Multiple Hypothesis Tracking Using Neutral-Atom Quantum Computers

Sanidhya Gupta[1] Natan Karaev[2] Marissa McMaster[3] Matt Prest[4]

[1]Durham University
[2]Technion – Israel Institute of Technology
[3]University of Maryland, College Park
[4]City University of New York

1 Introduction

Quantum computers are expected to efficiently handle complex computational tasks, especially for solving optimization and combinatorics problems. Many of these problems have well-understood classical solutions that can be handled by high-performance computers (HPCs) to some extent. Nonetheless, it is essential to implement these solutions on quantum computers, if possible, to check for speedups, disparities in resource allocation, or other interesting properties to inform the future directions for quantum hardware and applications.

Despite the recent explosive progress of quantum hardware, there does not exist a platform which is universally the best choice for all applications. As such, one must be thoughtful when choosing a quantum device to meet their objectives. Whereas some applications require balancing trade-offs between factors such as qubit count, connectivities, gate speeds, fidelities and error rates, and coherence times, there exist computational tasks that are well-suited to specific platforms for more fundamental reasons. A clear and popular example of this is that of the Maximum Weight Independent Set Problem (MWIS). MWIS is a combinatorial problem that takes a graph with weighted vertices and locates the independent set with the largest weight [3]. This maps nicely onto the neutral atom platform, on which atoms, representing graph vertices, can be easily arranged into arbitrary geometries and the Rydberg blockade mechanism can enforce independence of solutions. The tunability of Rydberg interactions also allows for the introduction of a weight to vertices that is not easily accessible in other quantum platforms. This report will focus on applying the MWIS problem to Multiple Hypothesis Tracking (MHT), a useful tool for constructing trajectories of multiple targets in fields ranging from defense to biology.

We propose a method to implement a MWIS problem, applied to MHT, on a neutral-atom quantum computer in a hybrid classical-quantum system. Section 2 outlines the problem in more detail, while Section 3 connects the mathematical structures in the previous section to physical elements of the neutral atom platform. In Section 4 we present our approach in concrete steps, starting with classically generating instructions for the quantum computer before executing the instructions on the device itself. Section 5 presents the results of running our algorithm on a classical computer and on Pasqal's Pulser quantum computer simulator [6] for some simple data sets. We conclude with an analysis of the of our results and comment on possible methods for improvement.

2 Mutliple-Hypothesis Tracking and the Maximum Weight Independent Set Problem

MHT is broadly applicable to problems in statistics involving competing subsets of dynamically evaluatable hypothesis in the presence of expanding data. One such common case is the application of MHT to data association in multiple (unlabeled) target tracking. This could include image frames of ballistic trajectories or slides of cell evolution over time. We define an *observation* as all of the observed quantities/measurements (e.g. kinematic data), while a *scan* is any set of input observations that are concurrent. We begin with S total scans of observations $Z(k)$, where each observation consists of m_k measurements for $k \in \{1, \ldots, S\}$. We then let $z^k_{i_k}$ denote the ith measurement within scan k, which gives us $Z(k) = \{z^k_{i_k}\}^{m_k}_{i_k=1}$, the set of all observations for a given scan k. Further, we introduce the dummy observation z^k_0 in each scan for later convenience. A *track* is a sequence of observations $(z^1_{i_1}, \ldots, z^S_{i_S})$. A *family* is all tracks descended from the same initiating detection, i.e. all competing stories about one object. A *hypothesis* is a set possible tracks with an associated likelihood score while a *global hypothesis* is a set of hypotheses that could all be true at once. Some assumptions implicit here

include that there is at most one observation per target within a given scan, and that any two tracks that claim the same observation are incompatible.

Given the realistic challenges of ambiguous object identification, missed observations and false detections, naively assigning a single hypothesis to every observation is insufficient. MHT instead stores collections of hypotheses that are then updated using Bayesian methods with new observations.

2.1 From MHT to MWISP

The Maximum Weighted Independent Set Problem (MWISP) takes an undirected Graph $G = (V, E)$ with vertices V and edges E. An independent set is a subset of vertices $v \subseteq V$ who's elements do not share an edge with one another. Each vertex is assigned a positive weight w. The MWISP is then finding the maximum total weight of any independent subset of V, which is NP-Hard.

We can translate an MHT problem into a MWISP problem by building a graph in which each vertex is a candidate track and is assigned a weight which is it's likelihood score, and an edge joins all pairs of tracks that are incompatible. The best global hypothesis is then

$$\max \sum_i w_i x_i \quad \text{s.t.} \quad x_i + x_j \leq 1 \;\; \forall (i,j) \in E, \tag{1}$$

with $x_i \in \{0, 1\}$ indicating whether track i is selected. The weight w_i is the track's log-likelihood ratio (LLR): how much more likely it is that this sequence of detections came from a real object than from independent false alarms. It is updated recursively,

$$\mathrm{LLR}(k) = \mathrm{LLR}(k-1) + \Delta\mathrm{LLR}(k), \tag{2}$$

$$\Delta\mathrm{LLR} = \begin{cases} \log(1 - P_D) & \text{no detection,} \\ \log \dfrac{P_D \, \mathcal{N}(\nu; 0, S)}{\beta_{\mathrm{FA}}} & \text{detection used,} \end{cases} \tag{3}$$

where P_D is the probability of detection, ν the innovation, S its covariance and β_{FA} the spatial density of false alarms. Two further quantities matter, because MHT wants more than the single best answer. The probability of a global hypothesis and of an individual track are

$$P\{H_j\} = \frac{e^{L_{H_j}}}{1 + \sum_i e^{L_{H_i}}}, \qquad P\{T_i\} = \sum_{j \,:\, T_i \in H_j} P\{H_j\}. \tag{4}$$

As the space of possible track grows it becomes infeasible to maintain every track, this is where we can use *pruning* to remove tracks whose likelihood score falls below a desired threshold t. Additionally, we can use *gating* which involves placing velocity constraints on objects so that tracks do not propose one object traveling too far between observations. Lastly, *clustering* is used decompose the problem into smaller independent problems. We define *clusters* as sets of families whose tracks share at least one observation.

3 Principles of Rydberg Physics and Applicability to MWISP

Rydberg atoms are excited to high principal quantum numbers ($n \sim 55 - 80$). As atomic radius scales as n^2, these atoms have significant spatial separation between the nucleus and the excited valence electron, generating a large dipole moment and increasing the atoms' sensitivity to external fields. One of the properties of Rydberg atoms that makes them such promising candidates for quantum computing devices is the Rydberg blockade, which can be used for highly-customizable operations. If one atom is excited to a Rydberg state, there is a finite radius called the Rydberg blockade in which contained atoms cannot be simultaneously excited to Rydberg levels (see Figure 1). This is because the presence of a Rydberg atom within the blockade radius generates a dipole-dipole interaction strong enough to light shift the non-excited atom off-resonance with excitation lasers.

To discuss exploiting this feature for the MWIS problem, it is important to establish the relevant expressions governing these interactions. Perhaps the most essential relationship is that of the blockade radius and hardware-tunable parameters such as excitation beam detuning and intensity (together modifying the frequency of Rabi oscillations between the Rydberg and ground states). The Hamiltonian of the system is represented by the Transverse Field Ising Model [1]:

$$H = \frac{1}{2} \sum_i \Omega_i \sigma_i^x - \sum_i \delta_i n_i + \sum_{i<j} \frac{C_6}{r_{ij}} n_i n_j \tag{5}$$

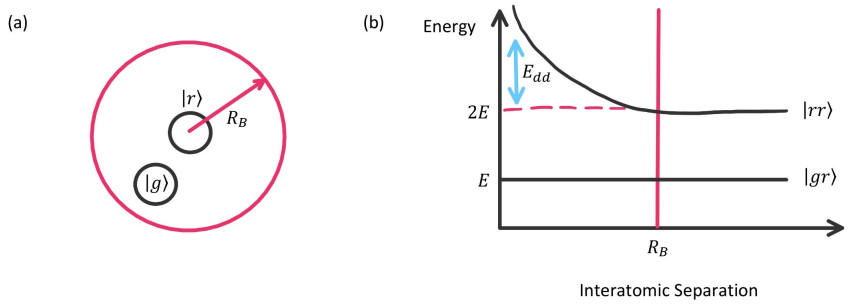


Figure 1: (a) Rydberg blockade in presence of a global Rydberg drive. The red arrow marks the blockade radius R_B and the red circle denotes the blockade area. The atom outside this region is excited to the Rydberg state, but the one contained in it is not. (b) Diagram of the mechanism of the Rydberg blockade effect. When two atoms are within the blockade radius, the dipole-dipole interaction with associated energy E_{dd} light-shifts the Rydberg level of the unexcited atom off-resonance. Once the separation is at least R_B, the light shift is eliminated, so the drive is resonant with both atoms and the double excitation is possible.

where Ω_i and δ_i are the Rabi frequency and the detuning for the $|g\rangle \rightarrow |R\rangle$ transition for the i-th atom, respectively, r_{ij} is the atomic separation, C_6 is the coefficient describing the strength of the Rydberg interaction, and n_i is the Rydberg occupancy for the i-th atom. An expression for the blockade radius comes straightforwardly from observation: above R_B, the system's energy should simply be

$$H = \frac{1}{2}\sum_i \Omega_i \sigma_i^x - \sum_i \delta_i n_i \tag{6}$$

Then, we can write that the blockade radius is given by

$$R_B = \left(\frac{C_6}{\Omega + \delta}\right)^{-1/6} \tag{7}$$

As predicted, we have great control over the blockade radius, which is inherited by flexibility of Ω and δ, and, depending on the choice of computer, C_6 via changing the target Rydberg state.

The MWIS problem requires the application of weights to vertices in the graph. There are a few ways to do this in neutral-atom system. The most experimentally simple way is to rearrange the atoms such that they are either contained in or excluded from their neighbors' blockade radii, with atomic separation taking the role of the weights. The ability to easily do this rearrangement is a hallmark of the neutral atom system. Then, the atoms would be excited to Rydberg states or blockaded by a uniform, global drive.

You can also consider a more experimentally challenging approach: applying local detunings to modify the blockade radius per atom rather than relying on varying their spacings. Figure 2 illustrates these possibilities graphically.

The two cases have different blockade radii than that of Equation 7: for the global drive, $R_B^G \approx 0.98 R_B$ and for the drive with local detunings, $R_B^L(\Omega_1, \Omega_2) \approx 0.49(R_B^G(\Omega_1) + R_B^G(\Omega_2))$, or approximately the average of the blockade radii of the atoms individually [1]. One the weights are applied with either method, set independence is guaranteed by ensuring that vertices connected by an edge are blockaded. For our proposal, we are choosing to use local drives, as these are predicted to perform better than the global drive for this application [1].

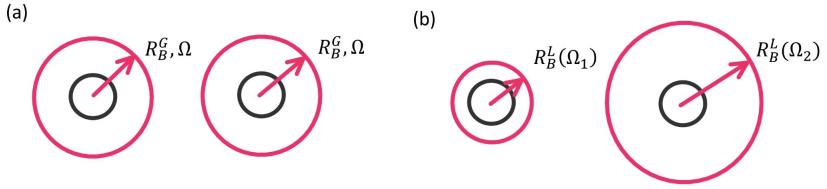


Figure 2: (a) Global addressing of qubits (b) local addressing of qubits

3.1 Quantum Adiabatic Algorithm

The physical principle permitting the application of neutral atom computers to this problem is the adiabatic theorem, which, in short, states that a system prepared in its ground state that is time-evolved adiabatically

will remain close to its ground state [2]. This requires a careful choice of parameters, which we will review in Section 4.2.

