

Quantum Optimised Beam Allocation with a Radar Network for Multiple Target Tracking

1st Benjamin Rodrigues de Miranda[✉]
*German Aerospace Center,
Friedrich-Alexander University
of Erlangen-Nuremberg
Munich, Germany
benjamin.rodriguesdemiranda@dlr.de*

2nd Sigurd Huber[✉]
*German Aerospace Center
Munich, Germany
sigurd.huber@dlr.de*

3rd Tobias Brosch
*Hensoldt Sensors GmbH
Ulm, Germany
tobias.brosch@hensoldt.net*

4th Johannes Pardey[✉]
*Hensoldt Sensors GmbH
Ulm, Germany
johannes.pardey@hensoldt.net*

5th Johannes Hauff
*Hensoldt Sensors GmbH
Ulm, Germany
johannes.hauff@hensoldt.net*

6th Gerhard Krieger[✉]
*German Aerospace Center,
Friedrich-Alexander University
of Erlangen-Nuremberg
Munich, Germany
gerhard.krieger@dlr.de*

Abstract—Optimally allocating a limited number of radar beams to actively track multiple targets is a computationally intensive task. The quantum approximate optimisation algorithm, a method for efficiently finding near optimal solutions to combinatorial optimisation problems, was assessed in its ability to find good beam allocations. A simulation, including sensors, targets, and an extended Kalman filter, was created specifically for this work and was used in tandem with a simulated quantum environment to this end. After reformulating the problem, two small scale test cases were implemented. In both cases, the returned beam allocations were optimal in that they resulted in as many targets as possible being tracked to an acceptable level based on predetermined thresholds. In addition, a complexity analysis was performed which suggests that a quantum advantage may be possible for this problem if the quantum circuit parameters can be efficiently calculated. Based on these results, the quantum approximate optimisation algorithm is promising for the beam allocation problem.

Index Terms—sensor fusion, state estimation, resource management, beam allocation, target tracking, radar, network, non-linear assignment, quantum computing, quantum approximate optimisation.

I. INTRODUCTION

Understanding the motion of enemy vehicles, missiles, and other malicious objects in a strategically significant airspace is crucial in defence scenarios. Radar technology enables the tracking of objects and is widely applied to safeguard areas from threats.

This work deals with allocating radar beams to targets for tracking, which becomes especially important when there are not enough sensors to track all targets adequately and deciding which targets will go unobserved must be done quickly. In large radar networks, calculating an optimal allocation with a classical computer is computationally intensive. As such, it makes sense to investigate whether alternative approaches can alleviate some of this computational load. The quantum approximate optimisation algorithm (QAOA) is suspected to

be able to approximately solve combinatorial optimisation problems faster than classical approximation methods with a comparable solution quality [1]. Here, the QAOA was applied to the beam allocation problem as a proof of concept. A single time step was considered, with optimisation over a time window being the goal of future work. From hereon out, a distinction has been made between a beam allocation, which refers to where beams are to be pointed during the next time step only, and a beam schedule, which refers to where these beams are to be pointed for each time step in a future time window.

The contributions of this work can be summarised as follows:

- reformulation of the beam allocation problem as an energy minimisation problem
- complexity analysis of the quantum circuit implementing the optimisation
- preliminary experimental results in a simulated environment

The remainder of this work is structured as follows. Section II provides an overview of typical approaches to beam allocation. Then Section III details the simulated scenario and states assumptions before Section IV formulates the problem mathematically. Results are provided in Section V before the quantum circuit complexity is analysed in Section VI. Finally, concluding remarks are made, limitations are discussed, and promising avenues for future research are suggested in Section VII.

II. BACKGROUND

There are many processes involved in target tracking. The focus here is on estimating the states of targets with an extended Kalman filter (EKF). This typical approach involves iteratively recomputing each target's state estimate $\hat{\mathbf{x}}$ and associated covariance $\mathbf{P}$ by first making predictions based on

internal motion models and then combining them with range and bearing measurements made by the sensors [2].

