

Quantum Multiple Rotation Averaging

Supplementary Material

This appendix supplements the main paper with expanded content, starting with a detailed background on MRA in Sec. A. Explanation of the Hamiltonian evolution during quantum annealing is provided in Sec. B. Sec. C explains the adiabatic theorem. Sec. D and F provide details of the propositions in the main text. The effect of penalization is explained in Sec. E. Sec. G details the transformation of the QUBO problem into the Ising form, a Hamiltonian representation more naturally aligned with the physical interactions implemented in quantum annealers. In Sec. H, we explain how Prop. 1 can be combined with [24], leading to an alternative direct MRA approach. Sec. I details the posterior protocol for improving sample quality. We introduce details of integrating IQARS, as a module for averaging multiple rotations, in a SfM pipeline; see Sec. J. In addition to synchronizing fully-connected camera rotations, Sec. K provides evaluations on sparsely-connected camera rotations. Visualizations of the logical coupling matrix and its embeddings on the hardware are provided in Sec. L and M. An additional comparative analysis between IQARS and a few other prominent local solvers is provided in Sec. N. Evaluations of IQARS with different initializations are provided in Sec. O.

A. Review: MRA

This section reviews the basic definition of MRA. Let $\{R_1, \dots, R_N\} \subset \text{SO}(3)$ represent a set of absolute 3D rotational matrices w.r.t. a global reference frame. The special orthogonal group $\text{SO}(3)$ is the set of matrix elements R that are orthogonal and have a unit determinant:

$$\text{SO}(3) := \{R \in \mathbb{R}^{3 \times 3}, RR^T = I, \det(R) = 1\}. \quad (23)$$

Given the set of absolute rotations $\{R_1, \dots, R_N\}$, the relative rotation R_{ij} between any two cameras i and j , where $i, j \in \{1, \dots, N\}$, can be computed as

$$R_{ij} = R_j R_i^{-1} \text{ or, equivalently, } R_{ij} R_i = R_j. \quad (24)$$

Assume we are given a set of noisy relative rotations $\{\tilde{R}_{ij} \mid i, j = 1, \dots, N\}$ between different input camera pairs i, j , the objective of MRA is to find the absolute rotations $\{R_i \mid i = 1, \dots, N\}$ that best explains the observed \tilde{R}_{ij} . MRA is an inverse optimization problem to minimize the discrepancy between the predicted and observed relative rotations:

$$\min_{R_1, \dots, R_N \in \text{SO}(3)} \sum_{(i,j)} \text{dist}(\tilde{R}_{ij} R_i, R_j), \quad (25)$$

where “ $\text{dist}(a, b)$ ” is a chosen distance metric that quantifies the difference between two rotations. With the widely used choice of squared chordal distance, i.e., $\text{dist}(a, b) = \|a - b\|_F^2$, we can re-formulate Eq. (25) as

$$\min_{R_1, \dots, R_N \in \text{SO}(3)} \sum_{(i,j)} \|\tilde{R}_{ij} R_i - R_j\|_F^2. \quad (26)$$

Eq. (26) provides the starting MRA formulation, which we map to a quantum-compatible form in the main text.

B. Review: Quantum System Evolution during Annealing

The evolution of quantum states during the Hamiltonian transition in the annealer is governed by the Schrödinger’s equation

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H(t) |\psi(t)\rangle, \quad (27)$$

where i is the imaginary unit, \hbar denotes the reduced Planck constant, $|\psi(t)\rangle$ represents the quantum state at time t , and $H(t)$ is the time-dependent Hamiltonian. Its analytical solution can be expressed as

$$|\psi(t)\rangle = \hat{T} \exp\left(-\frac{i}{\hbar} \int_0^t H(t') dt'\right) |\psi(0)\rangle, \quad (28)$$

where \hat{T} is the time ordering operator, ensuring chronological sequence in the exponential expansion. By substituting in the time-dependent Hamiltonian; see Eq. 1, we can, therefore, describe the evolution of quantum states during the annealing process as follows:

$$\begin{aligned} |\psi(t)\rangle = \hat{T} \exp \left(-\frac{i}{\hbar} \int_0^s \left[-A(s') \sum_i \sigma_i^x \right. \right. \\ \left. \left. + B(s') \left(\sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z + \sum_i h_i \sigma_i^z \right) \right] ds' \right) |\psi(0)\rangle. \end{aligned} \quad (29)$$

$A(s')$ and $B(s')$ are annealing schedules. σ_i^x and σ_i^z are Pauli spin operators. J_{ij} denotes coupling strengths, and h_i represents local magnetic fields. These collectively define the system’s quantum dynamics.