4 Hybrid Classical-Quantum Solution

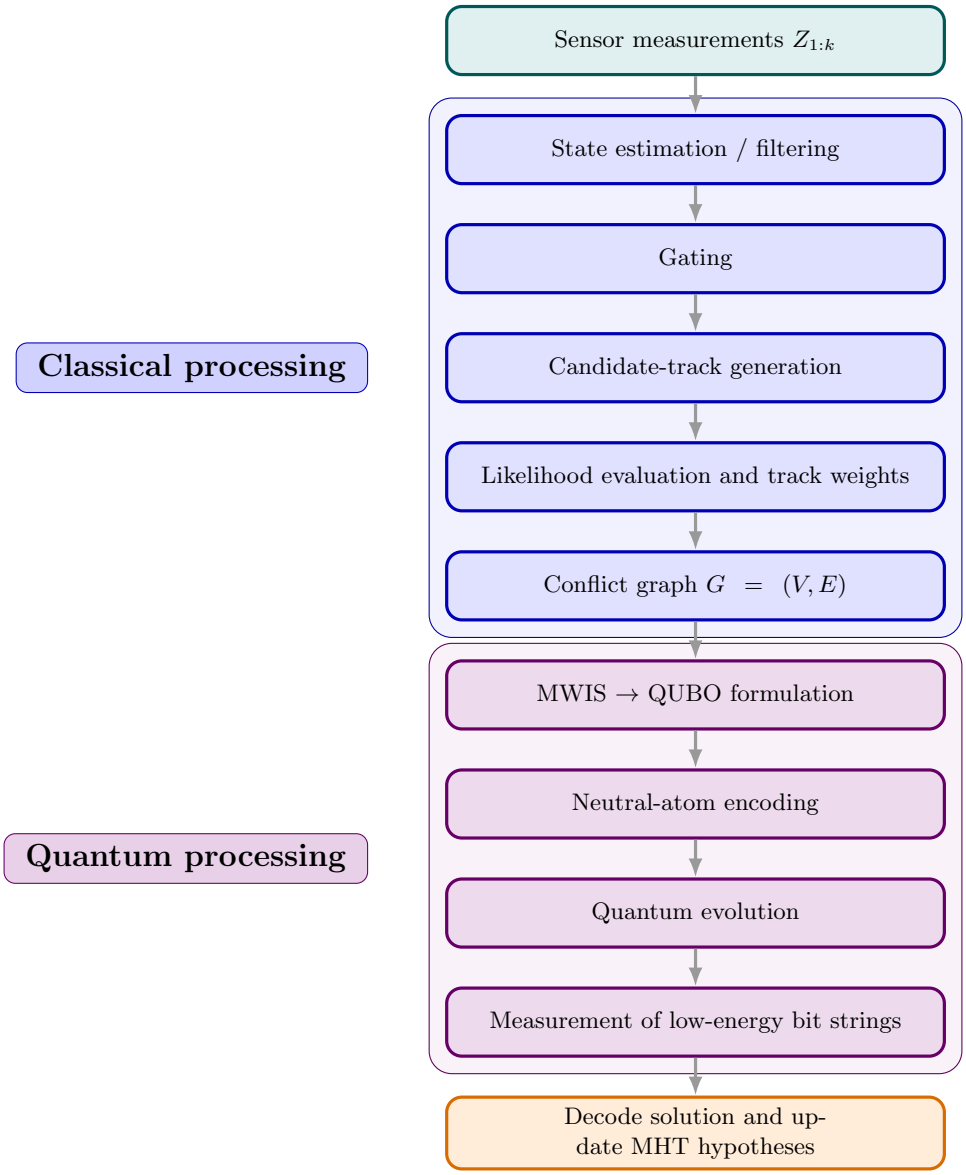


Figure 3: Hybrid classical–quantum workflow for multi-target tracking. Classical processing constructs and scores candidate track hypotheses, which are converted into a conflict graph and then mapped to an MWIS/QUBO problem. mm The quantum stage encodes and evolves the optimization problem, after which low-energy bit strings are measured and decoded to update the MHT hypotheses.

The proposed solution follows a hybrid classical-quantum architecture. The classical component performs the conventional tracking operations required to transform noisy sensor measurements into a weighted data-association problem. In particular, state estimation, gating, candidate-track generation, and likelihood evaluation are performed classically. The resulting data-association problem is represented as a Maximum Weight Independent Set (MWIS) problem and subsequently converted into a Quadratic Unconstrained Binary Optimization (QUBO) problem.

The quantum component then maps this QUBO formulation onto a programmable neutral-atom system. Each candidate track is represented by a neutral atom, candidate-track weights are encoded through local detunings of the Rydberg drive. The neutral-atom system is then evolved toward low-energy configurations, which are measured and decoded into candidate global tracking hypotheses.

The complete workflow is illustrated in Figure 3.

4.1 Classical Component

Let

$$Z_{1:k} = \{Z_1, Z_2, \ldots, Z_k\} \tag{8}$$

denote the sequence of sensor measurements available up to time step k. The objective of Multiple Hypothesis Tracking is to determine the most probable global association hypothesis,

$$H^\star = \arg\max_H P(H \mid Z_{1:k}), \tag{9}$$

where a global hypothesis H specifies a mutually consistent set of target tracks.

Direct evaluation of Eq. (9) becomes computationally difficult because the number of possible measurement-to-track associations grows combinatorially with the number of targets, measurements, and observation times. The classical part of the proposed approach therefore first performs state prediction, gating, and candidate-track generation in order to restrict the set of hypotheses passed to the optimization stage.

For every existing track i, we use a filter to produce a predicted state $\hat{\mathbf{x}}_i^-$ and corresponding covariance P_i^- before the arrival of the next measurement. For a new detection $\mathbf{z}_j$, the innovation is

$$\boldsymbol{\nu}_{ij} = \mathbf{z}_j - H\hat{\mathbf{x}}_i^-, \tag{10}$$

and the corresponding innovation covariance is

$$S_i = HP_i^-H^T + R, \tag{11}$$

where R denotes the sensor measurement-noise covariance.

A gating operation is used to remove associations that are inconsistent with the predicted target state. For each possible track–measurement pair, a distance d_{ij} is evaluated. The association between track i and measurement j is retained only when

$$d_{ij}^2 \le \gamma, \tag{12}$$

where γ is a chosen gating threshold. This procedure considerably reduces the number of candidate associations and consequently the size of the subsequent combinatorial optimization problem.

For each association surviving the gating stage, a likelihood, $p(\mathbf{z}_j \mid T_i)$, is assigned based on the agreement between the predicted state and the observed measurement.

Logarithmic scores are convenient to work with. A candidate-track weight can therefore be written generically as

$$w_i = \log L(T_i), \tag{13}$$

where $L(T_i)$ denotes the likelihood, or likelihood ratio, associated with candidate track T_i. In a complete MHT implementation, this score may also include factors associated with the probability of detection, missed detections, and clutter.

The surviving candidate tracks are represented by a weighted conflict graph

$$G = (V, E). \tag{14}$$

Each vertex $i \in V$ represents a candidate track T_i and is assigned a weight w_i. An edge $(i, j) \in E$ is introduced whenever candidate tracks T_i and T_j cannot simultaneously belong to the same global hypothesis, for example because they share the same sensor observation.

The data-association problem can consequently be formulated as a Maximum Weight Independent Set problem,

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{i=1}^{N} w_i x_i, \\ \text{subject to} \quad & x_i + x_j \le 1, \qquad \forall (i,j) \in E, \\ & x_i \in \{0,1\}. \end{aligned} \tag{15}$$

Here, $x_i = 1$ indicates that candidate track T_i is included in the selected global hypothesis, whereas $x_i = 0$ indicates that it is excluded.

For compatibility with the quantum optimization formulation, the constrained MWIS problem is converted into a QUBO problem by incorporating the incompatibility constraints as quadratic penalty terms:

$$E(\mathbf{x}) = -\sum_{i=1}^{N} w_i x_i + \lambda \sum_{(i,j)\in E} x_i x_j, \tag{16}$$

where $\lambda > 0$ is chosen sufficiently large such that the simultaneous selection of two incompatible tracks is energetically unfavorable.

The optimization is therefore

$$\mathbf{x}^\star = \arg\min_{\mathbf{x}\in\{0,1\}^N} E(\mathbf{x}). \tag{17}$$

Equation (16) may equivalently be expressed in the standard QUBO form

$$E(\mathbf{x}) = \mathbf{x}^T Q \mathbf{x}, \tag{18}$$

where, using a symmetric QUBO matrix,

$$Q_{ii} = -w_i \tag{19}$$

and

$$Q_{ij} = Q_{ji} = \begin{cases} \lambda/2, & (i,j) \in E, \\ 0, & \text{otherwise.} \end{cases} \tag{20}$$

The factor of $1/2$ appears because both $Q_{ij}x_ix_j$ and $Q_{ji}x_jx_i$ contribute to $\mathbf{x}^T Q\mathbf{x}$.

4.2 Quantum Encoding

First, we can verify that the QUBO Hamiltonian can be encoded in the quantum system. As mentioned in Section 3, each neutral atom in the Quantum Processing Unit (QPU) register will represent a vertex of the graph, and edges are created by the Rydberg blockade. Comparing the QUBO Hamiltonian (Equation 16) to the Rydberg Hamiltonian (Equation 5), setting $\delta = 0$ (driving on resonance), we have

$$H_c = -\sum_{i=1}^{N} w_i x_i + \lambda \sum_{(i,j)\in E} x_i x_j$$
$$H_q = \frac{1}{2}\sum_i \Omega_i \sigma_i^x + \sum_{i<j} \frac{C_6}{r_{ij}} n_i n_j$$

Clearly, these have the same structure, with half the Rabi frequency playing the role of the weight, and $x_i \sim n_i = |R_i\rangle\langle R_i|$.

The general steps taken by the QPU are as follows:

1. Receive graph instructions (vertex positions and weights) from classical computer; program atom positions (fix vertices) and detuning mask (apply weights)
2. Initialize: place atoms in the register, and intialize each into the ground state
3. Perform the time evolution: apply a slow Rydberg excitation pulse to guarantee adiabaticity
4. Read out the new ground state: report results with located MWIS to the CPU

We have made our best attempts to model this process using Pasqal's Pulser simultator. We figured it would be easier to run initial tests if we used Pulser's Mock Device, which doesn't enforce physical limits on the device such as minimum qubit spacing or maximum Rydberg sequence length. (However, we are still choosing a Rydberg sequence sufficiently short to be consistent with real hardware.) Figure 4 shows the Pulser workflow. With Pulser, we have the ability to create a DMM channel, which is analogous to a diffractive optic (DMD, SLM, etc.) that can locally vary Rydberg laser parameters at the atoms and force them to have different blockade radii. This is executed in the simulation by applying a global Rydberg pulse using the global channel and applying the detunings locally using a local channel and a map of desired detunings. This is once again analogous to the hardware, in which you would globally pulse the Rydberg light on to the diffractive optic, then rely on the diffracted pattern to apply the local effects.

