

# Optimal Persistent Monitoring of Mobile Targets in One Dimension

Jonas Hall<sup>1</sup>, Sean B. Andersson<sup>1,2</sup>, Christos G. Cassandras<sup>1,3</sup>

**Abstract**—This work shows the existence of optimal control laws for persistent monitoring of mobile targets in a one-dimensional mission space and derives explicit solutions. The underlying performance metric consists of minimizing the total uncertainty accumulated over a finite mission time. We first demonstrate that the corresponding optimal control problem can be reduced to a finite-dimensional optimization problem, and then establish existence of an optimal solution. Motivated by this result, we construct a parametric reformulation for which an event-based gradient descent method is utilized with the goal of deriving locally optimal solutions. We additionally provide a more practical parameterization that has attractive properties such as simplicity, flexibility, and robustness. Both parameterizations are validated through simulation.

## I. INTRODUCTION

In persistent monitoring problems, a team of autonomous agents is tasked to monitor a given environment. The problem is closely related to coverage control [1], the key difference being that the environment is assumed to change dynamically, so that all points of interest must be revisited persistently. The problem has received a lot of attention over the last decade due to its broad applicability across a range of applications, including environmental monitoring [2], data harvesting [3], and particle tracking [4]. Persistent monitoring problems are notoriously difficult to solve because of highly non-convex and non-smooth costs and dynamics along with the computational complexity due to the problem's combinatorial nature.

In this paper we consider the persistent monitoring problem applied to a finite number of targets in a one-dimensional connected mission space, similar to the formulations in [5], [6]. One-dimensional mission spaces are common in applications that involve rivers, roads in a transportation network, or power lines, to name a few. Unlike previous work, however, in addition to an internal state describing a measure of uncertainty, we allow *the targets themselves to be mobile* in the mission space. The control objective considered is the minimization of the overall uncertainty over a finite mission time. The uncertainty dynamics considered in this work can be viewed

as a flow model with constant input flow, whereas the output flow depends on the target-agent distance. This model is commonly used in the literature [5], [7], though other dynamics have been discussed as well. For example, in [8], the authors consider internal states with arbitrary linear time-invariant dynamics, with a control applied only when an agent is monitoring the target. In other applications, such as environmental monitoring tasks, the internal state is modeled as a Gaussian random field [9]. Similarly, the literature contains various performance metrics, e.g., that of maximizing the number of observed events [2], minimizing the inter-observation time [10], or minimizing estimation errors [9].

The persistent monitoring problem is often split into a higher and lower level. The higher level consists of path planning and scheduling target visiting sequences. Lan and Schwager proposed the use of rapidly-exploring random cycles in order to find an optimal periodic trajectory [9]. Other formulations abstract the mission space to a graph, where each node represents a region of interest and the edges capture the travel times [10], [11]. The graph-based formulation was further generalized in [12], in which targets were grouped into clusters and only a single target within each cluster must be visited. Due to the combinatorial nature of the scheduling task, the use of mixed-integer optimization is natural [13], and its similarity to the Traveling Salesman Problem has motivated solving the higher level scheduler with such methods [14]. On the lower level, the problem consists of optimizing the agents' motion control, either along a given path, or in order to realize a given schedule. For example, in [5] and [15], the authors optimize agent velocities along a given path, while in [16] a configuration-based solution is sought such that each point on a given path is revisited with a constant frequency.

The contribution of this paper consists of advances in the one-dimensional persistent monitoring problem of mobile targets. In particular, we prove that the considered infinite-dimensional Optimal Control Problem (OCP) can be reduced to an optimization problem of finite dimension, and that the OCP always admits a solution. Though similar results can be found for static targets, we provide novel (to the best of the authors' knowledge) results for mobile targets. Additionally, we provide two parametric reformulations and utilize the Infinitesimal Perturbation Analysis (IPA) method [17] in order to compute gradients and solve the parametric optimization problem locally.

The remainder of this paper is structured as follows.