C. Review: Adiabatic Theorem

Adiabatic quantum computation (AQC) represents a specialized framework within quantum computing that leverages the adiabatic theorem [1, 8, 28, 30]. Theorem 1 formalizes the key ideas of this principle:

Theorem 1. Adiabatic Theorem: *A physical system remains in its instantaneous eigenstate if a given perturbation is acting on it slowly enough and if there is a gap between the initial eigenvalue and the rest of the Hamiltonian's spectrum.*

D. Proof of Proposition 1

We provide proof of the Prop. 1 given in the main text:

$$\begin{aligned}
& \min_{R_1, \dots, R_N} \sum_{(i,j)} \|\tilde{R}_{ij} R_i - R_j\|_F^2 \\
&= \min_{R_1, \dots, R_N} \sum_{(i,j)} \|\tilde{R}_{ij}\|_F^2 \|R_i\|_F^2 + \|R_j\|_F^2 - 2\langle R_{ij} R_i, R_j \rangle_F \\
&= \min_{R_1, \dots, R_N} \sum_{(i,j)} 3\|R_i\|_F^2 + \|R_j\|_F^2 - 2\text{Tr}(R_i^\top R_{ij}^\top R_j) \\
&= \min_{R_1, \dots, R_N} \sum_{(i,j)} 3\|R_i\|_F^2 + \|R_j\|_F^2 - 2\text{Tr}(R_j R_i^\top R_{ij}^\top) \\
&= \min_{R_1, \dots, R_N} \sum_{(i,j)} 3\|R_i\|_F^2 + \|R_j\|_F^2 - \\
& \quad 2\text{vec}(R_i)^\top (I \otimes R_{ij}^\top) \text{vec}(R_j). \tag{30}
\end{aligned}$$

Given the orthonormal nature of rotation matrices $R_i \in \text{SO}(3)$, they satisfy the fundamental properties $R_i^\top = R_i^{-1}$ and $\det(R_i) = 1$. These constraints imply that the Frobenius norm $\|R_i\|_F$ is always $\sqrt{3}$ for any rotation matrix and can be eliminated from the optimization objective. We reformulate the quadratic alignment term $\sum_{(i,j)} -2\text{vec}(R_i)^\top (I \otimes R_{ij}^\top) \text{vec}(R_j)$ in compact matrix form by eliminating the explicit summation through Kronecker product identities and rewrite it as

$$-2 [\text{vec}(R_1)^\top \dots \text{vec}(R_N)^\top] P \begin{bmatrix} \text{vec}(R_1) \\ \vdots \\ \text{vec}(R_N) \end{bmatrix}, \tag{31}$$

with $\text{vec}(R_i)^\top \in \mathbb{R}^{1 \times 9}$. The corresponding P is

$$P = \begin{bmatrix} I \otimes R_{11}^\top & \dots & I \otimes R_{1N}^\top \\ \vdots & \ddots & \vdots \\ I \otimes R_{N1}^\top & \dots & I \otimes R_{NN}^\top \end{bmatrix} \in \mathbb{R}^{9N \times 9N}. \tag{32}$$

E. Penalization Effect

While each iterative update of IQARS preserves $R_i \in \text{SO}(3)$ through exponential map parameterization, the local search window—defined by first-order linearization of the tangent space—permits bounded exploration of points outside $\text{SO}(3)$. Within this region, solutions with larger $\|R_i\|_F$ are favored due to the negative coupling term (see Eq. 7) in the objective function. This exacerbates the tendency of solutions to drift away from the $\text{SO}(3)$ manifold constraints.