Our approach to this simulation follows Pasqal's own solution to this problem [4] relatively closely. We fix the Rydberg pulse time to be 40 us, which can be made to be sufficiently adiabatic. We bound the Rabi frequency from above with the maximum energy of two non-connected nodes, and we use a similar ramp of the Rabi frequency and detuning. A plot of these pulses is shown in Figure 5. The local detunings are normalized to the highest weight in the graph.

Stepping back to look at the whole quantum workflow, from start to finish, we first optimize atom placement to match graph couplings via constrained minimization with a configurable tolerance. Then we create a Pulser Register that places the atoms in the correct positions and verifies that there are no direct spatial conflicts (if we used a device other than the Mock Device, this would have a stricter failure condition in which atoms could

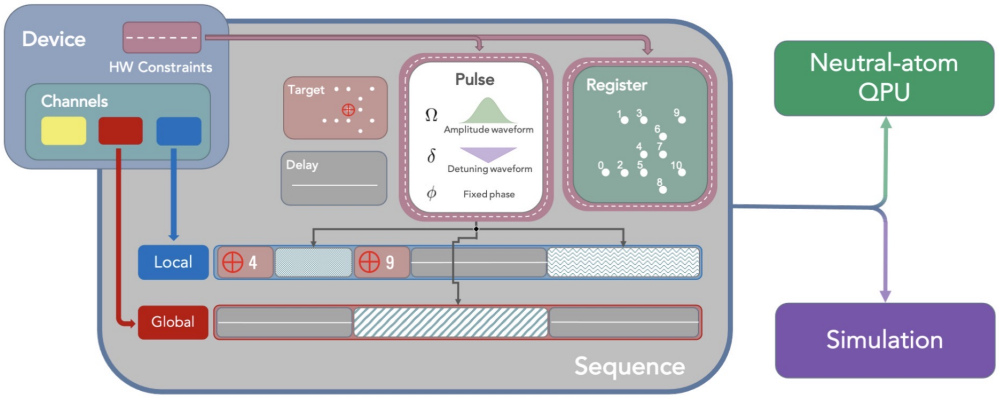


Figure 4: Pulser workflow, from [6]

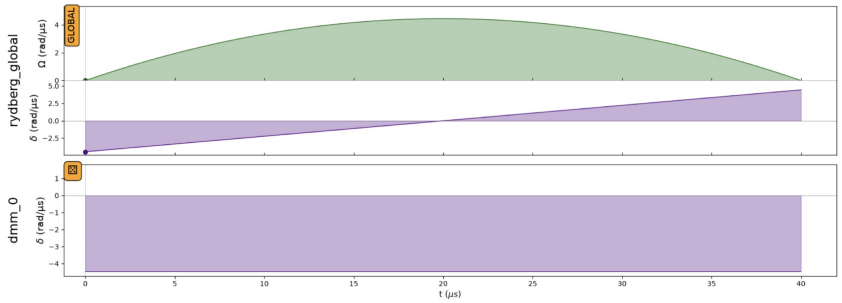


Figure 5: Example pulse sequence for an individual atom for our simulations. Top: global pulse parameters against pulse time. Bottom: local detuning parameters against pulse time. Note that the plot is showing a negative DMM detuning because this particular atom corresponds to a vertex with a small weight.

not be closer than 5 um.) Next, we construct a Pulser sequence that contains the desired pulse sequences and adds weight-biased detuning via the DMM. Then the program runs QuTiP simulation to estimate adiabatic evolution and collect bitstring samples. Finally, we decode the results by reordering bitstrings according to the physical atom layout, filtering out invalid or infeasible solutions by checking for edges between selected nodes, and selects the best feasible outcome.

5 Simulation Results

5.1 Controlled Synthetic Experiment

We first considered a controlled two-dimensional multi-target tracking problem for which the complete ground truth is known. This experiment was designed to test the full MHT–MWIS pipeline.

Synthetic data. Three targets, denoted by T_0, T_1, and T_2, were propagated over four scans, $k = 0, 1, 2, 3$, with sampling interval $\Delta t = 1$. The trajectories were chosen to approach and cross one another so that several measurement-to-track associations become plausible near the centre of the surveillance region. The noise-free positions were generated according to

$$\mathbf{p}_i(k) = \mathbf{p}_{0,i} + \mathbf{v}_i k\Delta t + \frac{1}{2}\mathbf{a}_i(k)(k\Delta t)^2, \tag{21}$$

where

$$\mathbf{a}_i(k) = \begin{bmatrix} 0.05\sin(0.8k + i) \\ 0.04\cos(0.6k + i) \end{bmatrix}. \tag{22}$$

Each target was independently detected with probability $P_D = 0.98$. A successful detection was corrupted by isotropic Gaussian position noise,

$$\mathbf{z}_i(k) = \mathbf{p}_i(k) + \boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 I), \tag{23}$$

with $\sigma_z = 0.55$. False detections were additionally drawn from a Poisson distribution with mean 0.6 per scan and placed uniformly within the simulated surveillance region. Finally, detections were randomly shuffled within

each scan so that their order contained no target-identity information. The resulting realization contained 4, 4, 3, and 4 detections in the four scans, respectively. Figure 6 shows the noise-free trajectories together with the unlabelled measurements supplied to the tracker.

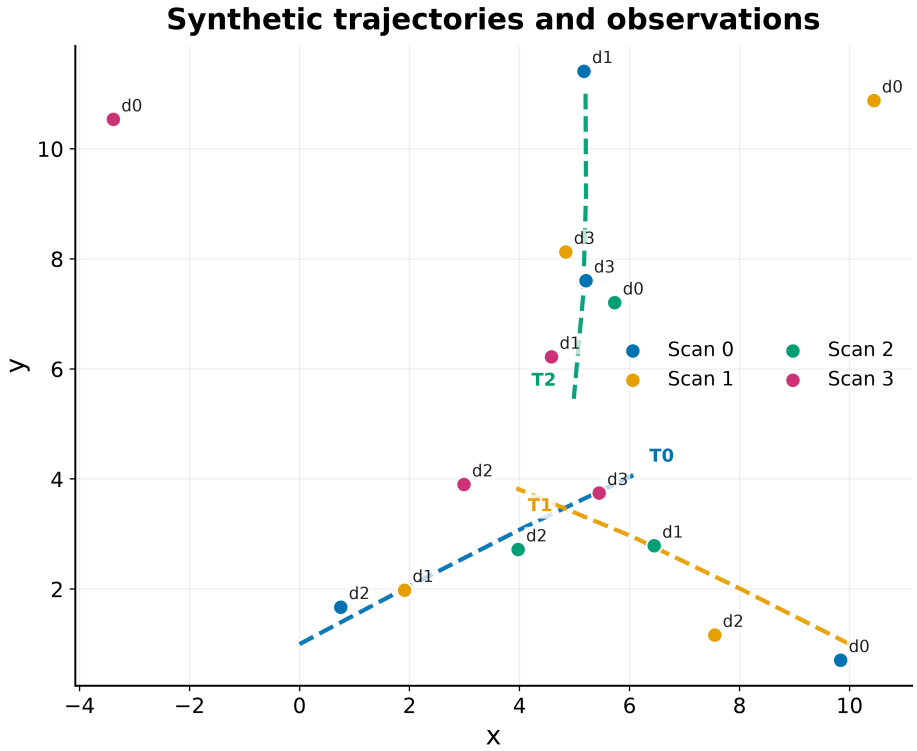


Figure 6: Controlled synthetic tracking scenario. Dashed curves denote the noise-free trajectories of targets T_0–T_2, while markers show the noisy, unlabelled measurements supplied to the tracker.

Kalman prediction and gating. Each first-scan detection initiated a candidate-track family. Track state was represented by

$$\mathbf{x}_k = [p_x, p_y, v_x, v_y]^{\mathrm{T}}, \tag{24}$$

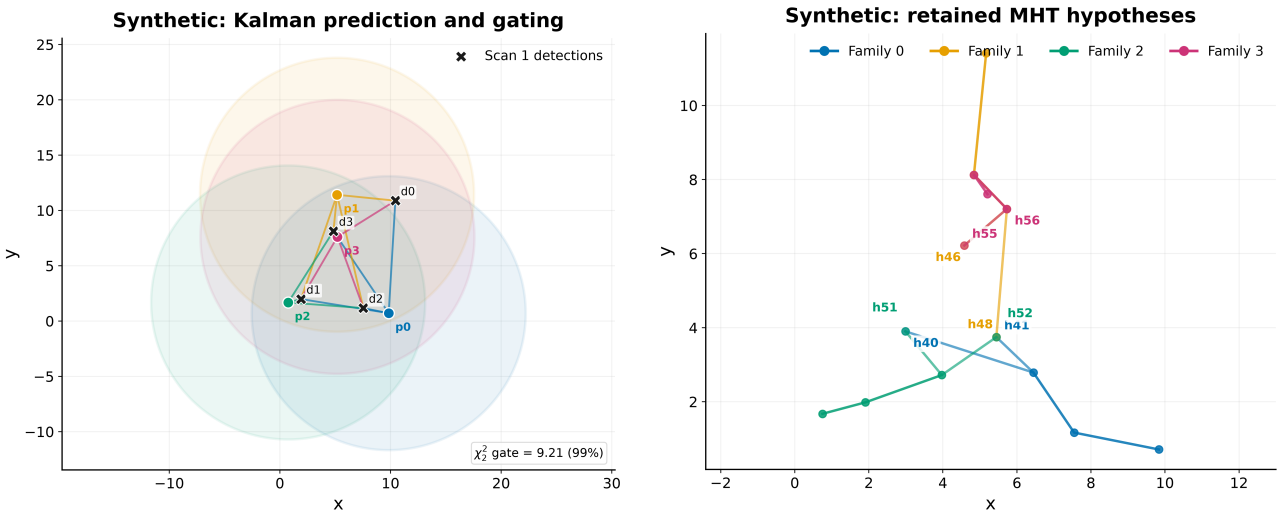
and propagated using the constant-velocity Kalman model described in Section 4.1. Since a single position measurement does not determine velocity, tracks were initialized with zero mean velocity but a large velocity uncertainty. Consequently, the first predicted position is not a confident point estimate: its covariance incorporates substantial uncertainty in the unknown initial velocity.

Candidate measurements were accepted when their squared Mahalanobis distance satisfied

$$d^2 = \boldsymbol{\nu}^{\mathrm{T}} S^{-1} \boldsymbol{\nu} \leq 9.21. \tag{25}$$

For a two-dimensional Gaussian innovation, 9.21 corresponds approximately to the 99% χ^2 confidence region. The resulting uncertainty regions are therefore ellipses determined by the innovation covariance rather than fixed Euclidean-radius circles. Figure 7a shows the first prediction step: scan-1 detections lying inside each ellipse form admissible associations and are connected to the corresponding prediction. The deliberately permissive initial gates preserve plausible motion hypotheses before sufficient measurements have accumulated to estimate target velocity reliably.

Candidate-track hypotheses. Every gated detection produces a continuation branch, while a missed-detection branch is also permitted. Candidate tracks are scored using the cumulative log-likelihood ratio introduced in Section 4.1. To prevent unrestricted exponential growth, only the two highest-scoring branches from each initiating family are retained after each scan. After four scans, this procedure produced eight positive-weight candidate tracks, shown in Figure 7b. Several tracks follow a single physical target consistently, whereas others switch between nearby targets around the crossing region. These alternatives are precisely the ambiguity that must subsequently be resolved at the global-hypothesis level.



