

Multi-Agent Persistent Monitoring of Targets with Uncertain States

Samuel C. Pinto, *Student Member, IEEE*, Sean B. Andersson, *Senior Member, IEEE*, Julien M. Hendrickx, *Member, IEEE* and Christos G. Cassandras, *Fellow, IEEE*

Abstract—We address the problem of persistent monitoring, where a finite set of mobile agents has to persistently visit a finite set of targets. Each of these targets has an internal state that evolves with linear stochastic dynamics. The agents can observe these states, and the observation quality is a function of the distance between the agent and a given target. The goal is then to minimize the mean squared estimation error of these target states. We approach the problem from an infinite horizon perspective, where we prove that, under some natural assumptions, the covariance matrix of each target converges to a limit cycle. The goal, therefore, becomes to minimize the steady state uncertainty. Assuming that the trajectory is parameterized, we provide tools for computing the steady state cost gradient. We show that, in one-dimensional (1D) environments with bounded control and non-overlapping targets, when an optimal control exists it can be represented using a finite number of parameters. We also propose an efficient parameterization of the agent trajectories for multidimensional settings using Fourier curves. Simulation results show the efficacy of the proposed technique in 1D, 2D and 3D scenarios.

I. INTRODUCTION

We consider the problem of multi-agent persistent monitoring. This problem consists of using a finite set of cooperative agents to monitor a finite set of targets, more numerous than agents, which have internal states that evolve over time with dynamics subject to uncertainty. Therefore, as time goes to infinity, in order to keep the uncertainty under control, the targets need to be visited not only a finite number of times, but persistently. The goal is to minimize the long-term uncertainty by designing movement policies that produce the best estimate possible of the target states. This paradigm finds applications across a wide range of domains, such as trajectory planning of underwater vehicles to measure ocean temperature [1]–[3], surveillance in smart cities [4] and tracking of multiple microparticles by an optical microscope [5].

This problem is closely related to the Multi Traveling Salesman Problem (MTSP) [6] and Multi-Vehicle Routing Problem (MVRP) [7], where, given a set of targets (possibly constrained to a graph-based structure), the goal is to find a cycle in which the agents efficiently visit all the targets in order

The authors are with the Dept. of Mechanical Engineering (Pinto and Andersson), the Division of Systems Engineering (Andersson, Cassandras) and the Dept. of Electrical and Computer Engineering (Cassandras), Boston University, Boston, MA, 02215, USA. (e-mail: {samcerq,sanderss,cgc}@bu.edu)

Julien M. Hendrickx is with the ICTEAM Institute, UCLouvain, Louvain-la-Neuve 1348, Belgium. (e-mail: julien.hendrickx@uclouvain.be)

This work was supported in part by NSF under grants ECCS-1931600, DMS-1664644, CNS-1645681, and CMMI-1562031, by ARPA-E's NEXTCAR program under grant DE-AR0000796, by AFOSR under grant FA9550-19-1-0158, and by the MathWorks. The work of J. Hendrickx was supported by the RevealFlight Concerted Research Action (ARC) of the Federation Wallonie-Bruxelles, by the Incentive Grant for Scientific Research (MIS) Learning from Pairwise Data of the F.R.S.-FNRS, and by a WBI World Excellence Fellowship.

to minimize the traveled distance or total travel time. These problems are proved to be computationally intractable (NP-hard) and most of the scalable solutions to these problems rely either on local optimization or heuristics [6]–[8]. The major differences between the MTSP and MVRP and the problem we are dealing with in this paper are that the optimization goal we consider is to minimize the uncertainty rather than distance or time between two consecutive observations of a given target and that in persistent monitoring, agents can also cooperate to observe the same node. The present work is also closely related to the sensor allocation problem [9], where a set of sensors can observe a set of targets, but due to various constraints not all the targets can be observed at the same time and, therefore, some of the sensors have to switch among the targets they observe. The sensor allocation problem, however, assumes that the sensors are fixed and therefore does not incorporate the effect of the agent movement (i.e. the mobile sensors) in the formulation.

In the realm of persistent monitoring, significant previous work has been done. Some of these works aim at planning agent trajectories in order to maximize the chance of detecting events happening randomly at specific locations in the mission space [8], [10]. Our work, however, considers a version of Persistent Monitoring which was introduced in [11], where the goal was to sense and estimate a process evolving at a fixed location in space rather than detecting events. Nonetheless, [11] does not discuss efficient solutions to the problem it proposed. [2], [12] introduced an optimal control approach that relied on a solution of the two-point boundary value problem resulting from a Hamiltonian analysis. The solution of the two-point value problem is numerically challenging and computationally expensive. In [13] the persistent monitoring problem is formulated using temporal logic to encode target visiting constraints rather than solving an optimization problem. Additionally, [1] introduces a variant of the Rapid-Exploring Random Tree (RRT) algorithm designed for cyclic Persistent Monitoring in discrete time. However, previous work [14] has shown that this algorithm can fail to converge within a reasonable time even in setups with few targets and a single agent, especially when the process can only be sensed from a finite range from the targets.

The present paper also builds up from significant previous work by the authors [15]–[17], where the problem of persistent monitoring was modelled using an uncertainty metric for each of the targets that either grew linearly with time when the agent was not observed or decreased linearly when an agent visited it. A common feature among these previous works and the present paper is the focus on scalable solutions with respect to the number of agents, targets and time horizon. Therefore, instead of looking for globally optimal visiting schedules, we

use a local optimization scheme (gradient descent) even though the obtained solution is not guaranteed to be globally optimal. One big challenge in order to use a gradient descent approach is to efficiently compute the gradients of the cost with respect to the parameters that define the trajectory.

The current work, unlike some previous work by the authors, considers each target as having an internal state that evolves with linear stochastic dynamics that can be observed with a linear observation model, **which is a much more natural model for performing estimation. In this case, the uncertainty metric does not need to be externally modeled, but rather arises naturally from the estimation process.** The signal to noise ratio of the observation is a function of the distance between the agent and the target. In this setting, the optimal estimator can be proven to be a Kalman-Bucy filter and the mean estimation error is directly related to the covariance matrix of this filter. The main contribution of this paper is to provide tools to efficiently represent and optimize the schedules for agents visiting targets. If we consider finite horizon schedules, as time grows to infinity, the number of parameters to represent a trajectory also tends to grow infinitely large. We, however, restrict ourselves to a periodic trajectory and approach the problem from an infinite horizon perspective. We show that under some very natural assumptions the estimation error converges to a limit cycle and we provide tools for optimizing one period of the limit cycle trajectory, which usually is represented by only a small number of parameters. We note that most of the previous work in the field [2], [11], [12] focuses on transient versions of the problem, and algorithms tend to scale with the time horizon. This present approach, however, by explicitly considering the *infinite-horizon* problem is able to efficiently optimize the long-term trajectory in a computationally efficient manner. We highlight that, since the goal is to monitor the process persistently, being able to efficiently handle long time horizons is essential. To the best of the authors knowledge, the only other work that addresses the multi-agent infinite horizon version of the problem is [1], but its unguided probabilistic exploration of the environment algorithm makes it unreliable even in simple setups [14].

Although the analysis introduced in this paper is independent of the particular parameterization chosen for the trajectory, we discuss two parameterizations that are particularly interesting. When the targets and agents are constrained to lie in a one-dimensional environment, we show that, under some assumptions, an optimal control can always be represented by a trajectory in which the agent is either moving with full speed or dwelling at a fixed position. This allows optimal trajectories to be described as a finite sequence of movement times and dwelling times, yielding a parameterization. On the other hand, when the agents and targets operate in a higher dimensional space (e.g. 2D and 3D), we cannot immediately extend such properties of an optimal control. We then parameterize the trajectories using Fourier curves, where the movement of an agent in each of the coordinates is described by a truncated Fourier series. Fourier curves are interesting because they are able to describe very general smooth movement policies with only a very small number of coefficients.

Recalling the goal of performing local optimization using a

gradient descent scheme, it is particularly important to provide good initial solutions for the optimization. We thus connect the persistent monitoring problem with the MTSP and use a heuristic solution to the MTSP as a basis for the initial trajectory of the agents in the optimization scheme. We benefit from the fact that efficient heuristic solutions of the MTSP are well studied in the scientific literature and that they always provide an initial trajectory where all the targets are visited. This is a very important feature for persistent monitoring, since it prevents the uncertainty of each target from becoming infinitely large.

Preliminary results of this work have appeared in previous publications. In [18], the target internal state dynamics and observation models, as well as the 1D transient analysis were introduced. The computation of steady state gradients and infinite horizon analysis was first introduced in [19]. The extension to multi dimensional environments using Fourier curves was initially presented in [20]. However, the approach described in our previous works was heavily dependent on the specific parameterization and in the present work we formulate the problem in a general framework that does not rely on the specific parameterization. Moreover, in Sec. III we provide a proof that guarantees the convergence and uniqueness of the steady state covariance matrices and also we show the soundness of our method to compute the steady state gradients. On top of that, we provide a stronger claim than we did in [18] about an optimal parameterization of 1D trajectories. Previously, we were only able to show that it can be parameterized, but now we provide an explicit bound on the number of parameters. Moreover, in Sec. V we include simulation results that significantly add to the results of our previous work [20].

The rest of this paper is organized as follows. Section II describes the models used for the agents and the target internal states, along with a formulation of the optimal joint control and estimation problem. Section III presents results on the convergence of the covariance matrix and the optimization procedure is given for the periodic, infinite horizon case. Section IV introduces the 1D parameterization, along with its properties, optimization initialization and some simulation results. In Sec. V, some features of the previous section are extended to higher dimensions using Fourier curves and 1D, 2D and 3D results are presented. Finally Sec. VI gives a conclusion and shares ideas for future works.

II. PROBLEM FORMULATION

Consider an environment with a set of M points of interest (targets) at fixed positions $x_i \in \mathbb{R}^P$, $i = 1, \dots, M$. Each of these targets has an internal state $\phi_i \in \mathbb{R}^{L_i}$ that needs to be monitored and that evolves according to linear time-invariant stochastic dynamics:

$$\dot{\phi}_i(t) = A_i \phi_i(t) + w_i(t), \quad (1)$$

where $w_i(t)$ is a white noise process distributed according to $w_i(t) \sim \mathcal{N}(0, Q_i)$, $i = 1, \dots, M$, and $w_i(t)$ and $w_j(t)$ are statistically independent if $i \neq j$. **This additive noise imposes an increase over time of the uncertainty in the estimate of the state $\phi_i(t)$.**

Suppose that there is a collection of N mobile agents at positions $s_i(t) \in \mathbb{R}^P$ that can **freely move over the mission space** with the following kinematic model:

$$\dot{s}_j(t) = u_j(t), \quad u_j(t) \in \mathcal{U}, \quad j = 1, \dots, N, \quad (2)$$

where u_j is an input, and \mathcal{U} is the set of admissible inputs. Even though we assume, for the sake of simplicity, first order dynamics and that only the speed may be bounded, the results in this paper could be extended to more complex dynamics and constraints. For example, [21] explored similar results in a simplified version of the persistent monitoring problem, considering double integrator agent dynamics with constraints both on the speed and the acceleration.

Each of these agents is equipped with sensors that can observe the targets, and thus reduce the estimator uncertainty. The observations are given according to the following model:

$$z_{i,j}(t) = \gamma_{i,j}(s_j(t)) H_i \phi_i(t) + v_{i,j}(t), \quad (3)$$

where $v_{i,j}(t)$ is a white noise process distributed according to $v_{i,j}(t) \sim \mathcal{N}(0, R_i)$ with $v_{i,j}(t)$ independent of $v_{k,l}$ if $i \neq k$ or $j \neq l$, and $\gamma_{i,j} : \mathbb{R}^N \mapsto \mathbb{R}$ is a function that captures the interdependence of measurement quality and agent position. The intuition behind this function is that the instantaneous signal to noise ratio (SNR) can be computed as:

$$\frac{E[\|z_{i,j}(t) - v_{i,j}(t)\|^2]}{E[\|v_{i,j}(t)\|^2]} = \gamma_{i,j}^2(s_j(t)) \frac{E[\phi_i^T H_i^T H_i \phi_i]}{\text{tr}(R_i)}, \quad (4)$$

where $\text{tr}(\cdot)$ is the trace of the matrix. Notice that the term $E[\phi_i^T H_i^T H_i \phi_i] (\text{tr}(R_i))^{-1}$ is a deterministic scalar that does not depend on the relative position between the target and the agent. Therefore, the function $\gamma_{i,j}$ captures entirely how the agent positions affect the quality of the measurement. It is worth noting that in most of the applications of mobile agents to sensing there is a limited sensing range or the quality of the measurement gets worse as the agent moves farther away from the target. The general model of $\gamma_{i,j}$ is capable of capturing both the finite range and the dependence between measurement quality and relative position of the target from the agent.

In this paper we approach the problem of persistent monitoring trajectory planning from a centralized perspective. However, we point out that state estimation can still be performed distributively online using any distributed version of Kalman filters (see e.g. [11], [22]), since this filter's equations only depend on the data acquired online and the agent position.

