results, and finally, Section V concludes and outlines directions for future work.

II. Background and Problem Formulation

A. Quantum Annealing

QA is a hardware-accelerated metaheuristic that exploits quantum tunneling and superposition to solve QUBO-formulated combinatorial optimization problems. The quantum processor evolves according to a time-dependent Hamiltonian that transitions from an initial uniform superposition to a final state encoding the problem solution. D-Wave's quantum annealers consist of thousands of physical qubits arranged in sparse, fixed topologies with limited connectivity: Chimera (degree 6), Pegasus (degree 15), and Zephyr (degree 20). Since arbitrary QUBO problems typically require substantially denser connectivity or all-to-all, a preprocessing step—namely ME—maps logical problem variables onto connected chains of physical qubits. This mapping determines problem-solving feasibility and directly impacts solution quality through chain-break effects and resource overhead, motivating our investigation of which structural properties govern ME difficulty.

B. Minor Embedding

Definition (Graph Minor)**.** *A graph G is a* minor *of a graph H, denoted $G \preceq_m H$, if G can be obtained from H by a sequence of vertex deletions, edge deletions, and edge contractions.*

The minor relation is central to the Graph Minor Theory of Robertson and Seymour [6], whose theorem establishes that the relation $\preceq_m$ is a *well-quasi-ordering* on finite graphs. In what follows, we present the ME problem via three perspectives: the formal graph-theoretic definition, the exact combinatorial optimization formulation, and the pragmatic heuristic problem-solving approach by MINORMINER.

Let $G = (V_G, E_G)$ be a *source* (logical) graph and $H = (V_H, E_H)$ a *target* (hardware) graph. We say that G is a *minor* of H, denoted $G \preceq_m H$, if there exists a mapping $\phi : V_G \to 2^{V_H}$, assigning each logical vertex to a subset of physical qubits in H, satisfying the following conditions:

(C1) Disjointness: For all $u \neq v \in V_G$ $\phi(u) \cap \phi(v) = \emptyset$ holds.

(C2) Connectivity: For all $v \in V_G$, the induced subgraph $H[\phi(v)]$ is connected.

(C3) Edge Preservation: For all $(u, v) \in E_G$, there exist $i \in \phi(u)$ and $j \in \phi(v)$ such that $(i, j) \in E_H$.

The set $\phi(v) \subseteq V_H$ is called the *chain* representing logical vertex v. Intuitively, each logical vertex is represented by a connected subtree of H, and logical edges are realized through physical edges connecting distinct chains.

Combinatorial Optimization Exact Formulation

The following description constitutes an exact formulation of the ME problem, which we pose for the Integer Programming (IP) framework. Given source graph $G = (V_G, E_G)$ and target graph $H = (V_H, E_H)$, we define binary decision variables $x_{v,i} \in \{0, 1\}$ indicating whether physical qubit $i \in V_H$ is part of the chain mapping to logical vertex $v \in V_G$. To enforce that each chain forms a connected subgraph, we employ a Miller-Tucker-Zemlin (MTZ) based formulation, which originated as a subtour elimination technique for the Traveling Salesman Problem [7]. For each logical node v, we assign a root physical node r_v and define directed edge variables $y^v_{ij} \in \{0, 1\}$ representing the spanning tree of the chain, alongside continuous ordering variables $u^v_i \in \mathbb{R}$. Let $n = |V_G|$, $m = |E_G|$, $N = |V_H|$, $M = |E_H|$. The formulation is:

$$
\begin{aligned}
\min_{\vec{x},\vec{y},\vec{u}} \quad & f_\ell(\vec{x}) \\
\text{s.t.} \quad & \sum_{v \in V_G} x_{v,i} \le 1 \quad \forall i \in V_H \\
& \sum_{i \in V_H} x_{v,i} \ge 1 \quad \forall v \in V_G \\
& \sum_{(i,j) \in E_H} y^v_{ij} = \sum_{i \in V_H} x_{v,i} - 1 \quad \forall v \in V_G \\
& y^v_{ij} \le x_{v,i}, \quad y^v_{ij} \le x_{v,j} \quad \forall v \in V_G, (i,j) \in E_H \\
& u^v_{r_v} = 0, \quad 1 \le u^v_i \le N - 1 \quad \forall v \in V_G, i \in V_H \setminus r_v \\
& u^v_j \ge u^v_i + 1 - N(1 - y^v_{ij}) \quad \forall v \in V_G, (i,j) \in E_H \\
& \sum_{(i,j) \in E_H} x_{u,i} \cdot x_{v,j} \ge 1 \quad \forall (u,v) \in E_G \ .
\end{aligned}
\tag{1}
$$

The objective function $f_\ell(\vec{x})$ is selected to prioritize either qubit economy or signal integrity. To minimize the *total chain length*, we set it as

$$
f_1(\vec{x}) = \sum_{v \in V_G} \sum_{i \in V_H} x_{v,i} \mapsto \min\,. \tag{2}
$$

