

Resilient Active Target Tracking with Multiple Robots

Lifeng Zhou,¹ Vasileios Tzoumas,² George J. Pappas,³ Pratap Tokekar¹

Abstract—The problem of target tracking with multiple robots consists of actively planning the motion of the robots to track the targets. A major challenge for practical deployments is to make the robots resilient to failures. In particular, robots may be attacked in adversarial scenarios, or their sensors may fail or get occluded. In this paper, we introduce planning algorithms for multi-target tracking that are resilient to such failures. In general, resilient target tracking is computationally hard. Contrary to the case where there are no failures, no scalable approximation algorithms are known for resilient target tracking when the targets are indistinguishable, or unknown in number, or with unknown motion model. In this paper we provide the first such algorithm, that also has the following properties: First, it achieves maximal resiliency, since the algorithm is valid for any number of failures. Second, it is scalable, as our algorithm terminates with the same running time as state-of-the-art algorithms for (non-resilient) target tracking. Third, it provides provable approximation bounds on the tracking performance, since our algorithm guarantees a solution that is guaranteed to be close to the optimal. We quantify our algorithm’s approximation performance using a novel notion of curvature for monotone set functions subject to matroid constraints. Finally, we demonstrate the efficacy of our algorithm through MATLAB and Gazebo simulations, and a sensitivity analysis; we focus on scenarios that involve a known number of distinguishable targets.

Index Terms—Multi-Robot Systems; Planning, Scheduling and Coordination; Robust/Adaptive Control of Robotic Systems.

I. INTRODUCTION

TASKS such as surveillance, exploration, and security often require the capability to detect, localize, and track targets within a prescribed area. For example, consider the tasks:

- (*Surveillance*) Detect and localize invasive fish in an ecosystem; [1]
- (*Area monitoring*) Detect and localize trapped people in a burning building; [2]

Manuscript received September, 10, 2018; Accepted October, 26, 2018. This paper was recommended for publication by Editor Nak Young Chong upon evaluation of the Associate Editor and Reviewers’ comments. This work was partially supported by the ARL DCIST CRA W911NF-17-2-0181 program, and by the NSF grants 1566247 and 1637915.

¹The authors are with the Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061 USA (email: {lfzhou, tokekar}@vt.edu).

²At the time the paper was written, the author was with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104-6228 USA. Currently, the author is with the Laboratory for Information & Decision Systems (LIDS), and the Department of Aeronautics and Astronautics (AeroAstro), Massachusetts Institute of Technology, Cambridge, MA 02139 USA (email: vtzoumas@mit.edu).

³The author is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (email: pappasg@seas.upenn.edu).

Digital Object Identifier (DOI): see top of this page.




Fig. 1. Aerial robots mounted with down-facing cameras to track multiple targets —depicted as dots— on the ground. The targets may be mobile; distinguishable or indistinguishable; known or unknown in number; or even with only partially known motion model. If a robot is under attack, its camera will be blocked.

- (*Patrolling*) Detect and localize adversarial agents that move in an urban environment. [3]

These tasks can greatly benefit by the use of robots that act as mobile sensors. Indeed, advancements in robotic mobility, sensing, and communication envision the deployment of collaborative robots to support target tracking [4]. The problem of planning the (joint) motion of robots for target tracking is known as *multi-robot active target tracking* in the literature [5]. This is a challenging problem due to the fact that the targets may be mobile and whose motion model may only be partially known. The targets may even be moving adversarially. There may be a large number of targets (more than the number of robots), even unknown in number, and may be indistinguishable from each other. Nevertheless, a number of algorithms have been designed that ensure near-optimal tracking for all of the aforementioned scenarios [5]–[12].

In this paper, we focus on scenarios where the robots operate in failure-prone or adversarial environments. In such environments, the robots may be subject to attacks leading to robotic failures [13], or the robots’ fields-of-view may become obstructed due to environmental hazards [14], or their sensors may fail completely [15] (see also Fig. 1). In particular, we consider failures and attacks that render robots unable to detect targets. Our goal is to provide planning and coordination algorithms that are resilient to such failures.

In this paper, we introduce a problem of resilient target tracking that guards against worst-case failures even when the targets are indistinguishable, unknown in number, or even with unknown motion model. By worst-case failures, we refer to

scenarios where at most α robots may fail, where α is known. Resilient target tracking is a computationally challenging problem since it needs to account for all possible failures, a problem of combinatorial complexity; and even in the presence of no failures, the problem is NP-hard [11]. This computational challenge motivates the main goal in this paper: to provide a scalable and provably close-to-optimal approximation algorithm. To this end, we capitalize on recent algorithmic results on resilient optimization subject to matroids [16], and present an approximation algorithm for resilient target tracking.

The algorithm runs in 2 steps: first, it approximates the worst-case subset of robots that could fail, and chooses their trajectories; and second, assuming that the approximated worst-case robots will actually fail, it jointly plans greedily the trajectories for the rest of the robots.

Contributions. In this paper, we make the contributions:

- *(Problem)* We formalize the problem of resilient active target tracking against *worst-case* failures even in the presence of targets that are (possibly) indistinguishable, unknown in number, or even of partially unknown motion model. This is the first work to formalize this problem.
- *(Solution)* We develop the first algorithm for the problem, and prove it has the following properties:
 - *maximal resiliency*: the algorithm is valid for any number of robot and/or sensor failures;
 - *minimal running time*: the algorithm terminates with the same running time as state-of-the-art algorithms for non-resilient target tracking;
 - *provable approximation performance*: the algorithm ensures a close-to-optimal solution for any target tracking objective function that is monotone and submodular (submodularity is a diminishing returns property [17]). Examples of such functions are the expected number of detected targets at a prescribed time and the mutual information between the predicted targets' location and the robots' sensor measurements [12].
- *(Empirical Evaluation)* We demonstrate with MATLAB and Gazebo simulations both the necessity for resilient target tracking against robot failures, and the efficacy and robustness of our approach. To this end, we focus on scenarios that involve distinguishable targets (known in number), and also conduct sensitivity analysis against non-worst-case attacks (random and greedy attacks).