We employ $\|R_i\|_F^2$ as a principled metric for quantifying geometric deviation from $\text{SO}(3)$ during local search, with an appropriately chosen regularization strength λ counteracting this tendency within the bounded exploration region.

F. Proof of Proposition 3

We provide proof of the Prop. 3 as given in the main text:

$$\begin{aligned}
& \arg \min_{\|\Delta \mathbf{v}\|_\infty \leq \delta^k} \Delta \mathbf{v}^\top \hat{Q} \Delta \mathbf{v} + \hat{\mathbf{c}}^\top \Delta \mathbf{v} \\
&= \arg \min_{\mathbf{q} \in \mathbb{B}^{3Nm}} (-\delta^k \mathbf{1}_{3N} + D\mathbf{q})^\top \hat{Q} (-\delta^k \mathbf{1}_{3N} + D\mathbf{q}) \\
& \quad + \hat{\mathbf{c}}^\top (-\delta^k \mathbf{1}_{3N} + D\mathbf{q}) \\
&= \arg \min_{\mathbf{q} \in \mathbb{B}^{3Nm}} \mathbf{q}^\top D^\top \hat{Q} D\mathbf{q} + \mathbf{q}^\top D^\top \hat{Q} (-\delta^k \mathbf{1}_{3N}) \\
& \quad + (-\delta^k \mathbf{1}_{3N})^\top \hat{Q} D\mathbf{q} + \hat{\mathbf{c}}^\top (D\mathbf{q}) \\
&= \arg \min_{\mathbf{q} \in \mathbb{B}^{3Nm}} \mathbf{q}^\top D^\top \hat{Q} D\mathbf{q} + \mathbf{q}^\top D^\top (\hat{\mathbf{c}} - 2\delta^k \hat{Q} \mathbf{1}_{3N}). \tag{33}
\end{aligned}$$

G. Transformation of QUBO Problems to Ising-Compatible Quadratic Forms

Given the general QUBO formulation

$$\min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^\top Q \mathbf{x} + \mathbf{b}^\top \mathbf{x}, \tag{34}$$

we require a transformation to the standard quadratic form for implementation on quantum annealers employing Ising Hamiltonians. Leveraging the Boolean constraint $\mathbf{x} \in \{0,1\}^n$ (which implies $x_i^2 = x_i$), we derive the equivalent pure quadratic form through the following sequence:

$$\begin{aligned}
\min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^\top Q \mathbf{x} + \mathbf{b}^\top \mathbf{x} &= \min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^\top Q \mathbf{x} + \mathbf{x}^\top \text{diag}(\mathbf{b}) \mathbf{x} \\
&= \min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^\top (Q + \text{diag}(\mathbf{b})) \mathbf{x} \\
&\triangleq \min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^\top Q' \mathbf{x}, \tag{35}
\end{aligned}$$

where the transformed matrix $Q' \in \mathbb{R}^{n \times n}$ is constructed as

$$Q' = Q + \text{diag}(\mathbf{b}). \tag{36}$$

This exact reformulation: (i) preserves the original optimization problem's solution space, (ii) maintains compatibility with physical Ising model implementations through the identity $x_i = x_i^2$, and (iii) enables efficient embedding on quantum annealers by consolidating all terms into quadratic couplings.

H. Alternative Approach for MRA

We also realized that in combination with Prop. 1, the approach proposed in [24] can be an alternative for MRA.

Specifically, they propose to directly approximate rotation matrices $R_i \in \text{SO}(3)$ through a linear combination of binary-activated basis matrices:

$$R_i \approx \sum_{\ell=0}^{m-1} q_{i,\ell} Q_\ell, \quad q_{i,\ell} \in \{0, 1\}, \quad (37)$$

where the basis matrices $\{Q_\ell\}$ quantize the components of the Rodrigues rotation formula; see Eq. (8). Specifically, the basis set comprises scaled versions of the identity matrix and generators of $\mathfrak{so}(3)$:

$$Q_\ell \in \{wC \mid w \in \{0.5, 0.2, 0.1, 0.1, 0.05\}, \\ C \in \{I, -I\} \cup \{M_k, -M_k\}_{k=1}^6\}, \quad (38)$$

with the generator matrices defined as

$$M_1 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}, M_2 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}, M_3 = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \\ M_4 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, M_5 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, M_6 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}. \quad (39)$$

