

Decentralized Multi-Agent Active Search and Tracking when Targets Outnumber Agents

Arundhati Banerjee¹ and Jeff Schneider¹

Abstract—Multi-agent multi-target tracking has a wide range of applications, including wildlife patrolling, security surveillance or environment monitoring. Such algorithms often make restrictive assumptions: the number of targets and/or their initial locations may be assumed known, or agents may be pre-assigned to monitor disjoint partitions of the environment, reducing the burden of exploration. This also limits applicability when there are fewer agents than targets, since agents are unable to continuously follow the targets in their fields of view. Multi-agent tracking algorithms additionally assume inter-agent synchronization of observations, or the presence of a central controller to coordinate joint actions. Instead, we focus on the setting of decentralized multi-agent, multi-target, simultaneous active search-and-tracking with asynchronous inter-agent communication. Our proposed algorithm DecSTER uses a sequential monte carlo implementation of the Probability Hypothesis Density filter for posterior inference combined with Thompson sampling for decentralized multi-agent decision making. We compare different action selection policies, focusing on scenarios where targets outnumber agents. In simulation, we demonstrate that DecSTER is robust to unreliable inter-agent communication and outperforms information-greedy baselines in terms of the Optimal Sub-Pattern Assignment (OSPA) metric for different numbers of targets and varying teamsizes.

I. INTRODUCTION

Searching for targets, detecting objects of interest (OOIs), localizing and following them are tasks integral to several robotics applications. Hence one or more of these sub-problems have been widely studied in a number of settings. For example, in informative path planning [1] or simultaneous localization and mapping (SLAM) [2], when the OOIs are fixed, an agent (or robot) adaptively selects actions to detect and localize the targets. Here the environment (or, the target distribution) being stationary, an agent tends to explore unseen parts of its surroundings more than exploit already observed viewpoints. In contrast, when targets are dynamic, the environment is non-stationary. Therefore, agents tracking an unknown number of moving targets should trade-off between exploring the possibly unobserved parts of the environment and exploiting its own posterior estimates to localize the previously detected targets at the current timestep. Unfortunately, prior work in multi-target tracking (MTT) has often assumed that the environment is known and exploration is not of primary concern [3]. Moreover, with multiple agents, existing MTT methods simplify the explore-exploit dilemma by separating search and tracking into sequential tasks where each agent is assigned to track a target as soon as it is found [4]. Another approach is to assign sub-teams

for executing these tasks separately [5]. Further, the majority of these multi-agent multi-target tracking (MAMTT) algorithms require either a central controller to coordinate joint tracking actions, facilitate target hand-offs among agents or they depend on synchronized inter-agent communication for distributed inference and decision making. Unfortunately, such conditions may not be feasible in practice: the environment may be unstructured and unknown, OOIs may need to be simultaneously detected and localized (i.e. without a separation between search and tracking phases), there may not be enough agents to continuously monitor all the targets and unreliable communication channels may hinder inter-agent synchronization at each timestep.

In this work, we aim to develop a more practical approach to the MAMTT setup. In particular, we focus on the setting where agents are outnumbered by targets, so that the multi-agent team is unable to continuously cover all targets in their fields of view. The number of targets and their initial locations are unknown. Therefore, agents need to interact with the environment to collect observations by adaptively making explore-exploit decisions. It is not feasible to continuously track all targets, but our goal is to produce an estimate of the number and positions of all targets with time. We assume that there is no central controller, and agents share their observations asynchronously with teammates, whenever possible. In other words, agents do not wait to receive observations, action selection policy or environment belief information from their teammates and can continue their online decision making when communication is unreliable or even unavailable.

Contributions. We propose a decentralized and asynchronous multi-agent algorithm, called DecSTER (*Decentralized Multi-Agent Active Search-and-Tracking without continuous coverage*) for simultaneous multi-target active search and tracking without continuously following the targets. In simulation, we compare a number of common decision making objectives from the tracking literature after adapting them to our simultaneous search-and-tracking setting. Our results show that DecSTER outperforms competitive baselines with different teamsizes and different target distributions in terms of the OSPA tracking performance.

II. RELATED WORKS

Target detection and tracking are both widely studied problems, typically considered as distinct tasks in various applications like search and rescue [6], security surveillance [7], computer games [8], etc. [3] is a detailed survey of the

¹ School of Computer Science, Carnegie Mellon University, Pittsburgh, USA. {arundhat, schneide}@cs.cmu.edu

many different approaches and taxonomy used in robotics and related fields for such scenarios. Here, we discuss some of the most relevant prior work in our context.