Alternatively, to minimize the *maximum chain length* (MCL)—critical for reducing chain breaks and preserving the energy scale—one may adopt a *minimax* objective. To this end, an auxiliary continuous variable $L_{\max}$ is introduced with an additional constraint (to ensure that it captures the cardinality of the largest logical chain):

$$L_{\max} \geq \sum_{i \in V_H} x_{v,i} \quad \forall\, v \in V_G, \tag{3}$$

leading to the following objective function:

$$f_2(\vec{x}) = L_{\max} \mapsto \min . \tag{4}$$

Next, we clarify the constraints in this model (1). The first constraint ensures *physical qubit disjointness*: each physical qubit $i \in V_H$ is assigned to at most one logical vertex. The second constraint guarantees *surjectivity*: every logical vertex $v \in V_G$ must be embedded in at least one physical qubit. The third through sixth constraints collectively enforce that each vertex model (chain) forms a *connected subgraph*. Specifically, the third constraint ensures that the number of active edges in a chain equals the number of active nodes minus one—a necessary condition for a tree. The fourth constraint links edge variables to vertex variables: a chain edge y_{ij}^v can only be active if both endpoints i and j belong to the vertex model of v. The fifth and sixth constraints are the Miller–Tucker–Zemlin (MTZ) subtour elimination constraints, which assign potential values u_i^v to enforce a directed path structure rooted at r_v, thereby guaranteeing connectivity.[1] The seventh constraint is the *pullback condition*: for each logical edge $(u, v) \in E_G$, there must exist at least one physical edge $(i, j) \in E_H$ connecting the vertex models of u and v. This constraint involves a quadratic term $x_{u,i} \cdot x_{v,j}$, which arises from the need to simultaneously verify the placement of two logical nodes. To incorporate this into a conventional *linear* IP framework, we linearize via the McCormick relaxation: auxiliary binary variables $z_{uvij} \in \{0, 1\}$ represent the product $x_{u,i} \cdot x_{v,j}$ through the linear constraints:

$$z_{uvij} \leq x_{u,i}, \quad z_{uvij} \leq x_{v,j}, \quad z_{uvij} \geq x_{u,i} + x_{v,j} - 1. \tag{5}$$

The pullback condition then becomes $\sum_{(i,j) \in E_H} z_{uvij} \geq 1$ for all $(u, v) \in E_G$. This transformation increases the variable count by $O(|E_G| \cdot |E_H|)$ but renders the formulation linear.
Addressing ME by IP was reported in [8], where a Monolithic IP approach was taken to tackle an equivalent of the aforementioned modeling.

QUBO Perspective: The ME problem can itself be cast as a QUBO by consolidating constraints into penalties:

$$Q = q_0 \sum_{v,i} x_{v,i} + q_1 \sum_i \left(\sum_v x_{v,i} - 1 \right)^2 + q_2 P_{\mathrm{e}} + q_3 P_{\mathrm{c}} \tag{6}$$

where P_{e} penalizes missing inter-chain edges and P_{c} penalizes disconnected chains; penalty weights must dominate q_0 to enforce feasibility. This creates a self-referential structure: embedding the embedding problem requires solving a QUBO of size $\mathcal{O}(nN + nM)$ with dense couplings.

MINORMINER *as a Fast Heuristic Approach*

While the IP formulation provides a mathematically exact framework for the ME problem, practical applications on hardware-scale graphs necessitate the use of fast heuristic approaches. The *de facto* standard in the field is the MINORMINER algorithm, introduced by Cai, Macready, and Roy [9]. It employs a randomized, iterative path-finding strategy characterized by local search refinements. It begins with an initial placement phase, where logical vertices V_G are mapped to seed qubits in V_H through layout-aware heuristics or random assignment. To satisfy the connectivity requirements of E_G, the algorithm iteratively identifies shortest paths in the target graph H to bridge disjoint chains. When the path-finding subroutine encounters a bottleneck—often due to physical qubit contention—it selectively removes problematic chains to re-route them through alternative, less congested regions of the hardware graph. As detailed in Algorithm 1, this loop continues until all logical edges are preserved or the maximum iteration threshold is reached.

Algorithm 1 MINORMINER$(G, H, \texttt{max_iter})$

Require: Source graph $G = (V_G, E_G)$, target graph $H = (V_H, E_H)$
Ensure: Embedding $\phi : V_G \to 2^{V_H}$ or FAILURE

- $\phi \leftarrow$ INITIALPLACEMENT(G, H)
- **for** $t = 1$ to `max_iter` **do**
  - $E_{\text{unsat}} \leftarrow \{(u, v) \in E_G : \phi(u) \not\sim_H \phi(v)\}$
  - **if** $E_{\text{unsat}} = \emptyset$ **then**
    - **return** ϕ
  - **end if**
  - **for** $(u, v) \in E_{\text{unsat}}$ **do**
    - $P \leftarrow$ SHORTESTPATH$_H(\phi(u), \phi(v))$
    - **if** $P \neq \emptyset$ **then**
      - Split P into P_u and P_v to establish adjacency
      - $\phi(u) \leftarrow \phi(u) \cup P_u$
      - $\phi(v) \leftarrow \phi(v) \cup P_v$
    - **else**
      - $\phi \leftarrow$ TEARANDREPAIR(ϕ, u, v, H)
    - **end if**
  - **end for**
  - $\phi \leftarrow$ CONTRACTCHAINS(ϕ, H) // Pruning qubits
- **end for**
- **return** FAILURE

The computational efficiency of MINORMINER is a primary advantage; with a per-iteration complexity of approximately $O(|E_G| \cdot (|V_H| + |E_H|))$ using shortest-path computations on weighted graphs it can embed moderate-sized instances ($n \lesssim 200$) in seconds. At the same time, since the algorithm

[1] In this formulation, the root $r_v \in V_H$ is a fixed parameter for each logical vertex $v \in V_G$. Consequently, the model finds the optimal embedding *given* a specific set of root locations. Finding the globally optimal embedding requires iterating over possible root assignments or heuristically pre-selecting them.

prioritizes feasibility over chain length, it offers no guarantees on the total number of qubits utilized. Furthermore, its performance is highly sensitive to the initial placement, often necessitating multiple restarts when the source graph density approaches the hardware's capacity limits.

