Adiabatic Quantum Computing for Multi Object Tracking

Jan-Nico Zaech[1]   Alexander Liniger[1]   Martin Danelljan[1]   Dengxin Dai[1,2]   Luc Van Gool[1,3]

[1]Computer Vision Laboratory, ETH Zurich, Switzerland,
[2]MPI for Informatics, Saarbrucken, Germany, [3]KU Leuven, Belgium

{zaechj,alex.liniger,martin.denelljan,dai,vangool}@vision.ee.ethz.ch

Abstract

Multi-Object Tracking (MOT) is most often approached in the tracking-by-detection paradigm, where object detections are associated through time. The association step naturally leads to discrete optimization problems. As these optimization problems are often NP-hard, they can only be solved exactly for small instances on current hardware. Adiabatic quantum computing (AQC) offers a solution for this, as it has the potential to provide a considerable speedup on a range of NP-hard optimization problems in the near future. However, current MOT formulations are unsuitable for quantum computing due to their scaling properties. In this work, we therefore propose the first MOT formulation designed to be solved with AQC. We employ an Ising model that represents the quantum mechanical system implemented on the AQC. We show that our approach is competitive compared with state-of-the-art optimization-based approaches, even when using of-the-shelf integer programming solvers. Finally, we demonstrate that our MOT problem is already solvable on the current generation of real quantum computers for small examples, and analyze the properties of the measured solutions.

1. Introduction

Multi-Object Tracking (MOT) is a task in computer vision that requires solving NP-hard assignment problems [29,30,50]. To make this feasible, the community proposed a range of different approaches: work on the problem formulation using domain knowledge helps to make it an easier to solve problem [29,50], approximate solvers extend the feasible problem size [30], and the combination of deep learning with simple heuristics can be seen as a data-driven approach to the problem [8, 12]. Nevertheless, integer assignment problems remain hard optimization tasks for any available solver. With the recent progress in quantum computing, a new way of solving such optimization problems becomes feasible in the near future [1, 36, 51].

Instead of iteratively exploring possible solutions, e.g.
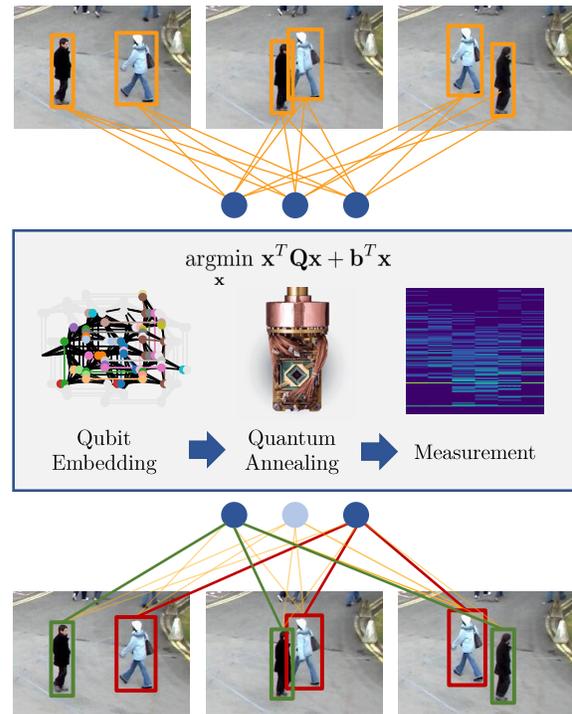
Figure 1. The proposed approach to MOT states the assignment problem between detections and a set of tracks as a quadratic unconstrained binary optimization task. We then represent the optimization problem as a quantum mechanical system that can be implemented on an AQC. Via quantum annealing, a minimum energy state is found that represent the best assignment.

via branch and bound, the problem is mapped to a quantum mechanical system, whose energy is equivalent to the cost of the optimization problem. Therefore, if it is possible to measure the lowest energy state of the system, a solution to the corresponding optimization problem is found. This is done with an adiabatic quantum computer (AQC), which implements a quantum mechanical system made from qubits and can be described by the Ising model [31]. Using this approach, a quantum speedup, which further

scales with system size and temperature, has already been shown for applications in physics [35, 36].

While there is a range of advantages that quantum computing can provide in the future, mapping a problem to an AQC is not trivial and often requires reformulating the problem from scratch, even for well investigated tasks [4, 6]. On the one hand, the problem needs to be matched to the Ising model, on the other hand, real quantum computers have a very limited number of qubits and are still prone to noise, which requires tuning of the model to handle the limitations.

In this work we present the first quantum computing approach to MOT. The number of required qubits in our formulation grows linearly in the number of detections, tracks and timesteps and only requires entanglement between qubits to model long-term relations. Our overall contributions are the following:

- A quantum computing formulation of MOT that is competitive with state-of-the-art methods.
- A method using few problem measurements to find Lagrange multipliers that considerably improve solution probability.
- Extensive MOT experiments on synthetic as well as real data using a D-wave AQC.

The remaining paper is structured as follows: After presenting related work, the basics of quantum computing are introduced. This is followed by our MOT formulation that is optimized to run on an AQC. We then show the changes required to make the problem solvable also with the classical computing paradigm. Finally, experiments on a D-wave quantum computer are presented together with results on larger problem instances.

2. Related Work

Quantum computing applications have recently started to emerge across a range of fields that rely on discrete optimization, as adiabatic quantum computers have become accessible. The applications include examples such as gene engineering [21], interaction reconstruction in particle physics [13], traffic flow optimization [44], or route selection in robotics [45]. In computer vision, discrete optimization is a ubiquitous part of many applications. While these applications frequently rely on heuristics today, quantum computing has the potential to provide an efficient way of directly solving them. In the area of 3D vision, quantum computing has been used by Feld *et al.* [19] for optimizing geometry compression. Benkner *et al.* [3] use adiabatic quantum computing to match 3D shapes and images with permutation matrices and investigate different constraint formulations to optimize the probability of finding a correct solution. By using an iterative approach, the same authors are able to scale the approach up to larger problem instances [4]. Closest to our work is the contribution of Birdal *et al.* [6]. They map the permutation synchronization task to an optimization problem solvable on a quantum computer and show results on small problem instances.

Multi-object tracking describes the problem of tracking all objects belonging to a predefined set of object types in 2D [14, 38, 41, 43] or 3D [10, 11, 23]. Most competitive trackers follow a tracking by detection approach, where a set of detections is given in every frame and the trackers perform association between frames, interpolation of occlusions, and rejection of false-positive detections. While most approaches use deep learning to generate appearance features [47, 52, 56, 57], two major groups of data assignment approaches exist. The first one maps the matching step to a deep learning task [8, 12, 55, 59] and uses simple heuristics to resolve the remaining inconsistencies. This allows for training the complete pipeline end-to-end, without the direct requirement to define a cost for data association. The second group directly performs data association using discrete optimization algorithms [29, 30, 39, 46, 48, 50], which is stated as a network flow optimization problem in most cases. These formulations allow to include long-term relations [29], and prior information about the nature of tracks in an intuitive and transparent way. Nevertheless, these properties come at a high computational cost. As most of the proposed optimization problems are NP-hard [22], a considerable effort was invested in finding heuristics and approximate solvers for them [30].