The single target state is commonly modeled assuming linear dynamics with additive Gaussian noise, using the Kalman filter [9] or using non-parametric particle filters [10]. In multi-target settings, alternative approaches like Multiple Hypothesis Tracker (MHT) [11], Joint Probabilistic Data Association (JPDA) [12] and Probability Hypothesis Density (PHD) filter [13] have been proposed, all of which differ in how they perform data association [14]. The PHD filter is particularly suited when unique identities for each target are not required, for example, in search and rescue tasks, where agents should detect and localize *all* survivors. In this work, we build on the Sequential Monte Carlo (SMC) formulation of the PHD filter presented in [15].

Prior work in MAMTT algorithms primarily considers centralized or distributed settings, the latter still necessitating synchronized communication among subgroups of agents at each time step. Coupled with a PHD filter, some of the common action selection methods previously proposed for tracking include mutual information and expected count based objectives [16], Renyi divergence maximization [4] and Lloyd’s algorithm for Voronoi-cell based control [17], [18]. [4] discusses the benefits of a simultaneous search-and-tracking algorithm but their proposed method requires that agents transition from searching to tracking upon target detection, foregoing further exploration. In similar spirit, [19] proposes solving an information gain based multi-objective optimization problem over a finite planning horizon to decide the best (greedy) joint action for a centralized search-and-tracking task. In contrast with these deterministic objectives, [20] demonstrates the superior performance of stochastic optimization methods like Particle Swarm Optimization (PSO) and Simulated Annealing (SA) for better coverage and localization in such settings. Unfortunately, none of these prior approaches are applied in the decentralized and asynchronous multi-agent setup when agents are unable to support continuous target coverage.

Recently, some learning based approaches have been proposed for tracking. [21] considers the single-agent setting, assuming a known number of (one or two) targets which only start moving after first observed by the agent. [22], [23] both propose GNN-based algorithms, trained by imitation from a centralized expert, and deployed in the distributed inference and decision making with synchronized communication setup. While [22] does not consider dynamic targets, [23] simplifies the problem to deterministic optimal control over a fixed horizon, differing from our more complex setting.

III. PROBLEM SETUP

Consider a team of J UAVs tasked with search and tracking of an unknown number of moving targets $\{1, \dots, k\}$ in a 2-dimensional (2D) region \mathcal{G} of length n_l and width n_w (Fig. 1a). The agents can self-localize in a global coordinate system. They are equipped with noisy sensors that provide location 2D coordinates of possible targets in their current

field of view (FOV). The targets can move in any direction in the search space at different speeds. We assume that agents typically move faster than targets. Each agent’s FOV includes a contiguous rectangular region of the search space, and agents may choose to observe a wider (smaller) area at a greater (lower) vertical height but with more (less) observation noise. We therefore consider a hierarchical region sensing action space for each agent. Agents can communicate asynchronously with their teammates. Over time T , agents observe different parts of the search space to detect and track all targets in the environment.




Fig. 1: **Problem setup.** (a) Agents sense different regions of the search space at different vertical heights, receiving noisy 2D location coordinates of the possible targets in their field of view, along with false positive measurements. The targets shown as black crosses move in the search space with different velocities shown by the red arrows. (b) The line at the top indicates the target’s continuous motion with time. In our asynchronous multi-agent setup, agents can collect observations without waiting for their teammates whereas in the synchronous setting, the solid boxes indicate the agents’ idle wait times.

A. Target and Measurement Representations

The state of each target is denoted by $\mathbf{x} = [l_x, l_y, v_x, v_y]^T$, where 2D coordinates $(l_x, l_y) \in [0, n_w] \times [0, n_l]$ and velocity $\mathbf{v} = [v_x, v_y]^T$, $\|\mathbf{v}\|_2 \leq v_{\max}$. Since both the number of targets and their true locations are unknown, we follow the Random Finite Set (RFS) representation for the multi-target state space $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{|\mathcal{X}|}\}$, where $|\mathcal{X}|$ follows a Poisson distribution and for a given cardinality, the set elements are sampled i.i.d from a uniform distribution [24]. Following prior work, we use the Probability Hypothesis Density (PHD) filter [13] to maintain a belief over the RFS \mathcal{X} .

