Mobile Radiation Source Interception by Aerial Robot Swarms

Indrajeet Yadav and Herbert G. Tanner

Abstract— **This paper presents a motion planning approach to track a moving target for the purpose of classifying its radioactivity properties, using a team of micro aerial vehicles (MAV). The team falls into an optimal, from a sensing perspective, formation around the target and collectively decide, within seconds, if the target is even weakly radioactive. The MAV formation is optimal in the sense that the vehicles minimize the detection error in their collective decision-making. Intuition from the analytical solution of the basic problem formulation, coupled with numerical simulations, guides the technical approach in which an intractable nonlinear optimization problem is converted into a quadratic program (QP). The QP solution then informs a motion planner based on navigation functions, to command the MAV into the desired formation. The complete motion planning and control architecture is tested in simulation.**

I. Introduction

Consider a vehicle moving in a cluttered, GPS-denied environment. This vehicle may be carrying sources of radioactivity, the emissions of which are weak enough to blend into background. With the right algorithmic and signal processing infrastructure, however, a small swarm of drones equipped with commercial off the shelf (COTS) lightweight gamma-ray or thermal neutron counters can unambiguously and confidently distinguish whether the vehicle is even slightly "hotter" than background in a matter of seconds.

Such a decision-making problem (hot vs. cold) can be formulated as a *binary hypothesis testing* problem: the mobile sensor in the vicinity of the target records counts (i.e., events where radioactivity particles are captured by the detector's surface), and then conducts a likelihood ratio test (LRT) to decide whether the target is indeed radioactive. The LRT will indicate whether the data collected up to that time instance is sufficient to determine radioactivity above a certain threshold [14]. The key to making this decision-making problem tractable and directly recasting it analytically as an optimal control problem is to replace the actual probability of missed detection (PMD) and probability of false alarm (PFA) associated with the detection problem, with analytically derived Chernoff bounds, and use the latter as proxies for the intractable true probabilities [14].

The inverse square law dependency of incident-to-detector radiation relative to the distance between sensor and source [12] dictates MAV as the platform of choice: an aerial platform can quickly close the distance with the target to spend more time in close proximity, resulting into more information-rich data and therefore more accurate

decision-making. If, in addition, multiple aerial detectors can be simultaneously deployed, and share the information they collect, then the accuracy of decision-making can be significantly boosted, especially in the low-radioactivity regime [19]. This paper takes a step in this direction, focusing on the final phases of the target interception scenario, where the mobile sensors have to coalesce around the moving target. The paper develops a methodology to generate an *optimal* formation around the target that would result into the best possible decision-making accuracy while satisfying a number of safety constraints.

Selected recent literature on aerial formations navigating in cluttered environments features multi-MAV formation trajectory planning and control [17], vision-based control [11], and distributed formation control with obstacle avoidance [1]. Along similar lines, one also sees *distributed* planning and control approaches where multiple MAV navigate a cluttered environment to similarly arrange themselves in a specified formation [2]. The MAV start from an initial configuration, collectively identify the largest obstacle-free convex region, and move in a way that the convex-hull of MAV positions always remain inside the safe region. The MAV desired configurations are obtained as a solution to a non-linear optimization problem that minimizes the error with the specified formation configuration. In all aforementioned cases, each MAV plans its motion in order to converge to a specified position in a given *static* formation. On the other hand, most of the target tracking literature treats a single MAV chasing the specified target [4], [16]. There is a dearth of work where formations of vehicles are coordinated to converge and track a moving target (cf. [18]) and there is even less reported work where tracking formations are optimized for a sensing task. Aerial radiation measurement has been considered in the context of mapping [5]; however, tracking radioactive targets in cluttered environments has not.

The methodology presented in this paper relies on no pre-specified or fixed desired formation. Based on knowledge of the moving target's position, it generates the desired formation configuration online and dynamically, maximizing an analytical metric of radiation detection decision-making accuracy. The solution approach first simplifies and reduces the initial complex nonlinear optimization problem associated with detection-optimal formation design into a convex problem, the solution of which may be sacrificing some optimality for very substantial gains in computational efficiency. A motion planner then generates a set of waypoints for the MAV to implement the solution, and assigns each MAV to a waypoint using a version of the Hungarian algorithm [8]. Then with knowledge of the workspace topology,

and assuming that the target does not come too close to the free workspace boundaries, a specialized dynamic navigation function on each MAV [10], combined with a low-level MAV thrust and moment controller [9], guides the mobile sensors to their desired formation while enabling them to avoid static obstacles in their workspace.