This work was supported in part by NSF under grants ECCS-1931600, DMS-1664644, CNS-1645681, CNS-2149511, by AFOSR under FA9550-19-1-0158, by ARPA-E under DE-AR0001282, by NIH under 1R01GM117039-01A1, and by the MathWorks.

<sup>1</sup>Division of Systems Engineering, Boston University, USA

<sup>2</sup>Division of Mechanical Engineering, Boston University, USA

<sup>3</sup>Department of Electrical and Computer Engineering  
{hallj, anderss, cg}@bu.edu

Sec. II introduces the problem in detail and establishes the OCP (4). In Sec. III, we show how the problem can be reduced to a finite-dimensional optimization problem. Two parametric formulations are provided and compared in Sec. IV. By applying the IPA calculus, we describe how perturbations of the parameters affect the cost function in Sec. V. We compare the parameterizations in numerical experiments in Sec. VI. Sec. VII concludes the paper and outlines ideas for future work.

## II. PROBLEM FORMULATION

Our problem formulation builds upon the foundation of [6], [7], extending it to include *mobile* targets. Let  $\mathcal{S} = \mathbb{R}$  be the work space in which we consider the persistent monitoring problem over a finite time interval  $[0, T]$ . We introduce a set of mobile agents  $\mathcal{A} = \{1, 2, \dots, N\}$ , and denote the position of agent  $j \in \mathcal{A}$  by  $s_j(t) \in \mathcal{S}$ . We assume that the agents follow first-order dynamics  $\dot{s}_j(t) = u_j(t)$ , where each  $u_j$  belongs to the admissible control set  $\mathcal{U} = \{[0, T] \rightarrow [-1, 1]\}$ , and the initial position  $s_j(0) \in \mathcal{S}$  is given. In addition to the agents, there is a set of mobile targets  $\mathcal{T} = \{1, 2, \dots, M\}$  with  $M > N$ , and we denote by  $\vartheta_i(t)$  and  $\dot{\vartheta}_i(t)$  the position and velocity of  $i \in \mathcal{T}$ , which are assumed to be known. The monitoring function of agent  $j$  for target  $i$  is given by

$$p_{ij}(\vartheta_i(t), s_j(t)) = \max \left\{ 0, 1 - \frac{|\vartheta_i(t) - s_j(t)|}{r_j} \right\}, \quad (1)$$

where  $r_j > 0$  is the sensing range of agent  $j$ . The joint monitoring function of target  $i$  from all agents is then

$$P(\vartheta_i(t), s(t)) = 1 - \prod_{j \in \mathcal{A}} (1 - p_{ij}(\vartheta_i(t), s_j(t))). \quad (2)$$

The question remains of what is being monitored. We associate each target with an internal state  $R_i(t) \geq 0$  describing a measure of uncertainty. This state is assumed to grow with a constant rate of  $A_i > 0$  when not monitored, and to change with a dynamic rate of  $A_i - B_i P(\vartheta_i(t), s(t))$  for a given  $B_i > A_i$  otherwise. To ensure that the uncertainty remains nonnegative, we limit the uncertainty reduction for targets in the set  $\mathcal{Z}(t) = \{i \in \mathcal{T} \mid R_i(t) = 0 \text{ and } A_i < B_i P(\vartheta_i(t), s(t))\}$ . Thus, the uncertainty dynamics  $\dot{R}_i(t)$  are captured by

$$f_{R_i}(t) = \begin{cases} 0, & i \in \mathcal{Z}(t), \\ A_i - B_i P(\vartheta_i(t), s(t)), & i \notin \mathcal{Z}(t). \end{cases} \quad (3)$$

The goal of the persistent monitoring problem is to minimize the overall uncertainty, i.e., we seek a control policy which solves the OCP

$$\underset{u \in \mathcal{U}^N}{\text{minimize}} \quad J(u) = \frac{1}{T} \int_0^T \sum_{i \in \mathcal{T}} R_i(t) dt \quad (4a)$$

$$\text{subject to} \quad \dot{s}(t) = u(t), \quad \dot{R}(t) = f_R(t). \quad (4b)$$

Note that the function (3) depends on time only through  $s(t)$ ,  $R(t)$ , and  $\vartheta(t)$ . However, to keep notation simple, we neglect the explicit dependence on the states.