The EKF equations can also be used to calculate expected state and covariance estimates. That is, the state and covariance estimates in the future. This is possible because the EKF can be propagated in the absence of measurements by setting the predicted state as the state estimate and the predicted covariance as the covariance estimate. This means that only the motion models are used to make state estimates while the elements of $\mathbf{P}$ increase each iteration due to the accumulation of noise, which would normally be counterbalanced by measurements.

The diagonal elements of the matrix $\mathbf{P}$ give the variance of each component of the state estimate. These values indicate how uncertain that estimate is. The traditional approach to beam scheduling uses this concept to predict when the variance will exceed a threshold, at which point the corresponding target should be observed to reduce the uncertainty in its state estimate [3].

A more recent approach also uses variance as a measure of state estimate uncertainty. Specifically, $\sqrt{\text{tr}(\mathbf{P}(k))}$ measures the uncertainty of a target at some time step k . This metric is combined with the idea of expected measurement updates, in which the radar equation is used to estimate the impact that scheduled observations will have on $\mathbf{P}$. The expected covariance across a time window covered by a given beam schedule can be calculated by using the predictive step of the EKF to update the covariance matrix when no observation is scheduled, which increases $\mathbf{P}$, and using an expected measurement update when an observation is scheduled, which decreases $\mathbf{P}$ [4]. When observations of a target are planned, this process lowers the uncertainty which reflects how the system would behave should those observations be made. When no observation is planned, the filter is propagated without an expected measurement update and the uncertainty increases, which is what would happen if the system was to ignore that target. This allows the expected covariance to be calculated across a future time window in which multiple observations of different targets are planned. Every possible combination of dwell length and revisit time, which together define a beam schedule, are evaluated for each sensor in the network from a set of possibilities and the one which minimises the accumulated uncertainty is selected [4].

The covariance associated with a target's state estimate decreases as more sensors observe it simultaneously through a process known as sequential sensor fusion. When multiple sensors are instructed to observe the same target, they make independent measurements which are combined into a single covariance estimate. The first sensor in the set estimates the covariance according to the EKF equations. This estimate is passed to the next sensor which uses it as its predicted covariance instead of calculating it from its motion model. Then this sensor, too, combines its own measurement with its predicted covariance, which is the previous sensor's covariance estimate, to get its own estimated covariance. The process continues like this until all sensors have fused their measurements [5]. For




Fig. 1. Sensors and targets in the area of interest. Fields of view are given by dashed lines and directions of motion are indicated by arrows.

TABLE I
SENSOR PARAMETERS

Parameter	Value
first sensor position	(-10km, -30km)
second sensor position	(10km, -40km)
boresight	north
field of view	60°
range resolution	70m
half power beam width	1.5°
unity signal-to-noise ratio range	200km
unity signal-to-noise ratio radar cross section	1m ²
unity signal-to-noise ratio dwell length	30ms
k_a	$\sqrt{2}$
k_b	1.81

an ordered set of sensors A which are instructed to observe a target n , the covariance matrix resulting from sequential fusion is calculated using (1), where $a_0, \dots, a_{|A|} \in A$ are sorted by distance of the corresponding sensor to n . The target in question has been added as a subscript for clarity.

$$\mathbf{P}_n(k; A) = \mathbf{P}_n(k; a_0) (\mathbf{I} - \mathbf{K}_n(k; a_1) \mathbf{H}_n(k; a_1)) \dots (\mathbf{I} - \mathbf{K}_n(k; a_{|A|}) \mathbf{H}_n(k; a_{|A|})) \quad (1)$$

III. SIMULATION DETAILS

A simulation, containing the sensor network, target dynamics, and EKF, was created specifically for this project. The experiments in this work made use of two north facing analog monostatic phased array radar devices. Their positions, as well as their capabilities, which were identical, are outlined in Table I. In the experiments, three targets manoeuvred in two dimensions, sharing a plane with the sensors. This setup is shown in Figure 1.