(a) Kalman prediction and Mahalanobis gating.

(b) Retained MHT candidate tracks.

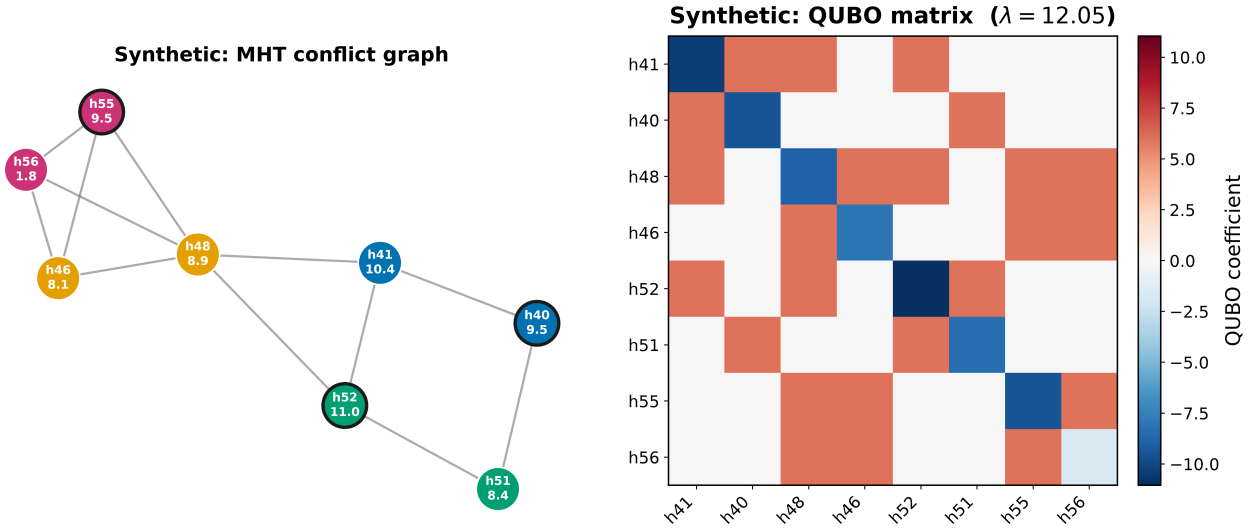
Figure 7: Classical MHT front end for the synthetic experiment. Gating retains statistically plausible measurement associations, from which competing candidate-track hypotheses are constructed.

Conflict graph. The retained candidate tracks produced the weighted conflict graph shown in Figure 8a. The synthetic instance contained eight vertices and twelve conflict edges. All eight candidate tracks belonged to a single connected component, meaning that the ambiguity could not be decomposed into smaller independent MWIS subproblems.

QUBO formulation. The resulting MWIS instance was converted to the QUBO described in Section 4.1,

$$E(\mathbf{x}) = -\sum_i w_i x_i + \lambda \sum_{(i,j)\in E} x_i x_j. \tag{26}$$

For this realization the penalty was chosen as $\lambda = 12.045$, larger than the largest positive candidate weight. In Figure 8b, the negative diagonal entries therefore encode the benefit of selecting high-scoring tracks, while positive off-diagonal entries penalize simultaneously selecting incompatible tracks. Exhaustive verification of this small QUBO instance gave the same optimum as the exact constrained MWIS problem, confirming that the graph-to-QUBO conversion itself did not alter the solution.



(a) Weighted MHT conflict graph.

(b) Corresponding QUBO matrix.

Figure 8: Conversion of the synthetic data-association problem into an MWIS/QUBO instance. Vertices represent candidate tracks and edges encode incompatibility; the QUBO rewards high-weight tracks while penalizing conflicting selections.

Selected global hypothesis. The exact classical MILP solver selected

$$H^{\star}_{\mathrm{syn}} = \{h_{40}, h_{52}, h_{55}\}, \tag{27}$$

with total weight

$$W^{\star}_{\mathrm{syn}} = 30.0824. \tag{28}$$

The three selected tracks correspond predominantly to T_1, T_0, and T_2, respectively, and no identity switches occur among their labelled detections. Figure 9 shows the resulting global tracking hypothesis. The optimum is not isolated from all alternatives: the next three classical global hypotheses have weights 28.6449, 28.3595, and 26.9220. This illustrates the original motivation for MHT—the data can support several plausible global explanations even when one has the largest score.

The reported identity purity should be interpreted carefully. Track h_{55} has purity 1 over its three labelled target detections, but one additional detection is unlabelled and may correspond to clutter. Thus the purity metric measures identity consistency among labelled detections rather than certifying that every point on a selected track originated from a true target.

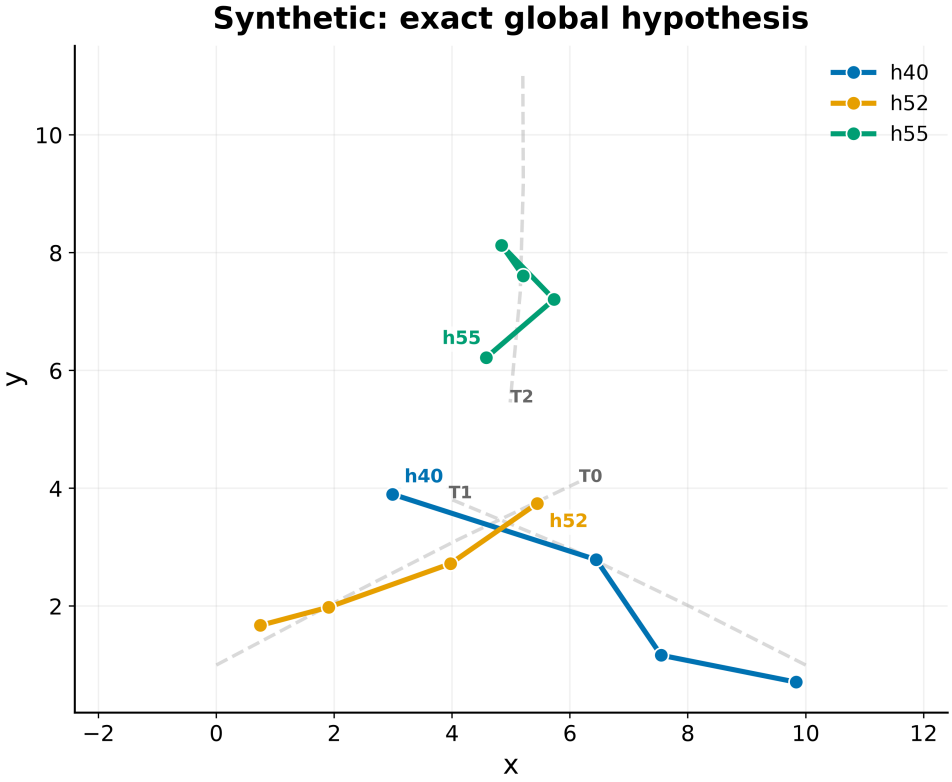


Figure 9: Exact classical MWIS solution for the synthetic experiment, selecting the mutually compatible candidate tracks h_{40}, h_{52} and h_{55}.

Neutral-atom embedding. The eight-vertex connected component was then mapped onto an eight-atom two-dimensional register. Candidate-track weights were encoded through the detuning modulation, while the desired conflict penalties were represented through the distance-dependent Rydberg interaction. The optimized register is shown in Figure 11a. Although the abstract QUBO was correct, the resulting geometric embedding was poor: the normalized embedding loss was approximately 0.534. This distinction is important because the neutral-atom device does not independently program an arbitrary interaction coefficient for every graph edge. Instead, every atom pair interacts according to its physical separation.

Desired and physical interactions. Figure 11b makes this discrepancy explicit by comparing the desired conflict adjacency with the interaction matrix generated by the optimized atomic coordinates. In the ideal graph, incompatible tracks should receive strong pairwise penalties and compatible tracks should receive negligible penalties. In the physical encoding, however, the interactions vary continuously as C_6/R_{ij}^6 for all atom pairs. The embedding must therefore satisfy many geometric constraints simultaneously. For this eight-vertex graph, some desired edges remain too weak while some desired non-edges acquire appreciable interactions. The very poor edge/non-edge separation observed for this component is consequently a direct warning that the physical Hamiltonian does not faithfully reproduce the intended abstract conflict graph.

Quantum sampling. The embedded instance was evolved using the Pulser neutral-atom emulator and the final state was sampled in the computational basis. Figure 10 shows the dominant output bit strings. In this case, none of the sampled probability mass corresponded to a feasible independent set: both the measured feasible probability and the probability assigned to the exact MWIS were zero. The dominant bit string was `11000110` with probability approximately 0.798, but it simultaneously selected mutually incompatible candidate tracks and was therefore not a valid MHT global hypothesis.

Consequently, the operational hybrid pipeline invoked its explicit classical fallback and returned the exact weight 30.0824. The resulting hybrid-to-classical objective ratio of 1 must therefore *not* be interpreted as a successful quantum solution. Rather, this experiment provides a useful negative result: the classical MHT, graph, MWIS, and QUBO stages agree, but the eight-vertex conflict graph is not represented sufficiently accurately by the chosen two-dimensional C_6/R^6 geometry. This failure motivates the embedding analysis in Section **??**.

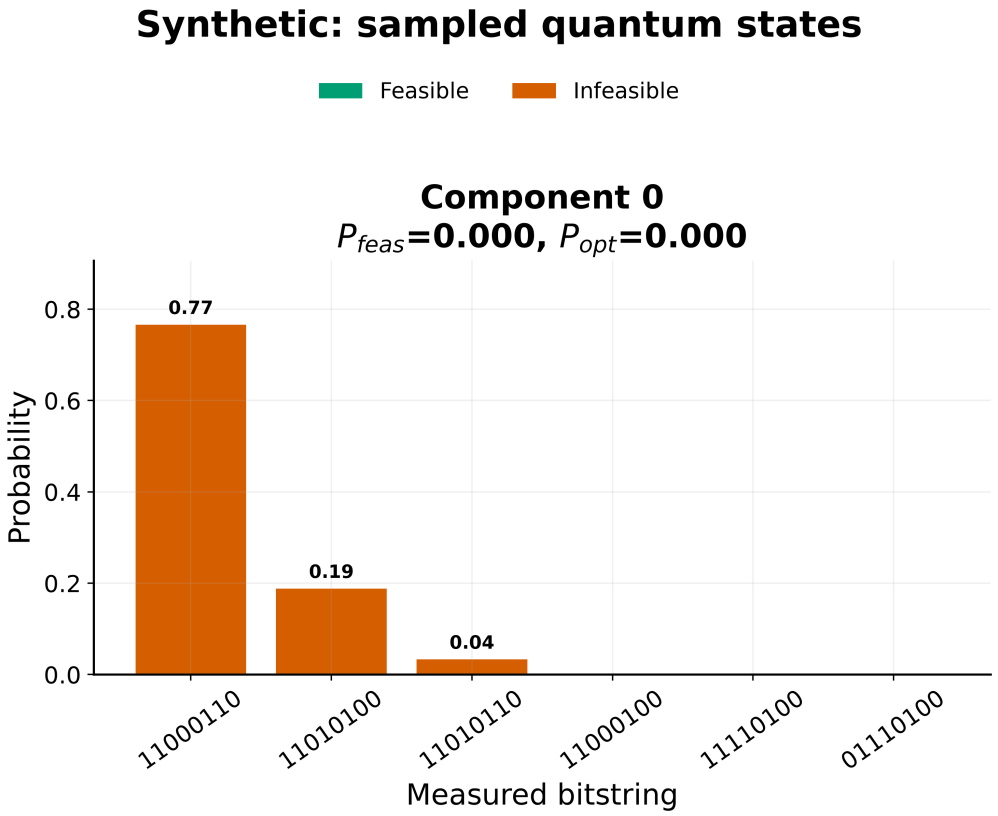


Figure 10: Dominant neutral-atom bit strings sampled for the synthetic instance. The highest-probability states are infeasible independent sets, and no sampled probability mass is assigned to the exact MWIS solution.