3. Preliminaries on Quantum Computing

Quantum computers are systems operating in a state that is described by its quantum properties, such as superposition and entanglement. By exploiting these properties, a range of problems that quickly grow in complexity on classical computers and thus, cannot be solved in any reasonable timeframe, could be solved considerably faster [20] by a quantum computer. Reaching such a point is widely referred to as quantum primacy. Even though implementations of quantum computers are still heavily experimental, some problems have already been shown to profit from them, including the sampling of pseudo-random quantum circuits [2, 53] and Gaussian boson sampling [58].

Qubits are two-state quantum-mechanical systems that form the basis of quantum computers. Like a bit, a qubit has two basis states $|0\rangle = [1\ 0]^T$ and $|1\rangle = [0\ 1]^T$ that in a superposition form the qubit's state. Qubits can be realized with a wide range of approaches, including superconducting circuits, ions trapped in an electromagnetic field, or photons.

Quantum Superposition refers to the property of a quantum system that it is not required to be in one of the basis states, but rather can be described by a linear combination of possible basis states. A qubit in a pure state $|\psi\rangle$ can be

described with its two basis states $\{|0\rangle, |1\rangle\}$ as

$$|\psi\rangle = c_1 |0\rangle + c_2 |1\rangle \qquad (1)$$

where c_1 and c_2 are complex numbers, called probability amplitudes, with $|c_1|^2 + |c_2|^2 = 1$.

Measurement of a qubit state results in one of the basis states $\{|0\rangle, |1\rangle\}$. The probability of measuring $|0\rangle$ and $|1\rangle$ evaluates to $|c_1|^2$ and $|c_2|^2$, respectively [28]. As a measurement corresponds to an observation of the qubit it leads to wave function collapse, which means that the qubit state is changed irreversibly [24].

Entanglement of qubits is at the very heart of quantum computing [32]. A system of entangled qubits is represented by a system state where each qubit cannot be described only with its own state but depends on the state of the remaining system [17, 27, 37, 40, 49]. Thus, measuring a single qubit can collapse the wave function of other entangled qubits, which alters their state and therefore, also their measurement outcome [24].

3.1. Adiabatic Quantum Computing

Adiabatic quantum computing [18, 51] is an approach that, instead of using gates as unitary operations on subsets of the available qubits [15], uses a problem Hamiltonian $\hat{H}_P$ that describes the operation applied on all qubits simultaneously. The problem Hamiltonian is designed such that its ground-state, which is the lowest energy configuration of the system, represents the result for a computational task [1]. In general, a Hamiltonian $\hat{H}(t)$ is an operator that represents the energy of a quantum system and can be used in the Schrödinger equation to describe the system's evolution over time as

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{H}(t) |\psi(t)\rangle, \qquad (2)$$

where i is the imaginary unit, and $\hbar$ the reduced Planck constant.

As the ground-state of the problem Hamiltonian is hard to find, the complete system is initialized with an initial Hamiltonian $\hat{H}_B$ that has an easy to prepare ground-state [18]. The system's Hamiltonian is then slowly evolved over an annealing time T to the problem Hamiltonian in an adiabatic transition [7, 34]

$$\hat{H}(t) = (1 - {}^t/T)\hat{H}_B + {}^t/T\hat{H}_P, \qquad (3)$$

which is a transition where the system stays in its basis state. This process is called quantum annealing and needs to be repeated for multiple measurements, as in a noisy system, not all solutions have the lowest energy. The condition for a sufficiently slow evolution depends mostly on two factors, the temperature of the environment and the spectral gap of the Hamiltonian, i.e. the difference between lowest

and second-lowest energy level or eigenvalue. While the first is a system property, the second can be influenced by choosing a suitable Hamiltonian [3].

The Hamiltonian describing current adiabatic quantum computers such as the *D-wave advantage*, is based on the Ising model [33]. The Ising model uses the Hamiltonian

$$\hat{H}_{ising} = \sum_{i,j} J_{i,j} \sigma_i \sigma_j + \sum_i h_i \sigma_i, \qquad (4)$$

where $\sigma \in \{-1, +1\}$ corresponds to the spin of a particle, $J_{i,j}$ represent the interaction between two particles and h_i is an external magnetic field. In an adiabatic quantum computer, the particles' spins are represented by the qubit states and the interactions and external field correspond to the couplings. The lowest energy of the Ising model is equivalent to solving the associated quadratic unconstrained binary optimization (QUBO)

$$\arg\min_{\mathbf{z}} \mathbf{z}^T \mathbf{Q} \mathbf{z} + b^T \mathbf{z}, \qquad (5)$$

which is NP-hard and known to be very challenging for classical solver. As this task can directly be implemented on an adiabatic quantum computer, a considerable speedup for large problem instances is expected in the future.

4. Quantum MOT

Most existing optimization-based approaches to MOT aim at finding feasible relaxations [30], implement efficient heuristics in the solution approach [29] or use deep learning together with post-processing [8] to solve the assignment problem. With the considerable amount of work invested into them, the problem became solvable for growing instances by now. Nevertheless, the assignment problem stays an NP-hard task to solve and growth is thus limited. Quantum computing with the associated speedup on hard problems can provide a solution to this challenge, even if the corresponding optimization problem is much harder to solve with classical approaches at the moment. However, representing tasks in a form suitable for quantum computing often requires a completely new formulation of the problem and MOT is not different in this aspect.

While widely used flow formulations [29, 30, 39] are suitable for exploiting sparsity, they come with a large set of inequality constraints, which makes them intractable on near-future quantum computers that are limited in the number of qubits. In this context, permutation matrices were shown to be a powerful tool for synchronization or shape matching [3, 4, 6]. In the following, we therefore propose a formulation based on assignment matrices that grows linearly in the number of required qubits for detections, tracks and frames. Furthermore, it allows to model long-term connections with terms in the cost-matrix that do not require additional qubits.

MOT Formulation. We approach the MOT problem following the tracking by detection paradigm and use a fixed set of available tracks. Given a set of detections in each frame of a video, appearance features are extracted for each detection. By using a multi-layer Perceptron, pairwise appearance similarities between detections at different timesteps are computed [30]. Starting with this, the goal of the tracking algorithm is to assign each detection to a track, such that the sum of the similarities of detections assigned to a single track is maximized. In this context, a track is defined by its track ID t and each detection in a frame f by its detection ID d.

