

Quantum Multiple Rotation Averaging

Shuteng Wang¹ Natacha Kuete Meli² Michael Möller² Vladislav Golyanik¹

¹Max Planck Institute for Informatics, Saarbrücken, Germany

²University of Siegen, Germany

Abstract

Multiple rotation averaging (MRA) is a fundamental optimization problem in 3D vision and robotics that aims to recover globally consistent absolute rotations from noisy relative measurements. Established classical methods, such as L1-IRLS and Shonan, face limitations including local minima susceptibility and reliance on convex relaxations that fail to preserve the exact manifold geometry, leading to reduced accuracy in high-noise scenarios. We introduce IQARS (Iterative Quantum Annealing for Rotation Synchronization), the first algorithm that reformulates MRA as a sequence of local quadratic non-convex sub-problems executable on quantum annealers after binarization, to leverage inherent hardware advantages. IQARS removes convex relaxation dependence and better preserves non-Euclidean rotation manifold geometry while leveraging quantum tunneling and parallelism for efficient solution space exploration. We evaluate IQARS’s performance on synthetic and real-world datasets. While current annealers remain in their nascent phase and only support solving problems of limited scale with constrained performance, we observed that IQARS on D-Wave annealers can already achieve $\approx 12\%$ higher accuracy than Shonan—the best-performing classical method evaluated empirically. Project page: <https://4dqv.mpi-inf.mpg.de/QMRA/>.

1. Introduction

Multiple rotation averaging (MRA) stands as a fundamental group synchronization problem in 3D computer vision, aiming to recover globally consistent absolute rotations from noisy relative measurements while satisfying compositional constraints on the $SO(3)$ manifold [10, 11, 25, 45, 48]. Applications and research fields that can benefit from MRA include structure-from-motion (SfM) [14, 15, 38, 43], simultaneous location and mapping (SLAM) [31, 39, 41, 55], multi-view reconstruction [23, 53], virtual reality [27, 44, 52] and robotics [29, 40], among others. A common approach to averaging multiple rotations involves enforcing




Figure 1. IQARS formulates MRA as QUBO sub-problems executable via quantum annealing to efficiently search for high-quality solutions. Estimated rotations are acquired from QUBO solutions once converged. Annealer’s image is taken from [26].

cycle consistency and distributing errors across all cameras to ensure that the composition of rotations between any two cameras closely aligns with their respective partial measurements, i.e., $R_{ij} = R_j R_i^T$ for all i and j .

MRA presents several unique challenges such as: (1) inherent non-convexity of $SO(3)$ ’s Lie group structure, whose Riemannian geometry creates complex optimization landscapes; (2) the inevitable presence of noise in pairwise measurements, arising from feature matching errors, which disrupt cycle consistency and introduce pathological local minima [51]. These challenges manifest mathematically as an increased susceptibility of gradient-based methods to convergence in sub-optimal local minima. Classical solvers minimize geodesic errors using iterative reweighted least squares [10, 11] like L1-IRLS, a dominating robust approach that efficiently rejects outliers *but can easily get stuck in local minimas*. More recent Shonan method provides certifiable optimality through semidefinite convex relaxation at $O(N^3)$ computational cost *but it suffers from increasing relaxation gap in presence of noise*; N is the number of cameras [17, 21, 50]. Developing accurate methods for MRA, therefore, remains an open challenge.

Quantum annealers are specialized quantum computing devices designed to sample high-quality solutions to NP-

hard quadratic unconstrained binary optimization (QUBO) problems. Leveraging quantum mechanical properties such as quantum parallelism and tunneling, such devices enable efficient exploration of complex energy landscapes compared to classical methods, and facilitate exploration of rugged energy landscapes with multiple low-energy basins. Empirical studies have substantiated their computational capabilities, demonstrating practical advantages of quantum effects in challenging optimization tasks [18, 32–34]. These advantages have further motivated exploring applications of quantum annealing to practical applications; see Sec. 2.

Motivated by recent advances in quantum-enhanced optimization, we identify that MRA can potentially benefit from quantum annealing to circumvent key limitations in classical approaches. In this work, we propose IQARS—a formulation of MRA designed for execution on Ising machines such as quantum annealers, which can globally synchronize multiple rotations with high accuracy; see Fig. 1 for an overview. Practical annealing process is often perturbed by hardware-induced noise such as thermal fluctuations and yields sub-optimal sample solutions in the vicinity of the global optima. We hence propose a posterior refinement framework to refine the solution quality that exploits the observation that, under moderate noise conditions maintaining quasi-equilibrium dynamics, the annealer’s output distribution approximates a Boltzmann distribution [54]. This framework effectively compensates for imperfections while preserving the sampling properties of QAs. In summary, the technical contributions of this work include:

- IQARS, the first end-to-end framework for solving MRA on Ising machines such as quantum annealers;
- Sequential QUBO formulations that strictly adhere to $SO(3)$ manifold geometry during optimization with monotonic residual decrease;

While quantum annealers remain in their nascent phase with limited problem-scale capacity, empirical evaluation of IQARS on D-Wave annealing hardware shows promising performance: $\approx 12\%$ in average lower residual than Shonan—the best-performing classical method evaluated empirically—under noisy measurements. In the MRA setting, our posterior analysis statistically recovers higher-fidelity solutions with 60% probability. Collectively, these advancements position IQARS as the first MRA method leveraging quantum annealing, with distinct characteristics arising from annealing-based optimization.

2. Related Work

Rotation Averaging. MRA became the de facto approach in contemporary 3D reconstruction pipelines for synchronizing locally estimated, noisy relative rotations into globally consistent absolute rotations [10–12, 14, 15, 19, 22, 35]. Chatterjee et al. [10] introduced a pioneering approach: L1-IRLS, an iterative approach for MRA with dynamic residual

weighting for robust outlier rejection. Eriksson et al. [19] proposed a rotation averaging solution through Lagrangian duality and semidefinite relaxation, while Dellaert et al. [17] formulated Shonan’s dimensional lifting approach, solving the problem in $SO(p)$ ($p > 3$) space before projecting to $SO(3)$. This guarantees global convergence under mild noise while maintaining geometric constraints, i.e. orthogonality and unit-determinant, inherent in rotation averaging. Gao et al. [22] presented an incremental estimation workflow for rotation averaging, circumventing the cubic-time complexity of simultaneous absolute rotation estimation. Lee et al. [35] proposed a robust Weiszfeld-based rotation averaging method, enhancing both computational efficiency and convergence reliability. This, however, only works for single-rotation estimation.