C. Related Work and Prior Art

The ME problem has been studied primarily in the context of D-Wave's quantum annealers. While D-Wave has accumulated significant expertise through years of hardware deployment and customer engagements, commercial considerations have limited the extent to which this knowledge has been shared with the academic community. As a result, published empirical studies on embedding quality and its impact on solver performance remain scarce, leaving important practical questions inadequately addressed in the open literature.

The primary line of ME-related research involves the development of heuristic approaches as alternatives to exact optimization or as extensions for MINORMINER. Goodrich et al. [10] introduced a graph-theoretic framework that exploits structural properties of the input graph, notably using odd cycle transversal decompositions to identify bipartite substructure and a biclique virtual hardware layer to simplify the interface to the physical topology. Sugie et al. [11] proposed using Simulated Annealing as a meta-heuristic for constructing graph minors, demonstrating embeddings that accommodate up to 50% more binary variables than state-of-the-art methods on sparse hardware graphs. Pinilla and Wilton [12] introduced *layout-aware embedding*, wherein geometric information about both source and target graphs guides qubit allocation; this approach reduces the search space, improves embedding success rates, and yields mappings that produce lower-energy samples on D-Wave hardware. More recently, Zbinden et al. [13] developed two topology-aware algorithms—Spring-Based MinorMiner (SPMM) and Clique-Based MinorMiner (CLMM)—targeting both Chimera and Pegasus architectures. Their evaluation on random graph ensembles revealed that CLMM outperforms competing methods for denser graphs (edge density > 0.08), while SPMM excels on sparser instances. Collectively, these heuristics trade optimality guarantees for scalability, addressing the practical need to embed problems rapidly on current and near-term quantum annealing hardware.

Most recently, a study by Gómez-Tejedor and co-workers [3] investigated aspects of the ME challenge partly aligned with our study by assessing the MINORMINER performance on a set of randomly generated graph instances. It provided an empirical analysis by focusing exclusively on the D-Wave Advantage system with Pegasus topology. Their study examines Erdős-Rényi random graphs across varying sizes and densities, demonstrating an exponential correlation between ACL and solution quality degradation. A key contribution is their comparative evaluation of MINORMINER against D-Wave's deterministic Clique Embedding algorithm, revealing significant instability in MINORMINER's output quality, with ACL variations of up to 4 qubits per chain for identical problem instances. While this work provides valuable insights into embedding quality metrics, it is limited to a **single hardware topology and generic random graph instances without consideration of problem-specific structural properties**.

D. The Max-Cut Problem

Given an undirected graph $G = (V, E)$ with edge weights $w_{ij} \geq 0$ for $(i, j) \in E$, the MAX-CUT PROBLEM seeks a bipartition of V that maximizes the total weight of edges crossing the partition. Due to space limitations, we omit its explicit formulation and refer the reader to [14]. Max-Cut holds a unique position in the QUBO landscape as one of the few $\mathcal{NP}$-hard problems admitting a direct formulation without auxiliary variables or penalty terms [4], [5]. Unlike problems requiring constraint penalization, which introduces variable overhead and parameter sensitivity [14], [15], Max-Cut's objective naturally aligns with the unconstrained quadratic binary form. This makes it a pristine benchmark for solver performance, isolating intrinsic optimization capability from the artifacts of reduction [5]. Theoretically, it anchors the hierarchy of pseudo-Boolean optimization [16], serving as the degree-2 baseline from which higher-order problems are derived via quadratization [15], [17].

E. Novelty and Concrete Contributions

In contrast to the reported study in [3], the present work investigates the ME problem across the full lineage of D-Wave hardware topologies—Chimera, Pegasus, and Zephyr. Furthermore, we transition from synthetic random graph ensembles to a comprehensive dataset of established Max-Cut benchmark instances. By quantitatively characterizing these instances through a high-dimensional suite of graph-theoretic metrics, we identify the structural invariants that govern embeddability. Our primary contributions are as follows:

1) **Formal Problem Characterization:** A complete ME formulation as a standalone CO challenge. This framing is designed to facilitate the development of novel heuristics within the Computational Intelligence community.
2) **Empirical Evaluation of MINORMINER:** Systematic study of ME across diverse graph structures, exposing the limits of density-centric performance prediction.
3) **Establishment of a Predictive Hierarchy:** A structural hierarchy of graph-theoretic predictors—led by the MIS ratio—that characterizes the transition between embedding feasibility (the 'gatekeeper' effect) and routing cost (the 'toll' effect).