We formulate the given task of assigning detections to a joint set of tracks using assignment matrices, which relax the assumptions of permutation matrices. The binary assignment matrix $\mathbf{X}_f$ for a frame f maps a vector of detection indices to a vector of tracks at every frame of a video. The elements $x_{dt} \in \{0,1\}$ of the assignment matrix represent the connections between detections d and tracks t. Given $D-1$ detections and $T-1$ tracks, the assignment matrix assigns a detection to a track if $x_{dt} = 1$. The requirement that a single detection is assigned to a track at one timestep, leads to the constraint

$$\sum_{d=1}^{D} x_{dt} = 1 \quad \forall t \in \{1, ..., T-1\}. \tag{6}$$

And reversely, Equation 7 asserts that every detection is assigned to a single track

$$\sum_{t=1}^{T} x_{dt} = 1 \quad \forall d \in \{1, ..., D-1\}. \tag{7}$$

To allow for false-positive detections as well as to handle the case of fewer detections than available tracks, one dummy-detection and one dummy-track, with the respective indices D and T, are introduced. A detection assigned to the dummy-track is treated as a false positive and a track that got the dummy-detection assigned to it is inactive or occluded. As the dummy-track and dummy-detection may be assigned multiple times, constraints 6 and 7 do not apply to them. To model tracks in a sequence consisting of F frames, a single assignment matrix $\mathbf{X}_f$ is required for each frame f, mapping the detections to tracks.

Quadratic Form. The basis for optimization-based trackers are costs between pairs of detections, where the cost is accounted for if two detections are connected by a common track. The goal of the tracker is to find a solution that minimizes the total cost associated with the assignment. Our approach using assignment matrices leads to a quadratic cost for a pair of frames i, j that reads

$$c_{ij} = \sum_{t} \sum_{d_i} \sum_{d_j} x_{id_i t} q_{d_i d_j} x_{jd_j t}, \tag{8}$$

with $x_{id_i t}$ and $x_{jd_j t}$ being entries from the assignment matrices $\mathbf{X}_i$ and $\mathbf{X}_j$ respectively and $q_{d_i d_j}$ as the corresponding similarity score. It is important to note that only detection pairs assigned to the same track incur a cost, which results in a single sum over the tracks t.

Equation 8 can be written in matrix form as

$$c_{ij} = \mathrm{vec}(\mathbf{X}_i)^T \mathbf{Q}_{ij} \mathrm{vec}(\mathbf{X}_j), \tag{9}$$

with $\mathrm{vec}(\mathbf{X})$ as a row-major vectorization of the corresponding assignment matrices and $\mathbf{Q}_{ij}$ as the cost matrix of the frame-pair. The maximum frame gap $\Delta f_{\max}$ that is modeled in our approach depends only on the density of the cost matrix. To include a connection between frames i and j, the matrix $\mathbf{Q}_{ij}$ needs to be filled with the corresponding similarity scores. The cost matrix $\mathbf{Q}_{ij}$ is sparse, as it also represents all terms that correspond to detection pairs matched to different tracks, which add no cost. Furthermore, no cost is associated with the mapping of a frame to itself, which includes the main diagonal of $\mathbf{Q}$.

A complete sequence consisting of F frames, can be represented with the stacked assignment matrix

$$\mathbf{z} = [\mathrm{vec}(\mathbf{X}_1)^T, ..., \mathrm{vec}(\mathbf{X}_F)^T]^T. \tag{10}$$

And the corresponding cost

$$c = \sum_{i=1}^{F} \sum_{j=1}^{F} c_{ij} = \mathbf{z}^T \mathbf{Q} \mathbf{z}, \tag{11}$$

where $\mathbf{Q}$ is a block-matrix made from all $\mathbf{Q}_{ij}$.

QUBO form. To solve the proposed MOT assignment problem with an adiabatic quantum computer it further needs to be represented as a QUBO task with $\{-1, +1\}$ spin states. This consists of two steps, firstly eliminating the constraints and secondly substituting the variables.
1) Constraints are represented using a Lagrangian multiplier λ. As our formulation does not include inequalities, no additional slack variables with corresponding qubits are required. Given the original quadratic program with constraints

$$\arg\min_{\mathbf{z}} \mathbf{z}^T \mathbf{Q} \mathbf{z} + \mathbf{b}^T \mathbf{z} \quad \text{s.t.} \quad \mathbf{G} \mathbf{z} = \mathbf{d}, \tag{12}$$

a QUBO can be formulated as

$$\arg\min_{\mathbf{z}} \mathbf{z}^T \mathbf{Q}' \mathbf{z} + \mathbf{b}'^T \mathbf{z} \tag{13}$$

with

$$\mathbf{Q}' = \mathbf{Q} + \lambda \mathbf{G}^T \mathbf{G} \tag{14}$$
$$\mathbf{b}' = -2\lambda \mathbf{G}^T \mathbf{b}. \tag{15}$$

2) Variables are substituted by replacing the optimization variables $z \in \{0,1\}$ with $s \in \{-1,1\}$ by using $z = 1/2(s+1)$ the resulting optimization problem reads

$$\arg\min_{\mathbf{s}} \mathbf{s}^T \mathbf{Q} \mathbf{s} + \mathbf{b}'^T \mathbf{s} \quad \text{with} \quad \mathbf{b}'^T = 2(\mathbf{b}^T + \mathbf{1}^T \mathbf{Q}). \quad (16)$$

Lagrangian Optimization. Solving the Lagrangian would require solving a problem in both discrete and continuous optimization variables (assignment, and Lagrangian multipliers, respectively). To solve the problem using AQC, we presented a constant penalty reformulation in the previous paragraph, which fixes the Lagrangian multipliers λ. In such an approach, if λ is large enough, constraint satisfaction is guaranteed. More precisely, a quadratic equality constraint reformulation of the form

$$\lambda ||\mathbf{G} \mathbf{z} - \mathbf{d}||_2^2, \quad (17)$$

is used in Equations 14 and 15, which allows to only consider positive Lagrangian multipliers λ. Even though λ needs to be just large enough from a theoretical perspective, in practice it should be as small as possible. This is especially relevant for AQC, as with a high λ the conditioning of the corresponding Hamiltonian in the AQC gets worse. This should be avoided as it results in a lower probability of finding the correct solution in each measurement.