*Remark 1:* The choice of the monitoring function is not restricted to (1). Any monitoring function can be chosen, as long as it *a)* has compact support, and *b)* is monotonically increasing with the distance between its arguments. Similarly, the results in this paper directly extend to other joint monitoring functions, as long as they are monotonically increasing with each of the individual monitoring functions, e.g.,  $\max_{j \in \mathcal{A}} p_{ij}(\vartheta_i(t), s_j(t))$ .

## III. OPTIMAL CONTROL CHARACTERIZATION

As in previous work [6], a standard Hamiltonian analysis can be applied to (4) to reveal that the optimal control can be expressed in a parametric form. To better understand the structure of the problem, however, here we show the existence of such a parametric reformulation directly. In order to prove the main result in Thm. 1 below, we first show that any control law can be adapted to one of the form given in (5) without increasing the cost. Applying this result repeatedly shows that any control can be replaced by one which can be partitioned into a finite number of intervals during which the control is either *a)* constant and maximal, or *b)* equal to the velocity of a target. From this we conclude that there exists a compact parameter space, with which we can ultimately show the existence of a parametrically represented optimal control. These theoretical results rely on the following assumptions for the motions of the targets.

*Assumption 1:* For all  $t \in [0, T]$ ,  $i, i_1 \neq i_2 \in \mathcal{T}$ , we assume that (i) target velocities are no larger than the maximal speed of the agents, i.e., their absolute velocities are bounded by one, and (ii) agents never sense two targets at the same time, i.e., there exists  $\Delta^{\min} > 0$  such that  $|\vartheta_{i_1}(t) - \vartheta_{i_2}(t)| \geq 2 \max_{j \in \mathcal{A}} \{r_j\} + \Delta^{\min}$ .

The first assumption is fundamental, whereas the second one is of a technical nature, without which it becomes increasingly difficult to characterize the optimal control. However, in Sec. VI we show that the considered parameterizations are adaptable to simultaneous sensing scenarios. Similarly, if a target's velocity exceeds the target control bounds, then the given parameterization with bounded controls provides a heuristic for the more general case. Our first lemma provides a characterization of the optimal control over a fixed interval.

Before we begin, let us denote by 'sgn' the sign function that maps all negative real numbers to  $-1$ , zero to zero, and all positive real numbers to  $+1$ . The proofs for all lemmas and theorems that follow are provided in [18].

*Lemma 1:* Let  $(u, s)$  be any feasible control-trajectory pair. Consider an interval  $(t_1, t_2)$  over which the agent  $j \in \mathcal{A}$  only ever senses one target  $i \in \mathcal{T}$ . Then, there exist  $\hat{t}, \hat{\hat{t}} \in [t_1, t_2]$  such that the modified control law

$$u'_j(t) = \begin{cases} u_j(t), & \text{if } t \notin (t_1, t_2), \\ \text{sgn}(\vartheta_i(t_1) - s_j(t_1)), & \text{if } t \in (t_1, \hat{t}), \\ \dot{\vartheta}_i(t), & \text{if } t \in (\hat{t}, \hat{\hat{t}}), \\ \text{sgn}(s_j(t_2) - \vartheta_i(t_2)), & \text{if } t \in (\hat{\hat{t}}, t_2), \end{cases} \quad (5)$$

satisfies  $J(u') \leq J(u)$ .

*Proof sketch:* We construct (5) such that  $s'_j$  coincides with that of  $s_j$  everywhere outside of  $(t_1, t_2)$ , and is at least as close to target  $i$  as  $s_j$  during  $(t_1, t_2)$ .