The key contributions of this paper are: a) a computationally efficient solution to the problem of determining the optimal for detection sensor placement around a moving radiation source, and b) analytic integration of differential-geometric quadrotor controllers [9] with 3D time-varying navigation function [10] for target tracking and interception in cluttered environments.

II. RADIATION DETECTION PRELIMINARIES

A COTS gamma-ray (neutron) detector essentially records a count every time a gamma-ray (neutron) hits its effective area. The main problem associated with typical counter-based radiation detection is the absence of information to distinguish upon reception a count due to the presence of a radiation source from a naturally occurring one due to background. Instead, what is done is that *over time*, statistics can indicate whether the *total* count can be confidently attributed to just background, or background plus an additional source. This question is mathematically codified as a binary hypothesis test: choosing between two conflicting hypothesis H_1 and H_0, corresponding to the source being present or not, respectively. A probability distribution is assigned to each hypothesis and the data is collected for a fixed time-interval T.[1] At the end of T the data are combined to a statistic, which is analyzed to determine, within some error margins, which of the two distributions is more likely to agree with the observations. For such a test, two types of errors can occur. A *false alarm* occurs if the source is erroneously considered present while a *missed detection* occurs when the outcome is the absence of the source while it is present. Probabilities for making any of these errors serve as metrics to assess the performance (accuracy) of the hypothesis test.

In a mobile detector and source scenario, the binary hypothesis testing can be formulated as a time-inhomogeneous Neyman-Pearson test [13]. Let vectors $x(t),\ x_t(t) \in \mathbb{R}^3$ denote the location of detector and source, respectively. Let also χ denote detector's radiation-sensitive cross-section, a the intensity of the source. Then the rate of counts (perceived intensity) observed at the detector varies inversely with its distance to the source (cf. [12])

$$\nu(t) = \frac{\chi a}{2\chi + \|x(t) - x_t(t)\|^2} \qquad (1)$$

The perceived (incident to detector) source intensity being a function of time, as the relative distance between detector and source changes, the detector's counting statistics are now expressed by a time-inhomogeneous Poisson process. The precise computation of the true error probabilities for such observation processes is mathematically intractable; however, it is possible to obtain reasonable Chernoff bounds on these error probabilities, and use *them* as surrogates for the unknonw true probabilities [13], [14]. Expressed explicitly in terms of relative distance between the detector and the source, these bounds have been incorporated in an optimal control formulation that prescribes a motion planning strategy that in a Neyman-Pearson sense optimizes the bound on the PMD while keeping the bound on PFA within an acceptable limit [15].

For a count rate generated due to background radiation (background intensity) $\beta(t)$; a non-dimensional number $\mu(t)$ indicating the SIGNAL-TO-NOISE-RATIO (SNR) of the detector; a given upper value for the bound on PFA $\alpha < 1$; and a constant $p \in (0, 1)$, the optimal motion control problem for a single detector is mathematically expressed[2] as [15]

$$\min \int_0^T [\mu^p \log \mu - \mu + 1]\beta \, \mathrm{d}s \qquad (2a)$$

$$\text{s.t.} \int_0^T [p\mu^p \log \mu - \mu^p + 1]\beta \, \mathrm{d}s = -\log \alpha \qquad (2b)$$

$$\mu \triangleq 1 + \frac{\nu}{\beta} = 1 + \frac{\chi a}{\beta[2\chi + \|x - x_t\|^2]} \qquad (2c)$$

The time dependency of quantities has been suppressed for presentation clarity. The expression in (2c) essentially expresses the increase in the number of counts from the source compared to background as a function of distance.

Intuitively, the solution to this problem calls for closing the gap between the source and the detector quickly as possible, and then keeping the sensor at the source for as long as the deadline T allows. The analytic determination of the error bounds also facilitates the analysis of the effects of information sharing among multiple detectors on the collective decision-making accuracy: it distinctly illustrates the advantages of utilizing multiple MAV for detection task [19]. The following lemmas describe important properties of the SNR and error probability bounds for a single detector case, which are at the center of the optimization problem.

Lemma 1. *Variable μ monotonically increases as the distance of the detector with the source reduces.*

Proof: Straightforward from (1) and (2c). ☐

Lemma 2. *With T and β fixed, a higher μ results in a lower value for the bound on PMD (2a).*

Proof: For a given bound α on PFA, the value of p decreases with μ [15, Lemma 3], while PMD is a strictly increasing function of p [14, Lemma 19]. Since (2b) is used to calculate p, a higher μ results into a lower p, which in tern result into a lower value of objective function (2a) as it is proportional to PMD. ☐

Thus a detector tracking a moving target at as close of a distance as practically possible, would provide the

---

[1]This distinguishes our approach from an sequential probability ratio test (SPRT), where one waits until the statistic computed builds to a point where a confident decision can be made. Here there is a decision-making deadline at which a decision is forced.