At a given instant, the combined observations from all the agents of a single target can be grouped in a vector $\tilde{z}(t)$ as:

$$\tilde{z}_i(t) = [z'_{i,1} \ \dots \ z'_{i,N}]' = \tilde{H}_i(s_1, \dots, s_n) \phi_i(t) + \tilde{v}_i(t), \quad (5)$$

where

$$\tilde{H}_i = [\gamma_{i,1}(s_1 - x_i) H'_1 \ \dots \ \gamma_{i,N}(s_N - x_i) H'_N]', \quad (6)$$

$$\tilde{v}_i(t) = [v'_{i,1}(t) \ \dots \ v'_{i,N}(t)]', \quad (7)$$

and, since $v_{i,j}(t)$ is independent of $v_{i,k}(t)$ if $k \neq j$,

$$E[\tilde{v}_i'(t) \tilde{v}_i(t)] = \tilde{R}_i = \begin{bmatrix} R_i & 0 & \dots & 0 \\ 0 & R_i & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & R_i \end{bmatrix}. \quad (8)$$

The overall goal is to obtain estimators $\hat{\phi}_i(t, z(t))$ and open-loop control inputs $u_j(t)$ to minimize the following cost function:

$$J = \frac{1}{t_f} \int_0^{t_f} \left(\sum_{i=1}^M E[e_i'(\zeta) e_i(\zeta)] + \beta \sum_{j=1}^N u_j'(\zeta) u_j(\zeta) \right) d\zeta, \quad (9)$$

where $e_i(t) = \hat{\phi}_i(t) - \phi_i(t)$ and t_f is the time horizon. This cost function represents a weighted sum of the mean squared estimation error and the control effort; thus, the weighting factor β is responsible for balancing the importance of these two optimization goals.

The models in (6) and (7) define a linear time-varying stochastic system. Based on a similar statement from [2], we have the following proposition:

Proposition 1. *The optimal unbiased estimator $\hat{\phi}_i$ for the the cost function (9), dynamics (1), and observation model (3), is the Kalman-Bucy filter, given by:*

$$\dot{\hat{\phi}}_i(t) = A_i \hat{\phi}_i(t) + \Omega(t) \tilde{H}_i'(t) \tilde{R}_i^{-1} (\tilde{z}_i(t) - \tilde{H}_i(t) \hat{\phi}_i(t)), \quad (10a)$$

$$\dot{\Omega}_i(t) = A_i \Omega_i(t) + \Omega_i(t) A_i' + Q_i - \Omega_i(t) \tilde{H}_i' \tilde{R}_i^{-1} \tilde{H}_i \Omega_i(t), \quad (10b)$$

where $\Omega_i(t)$ is the covariance matrix of the estimator.

Although our proof of this proposition is very similar to the well known Kalman-Bucy derivation in [23], it is given in Appendix A in order to highlight that the joint design of trajectories and estimator does not affect the optimality of this filter. Using (6) and (7), we can rewrite (10b) as:

$$\dot{\Omega}_i(t) = A_i \Omega_i(t) + \Omega_i(t) A_i' + Q_i - \Omega_i(t) G_i \Omega_i(t) \sum_{j=1}^N \gamma_{i,j}^2(t), \quad (11)$$

where $G_i = H_i' R_i^{-1} H_i$ and $\gamma_{i,j}(t) = \gamma_{i,j}(s_j(t) - x_i)$. Using the fact that

$$E[e_i'(t) e_i(t)] = \text{tr}(E[e_i(t) e_i'(t)]) = \text{tr}(\Omega_i(t)),$$

we can rewrite the cost function in (9) as

$$J = \frac{1}{t_f} \int_0^{t_f} \left(\sum_{i=1}^M \text{tr}(\Omega_i(\zeta)) + \beta \sum_{j=1}^N u_j'(\zeta) u_j(\zeta) \right) d\zeta. \quad (12)$$

The goal is then to minimize the cost (12) subject to the dynamics in (11) and (2). In other words, we aim to design a trajectory, with constrained controls $u_j \in \mathcal{U}$, and estimation error linked to the trajectory through the dynamics of the covariance matrix of the Kalman-Bucy Filter that minimizes a weighted sum of the total control effort and the mean estimation error.

III. OPTIMIZATION OF PARAMETERIZED TRAJECTORIES

A. Finite Horizon Trajectory Optimization

Even though we focus on the optimization of infinite horizon trajectories, we briefly review the procedure for optimizing trajectories with a finite time horizon in order to later extend to the infinite horizon setting. In this section, we establish a general formulation, where we assume that the agent trajectories can be fully defined by a finite set of parameters. In the following sections we approach specific settings that show that parameterizations tend to naturally fit the persistent monitoring problem. Our overall goal is to compute locally optimal solutions with respect to these parameters using gradient descent. Therefore, we initially discuss how to compute the gradients for the finite horizon version of the problem. We define the set of parameters that fully describe the trajectory for $t \in [0, t_f]$ as $\Theta = \{\theta_1, \dots, \theta_D\}$.

Recalling the expression for the cost (12), we can compute the partial derivative with respect to one of the parameters of the trajectory θ_d as:

$$\frac{\partial J}{\partial \theta_d} = \frac{1}{t_f} \int_0^{t_f} \left(\sum_{i=1}^M \text{tr} \left(\frac{\partial \Omega_i}{\partial \theta_d}(\zeta) \right) + \beta \sum_{j=1}^N \frac{\partial \|u_j\|^2}{\partial \theta_d}(\zeta) \right) d\zeta. \quad (13)$$

As a side note, we note that we can, under the Leibniz rule, interchange the order of integration and differentiation in the cost (12), leading to (13). The function $\Omega_i(t)$ is continuous with time, since it is the solution of an ODE. Moreover, $\frac{\partial \Omega_i}{\partial \theta_d}$ varies continuously with θ_d , as long as some regularity conditions are verified (please see Appendix B for more details).

Given the dynamics of the covariance matrix in (11), $\frac{\partial \Omega_i}{\partial \theta_d}$ is the solution of the following ODE:

$$\begin{aligned} \frac{d}{dt} \left(\frac{\partial \Omega_i(t)}{\partial \theta_d} \right) &= A_i \frac{\partial \Omega_i}{\partial \theta_d}(t) + \frac{\partial \Omega_i(t)}{\partial \theta_d} A'_i + Q_i \\ &- \left(\frac{\partial \Omega_i(t)}{\partial \theta_d} G_i \Omega_i(t) + \Omega_i(t) G_i \frac{\partial \Omega_i(t)}{\partial \theta_d} \right) \sum_{j=1}^N \gamma_{i,j}^2(t) \\ &- \Omega_i(t) G_i \Omega_i(t) \sum_{j=1}^N \frac{\partial \gamma_{i,j}^2}{\partial \theta_d}(t), \quad (14) \end{aligned}$$

with initial conditions $\frac{\partial \Omega_i}{\partial \theta_d}(0) = 0$. Also, we know that

$$\frac{\partial \gamma_{i,j}^2(t)}{\partial \theta} = \sum_{p=1}^P \frac{\partial \gamma_{i,j}^2(t)}{\partial s_j^{e_p}} \frac{\partial s_j^{e_p}(t)}{\partial \theta_d}, \quad (15)$$

where e_p , $p = 1, \dots, P$ is the p -th coordinate of the space where the agents move in and $\frac{\partial \gamma_{i,j}^2(t)}{\partial s_j^{e_p}}$ depends only on the function $\gamma_{i,j}^2$. The only terms that we have not yet given a procedure to compute are $\frac{\partial(u'_j u_j)}{\partial \theta_d}(t)$ and $\frac{\partial s_j^{e_p}}{\partial \theta_d}(t)$. The computation of both of these terms is related to the specific parameterization chosen and details of their computation are discussed in Secs. IV and V. Note that we use the partial derivatives of the covariance matrices in (13) in order to compute the gradient of the cost J . The complete procedure to compute the transient problem gradients is given in Alg. 1.

Algorithm 1 Transient Gradient Computation

```

1: procedure COMPUTETRANSIENTGRADIENT
2:   Input:  $\Theta$ 
3:   Compute  $s_1(t), \dots, s_N(t)$  from the parameterization
4:   for every  $\theta$  in  $\Theta$  do
5:     Compute  $\frac{\partial}{\partial \theta} \int_0^{t_f} \sum_{j=1}^N u'_j(\zeta) u_j(\zeta) d\zeta$  according to
       the parameterization
6:     Compute  $\frac{\partial s_j(t)}{\partial \theta}$  according to the parameterization
7:     for  $i$  ranging from 1 to  $M$  do
8:       Compute  $\frac{\partial \Omega_i(t)}{\partial \theta}$  by solving ODE (14)
9:     Compute  $\frac{\partial J}{\partial \theta}$  using (13).
10:   Output:  $\nabla J$ 

```

B. Steady State Persistent Monitoring

For a persistent monitoring task to be successful, it is necessary that targets are visited infinitely often as time goes to infinity, because otherwise their uncertainty can become unbounded. Periodicity naturally fits into the persistent monitoring paradigm, since targets need to be visited infinitely often and, although a periodic structure of the solution is not necessarily optimal, earlier work [24] shows that, in a similar setting, periodic policies can approximate arbitrarily well the optimal cost. Moreover, our own simulation results in the transient case show that the trajectories tend to converge to oscillatory behavior [18]. On top of that, periodicity provides an upper bound to the inter-visit time. Moreover, if periodicity is assumed, the infinite horizon trajectory is fully defined by the trajectory of a single period. This often leads to needing only a very small number of parameters to describe the infinite horizon trajectory and, as a consequence, only a small number of parameters have to be optimized in order to generate efficient trajectories. With that in mind, in this section we explore the properties of periodic solutions to the persistent monitoring problem when the system fulfills the following very natural assumptions.

Assumption 1. *The pair (A_i, H_i) is detectable, for every $i \in \{1, \dots, M\}$.*

Assumption 2. *Q_i and the initial covariance matrix $\Omega_i(0)$ are positive definite, for every $i \in \{1, \dots, M\}$.*

The intuition behind the first assumption is that it ensures that sensing can guarantee that the uncertainty of each target can be bounded even for long horizons, if targets are observed often enough. The second one ensures that the covariance matrix will always be positive definite, a fact that will be used to prove Prop. 3. The results in this paper would likely still hold if Assumption 2 was relaxed, even though the proof of Prop. 3 could become more complex. Under these assumptions, first we explore conditions under which the convergence of the covariance matrix is achieved. For the sake of notational conciseness, we define

$$\eta_i(t) = \sum_{j=1}^N \gamma_{i,j}^2(t), \quad (16)$$

which represents the instantaneous power level of the sensed signal, combining all the agents' observations of the same target i **under a centralized communication scheme**. Using a procedure similar to the one used in the proof of Lemma 9 in [9], we establish the following proposition:

Proposition 2. *If $\eta_i(t)$ is T -periodic and $\eta_i(t) > 0$ for some non-degenerate interval $[a, b] \in [0, T]$, then, under Assumption 1, there exists a unique non-negative stabilizing T -periodic solution to (11).*

Proof. According to [25, p. 130], a pair $(A_i, \eta_i(t)H_i)$ of a periodic system is detectable if and only if for every eigenpair (x, λ) with $x \neq 0$,

$$A_i x = \lambda x \implies \exists [a, b] \in [0, T] \text{ s.t. } \eta_i(t) e^{\lambda t} H_i x \neq 0, \quad (17)$$

$\forall t \in [a, b]$ and $[a, b]$ is non-degenerate. Notice that, due to Assumption 1, for any eigenvector x of A_i , $H_i x \neq 0$. Therefore, when $\eta_i(t) > 0$ (i.e. any $t \in [a, b]$), $\eta_i(t) e^{\lambda t} H_i x \neq 0$, which implies that $(A_i, \eta_i(t)H_i)$ is detectable. Therefore, the corollary to Theorem 3 in [26, p. 95] shows that there exists a non negative T -periodic solution to (11), $\bar{\Omega}_i(t)$, and

$$\lim_{t \rightarrow \infty} (\Omega_i(t) - \bar{\Omega}_i(t)) = 0$$

for any solution $\Omega_i(t)$ and initial condition $\Omega_i(0) \succ 0$. \square

Prop. 2 implies that, if $\eta_i(t)$ is periodic, given any initial covariance matrix $\Omega_i(0)$, the estimation covariance for target i converges to a T -periodic matrix $\bar{\Omega}_i(t)$, as long as target i is visited for some non-zero amount of time in the periodic trajectory. Therefore,

$$\forall \delta > 0, \exists t_0 \text{ s.t. } |\tilde{\Omega}_i(t) - \Omega_i(t)| \leq \delta, \quad \forall t \geq t_0,$$

which implies that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t |\text{tr}(\tilde{\Omega}_i(t') - \Omega_i(t'))| dt' \leq \delta. \quad (18)$$

This discussion entails that, if we run a periodic trajectory for long enough, the mean estimation error will become arbitrarily close to the mean steady state estimation error. Therefore, if we plan only (one period of) the steady state trajectory, the actual estimation error will be arbitrarily close to that of the planned trajectory as time goes to infinity. Even though Prop. 2 states that the solution of the periodic Riccati equation is globally attractive, it does not provide any convergence rate for its numerical computation. However, the problem of computing numerical solutions to this equation has been studied in other works and we refer the reader to [27] for a good review and discussion of these methods.

