

Adiabatic Quantum Graph Matching with Permutation Matrix Constraints

Marcel Seelbach Benkner¹

Vladislav Golyanik²

Christian Theobalt²

Michael Moeller¹

¹University of Siegen

²Max Planck Institute for Informatics, SIC

Abstract

Matching problems on 3D shapes and images are challenging as they are frequently formulated as combinatorial quadratic assignment problems (QAPs) with permutation matrix constraints, which are \mathcal{NP} -hard. In this work, we address such problems with emerging quantum computing technology and propose several reformulations of QAPs as unconstrained problems suitable for efficient execution on quantum hardware. We investigate several ways to inject permutation matrix constraints in a quadratic unconstrained binary optimization problem which can be mapped to quantum hardware. We focus on obtaining a sufficient spectral gap, which further increases the probability to measure optimal solutions and valid permutation matrices in a single run. We perform our experiments on the quantum computer D-Wave 2000Q (2¹¹ qubits, adiabatic). Despite the observed discrepancy between simulated adiabatic quantum computing and execution on real quantum hardware, our reformulation of permutation matrix constraints increases the robustness of the numerical computations over other penalty approaches in our experiments. The proposed algorithm has the potential to scale to higher dimensions on future quantum computing architectures, which opens up multiple new directions for solving matching problems in 3D computer vision and graphics¹.

1. Introduction

The question how quantum computing can be beneficial in solving problems in computer vision is still relatively unexplored. A quantum computer is a computing machine which takes advantage of quantum effects, *i.e.*, quantum superposition, entanglement, tunnelling and contextuality [55, 35]. Since late 1980s, the quantum computing paradigm attracts more and more attention of computer scientists. Multiple quantum methods were subsequently shown to improve the computational complexity [11, 61, 60, 31, 65] compared to their classical algorithmic

¹visit <http://gvv.mpi-inf.mpg.de/projects/QGM/> (the source code is available)




Figure 1: Example of 3D shape matching on quantum annealer D-Wave 2000Q. **(Left:)** 3D Shapes from the FAUST dataset [13] and the matches extracted with our QGM. **(Right:)** Histogram of solutions on D-Wave 2000Q. The peaks in the solution distribution arise from the proposed regularization which enhances the chance to measure a valid permutation matrix. The left (highest) bin corresponds to valid permutations. The solutions with the lowest energy, *i.e.*, the left-most measured values in the histogram, correspond to the recovered permutation in the picture.

counterparts.

Only recently the practicability of all these methods has been confirmed [7]. Thus, it is only a few years since quantum computing technology became mature enough and the first functional quantum computers suitable for real-world problems became available for research purposes for a broader community [1, 25, 50, 22, 30].

Our motivation in this paper is to investigate the applicability of modern and upcoming adiabatic quantum computers (AQC-ers) for computer vision and graphics tasks. Therefore, we choose the fundamental problem of visual computing *i.e.*, combinatorial graph matching (CGM). This problem allows for a general formulation and, on the other hand, is challenging enough to map on a quantum computer, for the reasons to become clear in the following.

Many matching problems on 2D, 3D and higher-dimensional structures such as point set and mesh alignment [40, 49] require optimizing a quadratic cost function over the set of permutation matrices, *i.e.*, solving a quadratic as-

signment problem (QAP)

$$\min_{X \in \mathbb{P}_n} f(\mathbf{x}) := \mathbf{x}^T W \mathbf{x} + \mathbf{c}^T \mathbf{x}, \quad (1)$$

where $\mathbf{x} = \text{vec}(X)$ is the n^2 -dimensional vector corresponding to the matrix $X \in \mathbb{R}^{n \times n}$, which is constrained to lie in the set of permutation matrices \mathbb{P}_n . While there exist several well-working heuristics and approximate solvers for (1) (see Sec. 3 for a summary), there are problem instances on which every such approach would fail to provide optimal results within reasonable computational time as to be expected due to the \mathcal{NP} -hard nature of QAPs.

Motivation and Contributions. With adiabatic quantum computing (AQC-ing) moving from theoretical considerations to actual implementations [57, 18], it becomes increasingly attractive for tackling \mathcal{NP} -hard problems in computer vision. Constrained problems like (1) can unfortunately not be solved with AQC-ers directly, and have to be converted to a quadratic *unconstrained* binary optimization problem (QUBO). We observe that a straightforward reformulation of the constrained problem as a QUBO via a quadratic penalty increases the difficulty of the resulting problem significantly. We conjecture that this difficulty is largely related to the problem’s property called the *spectral gap*, *i.e.*, the difference between the lowest and second-lowest energy state, and tailor our approaches at maximizing the latter while still provably solving the original constrained problem. To summarize, our **contributions** are:

- We show how quadratic assignment problems (1) with permutation matrix constraints can be mapped to QUBOs and efficiently solved with quantum annealing for small problem instances (Sec. 4), see Fig. 1 for an example of 3D shape registration. We call our method *Quantum Graph Matching* (QGM). It opens up new opportunities for multiple problem types in 3D computer vision, with a potentially high impact for the future generation of quantum annealers. We impose the permutation matrix constraints not in a post-processing step (*e.g.*, by projecting a relaxed solution to the space of permutation matrices), but directly through problem weights between the qubits, which leads to high probabilities of measuring solutions corresponding to valid permutation matrices.
- The probability to obtain globally-optimal solutions depends on 1) how the problem is mapped to a QUBO and 2) what is the spectral gap of the mapping. We thus propose different approaches to map our formulation to a QUBO and perform spectral gap analysis (Sec. 5). Our new way to impose soft permutation matrix constraints is controllable by a parameter λ per each row of the permutation matrix.
- Numerical verification in simulated experiments as well as on a real AQC-er, which shows that the proposed meth-

ods effectively increase the success rate of solving combinatorial optimization problems with permutation matrix constraints (Sec. 5). Despite several works related to machine learning and image classification [51, 2, 56, 2], where experiments on quantum hardware were performed, and some theoretical works in quantum computing for computer vision [52, 53, 66, 20, 42], we are not aware of previous AQC-ing experiments for image and shape matching.

In the rest of the paper, Sec. 2 discusses the foundations of modern adiabatic quantum computing. Sec. 3 reviews related works. Three variants of our QGM based on minimization of quadratic objectives over permutation matrices with AQC-ing are introduced in Sec. 4. Sec. 5 elaborates on numerical results with simulated and real data, followed by a Discussion 6 and Conclusion 7.

2. Modern Adiabatic Quantum Computing

We next review the basics of AQC-ing. We suppose that the reader is familiar with general concepts of linear and tensor algebra, calculus as well as operator theory, and no preliminary knowledge of quantum physics is assumed.

State Space of Quantum Systems. While the state of a quantum-mechanical system is generally described by a normalized vector in a Hilbert-space \mathcal{H} , a simple example for a quantum mechanical property is the spin of electrons, in which case the Hilbert-space \mathcal{H} is the two-dimensional complex Euclidean space, $\mathcal{H} = \mathbb{C}^2$. The corresponding quantum state, *i.e.*, an element of the Hilbert space with norm 1, commonly denoted as $|\psi\rangle \in \mathbb{C}^2$, is called a *qubit*. For vectors $|0\rangle, |1\rangle^2$ that form an orthonormal basis of \mathbb{C}^2 , any such quantum state can be represented as a linear combination of the basis vectors, *i.e.*,

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \quad (2)$$

Quantum mechanics dictates that measuring a qubit in the basis $|0\rangle, |1\rangle$ yields the states $|0\rangle$ and $|1\rangle$ with the probability $|\alpha|^2$ and $|\beta|^2$, respectively. The measurement itself influences the state such that it collapses to either $|\psi\rangle = |0\rangle$ or $|\psi\rangle = |1\rangle$. The state with $|\alpha|^2 = |\beta|^2 = 0.5$ is denoted by $|+\rangle$. The state $|\psi\rangle$ for $\alpha, \beta \neq 0$ is called a *superposition*.

Composite Quantum Systems. Currently, a common approach to realize a quantum computer is to build experiments with multiple qubits, in which case the Hilbert space for representing the overall system is expressed as the tensor product of the individual Hilbert spaces. If one now thinks of multiple spin- $\frac{1}{2}$ particles or n qubits, we obtain states $|\psi\rangle \in \mathcal{H}(\mathbb{C}^2)^{\otimes n} := \mathbb{C}^2 \otimes \dots \otimes \mathbb{C}^2$. Note that if the spins of the electrons were completely independent from each other, we would only need $2n$ many parameters to describe a state.

²we use Dirac notation; $|0\rangle$ and $|1\rangle$ represent an orthonormal basis

The reason we need the 2^n parameters for the state space is that for so called *entangled states* the spins of the individual systems cannot be treated separately.

Time Evolution of Quantum Systems. The time evolution of a quantum system, *i.e.*, description how a quantum state $|\psi\rangle$ changes over time under an external influence, is given by the Schrödinger equation:

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H(t) |\psi(t)\rangle, \quad (3)$$

where i is the imaginary unit, \hbar is the reduced Planck constant, t is the time and $H(t)$ is the so-called *Hamiltonian*, which is a self-adjoint, linear operator on the Hilbert space \mathcal{H} characterizing the energy of the physical system. If the Hamilton operator H in (3) is not time-dependent, and the initial state $|\psi(0)\rangle$ is an eigenvector of H to the eigenvalue $\lambda \in \mathbb{R}$, (3) has the close-form solution

$$|\psi(t)\rangle = \exp\left\{-\frac{i t}{\hbar \lambda}\right\} |\psi(0)\rangle, \quad (4)$$

such that $|\psi(t)\rangle$ remains in the same state modulo a phase.