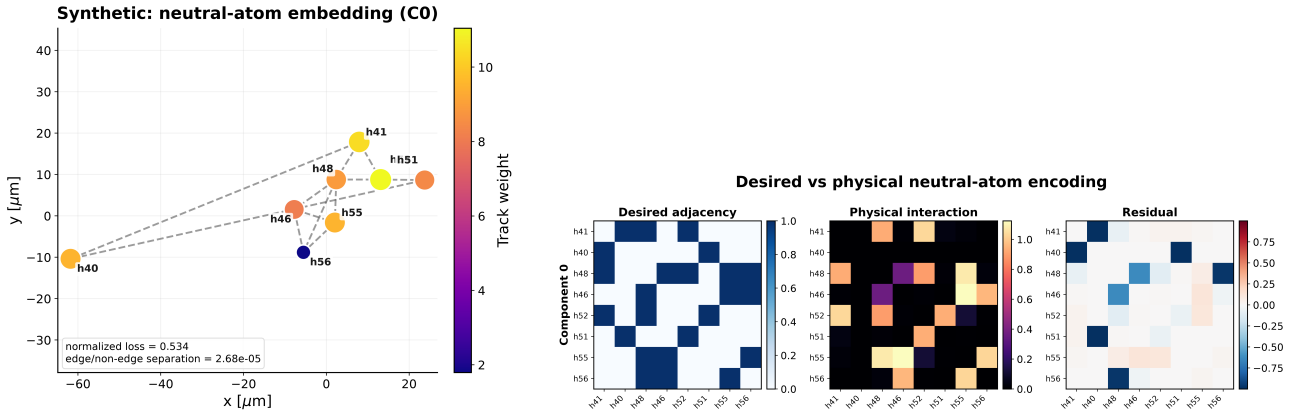
Sensitivity to measurement noise. Finally, we investigated how measurement uncertainty affects the difficulty of the classical data-association problem. To isolate this effect, the sweep used three targets over five scans with no clutter and unit detection probability. The measurement standard deviation was varied over

$$\sigma_z \in \{0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.5, 2.0\},$$

with twenty independent trials at each value. Figure 12 summarizes the resulting graph complexity and tracking quality.

At $\sigma_z = 0.2$, the conflict graph contained on average only 3.05 vertices and 0.15 edges, and all trials achieved perfect association. Increasing the measurement noise caused more measurements to remain plausible under competing tracks, rapidly increasing the number of retained hypotheses and conflict edges. By $\sigma_z = 0.6$, the graph contained on average 8.65 vertices and 21.55 edges. Above this point the vertex count approached a plateau near nine because the tracker retained at most three branches per family in this sweep; the number of conflict edges nevertheless continued to increase, reaching approximately 29.65 at $\sigma_z = 2.0$.

The same increase in ambiguity was reflected in association quality. Mean identity purity decreased from 1.00 at $\sigma_z \leq 0.6$ to approximately 0.943 at $\sigma_z = 0.8$, 0.868 at $\sigma_z = 1.0$, and 0.690 at $\sigma_z = 2.0$. Mean identity switches increased from zero to approximately 3.5 over the same range, while the fraction of perfectly associated trials fell from 1.0 to 0.05. The transition around $\sigma_z \approx 0.8$ is specific to the geometry and parameters of this synthetic experiment and should not be interpreted as a universal noise threshold. The broader result is that



(a) Optimized atom register.

(b) Desired and realized interactions.

Figure 11: Neutral-atom encoding of the synthetic conflict graph. The geometric C_6/R^6 interaction pattern does not accurately reproduce the desired conflict adjacency, explaining the poor quantum performance.

greater measurement uncertainty creates more competing MHT explanations, producing denser conflict graphs and a harder subsequent MWIS problem.

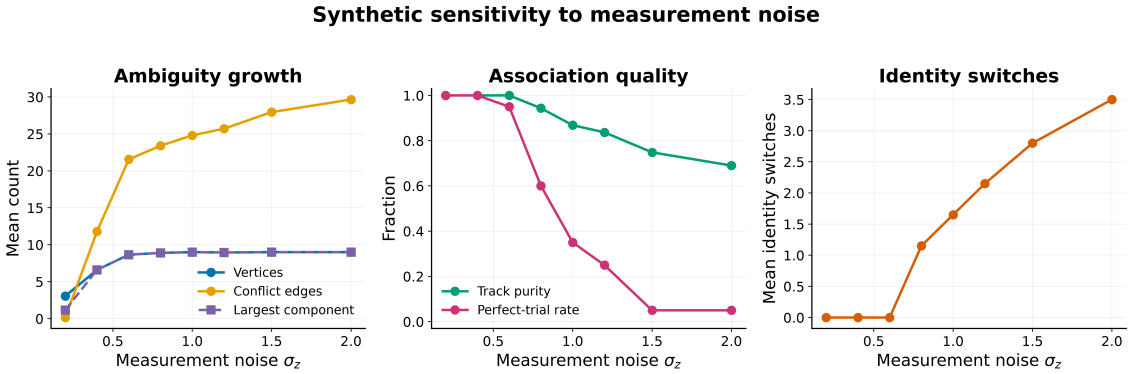


Figure 12: Sensitivity of the synthetic tracking problem to measurement noise. Increasing σ_z increases association ambiguity and conflict-graph density while degrading track purity and increasing identity switches.

5.2 Real-Data Experiment on RADIATE

We next applied the same MHT–MWIS pipeline to real automotive radar data from the RADIATE dataset [5]. The purpose of this experiment was not to evaluate object detection itself, but to test the data-association and optimization stages using measurements derived from a real sensor sequence.

Radar data and oracle detections. We used a short portion of the RADIATE `tiny_foggy` sequence recorded by the Navtech CTS350-X scanning radar. Four consecutive radar scans were selected at the native sampling interval of approximately 0.25 s. The four selected annotation frames were 10–13. Figure 13 presents the corresponding Cartesian radar images in a 2×2 grid.

The supplied RADIATE radar annotations were treated as *oracle detections*. For a bounding box with upper-left position (x, y), width w, and height h, its centre

$$(x_c, y_c) = \left(x + \frac{w}{2},\, y + \frac{h}{2} \right) \tag{29}$$

was converted from image pixels to an approximate two-dimensional radar-frame position in metres. Only vehicle-like classes were retained, and detections outside a 70 m radial range were discarded. At most four detections could be retained per scan; for the selected window, three vehicle detections were present in each of the four scans.

Persistent annotation identities were deliberately hidden from the Kalman filter, gating procedure, candidate-track construction, and MWIS solver. They were restored only after optimization to evaluate the selected tracks. Thus this experiment should be interpreted specifically as an *oracle-detection radar data-association*

experiment. The raw radar images provide the real sensor context, whereas the MHT algorithm itself operates on the annotation-derived bounding-box centres after conversion to unlabelled two-dimensional radar-frame positions.

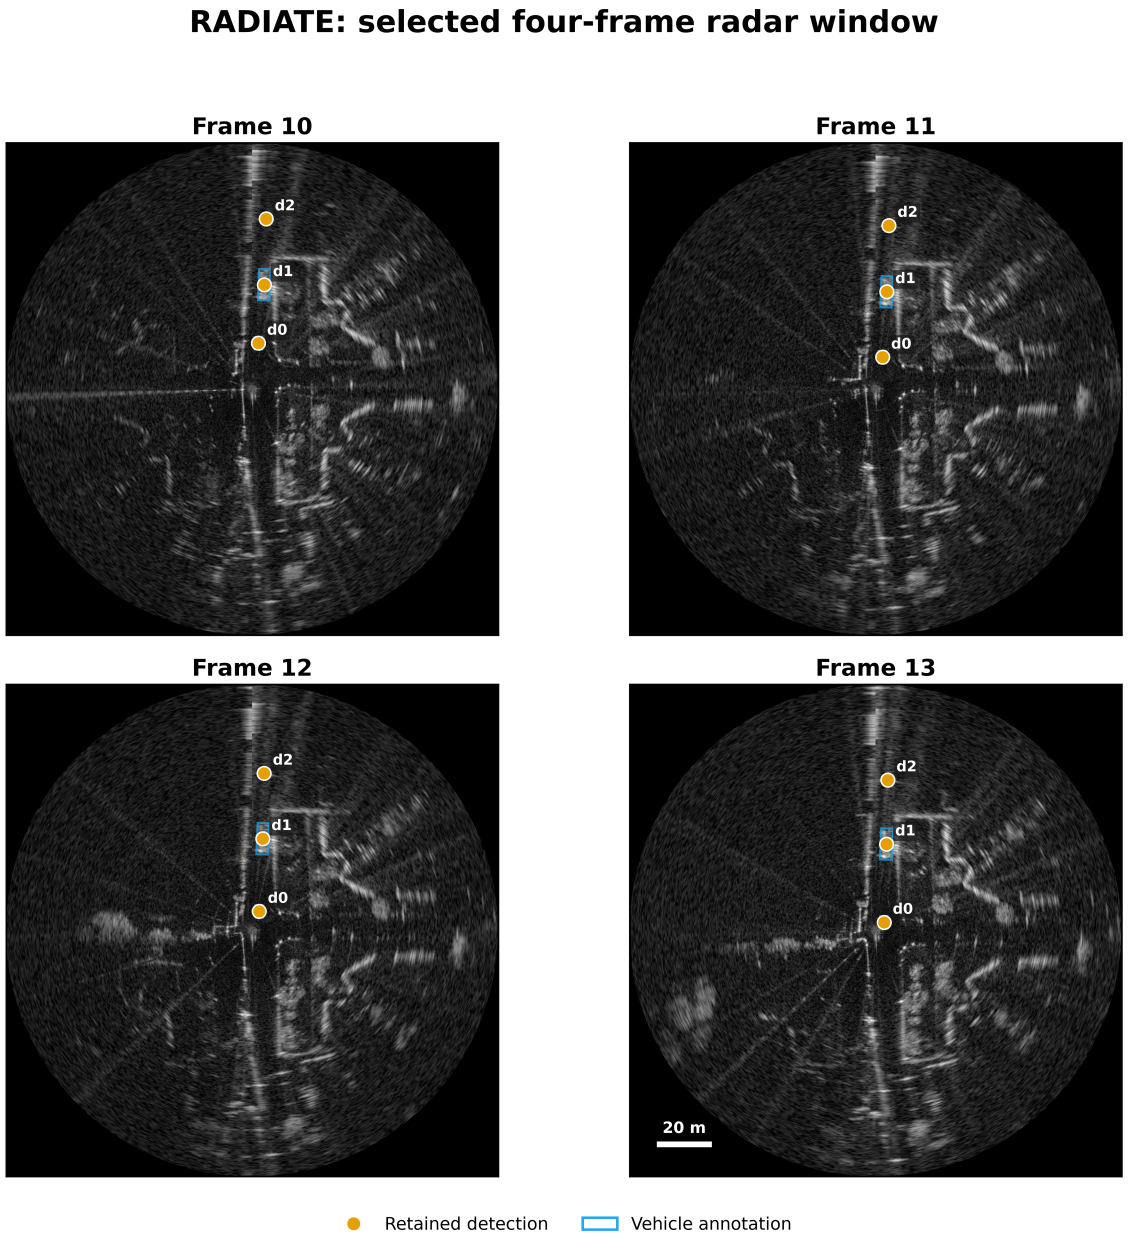


Figure 13: Four consecutive Cartesian radar frames from the selected RADIATE `tiny_foggy` window. Annotation-derived vehicle centres are used as oracle detections for the subsequent data-association experiment.