Similarly as in the transient case, we intend to optimize the trajectory of the agents using gradient descent. However, the computation of the steady state gradients of the covariance matrix is more challenging than the transient case discussed in Subsec. III-A. In the sequel, we provide the procedure to compute these gradients when they exist.

C. Steady State Gradients

Assuming that the trajectory is periodic and all the targets are visited, we introduce the change of variable $q = t/T$, where T is the period of the trajectory. The steady state cost \bar{J} can be rewritten as:

$$\bar{J} = \int_0^1 \left(\sum_{i=1}^M \text{tr}(\bar{\Omega}_i(q)) + \beta \sum_{j=1}^N \bar{u}'_j(q) \bar{u}_j(q) \right) dq, \quad (19)$$

where $\bar{u}(q) = u(qT)$. Similar to (13), we know that, given some parameter $\theta_d \in \Theta$:

$$\frac{\partial \bar{J}}{\partial \theta_d} = \int_0^1 \left(\sum_{i=1}^M \text{tr} \left(\frac{\partial \bar{\Omega}_i(q)}{\partial \theta} \right) + \beta \sum_{j=1}^N \frac{\partial (\bar{u}'_j \bar{u}_j)(q)}{\partial \theta_d} \right) dq. \quad (20)$$

$\bar{\Omega}_i(q)$, when it exists, is defined by the following dynamics

$$\dot{\bar{\Omega}}_i(q) = \frac{d\bar{\Omega}_i(q)}{dq} = T(A\bar{\Omega}_i(q) + \bar{\Omega}_i(q)A' + Q - \eta_i(q)\bar{\Omega}_i(q)G_i\bar{\Omega}_i(q)), \quad (21)$$

along with the periodicity condition $\bar{\Omega}_i(0) = \bar{\Omega}_i(1)$. As a side note, we highlight that in this section we abuse the notation and use the dot over a function to indicate the derivative with respect to the normalized time q (and not with respect to t , as we did in the previous sections). Now, suppose that the gradient of $\bar{\Omega}_i(q)$ with respect to a parameter θ_d exists. Then, this gradient is the solution of the following differential equation (note that the period may be a function of the parameters or a parameter itself):

$$\begin{aligned} \dot{\Sigma}(q) - T \left(A\Sigma(q) + \Sigma(q)A' - \eta_i(q)\bar{\Omega}_i(q)G_i\Sigma(q) \right. \\ \left. - \eta_i(q)\Sigma(q)G_i\bar{\Omega}_i(q) \right) = T \frac{\partial \eta_i(q)}{\partial \theta_d} \bar{\Omega}_i(q)G_i\bar{\Omega}_i(q) - \frac{\partial T}{\partial \theta_d} \frac{\dot{\bar{\Omega}}_i}{T}, \end{aligned} \quad (22)$$

with periodicity conditions $\Sigma(0) = \Sigma(1)$. Note that (22) is a periodic Lyapunov equation [28]. In order to computation of $\Sigma(q)$, we define the following auxiliary problems:

$$\dot{\Sigma}_H - T(A - \eta_i \bar{\Omega}_i G_i) \Sigma_H = 0, \quad \Sigma_H(0) = I, \quad (23)$$

$$\begin{aligned} \dot{\Sigma}_{ZI} - T(A - \eta_i \bar{\Omega}_i G_i) \Sigma_{ZI} - T\Sigma'_{ZI} (A - \eta_i \bar{\Omega}_i G_i)' \\ = T \frac{\partial \eta_i}{\partial \theta} \bar{\Omega}_i G_i \bar{\Omega}_i - \frac{\partial T}{\partial \theta_d} \frac{\dot{\bar{\Omega}}_i}{T}, \quad \Sigma_{ZI}(0) = 0, \end{aligned} \quad (24)$$

where the time dependence of $\eta_i(q)$, $\Omega_i(q)$, $\Sigma_{ZI}(q)$ and $\Sigma_H(q)$ was omitted for conciseness. Then, in the following Proposition we exploit these auxiliary problems for computing $\Sigma(q)$ and explores the relationship between $\Sigma(q)$ and $\partial \Omega_i(q)/\partial \theta_d$.

Proposition 3. *Suppose Σ_H is a solution of (23), Σ_{ZI} is a solution of (24), Assumptions 1 and 2 hold, and that target i is observed at least once in the period T . Then, the equation*

$$\Lambda = \Sigma_H(1)\Lambda\Sigma'_H(1) + \Sigma_{ZI}(1) \quad (25)$$

has a unique solution Λ . Additionally, when $\frac{\partial \bar{\Omega}_i(q)}{\partial \theta_d}$ exists,

$$\frac{\partial \bar{\Omega}_i(q)}{\partial \theta_d} = \Sigma(q) = \Sigma'_H(q)\Lambda\Sigma_H(q) + \Sigma_{ZI}(q). \quad (26)$$

Proof. Suppose Λ and $\tilde{\Lambda}$ are solutions of (23), then

$$\Lambda - \tilde{\Lambda} = \Sigma_H(1) (\Lambda - \tilde{\Lambda}) \Sigma'_H(1) \quad (27)$$

which is equivalent to

$$\text{vec}(\Lambda - \tilde{\Lambda}) = (\Sigma_H(1) \otimes \Sigma_H(1)) \text{vec}(\Lambda - \tilde{\Lambda}), \quad (28)$$

where $\text{vec}(\cdot)$ is the operator that performs the matrix vectorization and \otimes represents the matrix Kronecker product. Notice that $\Lambda = \tilde{\Lambda}$ is a solution of (28). This solution is the unique solution if and only if 1 is not an eigenvalue of $\Sigma_H(1) \otimes \Sigma_H(1)$. On the other hand, the eigenvalues of $\Sigma_H(1) \otimes \Sigma_H(1)$ are all in the form $\mu_1 \mu_2$, where μ_1 and μ_2 are distinct eigenvalues of $\Sigma_H(1)$ [29].

In the following we show that all the eigenvalues of $\Sigma_H(1)$ have absolute value lower than one. For that, first notice that since Q is positive definite, $\bar{\Omega}_i$ is also positive definite and hence, invertible. Define

$$\mathcal{W} = \bar{\Omega}_i^{-1},$$

and, since $\dot{\mathcal{W}} = -\bar{\Omega}_i^{-1} \dot{\bar{\Omega}}_i \bar{\Omega}_i^{-1} = -\mathcal{W} \dot{\bar{\Omega}}_i \mathcal{W}$, using (11) and (16), the dynamics of \mathcal{W} can be expressed as:

$$\dot{\mathcal{W}} = -T(\mathcal{W}A + A'\mathcal{W} + \mathcal{W}Q\mathcal{W} - \eta_i G). \quad (29)$$

Therefore, if we define the Lyapunov Function $V = \Sigma'_H \mathcal{W} \Sigma_H$, we have that:

$$\begin{aligned} \frac{d}{dq} (\Sigma'_H \mathcal{W} \Sigma_H) &= \Sigma'_H (T\mathcal{W}A + TA'\mathcal{W} + T\eta_i G + \dot{\mathcal{W}}) \Sigma_H \\ &= -T\Sigma'_H \mathcal{W} Q \mathcal{W} \Sigma_H. \end{aligned} \quad (30)$$

By integrating the previous relation, we have

$$\begin{aligned} \Sigma'_H(1)\mathcal{W}(1)\Sigma_H(1) - \Sigma_H(0)\mathcal{W}(0)\Sigma_H(0) &= \\ -T \int_0^1 \Phi(q, 0)' \mathcal{W} Q \mathcal{W} \Phi(q, 0) dq, \end{aligned} \quad (31)$$

where $\Phi(q_1, q_2)$ is the transition matrix of the system (23) between times q_1 and q_2 . Moreover, since $\bar{\Omega}_i(q)$ is periodic with period one and $\Sigma_H(0) = I$, we have that

$$\begin{aligned} \Sigma'_H(1)\mathcal{W}(0)\Sigma_H(1) - \mathcal{W}(0) &= \\ -T \int_0^1 \Phi(q, 0)' \mathcal{W} Q \mathcal{W} \Phi(q, 0) dq. \end{aligned} \quad (32)$$

Note that $\mathcal{W}\Phi(q, 0)$ is full rank on a nontrivial set, since \mathcal{W} is positive definite and $\Phi(q, 0)$ is full rank for at least a non-degenerate interval due to Assumption 1 and the fact that target i is observed at least once in a period. This, along with the fact that Q is positive definite, implies that the integral in (32) will be a positive definite matrix. Therefore,

$$\Sigma'_H(1)\mathcal{W}(0)\Sigma_H(1) - \mathcal{W}(0) \prec 0. \quad (33)$$

Consequently, one can see that

$$\frac{(\Sigma_H(1)x)' \mathcal{W}(0) (\Sigma_H(1)x)}{x' \mathcal{W}(0) x} < 1, \quad (34)$$

for every nonzero x . Since $\mathcal{W}(0)$ is positive definite, (34) shows that the norm of the matrix $\Sigma_H(1)$ induced by $\mathcal{W}(0)$ (i.e., $\|\Sigma_H(1)\|_{\mathcal{W}(0)}$) is less than 1, therefore its spectral radius is smaller than 1. This implies that the absolute value of all the eigenvalues of $\Sigma_H(1)$ are smaller than 1. Hence, $\Sigma_H(1) \otimes \Sigma_H(1)$ is stable, and $\Lambda = \tilde{\Lambda}$. Moreover, (25) has one solution given by

$$\Lambda = \sum_{j=1}^{\infty} (\Sigma_H(1))^j \Sigma_{ZI}(1) (\Sigma_H(1)')^j. \quad (35)$$

We point out that the sum in (35) converges, since the absolute value of the eigenvalues of $\Sigma_H(1)$ are all lower than 1.

Now, note that (22) is a first order linear matrix differential equation and its general solution is given by

$$\Sigma(q) = \Sigma'_H(q)\Sigma(0)\Sigma_H(q) + \Sigma_{ZI}(q). \quad (36)$$

Since there is a unique solution to (36), and when $\frac{\partial \bar{\Omega}_i(q)}{\partial \theta_d}$ exists it must satisfy (36), we know that $\Sigma(q) = \frac{\partial \bar{\Omega}_i(q)}{\partial \theta_d}$. \square

Remark 1. While the algebraic conditions for existence of solutions to generic Lyapunov equations are already well studied, the novelty in Prop. 3 consists in conditioning the existence of the solutions on physical characteristics of the persistent monitoring problem, namely Assumptions 1 and 2. Additionally, this proposition makes the connection between $\Sigma(q)$ and $\frac{\partial \bar{\Omega}_i(q)}{\partial \theta_d}$ and the existence of these partial derivatives is discussed in Appendix B.

The Lyapunov equation in (23) can be efficiently solved for low-dimensional systems using the algorithm proposed in [30]. Additionally, in order to compute the gradient, the partial derivatives of the steady state covariance matrices must be computed using the procedure in Prop. 3. Then, these partial derivatives are used along with (20) to compute the partial derivatives of the cost, which compose the gradient ∇J . Algorithm 2 summarizes this procedure.

Algorithm 2 Steady State Gradient Computation

```

1: procedure COMPUTESTEADYSTATEGRADIENT
2:   Input:  $\Theta$ 
3:   Compute  $s_1(q), \dots, s_N(q)$  from the parameterization
4:   for  $i$  ranging from 1 to  $M$  do
5:     Compute the steady state covariance  $\bar{\Omega}_i(q)$ 
6:     Compute  $\frac{\partial}{\partial \theta} \int_0^{t_f} \sum_{j=1}^N u_j'(\zeta) u_j(\zeta) d\zeta$  according to the
      parameterization
7:     Compute  $\frac{\partial s_j(t)}{\partial \theta}$  and  $\frac{\partial T}{\partial \theta}$  according to the parameterization
8:     for every  $\theta$  in  $\Theta$  do
9:       for  $i$  ranging from 1 to  $M$  do
10:        Compute  $\frac{\partial \bar{\Omega}_i(q)}{\partial \theta}$  as indicated in Prop. 3.
11:        Compute  $\frac{\partial J}{\partial \theta}$  using (20)
12:   Output:  $\nabla J$ 

```

In order to locally optimize the trajectories, the gradient computation needs to be used along with some gradient descent scheme. We describe the optimization procedure we used in Alg. 3, where κ_l is a scalar positive gain, and the proj operator projects the parameters into the set of feasible parameters ($u_j(t) \in \mathcal{U}$). As a side note, this projection might be difficult to compute in general and, therefore when choosing a parameterization it is important to make sure that there are efficient ways to compute this projection numerically.