III. METHODOLOGY

A. Experimental Setup

We benchmark the behavior of a standard ME heuristic on curated Max-Cut graph instances under three canonical D-Wave target topology families. Experiments do not involve any downstream Max-Cut solving; the output of this section is embedding feasibility and embedding-quality proxies (e.g., ACL). We vary target-topology family and source-graph size;

the dataset, graph-metric extraction, and embedding metrics are otherwise fixed.

Benchmark instances (MQLib subset). MQLib (Max-Cut and QUBO Library) [18] is an open-source C++ library providing a unified benchmarking environment for Max-Cut and QUBO heuristics and serving as a standard reference for validating new techniques against established baselines. Importantly, MQLib provides an extensive repository of Max-Cut benchmark instances (≈3500 graphs), out of which we extract our source graphs. To obtain a controlled size sweep, we construct a size-stratified benchmark set by partitioning instances into 13 blocks of ≈40 graphs (total $N = 548$), spanning $|V| \in [50, 540]$. Each instance is parsed as an undirected, weighted graph and carries a full vector of precomputed MQLib graph metrics as metadata.

Target topologies (idealized logical hardware graphs). Targets are instantiated as defect-free topology graphs using `dwave_networkx`. Pegasus corresponds to the Advantage family (typically > 5000 qubits and degree 15 at the topology level) [19], [20]. Zephyr corresponds to Advantage2[2]. (typically $4400+$ qubits and degree 20) [21]. For historical scale reference, Chimera systems are commonly associated with $C16$ (2048 qubits) [22]; in our experiments we use an extended Chimera instance ($C25$) to enable comparable families at a similar qubit-count scale.

Embedding heuristic and budgets. We use D-Wave's open-source heuristic `minorminer.find_embedding`. Unless stated otherwise, we run the default parameterization; to probe budget sensitivity, we also evaluate a higher-budget preset that increases the restart budget and time limit, and relaxes stopping criteria (conceptually: "try longer and restart more").

Protocol. For each topology family and each source-graph block, we:

1) Load instances together with their precomputed MQLib metrics and metadata;
2) Construct the idealized target topology graph ($C25$, $P16$, or $Z12$);
3) Run R independent randomized ME attempts per instance (in this study, $R = 10$), a value balancing statistical reliability against the substantial computational cost of each embedding attempt;
4) Log per-attempt outcomes and embedding statistics for downstream analysis.

Metrics and logging. Each embedding attempt yields a record containing: instance identifiers, graph-theoretic metrics, success flag, runtime, and embedding structural statistics. We report chain-quality primarily via ACL and also track maximum chain length, number of chains, and total physical qubits used. Failures are logged explicitly and included in success rate estimates; chain statistics are computed only for successful embeddings.

Execution environment. All experiments were executed on a Linux server equipped with an Intel Xeon E5-4669 v4 @ 3.0GHz (88 logical processors). The benchmark framework is implemented in Python 3.12 with D-Wave Ocean and MQLib libraries.[3]

B. Graph-Theoretic Metrics

MQLib additionally provides a broad set of graph-theoretic metrics for characterizing graph topology. The computed structural properties are listed in Table I. Two metrics warrant particular attention:

- The edge density ρ is defined as the ratio of the number of edges m to the number of possible edges in a complete graph with n vertices, given by $\rho = \frac{2m}{n(n-1)}$.
- The Maximum Independent Set (MIS): An independent set is defined as a subset of vertices where no two vertices share an edge. The independence number α is the cardinality of the largest independent set, and thus we consider α/n (MIS ratio) for cross-instance comparisons.

TABLE I: Graph-Theoretic Metrics Description

Metric Name	Description
Primary Structural Predictors (RQ1 & RQ2)	
`edge_density` (ρ)	Global edge density: $2m/(n(n-1))$
`mis_ratio`	Normalized Max Independent Set: $\alpha(G)/n$
Size & Degree Distribution	
`log_n`, `log_m`	Log of vertex/edge count
`deg_min`, `deg_max`	Min/Max vertex degree
`deg_mean`, `deg_stdev`	Mean ($2m/n$) and Std Dev of degrees
`deg_const`	Boolean: graph is regular
`deg_log_[kurt/skew]`	Kurtosis and Abs. Skew of log-degree
Clustering & Spectral Properties	
`clust_[min/max/mean]`	Min, Max, and Mean local clustering
`clust_stdev`, `clust_const`	Std Dev and Uniformity of clustering
`log_norm_ev[1/2]`	Log of 1st and 2nd normalized eigenvalues
`log_ev_ratio`	Log of eigenvalue ratio
Connectivity & Core Structure	
`disconnected`, `mis`	Multiple components (bool), MIS size
`core_[min/max/mean]`	Min, Max (degeneracy), and Mean core number
`core_stdev`, `core_const`	Std Dev and Uniformity of core membership
`core_log_[kurt/skew]`	Kurtosis and Abs. Skew of core numbers
Assortativity & Neighbors	
`assortativity`	Degree assortativity coefficient
`avg_neighbor_deg_*`	Stats for average neighbor degrees
`avg_deg_conn_*`	Stats for degree connectivity
Weight Distribution	
`weight_[min/max/mean]`	Min, Max, and Mean edge weight
`percent_pos`	Fraction of positive edge weights

C. Statistical Methodology

Quantitative analysis of the relationship between graph-theoretic features and ME performance measures was conducted using MATLAB 2024a. We employed Pearson correlation coefficients (r) and associated p-values to evaluate the linear dependencies between source metrics (e.g., edge density, MIS ratio) and outcome measures (success rate, ACL, MCL, and embedding time) across the Chimera, Pegasus, and Zephyr architectures. To maintain the distinction between embedding feasibility and routing quality, ACL, MCL, and computational time were computed exclusively for instances yielding a non-zero success rate.