Adiabatic Quantum Computing. AQC solves quadratic unconstrained binary optimization problems (QUBO) on quantum hardware relying on the *adiabatic quantum theorem* [14]. Assume one faces the following QUBO:

$$\arg \min_{\mathbf{s} \in \{-1,1\}^n} \mathbf{s}^T \mathbf{Q} \mathbf{s} + \mathbf{q}^T \mathbf{s}, \quad (5)$$

with an $n \times n$ -matrix \mathbf{Q} and an n -dimensional vector \mathbf{q} . Furthermore, assume one is able to set up a quantum system with a Hamilton operator $H(0) =: H_B$ in a state $|\psi(0)\rangle$ that is an eigenvector to the smallest eigenvalue of H_B , and one constructs a second Hamilton operator H_P whose eigenvector to the smallest eigenvalue is a solution to (5). Then, [14] dictates that the time evolution of (3) with

$$H(t) = \frac{t}{\tau} H_P + \left(1 - \frac{t}{\tau}\right) H_B, \quad (6)$$

for $t \in [0, \tau]$ with large τ to slowly transition between H_B and H_P , remains in the eigenvector with the smallest eigenvalue (provided that it remains unique and non-degenerate), so that measuring the corresponding state $|\psi(\tau)\rangle$ at $t = \tau$ yields the solution to (5). For the cases we investigate in the paper, H_P is a diagonal matrix and for $|\psi(0)\rangle$, all components are equal in the computational basis. To experimentally realize an evolution with the Hamiltonian H_P , one needs to provide couplings between the qubits in form of the matrix \mathbf{Q} and the biases as the vector \mathbf{q} upon (5).

Of course, an infinitely slow transition between the two Hamilton operators, *i.e.*, $\tau \rightarrow \infty$, is impossible (and would eliminate any computational advantages). Luckily, it has been shown that suitable finite choices of τ to remain in the

adiabatic case depend on the *spectral gap* of the problem, *i.e.*, the difference between the smallest and second lowest eigenvalue. For more details on how to estimate τ and spectral gap analysis, see [5, 36].

D-Wave 2000Q. D-Wave 2000Q is an AQC-er with 2048 qubits arranged on a *Chimera* graph, *i.e.*, a 16×16 set of cells with eight qubits each. The qubits inside a cell are fully connected, whereas there are only eight connections to qubits from two other neighbouring cells [25]. The computation always starts with the initialization state $|\psi(0)\rangle = |+\rangle^{\otimes n}$, where n is the number of physical qubits in the problem mapping to the Chimera graph, also known as *minor embedding*. Logical problem qubits are often mapped to chains of physical qubits, due to the restriction on the physical qubits connectivity. The D-Wave compiler performs the minor embedding automatically, and the programmer loads couplings Q and biases q for the target problem. Moreover, additional settings can be provided such as the chain strength, annealing time and the annealing path.

The default duration of an anneal is $20\mu s$. To solve a problem, usually multiple anneals of the same problem instance are performed, since there is no guarantee (due to the physical characteristics of current quantum hardware) that D-Wave will return optimal solutions.

The couplers are realized with loops of niobium. During the annealing, the processor imposes magnetic fluxes on the couplers corresponding to the provided \mathbf{Q} matrix.

3. Related Work

Our approach relates to multiple method classes proposed in the literature so far. In this section, we review the most related categories, *i.e.*, computer vision methods using quantum hardware including image matching and encoding of permutation matrices, as well as classical solutions to graph matching problems.

Quantum Computing in Computer Vision. Promising potential applications of quantum computers range from data fitting [66, 20] and image recognition [52, 51, 56] to training artificial neural networks [2, 42]. Classification-related problems were addressed in [15], to enhance the detection of vegetation zones in aerial images on D-Wave, in [56] to learn facial features and reproduce facial image collections, and in [54] which proposes a dictionary learning method for image classification. [67] covers various low-level quantum image processing topics such as quantum image encryption and segmentation. However, it does not address integration of permutation matrix constraints. To the best of our knowledge, [30] is the first work introducing quantum computing for computer vision at a major vision conference. The method finds a rotation that aligns two point sets and discretizes the space of affine transformation matrices to formulate the problem as (5).

Quantum Computing for Matching Problems. [38] finds

a small image embedded in a larger one by adapting the Grover’s algorithm [32]. Neven *et al.* [52] propose to match two images with AQC-ing. Their mapping requires finding the maximum independent vertex set of a conflict graph without a focus on permutation constraints.

Examples of solving quadratic assignment problems with AQC-ing are [62, 21]. Stollenwerk *et al.* [62] solve the flight gate assignment problem by using qubits to indicate whether resources (*e.g.*, flights) are assigned to facilities (*e.g.*, gates). The resulting matrix is indexed over the resources and facilities, and the quadratic constraints are converted to weights favoring valid configurations in the solution space. Choo [21] investigates several ways to map quadratic assignment problems to QUBO. Their analysis leads mainly to a quadratic penalty terms. They also propose a decomposition approach to mitigate the challenges.

[29, 70] addresses graph isomorphism problem with AQC-ing. In [29], permutation matrix constraints are enforced with a penalty term in the Hamiltonian. They first write the permutation in a table and then represent the entries in the binary system. Zick *et al.* [70] allocate a QUBO variable for every possible pair of vertices of the same degree and solve the problem with AQC-ing, but do not use a binary representation of the permutation table. In this paper, we investigate a different new way of injecting permutation matrix constraints.

Classical Related Methods. In terms of classical approaches to solve matching problems, it is well-known that problems of the form (1) can be solved exactly if the costs are linear ($W = 0$) using the Hungarian method [47] or the auction algorithm [12]. For truly quadratic costs, relaxations based on message passing [46] and Lagrange duality [63, 64], as well as general convex approximations [59, 3, 10, 27, 28, 16] have been proposed. Path following [68, 69, 37] and believe propagation [6] are popular heuristics to tackle (1) directly, and we refer to [58] for an overview of combinatorial approaches to the QAP.

A magnitude of the methods is inspired by quantum mechanics while not being meant for quantum hardware including *quantum cuts* for image segmentation [9], *wave kernel signatures* [8] for mesh matching and *genetic quantum algorithms* [41] for the knapsack problem. None of these methods operates on permutation matrices. A related probabilistic physics-based technique which can be applied to problems on graphs is simulated annealing [44]. In [34], (1) arises for finding the maximal pair of subgraphs between two graphs. The vertices of these graphs represent the interest points in the respective images one wants to match.

4. Minimization of Quadratic Objectives over Permutation Matrices

In this section, we derive three variants of our QGM approach, *i.e.*, baseline variant (Sec 4), the variant with row-

wise penalties which reduces the penalty-based influence on the problem in overall (Sec. 4.2) and the variant with eliminated variables (Sec. 4.3).

4.1. Baseline Quantum Graph Matching

To rephrase (1) in the form of (5), we have to do two conversions: 1) Switch from $x_i \in \{0, 1\}$ to $s_i \in \{-1, 1\}$, and 2) Switch from the constrained problem over permutation matrices to an unconstrained problem. The first one is straightforward by substituting $s_i = 2x_i - 1$, yielding

$$Q = \frac{1}{4}W, \quad q = \frac{1}{2}(\mathbb{1}^T W + \mathbf{c}), \quad (7)$$

between the costs associated with (1) and (5), respectively, up to a constant (assuming without restriction of generality that W is symmetric, and denoting by $\mathbb{1}$ a vector of ones).

The second conversion is less trivial. First of all, note that the set \mathbb{P}_n of permutation matrices can be written as

$$\mathbb{P}_n = \{X \in \mathbb{R}^{n \times n} \mid X_{ij} \in \{0, 1\}, \sum_i X_{ij} = 1 \forall j, \sum_j X_{ij} = 1 \forall i\}. \quad (8)$$

Note that – besides the binary constraint addressed with the above transformation – \mathbb{P}_n is now characterized by linear equality constraints. Thus, after the vectorization of $X \in \mathbb{P}_n$ to $\mathbf{x} \in \mathbb{R}^{n^2}$, one can phrase such constraints in matrix-vector form as $A\mathbf{x} = \mathbf{b}$, for a suitable matrix A and a vector \mathbf{b} . For $n = 2$, as an example, A and \mathbf{b} are given by:

$$A = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}. \quad (9)$$

For higher dimensions, \mathbf{b} still contains only ones. The matrix A for arbitrary n can be described with the formula:

$$A = \begin{pmatrix} \text{Id} \otimes \mathbb{1}^T \\ \mathbb{1}^T \otimes \text{Id} \end{pmatrix}, \quad (10)$$

where Id is the n -dimensional identity matrix. It is well known (see *e.g.*, [45]) that linear equality constraints can be incorporated in QUBO by adding a penalty term to the objective, *i.e.*,

$$\begin{aligned} & \arg \min_{\{\mathbf{x} \in \{0,1\}^{n^2} \mid A\mathbf{x} = \mathbf{b}\}} \mathbf{x}^T W \mathbf{x} + \mathbf{c}^T \mathbf{x} \\ & = \arg \min_{\mathbf{x} \in \{0,1\}^{n^2}} \mathbf{x}^T W \mathbf{x} + \mathbf{c}^T \mathbf{x} + \lambda \|A\mathbf{x} - \mathbf{b}\|^2, \end{aligned} \quad (11)$$

for sufficiently large λ . For the sake of completeness, we provide a proof of this statement, the specific appearance of

A and \mathbf{b} , and the lower bound

$$\lambda > \lambda^0 := \frac{1}{2} \left(\sum_{i,j=1}^N |W_{i,j}| + \sum_{i=1}^N |c_i| \right) \quad (12)$$

for (11) to provably hold in the supplementary material. After rewriting the second line of (11) to

$$\arg \min_{\mathbf{x} \in \{0,1\}^{n^2}} \mathbf{x}^T (W + \lambda A^T A) \mathbf{x} + (\mathbf{c}^T - 2\lambda \mathbf{b}^T A) \mathbf{x}, \quad (13)$$

we arrive at a form of the minimization problem (1) to which quantum computing is directly applicable.