Thus, in practice a problem dependent bound for the minimum penalty term λ_{min} should be used. One approach to reduce the spectral gap is to estimate an individual λ_i for each constraint $\mathbf{G}_i \mathbf{x} = \mathbf{d}_i$ using upper bounds. While such bounds can be computed, they are not tight in many cases. We, therefore, propose a heuristic to estimate the Lagrangian multipliers λ_i that closely match their minimal value $\lambda_{i,\text{min}}$. Each multiplier is modeled by

$$\lambda_i = \lambda_{\text{b}} + \lambda_i' + \lambda_{\text{off}}, \quad (18)$$

where λ_{b} is a small base value that resolves the easy to fulfill constraints, λ_i' is estimated during the optimization procedure and λ_{off} is an offset to increase the spectral gap.

Starting with $\lambda_i' = 0$ and $\lambda_{\text{off}} = 0$ for all constraints, the QUBO is solved using annealing. In general, this will result in a solution $\mathbf{z}_\lambda$ that does not fulfill the constraints. As in our formulation, only positive violations result in a cost improvement, i.e. $\mathbf{G} \mathbf{z} \geq \mathbf{d}$, the cost reduction of a constraint violation can be estimated as

$$a_i(\mathbf{z}_\lambda) = 2(\mathbf{z}_G^T \mathbf{Q} \mathbf{z}_\lambda - \min_j \mathbf{z}_{Gj}^T \mathbf{Q} \mathbf{z}_\lambda)/v_i^2, \quad (19)$$

$$\mathbf{z}_G^T = (\mathbf{G}_i \circ \mathbf{z}_\lambda^T) \quad (20)$$

$$v_i(\mathbf{z}_\lambda^0) = \mathbf{G}_i \mathbf{z}_\lambda - \mathbf{d}_i, \quad (21)$$

where $\mathbf{z}_G$ are the variables masked with $\mathbf{G}_i$ and v_i is the degree of violation. To fulfill the corresponding constraint, we set

$$\lambda_i'(\mathbf{z}_\lambda) = -a_i(\mathbf{z}_\lambda) - \lambda_{\text{b}} + \epsilon, \quad (22)$$

with a small ϵ to assert that constraint i is fulfilled in the current setting. While this can be evaluated for all constraints simultaneously, the full procedure needs to be performed iteratively, as not all constraints may be violated in the optimal solution. Nevertheless, the set of measurements returned by the AQC can be used to reduce the number of required iterations. Instead of taking a single best solution, all solutions $\mathbf{z}_j$ that are close to the optimal solution are evaluated and merged as $\lambda_i' = \max_j \lambda_i'(\mathbf{z}_j)$. In our formulation, these can be solutions where the track order is permuted.

After estimating the Lagrangian multipliers, the total cost matrix scale is small, nevertheless, the same also holds for the spectral gap, as the cost of not fulfilling constraints is small. Therefore, the additional offset λ_{off} is added to the Lagrangian multipliers.

Similarity Cost. We use the same approach for cost generation as AP-lift [30], where multi-layer Perceptrons are used to regress the similarity score between pairs of detections. Features used to compute this score are the intersection over union (IoU) of aligned boxes and the dot-product between DG-Net [57] appearance features. DG-Net features are generated with the network trained on the MOT15 dataset [38] together with [47,52,56]. To generate the MLP input vector, the features are normalized with a global context [29], which results in a total of 22 features [30]. Furthermore, assigning the dummy-detection to a track incurs no cost and assigning a detection to the dummy-track, i.e. labeling it as a false-positive, corresponds to a small negative value β. This is required to prevent the assignment of single detections to tracks.

Post Processing. Even in an offline setting, long sequences cannot be represented as a single optimization problem and need to be split into a set of overlapping subproblems. We set the overlap to the modeled frame gap, and match tracks using the common frames. Matching is stated as a linear sum problem that maximizes the number of detections that are jointly assigned to tracks in both subproblems. As multiple subsequent tracks can be modeled by a single track ID, tracks that are interrupted longer than the maximum modeled frame gap Δf_{max} are separated.

Problem Scaling. One important aspect when designing algorithms for current and near-future quantum computers is the required number of qubits. Many current formulations of MOT grow quickly in size w.r.t. the number of detections, tracks, frames and the length of the modeled frame gap. In contrast to this, the number of qubits in our approach only grows linearly in the number of detections, tracks and frames. Furthermore, by using a quadratic optimization problem, longer frame gaps can be modeled by

additional entries in the cost matrix, which correspond to additional couplings between qubits.

While on short sequences the number of possible tracks needs to be at least as high as the total number of tracks, long sequences can profit from a saturation of the required number of tracks. After a track has terminated, there is no cost associated with assigning new detections if they have a distance of more than the maximal frame gap $\Delta f_{\max}$ from the previous track. Therefore, multiple subsequent real tracks can be modeled by a single track ID and easily be separated in post-processing.

5. Traditional Solvers

While our formulation is advantageous when solved on an adiabatic quantum computer, publicly available real systems have not yet reached a scale where large experiments can be performed. We, therefore, use classical solvers to show the results of our approach on real-world tasks, even though a quadratic problem formulation is known to be hard in this context. A common requirement of solvers to perform quadratic binary optimization via branch and bound is the convexity of the continuous relaxation of the problem. This corresponds to a positive-definite cost matrix $\mathbf{Q}$, i.e. a matrix with only positive eigenvalues, which is not fulfilled for the given cost matrix in most cases.

5.1. Hessian Regularization

A common approach to enforce positive eigenvalues is adding an identity matrix scaled by ϵ. As this changes the cost function and thus the optimal solution, small values need to be used for ϵ, making this approach only suitable for compensating small negative eigenvalues. Nevertheless, investigating the constraints of our formulation leads to a sparse diagonal matrix $\mathbf{E}$ that can be added to the cost matrix $\mathbf{Q}$ without changing the optimal solution. With the same approach of grouping the total cost matrix into blocks between frames as in Equation (11), the following definition of $\mathbf{E}$ is provided in blocks between frames. As only diagonal elements are relevant, blocks between different frames are zero matrices $\mathbf{E}_{ij} = \mathbf{0}|i \neq j$. The blocks on the diagonal, which represent the mapping of a frame i to itself $\mathbf{E}_{ii}$ are diagonal matrices defined by the diagonal elements

$$e_{idt} = \begin{cases} e & d \in \{1, ..., D\}, t \in \{1, ..., T-1\} \\ 0 & t = T \end{cases}. \quad (23)$$

The indices refer to the position on the diagonal that correspond to detection d and track t. Given a block's assignment matrix $\mathbf{X}_i$, the total cost of the block after adding the diagonal term is

$$c_{ii} = \text{vec}(\mathbf{X}_i)^T (\mathbf{Q}_{ii} + \mathbf{E}_{ii}) \text{vec}(\mathbf{X}_i) = e(T-1), \quad (24)$$

with $\mathbf{Q}_{ii} = \mathbf{0}$ and T tracks in total. The intuition behind the definition is given in the following and the full proof is provided in the supplementary material.