Annealing-based Quantum-enhanced Computer Vision (QeCV). Quantum annealing has emerged as a promising approach to address some challenging problems in computer vision [33]. Pioneering works by Golyanik et al. [24] demonstrate its application to 2D/3D point set rotation estimation using binary-weighted basis matrices within a gravitational correspondence framework. Birdal et al. [7] leveraged quantum annealers for permutation synchronization problems solving for cycle-consistent discrete permutations. Farina et al. [20] introduced a quantum-enhanced feature matching via preference-consensus matrices under known model count constraints. Annealing-based approaches have since been adapted to and benefit diverse vision tasks, including shape matching [3, 6], motion segmentation [2], object detection [37], stereo matching [9], and super-resolution [13], collectively demonstrating quantum annealing’s potential, though challenges remain in scaling these methods to practical problem sizes.

Building upon prior work in transformation (mainly rotation) estimation [24, 42], which emphasizes estimating pairwise rotations between point sets, we propose a quantum-annealing-based synchronization paradigm for global multiple rotations. Our IQARS framework inherently preserves the $SO(3)$ manifold structure through the optimization process. This fundamental methodological advancement leads to significantly increased performance, as will be demonstrated empirically.

3. Review: Adiabatic Quantum Computing for Solving QUBO Problems

Adiabatic quantum computing (AQC) formulates combinatorial optimization problems as Hamiltonian ground-state searches and seeks to prepare low-energy states through slow adiabatic evolution; see App. C for details. The protocol involves an adiabatic transition between an initial Hamiltonian H_0 , with a computationally tractable ground state (e.g., $H_0 = -\sum_{i=1}^n X_i$, where X_i is the Pauli- X operator acting on the i -th qubit of an n -qubit quantum

system), and the problem Hamiltonian H_P , which encodes the solution to the target optimization problem. The time-dependent Hamiltonian is given by:

$$H(t) = (1 - s(t)) H_0 + s(t) H_P, \quad s(t) \in [0, 1], \quad (1)$$

where $s(t)$ is a scheduling function that governs the adiabatic evolution, smoothly transitioning the system from H_0 (at $s(t) = 0$) to H_P (at $s(t) = 1$).

QUBO as an AQC-Compatible Problem. QUBO problems, defined as:

$$\min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^\top Q \mathbf{x} + \mathbf{c}^\top \mathbf{x}, \quad (2)$$

where $Q \in \mathbb{R}^{n \times n}$ is a symmetric matrix and $\mathbf{c} \in \mathbb{R}^n$, are naturally suited for AQC due to their quadratic structure, which aligns with the interaction terms in quantum Hamiltonians. To embed a QUBO problem into AQC, each binary variable $x_i \in \{0, 1\}$ is mapped to a qubit, where the classical states 0 and 1 correspond to the eigenstates of the Pauli-Z operator Z_i , with eigenvalues $+1$ and -1 , respectively. This mapping is formalized via the transformation: $x_i \mapsto \frac{1}{2}(I - Z_i)$. Substituting this into the QUBO objective function $f(x)$ yields a quantum Hamiltonian:

$$H_P = \sum_{i=1}^n \sum_{j \geq i}^n Q_{ij} \frac{1}{4} (I - Z_i) (I - Z_j) + \sum_{i=1}^n c_i \frac{1}{2} (I - Z_i). \quad (3)$$

Upon simplification, constant terms are eliminated, resulting in an Ising-type Hamiltonian:

$$H_P = \sum_{i=1}^n h_i Z_i + \sum_{i < j}^n J_{ij} Z_i Z_j, \quad (4)$$

where h_i represents local magnetic fields and J_{ij} encodes inter-qubit-qubit couplings. The ground state of H_P encodes the optimal solution to the original QUBO problem (2), thereby establishing a connection between combinatorial optimization and quantum adiabatic evolution.

4. Quantum MRA

The classical MRA problem objective is:

$$\min_{R_1, \dots, R_N \in \text{SO}(3)} \sum_{(i,j)} \|\tilde{R}_{ij} R_i - R_j\|_F^2. \quad (5)$$

$\text{SO}(3)$ is the set of matrix elements R that are orthogonal and have a unit determinant. $\{R_1, \dots, R_N\} \subset \text{SO}(3)$ represent a set of absolute 3D rotational matrices w.r.t. a global reference frame. \tilde{R}_{ij} is the observed relative rotation; see App. A for detailed review. The process of transforming Eq. (5) into a QUBO form that can be sampled on a quantum annealer can be decomposed into two independent primary steps: 1) reformulating the original optimization problem (26) into an equivalent quadratic optimization problem

through algebraic manipulation and constraint embedding; 2) representing rotation-parameterizing variables as binary strings (e.g., via fixed-point quantization). We present intermediate steps below:

Proposition 1. Problem (26) can be formulated as a quadratic optimization problem in matrix form:

$$\min_{R_1, \dots, R_N} \sum_{(i,j)} \|\tilde{R}_{ij} R_i - R_j\|_F^2 = \min_R R^\top Q R, \quad (6)$$

with $R = [\dots \text{vec}(R_i)^\top \dots]^\top \in \mathbb{R}^{9N \times 1}$, where “ $\text{vec}(\cdot)$ ” denotes the vectorization of a matrix by stacking its columns, and $Q \in \mathbb{R}^{9N \times 9N}$ is the cost matrix:

$$Q = -2 \begin{bmatrix} I \otimes R_{11}^\top & \dots & I \otimes R_{1N}^\top \\ \vdots & \ddots & \vdots \\ I \otimes R_{N1}^\top & \dots & I \otimes R_{NN}^\top \end{bmatrix}. \quad (7)$$