Interestingly, applying quantum computing approaches to (13) often does not yield satisfactory results as we will illustrate in more detail in Sec. 5. We conjecture that this behavior is largely due to a significant reduction of the *spectral gap* of the coupling matrix $Q + \lambda \frac{1}{4} A^T A$ with increasing λ (after the global scale-normalization). As discussed in Sec. 2, the spectral gap *i.e.*, the smallest difference between the smallest and the second lowest eigenvalues of the Hamiltonian (6) over time, has a crucial influence on how slow the transition needs to be and how likely the system leaves its ground state during the time evolution.

The penalty in (13) is not the only reformulation of the constrained QAP as an unconstrained QUBO. We next derive two alternate formulations (Secs. 4.2–4.3) and evaluate them in terms of their resulting spectral gaps as well as behavior on real quantum hardware (Sec. 5).

4.2. The Formulation with a Row-Wise Penalty

As large values of λ seem to cause problems, we aim at reducing the penalty-based changes of the problem as little as possible by considering row-wise penalty parameters λ_i for enforcing each equation $(A\mathbf{x})_i = b_i$ separately, *i.e.*,

$$\min_{\mathbf{x} \in \{0,1\}^{n^2}} \mathbf{x}^T W \mathbf{x} + \mathbf{c}^T \mathbf{x} + \sum_{i=1}^{2n} \lambda_i ((A\mathbf{x})_i - b_i)^2. \quad (14)$$

While a change of variables from x to s along with the choice $\lambda_i = \lambda > \lambda^0$ according to (12) is obviously sufficient and reduces (21) to (13), the challenge is to find small *computable lower bounds* for λ_i , for which the unconstrained problem still *provably* coincides with the constrained problem (1).

Proposition 4.1. *The minimizers of (5) and (21) coincide provided that*

$$\lambda_i > \lambda_i^0 := D_{\mathcal{J}_i} + \frac{1}{2} D_{\{1, \dots, n^2\}}, \quad (15)$$

where \mathcal{J}_i denotes the indices that belong to a column or a row enumerated by i in (21), $D_{\{1, \dots, n^2\}} = D_{\mathcal{J}}$ for $\mathcal{J} = \{1, \dots, n^2\}$ and

$$D_{\mathcal{J}} := \max_{k \in \mathcal{J}} \left(\sum_i |W_{k,i} + W_{i,k}| + |W_{k,k}| + |c_k| \right). \quad (16)$$

4.3. Inserted Formulation with Eliminated Variables

In continuous optimization, linear equality constraints are frequently used to eliminate variables and reduce the complexity of the underlying problem. Even in our discrete setting, the sum-to-one constraints of permutation matrices in (8) naturally allow expressing $X_{1,j} = 1 - \sum_{i=2}^n X_{i,j}$ for all j as well as $X_{i,1} = 1 - \sum_{j=2}^n X_{i,j}$ as long as one ensures that the variables to be replaced remain in $\{0, 1\}$. In other words:

Lemma 4.1. *The set in \mathbb{P}_n of (8) can be written as*

$$\begin{aligned} \mathbb{P}_n = \{ & \begin{pmatrix} 2 - n + \sum_{i,j=2}^n x_{i,j} & 1 - \sum_{i=2}^n x_{i,2} & \dots & 1 - \sum_{i=2}^n x_{i,n} \\ 1 - \sum_{i=2}^n x_{2,i} & x_{2,2} & \dots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 1 - \sum_{i=2}^n x_{n,i} & x_{n,2} & \dots & x_{n,n} \end{pmatrix} \\ & \in \mathbb{R}^{n \times n} \mid \forall i \in \{2, \dots, n\}, j \in \{2, \dots, n\} x_{i,j} \in \{0, 1\}, \\ & \sum_{i,j=2}^n x_{i,j} \in \{n-2, n-1\}, \\ & \forall j, i, k \in \{2, \dots, n\}, i \neq k, x_{j,i} x_{j,k} = 0 \wedge x_{i,j} x_{k,j} = 0 \}. \end{aligned}$$

This reformulation of the constraint set allows us to reduce the amount of needed qubits from n^2 to $(n-1)^2$ by incorporating all remaining constraints via penalties as stated in the following proposition.

Proposition 4.2. *The problem (1) has the same solution as the following QUBO for sufficiently large λ_1 and λ_2 :*

$$\begin{aligned} \min_{\mathbf{x} \in \{0,1\}^{(n-1)^2}} \mathbf{x}^T \tilde{W} \mathbf{x} + \tilde{\mathbf{c}}^T \mathbf{x} + \lambda_1 \sum_{j,k=0}^{(n-1)^2} \chi(j,k) x_j x_k \\ + \lambda_2 \left(\sum_{i=1}^{(n-1)^2} x_i - (n-1) \right) \left(\sum_{i=1}^{(n-1)^2} x_i - (n-2) \right) \\ \chi(j,k) := \begin{cases} 1 & \text{for } \lfloor \frac{j}{n-1} \rfloor = \lfloor \frac{k}{n-1} \rfloor \\ & \forall j - (n-1) \lfloor \frac{j}{n-1} \rfloor = k - (n-1) \lfloor \frac{k}{n-1} \rfloor, \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

where the specific definitions of the reduced $(n-1) \times (n-1)$ matrix \tilde{W} , the vector $\tilde{\mathbf{c}} \in \mathbb{R}^{n-1}$ and the precise lower bounds for λ_1 and λ_2 are provided in the supplement.

The reduction in the number of required qubits is a clear advantage of the above *inserted* formulation, especially in the current early stage of AQC-ers, when the number of supported logical qubits is limited [25]. On the other hand, the asymptotic difference between the n^2 and $(n-1)^2$ (of the inserted formulation) vanishes, such that their spectral gaps, and, more importantly, their performance on real AQC-ers are the deciding factors we evaluate in the next section.

5. Experimental Results

We next test our methods for quantum graph matching on random problem instances in Sec. 5.1 and real data in Sec. 5.3. We also analyze in Sec. 5.2 the influence of experimental errors in the coupling parameters to the measured results. All in all, we investigate the following questions:

- Do the three equivalent reformulations of (1), *i.e.*, the baseline, the row-wise penalty and the inserted formulations lead to different spectral gaps?
- Do larger spectral gaps improve AQC-ing results?
- Do quantum computing approaches provide advantages over classical relaxation methods for solving (1)?
- Do AQC-ers – as well as their simulations, *i.e.*, numerical solutions of (3) – solve our problem?

We prepare the weight matrix W on a classical computer, and access D-Wave 2000Q remotely in the cloud via *Ocean Tools Leap 2* [26]. In all experiments, we report the solution distributions and analyse the probability and the energy of the most frequent solution. The ground-truth permutations are calculated by brute force and compared qualitatively with the expected outcomes on real data. Minor embedding of QGM results in a fully connected logical graph of qubits. For $n = 3$, the number of physical qubits in the minor embedding amounts to ≈ 27 .

5.1. Evaluation on Random Problem Instances

As the first step, we create random problems of the form (1) by drawing all entries in W and c uniformly from $[-1, 1]$. Those are representative for a broad class of problems arising in 3D vision such as mesh and point cloud matching under isometry, and rigid point set alignment. We compute the spectral gap for each of our three formulations via an iterative Lanczos method [48], since this calculation requires eigenvalue decomposition of large quadratic matrices with the dimension 2^{n^2} . Fig. 2 shows the spectral gaps of all three formulations averaged over ten different problem instances in dimensions $n = 3$ and $n = 4$, when the strength of the overall penalty is varied by a global scale factor (x-axis). A scale factor of 1 refers to the setting where the constrained and unconstrained problems are provably equivalent according to our derivations in Sec. 4. As we see, the spectral gap reduces significantly for all formulations and instances with the increasing strength of the constraint penalty. Among the three formulations, the row-wise penalty consistently results in the largest spectral gap which we expect to be favorable for real and simulated AQC-ing.

Fig. 3 shows the different results of our three reformulations run on D-Wave, and the recent relaxation-based algorithm [10] for $n \in \{3, 4\}$ over ten instances of (1). For each instance, we compute the global optimizer by brute force and normalize (shift) the energy by $-f_{opt}$ for each method.




Figure 2: Illustrating the spectral gap for our baseline, the row-wise penalty and inserted formulation as a function of the global scaling of the penalty term for $n = 3$ (left) and $n = 4$ (right). As we see, the row-wise penalty leads to the largest spectral gap which is favorable for AQC-ing.




Figure 3: Comparing quantum computing on DWave 2000Q for our three reformulations and the relaxation-based method DS* [10] for computing 3×3 (left) and 4×4 permutations (right) on ten random problem instances.

For the AQC-ing approach, we use 500 annealing cycles on the D-Wave 2000Q and select the most frequent solution among all annealing cycles as our final solution. On D-Wave, logical problem qubits are mapped to multiple physical qubits by chaining. For $n = 4$, we set the chain strength as the maximal element of the q vector in (7). This leads to an automatic adaption of the chain strength to the problem instance. For $n = 3$, this does not lead to further improvements over empirical fixed values of the chain strength. Here, we have the chain strengths 20, 30 and 35 for the inserted, the row-wise and the baseline method, respectively. Up to a $100\mu s$ break, we insert in the middle of the annealing path, the annealing time τ is $80\mu s$ for $n = 4$. For $n = 3$, $\tau = 20\mu s$ already yields desired result.

The simulation with QuTiP [39] showed that increasing the τ yields a narrower distribution. For $n = 3$, the simulation can find the correct optimum as the highest peak regardless of the method variant. The situation changes for $n = 4$ and the simulation with the *baseline* formulation does not yield satisfactory results. The method with inserted constraints, on the other hand, finds the minimum every time.

Table 1 summarizes the mean normalized energies over ten instances for $n = 3$ and $n = 4$, with the averages computed so that 0 is the best achievable value. Note that if

n	inserted	baseline	row-wise	DS* [10]	worst permutation
3	1.49	2.12	2.27	0.85	4.62
4	5.68	7.37	7.40	0.43	10.23

Table 1: Average energies for $n = 3, 4$.