A discrete Wiener process acceleration model [8] was used as the motion model for all targets. In this model, targets deviate from their existing trajectories due to white Gaussian noise in their acceleration. Specific details are provided in Table II.

TABLE II
TARGET PARAMETERS

Parameter	Value
number of targets	10
first target initial state	$[-20\text{km}, 300\text{ms}^{-1}, 0\text{ms}^{-2}, 20\text{km}, -300\text{ms}^{-1}, 0\text{ms}^{-2}]$
second target initial state	$[-15\text{km}, 300\text{ms}^{-1}, 0\text{ms}^{-2}, -8\text{km}, 100\text{ms}^{-1}, 0\text{ms}^{-2}]$
third target initial state	$[20\text{km}, -200\text{ms}^{-1}, 0\text{ms}^{-2}, 20\text{km}, -200\text{ms}^{-1}, 0\text{ms}^{-2}]$
initial covariance	$\text{diag}[60^2, 10^2, 5^2, 60^2, 10^2, 5^2]$
radar cross section	$20m$
process noise standard deviation	$5ms^{-2}$
time step length	$1s$




Fig. 2. Possible allocations of sensors to targets. Each solid edge indicates a possible observation and pairs of similarly coloured dashed edges are examples of valid allocations.

IV. PROBLEM FORMULATION

Recall that the goal is to select which radar looks where. A nonlinear assignment problem, which is $\mathcal{NP}$ complete [9], can be formulated to this end. Given a bipartite graph with vertex sets corresponding to sensors $m \in M$ and targets $n \in N$, each sensor can be assigned to a target. An instance of this problem is shown in Figure 2 for two sensors and three targets.

Candidate beam allocations are represented by a strings of binary variables, where each variable corresponds to an edge in the nonlinear assignment problem, of which there are MN . Whether or not the observation corresponding to each edge is made depends on whether that edge is labelled 1 or 0 respectively. The goal, then, is to label the edges such that those labelled 1 represent an optimal beam allocation. The form of the binary strings is given by (2), in which each variable has indices denoting which sensor and which target form the endpoints of the corresponding edge. This form is fitting because the QAOA optimises objective functions which accept binary strings as candidate solutions.

$$\mathbf{z} = [z_{0,0}, z_{1,0}, \dots, z_{M-1,N-1}] \quad (2)$$

The objective function needs to accept such a candidate beam allocation and return its cost. Then the QAOA can be used to find the string which minimises this value. There are

two components in this function, a penalty for high uncertainty and a penalty for infeasible allocations. First consider (3), the expected uncertainty of a single target n arising from its observation by a set of sensors A . Because only one time step is considered, k is not an argument of this function.

$$u_n(A) = \sqrt{\text{tr} \mathbf{P}_n(k; A)} \quad (3)$$

The expected covariance estimate $\mathbf{P}(k)$ is the result of sequential fusion described by (1), which multiplies information from each sensor in A . In other words, the expected uncertainty of a single target depends on the multiplication of information from incoming edges. If this information was instead summed, the problem would be one of linear assignment, for which efficient classical algorithms exist. Note that because the diagonal elements of $\mathbf{P}$ have different units, namely m , ms^{-1} , and ms^{-2} , the expected uncertainty is dimensionless.

Instead of directly using (3) as the cost for an individual target n , it is compared to a threshold δ_n which reflects the acceptable level of uncertainty associated with the target. Such a threshold is predetermined for each target individually. Ideally, a beam allocation results in the expected uncertainty of each target being lower than the corresponding threshold. Checking this can be done using (4), which evaluates to 1 if $u_n(A) \leq \delta_n$ and to 0 otherwise.

$$f_n(A; \delta_n) = \frac{\sqrt{(\delta_n - u_n(A))^2 + \delta_n - u_n(A)}}{2\sqrt{(\delta_n - u_n(A))^2}} \quad (4)$$