This lemma allows us to decompose an existing control into sequential phases (or modes) in which only one target is visited. We utilize this control modification for each individual interval to obtain a parameterized control law. For the next lemma we introduce the ceiling function  $\lceil \cdot \rceil$ , which maps each real number to the smallest integer greater than or equal to the number itself.

*Lemma 2:* Let  $(u, s)$  be any feasible control-trajectory pair. Then there exists an alternative control law  $u'$ , which satisfies  $J(u') \leq J(u)$  and is fully described by  $n = 5N \lceil T/\Delta^{\min} \rceil$  parameters.

*Proof sketch:* We repeatedly apply Lemma 1.

From this property, we can immediately conclude that if the OCP (4) admits a solution, then there exists a control described by a finite number of parameters. In order to show later that such a solution always exists, we now prove that the resulting parameter space is compact.

*Lemma 3:* There exists a finite-dimensional and compact parameter space  $\mathcal{Y}$  with functions

$$T : \mathcal{U}^N \rightarrow \mathcal{Y} \quad \text{and} \quad H : \mathcal{Y} \rightarrow \mathcal{U}^N, \quad (6)$$

such that  $J(H(T(u))) \leq J(u)$ . In addition, the function  $H$  can be chosen to be continuous.

*Proof sketch:* The first part is a corollary of Lemma 2. For the second part we prove sequential continuity.

The fact that the parameter space is compact leads to the existence of converging subsequences, which we utilize in the next theorem.

*Theorem 1:* Consider the OCP (4). Under Assumption 1 there exists an optimal control  $u^* \in \mathcal{U}^N$  and an optimal parameterization  $\pi^* \in \mathcal{Y}$  such that  $H(\pi^*) = u^*$ .

*Proof sketch:* The key is that the parameter space is compact, such that every sequence admits a convergent subsequence. Mapping an adequate sequence via  $H$  and taking the limit yields the result.

#### IV. PARAMETRIC REFORMULATION

The above analysis shows that the optimal control problem can be reduced to a much simpler *parametric* optimization problem. In what follows, we will introduce two different parameterizations. The first one provides a generalization of (5), and we refer to it as the *optimal parameterization*, since it contains an optimal control as established in Thm. 1. Its drawback is that it requires knowledge of the target velocities, which are difficult to estimate in practice. From a practical viewpoint, the policy (5) can be read as *stay as close to the target as possible*, which motivates the alternative *practical parameterization*. Though it loses optimality to some degree, it only requires positional feedback of the targets, and has additional advantages, such as a natural robustness towards noise, and simplicity of initialization; details are discussed later on. We now introduce both forms for a




Fig. 1. Comparison of the optimal (top) and practical (bottom) parameterizations. The variables below the intervals indicate their length and type, according to the parameterization. Dashed vertical lines represent foreseeable mode changes due to proceeding to the next interval, whereas dotted vertical lines depict observable mode changes due to sufficient reduction of the tracking error.

fixed agent  $j$ , while omitting the index for notational simplicity. For guidance, we refer the reader to Fig. 1, which illustrates both formulations.

**Optimal Parameterization.** Let us first define the set of convex combinations  $\mathcal{C} = \{\alpha \in [0, 1]^M \mid \sum_{i=1}^M \alpha_i = 1\}$ . Next, we introduce parameters that can be divided into switching points  $\psi_1, \dots, \psi_L \in \mathcal{S}$ , tracking combinations  $\alpha_1, \dots, \alpha_L \in \mathcal{C}$  and tracking durations  $\phi_1, \dots, \phi_L \in [0, T]$ , where  $L$  is some fixed integer. From this parameterization we obtain a control

$$u(t) = \begin{cases} \text{sgn}(\psi_\ell - s(\bar{t}_{\ell-1})), & \text{for } t \in [\bar{t}_{\ell-1}, \check{t}_\ell), \\ \alpha_\ell^\top \dot{v}(t), & \text{for } t \in [\check{t}_\ell, \bar{t}_\ell), \end{cases} \quad (7)$$

where  $\check{t}_\ell = \bar{t}_{\ell-1} + |\psi_\ell - s(\bar{t}_{\ell-1})|$  and  $\bar{t}_\ell = \check{t}_\ell + \phi_\ell$ , initialized with  $\bar{t}_0 = 0$ .