Substituting the approximation from Eq. (37) into the matrixized problem (see Eq. (6) in Prop. 1) yields a QUBO formulation that directly approximates the MRA problem. However, this approach presents two significant theoretical limitations: (1) the finite basis set $\{Q_\ell\}$ imposes fundamental approximation bounds due to its limited expressiveness, and (2) the resulting matrices are not guaranteed to satisfy the orthonormality conditions ($R_i^\top R_i = I$) or the determinant constraint ($\det(R_i) = 1$) required for proper $\text{SO}(3)$ membership, as these nonlinear constraints are not explicitly enforced in the binary formulation. In our IQARS, such constraints are resolved.

I. Posterior Analysis Protocol

Algorithmic details of performing posterior analysis to refine solution qualities are summarized in Alg. 2.

J. Application: IQARS within SfM

SfM is a fundamental technique that reconstructs 3D scene geometry and camera poses from a collection of 2D images. Within this framework, MRA serves as a critical synchronization step that enforces global consistency among relative orientation estimates. We implement a hybrid quantum-classical pipeline by replacing the conventional MRA module in Glomap [46]—a state-of-the-art SfM system—with our IQARS. The complete flow is described in Fig. 8 This hybrid pipeline maintains compatibility with conventional feature matching and downstream applications. Quantitative visualization of the reconstructed 3D scenes is provided in Sec. 5.

Algorithm 2 Posterior Refinement Protocol with Low-energy Binary Solution Space

- 1: **Input:** \mathbf{x} : Low-energy binary strings from annealers
 $E(\mathbf{x})$: Energies of prepared binary strings
 β : Inverse temperature
 L : Length of each binary string
 M : Number of binary strings
- 2: Construct an empty array $\mathbf{S} = [0, \dots, 0]$ of size L
- 3: Construct an empty array $\mathbf{m} = [0, \dots, 0]$ of size L
- 4: Calibrate energy spectrum to a standardized range
- 5: **For** $i = 1$ to M
- 6: Calculate Boltzmann weighting factor $P(\mathbf{x}_i)$ according to Eq. (22)
- 7: **For** $j = 1$ to L
- 8: **If** $\mathbf{x}_i[j] = 1$
- 9: $S[j] \leftarrow S[j] + P(\mathbf{x}_i)$
- 10: **If** $\mathbf{x}_i[j] = 0$
- 11: $S[j] \leftarrow S[j] - P(\mathbf{x}_i)$
- 12: **End**
- 13: **End**
- 14: **For** $k = 1$ to L
- 15: **If** $\mathbf{S}[j] > 0$
- 16: $\mathbf{m}[j] = 1$
- 17: **If** $\mathbf{S}[j] < 0$
- 18: $\mathbf{m}[j] = 0$
- 19: **return** \mathbf{m}, \mathbf{S}

K. MRA for Sparse Problems

Building upon our theoretical framework and experimental validation for complete camera graphs $G = (V, E)$ where vertices $v_i \in V$ correspond to absolute camera orientations $R_i \in \text{SO}(3)$, we generalize our analysis to sparse graph topologies characterized by incomplete edge sets

$E' \subset E$. The underlying mathematical formulation maintains consistency, where missing relative rotations \hat{R}_{ij} are treated as null observations and treated as zero terms during optimization. To preserve the minimum connectivity requirement, we ensure that each node maintains at least one outgoing edge. We leverage the same synthetic noisy dataset as in main experiments; see Sec. 5, and perform experiments by randomly switching off a certain percentage of relative observations. We progressively increase the sparsity of observations and visualize the results in Fig. 9.




Figure 9. Benchmark results of IQARS against other solvers for MRA on synthetic dataset with increasing camera graph sparsity.