AQC-ing yields a state that is not a valid permutation, the worst permutation is used for averaging. Compared to the simulation, where we were always able to find the global solution with the inserted or the row-wise method for large enough τ , AQC-ing on quantum hardware is significantly more difficult. As we can see in Fig. 3-(left), the most probable states have the lowest energy for the inserted and the row-wise method, and the results are comparable to DS* [10] in $n = 3$. For $n = 4$ the DS* performs clearly the best.

Interestingly, for $n = 3$, the row-wise penalties were able to substantiate their advantages with respect to the spectral gap in practice, yielding several problem instances on which it outperformed DS* and being on par with it on average. While the dimension of the underlying problems is of course small (for $n = 3$, there only exist six different permutation matrices), we consider these results to be very promising, as they illustrate the potential of AQC-ing beyond pure simulations with future generation of AQC-ers.

5.2. Ablative Study on the Coupling Parameters

On D-Wave, the input Q and q values are scaled so that $-1 < Q_{i,j} < 1$ and $-2 < q_i < 2$. Additionally, there are multiple further error sources for the couplings and biases [23]. Therefore, it can happen that the input $Q_{i,j} = 0.5$ is realized as 0.51, for example. As we saw in Sec. 4, Eqs. (11) and (21), the Hamiltonian has a part that depends on the problem instance W, c directly and includes the regularization term which only depends on W, c in the penalty parameters. The regularization part yields the same energy for every valid permutation. Let Q_{reg} and q_{reg} be the couplings and biases that stem from the regularization term.

If permutations only marginally differ in the qubit couplings, it is difficult to distinguish between them during a quantum anneal. This suggests that we have to choose the penalties as small as possible while already delimiting from random binary strings. Assuming that all values in W are of the same order of magnitude, λ would be n^4 times bigger than W values for the baseline method. Our calculations also confirm that for $n = 4$, the maximal value of c that is due to the penalty, is 500 times larger than the maximal value of c in the problem itself.

One easy way to have a larger (more beneficial) ratio between the range of the original W values and the range value in the penalty term is to consider problems with W entries randomly set to zero. The connectivity should stay the same, due to the regularization. For $n = 4$, we either

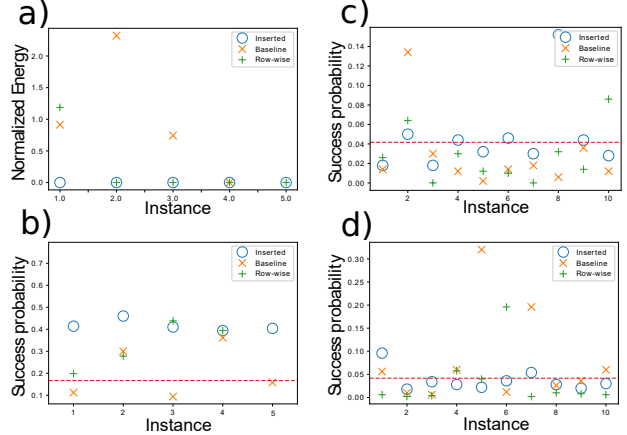


Figure 4: Problem instances, where multiple elements are randomly chosen to be zero. In **a)** and **b)**, the problem instances are for $n = 3$. Here, less than half of the entries in W and c are randomly set zero. **c)** and **d)** show the success probability for $n = 4$. For **c)**, exactly half of the entries are chosen zero and for **d)**, exactly 3/4 are chosen zero.

set a half or $\frac{3}{4}$ of the entries of the problem instance to zero. The results in Fig. 4 show that for the case $n = 4$, random guessing is still better most of the time.

5.3. Experiments with Real Data

We now solve several real problems of matching two graphs originating from two 3D shapes deformed in a near-isometric way. Assuming that we have an edge-weight or a notion of distance $d_1(v_i^1, v_k^1)$ between nodes v_i^1 and v_k^1 of our first graph, and, similarly, a distance $d_2(v_j^2, v_l^2)$ between nodes v_j^2 and v_l^2 of our second graph, a common choice for the quadratic costs in (1) is

$$\sum_{i,j,k,l} X_{i,j} X_{k,l} |d_1(v_i^1, v_k^1) - d_2(v_j^2, v_l^2)|. \quad (17)$$

The above formulation assures that if v_i^1 is matched to v_j^2 and v_k^1 is matched to v_l^2 , then the distance between v_i^1 and v_k^1 is similar to the distance between v_j^2 and v_l^2 .

We consider different instances of (17): In our first shape matching example, we aim to identify the corresponding points on two differently deformed instances of the same shape. In this case, d_1 and d_2 denote the geodesic distance on the respective shapes, and we introduce an additional linear term based on the Euclidean distance of the points after an initial rigid registration following [19]. Fig. 1 visualizes the matching of the two shapes along with the histogram we obtained from 500 runs of this problem instance on the D-Wave 2000Q. See Fig. 5 for the performance comparison of different QGM variants. We observe that the results with real 3D shapes match coarsely the distributions of solutions to experiments with random instances in Sec. 5.1.




Figure 5: Performance of the quantum annealer in the shape matching problem for $n = 3$.



Figure 6: Example of image matching for $n = 4$.




Figure 7: The probability to measure the optimum in one run compared to random guessing of permutations for $n = 3$ (left) and $n = 4$ (right). The problem instances are dense, *without* zero entries.

Fig. 6 illustrates another application of costs (17) for sorting images according to their mean color as considered in [27], see Fig. 6 for the visualization. In this setting, d_1 represents the difference in the mean image color, and d_2 is the Euclidean distance on the 2×2 image grid. The mean image color is sampled from the middle region of each image. As we can see, after the sorting on the right, the images are ordered from predominant blue on the top left (mountains with a lake), to the predominant yellow (desert with the sunset). This problem, likewise, has multiple possible minimizers, and the quantum annealer finds the optimum more often than the other (non-optimal) states.

6. Further Observations and Limitations

We were able to successfully confirm in our experiments the validity of our QGM approach for $n = 3$. On a wide range of randomly generated problem instances, our method outputs optimal valid permutations with high probability. On the current D-Wave hardware, the results for $n = 4$ are significantly worse, *i.e.*, neither do we obtain a conclusive best choice between the three different formulations, nor

do we get near the performance of the DS^* heuristic [10]. Fig. 7 shows the probabilities of obtaining the optimal solution based on 500 anneals per instance. For most instances, the different formulations performed even worse than random guessing, because they can also output vectors that do not correspond to permutations. These results cannot be due to chain breaks, because the D-Wave compiler did not report any. Another interesting quantity is the length of the chain. For $n = 2$, our chains have maximal length of two, whereas for $n = 3$, there are also chains of length three and four. For $n = 4$ the chains have length five and six.

Remark on IBM Q Experience. We also run tests via the IBM Q Experience [1]. They offer not an AQC-er but a circuit-based universal quantum computer, where the algorithm can be viewed as a circuit consisting of so-called *quantum gates*. To test our method on it, we discretize our Hamiltonian with constant time steps and apply trotterization [55] for the decomposition into quantum gates. For higher-dimensional problems, the number of gates rises, which leads to errors. We thus are only able to successfully realize an adiabatic transition for two qubits. By inserting two equality restrictions, we can optimize over 2×2 permutations. We provide more details on this alternative direction in the supplementary material.

7. Conclusions

The quadratic assignment problem which we tackle in this paper is a hard problem to be mapped and solved on a real quantum annealer, and we make a further successful step to solve it on real AQC-ers. The difficulties come from the high demand in connectivity between the qubits and the strong scaling in the qubit number ($n \log(n)$ in the best case). We achieve expected results for 3×3 permutations and demonstrate the potential of the new quantum method for solving quadratic assignment problems in 3D vision, whereas our formulation can be applied to problems in different dimensions which is demonstrated experimentally. The spectral gap analysis shows that it greatly matters how the problem is mapped to a QUBO and subsequently solved by an AQC-er. We expect a significant impact of the future quantum technology on 3D vision, while the accessibility of modern AQC-ers for research purposes allows us to lay the foundation already today. In future work, we will investigate other structures such as point clouds with outliers, more nodes and partial matches.

Supplementary Material. Our supplement contains the proofs of Lemma 4.1 and Prop. 4.1 as well as more experimental results.

Acknowledgement. This work was partially supported by the ERC Consolidator Grant 4DReply (770784).