Proof. see App. D. □

4.1. Taylor Linearization on $\text{SO}(3)$ Manifold

Instead of parametrizing rotation matrices on a binary-weighted matrix basis [24], we leverage the relationship between the orthogonal Lie group $\text{SO}(3)$ and its algebra $\mathfrak{so}(3)$ to ensure optimization strictly within the $\text{SO}(3)$ manifold [42]. Any matrix $R_i \in \text{SO}(3)$ can be generated from its tangent vector $\mathbf{v}_i = (v_i^1, v_i^2, v_i^3)^\top \in \mathbb{R}^3$ via the exponential map implemented through Rodrigues’ rotation formula:

$$R_i(\mathbf{v}_i) = \exp(\mathcal{R}(\mathbf{v}_i)) = I + \frac{\sin \theta}{\theta} \mathcal{R}(\mathbf{v}_i) + \frac{1 - \cos \theta}{\theta^2} \mathcal{R}^2(\mathbf{v}_i), \quad (8)$$

where $\mathcal{R} : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ denotes Lie algebra isomorphism:

$$\mathcal{R}(\mathbf{v}_i) := \begin{pmatrix} 0 & -v_i^3 & v_i^2 \\ v_i^3 & 0 & -v_i^1 \\ -v_i^2 & v_i^1 & 0 \end{pmatrix}, \quad \mathcal{R}(\mathbf{v}_i)^\top = -\mathcal{R}(\mathbf{v}_i). \quad (9)$$

The magnitude $\theta = \|\mathbf{v}_i\|_2 \in \mathbb{R}$ defines the rotation angle and the unit vector $\mathbf{x}_i = \mathbf{v}_i / \|\mathbf{v}_i\|_2$ defines the rotation axis. By aggregating all tangent vectors \mathbf{v}_i into $\mathbf{v} = (\dots, \mathbf{v}_i^\top, \dots)^\top \in \mathbb{R}^{3N \times 1}$ and vectorizing the rotations as $R(\mathbf{v}) = (\dots, \text{vec}(R_i(\mathbf{v}_i))^\top, \dots)^\top \in \mathbb{R}^{9N \times 1}$, we reformulate the right-hand side of Eq. (6) as:

$$\min_{\mathbf{v} \in \mathbb{R}^{3N}} R(\mathbf{v})^\top Q R(\mathbf{v}). \quad (10)$$

This transformation reduces the search space from a nonlinear $\text{SO}(3)$ manifold to Euclidean tangent vectors \mathbf{v}_i , where we solve for $\arg \min$ of problem (10). Due to the non-linearity of the exponential map, we expand R_i around its current guess \mathbf{v}_i^k up to the first order in the tangent space. The next iterate \mathbf{v}_i^{k+1} is obtained via solving the resulting

locally non-convex quadratic sub-problem, ensuring that updates always remain within the tangent space of $\text{SO}(3)$.

Specifically, we iteratively replace $R(\mathbf{v})$ in problem (10) with its first-order Taylor linearization:

$$\hat{R}^k(\Delta \mathbf{v}) := R(\mathbf{v}^k) + \nabla R(\mathbf{v}^k)^\top \Delta \mathbf{v}, \quad (11)$$

and optimize over $\Delta \mathbf{v}$, i.e.,

$$\begin{aligned} \Delta \mathbf{v}^{k+1} &= \arg \min_{\Delta \mathbf{v} \in \mathbb{R}^{3N}} \hat{R}^k(\Delta \mathbf{v})^\top Q \hat{R}^k(\Delta \mathbf{v}), \\ \mathbf{v}^{k+1} &= \mathbf{v}^k + \Delta \mathbf{v}^{k+1}, \end{aligned} \quad (12)$$

In Eq. (11), the gradient $\nabla R(\mathbf{v}^k) \in \mathbb{R}^{3N \times 9N}$ is a block diagonal matrix with partial derivative-wise vectorized $\nabla R_i(\mathbf{v}^k) \in \mathbb{R}^{3 \times 9}$ over the diagonal.

Remark 1. While the exact mapping $R(\mathbf{v}^{k+1})$ inherently preserves the $\text{SO}(3)$ manifold structure by construction, its first-order Taylor approximation $\hat{R}^k(\Delta \mathbf{v}^{k+1})$ can deviate from $\text{SO}(3)$ on the scale of $\mathcal{O}(\|\Delta \mathbf{v}^{k+1}\|^2)$.

To improve approximation fidelity while ensuring convergence, we propose: 1) iteration-adaptive box constraint in (12) (similar to the one used in Ref. [16, 42]), i.e. enforce

$$\|\Delta \mathbf{v}\|_\infty \leq \delta^k, \quad (13)$$

and 2) an orthogonality-promoting soft penalty term $\alpha \|\hat{R}^k(\Delta \mathbf{v})\|_F^2$ (see App. E on how the penalty encourages the (blocks of the) Taylor approximation to maintain $\text{SO}(3)$ proximity). We summarize our approach as follows:

Proposition 2. *Inserting the linearization (11) into the minimization Problem (12) and incorporating the softly-regularized box constraint, we obtain the QUBO problem:*

$$\begin{aligned} \Delta \mathbf{v}^{k+1} &= \arg \min_{\|\Delta \mathbf{v}\|_\infty \leq \delta^k} \Delta \mathbf{v}^\top \hat{Q} \Delta \mathbf{v} + \hat{c}^\top \Delta \mathbf{v}, \\ \mathbf{v}^{k+1} &= \mathbf{v}^k + \Delta \mathbf{v}^{k+1}, \end{aligned} \quad (14)$$

with

$$\begin{aligned} \hat{Q} &:= \nabla R(\mathbf{v}^k) (Q + \alpha N I_{9N}) \nabla R(\mathbf{v}^k)^\top, \\ \hat{c} &:= 2 \nabla R(\mathbf{v}^k) (Q + \alpha N I_{9N}) R(\mathbf{v}^k). \end{aligned} \quad (15)$$

We observe that aggressive regularization can enforce positive semi-definiteness (PSD) in \hat{Q} —thereby convexifying Eq. (14) and enabling classical efficient solvability. However, it simultaneously distorts the landscape by overwhelming the objective function and introduces artificial minima that fundamentally misrepresent the true optimization landscape. This marks calibration of α crucial for the optimal performance such that: 1) Taylor approximation’s deviation from $\text{SO}(3)$ within $\mathcal{O}(\|\Delta \mathbf{v}\|^2)$ bounds is constrained and 2) distortion of the original optimization landscape is minimized; see Fig. 2 for a schematic visualization.

Remark 2. The exponential map $\mathfrak{so}(3) \rightarrow \text{SO}(3)$ preserves Lie group traversal at each iteration $\mathbf{v}^{k+1} = \mathbf{v}^k + \Delta \mathbf{v}$ throughout the optimization.




Figure 2. Schematic local search behavior visualization under varying penalty strengths α .