The PHD $\nu(x)$ is the first statistical moment of a distribution over RFSs. In this case, it is a density over the state space of targets so that for any region $E \subseteq \mathcal{G}$, the expected cardinality of the target RFS in that region is $\int_E \nu(x) dx$. The PHD filter tracks the evolving target density over the search space using models of target motion and measurements gathered by the agents. The measurements \mathcal{Z} are also modeled as a (Poisson) RFS, as are clutter $\kappa(z)$ (false positives) and target births $b(x)$.

B. Sensing model

An agent with pose $\mathbf{q} = [q_x, q_y]^T$ executes a sensing action $\mathbf{a}_{\mathbf{q}}$, receiving a measurement set $\mathcal{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_m\}$. Any target \mathbf{x} within the agent’s FOV may generate a measurement \mathbf{z} , with a probability of detection $p_d(\mathbf{x}|\mathbf{q})$. In this

work, we assume a constant $p_d(\cdot)$ when the target \mathbf{x} is within the FOV at \mathbf{q} , and 0 otherwise. The agent follows a linear sensing model with additive i.i.d white noise: $\mathbf{z} = h(\mathbf{x}) + \boldsymbol{\omega}$, where $h(\mathbf{x}) = [l_x, l_y]^T$ and $\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{0}, \sigma_h^2 \mathbf{I})$. Additionally, \mathcal{Z} includes i.i.d false positives with clutter rate λ_q .

The (noisy) target dynamics from state $\boldsymbol{\xi}$ to \mathbf{x} is captured by the target motion model $f(\mathbf{x}|\boldsymbol{\xi})$. The survival probability $p_s(\mathbf{x})$ denotes the target's chances of persisting over successive time steps. The PHD filter formulates the following steps to propagate the posterior density over target states.¹

$$\text{Prediction: } \bar{\nu}_t(\mathbf{x}) = b(\mathbf{x}) + \int_E f(\mathbf{x}|\boldsymbol{\xi}) p_s(\boldsymbol{\xi}) \nu_{t-1}(\boldsymbol{\xi}) d\boldsymbol{\xi} \quad (1)$$

$$\text{Update: } \nu_t(\mathbf{x}) = (1 - p_d(\mathbf{x}|\mathbf{q})) \bar{\nu}_t(\mathbf{x}) + \sum_{\mathbf{z} \in \mathcal{Z}_t} \frac{\psi_{\mathbf{z}, \mathbf{q}}(\mathbf{x}) \bar{\nu}_t(\mathbf{x})}{\eta_{\mathbf{z}}(\bar{\nu}_t)} \quad (2)$$

$$\eta_{\mathbf{z}}(\nu) = \kappa(\mathbf{z}|\mathbf{q}) + \int_E \psi_{\mathbf{z}, \mathbf{q}}(\mathbf{x}) \nu(\mathbf{x}) d\mathbf{x} \quad (3)$$

$$\psi_{\mathbf{z}, \mathbf{q}}(\mathbf{x}) = g(\mathbf{z}|\mathbf{x}, \mathbf{q}) p_d(\mathbf{x}|\mathbf{q}) \quad (4)$$

Here, $\psi_{\mathbf{z}, \mathbf{q}}(\mathbf{x})$ is the probability that the agent at \mathbf{q} receives the measurement \mathbf{z} from a target \mathbf{x} and $g(\mathbf{z}|\mathbf{x}, \mathbf{q})$ is the measurement likelihood model. The PHD filter can handle appearing and disappearing targets by defining an appropriate birth density $b(\mathbf{x})$ over the search space, but in our experiments the number of ground truth targets k is fixed.

Remark 1. As [13] explains, using the first order moment to approximate the multi-target belief and deriving recursive PHD update equations to approximate the evolving posterior is justifiable when both sensor covariances and false alarm densities are small, so that (the distribution of) observations from true targets are centered around target states with negligible spread and there is lower noise due to false alarms. Besides, in our SMC-PHD implementation following [15], we also ensure we avoid particle impoverishment during update and propagation of the posterior density estimate.

Decentralized Asynchronous Multi-Agent Setup. In our setup, each agent $j \in \{1, \dots, J\}$ maintains its PHD estimate and decides its next sensing action in a decentralized manner. There is no central controller and inter-agent communication is asynchronous (Fig. 1b). Note that agents are not independent, since they share their own observations and incorporate the measurements received from their teammates in subsequent PHD filter update steps. Our setting is also different from the distributed computation setup in prior work [17] where each agent completes a part of the centralized update step and relies on inter-agent synchronized communication to maintain a global PHD estimate across all agents.