For the sake of sanity, the observation of a target by a set of sensors is referred to as a "beneficial" observation from hereon out if it results in that targets expected uncertainty being lower than its threshold. In other words, a combination of A and n that makes (4) evaluate to 1 is "beneficial". To check if a given beam allocation $\mathbf{z}$ contains a given beneficial observation, (4) is multiplied by a product of binary variables corresponding to the edges between sensors in A and the target n ; mathematically this looks like $f_n(A; \delta_n) \prod_{m \in A} z_{m,n}$, and works because even if the observation of n by sensors in A is beneficial, if any of the binary variables in $\mathbf{z}$ which correspond to that observation are 0, the whole expression evaluates to 0. Otherwise, it evaluates to 1 which indicates both that the beam allocation $\mathbf{z}$ specifies that sensors in A are to observe n and that doing so results in an expected uncertainty below the maximum amount specified by the threshold for that target.

It is now possible to write an expression demanding that at least one beneficial observation of a target n is contained in the candidate beam allocation. Such an expression should only evaluate to 1 if any set of sensors A causes $f_n(A; \delta_n) \prod_{m \in A} z_{m,n}$ to evaluate to 1. This is a series of logical or statements across different sensors combinations, which, by De Morgan's laws, can be implemented as in (5).

$$g_n(\mathbf{z}; \delta_n) = 1 - \prod_{A \subseteq \{i\}_{i=0}^{M-1}} \left[1 - f_n(A; \delta_n) \prod_{m \in A} z_{m,n} \right] \quad (5)$$

Good beam allocations are defined as those with many beneficial observations. Because the QAOA will preform minimisation, the objective function should assign low costs to good beam allocations and high costs to bad beam allocations. Currently, (5) does the opposite. For a given target, it adds 1 if the candidate beam allocation results in at least one beneficial allocation and does nothing otherwise. To instead penalise the lack of any beneficial observations of a target, $1 - g_n(\mathbf{z}; \delta_n)$ is used. The total cost of an allocation is the sum of the individual target costs,

$$\sum_{n=0}^{N-1} [1 - g_n(\mathbf{z}; \delta_n)],$$

however, now observing all targets with all sensors is optimal. This can be avoided with a penalty λ for each pair of observations by the same sensor, which is done by (6).

$$h(\mathbf{z}; \lambda) = \sum_{m=0}^{M-1} \sum_{i=0}^{N-1} \sum_{\substack{j=0, \\ j \neq i}}^{N-1} \lambda z_{m,i} z_{m,j} \quad (6)$$

The two innermost nested summations add λ whenever two edges share one endpoint at the same sensor and have their other endpoints at different targets. To do this for each sensor in turn, the outer summation is needed. This results in the final cost function shown in (7), which can be expanded to the form shown in (8).

$$C(\mathbf{z}) = \sum_{n=0}^{N-1} [1 - g_n(\mathbf{z}; \delta_n)] + h(\mathbf{z}; \lambda) \quad (7)$$

$$C(\mathbf{z}) = \sum_{n=0}^{N-1} \left[\prod_{A \subseteq \{i\}_{i=0}^{M-1}} \left[1 - f_n(A; \delta_n) \prod_{m \in A} z_{m,n} \right] \right] + \sum_{m=0}^{M-1} \sum_{i=0}^{N-1} \sum_{\substack{j=0, \\ j \neq i}}^{N-1} \lambda z_{m,i} z_{m,j} \quad (8)$$

From here, the objective function needs to be mapped to a tensor called the cost Hamiltonian. This tensor describes the efficient evolution implemented by a quantum circuit such that the objective function is minimised. This mapping is performed using the Hadamard transform. A description of the QAOA and this mapping are both beyond the scope of this work, but the reader is referred to the references [1] [10] for a background on these.

V. EXPERIMENTAL RESULTS

Using the setup shown in Figure (1), two experiments were run. Recall that the QAOA constructed here tries to allocate beams such that the uncertainty associated with each target




Fig. 3. Frequencies of returned solutions.

falls below a threshold, while restricting each sensor to observe no more than one target. In the first experiment, the thresholds were set such that each target's threshold could be satisfied if it was observed by any combination of sensors. In the second experiment, it was possible for the uncertainty of one of the targets to drop below its threshold only if it was observed by both sensors, while the thresholds for the other two targets were set to be impossible to satisfy. The first and second experiments are respectively referred to as "observe all" and "fusion" from hereon out.