[2]D-Wave announced its Advantage2 Quantum Computer in May 2025: https://support.dwavesys.com/hc/en-us/articles/32105885880087

[3]The source code is publicly available at our GitHub repository: https://github.com/notajourney/MinorEmbeddingCEC2026.

Double Dissociation via Model Selection: To test the hypothesis that distinct graph-theoretic features predict different aspects of ME outcomes, we used a double dissociation design with AIC-based model comparison. Two graph-theoretic metrics served as competing predictors: MIS ratio and edge density. ME outcomes were then categorized into two domains: Feasibility via success rate versus Cost via ACL. For each outcome and architecture, we fit four linear regression models with respect to the metrics α/n and ρ:

$$M_0 : y = \beta_0 + \varepsilon \quad \text{(null)}$$
$$M_1 : y = \beta_0 + \beta_1 \cdot \alpha/n + \varepsilon \quad (\alpha/n \text{ only})$$
$$M_2 : y = \beta_0 + \beta_1 \cdot \rho + \varepsilon \quad (\rho \text{ only})$$
$$M_3 : y = \beta_0 + \beta_1 \cdot \alpha/n + \beta_2 \cdot \rho + \varepsilon \quad \text{(full)}$$

Model comparison used the Akaike Information Criterion (AIC), AIC $= -2\ln(\hat{L}) + 2k$, where $\hat{L}$ is the maximized likelihood and k is the number of parameters. Lower AIC indicates a better balance of fit and model parsimony. We computed ΔAIC $=$ AIC(M_2) $-$ AIC(M_1), where positive values favor MIS and negative values favor density. Following [23], $|\Delta\text{AIC}| > 10$ constitutes very strong evidence. The double dissociation hypothesis is validated if the following holds:

1) Feasibility: ΔAIC > 0 (α/n preferred)
2) Cost: ΔAIC < 0 (ρ preferred)

Statistical Robustness: Modeling the feasibility outcome—a bounded proportion from $R = 10$ trials—via Ordinary Least Squares (OLS) technically violates the homoscedasticity and continuous-support assumptions of Gaussian regression. However, we employ OLS uniformly across both feasibility and cost metrics to establish a structurally identical baseline for ΔAIC comparisons. To verify that our findings are robust to this likelihood approximation, we refitted the feasibility models using binomial-logistic regression (logit link, $R = 10$). Across all three architectures, the ΔAIC sign favoring the MIS ratio is strictly preserved, with binomial magnitudes substantially exceeding the OLS values (Chimera: $+2414.6$ vs. $+118.3$; Pegasus: $+2979.8$ vs. $+96.7$; Zephyr: $+2925.1$ vs. $+96.5$). This confirms the observed double dissociation is a genuine structural property, not an artifact of the Gaussian assumption.

IV. Results

We systematically executed MINORMINER across 13 graph blocks and used the aforementioned statistical tests to answer our two primary questions: (RQ1) Does edge density (ρ) reliably predict ME performance? (RQ2) Which other structural properties affect the graph embeddability?

A. Validating the Edge Density Hypothesis

We calculated Pearson correlation coefficients between the graph-theoretic metrics and the resulting outcome measures. Figure 1 presents these correlations for the four blocks of graphs with $n = 100$ nodes (comprising 160 instances and 1600 individual ME runs). This specific subset was chosen because it features a full spectrum of ρ values and maintains a 100% success rate across all three architectures.

We consistently observe that global edge density ρ exhibits strong correlations with ME performance measures, provided all instances embed successfully.[4] However, these correlations weaken significantly when analyzing more challenging subsets with lower success rates, such as instances with $n = 200$ nodes. In these regimes, instances with similar ρ values yield highly dispersed ACL values, suggesting a decoupling of density from embedding cost. At the same time, the predictive power of other metrics increases, most notably the MIS ratio (α/n) and related structural independence metrics.

B. Graph-Theoretic Metrics: Correlation Heatmaps

To identify the primary structural drivers of embedding difficulty across the entire problem set ($N = 548$ graph instances altogether), we conducted a systematic exploratory analysis across the full feature space of Table I. Features were ranked by their mean absolute Pearson correlation coefficient ($|\bar{r}|$) specifically with ACL, our primary proxy for embedding quality. Figure 2 presents the resulting correlation heatmap computed over the consolidated dataset across all architecture-target pairs. A distinct hierarchy emerges: the top-ranked features are dominated by MIS-related metrics, followed by degree distribution statistics and core structure metrics; ρ no longer acts as the dominant predictor in this aggregate view. Given the observed impact of the MIS metric, we hypothesize that MIS ratio (α/n) predicts embedding *feasibility* (success rate), while edge density (ρ) predicts embedding *cost* (ACL).