There are two reasons why we propose the target tracking to be encoded in the form of convex combinations. First and foremost, the feasible set of the parametric description becomes convex. The only difficulty to overcome in an optimization scheme is that of the cost function's nonconvexity and nonsmoothness. Secondly, the resulting control law has some adaptability towards simultaneous sensing scenarios, as discussed in Sec. VI. Note that if Assumption 1 is satisfied, then the optimal control established in Thm. 1 is contained in the parameterization by choosing convex combinations that only select a single target at a time.

**Practical Parameterization.** As opposed to defining switching points in the mission space and then matching a convex combination of target velocities, we now propose a control law by directly tracking a convex combination of target positions.

The parameterization is thus reduced to tracking durations  $\phi_1, \dots, \phi_L \in [0, T]$  and tracking combinations  $\alpha_1, \dots, \alpha_L \in \mathcal{C}$ . The switching times are simply given

by  $\bar{t}_\ell = \bar{t}_{\ell-1} + \phi_\ell$ , which is again initialized with  $\bar{t}_0 = 0$ . For  $t \in [\bar{t}_{\ell-1}, \bar{t}_\ell)$  we then obtain a control

$$u(t) = \min(1, \max(-1, f_\ell^{\text{PI}}(t))), \quad (8)$$

via the PI-controller

$$f_\ell^{\text{PI}}(t) = K_p e_\ell(t) + K_i \int_{\bar{t}_\ell}^{\max(t, \bar{t}_\ell)} e_\ell(\sigma) d\sigma, \quad (9)$$

with tracking error  $e_\ell(t) = \alpha_\ell^\top \vartheta(t) - s(t)$ , and respective proportional and integral feedback gains  $K_p > 0$  and  $K_i > 0$ . Since the integral part is useful for driving a small proportional error to zero, we activate this part once the tracking error becomes small enough. Thus, the starting time of the integrator part can be written as

$$\bar{t}_\ell = \inf\{t \geq \bar{t}_{\ell-1} \mid |K_p e_\ell(t)| \leq \varepsilon_{\text{tol}}\}, \quad (10)$$

for some desired switching tolerance  $\varepsilon_{\text{tol}} \in (0, 1)$ . The motivation behind this strategy is that the beginning of a new tracking interval typically leads to a large proportional feedback gain due to the switched tracking combination  $\alpha_\ell$ . In order to prevent integrator windup and consequent control saturation during the PI tracking phase, under this strategy we first track via the P-controller, switching to a PI-controller once the tracking error becomes small enough. Note that this is a fairly simple anti-windup scheme, with the main goal of keeping the control law continuous and thereby the IPA analysis as tractable as possible; in practice more standard anti-windup algorithms may be preferred.

**Comparison of the parameterizations.** The practical parameterization clearly loses optimality, the degree of which depends on the position sampling rate and the chosen feedback gains. On the other hand, the practical implementation gives rise to attractive properties. One is the ease of initialization. Assume that we want to initialize the system with a given visiting schedule of targets. Then, all we have to do is decide how long we want to track each of the targets. Trying to initialize the optimal parameterization for mobile targets becomes tedious and complex; it requires estimates of where to intercept the targets and where agents will end up after each tracking phase. Such information is rarely available in practice. Additionally, the practical implementation is naturally more robust to perturbations due to its closed loop design.