References

- [1] IBM Q experience. <https://quantum-computing.ibm.com/> <http://www.research.ibm.com/quantum>. Accessed: 2020 February. 1, 8, 14
- [2] S. H. Adachi and M. P. Henderson. Application of quantum annealing to training of deep neural networks. *arXiv preprint arXiv:1510.06356*, 2015. 2, 3
- [3] Y. Aflalo, A. Bronstein, and R. Kimmel. On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, 112(10):2942–2947, 2015. 4
- [4] D. Aharonov et al. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM review*, 50(4):755–787, 2008. 14
- [5] M. H. S. Amin. Consistency of the adiabatic theorem. *Phys. Rev. Lett.*, 102:220401, Jun 2009. 3
- [6] D. Anguelov et al. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in Neural Information Processing Systems*, 2004. 4
- [7] F. Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 2019. 1
- [8] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *International Conference on Computer Vision Workshops*, 2011. 4
- [9] C. Aytekin, S. Kiranyaz, and M. Gabbouj. Automatic object segmentation by quantum cuts. In *International Conference on Pattern Recognition*, 2014. 4
- [10] F. Bernard, C. Theobalt, and M. Moeller. Ds*: Tighter lifting-free convex relaxations for quadratic matching problems. In *Computer Vision and Pattern Recognition*, 2018. 4, 6, 7, 8
- [11] E. Bernstein and U. Vazirani. Quantum complexity theory. In *Annual ACM Symposium on Theory of Computing*, 1993. 1
- [12] D. P. Bertsekas. A distributed algorithm for the assignment problem. Laboratory for Information and Decision Systems Working Paper, M.I.T., 1979. 4
- [13] F. Bogo, J. Romero, M. Loper, and M. J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Computer Vision and Pattern Recognition*, 2014. 1
- [14] M. Born and V. Fock. Beweis des adiabatensatzes. *Zeitschrift für Physik*, 51(3):165–180, 1928. 3
- [15] E. Boyda et al. Deploying a quantum annealing processor to detect tree cover in aerial imagery of california. *PLOS ONE*, 12(2):1–22, 02 2017. 3
- [16] O. Burghard and R. Klein. Efficient lifted relaxations of the quadratic assignment problem. In *Proceedings of the Conference on Vision, Modeling and Visualization*, 2017. 4
- [17] C. Cao, J. Xue, N. Shannon, and R. Joynt. Speedup of the quantum adiabatic algorithm using delocalization catalysis. *arXiv preprint arXiv:2007.11212*, 2020. 14
- [18] D. Castelvecchi. Quantum computers ready to leap out of the lab in 2017. *Nature News*, 541(7635):9, 2017. 2
- [19] Q. Chen and V. Koltun. Robust nonrigid registration by convex optimization. In *International Conference on Computer Vision*, 2015. 7
- [20] T.-J. Chin, D. Suter, S.-F. Chng, and J. Quach. Quantum robust fitting. *arXiv preprint arXiv:2006.06986*, 2020. 2, 3
- [21] J. Choo. Investigating the feasibility of solving the quadratic assignment problem using quantum computing. <https://joelchoo.com/research/qc-paper.pdf>. Accessed: 2020 July. 4
- [22] P. J. Coles et al. Quantum algorithm implementations for beginners. *arXiv preprint arXiv:1804.03719*, 2018. 1
- [23] Ice: Dynamic ranges in h and j values. https://docs.dwavesys.com/docs/latest/c_qpu_1.html. Accessed: 2020 July. 7
- [24] Performance tuning for d-wave quantum processors. https://www.dwavesys.com/sites/default/files/2_Wed_Am_PerfTips.pdf. Accessed: 2020 July. 18
- [25] D-Wave Systems. Technical Description of the D-Wave Quantum Processing Unit. https://docs.dwavesys.com/docs/latest/doc_qpu.html, 2019. online; accessed 5.9.2019 and 3.2.2020. 1, 3, 5
- [26] D-Wave Systems, Inc. D-Wave Leap: Take the Leap. <https://cloud.dwavesys.com/leap/>, 2020. online; accessed on the 8 April 2020. 6
- [27] N. Dym, H. Maron, and Y. Lipman. Ds++: A flexible, scalable and provably tight relaxation for matching problems. *ACM Trans. Graph.*, 36(6), 2017. 4, 8
- [28] F. Fogel, R. Jenatton, F. Bach, and A. d’Aspremont. Convex relaxations for permutation problems. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2013. 4
- [29] F. Gaitan and L. Clark. Graph isomorphism and adiabatic quantum computing. *Physical Review A*, 89(2):022342, 2014. 4
- [30] V. Golyanik and C. Theobalt. A quantum computational approach to correspondence problems on point sets. In *Computer Vision and Pattern Recognition*, 2020. 1, 3
- [31] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Annual ACM Symposium on Theory of Computing*, 1996. 1
- [32] L. K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical review letters*, 79(2):325, 1997. 4
- [33] M. Hebenstreit, D. Alsina, J. I. Latorre, and B. Kraus. Compressed quantum computation using a remote five-qubit quantum computer. *Phys. Rev. A*, 95:052339, May 2017. 14
- [34] L. Hérault, R. Horaud, F. Veillon, and J.-J. Niez. Symbolic image matching by simulated annealing. In *Proceedings of the British Machine Vision Conference*, 1990. 4
- [35] M. Howard, J. Wallman, V. Veitch, and J. Emerson. Contextuality supplies the ‘magic’ for quantum computation. *Nature*, 2014. 1
- [36] M. Jarret and S. P. Jordan. Adiabatic optimization without local minima. *arXiv preprint arXiv:1405.7552*, 2014. 3
- [37] B. Jiang, J. Tang, C. Ding, and B. Luo. Binary constraint preserving graph matching. In *Computer Vision and Pattern Recognition*, 2017. 4
- [38] N. Jiang, Y. Dang, and J. Wang. Quantum image matching. *Quantum Information Processing*, 15(9):3543–3572, 2016. 3

- [39] J. R. Johansson, P. Nation, and F. Nori. Qutip: An open-source python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 183(8):1760–1772, 2012. 6
- [40] I. Kezurer, S. Z. Kovalsky, R. Basri, and Y. Lipman. Tight relaxation of quadratic matching. In *Eurographics Symposium on Geometry Processing*, 2015. 1
- [41] K.H. Han and J.H. Kim. Genetic quantum algorithm and its application to combinatorial optimization problem. In *Congress on Evolutionary Computation*, 2000. 4
- [42] A. Khoshaman et al. Quantum variational autoencoder. *Quantum Science and Technology*, 4(1), 2018. 2, 3
- [43] A. D. King and C. C. McGeoch. Algorithm engineering for a quantum annealing platform. *arXiv preprint arXiv:1410.2628*, 2014. 18
- [44] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. 4
- [45] G. Kochenberger et al. The unconstrained binary quadratic programming problem: a survey. *Journal of Combinatorial Optimization*, 28(1):58–81, 2014. 4
- [46] V. Kolmogorov. A new look at reweighted message passing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(5):919–930, 2015. 4
- [47] H. W. Kuhn and B. Yaw. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955. 4
- [48] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, 1950. 6
- [49] H. Maron, N. Dym, I. Kezurer, S. Kovalsky, and Y. Lipman. Point registration via efficient convex relaxation. *ACM Trans. Graph.*, 35(4), 2016. 1
- [50] F. Neukart et al. Traffic flow optimization using a quantum annealer. *Frontiers in ICT*, 4:29, 2017. 1
- [51] H. Neven, V. S. Denchev, G. Rose, and W. G. Macready. Qboost: Large scale classifier training with adiabatic quantum optimization. In *Asian Conference on Machine Learning (ACML)*, volume 25, 2012. 2, 3
- [52] H. Neven, G. Rose, and W. G. Macready. Image recognition with an adiabatic quantum computer i. mapping to quadratic unconstrained binary optimization. *arXiv preprint arXiv:0804.4457*, 2008. 2, 3, 4
- [53] N. T. Nguyen and G. T. Kenyon. Image classification using quantum inference on the d-wave 2x. In *International Conference on Rebooting Computing*, 2018. 2
- [54] N. T. Nguyen, G. T. Kenyon, and B. Yoon. A regression algorithm for accelerated lattice qcd that exploits sparse inference on the d-wave quantum annealer. *Scientific Reports*, 2020. 3
- [55] M. A. Nielsen and I. Chuang. Quantum computation and quantum information, 2002. 1, 8, 14
- [56] D. O’Malley, V. V. Vesselinov, B. S. Alexandrov, and L. B. Alexandrov. Nonnegative/binary matrix factorization with a d-wave quantum annealer. *PLOS ONE*, 13(12):1–12, 12 2018. 2, 3
- [57] J. Palmer. Quantum devices: Here, there and everywhere. *The Economist, Technology Quarterly, Quarter*, 2, 2017. 2
- [58] P. M. Pardalos, F. Rendl, and H. Wolkowicz. The quadratic assignment problem: A survey and recent developments. In *DIMACS Workshop on Quadratic Assignment Problems*, volume 16. American Mathematical Society, 1994. 4
- [59] C. Schellewald and C. Schnörr. Probabilistic subgraph matching based on convex relaxation. In *Energy Minimization Methods in Computer Vision and Pattern Recognition*, 2005. 4
- [60] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, Nov 1994. 1
- [61] D. R. Simon. On the power of quantum computation. In *Annual Symposium on Foundations of Computer Science*, 1994. 1
- [62] T. Stollenwerk, E. Lobe, and M. Jung. Flight gate assignment with a quantum annealer. In *International Workshop on Quantum Technology and Optimization Problems*. Springer, 2019. 4
- [63] P. Swoboda, C. Rother, H. Abu Alhajja, D. Kainmuller, and B. Savchynsky. A study of lagrangean decompositions and dual ascent solvers for graph matching. In *Computer Vision and Pattern Recognition*, 2017. 4
- [64] L. Torresani, V. Kolmogorov, and C. Rother. A dual decomposition approach to feature correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(2):259–271, 2013. 4
- [65] W. van Dam and G. Seroussi. Efficient quantum algorithms for estimating gauss sums. *arXiv preprint arXiv:quant-ph/0207131*, 2002. 1
- [66] N. Wiebe, D. Braun, and S. Lloyd. Quantum algorithm for data fitting. *Phys. Rev. Lett.*, 109:050505, 2012. 2, 3
- [67] F. Yan and S. E. Venegas-Andraca. Quantum image processing. 3
- [68] M. Zaslavskiy, F. Bach, and J. Vert. A path following algorithm for the graph matching problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2227–2242, 2009. 4
- [69] F. Zhou and F. De la Torre. Factorized graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9):1774–1789, 2016. 4
- [70] K. M. Zick, O. Shehab, and M. French. Experimental quantum annealing: case study involving the graph isomorphism problem. *Scientific reports*, 5(1):1–11, 2015. 4

Adiabatic Quantum Graph Matching with Permutation Matrix Constraints

—Supplementary Material—

This supplementary material provides details on the derivations of the Quantum Graph Matching (QGM) approach as well as more experimental results. The proofs of Lemma (4.1) and Prop. (4.1) of the main matter can be found in Sec. A. Sec. B complements the description of our experiments with the quantum computers, accessible via IBM Quantum Experience. Next, we show additional results of QGM for point cloud alignment in Sec. C and nearly isometric non-rigid shape registration in Sec. D, followed by a discussion of the minor embedding and negative results in Secs. E and F, respectively. Note that D-Wave provides one minute of QPU time on their annealers for research purposes per month. We obtain three minutes in total during the current project and use 80% of this time for our experiments. One minute of time on the quantum annealer can suffice to run hundreds of problems.