4.2. Binary Encoding of Iterative Update

We discretize the search space $\Delta \mathbf{v}$ and leverage binary encodings to represent the discretized space for compatibility with the qubit-based Ising model supported by the annealer; see Sec. 3. Specifically, we represent $\Delta \mathbf{v}$ via a binary vector $\mathbf{q} \in \{0, 1\}^m$, where each binary variable q_i corresponds to the state of a logical qubit on the annealer. For hardware implementations (e.g., D-Wave systems), these binary variables are often mapped to spin states $\sigma_i \in \{-1, +1\}$ via the transformation $q_i = \frac{\sigma_i + 1}{2}$; see Sec. 3. To satisfy the constraint $\|\Delta \mathbf{v}\|_\infty \leq \delta^k$, we discretize the interval $[-\delta^k, \delta^k]$ uniformly using m qubits per dimension. Specifically, we express i -th discretized value of $\Delta \mathbf{v}$ as:

$$(\Delta \mathbf{v})_i = -\delta^k + \frac{2\delta^k}{s} \sum_{\ell=0}^{m-1} 2^\ell q_{i,\ell}, \quad \text{where } s = 2^m - 1. \quad (16)$$

where factor $2\delta^k/s$ with $s := 2^m - 1$ serves as a normalization factor and ensures that $\|\Delta \mathbf{v}\|_\infty \leq \delta^k$. Note that Eq. (16) discretizes the search space to

$$(\Delta \mathbf{v})_i = \left\{ -\delta^k + \frac{2T\delta^k}{s} : T \in \mathbb{N}, T \leq 2^m - 1 \right\} \quad (17)$$

representing *exponentially-many*, i.e. 2^m , equidistant points per dimension. We then write the variable vector $\Delta \mathbf{v}$ more compactly via a matrix-vector product as

$$\Delta \mathbf{v} = -\delta^k \mathbf{1}_{3N} + D \mathbf{q}, \quad (18)$$

where

$$\begin{aligned} \mathbf{q} &= \begin{bmatrix} \mathbf{q}_1 \\ \vdots \\ \mathbf{q}_{3N} \end{bmatrix} \in \mathbb{R}^{3Nm}, \quad \text{with } \mathbf{q}_i = \begin{bmatrix} q_{i,0} \\ \vdots \\ q_{i,m-1} \end{bmatrix} \in \mathbb{R}^m, \\ D &= \frac{2}{s} \begin{bmatrix} \delta^k & & 0 \\ & \ddots & \\ 0 & & \delta^k \end{bmatrix} \otimes [2^0 \ \dots \ 2^{m-1}] \in \mathbb{R}^{3N \times 3Nm}, \end{aligned} \quad (19)$$

$\mathbf{1}_{3N}$ is a vector of length $3N$ with all entries being one, and $\mathbf{q} \in \{0, 1\}^{3Nm}$ stacks all binary variables.

Remark 3. The binary encoding’s approximation precision when representing arbitrary floating-point values improves exponentially with scaling $\mathcal{O}(2^m)$.

Proposition 3. Binary encoding of the continuous search interval makes $\mathbf{q} \in \mathbb{B}^{3Nm}$ the new optimization variable, changing Problem (14) into

$$\arg \min_{\mathbf{q} \in \mathbb{B}^{3Nm}} \mathbf{q}^\top \tilde{Q} \mathbf{q} + \mathbf{q}^\top \tilde{\mathbf{c}}, \quad (20)$$

with

$$\begin{aligned} \tilde{Q} &:= D^\top \hat{Q} D, \\ \tilde{\mathbf{c}} &:= D^\top (\hat{\mathbf{c}} - 2\delta^k \hat{Q} \mathbf{1}_{3N}). \end{aligned} \quad (21)$$

Proof. See App. F. \square

Upon obtaining the optimal \mathbf{q} via Eq. (20), we update the solution iterate as \mathbf{v}^{k+1} via Eqs. (12) and (18). The iteration continues until the energy update converges.

4.3. Required Hardware Resource Analysis

Practical implementation of IQARS and its performance is constrained by available hardware resources. For synchronizing a camera graph of N nodes, our algorithm maintains favorable resource scaling properties and require $3Nm$ logical qubits (see Prop. 3) while achieving floating-point representation error on the order of $\mathcal{O}(\delta^{(k)}/2^m)$. The scaling behavior demonstrates: 1) *linear* dependence on problem size and 2) *logarithmic* dependence on solution precision, suggesting strong algorithmic potentials for large-scale MRA problems as hardware continues to evolve.

Remark 4. With limited quantum resource budgets, the multiplicative nature of resource requirement introduces trade-offs between problem scales and solution quality.

4.4. Algorithmic Protocol of IQARS

Large-scale MRA problems present complex optimization landscapes that typically require fine discretization, i.e. large m , for accurate solutions. However, current quantum hardware limitations impose strict constraints on available qubit resources for such a purpose. If discretization is coarse, the global optimum of Eq. (20) can easily lie outside the discretized search grid defined by Eq. (17), potentially leading to oscillatory convergence behavior. To ensure robust convergence behavior, we implement an adaptive search protocol inspired by [4, 5, 42] that dynamically adjusts the search radius δ based on local update magnitudes during optimization. While this heuristic approach does not provide formal convergence guarantees—particularly with respect to maintaining global optimality through successive radius contractions—our empirical results demonstrate its effectiveness in achieving stable convergence to high-quality solutions. This adaptive mechanism proves particularly valuable in resource-constrained scenarios where

coarse discretization would otherwise severely limit solution quality. We outline the algorithm in Alg. 1.