In Figure 3, each output state on the x -axis has an associated measurement frequency on the y -axis from running the learned quantum circuit multiple times. The frequencies for the observe all and fusion cases are shown in blue and green respectively. The states with the highest output probability were $|33\rangle$ for the observe all experiment and $|48\rangle$ for the fusion experiment. Decomposing these into binary vectors of length six leaves $\mathbf{z} = [1, 0, 0, 0, 0, 1]$ and $\mathbf{z} = [1, 1, 0, 0, 0, 0]$ respectively. Figure 2 marks these beam allocations in blue for the observe all experiment and in green for the fusion experiment. The next two most frequent outcomes for the observe all experiment were also optimal. For the fusion experiment, there were not other optimal solutions.

In the observe all case, there are thresholds demanding that all three targets are observed even though there are only two sensors. Therefore, optimal solutions are those in which two separate targets are observed. These solutions have an objective value of -2 . In the fusion case, two of the thresholds were unsatisfiable because they were set to be too low. The only way the uncertainty threshold for the remaining target could be met was by observing it with both sensors. This solution has an objective value of -1 . In both cases, the single solution with the highest frequency was optimal.

VI. COMPLEXITY ANALYSIS

In this work, the complexity of the quantum circuit was the length of the longest connected path through the quantum circuit as measured by the number of gates. The gates utilised here are the Hadamard gate H , the rotation gates R_X and R_Z , the controlled not gate $CNOT$, and measurements. The required number of qubits is MN , as each qubit is a binary variable from the objective function (8), which in turn has a binary variable per edge in the assignment problem graph.

The objective function (8) contains two components which are translated into quantum gates, the comparison of uncertainties to thresholds and the penalties. Both contribute to the quantum circuit depth. Starting with the former, the quantum circuit is split into registers, one for each target. The number of sequential gates required for a single target is $4(2^M)$. Because there are no operations acting across the registers, this can be done for all targets in parallel without increasing the depth.

Next, consider the depth incurred by the penalties. This requires $4(N - 1)$ gates per sensor. Similar to how there were independent registers for each target before, the operations implementing the constraints act on registers which are independent in terms of sensors. This means that the penalties for each sensor can be implemented in parallel. For p layers and some minor additional overhead, an upper bound on the final quantum circuit depth is

$$p(4(2^M) + 4(N - 1)) + 2 = 4p(2^M + N - 1) + 2,$$

which, upon first inspection, looks to be concerning due to the presence of an exponentiated M . Recall, however that M is the number of sensors in the network. In the real world, this number is fixed at a reasonable value and does not grow. What is of interest here is the asymptotic scaling in the number of targets, which not only varies in the real world but can be large enough to render classical computation infeasible. From the quantum circuit depth, it can be seen that the complexity is $\mathcal{O}(N)$.

Finding an exact solution classically means evaluating each solution independently and comparing these. This is $\mathcal{O}(N^M)$ where M is treated as constant. This means that, using QAOA, the complexity has been reduced from polynomial to linear in the total number of targets without requiring an exponential amount of space. It should be noted that, because QAOA is an approximation algorithm, a more fair comparison would be to a classical approximation algorithm, which is the goal of future work.

It is important to note that the improvement discussed above is a *potential* speedup because it hinges on the efficient computation of the quantum circuit parameters γ and β . This is by no means a small detail, and involves taking advantage of the locality of the problem as well as graph isomorphisms. Determining whether this is possible is another goal of upcoming work and may require modifications to how the constraints are implemented.