Since targets are in continuous motion, our agents must be able to deal with the uncertainty arising from observation noise as well as due to asynchronous communication of time-dependant observations in their posterior PHD updates. In order to enable time-ordered assimilation of received observations by all agents, we assume that any agent j communicates the tuple $(t, \mathbf{a}_t^j, \mathcal{Z}_t^j)$ where \mathbf{a}_t^j and \mathcal{Z}_t^j are

respectively the agent's sensing action and measurement set at time t .

IV. OUR APPROACH

We now describe our algorithm DecSTER for multi-agent active search and tracking without continuous coverage.

Notation. Agent j at time t has a history of available actions and observations $\mathbb{D}_t^j = \{(t', \mathbf{a}_{t'}^j, \mathcal{Z}_{t'}^j)\}_{t' < t, j' \in \{1, \dots, J\}}$. Using \mathbb{D}_t^j , it computes the PHD ν_t^j (Eqs. (1) and (2)) over the target RFS. In our SMC-PHD implementation, $\nu_t^j = \{(w_{t,1}^j, \mathbf{x}_{t,1}^j), \dots, (w_{t,\rho}^j, \mathbf{x}_{t,\rho}^j)\}$ where $\mathbf{x}_{t,1}^j, \dots, \mathbf{x}_{t,\rho}^j$ are the ρ particles with weights $w_{t,1}^j, \dots, w_{t,\rho}^j$.

The SMC-PHD filter propagation steps follow from [15]. In our decentralized setup, each agent maintains its own posterior PHD ν_t^j . Next, we will describe the decision making step executed by agent j at time t .

Thompson sampling for decision making. Prior work in multi-agent active search with static targets has demonstrated the effectiveness of Thompson sampling (TS) as a decentralized decision making algorithm, both in theory [25] and in practice [26]. TS ensures stochasticity in decision making by sampling different plausible realizations of the ground truth from the posterior belief and selecting the best action to maximize the desired reward for a particular sample. The uncertainty in the agent's belief over the state space is reflected in the posterior samples, which makes TS suitable for driving exploration and exploitation. [5] couples a TS-based active search strategy with the deterministic Lloyd's algorithm for tracking, but in their setup, agents are pre-assigned to only one of search or tracking tasks. Instead, we propose a TS strategy to enable agents to naturally trade-off exploratory sensing actions that might discover undetected targets, with exploitative sensing actions that help localize and track the previously detected dynamic targets in our simultaneous search-and-tracking setting.

To the best of our knowledge, prior work has not studied the problem of TS in a continuous (not discretized) search space with a PHD posterior. This is challenging because the PHD is not a distribution and does not include second order uncertainty information, whereas TS is typically applied in the Bayesian setting with the samples drawn from a posterior distribution for which both first and second order moment estimates are available [27]. Prior work in [28] has proposed particle Thompson sampling (PTS) and regenerative PTS (RPTS) algorithms for particle filters where particles are sampled proportional to their weights. Therefore, we adopt a similar principle in our first proposed TS strategy for the SMC-PHD posterior, called TS-PHD-I (Algorithm 1).

Algorithm 1 TS-PHD-I

- 1: **Input:** PHD $\nu = \{(w_1, \mathbf{x}_1), \dots, (w_\rho, \mathbf{x}_\rho)\}$.
 - 2: Sample $\tilde{\rho}$ particles $\{\tilde{\mathbf{x}}_i\}_{i=1}^{\tilde{\rho}}$ from ν , proportional to the weights $\{w_1, \dots, w_\rho\}$
 - 3: Cluster the $\tilde{\rho}$ particles using k-means with $\tilde{n} = \sum_{i=1}^{\tilde{\rho}} w_i$ centroids $(\tilde{\mathcal{X}} = \{\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_{\tilde{n}}\})$ which form the TS
-

¹For a detailed understanding of the PHD filter, please refer to [24].

We note that this method has drawbacks. It tends to sample more particles from the regions in the PHD where the agent already estimates targets might be present. The samples drawn are thus more likely to be biased against regions of the target state space where the agent might be less certain about its observations owing to false positives or missed detections. Furthermore, this method does a poor job of modeling the uncertainty about the number of true targets.