A. Proofs for Section 4

In order to ensure that we can incorporate a constraint of the form $A\mathbf{x} = \mathbf{b}$ as a penalty, the corresponding penalty parameter λ has to be chosen such that

$$\min_{\mathbf{x} \in \{0,1\}^n, A\mathbf{x}=\mathbf{b}} \mathbf{x}^T W \mathbf{x} < \min_{\mathbf{x} \in \{0,1\}^n, A\mathbf{x} \neq \mathbf{b}} \mathbf{x}^T W \mathbf{x} + \lambda \|A\mathbf{x} - \mathbf{b}\|^2. \quad (18)$$

A naive estimate for a sufficiently large λ is obtained by using

$$\min_{\mathbf{x} \in \{0,1\}^n, A\mathbf{x}=\mathbf{b}} \mathbf{x}^T W \mathbf{x} \leq \sum_{i,j} \max\{0, W_{i,j}\} \quad (19)$$

and

$$\begin{aligned} \min_{\mathbf{x} \in \{0,1\}^n, A\mathbf{x} \neq \mathbf{b}} \mathbf{x}^T W \mathbf{x} + \lambda \|A\mathbf{x} - \mathbf{b}\|^2 &\geq \min_{\mathbf{x} \in \{0,1\}^n, A\mathbf{x} \neq \mathbf{b}} \mathbf{x}^T W \mathbf{x} + \lambda \min_{\mathbf{x} \in \{0,1\}^n, A\mathbf{x} \neq \mathbf{b}} \|A\mathbf{x} - \mathbf{b}\|^2 \\ &\geq \min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^T W \mathbf{x} + \lambda \underbrace{\min_{\mathbf{s} \in \{-1,1\}^n, A\mathbf{s} \neq \mathbf{b}} \|A\mathbf{s} - \mathbf{b}\|^2}_{=:d} \geq - \sum_{i,j} \max\{0, -W_{i,j}\} + \lambda d. \end{aligned} \quad (20)$$

Thus a sufficient upper bound for λ is given by

$$\lambda > \frac{1}{d} \sum_{i,j} (\max\{0, -W_{i,j}\} + \max\{0, W_{i,j}\}) = \frac{1}{d} \sum_{i,j} |W_{i,j}|.$$

The influence of the linear term $\mathbf{c}^T \mathbf{x}$ can be calculated in the same way. Alternatively \mathbf{c} can be added to the diagonal of W , since $x_i^2 = x_i$ for $x_i \in \{0, 1\}$. The fact that $d = 2$ can be proven by the following reasoning. Consider that there is a violation for the rows so that one row does not sum up to one and $\sum_{k \text{ belongs to a row}} (Ax - b)_k^2 = 1$. Since the sum over all rows of a matrix is the same as the sum over all columns, one constraint for the columns is violated as well. Since the same argument can be made if we switch rows and columns d has to be 2.

A.1. Proof of Proposition (4.1)

Better estimates can be obtained by exploiting more about the specific structure of the constraint. We aim to formulate the optimization problem as an unconstrained optimization problem with row-wise penalty parameters:

$$\min_{\mathbf{x} \in \{0,1\}^{n^2}} \mathbf{x}^T W \mathbf{x} + \mathbf{c}^T \mathbf{x} + \sum_{i=1}^{2n} \lambda_i ((A\mathbf{x})_i - b_i)^2. \quad (21)$$

One can solve the following optimization problems

$$D_{\mathcal{J}_j} = \max_{k \in \mathcal{J}_j} \left(\sum_i |(W_{k,i} + W_{i,k})| + |W_{k,k}| + |c_k| \right) \quad (22)$$

for every row or column \mathcal{J}_j .

The following bound for λ_j can be used to obtain an unconstrained optimization problem as in (21):

$$\lambda_j = \max_{k \in \mathcal{C}_j} \left(\sum_i |(W_{k,i} + W_{i,k})| + |W_{k,k}| + |c_k| \right) + \frac{1}{2} \left(\max_k \sum_i |(W_{k,i} + W_{i,k})| + |W_{k,k}| + |c_k| \right) = D_{\mathcal{C}_j} + \frac{1}{2} D_{\{1, \dots, n^2\}}. \quad (23)$$

Proof. Let x be an arbitrary element in $\{0, 1\}^{n^2}$. We want to show that for the above choice of λ_j :

$$\exists \mathbf{p} \in \text{vec}(\mathbb{P}_n) : \mathbf{p}^T W \mathbf{p} + \mathbf{c}^T \mathbf{p} \leq \mathbf{x}^T W \mathbf{x} + \mathbf{c}^T \mathbf{x} + \sum_{i=1}^{2n} \lambda_i ((A\mathbf{x})_i - b_i)^2. \quad (24)$$

For the matrix X , which fulfills $x = \text{vec}(X)$, we can construct sets with the property:

$$\begin{aligned} (i, j) \in I \Rightarrow X_{i,j} = 1 \quad \wedge \quad \forall k \in \{1, \dots, n\} \setminus \{j\} (i, k) \notin I, \\ \wedge \quad \forall k \in \{1, \dots, n\} \setminus \{j\} (k, j) \notin I. \end{aligned} \quad (25)$$

We name one of these sets that has the maximal possible number of elements I_{\max} . The permutation matrix P with $\mathbf{p} = \text{vec}(P)$ that we want to construct can be any permutation with ones placed at the positions in I_{\max} . We can get from X to P by first erasing all ones that are not in the positions I_{\max} and then adding $n - |I_{\max}|$ ones.

Consider the set of matrices $(X^{(k)})_{0 \leq k \leq H}$ with:

$$\begin{aligned} X^{(0)} &= X, \\ X^{(H)} &= P, \\ \|X^{(k)} - X^{(k-1)}\| &= 1. \end{aligned} \quad (26)$$

These matrices can be constructed if we start from P and erase successively all ones that are not in I_{\max} . After that we can insert ones that have an index in common with an element in I_{\max} . This set we call B .

Inserting or erasing a one at the j -th column yields maximally to an energy difference of

$$D_{\mathcal{C}_j} = \max_{k \in \mathcal{C}_j} \left(\sum_i |(W_{k,i} + W_{i,k})| + |W_{k,k}| + |c_k| \right), \quad (27)$$

where \mathcal{C}_j describe the indices that belong to the j -th column and analogously \mathcal{R}_j describe the indices that belong to the j -th row. We define f as:

$$f(Y) = \mathbf{y}^T W \mathbf{y} + \mathbf{c}^T \mathbf{y}, \quad (28)$$

with $\mathbf{y} = \text{vec}(Y)$. To prove (24), we use the principle of a telescope sum:

$$\begin{aligned} f(P) - f(X) &= \sum_{k=0}^{H-1} f(X^{(k+1)}) - f(X^{(k)}) \leq \sum_{k=0}^{H-1} \left| f(X^{(k+1)}) - f(X^{(k)}) \right| \\ &= \sum_{\{k \in \{0, \dots, H-1\} \mid \text{pos}(X^{(k+1)} - X^{(k)}) \in B\}} \left| f(X^{(k+1)}) - f(X^{(k)}) \right| \\ &+ \sum_{\{k \in \{0, \dots, H-1\} \mid \text{pos}(X^{(k+1)} - X^{(k)}) \notin B\}} \left| f(X^{(k+1)}) - f(X^{(k)}) \right|. \end{aligned} \quad (29)$$

For the second sum we want to use:

$$\left| f(X^{(k+1)}) - f(X^{(k)}) \right| \leq D_{\{1, \dots, n^2\}}. \quad (30)$$

In the first sum we can make use of the specific column or row and use $D_{\mathcal{C}_j}$. This yields the following ansatz for $\lambda_{\mathcal{R}_j}$, which is the λ_k that belongs to the j -th row:

$$\lambda_{\mathcal{R}_j} = \alpha D_{\mathcal{R}_j} + \beta D_{\{1, \dots, n^2\}}. \quad (31)$$

To obtain the constant β we have to calculate the right side of the inequality:

$$\beta \geq \max_{\{\mathbf{x} \in \{0, 1\}^{n^2} \mid A\mathbf{x} \neq \mathbf{b}\}} \frac{|n - I_{\max}|}{\sum_j |A\mathbf{x}_j - b_j|}. \quad (32)$$

If we have a row, where no element of I_{\max} is present, there can be other ones placed there that share a column with a position in I_{\max} . If there is no one in that column, then $|Ax_j - b_j| \neq 0$ for the corresponding j . If there are ones then $\sum_j |Ax_j - b_j|$ also increases, since the ones are in places, where they share a column with a position in I_{\max} . This yields to the inequality:

$$\beta \geq \max_{\{\mathbf{x} \in \{0,1\}^{n^2} | A\mathbf{x} \neq \mathbf{b}\}} \frac{|n - I_{\max}|}{2 \cdot |n - I_{\max}|} = \frac{1}{2}. \quad (33)$$

To get an estimate for α we now consider a row that contains an element of I_{\max} . In the case, that there is only a single one we do not have to delete other ones. Every additional one increases $|Ax_j - b_j|$ also by one. Therefore:

$$\alpha \geq \max_{\{j, \mathbf{x} \in \{0,1\}^{n^2} | (A\mathbf{x})_j \neq b_j, \mathcal{J}_j \cap \text{vec}(I_{\max}) \neq \emptyset\}} \frac{|\text{vec}(B \setminus I_{\max}) \cap \mathcal{J}_j|}{|(A\mathbf{x})_j - b_j|} = 1. \quad (34)$$

A.2. Proofs of Lemma (4.1) and Proposition (4.2)

A.2.1 Proof of Lemma (4.1)