C. Double Dissociation: Predicting Feasibility versus Cost

We tested whether graph-theoretic features differentially predict ME outcomes using AIC-based model selection across three D-Wave architectures (Chimera, Pegasus, Zephyr; $N = 548$ graph instances each). Given our hypothesis that α/n predicts ME success rate, while ρ predicts ME ACL, we conducted the Double Dissociation analysis described in Section III-C. Table II presents this analysis' outcome, where the double dissociation hypothesis is confirmed in all 6/6 tests:

- **Feasibility:** α/n explains 17–19% of success rate variance; ρ explains $< 2\%$. All ΔAIC $> +96$.
- **Cost:** ρ explains 53–78% of ACL variance; α/n explains $\sim$24%. All ΔAIC < -184.

Our results suggest a mechanistic decoupling of the ME process into two distinct phases: a 'gatekeeper' and a 'toll'. The MIS ratio (α/n) acts as the *gatekeeper*: it determines the binary feasibility of the embedding; if a graph lacks sufficient structural independence ('slack'), it simply cannot fit onto the hardware graph, regardless of density. Once past the gate, edge density (ρ) determines the *toll*. For successfully embedded graphs, density dictates the cost of the solution (ACL). The *double dissociation* confirms that overcoming the 'gatekeeper' requires different structural properties than minimizing the 'toll'.

[4] Correlation plots for the remaining blocks are publicly available at our repository: https://github.com/notajourney/MinorEmbeddingCEC2026.

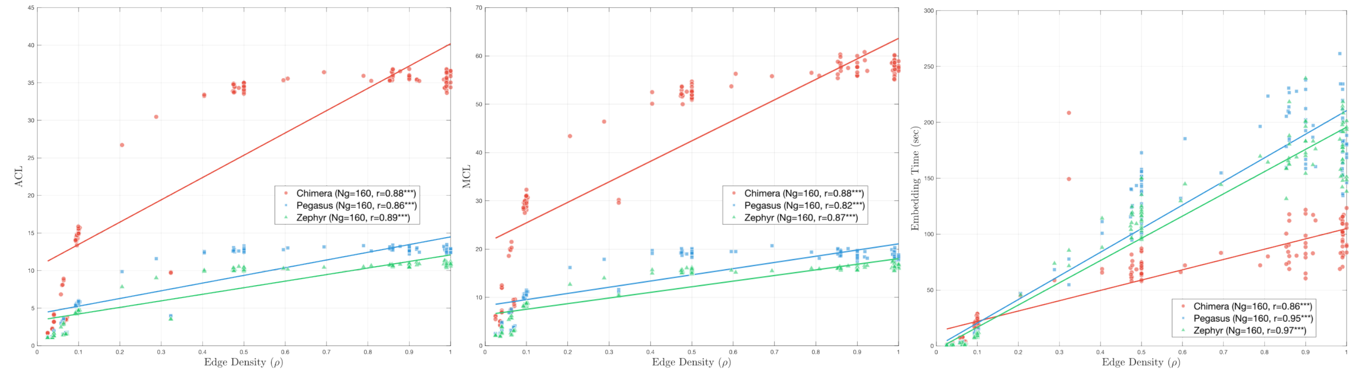

Fig. 1: Analyses of the ME performance measures over graphs with $n = 100$ nodes (160 instances): Pearson correlation tests against Edge Density (ρ) for ACL (left), MCL (middle) and Embedding Time (right), denoting r values in the legend boxes.

TABLE II: AIC-based *double dissociation analysis* (III-C).

Arch.	**Outcome**	N	$R^2_{\alpha/n}$	R^2_{ρ}	**ΔAIC**	**Winner**	**Hyp?**
Feasibility (expected: α/n)							
Chimera	success	548	**.194**	.000	+118.3	α/n	✓
Pegasus	success	548	**.173**	.013	+96.7	α/n	✓
Zephyr	success	548	**.173**	.014	+96.5	α/n	✓
Cost (expected: ρ)							
Chimera	ACL	335	.236	**.784**	−422.5	ρ	✓
Pegasus	ACL	374	.239	**.534**	−184.0	ρ	✓
Zephyr	ACL	378	.246	**.540**	−186.8	ρ	✓

V. Conclusion

In this work, we present the ME problem as a standalone CO challenge and provide a rigorous formulation of it. Using Max-Cut as our primary benchmark, we investigate the empirical behavior of the MINORMINER algorithm across diverse hardware topologies. Through an experimental study of 500+ graph instances, we recorded the outcomes of repeated heuristic deployments. While we corroborate the common wisdom regarding edge density (ρ) as a primary predictor for embedding cost, we identify a *functional dissociation* between feasibility and resource consumption. We revealed that while both metrics quantify graph compactness, they represent distinct structural dimensions: ρ measures global connectivity demand, which correlates with increased qubit consumption, whereas the Maximum Independent Set ratio (α/n) measures structural ‘slack’. While ρ is a mean-field property of total edge count, α/n is sensitive to the local distribution of edges, distinguishing the global pressure of satisfying connectivity from the local availability of independent placement sites.