To address the drawbacks of Algorithm 1, we now describe a second proposed approach to Thompson sampling from our SMC-PHD posterior (Algorithm 2). Recall that the expected cardinality of the target RFS \mathcal{X} over a region $E \subseteq \mathcal{G}$ is given by $\hat{n} = \int_E \nu(\mathbf{x}) d\mathbf{x}$. In case of the SMC-PHD representation, $\hat{n} = \sum_i w_i, \forall \mathbf{x}_i \in E$ [15] i.e. the sum of particle weights of the SMC-PHD in the region E is the expected cardinality of \mathcal{X} in that region. Further, [13] shows that the PHD is the best Poisson approximation of the multitarget posterior in terms of KL divergence. We therefore draw a sample \tilde{n} of the cardinality of the target RFS from a Poisson distribution with mean $\hat{n} = \sum_i w_i$ (i.e. $\tilde{n} \sim \text{Poisson}(\hat{n})$). Then we sample \tilde{n} locations of the possible targets by drawing from a mixture of already estimated target locations in the PHD and some locations drawn uniformly at random over the search space. These \tilde{n} particles $\tilde{\mathcal{X}} = \{\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_{\tilde{n}}\}$ form our TS.

Algorithm 2 TS-PHD-II

- 1: **Input:** PHD $\nu = \{(w_1, \mathbf{x}_1), \dots, (w_\rho, \mathbf{x}_\rho)\}$. $\hat{n}_{\mathcal{G}} = \sum_{i=1}^{\rho} w_i$. Target location estimates $\hat{\mathcal{X}} = \{\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_{\hat{n}_{\mathcal{G}}}\}$ from ν .
 - 2: Sample $\tilde{n} \sim \text{Poisson}(\hat{n}_{\mathcal{G}})$
 - 3: Sample uniformly random target locations $\tilde{\mathcal{X}}_R$
 - 4: Sample \tilde{n} target locations ($\tilde{\mathcal{X}}$) from $\hat{\mathcal{X}} \cup \tilde{\mathcal{X}}_R$ as the TS
-

Objective. The Optimal Sub-Pattern Assignment (OSPA) metric is typically used in the MTT literature for evaluating the tracking performance of an algorithm and is defined as the error between two sets. Given sets \mathcal{X} and \mathcal{Y} , where $|\mathcal{X}| = m \leq |\mathcal{Y}| = n$ without loss of generality,

$$\text{OSPA}(\mathcal{X}, \mathcal{Y}) = \left(\frac{1}{n} \min_{\pi \in \Pi_n} \sum_{i=1}^m d_c(x_i, y_{\pi(i)})^p + c^p(n-m) \right)^{\frac{1}{p}}$$

where c is the cut-off distance, $d_c(x, y) = \min(c, \|x - y\|)$ and Π_n is the set of all permutations of the set $\{1, \dots, n\}$. The distance error component of the OSPA computes the minimum cost assignment between \mathcal{X} and \mathcal{Y} , such that $x_i \in \mathcal{X}$ is matched to $y_{i'} \in \mathcal{Y}$ only when they are within a distance c of each other.

Given a true target set $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{|\mathcal{X}|}\}$ and an estimated set $\mathcal{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_{|\mathcal{Y}|}\}$ of possible target locations, our goal is to minimize $\text{OSPA}(\mathcal{X}, \mathcal{Y})$. Since the ground truth \mathcal{X} is unknown, each agent j instead draws a TS $\tilde{\mathcal{X}}_t^j$ from the predicted PHD $\tilde{\nu}_{t+1}^j$. Assuming observations are generated by $\tilde{\mathcal{X}}_t^j$ for any action \mathbf{a} and \mathcal{Y}_t^j is the estimated target set following the PHD filter update, agent j then selects:

$$\mathbf{a}_t^j = \arg \min_{\mathbf{a}} \mathbb{E}_{\mathcal{Y}_t^j | \tilde{\mathcal{X}}_t^j, \mathbf{a}} [\text{OSPA}(\tilde{\mathcal{X}}_t^j, \mathcal{Y}_t^j)] \quad (5)$$

Algorithm 3 outlines our proposed algorithm, called DecSTER. In our decentralized and asynchronous multi-agent setting, each agent individually runs DecSTER with its own sampled $\tilde{\mathcal{X}}_t^j$. Hence the stochasticity in the sampling procedure enables agents to make decentralized explore-exploit decisions for simultaneous search-and-tracking in their action space.