VII. CONCLUSION, LIMITATIONS, AND FUTURE WORK

Allocating the beams of a radar network to track targets in an airspace was framed as a nonlinear assignment problem and optimised with the QAOA. This involved working the EKF equations and the equations for expected measurements into a pseudo boolean objective function, and then simulating the radar network, target motion, and EKF. In a simulated quantum environment, two small scale test cases resulted in optimal allocations in terms of meeting predefined uncertainty thresholds. With this said, the quantum circuit actually had

a higher chance of returning a near optimal solution for any given measurement due to the presence of multiple suboptimal solutions having high measurement probabilities in both experiments. Despite this, taking the single most frequent solution over a number of shots resulted in optimal beam allocations. Whether or not a quantum speedup can be achieved depends on whether or not the quantum circuit parameters can be efficiently computed. In this work, these parameters were learned variationally, so while the quantum circuit complexity hints at a possible speedup, this was not achieved here and requires further attention. Should this be possible, the complexity reduces from polynomial to linear in the number of targets. This work is therefore an important first step in allocating beams for target tracking faster than is currently possible, something which is crucial in defence scenarios.

While the results in this work show promise, there are limitations. These include the simplistic setup and the small scale of the experiments. The intention is to address these in future work. Focus will also be placed on comparing of the quantum and classical approximation algorithms to determine the true extent of any quantum speedup and efficiently computing the quantum circuit parameters, which would allow the complexity improvements described in this work to be taken advantage of.

REFERENCES

- [1] E. Farhi and G. Jeffery and S. Gutmann, "A Quantum Approximate Optimisation Algorithm," 2014, arXiv:1411.4028.
- [2] Y. Bar-Shalom and X. R. Li and T. Kirubarajan, "The Extended Kalman Filter," in "Estimation with Applications to Tracking and Navigation," New York City, New York, USA: John Wiley & Sons, 2001, ch. 10.3, pp. 381-395.
- [3] G. Van Keuk and S. S. Blackman, "On Phased-Array Radar Tracking and Parameter Control," IEEE Trans. Aerosp. Electron. Syst., vol. 29, no. 1, pp. 186-194, Jan. 1993, 10.1109/7.249124.
- [4] R. Nadjasngar and A. Charlish, "A Performance Model for Target Tracking with a Radar Network," IEEE Radar Conference, pp. 238-243, Oct. 2015, 10.1109/RadarConf.2015.7411887.
- [5] Y. Bar-Shalom and X. Li, "Sequential Updating," in "Multitarget-Multisensor Tracking: Principals and Techniques," Storrs, Connecticut, USA: Y. Bar-Shalom, 1995, ch. 2.2.2, pp. 88.
- [6] P. B. Weichman, "Quantum-enhanced algorithms for classical target detection in complex environments," Phys. Rev. A, vol. 103, no. 4, Apr. 2021, 10.1103/PhysRevA.103.042424.
- [7] J. N. Zaeck and A. Liniger and M. Danelljan and D. Dai and L. V. Gool "Adiabatic Quantum Computing for Multi Object Tracking," IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8811-8822, Sep. 2022, 10.1109/CVPR52688.2022.00861.
- [8] Y. Bar-Shalom and X. Li and T. Kirubarajan, "Discrete Wiener Process Acceleration Model," in "Estimation with Applications to Tracking and Navigation," New York, USA: John Wiley & Sons, 2001, ch. 6.3.3, pp. 274.
- [9] R. A. Murphey, "Target-Based Weapon Target Assignment Problems," in "Nonlinear Assignment Problems: Algorithms and Applications," P. M. Pardalos and L. S. Pitsoulis, Ed., Combinatorial Optimization Series, New York, USA: Springer, 2010, ch. 3, pp. 39-41.
- [10] S. Hadfield, "On the Representation of Boolean and Real Functions as Hamiltonians for Quantum Computing," ACM T. Quantum Comput., vol. 2, no. 4, pp. 1-21, Dec. 2021, 10.1145/3478519.
- [11] D. Zuckerman, "On Unapproximable Versions of NP-Complete Problems," SIAM J. COMPUT., vol. 25, no. 6, p.1293-1304, Dec. 1996, 10.1137/S009753979426640.