Algorithm 1 Iterative Quantum Approach for Rotation Synchronization (IQARS)

- 1: **Input:** R_{ij} : Pairwise relative rotations
 Maxiter: Maximum number of iterations
 \mathbf{v}^0 : Initial rotation guess
 δ^0 : Search window radius
 κ : Threshold for decreasing search radius δ
 m : Qubits required for search discretization
 ϵ : Stopping criterion
 τ : Annealing rate
 - 2: **Initialization:** $j, e \leftarrow 0$; Construct Q according to Eq. (7); Initialized identity rotations: $R(\mathbf{v}^0) = I$.
 - 3: **while** $j \leq \text{Maxiter}$ **do:**
 - 4: Compute gradient $\nabla R(\mathbf{v}^j)$ according to Eq. (8)
 - 5: Construct \hat{Q} and $\hat{\mathbf{c}}$ according to Eq. (15)
 - 6: Construct D according to Eq. (19)
 - 7: Solve the QUBO problem (20) and measure binary solution vector \mathbf{q}
 - 8: $j \leftarrow j + 1$
 - 9: Update solution \mathbf{v}^j and $R(\mathbf{v}^j)$ using Eq. (12) and (8); update $e \leftarrow \text{Avg} \sum_{(i,j)} \|R_{ij} R_i - R_j\|_F^2$
 - 10: **If** $e < \epsilon$ **do:** break
 - 11: **Else:**
 - 12: **If** $\|R(\mathbf{v}^j) - R(\mathbf{v}^{j-1})\|_F < \kappa$ **do:**
 - 13: $\delta \leftarrow \frac{\delta}{\tau}$; $\kappa \leftarrow \frac{\kappa}{\tau}$
 - 14: **End**
 - 15: **End**
 - 16: **End**
 - 17: **return** Residual: e ; iteration count: j ; synchronized rotations: $\mathbf{v}^j, R(\mathbf{v}^j)$.
-

4.5. Solution Refinement via Posterior Analysis

Quantum annealers return solutions by sampling from low-energy states of a pre-defined QUBO energy landscape. Under ideal adiabatic conditions, the system will converge to the ground state, returning the global optimal solution. However, ensuring adiabatic evolution requires problem-specific annealing schedules, and in practice, environmental interactions such as thermal excitations could induce non-adiabatic effects that lead to a distribution of low-energy states well-approximated by classical Boltzmann distributions rather than deterministic ground-state [47, 54]. To improve the solution quality due to non-adiabatic effects leveraging the probabilistic outputs, we propose a *physics-inspired energy-weighted voting* approach that refines solution quality. The method assigns Boltzmann weights w_i :

$$w_i = \frac{e^{-\beta E(\mathbf{q}_i)}}{\mathcal{Z}}, \quad \mathcal{Z} = \sum_{j=1}^K e^{-\beta E(\mathbf{q}_j)}, \quad (22)$$




Figure 3. IQARS convergence behavior for synchronizing $N=10$ synthetic random noiseless rotations.

to each of the K -lowest energy solutions $\{\mathbf{q}_i\}_{i=1}^K$, where $\mathcal{Z} = \sum_{j=1}^K e^{-\beta E(\mathbf{q}_j)}$ serves as the partition function and $E(\mathbf{q}_i)$ denotes the solution energy. β represents an effective inverse temperature parameter. This creates a noise-robust aggregation scheme that naturally prioritizes lower-energy configurations. Details of the posterior refinement protocol are formally described in App. I. *While this protocol is evaluated specifically for MRA, we emphasize that it establishes a general methodology for enhancing solutions obtained from quantum annealing. The approach’s theoretical foundations in statistical mechanics and its demonstrated efficacy suggest significant potential for application to a range of other problems, which remain to be explored.*

5. Experiments

We empirically evaluate IQARS’s accuracy and robustness using synthetic and real-world datasets [49]. The QUBO parametrization—comprising coupling matrices \tilde{Q} and linear bias term \tilde{c} —is prepared on an Intel Xeon CPU @ 2.20 GHz with 12 GB RAM. Prepared QUBO formulations follow by execution on D-Wave *Advantage Series* 6.4 annealers operating at a cryogenic base temperature of 15 mK. All experimental configurations remain consistent unless explicitly mentioned otherwise. The annealing protocol maintains the default duration of 20 μ s apart from data transmission overheads. Each experimental trial is conducted with 10^2 reads per iteration. The search space is configured with an initial value of $\delta^0 = \pi/30$ discretized uniformly with $m = 3$ qubits, yielding $2^3 = 8$ equidistant sub-intervals—a coarse discretization scheme explicitly chosen to accommodate a larger problem scale N with constrained quantum computational budget; see details in Sec. 4.3.

Noise-free Synthetic Rotation Graphs. We construct a fully-connected synthetic rotation graph of N nodes. Each node i is initialized with a random tangent vector $\mathbf{v}_i \in \mathbb{R}^3$ uniformly sampled from the Lie algebra $\mathfrak{so}(3)$ and later mapped to its corresponding rotation matrix $R(\mathbf{v}_i)$ via the

N	Solvers	Avg. $\ R_{ij} - R_j R_i^T\ \downarrow$	Avg. $\ \theta^* - \theta\ \downarrow$
10	Ours (iterative)	$\{1.484 \pm 0.08\} \mathbf{e-17}$	$\{0.933 \pm 0.05\} \mathbf{e-17}$
	Ours (direct)	0.863 ± 0.07	0.826 ± 0.08
15	Ours (iterative)	$\{1.156 \pm 0.24\} \mathbf{e-17}$	$\{7.843 \pm 0.87\} \mathbf{e-18}$
	Ours (direct)	1.226 ± 0.18	0.975 ± 0.17
20	Ours (iterative)	$\{9.342 \pm 0.86\} \mathbf{e-17}$	$\{6.685 \pm 0.69\} \mathbf{e-17}$
	Ours (direct)	0.892 ± 0.065	0.793 ± 0.059

Table 1. Benchmark results of IQARS for MRA on synthetic noise-free rotations.

exponential mapping defined in Eq. (8). For each node pair $\{i, j\}$, we compute the relative rotation, constructing a dense symmetric matrix $\tilde{R}_{ij} \in R^{N \times N}$ that encapsulates the relative rotations within the camera graph. Leveraging the constructed dense relative rotation matrix, we implement Alg. 1 to recover synchronized absolute rotations R_i while systematically monitoring the convergence dynamics through quantitative metrics including evolution of energy spectrum E_m , neighboring energy gaps $E_m - E_n$, and rotation residuals throughout iterations, as visualized in Fig. 3. The observed monotonic decay of rotation residuals along with the absence of energy gap closures provides empirical validation for the theoretical convergence guarantees established in Theorem 1. In addition to the visualized convergence dynamics, the performance of synchronizing noise-free synthetic rotations of different scales with IQARS is recorded in Tab. 1. Besides IQARS—which strictly maintains the optimization path within the $SO(3)$ manifold through iterative exponential mapping; see Sec. 4.1—we also evaluate an alternative direct approach proposed in [24] that can be combined with Prop. 1 for MRA. This approach, however, lacks formal $SO(3)$ constraint guarantees; see App. H for implementation details. We denote IQARS as “ours (iterative)”, and the alternative approach as “ours (direct)”, with the experimental results summarized in Tab. 1 to enable a principled comparison of different approaches. We show that, compared to [24], our design is empirically valid and achieves promising performance on D-Wave’s annealing machines.