Algorithm 3 DecSTER for agent j at time t

- 1: **Input:** PHD $\nu_t^j = \{(w_1^j, \mathbf{x}_1^j), \dots, (w_\rho^j, \mathbf{x}_\rho^j)\}$
 - 2: Compute predicted PHD $\tilde{\nu}_{t+1}^j$ (Eq. (1)).
 - 3: Draw TS $\tilde{\mathcal{X}}_t^j \sim \tilde{\nu}_{t+1}^j$.
 - 4: Assuming pseudo-measurements at $\tilde{\mathcal{X}}_t^j$, estimate expected target set \mathcal{Y}_t^j and select action \mathbf{a}_t^j (Eq. (5)).
 - 5: Observe \mathcal{Z}_t^j . Update PHD ν_{t+1}^j (Eq. (2)).
 - 6: Estimate target set $\hat{\mathcal{X}}_{t+1}^j$ from ν_{t+1}^j [15].
 - 7: Asynchronously communicate $(t, \mathbf{a}_t^j, \mathcal{Z}_t^j)$ with team.
-

Remark 2. Prior work in search-and-tracking [4], [5] tends to separate the search and tracking phases of the task, and maintains either a visit count or dynamic occupancy grid to compute the action selection objective during the exploration phase. Such methods scale poorly with the size of the environment since agents need to maintain a discretization over the search space [19]. In contrast, our SMC inference for multi-target belief is parallelizable over particles in the posterior PHD, while our TS-based decision making is scalable with increasing teams size J .

V. RESULTS

We now describe our experimental setup. Consider a 2D search space with dimensions $n_l \times n_w = 16 \times 16$. There are k targets moving in this region, whose starting locations and velocities are chosen uniformly at random, such that $v_{\max} = 0.1$. A team of J agents are tasked with search-and-tracking of all the targets over $T = 150$ steps. The agents' action space \mathcal{A} consists of hierarchical region sensing actions of width $1 \times 1, 2 \times 2, 4 \times 4$ and 8×8 , $|\mathcal{A}| = 340$. Since actions with larger FOV receive noisier observations, we vary the false positive (clutter) rate as $\lambda \in \{0.005, 0.04, 1, 5\}$ for action widths 1, 2, 4 and 8 respectively. The survival probability in the PHD filter is set at $p_s = 1$ and the detection probability $p_d = 0.9$ for targets in the agent's FOV. In our SMC-PHD implementation following [15], we initialize 100 new (birth) particles per observation and re-sample 1000 particles per estimated target, following the low variance sampling method in [29]. We choose $\bar{\rho} = 100$ (Algorithm 1).

The agents assume the target motion model $\mathbf{x}_{t+1} = \mathbf{F}\mathbf{x}_t +$

$$\epsilon, \text{ where } \mathbf{F} = \begin{bmatrix} 1 & 0 & \Delta T & 0 \\ 0 & 1 & 0 & \Delta T \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \Delta T = 1 \text{ and } \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \mathbf{Q} = \begin{bmatrix} 0.03 & 0 & 0.05 & 0 \\ 0 & 0.03 & 0 & 0.05 \\ 0.05 & 0 & 0.1 & 0 \\ 0 & 0.05 & 0 & 0.1 \end{bmatrix}. \text{ The sensing}$$

model is $\mathbf{z} = \mathbf{H}\mathbf{x} + \boldsymbol{\omega}$, where $\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$ and $\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, $\sigma = 0.1$. \mathbf{I} is the identity matrix. For the OSPA metric, we set $c = 2$ and $p = 1$.

Remark 3. Our experimental setup is intended to illustrate the abilities of DecSTER for decentralized and asynchronous multi-agent multi-target search and tracking. The action space is chosen so that there is a non-trivial explore-exploit decision to be made by the agents. The maximum target speed is such that targets may cover a considerable distance in the search space over T steps.

In the following experiments, we measure performance in terms of the average OSPA for the entire team of agents. The plots show mean across 10 random trials with the shaded regions indicating standard error. Each trial differs in the initialization of the target locations and their velocities.

A. Comparing TS-PHD-I with TS-PHD-II

First, we compare the performance of DecSTER using the two proposed approaches for Thompson sampling from a PHD posterior. Fig. 2 compares OSPA with number of measurements per agent, for DecSTER-I and DecSTER-II using TS-PHD-I and TS-PHD-II respectively in Line 3 of Algorithm 3. We observe that decision making with TS-PHD-II consistently outperforms that with TS-PHD-I. Since TS-PHD-II samples both the cardinality and locations of the target RFS from the PHD, the samples for different agents are sufficiently diverse to capture the uncertainty regarding the true multi-target ground truth. In contrast, the samples from TS-PHD-I are generally clustered around the agent's current estimate of target locations. We also empirically observed an improvement in the OSPA performance with TS-PHD-II when, for a particular sampled cardinality \tilde{n} , an agent considers multiple samples of \tilde{n} target locations and averages the reward in Eq. (5) over them. Our results using DecSTER-II consider 10 such samples per action selection step for any agent j . We also allow DecSTER-I to similarly consider averaging over multiple samples, but this does not improve the sample diversity and does not lead to performance improvement.