Overall, in this paper we go beyond non-resilient target tracking [5]–[12] by proposing resilient target tracking; and beyond resilient tracking with distinguishable and known targets [18] by proposing resilient tracking with targets that are (possibly) indistinguishable, and/or unknown.

Organization of rest of the paper. Section II formulates resilient target tracking (Problem 1). Section III presents the first scalable algorithm for Problem 1. Section IV presents the main result in this paper: the scalability and performance guarantees of the proposed algorithm. Section V presents MATLAB and Gazebo simulations. Section VI concludes the paper.

Notation. Calligraphic fonts denote sets (e.g., \mathcal{A}). Given a set \mathcal{A} , then $2^{\mathcal{A}}$ denotes the power set of \mathcal{A} ; $|\mathcal{A}|$ denotes \mathcal{A} 's

cardinality; given another set \mathcal{B} , the set $\mathcal{A} \setminus \mathcal{B}$ denotes the set of elements in \mathcal{A} that are not in \mathcal{B} . Given a set \mathcal{V} , a set function $f : 2^{\mathcal{V}} \mapsto \mathbb{R}$, and an element $x \in \mathcal{V}$, $f(x)$ is a shorthand that denotes $f(\{x\})$.

II. PROBLEM FORMULATION

We formalize the problem of resilient multi-target tracking. In particular, the problem consists of planning the motion of the robots to optimally track targets despite robotic/sensor failures. The optimality of tracking is captured by an objective function such as the expected number of detected targets or the reduction in the uncertainty of the targets' positions.

A. Framework

Attacks: We assume that the maximum number of robotic/sensor failures are known, and denote it by α .¹ At any time at most α robots/sensor may fail. In addition, without loss of generality, the set of robots that fail may vary over time. A robot that fails at time t may be active at another time t' .

The rest of our problem formulation, e.g., assumptions about the targets, robots, and the objective function, follows the standard in the target tracking literature; see [12] and the references therein. Specifically:

Targets: Targets exist in an area of interest (environment). The targets can be ground or aerial vehicles, and can be mobile or immobile. They can be distinguishable [11] or indistinguishable [12]. Their number can be known [11] or unknown [12], fixed [11] or time-varying [12]. The target motion model can be known (e.g., a single integrator with known maximal speed [19]) or partially known; in the latter case, data-driven learning techniques may be employed [12].

Robots/sensors: We consider a team of mobile robots, and denote it by \mathcal{R} . The team is tasked to track targets in an area of interest. The robots can be ground or aerial vehicles (e.g., quad-rotors). We assume that the robots can communicate with each other at all times.

The robots carry onboard sensors (e.g., cameras or lidars), which enable the team's tracking capability. In particular, each robot $r \in \mathcal{R}$, at every time t , receives measurements from targets detected in the field-of-view of its sensors. Additional measurements may be obtained from off-board sensors in the environment. Given the measurements and a target model, each robot employs a detector and trajectory estimator. If the targets are distinguishable and their number is known, a Kalman or particle filter can be employed [20], whereas, if the targets are indistinguishable and their number is unknown, a Random Finite Sets (RFS) filter can be employed [12]. In both cases, the robots have only an estimate of the targets' true positions. The estimate is represented by a set of possible target locations in the environment. Given a target model, the robots propagate this set to obtain a predicted target position, by employing one of the above techniques.

The robots can also move in the environment, and detect multiple targets per motion step (Fig. 2). The robot trajectory generation framework is as follows: We assume that the robots

¹Henceforth, we refer to robotic and sensor failures interchangeably.




Fig. 2. Available robotic trajectories (denoted by τ_i^j for robot r_i) and their coverage region (denoted by $C(\tau_i^j)$ for robot r_i that chooses trajectory τ_i^j , and depicted as shaded regions).

have perfect localization (e.g., using GPS).² Time is divided into rounds of finite duration denoted by T (without loss of generality, we assume it fixed). At the beginning of each round, each robot generates a set of candidate trajectories, one of which will be followed in the current round. The trajectories can be generated by employing, for example, motion-space discretization [12] or spatial-sampling methods [21]. We denote the set of valid trajectories, for a round that starts at a time t , and for a robot $r \in \mathcal{R}$, by $\mathcal{T}_{r,t}$. We denote by $\mathcal{T}_{\mathcal{R},t}$, the set of all robots' valid trajectories, i.e., $\mathcal{T}_{\mathcal{R},t} \triangleq \cup_{r \in \mathcal{R}} \mathcal{T}_{r,t}$. Each trajectory in $\mathcal{T}_{\mathcal{R},t}$ is interpolated to yield a sequence of robot poses where the robot will take a measurement. That is, each trajectory corresponds to a sequence of robot poses. Without loss of generality, we assume the number of poses to be fixed across rounds, robots, and trajectories.

Target tracking objective function: Given a round that starts at a time t , the utility of each trajectory in $\mathcal{T}_{\mathcal{R},t}$ is captured by an objective function f . Two examples of f are the following [12]: the expected number of detected targets in the current round (time interval from t to $t+T$); and the mutual information between the predicted location of the targets at time $t+T$ and the collected measurements in the current round. Notably, both functions are *monotone* and *submodular* in the choice of the robots' trajectories [12, Lemma 1 and Lemma 2]. Submodularity is a diminishing returns property [17]; we provide its definition in Appendix A of the full version of this paper, which is found at the authors' websites. We henceforth focus on functions that satisfy these two properties.