[2]Note the equality in (2b). In general, minimizing the bound on PFA subjected to an inequality ($\leq$) condition in (2b) is a hard problem. A workaround is to use a equality constraint in (2b), to calculate the value of p in the cost functional that proportional to the bound on PMD and solve an unconstrained optimization [15]. (2a) represents unconstrained cost proportional to the exponent of bound on PMD.

best observation data for detection. The *specific question* addressed in this paper, however, is how does this principle generalize to the case of multiple MAV, when there are practical limitations on how close they can get to the target and to each other? What type of shape or *formation* do they have to fall into to achieve optimal decision-making accuracy? The next section seek answers to this question and presents a methodology to realize the desired formation.

III. FORMATION DESIGN

Low-cost COTS detectors generally have a low extrinsic and intrinsic efficiency; i.e., due to their small size, the solid angle they subtend to the source is small, preventing particles from reaching the detector and being registered. The efficiency of converting incident radiation into current through a photodetector also drops: for a miniature COTS thermal neutron counter this can be in the order of $0.04 - 0.1\%$ for high-energy incident neutrons. For these type of sensors, therefore, it can safely be assumed that the approach-the-target phase of their motion does not significantly contribute to the statistic computed for the LRT. On the other hand, once the MAV have fallen into formation around their moving target, we can assume that the distance between them and their target are not changing dramatically any more, suggesting that for each detector, $\dot{\mu} \approx 0$. One can thus expect that the optimal formation would be the one in which MAV arrange themselves to get the best possible value for their μ.

With μ, β, χ, a, x and x_t defined as before, let different MAV be indexed by i and time interval $[0, T]$ be partitioned into K intervals of equal length Δk. A discrete-time version of (2) can be formulated as a nonlinear optimization problem for N detectors as follows. Using standard combinatorial notation $[N] \triangleq \{1, 2, \ldots, N\}$, $[K] \triangleq \{0, 1, 2, \cdots K\}$,

$$\min_{x_i} \sum_{i=0}^{N} \sum_{k=0}^{K} [\mu_{ik}^p \log \mu_{ik} - \mu_{ik} + 1]\beta \Delta k \tag{3a}$$

$$\sum_{i=0}^{N} \sum_{k=0}^{K} [p\mu_{ik}^p \log \mu_{ik} - \mu_{ik}^p + 1]\beta \Delta k = -\log \alpha \tag{3b}$$

$$\mu_{ik} = 1 + \frac{\chi a}{\beta[2\chi + \|x_{ik} - x_{tk}\|^2]} \quad \forall i \in [N], \ \forall k \in [K] \tag{3c}$$

$$\|x_{ik} - x_{jk}\| \geq d_1 \qquad \forall i, j \in [N], \ \forall k \in [K] \tag{3d}$$

$$\|x_{ik} - x_{tk}\| \geq d_2 \qquad \forall i \in [N], \ \forall k \in [K] \tag{3e}$$

The value of parameter p is obtained through (3b). If (3b) is not satisfied, then the detection problem is not well-posed; there is not enough data to make any meaningful decision. Statements (3c), (3d) and (3e) serve to formally define μ, and place inter-MAV collision avoidance, and MAV-target safety distance constraints, respectively. Solution to (3) results in a minimum for the upper bound on PMD, should the detectors spend time T arranged in this configuration.

Figure 1 shows the results of a 2D numerical simulation in which (3) is solved for 3 detectors, constrained to stay at a radius of radius 3 m around their target. The values selected for the parameters, $\chi = 2.77 \times 10^{-6}\,\mathrm{m}^2$ and $\beta = 0.00583$ counts per second (CPS), are experimentally
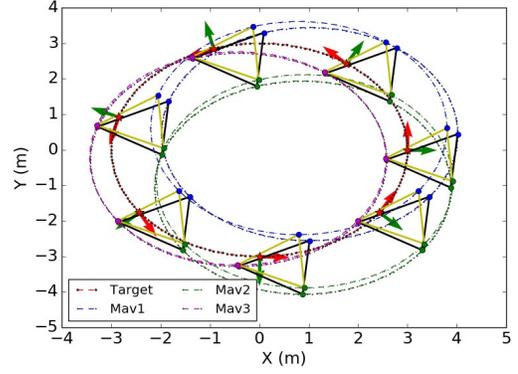
Fig. 1. Results of (3) for T = 50 (black) and 150 (yellow) seconds.