Kalman prediction and gating. The same constant-velocity state model used in the synthetic experiment was applied to the radar-frame positions, now with $\Delta t = 0.25\,\mathrm{s}$, measurement standard deviation 1.2 m, and process-noise standard deviation 3.0. A 99% Mahalanobis gate, $d^2 \leq 9.21$, was again used to reject implausible scan-to-scan associations. Figure 15a shows the first prediction and gating transition. The ellipses represent the complete predicted measurement uncertainty, while the retained links indicate measurements sufficiently consistent with each prediction to generate an MHT branch.

Track branching and LLR scoring were then applied exactly as in the synthetic case. The four-frame window generated six final candidate-track hypotheses. One candidate had negative LLR and was excluded from the positive-weight MWIS graph, leaving five optimization vertices. The retained candidate tracks are shown in Figure 14. Unlike the synthetic crossing scenario, this RADIATE window contained no close vehicle pairs under the window-selection criterion, so the remaining ambiguity was comparatively mild and arose mainly from competing continuation and missed-detection branches.

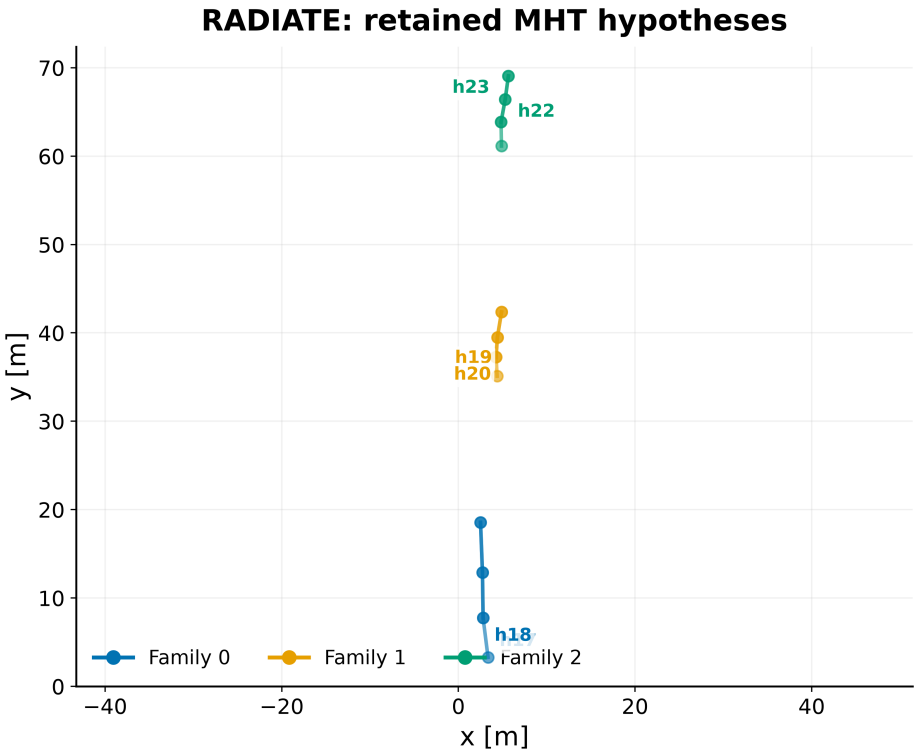


Figure 14: Retained candidate-track hypotheses for the RADIATE experiment, shown in radar-frame metric coordinates.

Conflict graph. The five positive-weight tracks produced only two conflict edges, as shown in Figure 15b. The graph decomposed into connected-component sizes

$$2,\ 2,\ 1. \tag{30}$$

The isolated vertex represents an unambiguous track and requires no quantum optimization. Each two-vertex component contains a pair of mutually exclusive candidate tracks. This decomposition is significant for the hybrid approach: only the genuinely ambiguous components need to be passed to the neutral-atom solver.

QUBO formulation. Each non-trivial connected component was independently converted to the same MWIS QUBO used in the synthetic experiment. Both non-trivial components contained only two mutually incompatible candidate tracks, so the resulting QUBO instances were considerably simpler than the eight-vertex synthetic problem.

Selected global hypothesis. Solving the full conflict graph exactly gave

$$H^{\star}_{\mathrm{rad}} = \{h_{17}, h_{19}, h_{22}\}, \tag{31}$$

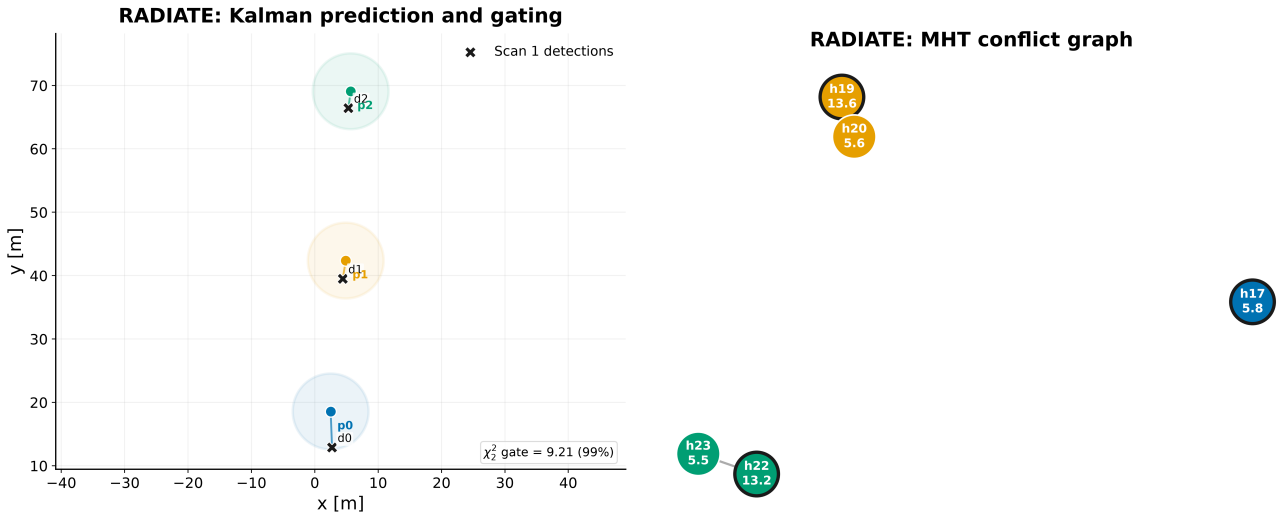
with total MWIS weight

$$W^{\star}_{\mathrm{rad}} = 32.6169. \tag{32}$$

All three selected tracks remained associated with a single persistent annotation identity across their labelled measurements, giving identity purity 1 and zero identity switches. Figure 16 shows the detections belonging to the selected global hypothesis overlaid on the original radar frames, providing a direct visual comparison between the optimization result and the underlying sensor observations.

The optimum was also substantially better separated from the alternatives than in the synthetic experiment. The normalized score mass assigned to the highest-scoring classical global hypothesis was approximately 0.996, showing that this particular four-frame radar window presented a relatively unambiguous association problem.

Neutral-atom embeddings. The isolated one-vertex component was solved classically as a trivial case, whereas the two non-trivial two-vertex components were mapped to separate two-atom Pulser instances. Figure 17a shows the corresponding atomic registers. Both mappings achieved normalized embedding losses of approximately 9.9×10^{-9}, several orders of magnitude below that of the eight-atom synthetic mapping. For a two-vertex graph containing a single conflict edge, the desired interaction can be reproduced almost exactly by placing the two atoms at an appropriate separation.



(a) Prediction and gating.

(b) Resulting conflict graph.

Figure 15: RADIATE tracking ambiguity after classical preprocessing. Gating generates the candidate associations, while the final conflict graph decomposes into two two-track ambiguity components and one isolated track.

Quantum sampling. Figure 17b shows the sampled bit-string distributions for the two non-trivial components. The improved geometric encoding was reflected in the neutral-atom output. For the first two-track ambiguity component, the bit string selecting the classically optimal track was sampled with probability approximately 0.998, while the competing one-track solution received approximately 0.002. For the second component, the optimal bit string was sampled with probability approximately 0.992, while the principal alternative received approximately 0.007. In both components, the total feasible probability was 1.0.

Combining the two quantum components with the isolated deterministic track produced the same global set

$$\{h_{17}, h_{19}, h_{22}\}$$

as the exact classical solver, with total weight 32.6169. Unlike the synthetic experiment, this result did not require a classical fallback for either non-trivial component. The dominant combined sampled global hypothesis had probability approximately 0.990.

These sampling probabilities should not be interpreted as Bayesian MHT posterior probabilities. They describe the output distribution of the simulated neutral-atom optimization and are used here to quantify how reliably the encoded Hamiltonian produces the desired MWIS solution.

Sensitivity and scope. A separate parameter sweep was not performed for the RADIATE example; the controlled measurement-noise sensitivity analysis was intentionally confined to the synthetic data, where measurement uncertainty could be varied independently of clutter, missed detections, and sensor-scene changes. The selected RADIATE window should also be regarded as a compact proof-of-concept rather than a large-scale tracking benchmark: it contains only three retained vehicles per scan and decomposes into two very small ambiguity components.

Nevertheless, the experiment provides a useful complement to the more difficult synthetic instance. When the MHT conflict graph decomposes into small components whose adjacency can be represented accurately by the two-dimensional Rydberg interaction geometry, the neutral-atom simulation recovers the exact MWIS with high probability. The contrast with the synthetic case therefore suggests that geometric embeddability, rather than graph size alone, is a central consideration for this approach.