Algorithm 3 Gradient Descent

```

1: procedure GRADIENT DESCENT
2:   Input:  $\Theta^0$ ,
3:    $||\nabla J|| \leftarrow \infty$ 
4:    $l \leftarrow 0$ 
5:   while  $||\nabla J|| > \epsilon$  do
6:      $\nabla J \leftarrow \text{ComputeGradient}(\Theta^l)$ 
7:      $\Theta^{l+1} \leftarrow \text{proj}(\Theta^l - \kappa_l \nabla J)$ 
8:      $l \leftarrow l + 1$ 
9:   Output:  $\underline{\Theta}$ 

```

IV. PARAMETERIZATION OF AN OPTIMAL TRAJECTORY IN 1-D WITH SPEED BOUNDS

When the agents and targets are constrained to a line, a particularly interesting case is the one where the absolute value of controls is bounded ($\mathcal{U} = \{u \in \mathbb{R} \mid |u| < u_{\max}\}$) and there is no penalty for control effort in the optimization cost J (i.e. $\beta = 0$). In this case we can represent optimized controls using a simple parameterization that could even lead to global optimality. It is worth noticing that in many real-world applications of persistent monitoring agents are constrained to (possibly multiple) uni-dimensional mobility paths, such as powerline inspection agents, cars on streets, and autonomous vehicles in rivers.

Assuming proper rescaling, we can consider $-1 \leq u_j \leq 1$, i.e., $\mathcal{U} = [-1, 1]$. In the remainder of this section, we derive properties of the optimal control, establish a parameterization that is able to represent an optimal control, and then compute the gradients necessary in order to optimize the trajectories.

A. Properties of an Optimal Control

In order to derive the properties of an optimal control, we first introduce the following lemma. The intuition behind it is that if a target is observed for a longer time (or with better quality), its uncertainty will be lower. We note that, although this lemma is introduced in this Section, it is not restricted to the 1D setting with bounded input.

Lemma 1. *Given $\Omega_1(t)$ and $\Omega_2(t)$, two bounded covariance matrices under the dynamics in (11) with $A = A_1 = A_2$, $G = G_1 = G_2$, $Q = Q_1 = Q_2$, then if $\Omega_1(0) - \Omega_2(0)$ is negative semi-definite and $\eta_1(t) \geq \eta_2(t) \forall t$, then $\Omega_1(t) - \Omega_2(t)$ is a negative semi definite matrix for all $t \geq 0$.*

Proof. Define $\beta = \Omega_1(t) - \Omega_2(t)$. The dynamics of β is described by the following equation.

$$\begin{aligned} \dot{\beta}(t) &= A\beta(t) + \beta A' - \eta_1(t)\Omega_1(t)G\Omega_1(t) \\ &\quad + \eta_2(t)\Omega_2(t)G\Omega_2(t). \end{aligned} \quad (37)$$

Adding and subtracting the terms $\eta_1(t)\Omega_2(t)G\Omega_2(t)$ and $\eta_1(t)\Omega_1(t)G\Omega_2(t)$ to the equation, we can rewrite (37) as:

$$\begin{aligned} \dot{\beta}(t) &= A\beta(t) + \beta A' - \eta_1(t)[\Omega_1(t)G\beta(t) + \beta(t)G\Omega_2(t)] \\ &\quad + [\eta_2(t) - \eta_1(t)]\Omega_2(t)G\Omega_2(t). \end{aligned} \quad (38)$$

From Thm. 1.e in [31], since $\beta(t)$ is a C^1 matrix, its eigenvalues can be C^1 time parameterized. Let μ_n denote the n^{th} eigenvalue of $\beta(t)$ and $x_n(t)$ the corresponding unit norm eigenvector. Then, from Thm. 5 in [32] we have that

$$\dot{\mu}_n = x_n' \dot{\beta} x_n.$$

Also, notice that by using (38) and the fact that $\lambda_{\min}\left(\frac{D+D'}{2}\right) \leq \frac{x'Dx}{\|x\|} \leq \lambda_{\max}\left(\frac{D+D'}{2}\right) = \|D\|$, for any square matrix D ,

$$\begin{aligned} \dot{\mu}_n &\leq \|A\| \mu_n - \eta_1 \beta \mu_n + [\eta_2 - \eta_1] x_n' \Omega_2 G \Omega_2 x_n \\ &\leq \|A\| \mu_n - \eta_1 \beta \mu_n, \end{aligned}$$

where $\beta = \lambda_{\min}((\Omega_1 + \Omega_2)G + G(\Omega_1 + \Omega_2))$. Using Gronwall's inequality [33] and the fact that the solution of a first order linear homogeneous ODE does not change sign, we conclude that $\mu_n(t) \leq 0, \forall t \in [0, T]$ and, therefore, $\beta(t)$ is negative semidefinite. \square

In Lemma 1, Ω_1 and Ω_2 can also be understood as covariance matrices for the same target but under different agent trajectories.

Before proceeding to the proposition about the structure of an optimal control, we first make the following assumption:

Assumption 3. *The function $\gamma_{i,j}(s_j)$ is unimodal and has a finite support equal to $r_{i,j}$.*

And we define an *isolated target* i as a target such that

$$\min_{k \neq i} |x_i - x_k| > 2r_{\max}, \quad r_{\max} = \max_{i,j} \{r_{i,j}\}.$$

Therefore, an isolated target is a target for which an agent cannot see another target when visiting it. Referring to the regions in space where an agent can sense a target as “visible areas”, the minimum distance between visible areas d_{\min} is defined as:

$$d_{\min} = \min_{i,k} |x_i - x_k| - 2r_{\max} > 0,$$

and the finite time cost is defined as

$$J(u_1, \dots, u_N, t) = \frac{1}{t} \int_0^t \left(\sum_{i=1}^M \text{tr}(\Omega_i(\beta)) \right) d\beta. \quad (39)$$

We can then claim the following proposition.

Proposition 4. *In an environment where all the targets are isolated, given any policy $u_j(\beta)$, $j = 1, \dots, N$, then there is a policy $\tilde{u}_j(\beta)$ with $\tilde{u}_j(\beta) \in \{-1, 0, 1\} \forall \beta \in [0, t]$ and with the number of control switches for each agent (i.e.*

discontinuities in $\tilde{u}_j(\beta)$) upper bounded by $2\frac{t}{d_{\min}} + 4$ such that $J(u_1, \dots, u_N, t) \geq J(\tilde{u}_1, \dots, \tilde{u}_N, t)$.

Proof. We prove this result by construction: given a policy $u_j(t')$ with $\eta_i(t')$ associated to it (as defined by (16)), we will construct an alternative policy $\tilde{u}_j(t')$ associated with $\tilde{\eta}_i(t')$ such that $\tilde{\eta}_i(t') \geq \eta_i(t') \forall t' \in [0, t]$ and $i = 1, \dots, M$, and then use Prop. 1, along with the definition of the cost (39), to show that the alternative policy has lower or equal cost than the original one.

Initially, we focus on the policy $u_j(t')$. We say that an agent j “visits” a target i if at some time t' , $|s_j(t') - x_i(t')| < r_j$. For every agent in the policy $u_j(t')$, there is an ordered collection of targets it visits in $[0, t]$. Therefore, there must exist a set of indices of all the targets visited by agent j : $\{y_0^j, \dots, y_{K_j}^j\} \subseteq \{1, \dots, M\}$, such that $y_p^j \neq y_{p-1}^j$ and agent j visited no other target in the time between visiting targets y_p^j and y_{p-1}^j . This is the sequence of all the targets that agent j visited over $[0, t]$, not considering consecutive visits to the same target. In other words, the same target can be present more than once in the sequence $\{y_0^j, \dots, y_{K_j}^j\}$ but, if that is the case, it will not be in consecutive positions.

For each of these visits, we can define the initial visiting time t_p^j for $p = 1, \dots, K_j$ as

$$t_p^j = \inf\{t' | t' > t_{p-1}^j \text{ and agent } j \text{ visits target } y_p^j \text{ at time } t'\},$$

and $t_0^j = 0$ and $t_{K_j+1}^j = t$. Also note that while $t' \in [t_{p-1}^j, t_p^j]$, agent j only influences the value of $\eta_i(j)$ of the target it is currently visiting. We propose the following alternative policy, where $\tilde{u}_j(t')$ for $t' \in [t_{p-1}^j, t_p^j]$ is such that:

$$\tilde{u}_j(t') = \begin{cases} \frac{s_j(t_p^j) - s_j(t')}{|s_j(t_p^j) - s_j(t')|}, & \text{if } \frac{|s_j(t_p^j) - s_j(t')|}{t_p^j - t'} \leq 1, \\ \frac{x_{y_p^j} - s_j(t')}{|x_{y_p^j} - s_j(t')|}, & \text{if } \frac{|s_j(t_p^j) - s_j(t')|}{t_p^j - t'} > 1 \text{ and} \\ & s_j(t') \neq x_{y_p^j}. \\ 0, & \text{otherwise.} \end{cases}$$

Notice that this construction provides a feasible trajectory, since the original trajectory is assumed feasible. Also, in the alternative policy $\tilde{u}_j(t') \in \{-1, 0, 1\} \forall t' \in [0, t]$, since the speed is either zero or a scalar divided by its absolute value.

The intuition behind the proposed alternative policy is that at the beginning of each visit, the agent moves with maximum speed towards the target y_p^j and if it reaches the target, it dwells on top of it. However, it must move in a way such that it begins the next visit at the same time as in the original policy, i.e., the positions of agent j associated to the alternative policy $\tilde{s}_j(t')$ is such that $\tilde{s}_j(t_p^j) = s_j(t_p^j)$.

Also, for time $t' \in [t_p^j, t_{p+1}^j]$ both the original and the alternative policies only influence the value of η_i for $i = y_p^j$, since in the alternative policy the agent is closer (or at least as close) to the currently visited target. Thus, from (16) we have that

$$\tilde{\eta}_i(t') \geq \eta_i(t'), \quad \forall t' \in [0, t], \quad i \in \{1, \dots, M\}.$$

Therefore, using Lemma 1 and the cost definition (39), we get that

$$\begin{aligned} J(\tilde{u}_1, \dots, \tilde{u}_N, t) - J(u_1, \dots, u_N, t) = \\ \frac{1}{t} \int_0^t \sum_{i=1}^M \text{tr}(\tilde{\Omega}_i(t') - \Omega_i(t')) \leq 0. \end{aligned}$$

which shows that the alternative policy has a lower or equal cost compared to the original one. Note that, due to velocity constraints, in both the original and the alternative policy there is a maximum of $\frac{t}{d_{\min}} + 1$ visits to targets per agent. Moreover, in the alternative policy, an agent has at most 2 velocity switches at each target visit. Therefore, at most $2\frac{t}{d_{\min}} + 4$ velocity switches can happen due to target visits, plus one switch to match the initial position of the original policy and another to match the terminal position of the original policy. This implies that the maximum number of velocity switches in the alternative policy is $2\frac{t}{d_{\min}} + 4$. \square

One way to interpret this proposition is that if one looks ahead at the next T units of time (where T is the period of a periodic solution or the prediction horizon, in the transient case), any trajectory can be improved (or at least, maintain same cost) by adequately selecting its controls $u_j(t)$ in the set $\{-1, 0, 1\}$. Also, notice that even though we were not able so far to prove that the same result holds when the targets are not necessarily isolated, the same structure can still be used but without the guarantee of optimality.

B. Parameterization of an Optimal 1D Trajectory

The result in Prop. 4 implies that when the targets are isolated, there is no loss of performance if we restrict ourselves to controls of the form $u_j(t) \in \{-1, 0, 1\} \forall t > 0$, with a bounded number of control switches. This property allows the optimal trajectory to be described by a finite set of parameters, similar to optimal control results in previous work by the authors [15], [18]. Here, in particular, we are looking into periodic trajectories and, hence, this property implies that the movement in each period of agent j consists of a sequence of dwelling at the same position for some duration of time followed by moving at maximum speed to another location. Therefore, one period of the trajectory of an agent j can fully be described by the following set of parameters:

- 1) T , the period of the trajectory.
- 2) $s_j(0)$, the initial position.
- 3) $\omega_{j,p}$, $p = 1, \dots, P_j$, the normalized dwelling times for agent j , i.e., the agent dwells for $\omega_{j,p}T$ units of time before it moves with maximum speed for the p -th time in the cycle.
- 4) $\tau_{j,p}$, $p = 1, \dots, P_j$, the normalized movement times for agent j , i.e., the agent j moves for $\tau_{j,p}T$ units of time to the right (if p is odd) or to the left (if p is even) after dwelling for $\omega_{j,p}T$ units of time in the same position.