Given a binary problem, any diagonal entry adds cost if a variable is active. In the detection track assignment problem, this corresponds to adding a constant if a detection is assigned to a track. As constraint 6 asserts that exactly one detection (real- or dummy-detection) is assigned to every real track each time-step, having a cost e for the assignment adds this cost for each of the $T - 1$ real tracks. As the constraint does not apply for the dummy-track with index T and an arbitrary number of detections may be assigned to it. Therefore, the same argument would not hold and we can not add an additional cost to these entries ($e_{ikl} = 0|t = T$), without influencing the total cost function.

6. Experiments and Results

AQC experiments are performed on a D-wave Advantage 4.1 [42]. The system contains at least 5000 qubits and 35,000 couplers implemented as superconducting qubits [9] and Josephson-junctions [26] respectively. Every qubit of the D-wave Advantage is connected to 15 other qubits, which needs to be reflected in the sparsity pattern of the cost matrix. If a denser matrix is required, chains of qubits are formed that represent a single state. The actual parameters can vary due to defective qubits and couplers. All experiments are performed using an annealing time of $1600\,\mu s$ and an additional delay between measurements to reduce the inter-sample correlation. In the following, a single measurement refers to the combination of an annealing cycle and the subsequent measurement.

Simulated annealing is used to evaluate our approach in a noise-free setting. We use the simulation provided by D-wave for this purpose.

Classical solvers are used to demonstrate the performance of the proposed algorithm on the full MOT15 dataset. All experiments using classical solvers are performed using Gurobi [25] with CVXPY [16] as a modeling language.

6.1. Lagrangian Multiplier

Fixed Lagrangian multipliers represent the basic approach to include constraints in the QUBO. We run experiments with synthetic tracking sequences where object detections are in random order. The scenarios are defined by their similarity scores, which we set to 0.8 for a match and -0.8 for different objects. Furthermore, we add Gaussian noise with variance σ^2 to the similarity scores and subsequently truncate them to $[-1, 1]$. In the experiments 3 detections over 5 frames and a noise level between $\sigma = 0.2$ to $\sigma = 1.0$ is used. The tracking parameters are set to 4 tracks and a maximal frame-gap of $\Delta f_{\max} = 3$ frames.

Results generated with simulated annealing are shown in Figure 2, where the top plot shows the solution probabil-
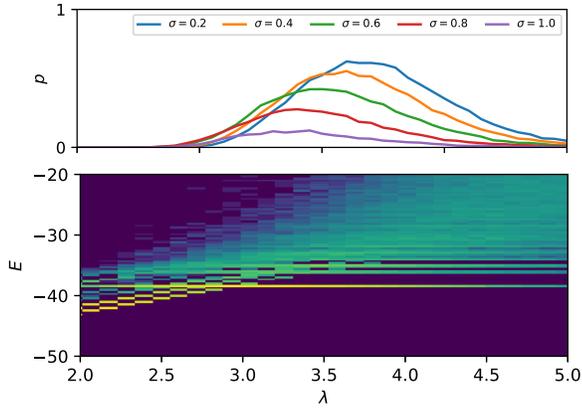
Figure 2. Solution probability and energy levels using simulated annealing for different noise levels and changing λ.

Figure 3. Solution probability and energy levels using quantum annealing for different noise levels and changing λ.

ity for different noise levels over an increasing Lagrangian multiplier. For each λ, 4096 measurements are performed. The lower plot shows the histogram over the energy of the returned solutions for a noise level of $\sigma = 0.6$. The correct solution can be seen at an energy level of -38.6.

With increasing noise level, the solution probability for the best value of λ reduces considerably, which can be explained by the energy histogram. As described in Section 3.1, a low spectral gap, i.e. the difference between the lowest and second-lowest energy level, reduces the probability of the AQC staying in its ground state and thus, the probability of finding the correct solution. In the energy plot, the spectral gap is visible as the distance between the energy band of the correct solution and the next higher energy band, given a sufficiently high λ, such that the correct solution has the lowest energy.

Tracking with the D-wave advantage is performed on a problem with 3 detections over 4 frames and noise levels $\sigma \in \{0.0, 0.1, 0.2\}$. Results using 4000 measurements for each setting are shown in Figure 3. Solution probabilities are lower compared to simulated annealing and high energy solutions are returned more often. This can be explained by the high noise of current AQCs.

Optimized Lagrangian multipliers are introduced to improve the spectral gap of the normalized cost matrix. We perform the same tracking tasks as for fixed Lagrangian multipliers, but evaluate the results w.r.t. the offset term λ_{off}. Results generated with simulated annealing are shown in Figure 4. Optimization of the Lagrangian multipliers is initialized with a base value of $\lambda_b = 0.5$. The probability of finding the right solution is increased and stays high over a large range of λ_{off} compared to only using a single λ. Furthermore, the best solution probability for each of the noise levels is better than the optimum for a fixed La-
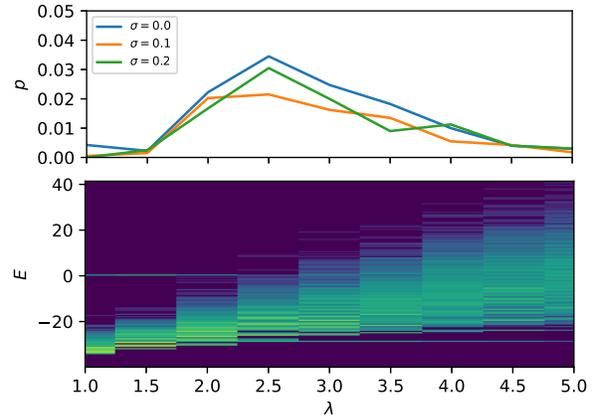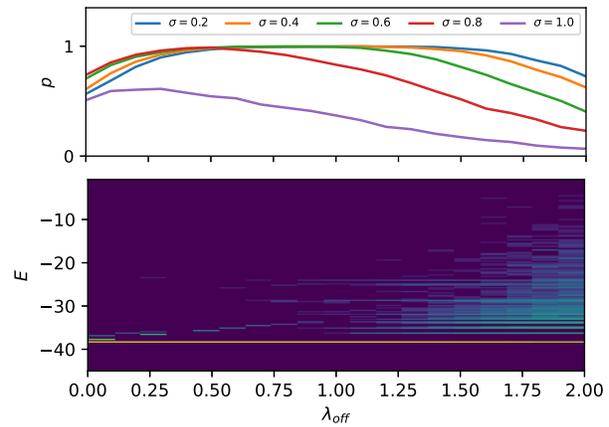
Figure 4. Solution probability and energy levels using simulated annealing and optimized λ_i for different noise levels over λ_{off}.

grangian multiplier. This has two advantages: first, fewer measurements are needed to find the correct solution and secondly, less effort needs to be invested to find a good setting for λ. Results for the problem with an optimized Lagrangian multiplier with $\lambda_b = 1.0$ solved on the AQC are shown in Figure 5. When optimally tuned for $\sigma = 0$, our method returns the best solution in 4.8% of the measurements, compared to 3.5% when using a fixed multiplier. Furthermore, even without an additional offset $\lambda_{\text{off}} = 0$, the best solution is returned in 0.8% of the measurements.