Proof. We express the set of permutation matrices in terms of the coefficients $x_{i,j}$:

$$\begin{aligned} \mathbb{P}_n &= \{X \in \mathbb{R}^{n \times n} | \forall i, j \in \{1, \dots, n\} x_{i,j} \in \{0, 1\} \quad \forall j \in \{1, \dots, n\} \sum_{i=1}^n x_{i,j} = \sum_{i=1}^n x_{j,i} = 1\} \\ &= \left\{ \begin{pmatrix} 1 - \sum_{i=2}^n x_{i,1} & 1 - \sum_{i=2}^n x_{i,2} & \dots & 1 - \sum_{i=2}^n x_{i,n} \\ x_{2,1} & x_{2,2} & \dots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n,1} & x_{n,2} & \dots & x_{n,n} \end{pmatrix} \in \mathbb{R}^{n \times n} | \forall i \in \{2, \dots, n\}, j \in \{1, \dots, n\} x_{i,j} \in \{0, 1\}, \right. \\ &\quad \left. \forall j \in \{1, \dots, n\} \sum_{i=1}^n x_{j,i} = 1 \wedge \sum_{i=2}^n x_{i,j} \leq 1 \right\} \\ &= \left\{ \begin{pmatrix} 1 - \sum_{i=2}^n x_{i,1} & 1 - \sum_{i=2}^n x_{i,2} & \dots & 1 - \sum_{i=2}^n x_{i,n} \\ 1 - \sum_{i=2}^n x_{2,i} & x_{2,2} & \dots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 1 - \sum_{i=2}^n x_{n,i} & x_{n,2} & \dots & x_{n,n} \end{pmatrix} \in \mathbb{R}^{n \times n} | \forall i \in \{2, \dots, n\}, j \in \{1, \dots, n\} x_{i,j} \in \{0, 1\}, \right. \\ &\quad \left. \forall j \in \{1, \dots, n\} \sum_{i=1}^n x_{j,i} = 1 \wedge \sum_{i=2}^n x_{i,j} \leq 1 \right\} \\ &= \left\{ \begin{pmatrix} 1 - (n-1) + \sum_{i,j=2}^n x_{i,j} & 1 - \sum_{i=2}^n x_{i,2} & \dots & 1 - \sum_{i=2}^n x_{i,n} \\ 1 - \sum_{i=2}^n x_{2,i} & x_{2,2} & \dots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 1 - \sum_{i=2}^n x_{n,i} & x_{n,2} & \dots & x_{n,n} \end{pmatrix} \in \mathbb{R}^{n \times n} | \forall i \in \{2, \dots, n\}, j \in \{2, \dots, n\} x_{i,j} \in \{0, 1\}, \right. \\ &\quad \left. \sum_{i,j=2}^n x_{i,j} \in \{n-2, n-1\} \forall j \in \{1, \dots, n\}, \sum_{i=2}^n x_{j,i} \leq 1 \wedge \sum_{i=2}^n x_{i,j} \leq 1 \right\} \\ &= \left\{ \begin{pmatrix} 2 - n + \sum_{i,j=2}^n x_{i,j} & 1 - \sum_{i=2}^n x_{i,2} & \dots & 1 - \sum_{i=2}^n x_{i,n} \\ 1 - \sum_{i=2}^n x_{2,i} & x_{2,2} & \dots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 1 - \sum_{i=2}^n x_{n,i} & x_{n,2} & \dots & x_{n,n} \end{pmatrix} \in \mathbb{R}^{n \times n} | \forall i \in \{2, \dots, n\}, j \in \{2, \dots, n\} x_{i,j} \in \{0, 1\}, \right. \\ &\quad \left. \sum_{i,j=2}^n x_{i,j} \in \{n-2, n-1\} \forall j \in \{2, \dots, n\} \sum_{i=2}^n x_{j,i} \leq 1 \wedge \sum_{i=2}^n x_{i,j} \leq 1 \right\} \\ &= \left\{ \begin{pmatrix} 2 - n + \sum_{i,j=2}^n x_{i,j} & 1 - \sum_{i=2}^n x_{i,2} & \dots & 1 - \sum_{i=2}^n x_{i,n} \\ 1 - \sum_{i=2}^n x_{2,i} & x_{2,2} & \dots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 1 - \sum_{i=2}^n x_{n,i} & x_{n,2} & \dots & x_{n,n} \end{pmatrix} \in \mathbb{R}^{n \times n} | \forall i \in \{2, \dots, n\}, j \in \{2, \dots, n\} x_{i,j} \in \{0, 1\}, \right. \\ &\quad \left. \sum_{i,j=2}^n x_{i,j} \in \{n-2, n-1\} \forall j, i, k \in \{2, \dots, n\} \quad i \neq k, \quad x_{j,i} x_{j,k} = 0 \wedge x_{i,j} x_{k,j} = 0 \right\}. \end{aligned}$$

□

A.2.2 Proof of Proposition (4.2)

We want to find sufficient lower bounds for λ_1^j and λ_2 in

$$\min_{\mathbf{x} \in \{0,1\}^{(n-1)^2}} \mathbf{x}^T \tilde{W} \mathbf{x} + \tilde{c}^T \mathbf{x} + \sum_{j=1}^{2(n-1)} \lambda_1^j \left(\sum_{k \in \mathcal{J}_j} x_k \right) \left(\sum_{k \in \mathcal{J}_j} x_k - 1 \right) + \lambda_2 \left(\sum_{i=1}^{(n-1)^2} x_i - (n-1) \right) \left(\sum_{i=1}^{(n-1)^2} x_i - (n-2) \right), \quad (35)$$

so that it coincides with the constrained optimization problem. We make the following ansatz for the λ_2 parameter:

$$\lambda_2 = \frac{1}{2} D_{\{1, \dots, (n-1)^2\}}, \quad (36)$$

since the function

$$\left(\sum_{i=1}^{(n-1)^2} x_i - (n-1) \right) \left(\sum_{i=1}^{(n-1)^2} x_i - (n-2) \right) \quad (37)$$

increases faster than $2 \cdot H$, when H is the number of entries that need to be switched in order to have $n-1$ - or $n-2$ -many entries equal to 1.

This choice for λ_2 allows us to only investigate the case, where we have $n-1$ or n ones placed anywhere. Similar to the prior proof we ask, what the worst ratio between ones we have to insert anywhere and ones we have to insert or delete in a particular column is. It can be easily seen that in the worst case scenario all ones are in one column or row. Therefore the ratio is one half. This shows that we can choose:

$$\lambda_1^j = \frac{1}{2} D_{\mathcal{J}_j} + \frac{1}{2} D_{\{1, \dots, (n-1)^2\}}. \quad (38)$$

□

B. Experiments on IBM Burlington

Several of our experiments are conducted on the IBM five qubit QC-er in Burlington [1], which is accessible per cloud. In contrast to D-Wave quantum annealers, the circuit-model is used for the design of this quantum computer. From a purely theoretical perspective there is some equivalence [4] between adiabatic quantum computing and quantum computing with the circuit model, though due to the technical challenges, significant differences remain in practice. One way to execute a time-dependent Hamiltonian $H(t)$ (6) on the IBM machine is to approximate it with L constant Hamiltonians. The Hamiltonians, which are constant in time, can be decomposed and expressed with quantum gates via *trotterization* [55].

More specifically, choosing the value of the parameter L in the piecewise constant approximation of the Hamiltonian is crucial as we loose the (pseudo-)adiabaticity for small L , but – as each gate has a certain chance of introducing an error and a long execution time of the circuit makes errors due to decoherence more probable – large L are also prone to fail.

Because of the above-mentioned reasons, we are able to perform the adiabatic evolution for only two qubits on the 5 qubit processor until now. Fortunately, using the trick from Sec. 4.3, we can optimise over the 2×2 permutation matrices. In our experiment, we performed 50 time steps. We choose the evolution time 0.1 and use Suzuki expansion of order 2, similar to Kraus *et al.* [33]. The matrix W and the vector c were generated randomly for every of the 20 iterations. Every circuit was executed 8192 times, which is the maximal amount. An exemplary result of an execution is summarized in the histogram in Fig. 8. Although we obtain slightly different distributions each time, the highest peak always coincides with the second column of the optimal permutation, which is $|01\rangle$ here.

The recent publication [17] about adiabatic quantum computing on machines from IBM Quantum Experience proposes catalyst Hamiltonian, which could improve the results in future and make it possible to succeed in higher dimensions.

C. Point Cloud Matching Example

To illustrate another application of the studied matching problems, we consider the registration of two 3D point clouds by finding correspondences of four pre-selected points in each scene as illustrated in Fig. 9. While this pre-selection is, of course, a very challenging part of the overall solution, our main goal is to illustrate the solution of a (low dimensional) matching problem via quantum computing – possibly as the solution to a subproblem in an iterative registration algorithm.




Figure 8: QGM execution histogram over 8192 runs using trotterization on a five qubit processor from IBM Quantum Experience. The states $|00\rangle$ and $|11\rangle$ are suppressed, because of the proposed regularization terms, *i.e.*, since these states violate that the sum of all elements in a column of a permutation matrix equals to 1.

To set up the matching problem, we select four arbitrary points in one scene and use the ground truth transformation between the two frames to identify the corresponding points in the other scene. We then set up a matching problem with costs as in (17) by simply using Euclidean distances between the points.

The histogram of energies obtained by all three quantum graph matching formulations over 500 anneals is shown in Fig. 10. The top row illustrates the overall histogram with strong peaks at low energies. As these peaks correspond to permutation matrices, we can conclude that all penalty terms were successful in strongly promoting permutations, with the inserted formulation yielding the best results. Zooming into the leftmost peak (illustrated in the bottom row of Fig. 10), however, reveals that none of the three formulations was successful in consistently predicting the global optimum among the permutation matrices. Considering the probabilities of less than 1% for the row-wise and baseline, and about 5.5% for the inserted formulations to predict the ground truth solution, one must conclude that all algorithms do not provide significantly better solutions than random guessing (which has a success probability of 4.17% for $n = 4$).