determined for the Domino Neutron detector,[3] while $a = 1.7 \times 10^4$ approximates a $5\,\mu$Ci source[4] available. The upper bound on PFA is set at $\alpha = 10^{-3}$, and distances d_1 and d_2 are set at $1.5\,\mathrm{m}$ and $0.5\,\mathrm{m}$, respectively. The standard solver NLOPT had been used to solve the nonlinear optimization.[5]

Decision deadline T was picked at 50 and 100 seconds in two different solutions that are compared against each other in Figure 1. Interestingly, the optimal solution does not distribute all the detectors symmetrically around the target; rather, it places one very close to the target while keeping the other two trailing while satisfying all constraints. That behavior is observed consistently for different decision deadlines and source intensities. The nonlinear dependence of SNR to the distance between sensor and source dictates that *having at least one sensor as close as possible is preferable to uniformly minimizing all sensor-source distances.*

We hypothesize that one of the reasons this solution arises as optimal is the need to satisfy the constraint (3b) on PFA; without it, the decision problem is infeasible. Thus at least one sensor should obtain the absolute best information possible about the target. The insight on the nature of the optimal solution motivates a relaxation of the formation optimization problem that allows efficient and fast, real-time solutions. We can weigh the different sensors according to their sensitivity, and then recast the problem as a convex weighted minimum-distance problem, where the most sensitive sensor is weighted higher.

The *convex restriction* of the non-convex minimum distance constraints, is now a QP of the form

$$\min_{x_i} \sum_{i=0}^{N} w_i \|x_i - x_t\|^2 \tag{4a}$$

$$\text{s.t. } a_{ij}^T \cdot (x_i - x_j) \geq d_1 \qquad \forall i, \ j \in [N] \tag{4b}$$

$$b_{it}^T \cdot (x_i - x_t) \geq d_2 \qquad \forall i \in [N] \tag{4c}$$

[3]Radiation Detection Technologies Inc http://radectech.com/products/rdt-domino-v5-4

[4]The value of source activity is calculated assuming 4.4×10^9 neutrons per Curie per second emitted by Cf-252 source. (See http://www.logwell.com/tech/nuclear/Californium-252.html).

[5]Steven G. Johnson, The NLopt nonlinear-optimization package, http://ab-initio.mit.edu/nlopt
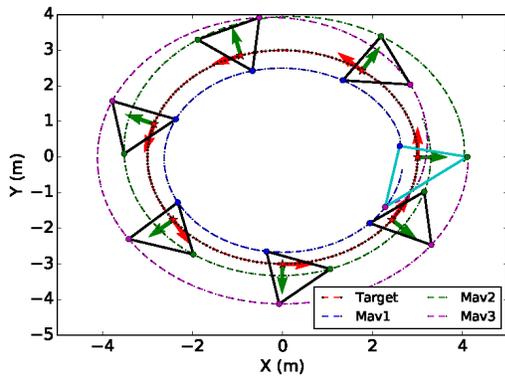
Fig. 2. Results of (4) for similar target motion and d_1, d_2 values.

$$x_{iz} - x_{tz} \geq h_i \tag{4d}$$

Constraints (4b) and (4c) now form the convex relaxation of the hard non-convex minimum distance constraints [3]. The direction vectors $a_{ij} \in \mathbb{R}^3$ and $b_{it} \in \mathbb{R}^3$ for $i, j \in [N]$ are unitary. Problem (4) can now be solved at each time step and provides the direction vectors for the following time step [3] in the form $a_{ij} = \frac{(\hat{x}_i - \hat{x}_j)}{\|(\hat{x}_i - \hat{x}_j)\|}$ and $b_{it} = \frac{(\hat{x}_i - x_t)}{\|(\hat{x}_i - x_t)\|}$. The extra constraint (4d) forces each MAV to hold a vertical distance h above the target at all times during tracking.

Figure 2 shows the solution of (4) obtained using the GUROBI optimizer [7], under the same settings for target motion and parameters d_1, d_2. The results are similar to the original nonlinear optimization, except for the (cyan) triangle formation at the beginning of the motion, when there is no previous solution available to inform on a_{ij} and b_{it}. The QP formulation is more computationally efficient, allows the possible incorporation of additional constraints, and is not sensitive to the choice for the initial point in the nonlinear optimization process.