B. Problem definition

Problem 1 (Resilient Multi-Target Tracking with Multiple Robots). *In reference to Section II-A's framework, consider: a set of targets; a set \mathcal{R} of mobile robots/sensors; a division of time into rounds of finite duration. Moreover, consider the beginning of a round, and the corresponding set of valid robot trajectories $\mathcal{T}_{\mathcal{R}} \triangleq \cup_{r \in \mathcal{R}} \mathcal{T}_r$, where \mathcal{T}_r is the set of valid trajectories for the robot r . Finally, consider a target tracking objective function f that is monotone and submodular (e.g., the expected number of detected targets in the current round).*

The problem of resilient multi-target tracking with mobile robots is to achieve a maximal value for f , by selecting the robot trajectories throughout the round, despite a worst-case failure of at most α robots/sensors. Formally:

$$\begin{aligned} \max_{\mathcal{S} \subseteq \mathcal{T}_{\mathcal{R}}} \min_{\mathcal{A} \subseteq \mathcal{S}} f(\mathcal{S} \setminus \mathcal{A}) : \\ |\mathcal{S} \cap \mathcal{T}_r| = 1, \quad r \in \mathcal{R}; \\ |\mathcal{A}| \leq \alpha, \end{aligned} \quad (1)$$

where: \mathcal{S} is the set of selected trajectories for all robots; the constraint $|\mathcal{S} \cap \mathcal{T}_r| \leq 1$ for each robot $r \in \mathcal{R}$ represents the natural constraint that each robot r can follow one trajectory;³ and the constraint $|\mathcal{A}| \leq \alpha$ captures the problem assumption that at most α robots/sensors can fail.

Problem 1 may be interpreted as a 2-stage perfect information sequential game between two players [22, Chapter 4], namely, a "maximization" player (who aims for optimal target tracking performance), and a "minimization" player (who aims to compromise the target tracking performance). In particular, the "maximization" player plays first by selecting the set \mathcal{S} , and, then, the "minimization" player observes \mathcal{S} , and plays second by selecting a worst-case attack/removal \mathcal{A} from \mathcal{S} . Evidently, this is a stricter (worse) version of the problem where the "minimization" player cannot observe \mathcal{S} .

Problem 1 goes beyond the traditional objective of target tracking with mobile robots, by protecting (in a receding horizon fashion) the robots' motion plan against failures.

III. ALGORITHM FOR PROBLEM 1

We present the first scalable algorithm for Problem 1, by capitalizing on the algorithmic results in [16]. The pseudo-code of the algorithm is described in Algorithm 1.

A. Intuition behind Algorithm 1

Problem 1 selects trajectories for all robots, denoted by the set \mathcal{S} in eq. (1), so to maximize the value of the objective function f despite that \mathcal{S} can incur a removal \mathcal{A} of α elements due to robotic failures. In this context, Algorithm 1 aims to maximize f by constructing \mathcal{S} as the union of two sets, the \mathcal{S}_1 and \mathcal{S}_2 (line 16), whose role we describe below.

Set \mathcal{S}_1 approximates a worst-case set removal from \mathcal{S} : Algorithm 1 aims with the trajectory set \mathcal{S}_1 to capture a worst-case removal of α trajectories among the trajectories Algorithm 1 will select in \mathcal{S} . Equivalently, \mathcal{S}_1 is aimed to act as a "bait" to an attacker that selects to remove the *best* α trajectories from \mathcal{S} (*best* with respect to the trajectories' contribution towards maximizing the function f). However, the problem of selecting the *best* trajectories in $\mathcal{T}_{\mathcal{R}}$ per Problem 1 is combinatorial and, in general, intractable [23]. For this reason, Algorithm 1 aims to *approximate* the best set of α trajectories, by letting \mathcal{S}_1 be the trajectories with the largest contributions to the value of f (lines 3). In addition, since \mathcal{S} needs to satisfy the constraint that each robot $r \in \mathcal{R}$ can be assigned one trajectory, Algorithm 1 constructs \mathcal{S}_1 so that not only $|\mathcal{S}_1| \leq \alpha$ but also $|\mathcal{S}_1 \cap \mathcal{T}_r| \leq 1$ for all $r \in \mathcal{R}$ (lines 4-6).

²The uncertainty in the robot's position can be incorporated in the uncertainty in the targets' estimates [11].

³This type of constraint is called a *partition matroid* in the literature of combinatorial optimization [17].

Algorithm 1 Scalable algorithm for Problem 1.**Input:** Per Problem 1, Algorithm 1 receives the inputs:

- set of robots \mathcal{R} ;
- robot trajectories \mathcal{T}_r , for all robot $r \in \mathcal{R}$;
- target tracking objective function f ;
- maximum number of failures α .

Output: Robots' trajectories \mathcal{S} .

```

1:  $\mathcal{S}_1 \leftarrow \emptyset$ ;  $\mathcal{M}_1 \leftarrow \emptyset$ ;  $\mathcal{S}_2 \leftarrow \emptyset$ ;  $\mathcal{M}_2 \leftarrow \emptyset$ ;
2: while  $\mathcal{M}_1 \neq \mathcal{T}_{\mathcal{R}}$  do
3:    $s \in \arg \max_{y \in \mathcal{T}_{\mathcal{R}} \setminus \mathcal{M}_1} f(y)$ ;
4:   if for all  $r \in \mathcal{R}$  it is  $|(\mathcal{S}_1 \cup \{s\}) \cap \mathcal{T}_r| \leq 1$ , and  $|\mathcal{S}_1 \cup \{s\}| \leq \alpha$  then
5:      $\mathcal{S}_1 \leftarrow \mathcal{S}_1 \cup \{s\}$ ;
6:   end if
7:    $\mathcal{M}_1 \leftarrow \mathcal{M}_1 \cup \{s\}$ ;
8: end while
9: while  $\mathcal{M}_2 \neq \mathcal{T}_{\mathcal{R}} \setminus \mathcal{S}_1$  do
10:   $s \in \arg \max_{y \in \mathcal{T}_{\mathcal{R}} \setminus (\mathcal{S}_1 \cup \mathcal{M}_2)} f(\mathcal{S}_2 \cup \{y\}) - f(\mathcal{S}_2)$ ;
11:  if for all  $r \in \mathcal{R}$  it is  $|(\mathcal{S}_1 \cup \mathcal{S}_2 \cup \{s\}) \cap \mathcal{T}_r| \leq 1$  then
12:     $\mathcal{S}_2 \leftarrow \mathcal{S}_2 \cup \{s\}$ ;
13:  end if
14:   $\mathcal{M}_2 \leftarrow \mathcal{M}_2 \cup \{s\}$ ;
15: end while
16:  $\mathcal{S} \leftarrow \mathcal{S}_1 \cup \mathcal{S}_2$ ;