To enforce consistency of the trajectory, we add the following constraints:

$$\tau_{j,m} \geq 0, \quad \omega_{j,m} \geq 0, \quad T \geq 0. \quad (40)$$

Notice that this description does not exclude transitions of u_j of the kind $\pm 1 \rightarrow \mp 1$ and $\pm 1 \rightarrow 0 \rightarrow \pm 1$, since it allows $\omega_{j,m} = 0$ and $\tau_{j,m} = 0$. In addition to the constraints in (40), in order to ensure periodicity, we need to make sure that the sum of the movement times and dwelling times does not exceed one period and that the total time spent moving to the left is equal to the total time spent moving to the right over one period (i.e. the agent returns to its initial position at the end of the period). Therefore, we have the additional constraints:

$$\sum_{m=1}^{P_j} (\tau_{j,m} + \omega_{j,m}) \leq 1, \quad \sum_{m=1}^{P_j} (-1)^m \tau_{j,m} = 0. \quad (41)$$

This parameterization defines a hybrid system in which the dynamics of the agents remain unchanged between events and abruptly switch when an event occurs. Events are given by a change in control value at completion of movement and dwell times. These may occur simultaneously, for instance, if the dwell time is zero (representing a switch of control from ± 1 to ∓ 1). This parameterization also applies to the aperiodic transient case, with minor modifications. Although we do not explore all the details for the sake of readability, we refer the interested reader to [18].

C. Position Gradients

Given this parameterization, we use the procedure given in Sec. III to optimize the cost. However, one item missing in Sec. III for computing the gradient of the covariance matrix was the gradient of the agent position with respect to the parameters defining the trajectory.

The movement and dwelling time parameterization defines, along with the uncertainty metric, a hybrid system. For such systems, Infinitesimal Perturbation Analysis (IPA) can be used to compute an event-driven online estimate of the stochastic gradient of the system. An important feature of IPA is that the unbiased gradient estimate can be computed online using only the data observed along the trajectory. Even though we do not discuss in this paper the details of the IPA interpretation of the equations in this subsection, we refer the reader to [15], [34] for more information about IPA.

One can see that the position of agent j at normalized time q , after the k -th event and before the $k+1$ -th is

$$s_j(q) - s_j(0) = \begin{cases} T \left((-1)^{k/2+1} \left(q - \sum_{p=1}^{k/2-1} (\tau_{j,p} + \omega_{j,p}) \right. \right. \\ \left. \left. + \omega_{j,\frac{k}{2}} \right) + \sum_{p=1}^{k/2} (-1)^{p+1} \tau_p \right), & k \text{ even}, \\ T \sum_{p=1}^{\frac{k-1}{2}} (-1)^{p+1} \tau_{j,p}, & k \text{ odd}. \end{cases}$$

Therefore, we can compute the following gradients,

$$\begin{aligned} \frac{\partial s_j}{\partial \tau_{j,m}} &= \begin{cases} \left((-1)^{\frac{k}{2}} + (-1)^{m+1} \right) T, & m < \frac{k}{2}, k \text{ even}, \\ (-1)^{m+1} T, & m \leq \frac{k-1}{2}, k \text{ odd}, \\ 0, & \text{otherwise}, \end{cases} \\ \frac{\partial s_j}{\partial \omega_{j,m}} &= \begin{cases} T, & m \leq \frac{k}{2}, k \text{ even}, \\ 0 & \text{otherwise}, \end{cases} \\ \frac{\partial s_j(q)}{\partial T} &= \frac{s_j(q) - s_j(0)}{T}, \\ \frac{\partial s_j}{\partial s_j(0)} &= 1. \end{aligned}$$

D. Initial Trajectory for the Optimization

While we use a gradient descent approach in Alg. 3 to locally minimize the cost function, it is necessary to find an initial parameter configuration. Therefore, we propose a method to efficiently compute a starting point for the optimization.

Proposition 2 states that if every target is visited at least once in a periodic trajectory, then the steady-state covariance matrix exists. However, if in a periodic trajectory one of the targets is never visited and its internal state dynamics is unstable, then the estimation error will grow without bound as time goes to infinity. Also, when a target is not visited in the initial trajectory, the gradient descent optimization may converge to undesired solutions, a problem known as the ‘lack of event excitation’ and discussed in depth in [35]. Therefore, this kind of initial trajectory will not be considered in this work.

In this section, we discuss a method for finding these initial trajectories that will always lead to a feasible initial configuration. Due to the local nature of our optimization procedure, different initial conditions can lead to different local optima. We, therefore, leverage intuition about the problem to provide reasonable initial solutions with the hope that they will converge to good local optima.

The idea of finding a schedule where all the targets are visited fits naturally into a graph search paradigm, where the targets are modelled as nodes and the edge weights between nodes are the distances between the targets. The problem of finding a feasible schedule can be translated to one of finding N sequences (that represent the schedule of each agent) of nodes where each target belongs to at least one of these sequences. One can add to that a cost function that guides the way in which these sequences are created. A goal that intuitively will lead to reasonable initial solutions is to minimize the distance of the agent that has the longest travel path. This is the well known MTSP (see [6] for a good overview of this problem and approaches to solve it). It is worth mentioning that the MTSP is in general NP-hard, thus computationally intractable. However, in 1D environments, it can be efficiently solved using the algorithm described in [8].

The MTSP problem finds a minimal length cycle and therefore can be immediately converted to parameters that represent one period of the steady state solution. We choose the dwelling times to be initially zero.

Remark 2. *Although we use MTSP in the generation of the initial trajectory, this is not an integral part of our approach and any initialization technique could be integrated to our optimization scheme, as long as it guarantees that each target is visited. Alternative initialization procedures include variations of MTSP where the optimization goal goes beyond travel distance minimization (see e.g. [36]). However, to the authors' best knowledge, none of the MTSP variations available in the literature can directly minimize the estimation uncertainty, therefore using these different MTSP variations into our framework is not trivial and we intend to pursue this in future work.*

E. 1D Simulation Results

In the simulations, we have chosen to highlight interesting aspects of the solution, rather than simply give an example of the techniques discussed in this paper. We have analyzed a steady state problem with 2 agents and 5 targets. We used the following matrices in the state evolution model

$$A_i = \begin{bmatrix} -1 & -0.1 \\ -0.1 & 0.01 \end{bmatrix}, \quad Q_i = \text{diag}(1, 1),$$

and the following parameters for the observation model

$$H_i = \text{diag}(1, 1), \quad R_i = \text{diag}(1, 1).$$

Moreover, we considered the following expression for $\gamma_{i,j}$:

$$\gamma_{i,j}(\alpha) = \begin{cases} 0, & \|\alpha\| > r_{i,j}, \\ \sqrt{1 - \frac{\|\alpha\|}{r_{i,j}}}, & \|\alpha\| \leq r_{i,j}. \end{cases} \quad (43)$$

with $r_{i,j} = 0.9$. The intuition behind this specific form is that the best measurement quality is achieved when the agent's location coincides with that of the target, with the SNR decaying linearly as the agent moves away. When the agent is at a distance larger or equal to its sensing radius $r_{i,j}$, only noise is observed.

Instead of using the initialization method proposed in Subsec. IV-D, we used the following set of parameters:

$$s_1^0(0) = 2.7, \quad s_2(0) = 6.8, \quad T^0 = 6,$$

$$\tau_1^0 = \tau_2^0 = 0.1[1, 0.1, 1, 1, 0.1, 1, 0.1, 1, 1, 0.1, 1],$$

$$\omega_1^0 = \omega_2^0 = 0.0125[1, 1, 1, 1, 1, 1, 1, 1, 1, 1].$$

The goal of using these initialization parameters was to have both agents share one target in the first iteration of the optimization process and then explore whether or not they would remain sharing the target after the local optimization procedure. The gradient descent step size was set to be constant, $\kappa_0 = \kappa_l = 0.02$.

Figure 1 shows the results of the optimization in this scenario. Although both agents and all the targets have the same dynamic models, the solution at the last iteration of the optimization was such that one of the agents visits three of the targets and the other two of them. Note that on the trajectories in Fig. 1b, in the period between times 6 and 8, agent 1 moves with a small amplitude around target 1.

The effects of this oscillatory movement are hard to notice in the trace of the covariance of target 1 in Fig. 1c, since the difference in performance is negligible. Therefore, even though it is intuitively clear that staying still rather than moving with this oscillatory behavior will lead to a lower cost solution, the difference in terms of cost is minor. Also, the solution has not yet fully converged, as can be seen in Fig. 1a and further iterations would remove this small oscillatory behavior. Finally, we point out that while the maximum number of switches in a direction allowed to each agent was set to 11, the final solution appears to have fewer because some of the movement and dwelling times in the final solution are zero.

V. FOURIER CURVES FOR MULTI-DIMENSIONAL PERSISTENT MONITORING WITH UNBOUNDED SPEED

For the 1D case we derived a parameterization with a finite number of parameters of the optimal solution. Unfortunately, the same result does not extend to multi-dimensional persistent monitoring problems. Therefore, instead of looking for an exact representation of the optimal trajectory, we focus on a family of parameterized curves that can approximate very general curves. We pick as an illustration the case where speed is not bounded, in part because the projection operation in line 7 of Alg. 3 becomes trivial. Note that whenever the constant that weights the control effort penalization is not zero, i.e. $\beta \neq 0$ as defined in (19), the fact that the control effort is considered in the total cost will not allow the control to be unbounded. An appropriate choice of β can provide adequate speed bounds for any given dynamics of the system. As a side note, we highlight that bounded speeds can also be handled in this framework, however the projection operator in the gradient descent optimization becomes more complex.

Since periodicity is an essential feature of the steady-state analysis discussed in this work, a natural choice is to use a truncated Fourier series to represent the movement of the agents in each of the coordinates e_p , $p = 1, \dots, P$, i.e.

$$s_j^{e_p}(q) = s_{j,0}^{e_p} + \sum_{k=1}^K a_{j,k}^{e_p} \sin(2\pi f_k q) + b_{j,k}^{e_p} (\cos(2\pi f_k q) - 1),$$

where f_k are integer frequencies and, therefore, $s_j^{e_p}(q)$ is periodic with period 1. The set of parameters that fully characterize all the agents trajectories is $\Theta = \{ \{a_{j,k}^{e_p}\}, \{b_{j,k}^{e_p}\}, \{s_{j,0}^{e_p}\}, T \}$, $j = 1, \dots, N$, $p = 1, \dots, P$, $k = 1, \dots, K$. As in the 1D case, in order to compute the derivative of the covariance matrix, we need to give a procedure to compute $\frac{\partial s_k}{\partial \theta}$. For any parameter $\theta \in \Theta$,

$$\frac{\partial s_j^{e_p}}{\partial a_{m,k}^{e_r}} = \begin{cases} \sin(2\pi f_k q), & \text{if } j = m \text{ and } p = r, \\ 0, & \text{otherwise,} \end{cases} \quad (44a)$$

$$\frac{\partial s_j^{e_p}}{\partial b_{m,k}^{e_r}} = \begin{cases} \cos(2\pi f_k q) - 1, & \text{if } j = m \text{ and } p = r, \\ 0, & \text{otherwise,} \end{cases} \quad (44b)$$

$$\frac{\partial s_j^{e_p}}{\partial s_{m,0}^{e_r}} = \begin{cases} 1, & \text{if } j = m \text{ and } p = r, \\ 0, & \text{otherwise,} \end{cases} \quad (44c)$$

$$\frac{\partial s_j^{e_p}}{\partial T} = 0. \quad (44d)$$

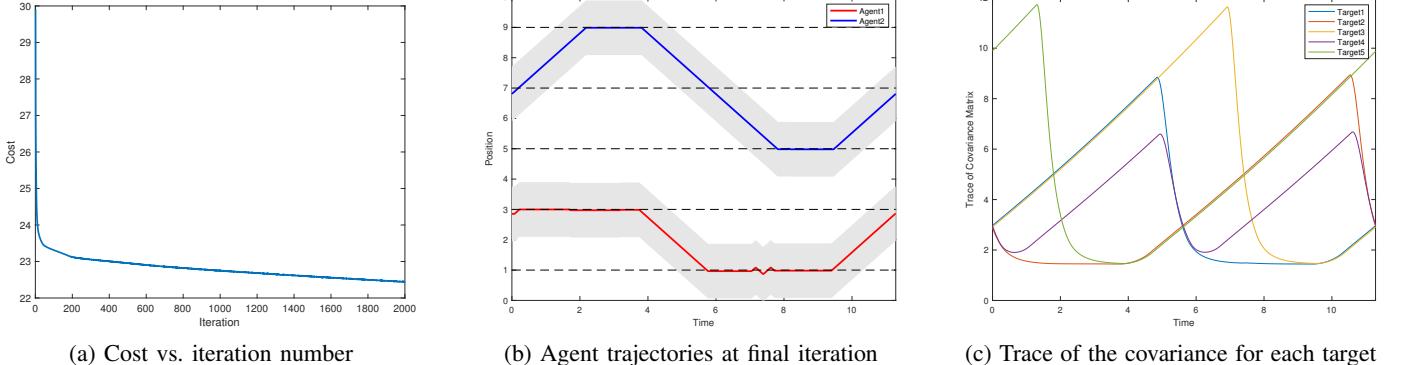


Fig. 1: Results of a simulation with two agents and five targets. (a) Evolution of the overall cost as a function of iteration number on the gradient descent. (b) Trajectories of the agents at the final iteration. The dashed lines indicate the positions of the targets and the grey shaded area the visibility region of the agent. (c) Evolution of the trace of the estimation covariance matrices of the five targets.