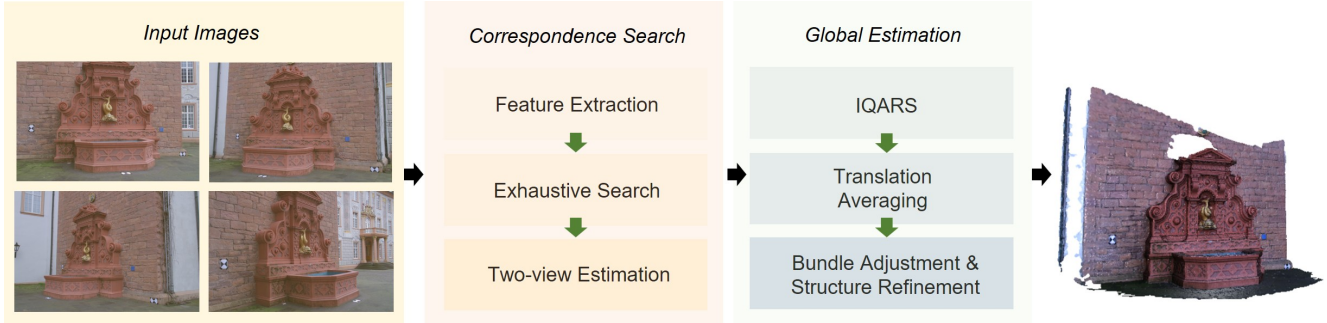


Figure 8. Integration of our *IQARS* protocol within a traditional SfM pipeline. Our *IQARS* replaces the traditional MRA component while maintaining compatibility with the existing pipeline.




Figure 10. Qualitative coupling matrix sparsity pattern visualization with increasing problem scale N for MRA.

Results confirm that our algorithm maintains robustness across different sparsity ratios.

L. QUBO Coupling Sparsity Visualization

We visualized the sparsity of the coupling matrix in the QUBO formulation before it gets uploaded to the annealer for execution. The sparsity of the problem formulation determines how logical qubits—which represent the MRA solution update in *IQARS*—should interact with each other during the annealing process. It also concerns the problem’s embeddability on quantum hardware. Typically, problems with a certain sparsity can admit efficient hardware mappings that reduce physical resource requirements and operational overhead. We visualized the sparsity pattern of MRA that is employed in our algorithm in Fig. 10 with increasing problem sizes. This can provide valuable guidelines for hardware embedding analysis.

M. Visualizations of QPU (Pegasus Topology) Embedding

For completeness, we provide a qualitative visualization of the problem embeddings on the annealer of Pegasus topology; see Fig. 11. Specifically, Fig. 11-(A) visualizes the hardware topology of the annealer. Fig. 11-(B) to (D) visualizes the embedded MRA problems under the specific hardware topology. The process of finding such problem embedding is called *minor embedding*. Compared to

Noise Levels	$\pi/10$	$\pi/5$	$\pi/3$	$\pi/2$
LM	0.2209	0.4593	0.7443	1.1179
Trust-Region	0.2341	0.4793	0.7839	1.180
Ours	0.1929	0.3932	0.6388	0.9235

Table 3. Performance comparison with local solvers on synthetic noisy dataset across different noise levels σ .

the previous-generation Chimera architectures, the Pegasus topology enables more efficient embeddings of complex optimization problems and supports solving problems of larger scales.

N. Complementary Comparison with Classical Local Solvers

While both Levenberg-Marquardt (LM) and trust-region methods leverage iterative frameworks analogous to *IQARS*, their implementations fundamentally differ in handling the $SO(3)$ manifold’s non-convexity. We give further explanations as a supplement. LM approximates second-order behavior through a damped Gauss-Newton scheme with Hessian approximation, while trust-region methods solve constrained quadratic subproblems within a radii. Crucially, both approaches consistently convexify the $SO(3)$ Riemannian geometry through Euclidean projections and can generate artificial critical points. In contrast, *IQARS* preserves geometric integrity without convexification through: (1) exact $SO(3)$ constraints via Prop. 1’s exponential maps, (2) adaptive trust-region optimization and orthogonality-promoting penalty, and (3) quantum annealing’s Hamiltonian evolution enabling efficient exploration of a non-convex landscape via tunneling. We perform additional quantitative empirical comparison with the mentioned solvers on noisy datasets with $N = 20$ and record the results in Tab. 3; the results confirm that *IQARS* maintains a clear advantage and further validates the theoretical claims.