## V. OPTIMIZING THE SYSTEM PARAMETERS

The parametric controllers introduced in the preceding section define a hybrid system for each agent. We denote the feasible set of parameters as  $\Theta$ , the precise form of which depends on the chosen parameterization. The OCP (4) is now reduced to the parametric optimization problem

$$\underset{\theta \in \Theta}{\text{minimize}} \quad J(\theta),$$

which we solve offline. In order to perform a gradient descent method we need the gradient of  $J$  with respect

to  $\theta$ ; we find this using IPA [17]. It was shown in [6], [7] that the derivative of the cost function with respect to the parameters is given by

$$\frac{\partial J(\theta)}{\partial \theta} = \frac{1}{T} \sum_{k=1}^K \sum_{i=1}^M \int_{\tau_{k-1}(\theta)}^{\tau_k(\theta)} \frac{\partial R_i(\sigma)}{\partial \theta} d\sigma,$$

where  $\tau_0, \tau_1, \dots, \tau_K$  describe the event times when mode switches occur in the hybrid system, and  $K$  denotes the total number of events (note  $\tau_0 = 0$  and  $\tau_K = T$ ). In order to calculate the gradients of  $R_i$  we must consider the various mode changes. We omit the explicit calculation here and refer to the extended version [18].

**Finding a feasible search direction.** With all gradients determined via IPA, we are almost ready to perform a gradient descent method. However, in order for the tracking combination constraint to remain satisfied, the step direction with respect to the tracking constraints, denoted by  $p_\alpha$ , must be chosen appropriately, i.e., its components must add up to zero for each tracking period  $\ell$  and agent  $j$ . We do so by computing a feasible step direction via the least squares problem

$$\underset{p_\alpha}{\text{minimize}} \quad \left\| \frac{\partial J}{\partial \alpha} - p_\alpha \right\|_2^2 \quad (11a)$$

$$\text{subject to} \quad 0 \leq \alpha_{\ell ij} + p_{\alpha_{\ell ij}} \leq 1 \quad \forall \ell, i, j, \quad (11b)$$

$$\sum_{m=1}^M p_{\alpha_{\ell m j}} = 0 \quad \forall \ell, j, \quad (11c)$$

where  $p_{\alpha_{\ell ij}}$  is the  $i$ th component of the step direction for agent  $j$  in the  $\ell$ th tracking period. The global solution of this convex optimization problem yields the feasible direction closest to the gradient of  $J$ . For the remaining parameters, we directly use the IPA gradients. In order to achieve cost function descent at each iterate, we choose a step length via the Armijo backtracking technique [19]. **Challenges in finding good local solutions.** Though we are now able to optimize the system parameters over a convex set, and gradients are available for suitable initializations, the problem remains difficult to solve due to the local nature of the gradient descent procedure. For the one-dimensional persistent monitoring problem for static targets, it was demonstrated that finding the global solution of a problem with fixed target visiting sequence is tractable [6]. Fig. 2 demonstrates that this changes drastically in the mobile target case. The cost, even with respect to a single parameter, can be highly nonconvex. Moreover, the cost function is in parts extremely steep, whereas for other segments it is very flat. This combination presents a challenge for gradient descent methods and highlights the need to select appropriate step sizes in order to avoid stagnation and remain in the contraction area of a desirable solution.

Another difficulty, which is typical for event-based algorithms, is the potential lack of event excitation [20]. This occurs in the persistent monitoring problem if the agents change their modes while not sensing targets. In such cases the IPA gradients fail to estimate the true cost




Fig. 2. Depiction of the highly nonconvex cost (top) together with respective local (purple) and global (green) solutions (bottom) when fixing all but one parameter. Note that the sensing range (red) is shown around the targets as opposed to around the agents.

gradient resulting in a failure of the optimization scheme. Though beyond the scope of this paper, the methods discussed in [20] could help overcome this problem.

## VI. NUMERICAL EXPERIMENTS

**Comparison to existing methods.** We first compare the proposed methods to the related algorithm from [6] that was designed for static targets. We apply the new parameterizations to a static target experiment, the globally optimal cost of which was shown to be  $J^* = 25.07$  [6]. Fig. 3 shows the results for both the optimal and practical parameterizations. The solution quality of the practical (25.54) and the optimal (25.61) are comparable to the local solution of the previous method (25.54), which indicates that the generalized methods remain applicable to static targets as well. The fact that the practical parameterization slightly outperforms the optimal one in this specific case is due to local convergence behaviors, and other initializations may lead to different results.