The derivatives in (44) give enough information to compute the partial derivatives of the steady state covariance matrix as indicated in Prop. 3. In order to compute the gradient of the cost function, the following expression can be used:

$$\frac{\partial J}{\partial \theta} = \int_0^1 \sum_{i=1}^N \text{tr} \left(\frac{\partial \Omega_i}{\partial \theta} \right) dq + \beta \frac{\partial}{\partial \theta} \sum_{j=1}^N \int_0^1 \left\| \frac{ds_j}{dt} \right\|^2 dq. \quad (45)$$

Note that

$$\frac{ds_j}{dq} = T \frac{ds_j}{dt}. \quad (46)$$

Using (V), we can compute

$$\begin{aligned} \sum_{j=1}^N \int_0^1 \left\| \frac{ds_j}{dt} \right\|^2 dq = \\ \sum_{j=1}^N \sum_{p=1}^P \sum_{k=1}^K \frac{(2\pi f_k)^2}{2T^2} \left(\left(a_{j,k}^{e_p} \right)^2 + \left(b_{j,k}^{e_p} \right)^2 \right), \end{aligned} \quad (47)$$

and, therefore,

$$\frac{\partial}{\partial a_{j,k}^{e_p}} \sum_{j=1}^N \int_0^1 \left\| \frac{ds_j}{dt} \right\|^2 dq = \frac{(2\pi f_k)^2}{2T^2} a_{j,k}^{e_p}, \quad (48a)$$

$$\frac{\partial}{\partial b_{j,k}^{e_p}} \sum_{j=1}^N \int_0^1 \left\| \frac{ds_j}{dt} \right\|^2 dq = \frac{(2\pi f_k)^2}{2T^2} b_{j,k}^{e_p}, \quad (48b)$$

$$\frac{\partial}{\partial s_{j,0}^{e_p}} \sum_{j=1}^N \int_0^1 \left\| \frac{ds_j}{dt} \right\|^2 dq = 0, \quad (48c)$$

$$\begin{aligned} \frac{\partial}{\partial T} \sum_{j=1}^N \int_0^1 \left\| \frac{ds_j}{dt} \right\|^2 dq = \\ \sum_{j=1}^N \sum_{p=1}^P \sum_{k=1}^K \frac{-(2\pi f_k)^2}{T^3} \left(\left(a_{j,k}^{e_p} \right)^2 + \left(b_{j,k}^{e_p} \right)^2 \right). \end{aligned} \quad (48d)$$

A. Optimization Initialization

In the multi-dimensional optimization, we still use the sub-optimal solution of the MTSP problem as a starting point.

However, unlike the 1-D scenario with the movement and dwelling time parameterization, the heuristic solution of the MTSP problem cannot be directly converted to a Fourier Curve trajectory. The solution of the MTSP problem gives, for each agent j , a cyclic schedule of targets $\mathcal{S}_j = \{y_j^1, \dots, y_j^{Y_j}, y_j^1\}$ and, therefore, it is still necessary to obtain the parameters $\Theta = \{\{a_{j,k}^{e_p}\}, \{b_{j,k}^{e_p}\}, \{s_{j,0}^{e_p}\}, T\}$ from this schedule. We define d_j^m as the cumulative distance that the agent has traveled when it reaches the m -th target in the schedule \mathcal{S}_j , and D_j as the total distance traveled by an agent in one cycle. We then look for a feasible truncated Fourier series trajectory such that at the normalized time $q = d_j^m / (D_j T)$, the agent is at a distance lower or equal to the sensing radius (multiplied by a factor $1 - \delta$, $0 < \delta < 1$, in order to give some distance margin) from the target. The position of the agent at the beginning of the cycle is set to be the position of the first target in the schedule \mathcal{S}_j . The period T can be set to any positive number. For each of the agents, the following optimization problem gives a set of feasible $\{a_{j,k}^{e_p}\}, \{b_{j,k}^{e_p}\}$.

$$\begin{aligned} \min_{a_{j,k}^{e_p}, b_{j,k}^{e_p}} \quad & \sum_{p=1}^P \sum_{k=1}^K f_k |a_{j,k}^{e_p}| + f_k |b_{j,k}^{e_p}| \\ \text{s.t.} \quad & \left\| s_j \left(\frac{d_j^m}{D_j} \right) - x_{y_j^m} \right\|_2 \leq (1 - \delta) r_j, \quad m = 1, \dots, Y_j \end{aligned} \quad (49)$$

If we substitute the definition (V) into the constraint (49), this optimization can be formulated as a Quadratically Constrained Program, which is a convex optimization problem that can be solved efficiently. From our experience, minimizing a weighted sum of absolute values in the objective function of (49) has led to smooth initial trajectories. However, other optimization objectives could be used.

It is worth observing that for each of the agents, the trajectory generated by the heuristic solution of the MTSP problem consists of segments of straight lines that visit each of the targets in the schedule \mathcal{S}_j . Note that this trajectory, as a function of time, composed by sequence of straight lines can be projected in each of the axis e_p and the projection in

that axis will still be a sequence of segments of straight lines. Since piecewise linear functions can be represented by Fourier series, there always exist a K large enough such that there is a solution to (49) because for that K there is a representation of the trajectory that would be close enough to the original MTSP solution such that it is able to satisfy the constraint in (49). Therefore, we can always find feasible solutions to (49) if we have a MTSP solution.

In addition, we highlight that although in 1D the TSP can be efficiently solved, this is not the case for higher dimensional environments. However, there are meta-heuristic approaches that can provide feasible, though not necessarily optimal, solutions. In this work, we use the genetic algorithm described in [37] to find heuristic solutions. This approach finds a feasible solution in the first iteration and refines it as the number of iterations increases. Therefore, one can decide how much computation time to spend, leveraging the tradeoff between optimality and computation effort.

B. 2D Simulation Results

In this section, we demonstrate the results of the algorithm in two simulated 2D scenarios, one with one agent and three targets and the other one with three agents and 15 targets. All the internal states of the targets have the same state dynamics, evolving according to (1) with

$$A_i = \begin{bmatrix} -1 & -0.1 \\ -0.1 & 0.01 \end{bmatrix}, \quad Q_i = \text{diag}(1, 1),$$

and the agent observation models are given by (3) and the expression of $\gamma_{i,j}$ is given in (43) with

$$H_i = R_i = \text{diag}(1, 1), \quad r_{i,j} = 0.5.$$

and objective weighting, as given in (12), $\eta = 10^{-3}$. For each of the agents, their trajectories had the first five harmonics in each axis, i.e., $f_k = k$, $k = 1, \dots, 5$, $\forall j$. In the initial step of the optimization, the period T was set to 1. The initial coefficients $a_{j,k}^{ep}, b_{j,k}^{ep}$ were obtained by solving the optimization problem in (49). The MTSP solution was obtained after 3000 iterations of the genetic algorithm proposed in [37] for solving the associated MTSP. The initial position of each agent was set to coincide with the position of the first target in the solution of the MTSP. A constant descent stepsize $\kappa_l = 10^{-4}$ was used in the gradient descent.

In the first scenario (with one agent and three targets), targets were located at positions $x_1 = (0, 0.5)$, $x_2 = (0.5, 0)$ and $x_3 = (-0.5, 0)$. Figures 2-4 show the results we obtained. Figure 2a highlights how the trajectory changed from the initial one (an ellipse) to one with an almost triangular shape. Note, however, that not only the geometry of the trajectory is being optimized, but also the speed of the agent along the trajectory. From Fig. 2b we can see that the agent moves with higher speed when it is not visiting any target and at reduced speed (and the speed even completely vanishes) when it is close to the targets. Also, we can note that the trajectory in the last step of the optimization had a period lower than 1, which was the period on the initial optimization step. The mean estimation error over time for each of the targets is

displayed in Fig. 3 and the cost along the optimization process is shown in Fig. 4.

In the second scenario, the positions of the targets were generated randomly from independent uniform distributions ranging from -5 to 5 in both axes. Fig. 5 compares the trajectories of the agents in the first and last step of the gradient descent optimization, while Fig. 6 shows the evolution of the cost as a function of the gradient descent step. The results of the optimization show that the solution of (49) led to smoother trajectories that still visited all the targets. The gradient descent changed the geometry of the trajectories but did not change the visiting order. As can be observed in Fig. 6, the cost has an abrupt reduction in the beginning of the optimization and then the convergence speed reduces significantly. The optimization process leads to very significant reductions of the cost, reducing it to less than one third of its initial value.

C. 3D Simulations Results

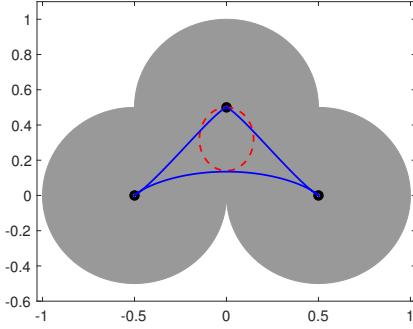
In order to illustrate the extension of techniques proposed in this paper to higher dimensions, we present a result in a 3D environment, with 2 agents and 10 targets. The A_i , Q_i , H_i , R_i matrices and r_j are the same as in the 2D simulations. A constant gradient descent stepsize $\kappa_l = 10^{-2}$ was used. The target locations were drawn from a uniform distribution in the cube with coordinates ranging from $[-5, 5]$ in each axis. The trajectories after 4000 gradient descent iterations are shown in Fig. 7 and the evolution of the cost is displayed in Fig. 8.

The 3D results follow a very similar trend of the 2D ones. The trajectories provided by the initialization procedure tend to be smoother, while the shape of the optimized ones are stiffer.

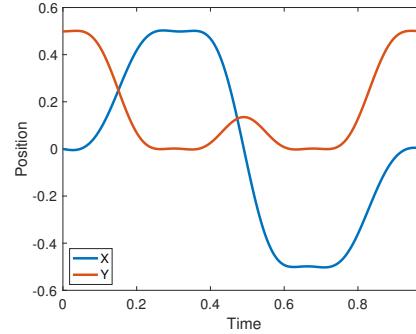
VI. CONCLUSION AND FUTURE WORK

We have addressed the problem of persistent monitoring from an infinite horizon perspective. We used a model that captures internal states of the targets evolving with linear stochastic dynamics and an observation model where the observation quality varies with distance. We derived necessary conditions for the convergence of the covariance matrix to a limit cycle as time goes to infinity. We also provided an algorithm for computing the cost gradient with respect to the parameters that define the trajectory. For a 1D environment, we showed that under some assumptions it is possible to fully characterize an optimal control by a finite set of parameters and used this as a basis for constructing an efficient parameterization. In higher dimensions, we proposed the use of Fourier curves for representing the trajectory. Our simulations illustrated the application of the proposed techniques in 1D, 2D and 3D scenarios, considering finite and infinite horizons for the cost.

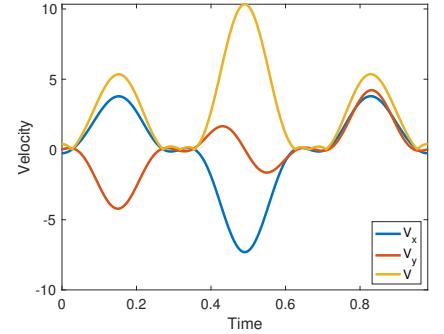
Some challenges still remain for the framework presented in this paper. In ongoing research, we are studying how to efficiently select the gradient descent stepsize and also the feasibility and efficiency of local optimization methods other than gradient descent. We intend to study initialization methods that directly use uncertainty (instead of distance) as a criterion for generating initial schedules. We also plan



(a) 2D trajectory of the agent.



(b) Agent positions as a function of time.



(c) Agent velocity.

Fig. 2: Simulation results with one agent and three targets. (a) Comparison of the 2-D trajectory of the agents in the initial (red) and final (blue) trajectories. The targets are marked in black and the gray area is the region where an agent can sense that target. (b) Agent trajectories at the final iteration in the x (blue) and y (red) directions. (c) Agent velocities at the final iteration in the x (blue) and y (red) directions and the resulting agent speed (yellow).

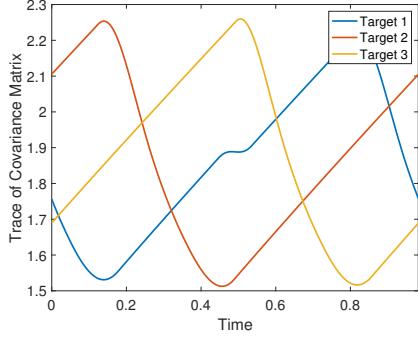


Fig. 3: Trace of the covariance for each target at the final step of the optimization in the scenario with one agent and three targets.