Fig. 2 further demonstrates the scalability of TS in the decentralized multi-agent active search-and-tracking setting. When teams size increases n times, agents achieve similar OSPA with n times fewer measurements per agent [25].

B. Baseline comparisons

We compare DecSTER-I and DecSTER-II with the following baselines. Note that all of them use the same PHD filter inference method, but differ in the action selection policy.

1) RANDOM. Each agent j selects its next sensing action uniformly at random. **2) RENYI.** At t , agent j computes the predicted PHD $\bar{\nu}_{t+1}^j$ and generates a (pseudo) measurement set $\tilde{\mathcal{Z}}_t^j$ for any action $\mathbf{a} \in \mathcal{A}$ assuming the estimated target set $\tilde{\mathcal{X}}_t^j$ from ν_t^j as ground truth. It then selects the action \mathbf{a}_t^j that maximizes the Renyi divergence (with $\alpha = 0.5$) between




Fig. 2: **Comparing the proposed TS methods.** DecSTER-II using TS-PHD-II consistently outperforms DecSTER-I using TS-PHD-I. Increasing team size J reduces the number of measurements per agent required to achieve similar OSPA.

$\bar{\nu}_{t+1}^j$ and its expected updated PHD ν_{t+1}^j (Eq. (2)). With the SMC-PHD formulation, the Renyi divergence is [30]:

$$\sum_{i=1}^{\rho} \bar{w}_i + \frac{\alpha}{1-\alpha} \sum_{i=1}^{\rho} w'_i - \frac{1}{1-\alpha} \sum_{i=1}^{\rho} w'_i \bar{w}_i^{1-\alpha} \quad (6)$$

where \bar{w}_i and w'_i are the weights of the particle i in $\bar{\nu}_{t+1}^j$ and ν_{t+1}^j respectively. **3) TS-RENYI.** We modify RENYI to use $\tilde{\mathcal{X}}_t^j \sim \bar{\nu}_{t+1}^j$ (with TS-PHD-II) for computing the (pseudo) measurement set $\tilde{\mathcal{Z}}_t^j$ and the updated weights w'_i .

Fig. 3 shows that our proposed method outperforms all the baselines for different number of targets k and team sizes J . We observe that RENYI agents are information-greedy, therefore the lack of stochasticity in their decision making objective leads different agents to select the same action in the decentralized asynchronous multi-agent setting. Moreover, the computation in Eq. (6) depends only on the particles in $\bar{\nu}_{t+1}^j$ and does not account for previously undetected targets. This highlights the drawback of using Renyi divergence as an optimization objective for explore-exploit decisions in the search-and-tracking setting, in contrast with its success in the tracking setting where exploration is not a concern [4]. This motivated us to propose the TS-RENYI baseline in order to encourage exploration with samples drawn from TS-PHD-II. We observe that TS-RENYI still does not perform noticeably better than RENYI. This is because the weights of the particles in the SMC-PHD filter relate to the expected cardinality of the target set, therefore Eq. (6) does not account for any measure of the distance error between $\hat{\mathcal{X}}_t^j$ (or $\tilde{\mathcal{X}}_t^j$) and the estimate $\hat{\mathcal{X}}_{t+1}^j$ from ν_{t+1}^j . In contrast, the OSPA objective accounts for both localization error as well as cardinality error in the estimated target set. Thus we observe that our algorithm DecSTER-I is competitive with or outperforms random sensing and information-greedy baselines, and DecSTER-II consistently achieves the lowest OSPA among all with the same number of measurements per agent. Based on these results, we consider DecSTER-II as our best approach in this setting, labeled DecSTER in the following experiments.