**Simultaneous sensing scenarios.** The theoretical results in this paper were proven under the assumption that only one target could be within an agent’s sensing range at a time. In Fig. 4 we demonstrate the fact that describing the target tracking via convex combinations allows the proposed parameterizations to adapt to scenarios where this assumption is violated. This allows agents to place themselves between targets, which is particularly useful if targets are so close to each other that they share a *deadzone*, i.e., an area in which an agent can be positioned such that the uncertainty of both targets becomes non-increasing. In such a case, the uncertainty of both targets can be driven to zero and kept at that level, which is only possible if the agent




Fig. 3. Results for the same experiment as conducted in [6]. The top plot shows the comparison of the agent trajectories obtained from the optimal (green) and practical (purple) parameterizations. The targets are depicted in red. The second plot shows the evolution of the target uncertainties, where the line style specifies the target and the color specifies the used parameterization. The third plot shows the cost reduction for the IPA iterates.

stays within that deadzone.

**Robustness towards noise.** As previously discussed in [6], the IPA analysis provides gradients that are naturally robust towards errors in the estimation of the uncertainty growth rate  $A$ . This is due to the fact that  $A$  only implicitly affects the switching times of the uncertainty dynamics without explicitly appearing in the gradient evaluations.

In this section we want to show that the practical parameterization is additionally robust towards noisy measurements. In particular, we assume that each target’s positional measurement is given by  $\hat{\vartheta}_i^{\text{est}}(t) = \vartheta_i(t) + \kappa_1 \nu_i(t)$  for  $\nu_i(t) \sim \mathcal{N}(0, 1)$  and some scaling factor  $\kappa_1 > 0$ . We ran 1000 repetitions of a fixed persistent monitoring problem with  $M = 4, N = 2, T = 50, L = 4$ , and randomized initial visiting sequences. We repeated the same experiment for the optimal parameterization with noisy target velocity estimates  $\hat{\vartheta}_i^{\text{est}}(t) = \dot{\vartheta}_i(t) + \kappa_2 \nu_i(t)$ , where  $\mu_i(t) \sim \mathcal{N}(0, 1)$  and  $\kappa_2 > 0$ . Fig. 5 compares the cost distributions of the initializations to the resulting trajectories obtained through the respective gradient descent methods. Both methods are able to reduce the cost drastically, though the practical parameterization achieves notably better results. This is due to the fact that the practical version incorporates tracking error feedback into its control law. On the other hand, the




Fig. 4. Simultaneous sensing experiment with two targets (red solid and dotted) and one agent (blue solid). The sensing range (blue shaded) is shown on the agent.




Fig. 5. Depiction of the randomized experiment. The cost distribution of the initializations is depicted in orange (optimal) and red (practical), whereas the optimized cost distributions are shown in blue (optimal) and green (practical).

optimal version simply tries to copy a target's motions, though without validating its proximity.

## VII. CONCLUSION AND FUTURE WORK

In this paper, we showed the existence of optimal solutions for the one-dimensional persistent monitoring problem of mobile targets by reducing the infinite-dimensional OCP (4) to a parametric optimization problem. We characterized the optimal control, which motivated the design of an event-based optimization scheme. We then proposed a suboptimal parametric description of the system that proved to perform robustly in various numerical experiments.

In future work, we will consider the extension to more realistic second-order agent dynamics, extending the problem to the infinite horizon case through the use of periodic patterns, and the consideration of problem formulations in two-dimensional settings.