```

Set \mathcal{S}_2 is such that the set $\mathcal{S}_1 \cup \mathcal{S}_2$ approximates solution to Problem 1: Assuming \mathcal{S}_1 is the set to be removed from Algorithm 1's selection \mathcal{S} , Algorithm 1 needs to select a trajectory set \mathcal{S}_2 to complete the construction of \mathcal{S} . In particular, for $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$ to be a solution to Problem 1, Algorithm 1 needs to select \mathcal{S}_2 as a *best* set of trajectories from $\mathcal{T}_{\mathcal{R}} \setminus \mathcal{S}_1$ subject to the natural constraint that one trajectory is assigned to each robot (lines 11-13). Nevertheless, the problem of selecting a *best* set of elements subject to such a constraint is combinatorial and, in general, intractable [23]. Hence, Algorithm 1 aims to *approximate* such a best set, using the greedy procedure in lines 9-15.

Overall, Algorithm 1 constructs the sets \mathcal{S}_1 and \mathcal{S}_2 to approximate with their union \mathcal{S} an optimal solution to Problem 1.

We next describe the steps in Algorithm 1 in more detail.

B. Description of steps in Algorithm 1

Algorithm 1 executes four steps:

a) *Initialization (line 1)*: Algorithm 1 defines 4 sets, the \mathcal{S}_1 , \mathcal{M}_1 , \mathcal{S}_2 , and \mathcal{M}_2 , and initializes each of them with the empty set (line 1). The purpose of \mathcal{S}_1 and \mathcal{S}_2 is to construct the set \mathcal{S} . Specifically, the union of \mathcal{S}_1 and \mathcal{S}_2 constructs \mathcal{S} by the end of Algorithm 1 (line 16). The purpose of \mathcal{M}_1 and of \mathcal{M}_2 is to support the construction of \mathcal{S}_1 and \mathcal{S}_2 . During the construction of \mathcal{S}_1 , Algorithm 1 stores in \mathcal{M}_1 the trajectories in $\mathcal{T}_{\mathcal{R}}$ that have either been included already or cannot be included in \mathcal{S}_1 (line 7); that way, Algorithm 1 keeps track of which trajectories remain to be checked if they could be added in \mathcal{S}_1 (line 5). During the construction of \mathcal{S}_2 , Algorithm 1 stores in \mathcal{M}_2 the trajectories of $\mathcal{T}_{\mathcal{R}} \setminus \mathcal{S}_1$ that have either been included already or cannot be included in \mathcal{S}_2 (line 14); that

way, Algorithm 1 keeps track of which trajectories remain to be checked if they could be added in \mathcal{S}_2 (line 12).

b) *Construction of set \mathcal{S}_1 (lines 2-8)*: Algorithm 1 constructs \mathcal{S}_1 sequentially by adding one trajectory at a time from $\mathcal{T}_{\mathcal{R}}$ to \mathcal{S}_1 . Specifically, \mathcal{S}_1 , being the “bait” set, is constructed such that it satisfies both the trajectory assignment constraint (one trajectory per robot) and the failures cardinality constraint (line 4). Also, \mathcal{S}_1 is constructed such that each trajectory $s \in \mathcal{T}_{\mathcal{R}}$ added in \mathcal{S}_1 achieves the highest value of $f(s)$ among all the trajectories in $\mathcal{T}_{\mathcal{R}}$ that have not been yet added in \mathcal{S}_1 and can be added in \mathcal{S}_1 (line 5).

c) *Construction of set \mathcal{S}_2 (lines 9-15)*: Algorithm 1 constructs the set \mathcal{S}_2 sequentially, by picking greedily trajectories from the set $\mathcal{T}_{\mathcal{R}} \setminus \mathcal{S}_1$ such that $\mathcal{S}_1 \cup \mathcal{S}_2$ satisfies the trajectory assignment constraint in Problem 1 (one trajectory per robot). Specifically, the greedy procedure in Algorithm 1's “while loop” (lines 9-15) selects a trajectory $y \in \mathcal{T}_{\mathcal{R}} \setminus (\mathcal{S}_1 \cup \mathcal{M}_2)$ to add in \mathcal{S}_2 only if y maximizes the value of $f(\mathcal{S}_2 \cup \{y\}) - f(\mathcal{S}_2)$, where the set \mathcal{M}_2 stores the trajectories that either have already been added to \mathcal{S}_2 or have been considered to be added to \mathcal{S}_2 but they were not since the resultant set $\mathcal{S}_1 \cup \mathcal{S}_2$ would not satisfy the trajectory assignment constraint.

d) *Construction of set \mathcal{S} (line 16)*: Algorithm 1 constructs the set \mathcal{S} as the union of the previously constructed sets \mathcal{S}_1 and \mathcal{S}_2 (lines 16).

In summary, Algorithm 1 proposes a trajectory assignment \mathcal{S} as solution to Problem 1. In particular, Algorithm 1 constructs \mathcal{S} to withstand any compromising robotic/sensor failure.