Let $g_i(t) \in \mathbb{R}^3$, $\forall i \in [N]$ be the set of goal positions for each MAV, obtained by (4). Each MAV is assigned a specific goal through the Hungarian algorithm [8] (with cost defined as square of the distance) that finds a permutation matrix $\xi : \mathcal{I} \to \mathcal{I}$ such that

$$\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} \xi_{ij} \|x_i(t) - g_j(t)\|^2 \tag{5}$$

is minimized with respect to ξ_{ij}.

IV. UAV MOTION PLANNING AND CONTROL

With knowledge of the workspace each MAV uses its own motion planner that employs a time-varying, at least C^2 navigation function [10] $\varphi \to \mathbb{R}_+$, constructed on the MAV's free configuration space $\mathcal{F}$, which can be tuned appropriately to have a unique minimum at the desired formation configuration and be uniformly maximal over the boundary of $\mathcal{F}$. Assuming both $x(t)$ and $g(t)$ remain in the interior of $\mathcal{F}$,[6] the navigation goal becomes the minimum of

$$J(x, g) = \|x(t) - g(t)\|^2 \tag{6}$$

[6]In fact, there are topological constraints [10] that constraint this invariant to a subset of the interior of $\mathcal{F}$.

It has been shown [15] that for a suitably selected obstacle function $\beta(x)$ and a suitably large parameter $\lambda \in \mathbb{R}_+$, there exists a positive number N such that $\forall \kappa \geq N$,

$$\tilde{\varphi}(x, g) = \frac{J(x, g)}{\left[J(x, g)^\kappa + \lambda \beta(x) \right]^{1/\kappa}} \tag{7}$$

is a navigation function when the free configuration space of the robot is a sphere world $\mathcal{S}$. This result has been extended to star worlds [10]. Then a diffeomorphism $h_{\lambda_{sq}} : \mathcal{F} \to \mathcal{S}$ parameterized by a suitably chosen positive parameter $\lambda_{sq} \in \mathbb{R}_+$ can give $\varphi = \tilde{\varphi} \circ h_{\lambda_{sq}}(x, g)$ navigation function properties on $\mathcal{F}$, in the sense that for any position of the target satisfying some reasonable topological conditions, all (unstable) critical points outside the destination manifold are nondegenerate with attraction region of measure zero.

The *desired* velocity for each MAV in the formation is thus determined using φ. Specifically, if $v_{\max}$ denotes the vehicle's maximum speed given the capabilities of its actuators or safety specifications, and k_φ is a positive control gain, then the desired velocity relative to an inertial frame would be

$$\dot{x}_d \triangleq -\operatorname{erf}\left(k_\varphi(\|x - g\|) \right) \cdot \frac{\nabla_x \varphi}{\|\nabla_x \varphi\|} \cdot v_{\max} \tag{8}$$

with ∇_x denoting the gradient with respect to variable x. The desired position x_d and desired acceleration $\ddot{x}_d$ can be obtained by integrating and differentiating the desired velocity, respectively. A suitable value of k_φ in erf regulates the rate at which the MAV slows down to to match the velocity of the target in its vicinity.

A. UAV *Control*

Let m denote the mass of a MAV and $J \in \mathbb{R}^{3 \times 3}$ its moment of inertia about principal axes attached at the center of mass (hereafter referred to as body fixed frame). The relative orientation between the inertial frame and the body fixed frame is encoded in the rotation matrix $R \in SO(3)$. Let $x \in \mathbb{R}^3$ denote the position of the MAV in the inertial frame, and Ω its angular velocity in body fixed frame. The operator $\widehat{\cdot}$ denotes the (wedge) operation that maps a vector in $\mathbb{R}^3$ to a member of the Lie algebra $\mathfrak{so}(3)$; f and M denote the magnitude of the total actuator-generated thrust and moment acting on body fixed frame. The dynamics of the MAV is

$$m\ddot{x} = -mge_3 + fRe_3 \tag{9a}$$

$$\dot{R} = R\hat{\Omega} \tag{9b}$$

$$M = J\dot{\Omega} + \Omega \times J\Omega \tag{9c}$$

Given the MAV desired velocity and acceleration derived from φ, we define position, velocity, orientation and angular velocity errors, and determine the MAV control inputs (for suitably selected gain vectors k_x, k_v, k_R and k_Ω) as

$$f = \left(-k_x e_x - k_v e_v + mge_3 + m\ddot{x}_d \right) \cdot Re_3 \tag{10a}$$

$$M = -k_R e_R - k_\Omega e_\Omega + \Omega \times J\Omega \tag{10b}$$

which have been shown [9] to establish almost global exponential convergence of the error dynamics to the origin for

all initial conditions —and exponential convergence under some restriction on the initial conditions.