5.3 Synthetic Data Analysis and Quantum Reliability

While gating, pruning and clustering are effective classical pre-processing techniques to induce sparsity in the MWISP graph, translating this graph to real neutral atom architecture does not follow straightforwardly. For an undirected graph G to be sparse it needs to have much fewer edges E than the maximum possible number of edges $V(V-1)/2$, typically $\mathcal{O}(V)$. The inspiration for this problem is to use the proximity of atoms in physical space to implement edges through the Rydberg blockage, however a packing issue persists that sparsity does not prevent. Consider a star graph $G*$ with a single central vertex and many edge vertices that are only connected to the central vertex. Such a graph is sparse, however the spacing between neighboring vertices at a

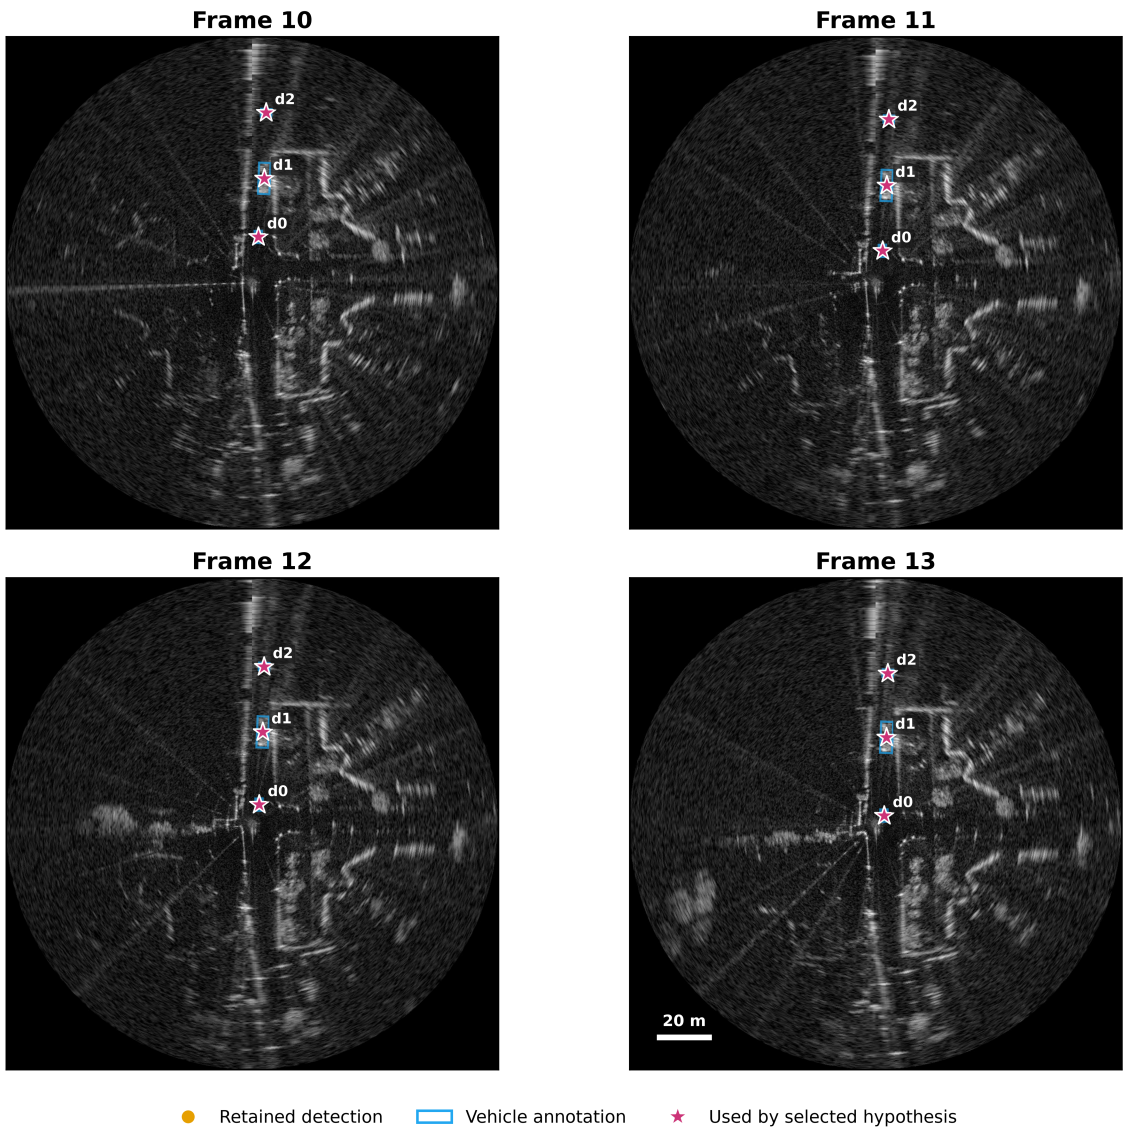
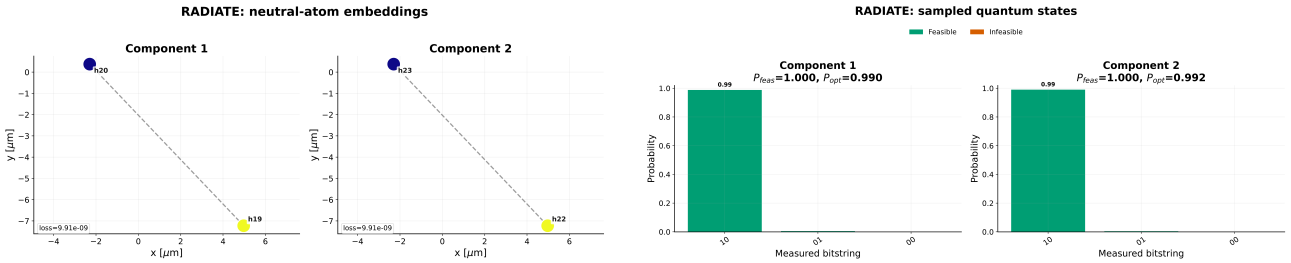


Figure 16: Selected global tracking hypothesis overlaid on the four RADIATE radar frames. Persistent annotation identities are shown only for evaluation and were not supplied to the tracker.

fixed radius decreases monotonically with increasing vertex count. What this means for embedding our MWISP graph into the neutral atom architecture is that atoms representing certain vertices that do not share an edge may be forced to be close enough to neighbors that they interfere with one another via Rydberg interactions. While the Rydberg interaction can be varied it is not a completely free parameter and the problem ultimately persists. This leads to graphs that cannot be completely embedding in the neutral atom architecture. There may be workarounds that utilize additional qubits and sophisticated radius variability that we leave for future work. We do however implement basic location optimization of the neutral atoms to reduce the likelihood of conflicting with the ideal graph embedding where possible. An example of the physical constraints is visible in Fig. 19.



(a) Two-atom embedding.

(b) Quantum sampled solutions.

Figure 17: Neutral-atom solution of a RADIATE ambiguity component. The two-vertex conflict graph admits an essentially exact geometric embedding, and the optimal independent set dominates the sampled distribution.

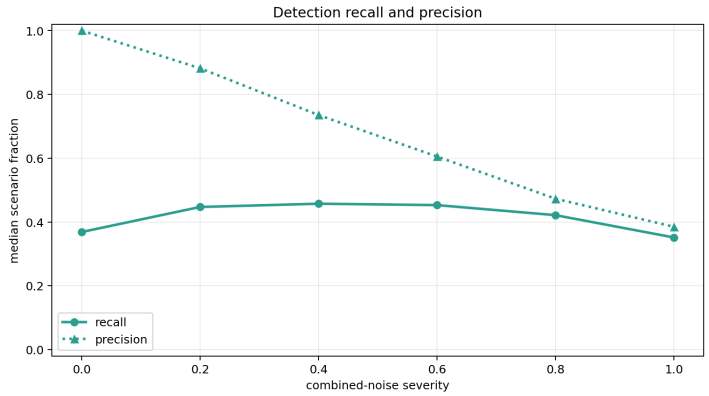


Figure 18: The median precision and recall of the image detection step at varying noise levels of the simulated data.

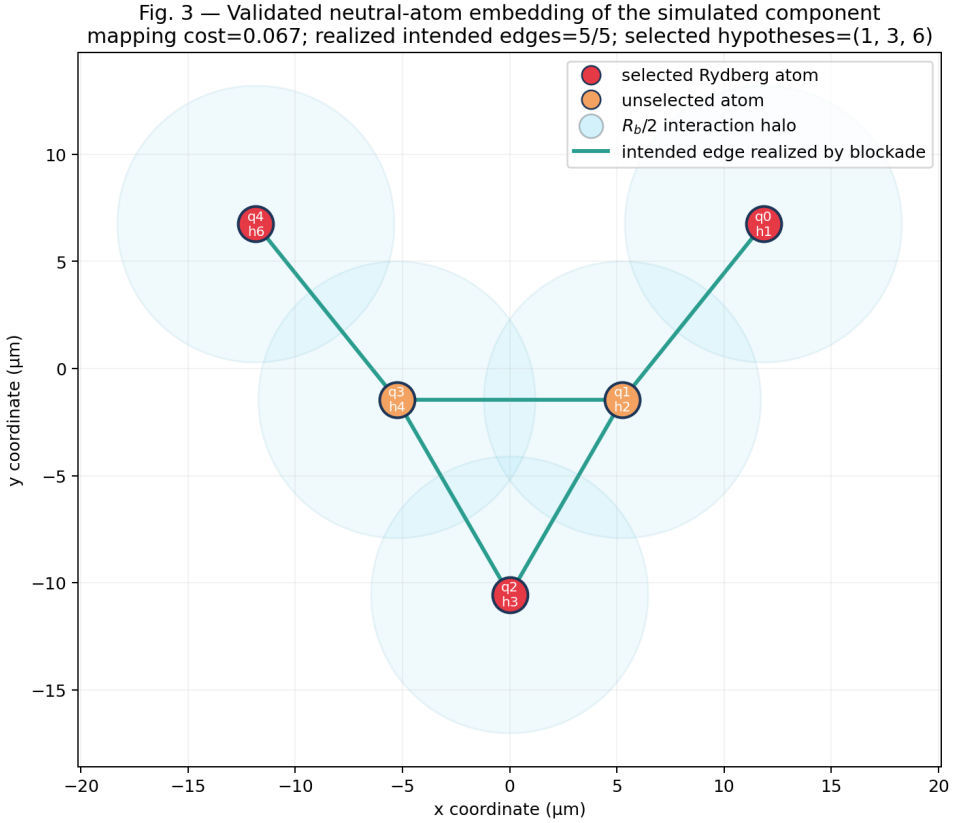


Figure 19: An example embedding of a particular frame of the MWISP conflict graph into the target neutral atom architecture.

We validate the image detection and synthetic data generation by demonstrating the precision and recall metrics against noise, as well as by visually inspect the output of the image detection step in Fig. 22. In Fig. 18 we see that precision, which is the fraction of positive detections that are true positives decreases steadily as noise increases. This is consistent with expectation, essentially the image detection process is interpreting noise as signal. The recall which is defined as the number of true positives over the number of true positives and false negatives is a target for minimizing false negatives. The recall shows a less clear trend, though the sample size is only over 100 iterations of 30 objects at each noise level. The general picture here is that the object detection step is consistently missing certain objects.

Using simulated data we were able to vary parameters such as the number of nodes, and image noise to explore how frequently the graph embedding was incompatible with the real neutral atom architecture. We demonstrate this in Fig. 20. In this figure we see that in both plots that when the graph is successfully implemented the optimal solution is found, as no completed suboptimal cases were detected in either subplot. Additionally, we see on the left subplot that increased noise-severity decreases the number of successful graph embedding at a fixed size $n = 18$. This fits with our expectations as noisy data is more likely to contain many lower likelihood competing hypotheses where geometric constrains could interfere. The right subplot demonstrates a similar trend but with a fixed low noise level and increasing numbers of vertices, which similarly increases the likelihood of breaking geometric constraints.