Hyperparameter Study. We perform a rigorous sensitivity analysis of hyperparameters in Alg. 1 on synthetic datasets via controlled experiments. Specifically, we evaluate: (i) annealing threshold κ , (ii) annealing rate τ , (iii) qubit count m for discretization, and (iv) penalty coefficient α . All hyperparameters remained fixed post-optimization for controlled studies. Our experimental results as presented in Fig. 4 (log scale) demonstrate that increasing m at the cost of more quantum resource budget leads to monotonic improvements in convergence properties due to increased search interval representation fidelity; see Sec. 4.3. Reducing the annealing rate τ results in more stable optimization trajectories. As the problem scales up, the con-




Figure 4. Visualization of IQARS’s hyperparameter choice on the convergence behavior (MRA residual on log scale on y-axis).

vergence behavior tends to get progressively slower under otherwise identical experimental conditions. Notably, the penalty factor α exhibits a Goldilocks effect—both excessive and insufficient penalization degrade performance, leaving the optimal α worth exploring. This empirical observation is consistent with our theoretical analysis earlier in Sec. 4.1. These findings collectively establish principled guidelines for parameter selection in leveraging IQARS.

Noise-Corrupted Synthetic Rotation Graphs. Practically, observed real-world rotation measurements are inherently noisy. We model this scenario and generate perturbed measurements \tilde{R}_{ij} by applying controlled multiplicative distortions on the $\text{SO}(3)$ manifold. Specifically, we simulate such noise corruption by: (i) sampling a random tangent vector $\mathbf{v} \in \mathbb{R}^3$ uniformly from $[0, 1]^3$, (ii) scaling it by a prescribed noise level σ , and (iii) mapping it to $\text{SO}(3)$ via the exponential map $\exp(\sigma\mathbf{v})$. This noise construction is consistent with previous approaches [36, 51], ensuring that the perturbed rotations remain properly constrained to the manifold while preserving the geometric Lie group structure. We evaluate the performance of IQARS and provide additional experiments with established noise-robust classical solvers, specifically the Riemannian optimization framework of Shonan [17] and the robust L1-IRLS algorithm [10]. L1-IRLS leverages minimum spanning tree (MST) initialization [10, 25], while Shonan employs random $\text{SO}(p)$ initialization; $p \geq 1$ [17]. We visualize in Fig. 5 the performance of different MRA solvers on synchronizing $N = 20$ rotations across increasing noise ratios σ . While primarily designed to pioneer a novel quantum solver for MRA rather than challenging classical MRA solvers especially with available hardware of limited resource budget resulting in small problem sizes and coarse search space discretization, empirical results highlight that IQARS on D-Wave machines can achieve residual reduction by approximately 12% compared to the second best—Shonan, across a wide range of noise regimes, showing




Figure 5. Benchmark results of IQARS against other solvers for MRA on a synthetic noisy dataset across noise levels σ ; R_{iGT} represents the ground truth of R_i .

great potentials of our approach in future. For completeness, we also: i) benchmark IQARS against two prominent local optimization methods—trust-region and Levenberg-Marquardt (LM)—both theoretically and empirically; see details in App. N and ii) evaluate IQARS (identity initialization by default to avoid heuristic bias) with other initializations; see App. O.

Real-World Noisy Dataset Evaluation. To validate the practical utility of IQARS beyond synthetic experiments, we evaluate its performance alongside Shonan and L1-IRLS on established real-world benchmark datasets [49]. Results are documented in Tab. 2, demonstrating that IQARS maintains robustness and has performance on par with or outperforms other classical solvers on real-world datasets. This observation is consistent with our previous experiments on synthetic datasets. The chain break ratio, as a proxy for IQARS embedding quality within the quantum hardware topology, is also monitored and reported in Tab. 2. As the

N	Dataset	Evaluation Metrics	Ours (iterative)	Shonan	L1-IRLS	Ours (direct)
11	Fountain	Avg. $\ R_{ij} - R_j R_i^T\ \downarrow$	3.72e-3	4.18e-3	4.39e-3	8.17e-2
		Avg. $\ \theta^* - \theta\ \downarrow$	3.51e-3	4.04e-3	4.23e-3	6.73e-2
		Avg. chain break (%) \downarrow	6.2e-2	/	/	5.81e-2
15	Castle	Avg. $\ R_{ij} - R_j R_i^T\ \downarrow$	1.12e-3	1.33e-3	1.65e-3	7.65e-2
		Avg. $\ \theta^* - \theta\ \downarrow$	1.02e-3	1.19e-3	1.48e-3	5.65e-2
		Avg. chain break (%) \downarrow	1.13e-1	/	/	1.02e-1
8	Herz-Jesus	Avg. $\ R_{ij} - R_j R_i^T\ \downarrow$	3.26e-3	3.89e-3	4.06e-3	9.86e-2
		Avg. $\ \theta^* - \theta\ \downarrow$	3.08e-3	3.62e-3	3.88e-3	8.25e-2
		Avg. chain break (%) \downarrow	3.8e-2	/	/	4.2e-2

Table 2. Benchmark results of IQARS against other solvers for MRA on real-world datasets [49]; hardware statistics such as chain break ratios reflecting machine embedding quality are reported.

problem scales up, i.e. larger N , the chain break ratio is observed to slightly increase. While this (at a low level) does not seem to significantly impact the performance, this reflects an inherent limitation of embedding larger problems into existing hardware topology. Lastly, we also integrated IQARS within the Glomap, a SfM pipeline [46]; see App. J, for a qualitative visual analysis of the Poisson-reconstructed meshes. The visualization in Fig. 6 shows accurate 3D reconstructions, establishing IQARS as a robust MRA method with potentials for applications such as SfM.