In summary, this experiment underlines the great difficulty current quantum hardware still has with problem instances of $n = 4$, *i.e.*, looking for the values of the best 16 qubits that are constrained to representing a permutation matrix. The performance of our three QUBO formulations on this problem can be seen in Fig. 10.

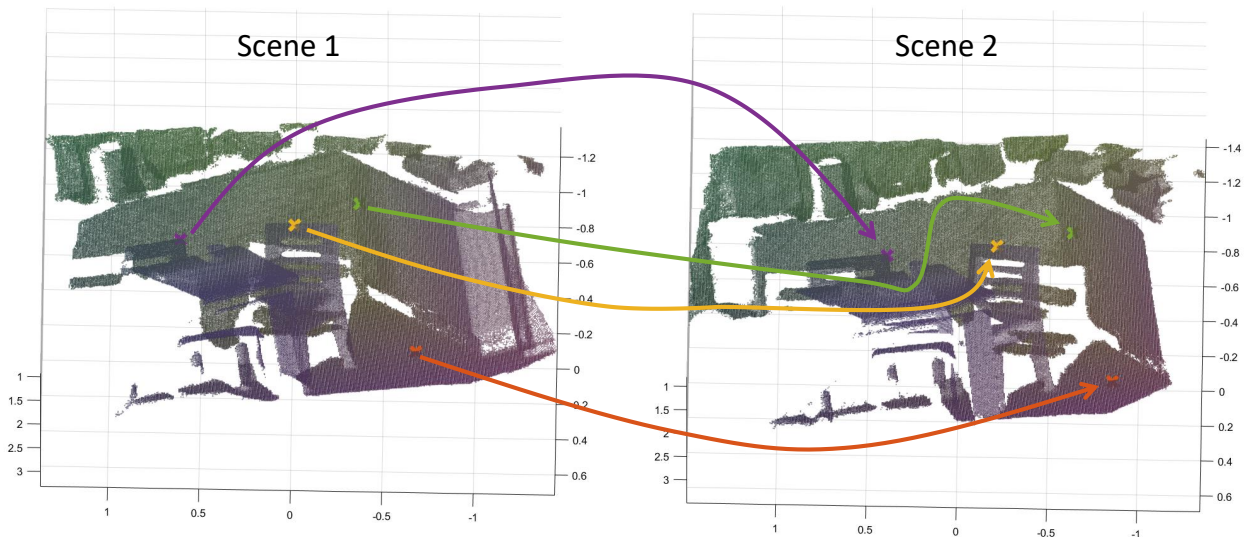


Figure 9: Illustrating the correct matching between four selected keypoints in two frames of the 7-Scenes 'Redkitchen' dataset, as provided at <https://3dmatch.cs.princeton.edu/>




Figure 10: The histograms show the states obtained for the point matching problem in Fig. 9. The left triple shows the overall histogram obtained over 500 anneals while the right triple is a (rescaled) zoom into the leftmost peaks of the left triple, which corresponds to actual permutation matrices.

D. Near-Isometric Matching

Fig. 11 shows an example of a matching problem between two shapes that are only approximately related by an isometry, *i.e.*, matching a wolf to a cat in this particular instance. Still modelling the problem as an isometric deformation using costs determined by (17) yields an instance, where the true matching still is the global optimum of the quadratic assignment problem, but where wrong permutations have considerably more similar energies than in the point cloud and isometric shape matching examples.




Figure 11: Illustrating a matching problem in which the shapes to be matched are only approximately related by an isometry.

The success probabilities for the *inserted*, the *baseline* and the *row-wise* QGM variants are 3.8%, 4% and 9.2%, respectively. As we can see in Fig. 12, the energy landscape has a wide range and the energy values corresponding to permutations are closer to each other compared to the case when the isometry assumption is strictly fulfilled. If one looks at the histograms with only permutations, only the *row-wise* QGM shows some trend towards the permutation with the lowest energy. Although its success probability is more than twice as good as random guessing, such a factor to random guessing would still not be sufficient to scale such an algorithm to large n .

E. Embedding to the Chimera Graph

The embeddings to the Chimera graph for different dimensions with the row-wise formulation can be seen in Fig. 13. While the number of logical qubits grows quadratically with n , the number of physical qubits required to embed those to the Chimera graph grows as n^4 . For $n = 2$, the length of the longest chain is two physical qubits. For $n = 3$ and $n = 4$, the chain length does not exceed four and six, respectively.

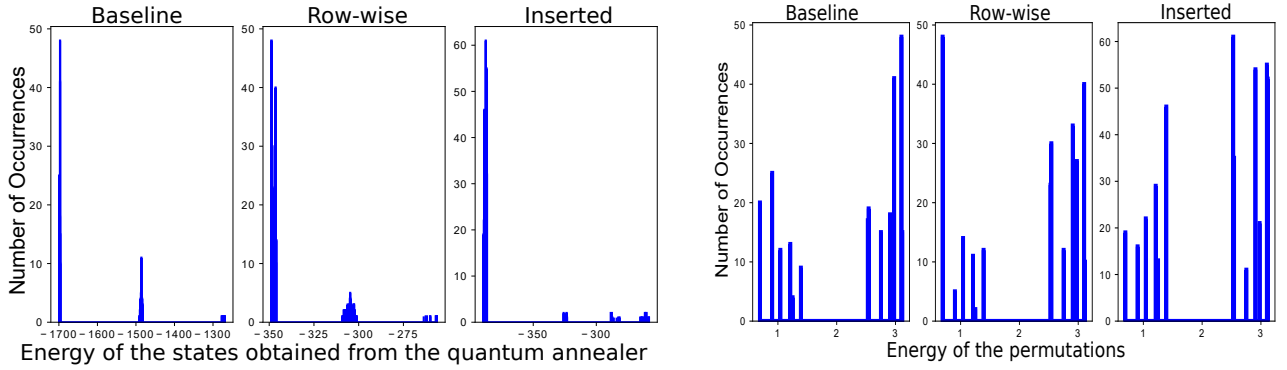


Figure 12: The histograms show the states obtained for the point matching problem in Fig. 11. The left triple shows the overall histogram obtained over 500 anneals while the right triple is a (rescaled) zoom into the leftmost peaks of the left triple, which corresponds to actual permutation matrices.




Figure 13: QGM minor embeddings to the chimera graph, for $n = 2$ (left), $n = 3$ (middle) and $n = 4$ (right). The number of logical qubits grows quadratically with n , and the number of physical qubits required to embed the logical qubits to the Chimera graph grows as n^4 on the current generation of D-Wave annealers. The white and blue circles denote measured values zero and one, respectively, and the lines between the qubits denote couplers. Grey lines connect the physical qubits with the same resulting values. A subset of those build chains and represent a single logical qubit. Blue lines connect different measured values. (Graphic generated using D-Wave leap problem inspector)

F. Discussion of Negative Results

Since physical quantum computing is an emerging technology, reporting and discussing negative results on the early stage is of high relevance for the community. Insights of this section can help in choosing a promising direction for improvements and future research.

As discussed in Sec. 5.2 we claim that the failure to provide the ground state with more probability than random guessing is due to the experimental errors in the coupling parameters. To back this up we provide the smallest and largest values of couplings and biases of the regularization and data term in Table 2.

The couplings and the biases are scaled, so that they fit the feasible region of the annealer. $Q_{\text{reg}}, q_{\text{reg}}$ do yield constant energies for all permutations and therefore $Q_{\text{prob}}, q_{\text{prob}}$ contains the information, which permutation is optimal. The sum are the real, physical couplings

$$Q = Q_{\text{reg}} + Q_{\text{prob}}$$

and biases

$$q = q_{\text{reg}} + q_{\text{prob}}$$

One can see that most of the accessible range of the coupling parameters has to make sure that the output is a permutation.

Table 2: Illustrating the range of values arising from the penalty to constrain each formulation to permutation matrices (Q_{reg^*} and q_{reg^*}) and from the actual problem costs (Q_{prob^*} and q_{prob^*}). As we can see the constraints contribute more to the quadratic coupling matrix by a factor of around 6 for the inserted, 13 for the row-wise, and almost 55 for the baseline.

	Row-wise	Inserted	Baseline
Q_{max}	-0.887	-0.987	-0.973
Q_{min}	-1.037	-1.295	-1.009
q_{max}	-1.937	-8.884	-116.772
q_{min}	-2.277	-9.753	-120.772
$Q_{\text{reg max}}$	0.962	1.121	0.983
$Q_{\text{reg min}}$	-0.962	-0.705	-0.991
$q_{\text{reg max}}$	4.017	10.885	118.772
$q_{\text{reg min}}$	0.140	7.753	118.772




Figure 14: Percentage of the 5000 runs that found the optimum with simulated annealing averaged over 10 random problem instances.

There exists already ways to deal with these problems, which we want to try out in further experiments. One possibility would be to use extended J -range parameter [24]. The easiest way for this would be to use the virtual graph embedding instead of the default one (EmbeddingComposite). First attempts in this direction show that using the virtual graph embedding requires a lot of computation time. For one instance, where we got an error warning we had to invest 12% of our access time.

As reported in the main paper we first used a long annealing path with a break and at some point switched to using $20\mu s$. Although a longer annealing time can often be used to enhance the success probability, if we additionally look at the time it takes to perform the experiment until one gets the optimum with for example 99% certainty then according to [43] $20\mu s$ seems to be the better choice, if less than 512 qubits are used.

G. Beyond Quantum Computing

In addition to the numerical experiments using quantum computing, we also briefly tested the effect of our three reformulations on methods that are inspired by physical systems. In Fig. 14, we compare the success probability of our three formulations for simulated annealing using random instances of graph matching problems. As we can see, the row-wise and—even more so—the inserted formulations yield results clearly superior to the baseline method, indicating that our analysis might be of use beyond quantum computing. Due to the \mathcal{NP} -hard nature of the underlying problem, it is to be expected that the overall success probability still decreases exponentially with increasing n (for a fixed number of simulated annealing runs).