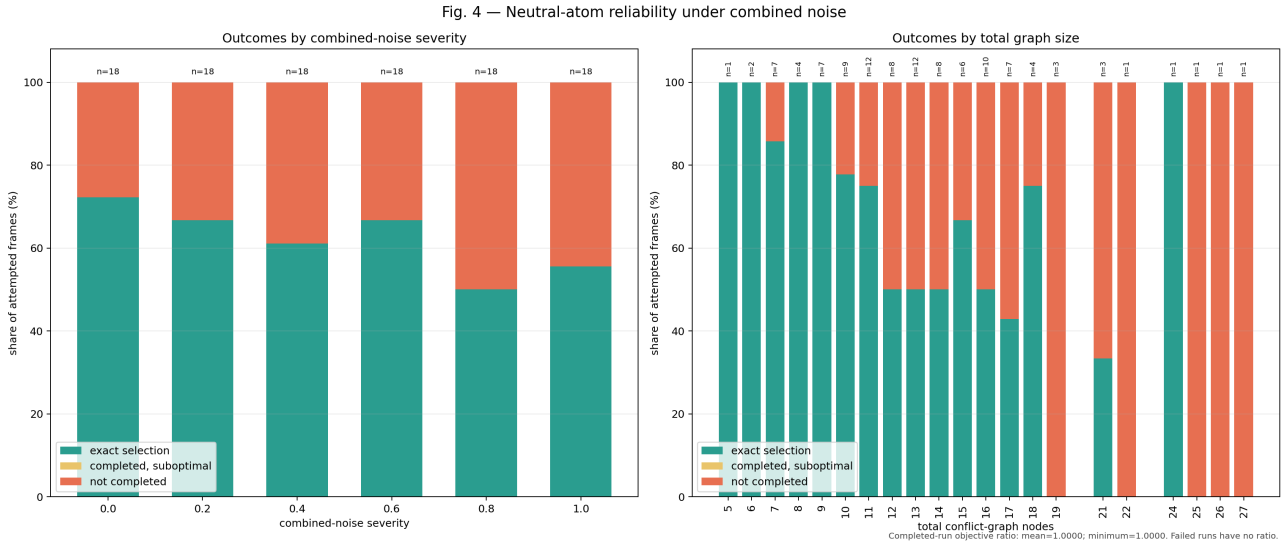


Figure 20: Simulation outcomes using the neutral atom adiabatic approach on synthetic data of varying noise level (left) and size (right).

Lastly, we explored the wall-clock runtime of both the quantum version and the alternative classical simulated annealing version of the MWISP solver. The quantum runtime was estimated from simulation runtime data translated into hardware runtime using Pasqal hardware parameters for pulse and measurement times. Even with limited classical hardware the classical device vastly outperform the quantum process. To explore scaling with graph size we adjusted the classical power defined as a fraction of the benchmark CPU performance, to match the runtime for a graph of 5 vertices. No scaling advantage is discernible for the quantum adiabatic algorithm.

6 Conclusion

The MHT problem via the MWISP formulation at first glance is well suited to a neutral atom quantum adiabatic algorithm, however there are significant complications. Primarily the geometry problem of embedding the MWISP graph into the neutral atom architecture without introducing unintended edges requires further exploration. Other geometric constraints such as residual cross-talk between non-adjacent atoms and lingering distant interactions also limit the accuracy. Additionally, every frame of data requires both significant classical pre-processing as well as additional time to optimize the graph embedding in the neutral atom architecture and to perform repeated shots of the quantum adiabatic algorithm, which then have to interface with the classical device for the next frame. This feedback loop is impractical on real devices when the time steps are small. Lastly, utilizing gating, pruning, and clustering allows approximate classical methods such as simulated annealing to work successfully for the scales we have explored here. It is possible that such classical approaches begin to perform less successfully at larger scales that we have not been able to simulate as of yet.

H

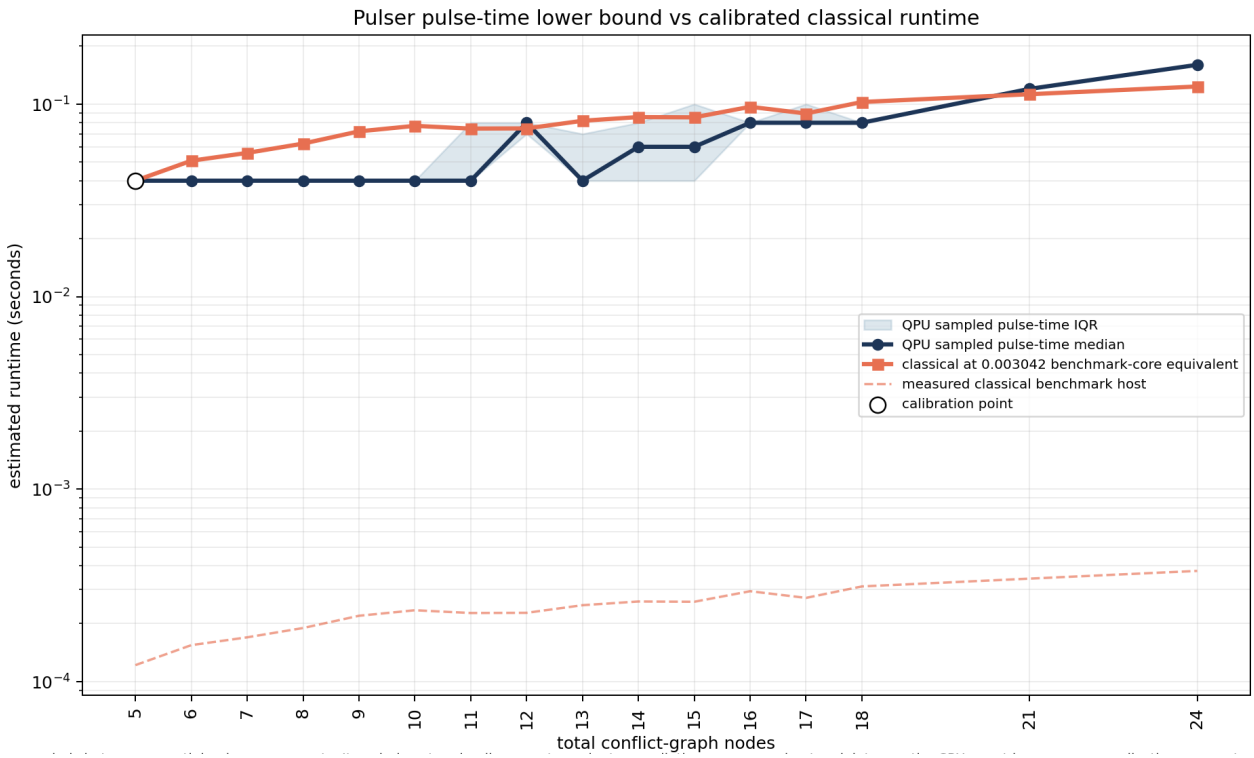


Figure 21: Runtime of both the classical MWISP solver and the estimated quantum runtime based on simulations and real hardware values. A CPU with performance scaled to match the quantum runtime for a graph of 5 vertices is plotted to explore scaling times.

H

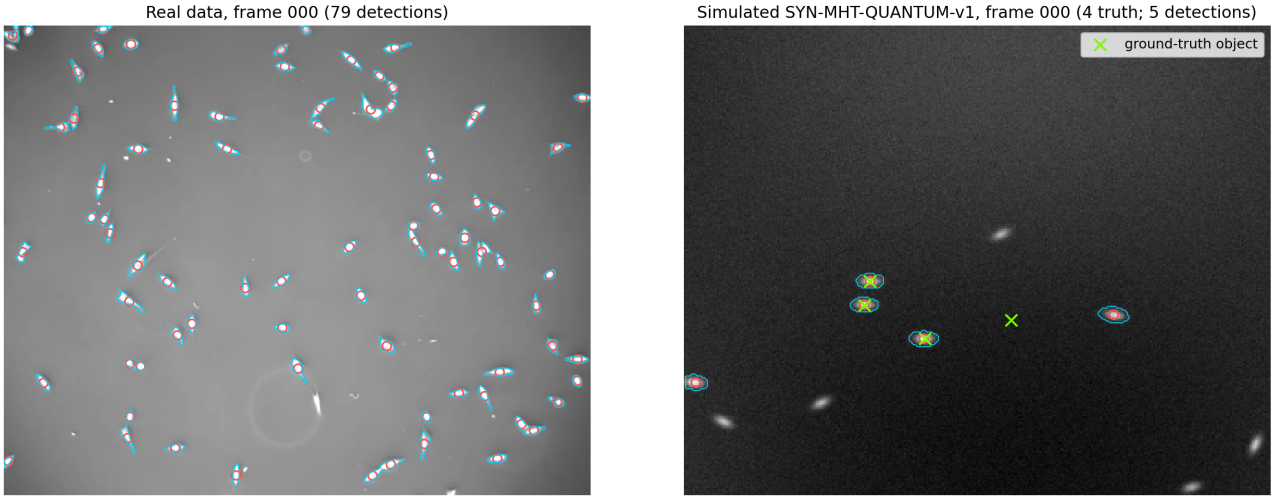


Figure 22: An example of the image detection step overlayed with both cell image data (left) and synthetic data (right).

Code availability. The code and experiment notebooks used in this work are available at `https://github.com/SAN-G-8055/Neutral-Atom-QC-for-MHT`.

AI Acknowledgments AI tools were used to support parts of the project, including the organization and development of the simulation code, generation of synthetic data, initial drafting of text describing the simulation results, and preparation of an initial poster layout. All AI-assisted material was reviewed and refined by the authors.

References

[1] Elie Bermot et al. *Local Rydberg blockade regimes for disk graph embedding and quantum optimization*. en. arXiv:2506.13228 [quant-ph]. Dec. 2025. DOI: 10.48550/arXiv.2506.13228. URL: http://arxiv.org/abs/2506.13228 (visited on 08/18/2026).

[2] Edward Farhi et al. *Quantum Computation by Adiabatic Evolution*. en. arXiv:quant-ph/0001106. Jan. 2000. DOI: 10.48550/arXiv.quant-ph/0001106. URL: http://arxiv.org/abs/quant-ph/0001106 (visited on 08/19/2026).

[3] Dimitri J. Papageorgiou and Michael R. Salpukas. "The Maximum Weight Independent Set Problem for Data Association in Multiple Hypothesis Tracking". en. In: *Optimization and Cooperative Control Strategies*. Ed. by Manfred Morari et al. Vol. 381. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 235–255. ISBN: 9783540880622 9783540880639. DOI: 10.1007/978-3-540-88063-9_15. URL: http://link.springer.com/10.1007/978-3-540-88063-9_15 (visited on 08/18/2026).

[4] *QAA to solve a MWIS problem*. URL: https://docs.pasqal.com/pulser/tutorials/mwis/ (visited on 08/18/2026).

[5] Marcel Sheeny et al. "RADIATE: A Radar Dataset for Automotive Perception in Bad Weather". In: Xi'an, China: IEEE, May 2021, pp. 1–7. ISBN: 9781728190778. DOI: 10.1109/ICRA48506.2021.9562089. URL: https://ieeexplore.ieee.org/document/9562089/ (visited on 08/19/2026).

[6] Henrique Silvério et al. "Pulser: An open-source package for the design of pulse sequences in programmable neutral-atom arrays". en. In: *Quantum* 6 (Jan. 2022). arXiv:2104.15044 [quant-ph], p. 629. ISSN: 2521-327X. DOI: 10.22331/q-2022-01-24-629. URL: http://arxiv.org/abs/2104.15044 (visited on 08/18/2026).