The structure of the overall architecture for planning, control and formation generation is shown in Fig. 3.
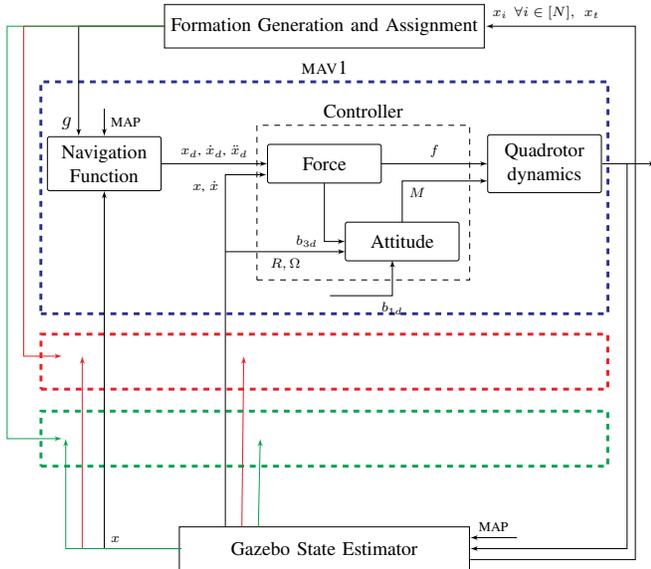
Fig. 3. Current position of the MAVs and the target is utilized in (4) to generate the formation and the assigned goal points are fed into the planner running on each MAV. The navigation function also takes the states of the MAV and the target from from GAZEBO and feeds the smooth position, velocity and acceleration trajectories to the force controller that generates the required force and feeds the stabilizing direction b_{3d} [9] to the attitude controller. Given a suitable heading direction (assume constant here) b_{1d} [9] and utilizing current angular velocity and orientation, the attitude controller generates required moments. The workspace map is assumed known. Red and Green boxes and lines indicate similar methodology for other MAVs.

V. SIMULATIONS

The $15 \times 15\,\text{m}$ simulated workspace along with three Asctec Firefly quadrotors and a Turtlebot in the role of the target is shown in Fig. 4. State estimation for each MAV and the target is taken from the simulated odometry sensors in GAZEBO. The motion of the target and the MAVs is shown in Fig. 5. The target movement (black curve) has multiple phases with different linear and angular velocity depicted in different shades of black. Three MAVs, each running its own planner and low-level controller chase the target taking goal point from the centralized algorithm that generate the formation. All the algorithms run at $100\,\text{Hz}$ while the optimization is solved using GUROBI optimizer [7].

Figures 6(a), 6(b) and 6(c) show the time history of inter-MAV distances, the distance of each MAV to the target, and the value of navigation function for each MAV, respectively. Note that the distance between MAV1 and MAV3 falls below the safety distance for few seconds when they approach the target; this is technically permitted since during the approach phase the navigation functions on each MAV do not enforce collision avoidance between the vehicles.[7] We stress that the focus of this paper (see Section I) is the integration of the differential geometric controller [9] with methodologies

[7] The introduction of this additional feature is part of ongoing work.
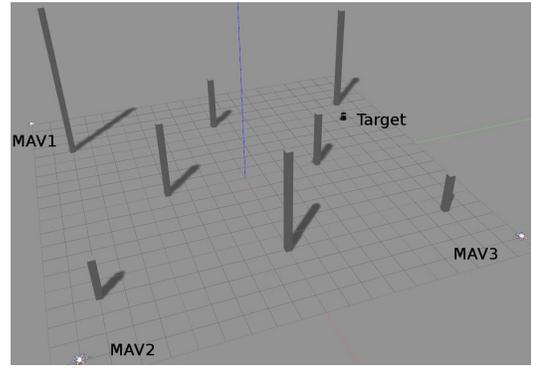
Fig. 4. Simulated workspace. Three Firefly quadrotors start from different locations and intercept the moving Turtlebot.

for motion planning based on navigation functions, and the determination of the optimal for radiation detection mobile sensor configuration around a moving source. The distances to the target, however, being enforced in the optimal formation configuration, always remain above the safety distances (Fig. 6(b)). The decreasing values of navigation function for each MAV indicates that a MAV does not collide with any obstacle during the target interception phase.