Posterior Analysis. We rigorously evaluate our posterior




Figure 6. Reconstructed 3D geometries from multi-view imagery via Glomap [46] integrated with IQARS. Both color-textured (top) and geometry-only (bottom) representations of three benchmark datasets—*Fountain*, *Castle*, and *Herz-Jesus* (left-to-right) are visualized. Quantitative performance metrics are detailed in Tab. 2.

protocol that aims to relieve the thermal noise effect during annealing and improve the solution quality; see Sec. 4.5. The protocol incorporates low-energy solution spaces of various spans. We define and maintain inverse temperature $\beta = 2$ —carefully chosen to empirically balance exploration/exploitation trade-offs—and systematically vary the span K of low-energy solution spaces. We observe that the proposed posterior protocol consistently improves solution quality and yields statistically significant gains over the unrefined annealer outputs. Detailed statistical results with different variable spans K of the solution space are recorded and visualized in Fig. 7. The visualization additionally shows a characteristic saturation curve: with initial increases in K yielding substantially improved performance, the marginal improvement diminishes asymptotically beyond an optimal value, i.e. a sweet spot, for low-energy sample space selection. This sweet spot for the MRA is empirically observed to be around $K = 30$, suggesting optimal sample space selection.

Embedding on Contemporary Annealers. Alongside theoretical scalability analysis in Sec. 4.3 that characterize *logical qubit* requirements of IQARS, we evaluate its *physical resource* demands by considering problem embeddings on the hardware topology of D-Wave Advantage System (Pegasus architecture). Specifically, we measure the physical resource consumption for MRA across varying problem sizes N and visualized it in Fig. 13 in App.; it also includes other machine versions. As an example, for synchronizing $N = 15$ rotations requiring $135 = 3Nm$ logical qubits, approximately 2500 physical qubits are required to embed the logical problem on the hardware topology of the Advantage 6.4 annealer. This accounts for nearly half of the processor’s total capacity (5000 physical qubits). *While IQARS is designed to be resource-efficient, its practical implementation for real-world tasks on contemporary annealer hardware can present challenges that constrain its performance.*




Figure 7. Statistical improvements due to our refinement protocol with variable span K . Top- K refers to posterior aggregation performed over the K lowest-energy annealer samples.

6. Discussion and Conclusion

We introduced IQARS, the first framework for MRA leveraging quantum hardware. While not positioned as a replacement for classical approaches in the near future, IQARS establishes a novel computational paradigm exploration for MRA. Despite limited quantum resource budget resulting in constrained performance, extensive experiments demonstrate that: i) IQARS can recover ground-truth solutions for MRA with a noise-free synthetic dataset; ii) for noisy datasets (synthetic and real-world), a consistent 10–15% residual reduction is observed for IQARS compared to Shonan. With more resources accessible leading to e.g. finer search discretization, the performance of IQARS is expected to improve significantly as per our hyperparameter study. Our posterior refinement protocol statistically improves sampled solution quality, achieving up to a 60% probability of producing lower-residual solutions compared to the unrefined annealer outputs. We aim to establish foundations for advances in 3D vision that builds upon ours in Quantum-enhanced Computer Vision (QeCV), and will release source codes for reproducibility.

Limitations. While IQARS, as a resource-efficient quantum MRA algorithm, demonstrates considerable promise, the experimental scale and its potential for large-scale real-world problems remain constrained with limited resource budget from contemporary annealer hardware.

Future Work. IQARS holds significant potential for diverse downstream applications. Beyond SfM, promising tasks include robotic navigation, medical imaging and augmented reality, among others. Further investigation into alternative MRA formulations with different hardware support presents a compelling avenue for future research.

Acknowledgement. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), project number 534951134. NKM and MM acknowledge support by the Lamarr Institute for Machine Learning and Artificial Intelligence.

References

- [1] Tameem Albash and Daniel A Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1):015002, 2018. **11**
- [2] Federica Arrigoni, Willi Menapace, Marcel Seelbach Benkner, Elisa Ricci, and Vladislav Golyanik. Quantum motion segmentation. In *European Conference on Computer Vision*, 2022. **2**
- [3] Marcel Seelbach Benkner, Zorah Löhner, Vladislav Golyanik, Christof Wunderlich, Christian Theobalt, and Michael Moeller. Q-match: Iterative shape matching via quantum annealing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7586–7596, 2021. **2**
- [4] Marsha J Berger and Phillip Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of computational Physics*, 82(1):64–84, 1989. **5**
- [5] Marsha J Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of computational Physics*, 53(3):484–512, 1984. **5**
- [6] Harshil Bhatia, Edith Tretschk, Zorah Löhner, Marcel Seelbach Benkner, Michael Moeller, Christian Theobalt, and Vladislav Golyanik. Ccuantumm: Cycle-consistent quantum-hybrid matching of multiple shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. **2**
- [7] Tolga Birdal, Vladislav Golyanik, Christian Theobalt, and Leonidas J Guibas. Quantum permutation synchronization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. **2**
- [8] Max Born and Vladimir Fock. Beweis des adiabatenatzes. *Zeitschrift für Physik*, 51(3):165–180, 1928. **11**
- [9] Cameron Braunstein, Eddy Ilg, and Vladislav Golyanik. Quantum-hybrid stereo matching with nonlinear regularization and spatial pyramids. In *International Conference on 3D Vision (3DV)*, 2024. **2**
- [10] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *Proceedings of the IEEE international conference on computer vision*, 2013. **1, 2, 7**
- [11] Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):958–972, 2017. **1**
- [12] Yu Chen, Ji Zhao, and Laurent Kneip. Hybrid rotation averaging: A fast and robust rotation averaging approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10358–10367, 2021. **2**
- [13] Han Yao Choong, Suryansh Kumar, and Luc Van Gool. Quantum annealing for single image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. **2**
- [14] Hainan Cui, Xiang Gao, Shuhan Shen, and Zhanyi Hu. Hsfm: Hybrid structure-from-motion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. **1, 2**
- [15] Zhaopeng Cui and Ping Tan. Global structure-from-motion by similarity averaging. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015. **1, 2**
- [16] Amer Delilbasic, Bertrand Le Saux, Morris Riedel, Kristel Michielsen, and Gabriele Cavallaro. A single-step multiclass svm based on quantum annealing for remote sensing data classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2023. **4**
- [17] Frank Dellaert, David M Rosen, Jing Wu, Robert Mahony, and Luca Carlone. Shonan rotation averaging: Global optimality by surfing so (p) n so (p) n. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, 2020. **1, 2, 7**
- [18] Vasil S Denchev, Sergio Boixo, Sergei V Isakov, Nan Ding, Ryan Babbush, Vadim Smelyanskiy, John Martinis, and Hartmut Neven. What is the computational value of finite-range tunneling? *Physical Review X*, 6(3):031015, 2016. **2**
- [19] Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. **2**
- [20] Matteo Farina, Luca Magri, Willi Menapace, Elisa Ricci, Vladislav Golyanik, and Federica Arrigoni. Quantum multi-model fitting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. **2**
- [21] Johan Fredriksson and Carl Olsson. Simultaneous multiple rotation averaging using lagrangian duality. In *Asian Conference on Computer Vision*, 2012. **1**
- [22] Xiang Gao, Lingjie Zhu, Zexiao Xie, Hongmin Liu, and Shuhan Shen. Incremental rotation averaging. *International Journal of Computer Vision*, 129:1202–1216, 2021. **2**
- [23] Michael Goesele, Brian Curless, and Steven M Seitz. Multi-view stereo revisited. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, 2006. **1**
- [24] Vladislav Golyanik and Christian Theobalt. A quantum computational approach to correspondence problems on point sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. **2, 3, 6, 11, 12**
- [25] Richard Hartley, Jochen Trumpf, Yuchao Dai, and Hongdong Li. Rotation averaging. *International journal of computer vision*, 103:267–305, 2013. **1, 7, 15**
- [26] Jeremy Hsu. D-wave’s year of computing dangerously, 2013. **1**
- [27] Wolfgang Hürst and Matthias Helder. Mobile 3d graphics and virtual reality interaction. In *Proceedings of the 8th international conference on advances in computer entertainment technology*, pages 1–8, 2011. **1**
- [28] Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler. Bounds for the adiabatic approximation with applications to quantum computation. *Journal of Mathematical Physics*, 48(10), 2007. **11**
- [29] Saike Jiang, Shilin Wang, Zhongyi Yi, Meina Zhang, and Xiaolan Lv. Autonomous navigation system of greenhouse mobile robot based on 3d lidar and 2d lidar slam. *Frontiers in plant science*, 13:815218, 2022. **1**
- [30] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Physical Review E*, 58(5):5355, 1998. **11**