Figure 11. (A) Pegasus topology structure; (B)–(D) Embedding visualization on QPUs (Pegasus topology) for MRA of problem sizes $N = 3, 5$ and 10 . The corresponding longest chain for a logical qubit during the embedding is highlighted, respectively.

O. Alternative Initialization Strategy

Our IQARS is configured with identity initialization by default to provide a neural and unbiased way for MRA and establish a baseline configuration free of initialization bias. Notably, it also naturally accommodates alternative initialization schemes such as minimum spanning tree (MST) [25] which was originally integrated into L1-IRLS. MST initialization provides a principled approach for generating initial rotation estimates that are geometrically consistent with relative measurements. This especially benefit iterative algorithms such as ours under limited quantum resource budget. We replace the original identity initialization protocol in IQARS with MST, and perform additional evaluations. Other settings and configurations remain consistent as in primary experiments. Quantitative results, as presented in Tab. 4, show that MST initialization under our configuration can indeed lead to significantly improved MRA performance.

P. Quantum vs. Simulated Annealing Performance

Besides quantum annealers, D-Wave also provides a classical simulated annealing (SA) solver. It serves as a classical benchmark for evaluating quantum annealing perfor-

Noise Ratios	$\pi/10$	$\pi/5$	$\pi/3$	$\pi/2$
Ours (Identity init)	0.193	0.393	0.639	0.923
Ours (MST)	0.177	0.352	0.573	0.845

Table 4. IQARS performance under different initialization schemes on synthetic noisy rotations of $N = 20$.

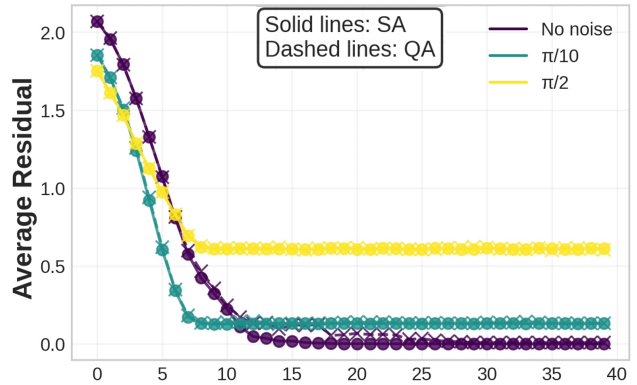


Figure 12. MRA performance comparison between quantum and simulated annealing for dataset across different noise levels.

mance, particularly in assessing the hardware noise effect on the solution quality. SA, inspired by thermodynamic annealing in metallurgy, utilizes thermal fluctuations gov-

erned by the Metropolis-Hastings algorithm to escape local minima with state transitions following the Boltzmann acceptance criterion. It operates classically, requiring $\mathcal{O}(n^2)$ operations per iteration for n -variable QUBO problems due to pairwise interaction calculations. In contrast, while the theoretical time complexity of quantum annealing is governed by the adiabatic theorem, annealing time in practice is by default set to $20\mu s$ on D-Wave machines; see set-up in Sec. 5. We performed an empirical comparison with MRA on dataset across various noise levels, and benchmarked QA against the SA solver; results are visualized in Fig. 12. It is noticeable that SA exhibits marginally better convergence than quantum annealing executed on real noisy machines for noiseless dataset. With increasing noise levels, quantum annealers can achieve performance comparable to classical SA implementations. We also record the clock time statistics. During execution, SA takes around 1.32s while QA, including data transmission overhead, takes 0.38s per iteration in the experiment. Considering SA’s quadratic scaling, it is expected that the computational demand for SA becomes more pronounced for larger problem sizes. These findings suggest that while QA algorithms leveraging practical quantum machines yet may still underperform SA in terms of performance, the inherent time efficiency and scalability suggest a compelling pathway for addressing computationally hard optimization problems in the future.

Q. Additional Visualizations

We provide additional visualizations to complement the main text. Fig. 13 visualizes the physical qubit requirements for embedding IQARS QUBO instances on different generations of D-Wave annealers across different problem sizes N .




Figure 13. Physical hardware resource occupation for logical problems of various sizes.