## REFERENCES

- [1] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.
- [2] J. Yu, S. Karaman, and D. Rus, "Persistent monitoring of events with stochastic arrivals at multiple stations," *IEEE Transactions on Robotics*, vol. 31, no. 3, pp. 521–535, 2015.
- [3] Y. Khazaeni and C. G. Cassandras, "Event-driven trajectory optimization for data harvesting in multiagent systems," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 3, pp. 1335–1348, 2017.
- [4] Z. Shen and S. B. Andersson, "Tracking nanometer-scale fluorescent particles in two dimensions with a confocal microscope," *IEEE Transactions on Control Systems Technology*, vol. 19, no. 5, pp. 1269–1278, 2010.
- [5] S. L. Smith, M. Schwager, and D. Rus, "Persistent monitoring of changing environments using a robot with limited range sensing," in *2011 IEEE International Conference on Robotics and Automation*, pp. 5448–5455, IEEE, 2011.
- [6] N. Zhou, X. Yu, S. B. Andersson, and C. G. Cassandras, "Optimal event-driven multiagent persistent monitoring of a finite set of data sources," *IEEE Transactions on Automatic Control*, vol. 63, no. 12, pp. 4204–4217, 2018.
- [7] C. G. Cassandras, X. Lin, and X. Ding, "An optimal control approach to the multi-agent persistent monitoring problem," *IEEE Transactions on Automatic Control*, vol. 58, no. 4, pp. 947–961, 2012.
- [8] X. Yu, S. B. Andersson, N. Zhou, and C. G. Cassandras, "Scheduling multiple agents in a persistent monitoring task using reachability analysis," *IEEE Transactions on Automatic Control*, vol. 65, no. 4, pp. 1499–1513, 2019.
- [9] X. Lan and M. Schwager, "Planning periodic persistent monitoring trajectories for sensing robots in gaussian random fields," in *2013 IEEE International Conference on Robotics and Automation*, pp. 2415–2420, IEEE, 2013.
- [10] S. Alamdari, E. Fata, and S. L. Smith, "Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 138–154, 2014.
- [11] S. Welikala and C. G. Cassandras, "Event-driven receding horizon control for distributed persistent monitoring in network systems," *Automatica*, vol. 127, p. 109519, 2021.
- [12] S. K. K. Hari, S. Rathinam, S. Darbha, K. Kalyanam, S. G. Manyam, and D. Casbeer, "The generalized persistent monitoring problem," in *2019 American Control Conference (ACC)*, pp. 2783–2788, IEEE, 2019.
- [13] P. Maini, K. Yu, P. Sujit, and P. Tokekar, "Persistent monitoring with refueling on a terrain using a team of aerial and ground robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8493–8498, IEEE, 2018.
- [14] M. Ostertag, N. Atanasov, and T. Rosing, "Trajectory planning and optimization for minimizing uncertainty in persistent monitoring applications," *Journal of Intelligent & Robotic Systems*, vol. 106, no. 1, pp. 1–19, 2022.
- [15] C. Song, L. Liu, G. Feng, and S. Xu, "Optimal control for multi-agent persistent monitoring," *Automatica*, vol. 50, no. 6, pp. 1663–1668, 2014.
- [16] M. Boldrer, F. Pasqualetti, L. Palopoli, and D. Fontanelli, "Multi-agent persistent monitoring via time-inverted kuramoto dynamics," *IEEE Control Systems Letters*, 2022.
- [17] C. G. Cassandras, Y. Wardi, C. G. Panayiotou, and C. Yao, "Perturbation analysis and optimization of stochastic hybrid systems," *European Journal of Control*, vol. 16, no. 6, pp. 642–661, 2010.
- [18] J. Hall, S. B. Andersson, and C. G. Cassandras, "Optimal persistent monitoring of mobile targets in one dimension," 2022. pre-print arXiv:2210.01294.
- [19] L. Armijo, "Minimization of functions having lipschitz continuous first partial derivatives," *Pacific Journal of mathematics*, vol. 16, no. 1, pp. 1–3, 1966.
- [20] Y. Khazaeni and C. G. Cassandras, "Event excitation for event-driven control and optimization of multi-agent systems," in *2016 13th International Workshop on Discrete Event Systems (WODES)*, pp. 197–202, IEEE, 2016.