6.2. MOT15

We use the MOT15 dataset [38] to show that our method performs on par with state-of-the-art tracking methods. For this dataset, GUROBI [25] is used to find a solution for the optimization problem. The sequence is evaluated in segments of 20 frames using a maximum frame gap of $\Delta f_{\text{max}} = 10$. As binary quadratic problems are very hard to
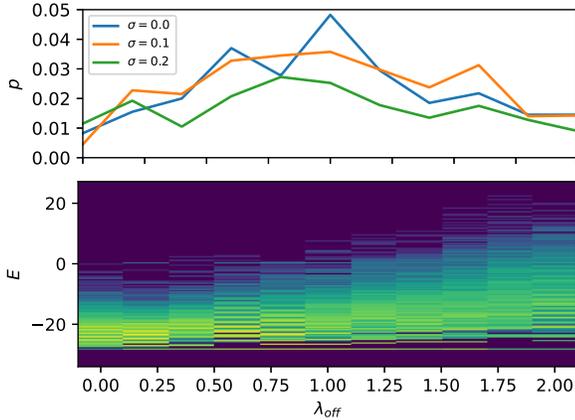
Figure 5. Solution probability and energy levels using quantum annealing and optimized λ_i for different noise levels over λ_{off}.

	Method	MOTA	IDF1	MT	ML	FP	FN	IDs
Test	Lif_T [29]	**52.5**	**60.0**	244	186	6837	21610	730
	MPNTrack [8]	51.5	58.6	225	187	7260	21780	**375**
	ApLift [30]	51.1	59.0	**284**	**163**	10070	**19288**	677
	MFI_TST [54]	49.2	52.4	210	176	8707	21594	912
	Tracktor [5]	44.1	46.7	130	189	6477	26577	1318
	Ours	49.9	53.5	187	179	**5924**	24032	1689
X-val	AP lift	59.6	**67.8**	237	**133**	8897	**10150**	**283**
	Ours	**59.7**	67.6	234	134	**8720**	10214	370

Table 1. Results on MOT15 [38]. X-val refers to results on the training set using leave-one-out cross validation.

solve with classical approaches, it is not possible to find an optimum solution for segments that contain a high number of tracks. In these cases, we terminate the optimization after $900\,\text{s}$ on a single segment and use the best solution found.

For comparisons, ApLift [30] is closest to our method, as it uses the same set of similarity features. On the test set, we achieve a MOTA-score of 49.9% and perform only 1.2% below ApLift, even though it models gaps up to 50 frames.

For a comparison under similar settings, we evaluate our method and ApLift [30] with the same frame gap of $\Delta f_{\text{max}} = 10$. As MOT15 does not contain a validation set, we use leave one out cross-validation on all samples of the training set for a fair comparison. In this scenario, our method improves by 0.2% over ApLift in $MOTA$ score. An explanation for this is that the MOT15 test set contains more detections in each frame on average (10.6 vs. 7.3) than the training set. In this case, there are more sequences where the classical solver does not find a solution and thus, generates a non-optimal result.

MOT15 with AQC. To show that tracking with an AQC already scales to small real-world examples, a part of the PETS09-S2L1 sequence is used. As the problem size has to be limited, three tracks that contain two occlusions, are extracted between frames 121 to 155. We execute our pipeline

Figure 6. Frames from the extracted sequence tracked on the AQC.
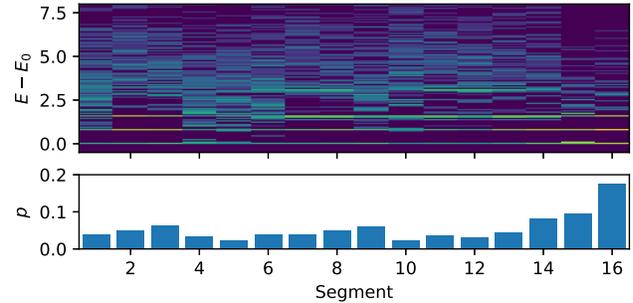
Figure 7. Energy of measurements returned by performing tracking of the PETS09-S2L1 sequence on the D-wave Advantage. The bar-plot shows the probability of measuring the optimal solution.

on segments of 5 frames with 3 tracks, a maximum frame-gap of 3, and optimized Lagrangian multipliers. The subproblems are solved on the D-wave Advantage with $1600\,\mu s$ annealing time and 500 measurements per segment. The most relevant frames that highlight occlusions are shown in Figure 6. The normalized energy $E - E_0$ levels of the measurements for each subproblem are shown at the top of Figure 7 and the corresponding probabilities p of measuring the right solutions are plotted in the lower one. The subproblems 5 and 10 correspond to the two occlusions highlighted in Figure 6. These are harder to solve problems, as multiple solutions with small differences in their energy exist and thus, they have a lower solution probability.

7. Conclusion

In this work, we proposed the first quantum computing formulation of MOT. We demonstrated that current AQCs can solve small real-world tracking problems, and that our approach closely matches state-of-the-art MOT methods. Current limitations stem from the proposed formulation being optimized to run on an AQC. As QUBO is know to be hard using classical approaches and as current AQCs are still at an experimental stage, problems are limited to a small scale. Nevertheless, quantum computing has the potential to make much larger problems feasible in the future.

References

[1] B. Apolloni, C. Carvalho, and D. de Falco. Quantum stochastic optimization. *Stochastic Processes and their Applications*, 33(2):233–244, Dec. 1989.

[2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, Oct. 2019.

[3] Marcel Seelbach Benkner, Vladislav Golyanik, Christian Theobalt, and Michael Moeller. Adiabatic Quantum Graph Matching with Permutation Matrix Constraints. In *2020 International Conference on 3D Vision (3DV)*, pages 583–592, Fukuoka, Japan, Nov. 2020. IEEE.

[4] Marcel Seelbach Benkner, Zorah Lahner, 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 (ICCV)*, page 11, 2021.

[5] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking Without Bells and Whistles. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 941–951, Seoul, Korea (South), Oct. 2019. IEEE.