The video submission visualizes the scenario described: experiment 1 presents GAZEBO simulations using quadrotor models and state estimation from ROTORS simulation package [6].[8] Experiment 2 depicts a more obstacle-dense environment (with more than double the number of obstacles) run with the same MAV motion parameters, suggesting that the method's can be robust to variations in the structure of the environment; the tuning of the navigation function, however, can still present some challenges as the density of obstacles increases. As expected, the tuning of the navigation function influences the nature of the MAV paths generated, with those produced by a well-tuned function resembling geodecics. A detailed probabilistic analysis can help in determining the threshold in the density of obstacles at which solutions suffer, but it is beyond the scope of the current work.

Our formation design methodology (Section III) works adequately up to six agents and have been tested and validated with five agents (Experiment 3). That latter case study is performed in the same workspace as Experiment 1 only with more MAVs. Although an upper limit on the number of deployable MAVs could only be established through a detailed parametric analysis —beyond the scope of the present work— our belief is that cluttered environments are better served by relatively small radiation sensor MAV teams. When flight safety constraints are considered, a trade-off emerges between formation size (information quantity) and individual MAV SNR (information quality), which has an impact on the accuracy of cooperative radiation detection [19].

[8] The ROS implementation of the navigation function planner with the low-level controller and formation generation scripts can be found at https://github.com/indsy123/Navigation-function-based-planning-and-control.
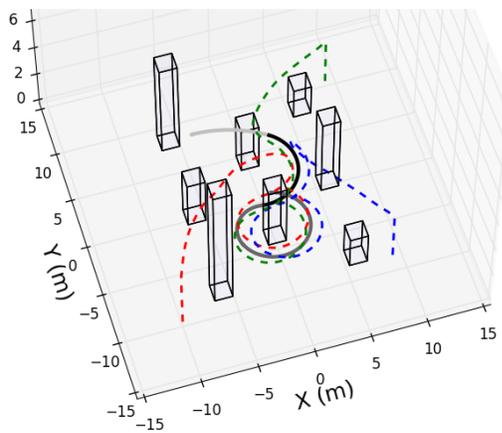
Fig. 5. Black and grey curve depicts the motion of the target. Initial light grey period corresponds to a linear and angular velocity of 1.1ms^{-1} and -0.08rads^{-1}. This is followed by two different phases (solid black and dark grey) of target motion in which the linear and angular velocities were $(1.8\text{ms}^{-1}, -0.43\text{rads}^{-1})$ and $(1.5\text{ms}^{-1}, 0.65\text{rads}^{-1})$ respectively. Towards the end velocity remains the same but angular velocity reduces to 0.5rads^{-1}. Dashed red, blue and green curves shows the movement of the three MAVs chasing the target at 3ms^{-1}.

VI. CONCLUSIONS

The strong dependence on the quality of collected data [19] of the detection accuracy of a swarm of mobile radiation sensors, and the nonlinear dependence of the radiation measurement SNR on the distance between the source and the detector, dictates that the collective decision-making is optimal not when individual detector selfishly improves its own chance to decide in a greedy fashion, but when the swarm allows one of its members to obtain the best possible placement while the rest share the next-best views. This strategy gives priority to satisfying the constraint on PFA, ensuring first that the decision-making problem becomes feasible as soon as possible. This insight is utilized to relax the formation design problem into a relatively simple QP.

The paper demonstrates that navigation functions can be coupled with differential-geometric quadrotor controllers to provide a truly feedback-based strategy for MAV navigation and target tracking in (known) cluttered environments. It is noteworthy, however, that the two key paper contributions, namely the optimal formation design for radiation detection, and the integration of navigation functions in geometric MAV control, are each useful and applicable in their own right.

REFERENCES

[1] J. Alonso-Mora, E. Montijano, M. Schwager, and D. Rus. Distributed multi-robot formation control among obstacles: A geometric and optimization approach with consensus. In *2016 IEEE International Conference on Robotics and Automation*, pages 5356–5363, May 2016.
[2] Javier Alonso-Mora, Eduardo Montijano, Tobias Nägeli, Otmar Hilliges, Mac Schwager, and Daniela Rus. Distributed multi-robot formation control in dynamic environments. *Autonomous Robots*, Jul 2018.
[3] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
[4] J. Chen, T. Liu, and S. Shen. Tracking a moving target in cluttered environments using a quadrotor. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 446–453, Oct 2016.
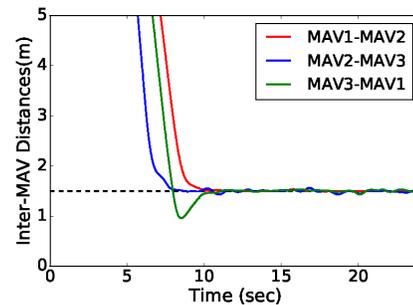
(a) Inter-MAV Distances(m)
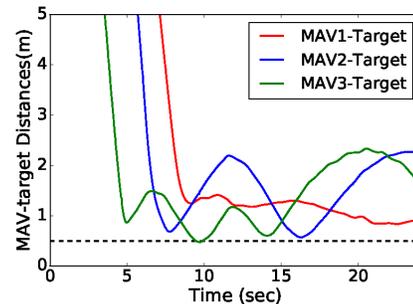
(b) MAV-Target Distances(m)
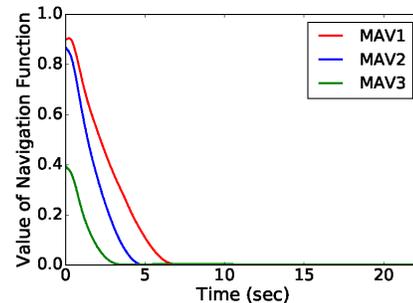
(c) Value of Navigation Function