We stress that ME is a purely graph-theoretic problem: given a source graph G and a target hardware graph H, the task is to find a subgraph of H that is contractible to G, irrespective of the combinatorial origin of G. While our benchmark instances originate from Max-Cut problem libraries, the embedding algorithm receives only the abstract graph as input and is agnostic to whether G encodes a penalty-free Max-Cut instance, a constrained QUBO with auxiliary variables, or a randomly generated topology. Consequently, our structural predictors—MIS ratio for embedding feasibility, edge density for chain length—characterize the graph-theoretic difficulty of embedding *per se*, not properties specific to the Max-Cut problem class. The scope limitation of this study thus lies in benchmark coverage rather than in methodology: the Max-Cut-derived instances sample a particular region of graph-structure space that may underrepresent motifs typical of penalized QUBO formulations, such as dense penalty-induced cliques or auxiliary-variable substructures. Extending the analysis to such structural regimes is a natural direction for future work. Moreover, building on the functional dissociation identified in this study, another direction for future research is condensing the high-dimensional graph-metric space into a manageable, interpretable ontology to enhance predictive accuracy across diverse problem classes. This involves a multi-stage framework—comprising latent structure estimation via eigendecomposition to expose dominant co-variation axes, deduplication of redundant features, and semantic family clustering—to identify a compact set of *medoid* representatives that capture the latent drivers of embedding difficulty. By mapping embedding outcomes such as success rate and resource cost against these standardized metric families, we aim to develop a robust, architecture-agnostic modeling framework. This systematic approach will allow for more nuanced partial and residual analyses, ultimately providing practitioners with a principled method for predicting QA performance based on the intrinsic topological ‘fingerprint’ of a problem.

Acknowledgments

R.L. and O.M.S. gratefully acknowledge support from the Israeli Ministry of Science and Technology (MOST). A.M. gratefully acknowledges support from the UK’s International Science Partnerships Fund (ISPF) UK–Israel Research Collaboration programme. The authors used Claude (Anthropic) for text editing assistance.

References

[1] H. Munoz-Bauza and D. Lidar, “Scaling advantage in approximate optimization with quantum annealing,” *Phys. Rev. Lett.*, vol. 134, p. 160601, Apr 2025.

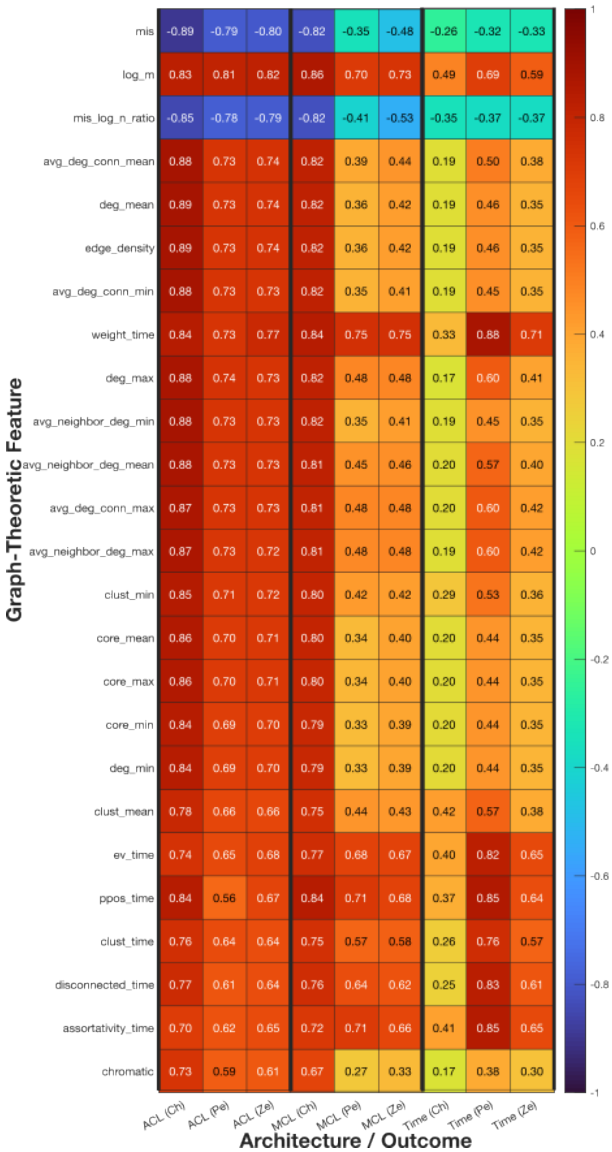

Fig. 2: Predictive Hierarchy of Graph-Theoretic Metrics. Features are ranked by their mean absolute correlation with ACL. The first three columns – ACL for Chimera (Ch), Pegasus (Pe), and Zephyr (Ze) – represent the primary measure of embedding quality. Subsequent columns show secondary resource metrics: MCL and Embedding Time. Pearson coefficients r range from -1 (blue) to $+1$ (red). The strong negative correlation of MIS-related metrics is evident.

[2] V. Choi, "Minor-embedding in adiabatic quantum computation: I. The parameter setting problem," *Quantum Information Processing*, vol. 7, no. 5, pp. 193–209, 2008. [Online]. Available: https://doi.org/10.1007/s11128-008-0082-9

[3] A. Gomez-Tejedor and E. Osaba, "Addressing the minor-embedding problem in quantum annealing and evaluating state-of-the-art algorithm performance," *arXiv preprint*, 2025, arXiv:2504.13376 [quant-ph].