[6] Tolga Birdal, Vladislav Golyanik, Christian Theobalt, and Leonidas Guibas. Quantum Permutation Synchronization. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13117–13128, Nashville, TN, USA, June 2021. IEEE.

[7] M. Born and V. Fock. Beweis des Adiabatensatzes. *Zeitschrift für Physik*, 51(3):165–180, Mar. 1928.

[8] Guillem Braso and Laura Leal-Taixe. Learning a Neural Solver for Multiple Object Tracking. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6246–6256, Seattle, WA, USA, June 2020. IEEE.

[9] P. I. Bunyk, E. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, and J. Whittaker. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Transactions on Applied Superconductivity*, 24(4):1–10, Aug. 2014.

[10] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019.

[11] Ming-Fang Chang, John W Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3D tracking and forecasting with rich maps. In *Conference on computer vision and pattern recognition (CVPR)*, 2019.

[12] Peng Dai, Renliang Weng, Wongun Choi, Changshui Zhang, Zhangping He, and Wei Ding. Learning a Proposal Classifier for Multiple Object Tracking. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2443–2452, Nashville, TN, USA, June 2021. IEEE.

[13] Souvik Das, Andrew J. Wildridge, Sachin B. Vaidya, and Andreas Jung. Track clustering with a quantum annealer for primary vertex reconstruction at hadron colliders. *arXiv:1903.08879 [hep-ex, physics:quant-ph]*, Aug. 2020. arXiv: 1903.08879.

[14] Patrick Dendorfer, Hamid Rezatofighi, Anton Milan, Javen Shi, Daniel Cremers, Ian Reid, Stefan Roth, Konrad Schindler, and Laura Leal-Taixé. MOT20: A benchmark for multi object tracking in crowded scenes. *arXiv:2003.09003 [cs]*, Mar. 2020. arXiv: 2003.09003.

[15] David Elieser Deutsch and Roger Penrose. Quantum computational networks. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 425(1868):73–90, Sept. 1989. Publisher: Royal Society.

[16] Steven Diamond and Stephen Boyd. CVXPY: a python-embedded modeling language for convex opti-

mization. *The Journal of Machine Learning Research*, 17(1):2909–2913, Jan. 2016.

[17] A. Einstein, B. Podolsky, and N. Rosen. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Physical Review*, 47(10):777–780, May 1935. Publisher: American Physical Society.

[18] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum Computation by Adiabatic Evolution. *arXiv:quant-ph/0001106*, Jan. 2000. arXiv: quant-ph/0001106.

[19] Sebastian Feld, Markus Friedrich, and Claudia Linnhoff-Popien. Optimizing Geometry Compression Using Quantum Annealing. In *2018 IEEE Globecom Workshops (GC Wkshps)*, pages 1–6, Dec. 2018.

[20] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, June 1982.

[21] Dillion M. Fox, Kim M. Branson, and Ross C. Walker. mRNA codon optimization with quantum computers. *PLOS ONE*, 16(10):e0259101, Oct. 2021. Publisher: Public Library of Science.

[22] Robert Ganian, Thekla Hamm, and Sebastian Ordyniak. The Complexity of Object Association in Multiple Object Tracking. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(2):1388–1396, May 2021. Number: 2.

[23] A Geiger, P Lenz, C Stiller, and R Urtasun. Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, Sept. 2013.

[24] David J. Griffiths and Darrell F. Schroeter. *Introduction to quantum mechanics*. Cambridge University Press, Cambridge, 3 edition, 2018.

[25] Gurobi Optimization, LLC. Gurobi optimizer reference manual, 2021.

[26] R. Harris, T. Lanting, A. J. Berkley, J. Johansson, M. W. Johnson, P. Bunyk, E. Ladizinsky, N. Ladizinsky, T. Oh, and S. Han. Compound Josephson-junction coupler for flux qubits with minimal crosstalk. *Physical Review B*, 80(5):052506, Aug. 2009. Publisher: American Physical Society.

[27] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson. Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682–686, Oct. 2015.

[28] Alexander S. Holevo. Repeated and Continuous Measurement Processes. In Alexander S. Holevo, editor, *Statistical Structure of Quantum Theory*, Lecture Notes in Physics Monographs, pages 97–118. Springer, Berlin, Heidelberg, 2001.

[29] Andrea Hornakova, Roberto Henschel, Bodo Rosenhahn, and Paul Swoboda. Lifted Disjoint Paths with Application in Multiple Object Tracking. In *Proceedings of the 37th International Conference on Machine Learning*, pages 4364–4375. PMLR, Nov. 2020. ISSN: 2640-3498.

[30] Andrea Hornakova, Timo Kaiser, Paul Swoboda, Michal Rolinek, Bodo Rosenhahn, and Roberto Henschel. Making Higher Order MOT Scalable: An Efficient Approximate Solver for Lifted Disjoint Paths. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, page 11, 2021.

[31] Ernst Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, Feb. 1925.

[32] Richard Jozsa and Noah Linden. On the role of entanglement in quantum-computational speed-up. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 459(2036):2011–2032, Aug. 2003. Publisher: Royal Society.

[33] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse Ising model. *Physical Review E*, 58(5):5355–5363, Nov. 1998.

[34] Tosio Kato. On the Adiabatic Theorem of Quantum Mechanics. *Journal of the Physical Society of Japan*, 5(6):435–439, Nov. 1950. Publisher: The Physical Society of Japan.

[35] Andrew D. King, Juan Carrasquilla, Jack Raymond, Isil Ozfidan, Evgeny Andriyash, Andrew Berkley, Mauricio Reis, Trevor Lanting, Richard Harris, Fabio Altomare, Kelly Boothby, Paul I. Bunyk, Colin Enderud, Alexandre Fréchette, Emile Hoskinson, Nicolas Ladizinsky, Travis Oh, Gabriel Poulin-Lamarre, Christopher Rich, Yuki Sato, Anatoly Yu Smirnov, Loren J. Swenson, Mark H. Volkmann, Jed Whittaker, Jason Yao, Eric Ladizinsky, Mark W. Johnson, Jeremy Hilton, and Mohammad H. Amin. Observation of topological phenomena in a programmable lattice of 1,800 qubits. *Nature*, 560(7719):456–460, Aug. 2018.