IV. PERFORMANCE ANALYSIS OF ALGORITHM 1

We quantify the performance of Algorithm 1, by bounding its running time, and its approximation performance. To this end, we use the following notion of curvature for set functions.

A. Constrained curvature of monotone submodular functions

Definition 1 (Constrained curvature). Consider a set \mathcal{T} , and a non-decreasing submodular set function $f : 2^{\mathcal{T}} \mapsto \mathbb{R}$ such that (without loss of generality) for any element $s \in \mathcal{T}$, it is $f(s) \neq 0$. Moreover, consider a collection of subsets of \mathcal{T} , denoted by \mathcal{I} ; e.g., \mathcal{I} represents admissible sets where f can be evaluated at. Then, the constrained curvature of f over \mathcal{I} is:

$$\nu_f(\mathcal{I}) \triangleq 1 - \min_{S \in \mathcal{I}} \min_{s \in S} \frac{f(S) - f(S \setminus \{s\})}{f(s)}. \quad (2)$$

The curvature ν_f measures how far f is from being additive. In particular, Definition 1 implies $0 \leq \nu_f \leq 1$: If $\nu_f = 0$, then for all sets $S \in \mathcal{I}$ it holds $f(S) = \sum_{s \in S} f(s)$. In contrast, if $\nu_f = 1$, then there exist a set $S \in \mathcal{I}$ and an element $s \in \mathcal{T}$ such that $f(S) = f(S \setminus \{s\})$; that is, in the presence of $S \setminus \{s\}$, the element s loses all its contribution to the value of $f(S)$. Notably, Definition 1 adapts the notion of curvature discussed in [24] to the case where the set \mathcal{S} is constrained in an \mathcal{I} , instead of \mathcal{S} being able to be any subset of \mathcal{T} .

For example, in reference to the target tracking framework of Section II, consider the expected number of detected targets as a function of the robot trajectories. Then, this function has




Fig. 3. Given a natural number n , plot of $h(n, \alpha)$ versus α . Given a finite n , then $h(n, \alpha)$ is always non-zero, with minimum value $2/(n+2)$, and maximum value 1.

curvature zero if each robot detects different targets from the rest of the robots. In contrast, it has curvature one if, for example, at least two robots by following their trajectories receive the exact same measurements.

B. Performance Analysis for Algorithm 1

Theorem 1 (Performance of Algorithm 1). Consider an instance of Problem 1, the notation therein, the notation in Algorithm 1, and the definitions:

- let the number f^* be the (optimal) value to Problem 1;
- given a set \mathcal{S} as solution to Problem 1, let $\mathcal{A}^*(\mathcal{S})$ be a worst-case set removal from \mathcal{S} , per Problem 1, that is: $\mathcal{A}^*(\mathcal{S}) \in \arg \min_{\mathcal{A} \subseteq \mathcal{S}, |\mathcal{A}| \leq \alpha} f(\mathcal{S} \setminus \mathcal{A})$. Evidently, a removal from \mathcal{S} corresponds to a set of robot/sensor failures;
- define $h(|\mathcal{R}|, \alpha) \triangleq \max[1/(1+\alpha), 1/(|\mathcal{R}|-\alpha)]$.⁴

Finally, without loss of generality, consider that $f(\emptyset) = 0$.

The performance of Algorithm 1 is bounded as follows:

- 1) (Approximation performance) Algorithm 1 returns a trajectory set \mathcal{S} such that each robot is assigned a single trajectory, and:

$$\frac{f(\mathcal{S} \setminus \mathcal{A}^*(\mathcal{S}))}{f^*} \geq \frac{\max[1 - \nu_f(\mathcal{I}), h(|\mathcal{R}|, \alpha)]}{2}, \quad (3)$$

where \mathcal{I} is the collection of valid trajectory assignments to robots per Problem 1 (one trajectory per robot), that is: $\mathcal{I} \triangleq \{\mathcal{S} : \mathcal{S} \subseteq \mathcal{T}_{\mathcal{R}}, |\mathcal{S} \cap \mathcal{T}_r| = 1 \text{ for all } r \in \mathcal{R}\}$.

- 2) (Running time) Algorithm 1 constructs the trajectory set \mathcal{S} as a solution to Problem 1 with $O(|\mathcal{T}_{\mathcal{R}}|^2)$ evaluations of f .

The proof of Theorem 1 is the same as the proof of [16, eq. (7) in Theorem 1].

Provable approximation performance. Theorem 1 implies on the approximation performance of Algorithm 1:

Near-optimality: Algorithm 1 guarantees a value finitely close to the optimal, for any monotone submodular objective function f : per ineq. (3), Algorithm 1’s approximation factor is bounded by $h(|\mathcal{R}|, \alpha)/2$, which is non-zero for any finite number of robots $|\mathcal{R}|$ (see also Fig. 3). Similarly, the approximation factor is also bounded by $(1 - \nu_f)/2$, which is also non-zero for any monotone submodular f with $\nu_f < 1$.

⁴A plot of $h(|\mathcal{R}|, \alpha)$ is found in Fig. 3.




Fig. 4. MATLAB simulation setup: Each robot r_i has 4 possible trajectories (forward, backward, left, and right, denoted by τ_i^j for $j = 1, 2, 3, 4$, respectively). The tracking region of each trajectory is rectangular, is denoted by $\mathcal{C}(\tau_i^j)$ for the trajectory τ_i^j , and has the same dimensions across all 4 trajectories; in particular, the lengths l_t and l_o define the dimension of each rectangular region for each trajectory; and l_f defines the fly length for the robot. We set $l_t = l_f + l_o$.

Approximation performance for no failures: When the number of failures is zero ($\alpha = 0$), Algorithm 1’s approximation performance is the same as that of the state-of-the-art algorithms for (non-resilient) target tracking. In particular, these algorithms have approximation performance at least $1/2$ [11], [12]; at the same time, Algorithm 1 also has performance at least $1/2$ for $\alpha = 0$, since $h(|\mathcal{R}|, 0) = 1$ per ineq. (3).