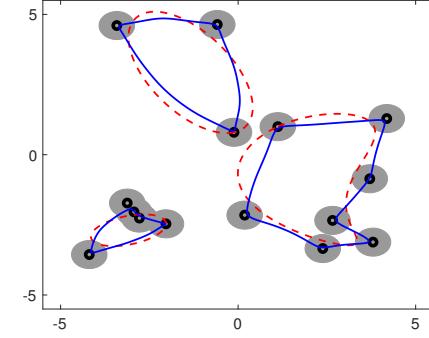


Fig. 5: Trajectories of the targets in the first (red dashed line) and last (blue continuous line) iterations of the gradient descent optimization on the scenario with 15 targets and 3 agents. The target's locations are marked in black and the grey shaded are represent the regions where the target can be sensed by an agent.

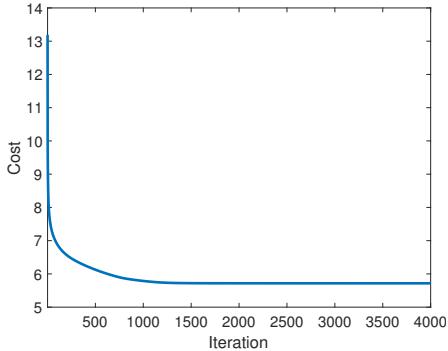


Fig. 4: Evolution of the cost function in the gradient descent optimization in the scenario with 3 targets and 1 agent.

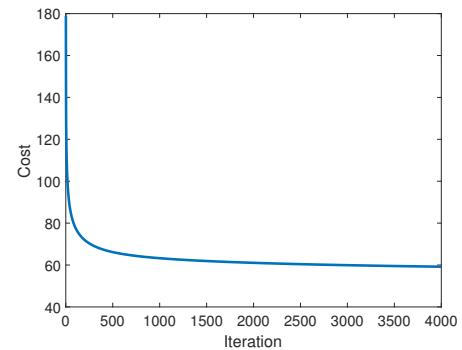


Fig. 6: Evolution of the cost function in the gradient descent optimization in the scenario with 15 targets and 3 agents.

to extend this paradigm to discrete time formulations and to investigate the feasibility of distributed solutions. Lastly, we plan to study this problem when targets can also move, including movement models that are not fully deterministic.

REFERENCES

[1] 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*. IEEE, 2013, pp. 2415–2420.

[2] ——, “A Variational Approach to Trajectory Planning for Persistent Monitoring of Spatiotemporal Fields,” in *2014 American Control Conference*. IEEE, 2014, pp. 5627–5632.

[3] T. Alam, G. M. Reis, L. Bobadilla, and R. N. Smith, “A Data-Driven Deployment Approach for Persistent Monitoring in Aquatic Environments,” in *IEEE International Conference on Robotic Computing*, 2018, pp. 147–154.

[4] H. Kim, L. Mokdad, and J. Ben-Othman, “Designing UAV Surveillance Frameworks for Smart City and Extensive Ocean with Differential

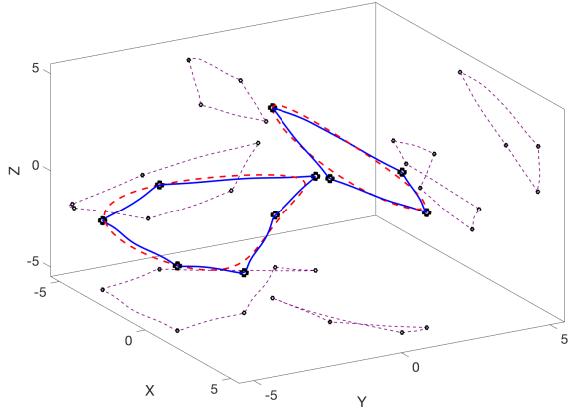


Fig. 7: Simulation results in a 3D environment with two targets and ten agents. In red, the initial trajectory in the gradient descent optimization, in blue, the trajectory at the end of the optimization. The projection of the final agent trajectories in three planes is plotted in dashed purple.

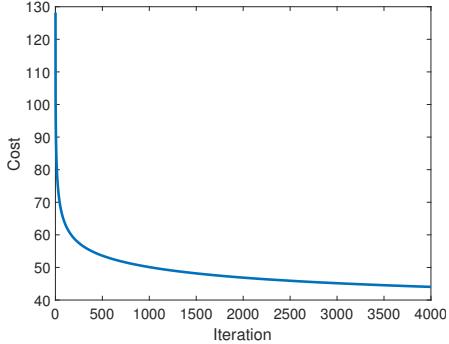


Fig. 8: Evolution of the cost function in the gradient descent optimization in the 3D scenario.

Perspectives,” *IEEE Communications Magazine*, vol. 56, no. 4, pp. 98–104, 2018.

- [5] 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.
- [6] T. Bektas, “The Multiple Traveling Salesman Problem: an Overview of Formulations and Solution Procedures,” *Omega*, vol. 34, no. 3, pp. 209 – 219, 2006.
- [7] G. Laporte, “Fifty years of Vehicle Routing,” *Transportation Science*, vol. 43, no. 4, pp. 408–416, 2009.
- [8] F. Pasqualetti, A. Franchi, and F. Bullo, “On Cooperative Patrolling: Optimal Trajectories, Complexity analysis, and Approximation Algorithms,” *IEEE Transactions on Robotics*, vol. 28, no. 3, pp. 592–606, 2012.
- [9] J. Le Ny, E. Feron, and M. A. Dahleh, “Scheduling Continuous-Time Kalman Filters,” *IEEE Transactions on Automatic Control*, vol. 56, no. 6, pp. 1381–1394, 2010.
- [10] C. Baykal, G. Rosman, K. Kotowick, M. Donahue, and D. Rus, “Persistent Surveillance of Events with Unknown Rate Statistics,” in *Algorithmic Foundations of Robotics XII*. Springer, 2020, pp. 736–751.
- [11] B. Grocholsky, A. Makarenko, and H. Durrant-Whyte, “Information-theoretic Coordinated Control of Multiple Sensor Platforms,” in *2003 IEEE International Conference on Robotics and Automation*. IEEE, 2003, pp. 1521–1526.
- [12] I. I. Hussein, “Kalman Filtering with Optimal Sensor Motion Planning,” in *2008 American Control Conference*. IEEE, 2008, pp. 3548–3553.
- [13] A. Jones, M. Schwager, and C. Belta, “Information-guided Persistent Monitoring Under Temporal Logic Constraints,” in *American Control Conference (ACC)*, 2015, pp. 1911–1916.
- [14] S. C. Pinto, S. B. Andersson, J. M. Hendrickx, and C. G. Cassandras, “A Semidefinite Approach to Discrete-time Persistent Monitoring,” in *European Control Conference (to appear)*, 2021.
- [15] 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, 2013.
- [16] 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.
- [17] X. Yu, S. B. Andersson, N. Zhou, and C. G. Cassandras, “Optimal Visiting Schedule Search for Persistent Monitoring of a Finite Set of Targets,” in *American Control Conference (ACC)*. IEEE, 2018, pp. 4032–4037.
- [18] S. C. Pinto, S. B. Andersson, J. M. Hendrickx, and C. G. Cassandras, “Optimal Multi-Agent Persistent Monitoring of the Uncertain State of a Finite Set of Targets,” in *Control and Decision Conference*, 2019.
- [19] —, “Optimal Periodic Multi-Agent Persistent Monitoring of a Finite Set of Targets with Uncertain States,” in *American Control Conference*, 2020.
- [20] —, “Multi-Agent Infinite Horizon Persistent Monitoring of Targets with Uncertain States in Multi-Dimensional Environments,” in *IFAC World Congress*, 2020.
- [21] Y. Wang, Y. Wei, X. Liu, N. Zhou, and C. G. Cassandras, “Optimal Persistent Monitoring Using Second-Order Agents with Physical Constraints,” *IEEE Transactions on Automatic Control*, 2018.
- [22] M. S. Mahmoud and H. M. Khalid, “Distributed Kalman Filtering: a Bibliographic Review,” *IET Control Theory & Applications*, vol. 7, no. 4, pp. 483–501, 2013.
- [23] M. Athans and E. Tse, “A Direct Derivation of the Optimal Linear Filter Using the Maximum Principle,” *IEEE Transactions on Automatic Control*, vol. 12, no. 6, pp. 690–698, 1967.
- [24] L. Zhao, W. Zhang, J. Hu, A. Abate, and C. J. Tomlin, “On the Optimal Solutions of the Infinite-horizon Linear Sensor Scheduling Problem,” *IEEE Transactions on Automatic Control*, vol. 59, no. 10, pp. 2825–2830, 2014.
- [25] S. Bittanti, P. Colaneri, and G. Guardabassi, “Periodic Solutions of Periodic Riccati Equations,” *IEEE Transactions on Automatic Control*, vol. 29, no. 7, pp. 665–667, 1984.
- [26] G. Nicolao, “On the Convergence to the Strong Solution of Periodic Riccati Equations,” *International Journal of Control*, vol. 56, no. 1, pp. 87–97, 1992.
- [27] A. Varga, “Computational Issues for Linear Periodic Systems: Paradigms, Algorithms, Open Problems,” *International Journal of Control*, vol. 86, no. 7, pp. 1227–1239, 2013.
- [28] P. Bolzern and P. Colaneri, “The Periodic Lyapunov Equation,” *SIAM Journal on Matrix Analysis and Applications*, vol. 9, no. 4, pp. 499–512, 1988.
- [29] F. Zhang, *Matrix Theory: Basic Results and Techniques*. Springer Science & Business Media, 2011.
- [30] A. Barraud, “A Numerical Algorithm to Solve AT XA-X=Q,” in *1977 IEEE Conference on Decision and Control including the 16th Symposium on Adaptive Processes and A Special Symposium on Fuzzy Set Theory and Applications*, no. 16, 1977, pp. 420–423.
- [31] A. Kriegel, P. W. Michor, and A. Rainer, “Denjoy–Carleman Differentiable Perturbation of Polynomials and Unbounded Operators,” *Integral Equations and Operator Theory*, vol. 71, no. 3, p. 407, 2011.
- [32] P. Lancaster, “On Eigenvalues of Matrices Dependent on a Parameter,” *Numerische Mathematik*, vol. 6, no. 1, pp. 377–387, 1964.
- [33] T. H. Gronwall, “Note on the Derivatives with Respect to a Parameter of the Solutions of a system of Differential Equations,” *Annals of Mathematics*, pp. 292–296, 1919.
- [34] 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.
- [35] Y. Khazaeni and C. G. Cassandras, “Event Excitation for Event-driven Control and Optimization of Multi-agent Systems,” in *13th International Workshop on Discrete Event Systems*, 2016, pp. 197–202.
- [36] N. Jozefowicz, F. Semet, and E.-G. Talbi, “Multi-objective Vehicle Routing Problems,” *European journal of operational research*, vol. 189, no. 2, pp. 293–309, 2008.
- [37] L. Tang, J. Liu, A. Rong, and Z. Yang, “A Multiple Traveling Salesman Problem Model for Hot Rolling Scheduling in Shanghai Baoshan Iron &

APPENDIX A

PROOF OF OPTIMALITY OF KALMAN BUCY FILTER ON THE PERSISTENT MONITORING PROBLEM WITH UNCERTAIN STATES

The set of all unbiased estimators $\hat{\phi}_i(t)$ of $\phi_i(t)$, as discussed in Sec. IV of [23], is:

$$\dot{\hat{\phi}}_i(t) = \left(A_i - G_i(t)\tilde{H}_i(t) \right) \hat{\phi}_i(t) + G_i(t)\tilde{z}_i(t), \quad (50)$$

with $E[\hat{\phi}_i](0) = E[\phi_i(0)]$ and $G(t)$ a gain function that should be considered an input for the sake of optimality analysis. If $\Omega_i(t) = E[e_i(t)e_i'(t)]$, where $e_i = \hat{\phi}_i(t) - \phi_i(t)$, then

$$\begin{aligned} \dot{\Omega}_i(t) &= \left(A_i - G_i(t)\tilde{H}_i(t) \right) \Omega_i(t) + G_i(t)\tilde{R}_i G_i'(t) \\ &\quad + Q_i + \Omega_i(t) \left(A_i' - \tilde{H}_i(t)' G_i'(t) \right) \end{aligned} \quad (51)$$

and $\Omega_i(0) = \Omega_{i,0}$. Defining the following cost:

$$J = \int_0^{t_f} \left(\sum_{i=1}^M \text{tr}(\Omega_i(t')) + \beta u'(t')u(t') \right) dt' \quad (52)$$