[4] A. Lucas, "Ising formulations of many NP problems," *Frontiers in Physics*, vol. 2, p. 5, 2014. [Online]. Available: https://dx.doi.org/10.3389/fphy.2014.00005

[5] G. Kochenberger, J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, and Y. Wang, "The unconstrained binary quadratic programming problem: A survey," *Journal of Combinatorial Optimization*, vol. 28, no. 1, pp. 58–81, 2014. [Online]. Available: https://dx.doi.org/10.1007/s10878-014-9734-0

[6] N. Robertson and P. D. Seymour, "Graph minors. XIII. The disjoint paths problem," *Journal of Combinatorial Theory, Series B*, vol. 63, no. 1, pp. 65–110, 1995. [Online]. Available: https://dx.doi.org/10.1006/jctb.1995.1006

[7] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*, ser. Dover Books on Computer Science. Mineola, NY, USA: Dover Publications, 1998.

[8] D. E. Bernal, K. E. C. Booth, R. Dridi, H. Alghassi, S. Tayur, and D. Venturelli, "Integer programming techniques for minor-embedding in quantum annealers," in *Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR 2020)*. Springer International Publishing, 2020, pp. 112–129. [Online]. Available: https://doi.org/10.1007/978-3-030-58942-4_8

[9] J. Cai, W. G. Macready, and A. Roy, "A practical heuristic for finding graph minors," *arXiv preprint*, 2014, arXiv:1406.2741 [quant-ph].

[10] T. D. Goodrich, B. D. Sullivan, and T. S. Humble, "Optimizing adiabatic quantum program compilation using a graph-theoretic framework," *Quantum Information Processing*, vol. 17, no. 5, p. 1–26, May 2018. [Online]. Available: https://doi.org/10.1007/s11128-018-1863-4

[11] Y. Sugie, Y. Yoshida, N. Mertig, T. Takemoto, H. Teramoto, A. Nakamura, I. Takigawa, S.-I. Minato, M. Yamaoka, and T. Komatsuzaki, "Graph minors from simulated annealing for annealing machines with sparse connectivity," in *Proceedings of TPNC 2018*. Berlin, Heidelberg: Springer-Verlag, 2018, pp. 111–123.

[12] J. P. Pinilla and S. J. E. Wilton, "Layout-aware embedding for quantum annealing processors," in *High Performance Computing*, M. Weiland, G. Juckeland, C. Trinitis, and P. Sadayappan, Eds. Cham: Springer International Publishing, 2019, pp. 121–139. [Online]. Available: https://doi.org/10.1007/978-3-030-20656-7_7

[13] S. Zbinden, A. Bärtschi, H. Djidjev, and S. Eidenbenz, "Embedding algorithms for quantum annealers with chimera and pegasus connection topologies," in *High Performance Computing*, P. Sadayappan, B. L. Chamberlain, G. Juckeland, and H. Ltaief, Eds. Cham: Springer International Publishing, 2020, pp. 187–206. [Online]. Available: https://doi.org/10.1007/978-3-030-50743-5_10

[14] F. Glover, G. Kochenberger, R. Hennig, and Y. Du, "Quantum bridge analytics I: A tutorial on formulating and using QUBO models," *Annals of Operations Research*, vol. 314, no. 1, pp. 141–183, 2022. [Online]. Available: https://dx.doi.org/10.1007/s10479-022-04634-2

[15] M. Anthony, E. Boros, Y. Crama, and A. Gruber, "Quadratic reformulations of nonlinear binary optimization problems," *Mathematical Programming*, vol. 162, no. 1–2, pp. 115–144, 2017. [Online]. Available: https://doi.org/10.1007/s10107-016-1032-4

[16] E. Boros and P. L. Hammer, "Pseudo-Boolean optimization," *Discrete Applied Mathematics*, vol. 123, no. 1–3, pp. 155–225, 2002. [Online]. Available: https://dx.doi.org/10.1016/S0166-218X(01)00341-9

[17] E. Boros and A. Gruber, "On quadratization of pseudo-Boolean functions," *arXiv preprint arXiv:1404.6538*, 2014.

[18] I. Dunning, S. Gupta, and J. Silberholz, "What works best when? a systematic evaluation of heuristics for max-cut and QUBO," *INFORMS Journal on Computing*, vol. 30, no. 3, 2018. [Online]. Available: https://doi.org/10.1287/ijoc.2017.0798

[19] D-Wave Systems Inc., "The D-Wave Advantage System: An Overview," D-Wave Systems Inc., Tech. Rep. 14-1049A-A, 2020. [Online]. Available: https://www.dwavequantum.com/media/s3qbjp3s/14-1049a-a_the_d-wave_advantage_system_an_overview.pdf

[20] ——, "Next-Generation Topology of D-Wave Quantum Processors," D-Wave Systems Inc., Tech. Rep. 14-1026A-C, 2019. [Online]. Available: https://www.dwavequantum.com/media/jwwj5z3z/14-1026a-c_next-generation-topology-of-dw-quantum-processors.pdf

[21] K. Boothby, A. D. King, and J. Raymond, "Zephyr topology of d-wave quantum processors," D-Wave Systems Inc., Tech. Rep. 14-1056A-A, 2021. [Online]. Available: https://www.dwavequantum.com/media/2uznec4s/14-1056a-a_zephyr_topology_of_d-wave_quantum_processors.pdf

[22] D-Wave Systems Inc., "What is the chimera topology?" D-Wave Support Article, 2024. [Online]. Available: https://support.dwavesys.com/hc/en-us/articles/360003695354-What-Is-the-Chimera-Topology

[23] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer, 2002.