Minimal running time. Theorem 1 implies that Algorithm 1, even though it goes beyond the objective of (non-resilient) target tracking, has the same order of running time as state-of-the-art algorithms for (non-resilient) target tracking. In particular, these algorithms terminate with $O(|\mathcal{T}_{\mathcal{R}}|^2)$ evaluations of the function f [11], [12], and Algorithm 1 also terminates with the same time.

V. NUMERICAL EVALUATION

We present MATLAB and Gazebo evaluations of our algorithm (Algorithm 1) that demonstrate both the necessity for resilient target tracking and the benefits of our approach. In particular, both evaluations demonstrate: (i) the near-optimal performance of Algorithm 1, since it performs close to the brute-force algorithm (which is viable only in small-scale scenarios); also, it is superior to the greedy and random heuristics; and (ii) the superior robustness of Algorithm 1 to scenarios where non-worst-case or even no attacks occur. Our MATLAB and Gazebo implementations are available online.⁵

Compared algorithms. We compare Algorithm 1 with three other algorithms. The algorithms differ in how they select the robot trajectories. The first algorithm is an optimal, brute-force algorithm, and it attains the optimal value for Problem 1. Evidently, the brute-force approach is viable only when the number of available robots is small. We refer to this algorithm by “brute-force.” The second algorithm is a greedy algorithm that ignores the possibility of robotic/sensor attacks, and picks greedily the robot trajectories per the algorithm

⁵https://github.com/raaslab/resilient_target_tracking.git




Fig. 5. MATLAB evaluation results: Performance comparison (average and standard deviation over 30 trials) of Algorithm 1 (called “resilient” in this figure) with the “brute-force” algorithm; the “greedy” algorithm; and the “random” algorithm. Performance is measured as the number of tracked targets, and is compared across 3 settings that differ on how an attacker would select to attack α robots so to minimize the number of tracked targets given the selected robot trajectories: in Fig. 5(a) and Fig. 5(b), the attacker uses a brute-force algorithm to find an optimal robot attack; in Fig. 5(c), the attacker uses the greedy algorithm in [17]; and in Fig. 5(d), the attacker chooses randomly (uniformly across all robots).

proposed in [17]; we refer to this algorithm by “greedy.” The third algorithm is a random algorithm that picks randomly (uniformly) the robot trajectories; we refer to this algorithm by “random.” Finally, we refer to Algorithm 1 by “resilient.”

Scalability of our algorithm. Since we compare our algorithm with the optimal brute-force algorithm (which is viable for only small-scale instances of Problem 1), in the following numerical evaluations we focus on small-scale problem instances of up to 6 robots and 60 targets. However, the MATLAB implementation of our algorithm runs within 1 second even for up to 100 robots with 10 failures. Our code is available online (see Footnote 5).

A. MATLAB evaluation over one step with static targets

We study the effect of the number of targets and of the attack strategy by running the algorithms over random instances of Problem 1 for a single round (one-step time horizon).

Simulation setup. We consider 6 robots and a number of targets m that varies from 30 to 60. We set the number of attacks α equal to 3 and 4. A top view of the robots and targets is shown in Fig. 4. We assume that each robot $r_i \in \mathcal{R}$ flies on a fixed plane and has 4 trajectories: forward, backward, left, and right, denoted by τ_i^j for $j = 1, 2, 3, 4$, respectively. Each robot r_i has a square field-of-view, centered at the planar position of robot r_i , and is illustrated in Fig. 4 by the darker blue square of dimension $l_o \times l_o$. Once each robot selects a trajectory, it flies a distance l_f along that trajectory. Thus, each trajectory τ_i^j has a rectangular tracking region with length $l_t \triangleq l_f + l_o$ and width l_o ; we set $l_t = 10$ and $l_o = 3$ for all robots. For each number of targets $m = 30, 31, \dots, 60$, the planar positions of the robots and targets are randomly generated in the 2D space $[0, 10] \times [0, 10] \in \mathbb{R}^2$, across 30 trials. We consider that the robots have already available an estimate of the targets position. For each trial, all algorithms are executed with the same initialization, i.e., the same positions of targets and robots. All algorithms are executed for one round.

The algorithms’ performance is captured as the number of tracked targets given the selected robot trajectories. We examine the performance across 3 settings that differ on how an attacker would select to attack α robots so to minimize

the performance of the remaining robots: we first consider an attacker that uses a brute-force algorithm to find an optimal robot attack; this scenario is in agreement with the definition of Problem 1, where the attacks are indeed worst-case attacks. Then, we consider an attacker that uses the greedy algorithm in [17] to approximate an optimal robot attack; and finally, we consider an attacker that chooses randomly a robot attack (uniformly across all robots). We examine the last two cases (Fig. 5(c) and Fig. 5(d)) as part of a sensitivity analysis of Algorithm 1’s performance against non-worst-case attacks.

Results. The comparison results are reported in Fig. 5. The following observations from Fig. 5 are due:

a) *Close-to-optimality of Algorithm 1:* Algorithm 1 is designed to guarantee superior performance in the presence of worst-case attacks; indeed, per Fig. 5(a) (and per Fig. 5(b)), Algorithm 1 —colored blue in Fig. 5— has on average superior performance to the greedy and random heuristics. In particular, Algorithm 1’s performance is close to the optimal achieved by the brute force algorithm (red in Fig. 5).

b) *Robustness of Algorithm 1’s performance to non-worst-case attacks:* Although Algorithm 1 is designed to guarantee superior performance for worst-case attacks, in practice, the attack of robots may not necessarily be the worst-case one. For example, from the perspective of an attacker, finding the optimal robot attack is also an NP-hard problem, since it constitutes a cardinality constrained submodular minimization problem [25]. It is therefore relevant to ask whether Algorithm 1, being an approximation algorithm, will indeed have a better target tracking performance when the attacks are non-worst-case. By comparing Fig. 5(a) with both Fig. 5(c) and Fig. 5(d), we observe that for each given number of targets (horizontal axes in each plot in Fig. 5) the performance of Algorithm 1 increases for non-worst-case attacks. For example, for 30 targets, when the attack is worst-case (Fig. 5(a)) Algorithm 1 achieves 14 tracked targets, whereas: when the attack is greedy the performance increases to 17 (Fig. 5(c)). When the attacks are random, the performance increases to 18 (Fig. 5(d)). Overall, since Algorithm 1 is designed to protect against worst-case attacks, it also protects at least equally well against non-worst-case attacks (as it would be expected).