The Hamiltonian is then

$$\begin{aligned} \mathcal{H} &= \sum_{i=1}^M \text{tr}(\Omega_i(t)) + \beta u'(t)u(t) \\ &\quad + \sum_{i=1}^M \text{tr} \left(\Gamma_i(t)\dot{\Omega}_i(t) \right) + \sum_{j=1}^N \alpha_j(r)s_j(t), \end{aligned} \quad (53)$$

where Γ_i is the costate of Ω_i . Using Pontryagin's minimum principle, at an optimal trajectory, since G_i is unconstrained, we have

$$\frac{\partial \mathcal{H}^*}{\partial G_i} = 0. \quad (54)$$

Substituting the dynamics of the covariance matrix (51) on (54), we get

$$-\Gamma_i \Omega_i \tilde{H}'_i - \Gamma'_i \Omega_i \tilde{H}'_i + \Gamma'_i G_i \tilde{R}_i + \Gamma_i G_i \tilde{R}_i = 0. \quad (55)$$

Now, again from the minimum principle,

$$\dot{\Gamma}_i = -\frac{\partial \mathcal{H}}{\partial \Omega_i} - (A_i - G_i \tilde{H}_i)' \Gamma_i - \Gamma_i (A_i - G_i \tilde{H}_i) - I. \quad (56)$$

Since $\Gamma_i(t_f) = 0$ due to the boundary conditions of Pontryagin's minimum principle, the symmetric nature of this ODE allow us to see that Γ_i will be symmetric for $t \in [0, t_f]$. Moreover, note that the ODE is linear and the single non-homogeneous term is $-I$. Since $\Gamma_i(t_f) = 0$,

$$\begin{aligned} \Gamma_i(t) &= - \int_{t_f}^t \Phi'(t, t_f) \Phi(t, t_f) dt, \\ \Phi(a, b) &= \exp \left(\int_a^b (A_i - G(\beta) \tilde{H}_i(\beta)) d\beta \right). \end{aligned} \quad (57)$$

This implies that $\Gamma_i(t) \succ 0$ for $t \in [0, t_f]$. Therefore, since $\Gamma_i(t)$ is invertible and symmetric, Eq. (55) can be reduced to

$$\Omega_i \tilde{H}'_i + \Omega'_i \tilde{H}'_i = 2G_i \tilde{R}_i. \quad (58)$$

Since the covariance matrix Ω_i is also symmetric,

$$G_i(t) = \Omega_i(t) \tilde{H}_i(t) \tilde{R}_i^{-1}(t) \quad (59)$$

Plugging in this expression on (51) and (50), we get the usual Kalman-Bucy filter equations, which along with the initial conditions $\Omega_i(0) = \Omega_{i,0}$ and $\hat{\phi}_i(0) = E[\phi_i(0)]$, have unique solutions.

APPENDIX B

EXISTENCE OF STEADY STATE COVARIANCE DERIVATIVES

In this appendix, we discuss the existence of the gradients of the steady state covariance matrix. Note that, if in a periodic trajectory $\eta_i(q) = 0 \forall q \in [0, 1]$ (i.e., target i is never visited), the existence of the steady state covariance matrix is not guaranteed by Prop. 2. Obviously, if the steady state covariance does not exist, its derivative will also not exist. This illustrates the fact that the existence $\frac{\partial \bar{\Omega}_i}{\partial \theta_d}$ is not guaranteed. What we show in this appendix is that, under very natural assumptions, the derivative $\frac{\partial \bar{\Omega}_i}{\partial \theta_d}$ exists for the parameters that belong to the interior of the set of parameters that will lead to convergence of the steady state covariance, except for a set of zero measure.

Since here we analyze the behavior of the steady state covariance with respect to parameter variations, we will use a notation that explicitly shows the dependence of the variables with the parameters. For example, $\bar{\Omega}_i$ is a function of q and of the parameters Θ and, hence, it will be denoted as $\bar{\Omega}_i(q; \Theta)$.

We define the set of parameters for which the steady state covariance is guaranteed to exist as:

$$\vartheta = \{ \Theta \mid \eta_i(q, \tilde{\Theta}) > 0$$

for some non-degenerate interval $q \in [a, b]$, (60)

and Ψ as the interior of the set ϑ .

Our goal is to show that, for any $\Theta \in \Psi$, the partial derivatives $\frac{\partial \bar{\Omega}_i(q; \Theta)}{\partial \theta_d}$ exist locally. From Prop. 3, we know that, when this partial derivative exists, it is equal to $\Sigma(q; \Theta)$. We also know that $\Sigma(q; \Theta)$ is well defined for any $\theta \in \Psi$. We now make the following assumption about the regularity of Σ :

Assumption 4. $\Sigma(q; \Theta)$ is locally Riemann integrable for $\Theta \in \Psi$.

In light of Prop. 3, Assumption 3 means that the parameterizations that we consider do not allow for an infinite number of discontinuities of $\Sigma_h(q; \Theta)$ and $\Sigma_{ZI}(q; \Theta)$. Note that, due to the linear nature of their underlying differential equations, $\Sigma_h(q; \Theta)$ and $\Sigma_{ZI}(q; \Theta)$ are bounded for any $\Theta \in \Psi$. Therefore, $\Sigma(q; \Theta)$ is also bounded.

Proposition 5. *Under Assumptions 1, 2 and 3, the partial derivative $\frac{\partial \bar{\Omega}_i(q; \Theta)}{\partial \theta_d}$, $q \in [0, 1]$ and $\Theta \in \Psi$, exists almost everywhere in $[0, 1] \times \Psi$.*

Proof. By construction, we pick two parameter sets Θ_1 and Θ_2 , such that any convex combination of Θ_1 and Θ_2 belongs to Ψ . Additionally, since our goal is to compute the partial derivative with respect to θ_d , we pick Θ_2 such that it differs from Θ_1 only in its d -th coordinate. Since the set Ψ is open, if we pick any $\Theta_1 \in \Psi$, we can always find a Θ_2 that fulfills the aforementioned properties.

We define the function $\Upsilon(q; \Theta_2)$ (which later we will show $\Upsilon(q; \Theta_2) = \bar{\Omega}_i(q; \Theta_2)$) as:

$$\Upsilon(q; \Theta_2) = \bar{\Omega}_i(q; \Theta_1) + \int_0^1 \Sigma(q; \Theta_1 + \xi(\Theta_2 - \Theta_1)) d\xi. \quad (61)$$

Note that, if $\Upsilon(q; \Theta_2) = \bar{\Omega}_i(q; \Theta_2)$ for generic Θ_1, Θ_2 , then $\Sigma(q; \Theta) = \frac{\partial \bar{\Omega}_i(q; \Theta)}{\partial \theta_d}$ almost everywhere, since $\Sigma(q; \Theta)$ plays the role of a partial derivative in Eq. (61).

$\bar{\Omega}_i(q; \Theta_2)$ is uniquely defined by satisfying the differential equation (21) and being periodic with period one. We then show that $\Upsilon(q, \Theta_2)$ also satisfies both of these properties, which imply that indeed $\Upsilon(q, \Theta_2) = \bar{\Omega}_i(q; \Theta_2)$.

First, notice that $\Upsilon(0; \Theta_2) = \Upsilon(1; \Theta_2)$ since $\bar{\Omega}_i(0; \Theta_1) = \bar{\Omega}_i(1; \Theta_1)$ and $\Sigma(0, \Theta) = \Sigma(1, \Theta)$, for any $\Theta \in \Psi$. Also, since $\Sigma(q; \Theta)$ is a solution of (22),

$$\int_0^1 \dot{\Sigma}(q; \Theta_1 + \xi(\Theta_2 - \Theta_1)) d\xi = \dot{\bar{\Omega}}_i(q, \Theta_2) - \dot{\bar{\Omega}}_i(q, \Theta_1). \quad (62)$$

Therefore, taking the derivative of (61) with respect to q and substituting (62), we get

$$\dot{\Upsilon}(q, \Theta_2) = \dot{\bar{\Omega}}_i(q, \Theta_2). \quad (63)$$

Hence we conclude that $\Upsilon(q, \Theta_2) = \bar{\Omega}_i(q; \Theta_2)$, and, as a consequence, $\frac{\partial \bar{\Omega}_i(q; \Theta)}{\partial \theta_d}$ exists almost everywhere in Ψ . Additionally, as already stated in Prop. 3, $\frac{\partial \bar{\Omega}_i(q; \Theta)}{\partial \theta_d} = \Sigma(q, \Theta)$ wherever it exists. \square



Julien M. Hendrickx is professor of mathematical engineering at Universit catholique de Louvain, in the Ecole Polytechnique de Louvain since 2010.

He obtained an engineering degree in applied mathematics (2004) and a PhD in mathematical engineering (2008) from the same university. He has been a visiting researcher at the University of Illinois at Urbana Champaign in 2003-2004, at the National ICT Australia in 2005 and 2006, and at the Massachusetts Institute of Technology in 2006 and 2008. He was a postdoctoral fellow at the Laboratory

for Information and Decision Systems of the Massachusetts Institute of Technology 2009 and 2010, holding postdoctoral fellowships of the F.R.S.-FNRS (Fund for Scientific Research) and of Belgian American Education Foundation. He was also resident scholar at the Center for Information and Systems Engineering (Boston University) in 2018-2019, holding a WBI.World excellence fellowship.

Doctor Hendrickx is the recipient of the 2008 EECI award for the best PhD thesis in Europe in the field of Embedded and Networked Control, and of the Alcatel-Lucent-Bell 2009 award for a PhD thesis on original new concepts or application in the domain of information or communication technologies.



Christos G. Cassandras (F'96) Christos G. Cassandras (F96) is Distinguished Professor of Engineering at Boston University. He is Head of the Division of Systems Engineering, Professor of Electrical and Computer Engineering, and co-founder of Boston University's Center for Information and Systems Engineering (CISE). He received degrees from Yale University, Stanford University, and Harvard University. In 1982-1984 he was with ITP Boston, Inc. where he worked on the design of automated manufacturing systems. In 1984-1996 he was a

faculty member at the Department of Electrical and Computer Engineering, University of Massachusetts/Amherst. He specializes in the areas of discrete event and hybrid systems, cooperative control, stochastic optimization, and computer simulation, with applications to computer and sensor networks, manufacturing systems, and transportation systems. He has published over 450 refereed papers in these areas, and six books. He has guest-edited several technical journal issues and currently serves on several journal Editorial Boards, including Editor of Automatica. In addition to his academic activities, he has worked extensively with industrial organizations on various systems integration projects and the development of decision support software. He has most recently collaborated with The MathWorks, Inc. in the development of the discrete event and hybrid system simulator SimEvents.

Dr. Cassandras was Editor-in-Chief of the IEEE Transactions on Automatic Control from 1998 through 2009 and has also served as Editor for Technical Notes and Correspondence and Associate Editor. He was the 2012 President of the IEEE Control Systems Society (CSS). He has also served as Vice President for Publications and on the Board of Governors of the CSS, as well as on several IEEE committees, and has chaired several conferences. He has been a plenary/keynote speaker at numerous international conferences, including the 2017 IFAC World Congress, the American Control Conference in 2001 and the IEEE Conference on Decision and Control in 2002 and 2016, and has also been an IEEE Distinguished Lecturer.

He is the recipient of several awards, including the 2011 IEEE Control Systems Technology Award, the Distinguished Member Award of the IEEE Control Systems Society (2006), the 1999 Harold Chestnut Prize (IFAC Best Control Engineering Textbook) for Discrete Event Systems: Modeling and Performance Analysis, a 2011 prize and a 2014 prize for the IBM/IEEE Smarter Planet Challenge competition, the 2014 Engineering Distinguished Scholar Award at Boston University, several honorary professorships, a 1991 Lilly Fellowship and a 2012 Kern Fellowship. He is a member of Phi Beta Kappa and Tau Beta Pi. He is also a Fellow of the IEEE and a Fellow of the IFAC.



Samuel C. Pinto (S'19) received the B.S. and M.S. degrees in electronic engineering from Instituto Tecnológico de Aeronáutica, São Paulo, Brazil, in 2017 and 2018, respectively. He is currently a Ph.D. candidate in the department of Mechanical Engineering at Boston University, Boston, MA, USA. His control interests include robotics, multiagent systems and the intersection between estimation and control. Mr. Pinto is a student member of the IEEE Control Systems Society.



Sean B. Andersson (SM'14) received a B.S. degree in engineering and applied physics from Cornell University, Ithaca, NY, USA, in 1994, an M.S. degree in mechanical engineering from Stanford University, Stanford, CA, USA, in 1995, and a Ph.D. degree in electrical and computer engineering from the University of Maryland, College Park, MD, USA, in 2003. He has worked at AlliedSignal Aerospace and Aeroenvironment, Inc. and is currently a Professor of mechanical engineering and of systems engineering with Boston University, Boston, MA, USA. He serves as the director of the Masters of Science in Robotics and Autonomous Systems program. His research interests include systems and control theory with applications in scanning probe microscopy, dynamics in molecular systems, and robotics. He received an NSF CAREER award in 2009, is a senior member of the IEEE, and was an associate editor for the IEEE Transactions on Automatic Control (2014-2018) and for the SIAM Journal on Control and Optimization (2013-2018).