Fig. 6. Simulation Results: (a) Variation of inter-MAV distance with time and its comparison with safety distance (Dashed Black line). (b) Similar comparison of MAV-target distances with the safety distance. (c) Value of Navigation Function over the entire trajectory of each MAV.

[5] Z. Cook, M. Kazemeini, and A. Barzilov. Low-altitude contour mapping of radiation fields using uas swarm. *Intelligent Service Robotics*, pages 1–12, April 2019.
[6] Fadri Furrer, Michael Burri, Markus Achtelik, and Roland Siegwart. *Robot Operating System (ROS): The Complete Reference (Volume 1)*, chapter RotorS—A Modular Gazebo MAV Simulator Framework, pages 595–625. Springer International Publishing, Cham, 2016.
[7] Gurobi Optimization Inc. Gurobi Optimizer Reference Manual, 2014. http://www.gurobi.com.
[8] H. W. Kuhn and Bryn Yaw. The hungarian method for the assignment problem. *Naval Res. Logist. Quart*, pages 83–97, 1955.
[9] T. Lee, M. Leoky, and N. H. McClamroch. Geometric tracking control of a quadrotor uav on se(3). In *Proceedings of 49th IEEE Conference on Decision and Control*, pages 5420–5425, 2010.
[10] C. Li and H. G. Tanner. Navigation functions with time-varying destination manifolds in star worlds. *IEEE Transactions on Robotics*, 35(1):35–48, Feb 2019.
[11] E. Montijano, E. Cristofalo, D. Zhou, M. Schwager, and C. Sags. Vision-based distributed formation control without an external positioning system. *IEEE Transactions on Robotics*, 32(2):339–351, April 2016.
[12] R. J. Nemzek, J. S. Dreicer, D. C. Torney, and T. T. Warnock. Distributed sensor networks for detection of mobile radioactive sources. *IEEE Transactions on Nuclear Science*, 51(4):1693–1700, Aug 2004.

[13] Chetan D. Pahlajani, Ioannis Poulakakis, and Herbert G. Tanner. Networked decision making for Poisson processes with applications to nuclear detection. *IEEE Transactions on Automatic Control*, 59(1):193–198, 2014.

[14] Chetan D. Pahlajani, Jianxin Sun, Ioannis Poulakakis, and Herbert G. Tanner. Error probability bounds for nuclear detection: Improving accuracy through controlled mobility. *Automatica*, 50(10):2470–2481, 2014.

[15] Jianxin Sun and Herbert G. Tanner. Constrained decision-making for low-count radiation detection by mobile sensors. *Autonomous Robots*, 39(4):519–536, 2015.

[16] J. Thomas, J. Welde, G. Loianno, K. Daniilidis, and V. Kumar. Autonomous flight for detection, localization, and tracking of moving targets with a small quadrotor. *IEEE Robotics and Automation Letters*, 2(3):1762–1769, July 2017.

[17] Matthew Turpin, Nathan Michael, and Vijay Kumar. Trajectory design and control for aggressive formation flight with quadrotors. *Autonomous Robots*, 33(1):143–156, Aug 2012.

[18] Luis R. Valbuena and Herbert G. Tanner. Flocking, formation control and path following for a group of mobile robots. *IEEE Transactions on Control Systems Technology*, 23(4):1268–1282, 2015.

[19] Indrajeet Yadav, Chetan D. Pahlajani, Herbert G. Tanner, and Ioannis Poulakakis. Information-sharing and decision-making in networks of radiation detectors. *Autonomous Robots*, 42(8):1715–1730, Dec 2018.