Fig. 6. Gazebo simulation setup: 4 aerial robots and 30 ground mobile targets: (a) setup in Gazebo environment; and (b) setup in Rviz environment, where: each aerial robot is color-coded, and its coverage region is depicted with the same color. The targets are depicted as white cylindrical markers.

Importantly, that way an attacker is forced by Algorithm 1 to deploy a worst-case attack, which is exactly the scenario that Algorithm 1 guarantees protection from. That conclusion makes also irrelevant the observation that, for example, in the case of random attacks (Fig. 5(d)) both Algorithm 1 and the greedy heuristic perform similarly. Notably, in the case of greedy attacks (Fig. 5(c)), Algorithm 1 is again superior to both the greedy and the random heuristics. In general, in the above simulation setup, Algorithm 1 achieves a superior and close-to-optimal performance, and remains superior even against non-worst-case attacks.


B. Gazebo evaluation over multiple steps with mobile targets

We study the effect of the number of targets and of the attack strategy by running the algorithms across multiple rounds (multi-step time horizon). That way, we take into account the kinematics and dynamics of the robots, as well as the fact that the kinematics and dynamics of the robot, the actual trajectories of the targets, and the sensing noise may force the robots to track fewer targets than expected.


Simulation setup. We consider a scenario of 4 aerial robots tasked to track 30 ground mobile targets (Fig. 6(a)). We set the number of attacks α equal to 2. We also visualize the robots, their field-of-view, and the targets using the Rviz environment (Fig. 6(b)): in particular, we visualize the robots as spherical markers, their field-of-views as colored areas with the same color as their corresponding robot, and the targets as white cylindrical markers. Similarly to the MATLAB simulation setup above, each robot has 4 trajectories (forward, backward, left, and right), and flies on a different fixed plane (to avoid collision with other robots). Moreover, we set the tracking length $l_t = 6$ and width $l_o = 3$ for all robots. We assume each target has the single integrator motion model

$$p_t^j(k+1) = p_t^j(k) + v_t^j(k),$$

where p_t^j and v_t^j denote the position and the velocity of target $j = 1, \dots, 30$, respectively. The robots obtain noisy position measurements of all targets. They use a Kalman filter for updating the estimated position of the target at the next round. The targets' velocity is initialized to zero and is updated by



(a) Performance comparison against worst-case attacks



(b) Attack rate comparison

Fig. 7. Gazebo evaluation results: Comparison (average and standard deviation across the 50 rounds) of Algorithm 1 (called “resilient” in this figure) with the “brute-force” algorithm; the “greedy” algorithm; and the “random” algorithm. Performance is measured as the expected number of tracked targets. Fig. 7(b) compares the sensitivity of the algorithms’ performance against the case where no attacks are present per the definition of the attack rate in eq. (4)—that is, the smaller the attack rate, the better.

using two consecutive measurements and the time interval of these two measurements, as follows:

$$v_t^j(k') = (\tilde{p}_t^j(k') - \tilde{p}_t^j(k)) / (k' - k),$$

where $\tilde{p}_t^j(k')$ and $\tilde{p}_t^j(k)$ are two consecutive position measurements of target j at round k' and round k with $k' > k$.

For each algorithm (see “Compared algorithms” at the beginning of the section), at each round each robot selects one of its 4 trajectories. Then, the robots fly a $l_f = 3$ distance along their selected trajectory. When an attack happens, we assume that the attacked robot’s camera is turned-off; nevertheless, we assume that it can be active again at the next round, so that at each round the worst-case set of α robots is considered failed. We repeat this process for 50 rounds.

At each round, we capture the performance of each algorithm with the expected number of targets tracked. We first compare the algorithms with respect to the average and the standard deviation of the expected number of targets tracked. Moreover, we compare the sensitivity of the algorithms’ performance against the case where no attacks are present: specifically, we compare the average and the standard

deviation of their attack rate per round, which is defined—for an algorithm that selects a set of trajectories \mathcal{S} —by

$$\frac{f(\mathcal{S}) - f(\mathcal{S} \setminus \mathcal{A}^*(\mathcal{S}))}{f(\mathcal{S})}, \quad (4)$$

where $f(\mathcal{S})$ is the expected number of targets tracked in the presence of no attacks, and $f(\mathcal{S} \setminus \mathcal{A}^*(\mathcal{S}))$ is the expected number of targets tracked in the presence of an optimal attack $\mathcal{A}^*(\mathcal{S})$. All in all, the above definition of attack rate captures how much worse is the performance of an algorithm in the presence of attacks than in the absence of attacks. A video for this implementation is available online.⁶

Results. The comparison results are reported Fig. 7. The following observations from Fig. 7 are due:

a) *Close-to-optimality of Algorithm 1:* Fig. 7(a) suggests that Algorithm 1 has on average superior performance than the current heuristics (the greedy and the random). In particular, Algorithm 1’s performance is close to the optimal, as it is achieved by a brute force algorithm.

b) *Robustness of Algorithm 1’s performance to no-attacks:* Per Fig. 7(b), Algorithm 1 exhibits superior attack rate than the greedy and random heuristics, and as a result, when for example the scenario “at most α attacks per round” and the scenario “no attacks per round” happen with equal probability, Algorithm 1 still guarantees superior average performance.