[36] Andrew D. King, Jack Raymond, Trevor Lanting, Sergei V. Isakov, Masoud Mohseni, Gabriel Poulin-Lamarre, Sara Ejtemaee, William Bernoudy, Isil Ozfidan, Anatoly Yu Smirnov, Mauricio Reis, Fabio Altomare, Michael Babcock, Catia Baron, Andrew J. Berkley, Kelly Boothby, Paul I. Bunyk, Holly Chris-

tiani, Colin Enderud, Bram Evert, Richard Harris, Emile Hoskinson, Shuiyuan Huang, Kais Jooya, Ali Khodabandelou, Nicolas Ladizinsky, Ryan Li, P. Aaron Lott, Allison J. R. MacDonald, Danica Marsden, Gaelen Marsden, Teresa Medina, Reza Molavi, Richard Neufeld, Mana Norouzpour, Travis Oh, Igor Pavlov, Ilya Perminov, Thomas Prescott, Chris Rich, Yuki Sato, Benjamin Sheldan, George Sterling, Loren J. Swenson, Nicholas Tsai, Mark H. Volkmann, Jed D. Whittaker, Warren Wilkinson, Jason Yao, Hartmut Neven, Jeremy P. Hilton, Eric Ladizinsky, Mark W. Johnson, and Mohammad H. Amin. Scaling advantage over path-integral Monte Carlo in quantum simulation of geometrically frustrated magnets. *Nature Communications*, 12(1):1113, Feb. 2021.

[37] Carl A. Kocher and Eugene D. Commins. Polarization Correlation of Photons Emitted in an Atomic Cascade. *Physical Review Letters*, 18(15):575–577, Apr. 1967. Publisher: American Physical Society.

[38] Laura Leal-Taixé, Anton Milan, Ian Reid, Stefan Roth, and Konrad Schindler. MOTChallenge 2015: Towards a Benchmark for Multi-Target Tracking. *arXiv:1504.01942 [cs]*, Apr. 2015. arXiv: 1504.01942.

[39] Zhang Li, Li Yuan, and R. Nevatia. Global data association for multi-object tracking using network flows. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2008. ISSN: 1063-6919.

[40] Xiao-Song Ma, Thomas Herbst, Thomas Scheidl, Daqing Wang, Sebastian Kropatschek, William Naylor, Bernhard Wittmann, Alexandra Mech, Johannes Kofler, Elena Anisimova, Vadim Makarov, Thomas Jennewein, Rupert Ursin, and Anton Zeilinger. Quantum teleportation over 143 kilometres using active feed-forward. *Nature*, 489(7415):269–273, Sept. 2012.

[41] Santiago Manen, Michael Gygli, Dengxin Dai, and Luc Van Gool. PathTrack: Fast Trajectory Annotation with Path Supervision. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 290–299, Venice, Oct. 2017. IEEE.

[42] Catherine McGeoch and Pau Farré. The Advantage System: Performance Update. page 31, Oct. 2021.

[43] Anton Milan, Laura Leal-Taixe, Ian Reid, Stefan Roth, and Konrad Schindler. MOT16: A Benchmark for Multi-Object Tracking. *arXiv:1603.00831 [cs]*, May 2016. arXiv: 1603.00831.

[44] Florian Neukart, Gabriele Compostella, Christian Seidel, David von Dollen, Sheir Yarkoni, and Bob Parney. Traffic Flow Optimization Using a Quantum Annealer. *Frontiers in ICT*, 4:29, 2017.

[45] Masayuki Ohzeki, Akira Miki, Masamichi J. Miyama, and Masayoshi Terabe. Control of Automated Guided Vehicles Without Collision by Quantum Annealer and Digital Devices. *Frontiers in Computer Science*, 1:9, 2019.

[46] Weihong Ren, Xinchao Wang, Jiandong Tian, Yandong Tang, and Antoni B. Chan. Tracking-by-Counting: Using Network Flows on Crowd Density Maps for Tracking Multiple Targets. *IEEE Transactions on Image Processing*, 30:1439–1452, 2021. Conference Name: IEEE Transactions on Image Processing.

[47] Ergys Ristani, Francesco Solera, Roger S. Zou, Rita Cucchiara, and Carlo Tomasi. Performance Measures and a Data Set for Multi-Target, Multi-Camera Tracking. *arXiv:1609.01775 [cs]*, Sept. 2016. arXiv: 1609.01775.

[48] Amir Roshan Zamir, Afshin Dehghan, and Mubarak Shah. GMCP-Tracker: Global Multi-object Tracking Using Generalized Minimum Clique Graphs. In Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid, editors, *Computer Vision – ECCV 2012*, Lecture Notes in Computer Science, pages 343–356, Berlin, Heidelberg, 2012. Springer.

[49] E. Schrödinger. Discussion of Probability Relations between Separated Systems. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):555–563, Oct. 1935. Publisher: Cambridge University Press.

[50] Siyu Tang, Mykhaylo Andriluka, Bjoern Andres, and Bernt Schiele. Multiple People Tracking by Lifted Multicut and Person Re-identification. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3701–3710, Honolulu, HI, July 2017. IEEE.

[51] W. van Dam, M. Mosca, and U. Vazirani. How powerful is adiabatic quantum computation? In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 279–287, Oct. 2001. ISSN: 1552-5244.

[52] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person Transfer GAN to Bridge Domain Gap for Person Re-identification. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 79–88, Salt Lake City, UT, USA, June 2018. IEEE.

[53] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun

Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang, Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Qingling Zhu, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. Strong Quantum Computational Advantage Using a Superconducting Quantum Processor. *Physical Review Letters*, 127(18):180501, Oct. 2021. Publisher: American Physical Society.

[54] Jieming Yang, Hongwei Ge, Jinlong Yang, Yubing Tong, and Shuzhi Su. Online multi-object tracking using multi-function integration and tracking simulation training. *Applied Intelligence*, May 2021.

[55] Jan-Nico Zaech, Dengxin Dai, Alexander Liniger, Martin Danelljan, and Luc Van Gool. Learnable Online Graph Representations for 3D Multi-Object Tracking. *arXiv:2104.11747 [cs]*, Apr. 2021. arXiv: 2104.11747.

[56] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable Person Re-identification: A Benchmark. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1116–1124, Santiago, Chile, Dec. 2015. IEEE.

[57] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. Joint Discriminative and Generative Learning for Person Re-Identification. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2133–2142, Long Beach, CA, USA, June 2019. IEEE.

[58] Han-Sen Zhong, Yu-Hao Deng, Jian Qin, Hui Wang, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Dian Wu, Si-Qiu Gong, Hao Su, Yi Hu, Peng Hu, Xiao-Yan Yang, Wei-Jun Zhang, Hao Li, Yuxuan Li, Xiao Jiang, Lin Gan, Guangwen Yang, Lixing You, Zhen Wang, Li Li, Nai-Le Liu, Jelmer J. Renema, Chao-Yang Lu, and Jian-Wei Pan. Phase-Programmable Gaussian Boson Sampling Using Stimulated Squeezed Light. *Physical Review Letters*, 127(18):180502, Oct. 2021. Publisher: American Physical Society.

[59] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking Objects as Points. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, Lecture Notes in Computer Science, pages 474–490, Cham, 2020. Springer International Publishing.