- [31] Iman Abasapur Kazerouni, Luke Fitzgerald, Gerard Dooly, and Daniel Toal. A survey of state-of-the-art on visual slam. *Expert Systems with Applications*, 205:117734, 2022. [1](#)
- [32] Andrew D. King et al. Computational supremacy in quantum simulation. *arXiv e-prints*, 2024. [2](#)
- [33] Natacha Kuete Meli, Shuteng Wang, Marcel Seelbach Benkner, Michele Sasdelli, Tat-Jun Chin, Tolga Birdal, Michael Moeller, and Vladislav Golyanik. Quantum-enhanced computer vision: Going beyond classical algorithms. *arXiv e-prints*, 2025. [2](#)
- [34] T. Lanting et al. Entanglement in a quantum annealing processor. *Phys. Rev. X*, 4:021041, 2014. [2](#)
- [35] Seong Hun Lee and Javier Civera. Robust single rotation averaging. *arXiv preprint arXiv:2004.00732*, 2020. [2](#)
- [36] Seong Hun Lee and Javier Civera. Robust single rotation averaging revisited, 2024. [7](#)
- [37] Junde Li and Swaroop Ghosh. Quantum-soft qubo suppression for accurate object detection. In *European Conference on Computer Vision*, 2020. [2](#)
- [38] Philipp Lindenberger, Paul-Edouard Sarlin, Viktor Larsson, and Marc Pollefeys. Pixel-perfect structure-from-motion with featuremetric refinement. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5987–5997, 2021. [1](#)
- [39] Andréa Macario Barros, Maugan Michel, Yoann Moline, Gwenolé Corre, and Frédéric Carrel. A comprehensive survey of visual slam algorithms. *Robotics*, 11(1):24, 2022. [1](#)
- [40] Daniel Maier, Armin Hornung, and Maren Bennewitz. Real-time navigation in 3d environments based on depth camera data. In *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)*, pages 692–697. IEEE, 2012. [1](#)
- [41] Hidenobu Matsuki, Riku Murai, Paul HJ Kelly, and Andrew J Davison. Gaussian splatting slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [1](#)
- [42] Natacha Kuete Meli, Florian Mannel, and Jan Lellmann. An iterative quantum approach for transformation estimation from point sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 529–537, 2022. [2](#), [3](#), [4](#), [5](#)
- [43] Pierre Moulon, Pascal Monasse, and Renaud Marlet. Global fusion of relative motions for robust, accurate and scalable structure from motion. In *Proceedings of the IEEE international conference on computer vision*, pages 3248–3255, 2013. [1](#)
- [44] Yana Nehmé, Jean-Philippe Farrugia, Florent Dupont, Patrick Le Callet, and Guillaume Lavoué. Comparison of subjective methods for quality assessment of 3d graphics in virtual reality. *ACM Transactions on Applied Perception (TAP)*, 18(1):1–23, 2020. [1](#)
- [45] Onur Özyeşil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion*. *Acta Numerica*, 26:305–364, 2017. [1](#)
- [46] Linfei Pan, Daniel Barath, Marc Pollefeys, and Johannes Lutz Schönberger. Global Structure-from-Motion Revisited. In *European Conference on Computer Vision (ECCV)*, 2024. [7](#), [8](#), [13](#)
- [47] Maria Schuld and Francesco Petruccione. *Machine learning with quantum computers*. Springer, 2021. [5](#)
- [48] Chitturi Sidhartha and Venu Madhav Govindu. It is all in the weights: Robust rotation averaging revisited. In *2021 International Conference on 3D Vision (3DV)*, 2021. [1](#)
- [49] Christoph Strecha, Wolfgang Von Hansen, Luc Van Gool, Pascal Fua, and Ulrich Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *2008 IEEE conference on computer vision and pattern recognition*, 2008. [6](#), [7](#)
- [50] Lanhui Wang and Amit Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2(2):145–193, 2013. [1](#)
- [51] Kyle Wilson, David Bindel, and Noah Snavely. When is rotations averaging hard? In *European Conference on Computer Vision*, pages 255–270. Springer, 2016. [1](#), [7](#)
- [52] Isabell Wohlgenannt, Alexander Simons, and Stefan Stieglitz. Virtual reality. *Business & Information Systems Engineering*, 62:455–461, 2020. [1](#)
- [53] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. [1](#)
- [54] Jan-Nico Zaeck, Martin Danelljan, Tolga Birdal, and Luc Van Gool. Probabilistic sampling of balanced k-means using adiabatic quantum computing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [2](#), [5](#)
- [55] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022. [1](#)

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

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

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.