All in all, in the above simulation setup, Algorithm 1 achieves a superior and close-to-optimal performance, and remains superior even when no-attacks happen.

VI. CONCLUDING REMARKS & FUTURE WORK

We take the first steps to protect critical target tracking tasks from robot failures (Problem 1). In particular, we provide the first algorithm for Problem 1, and proved its guaranteed performance against any number of failures, and even for targets that are indistinguishable and/or unknown. We demonstrate the need for resilient target tracking and the robustness of our algorithm with MATLAB and Gazebo evaluations. Notably, the results of this paper extend to any active information gathering problem where the goal is to choose from a set of admissible trajectories for each robot (at each motion step), similar to [12], [21].

This work opens a number of avenues for future research, both theoretical and experimental. Future theoretical work in theory includes the decentralized design of the robots’ motion plan. Moreover, online extensions of Algorithm 1, that guarantee near-optimality across multiple rounds, are natural next steps. Future experimental work includes real-world testing of our resilient target tracking framework in the context of practical applications of surveillance and patrolling.

REFERENCES

- [1] P. Tokekar, E. Branson, J. Vander Hook, and V. Isler, “Tracking aquatic invaders: Autonomous robots for monitoring invasive fish,” *IEEE Robotics & Automation Magazine*, vol. 20, no. 3, pp. 33–41, 2013.
- [2] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, “Cooperative air and ground surveillance,” *IEEE Robotics & Automation Magazine*, vol. 13, no. 3, pp. 16–25, 2006.
- [3] U. Zengin and A. Dogan, “Real-time target tracking for autonomous uavs in adversarial environments: A gradient search algorithm,” *IEEE Transactions on Robotics*, vol. 23, no. 2, pp. 294–307, 2007.
- [4] V. Kumar and N. Michael, “Opportunities and challenges with autonomous micro aerial vehicles,” *The International Journal of Robotics Research*, vol. 31, no. 11, pp. 1279–1291, 2012.
- [5] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, “Information acquisition with sensing robots,” in *IEEE International Conference on Robotics and Automation*, 2014, pp. 6447–6454.
- [6] C. Robin and S. Lacroix, “Multi-robot target detection and tracking: taxonomy and survey,” *Autonomous Robots*, vol. 40, no. 4, pp. 729–760, 2016.
- [7] J. R. Spletzer and C. J. Taylor, “Dynamic sensor planning and control for optimally tracking targets,” *The International Journal of Robotics Research*, vol. 22, no. 1, pp. 7–20, 2003.
- [8] E. W. Frew, *Observer trajectory generation for target-motion estimation using monocular vision*. Stanford University, 2003.
- [9] M. Schwager, B. J. Julian, M. Angermann, and D. Rus, “Eyes in the sky: Decentralized control for the deployment of robotic camera networks,” *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1541–1561, 2011.
- [10] A. Pierson, Z. Wang, and M. Schwager, “Intercepting rogue robots: An algorithm for capturing multiple evaders with multiple pursuers,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 530–537, 2017.
- [11] P. Tokekar, V. Isler, and A. Franchi, “Multi-target visual tracking with aerial robots,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 3067–3072.
- [12] P. Dames, P. Tokekar, and V. Kumar, “Detecting, localizing, and tracking an unknown number of moving targets using a team of mobile robots,” *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1540–1553, 2017.
- [13] E. Sless, N. Agmon, and S. Kraus, “Multi-robot adversarial patrolling: Facing coordinated attacks,” in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2014, pp. 1093–1100.
- [14] H. H. González-Banos, C.-Y. Lee, and J.-C. Latombe, “Real-time combinatorial tracking of a target moving unpredictably among obstacles,” in *Robotics and Automation, 2002. Proceedings. ICRA’02. IEEE International Conference on*, vol. 2. IEEE, 2002, pp. 1683–1690.
- [15] S. I. Roumeliotis, G. S. Sukhatme, and G. A. Bekey, “Sensor fault detection and identification in a mobile robot,” in *Intelligent Robots and Systems, 1998. Proceedings., 1998 IEEE/RSJ International Conference on*, vol. 3. IEEE, 1998, pp. 1383–1388.
- [16] V. Tzoumas, A. Jadbabaie, and G. J. Pappas, “Resilient Non-Submodular Maximization over Matroid Constraints,” *arXiv e-prints: 1804.01013*, 2018.
- [17] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, “An analysis of approximations for maximizing submodular set functions–II,” *Polyhedral combinatorics*, pp. 73–87, 1978.
- [18] B. Schlotfeldt, V. Tzoumas, D. Thakur, and G. J. Pappas, “Resilient active information gathering with mobile robots,” *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- [19] S. Martinez, J. Cortes, and F. Bullo, “Motion coordination with distributed information,” *IEEE Contr. Sys.*, vol. 27, no. 4, pp. 75–88, 2007.
- [20] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [21] C. J. Green and A. Kelly, “Toward optimal sampling in the space of paths,” in *Robotics Research*, 2010, pp. 281–292.
- [22] R. B. Myerson, *Game theory*. Harvard University Press, 2013.
- [23] D. S. Hochbaum and A. Pathria, “Analysis of the greedy approach in problems of maximum k-coverage,” *Naval Research Logistics*, vol. 45, no. 6, pp. 615–627, 1998.
- [24] M. Conforti and G. Cornuéjols, “Submodular set functions, matroids and the greedy algorithm,” *Discrete Applied Mathematics*, vol. 7, no. 3, pp. 251–274, 1984.
- [25] R. Iyer, S. Jegelka, and J. Bilmes, “Fast semidifferential-based submodular function optimization,” in *International Conference on Machine Learning*, 2013, pp. 855–863.

⁶https://youtu.be/0pGrg514_eg