DecSTER vs. DecSTER-C. Motivated by prior work [16] that showed the effectiveness of maximizing the expected number of target detections for action selection in multi-




Fig. 3: **Baseline comparisons.** For different numbers of targets and with fewer agents than targets, DecSTER with TS-PHD-II (denoted DecSTER-II) outperforms random sensing (RANDOM) and information greedy baselines (RENYI, TS-RENYI) by achieving a lower OSPA for the same number of measurements per agent.

agent tracking, we introduce the DecSTER-C baseline where agents select actions minimizing only the cardinality error term of the OSPA. Unlike DecSTER, each agent with DecSTER-C draws multiple samples of cardinality $\hat{n} \sim \text{Poisson}(\hat{n})$ to consider the uncertainty about the real number of targets in its objective. Fig. 4 shows that DecSTER still outperforms DecSTER-C, indicating that jointly considering detection and localization error in the decision making objective is more advantageous for TS-guided action selection in our search-and-tracking setting. This further supports our earlier observations regarding the drawbacks of using the particle weight based Renyi divergence optimization objective in this setting.




Fig. 4: **DecSTER vs. DecSTER-C.** DecSTER-C optimizes only for the cardinality error in the OSPA objective. DecSTER outperforms DecSTER-C indicating the advantage of jointly optimizing detection and localization errors with TS-guided explore-exploit decisions.

C. Robustness to communication delays

Multi-agent systems benefit from leveraging observations shared by their teammates. Agents in our decentralized and asynchronous multi-agent setting benefit from any information shared by their teammates but they can continue searching for and tracking targets without waiting for such communication. Therefore, we now analyze the robustness of DecSTER under unreliable inter-agent communication. In simulation, we consider each agent chooses to communicate its own observation at time t , along with any prior observations it had not shared with its teammates, with a probability $p \in \{0.05, 0.25, 0.50, 0.75, 1\}$. The $p = 1$ setting corresponds to our description and analysis of DecSTER

in Fig. 3. We observe a graceful decay in the OSPA performance with decreasing rates of inter-agent communication in Fig. 5, both when targets outnumber agents and vice versa. Compared to prior work in the centralized or distributed multi-agent tracking setting [3], DecSTER does not depend on synchronized communication within the team, thus agents can adapt and continue their search-and-tracking tasks even when communication is unreliable or unavailable.




Fig. 5: **Robustness to unreliable communication.** When agents communicate their actions and observations with decreasing probability p , DecSTER experiences a graceful deterioration in OSPA performance and agents require increasingly more measurements to estimate the number and locations of true targets in the search space.

VI. CONCLUSION

We introduce DecSTER, a novel decentralized and asynchronous algorithm for multi-agent multi-target active search-and-tracking that relaxes the restrictive assumption of requiring continuous target coverage.² In simulation, DecSTER outperforms competitive baselines that greedily optimize for information gain or expected target detections. A key contribution is adapting TS to effectively drive exploration and exploitation using the SMC-PHD filter. Future work includes theoretical analysis of the proposed TS methods and learning improved models of environment uncertainty for non-stationary multi-target tracking. Building on recent success of TS-based multi-agent active search [26], we also aim to validate DecSTER's performance when deployed on teams of physical robots in the real world.

²Code and supplementary video: <https://sites.google.com/view/decster-icra24>.

REFERENCES

- [1] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for uav-based terrain monitoring," *Autonomous Robots*, vol. 44, pp. 889–911, 2020.
- [2] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Transactions on Robotics*, 2023.
- [3] C. Robin and S. Lacroix, "Multi-robot target detection and tracking: taxonomy and survey," *Autonomous Robots*, vol. 40, pp. 729–760, 2016.
- [4] S. Papaioannou, P. Kolios, T. Theocharides, C. G. Panayiotou, and M. M. Polycarpou, "A cooperative multiagent probabilistic framework for search and track missions," *IEEE Transactions on Control of Network Systems*, vol. 8, no. 2, pp. 847–858, 2020.
- [5] J. Chen and P. Dames, "Active multi-target search using distributed thompson sampling," 2022.
- [6] R. R. Murphy, "Human-robot interaction in rescue robotics," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 34, no. 2, pp. 138–153, 2004.
- [7] L. Doitsidis, S. Weiss, A. Renzaglia, M. W. Achtelik, E. Kosmatopoulos, R. Siegwart, and D. Scaramuzza, "Optimal surveillance coverage for teams of micro aerial vehicles in gps-denied environments using onboard vision," *Autonomous Robots*, vol. 33, pp. 173–188, 2012.
- [8] T. Oskam, R. W. Sumner, N. Thuerey, and M. Gross, "Visibility transition planning for dynamic camera control," in *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2009, pp. 55–65.
- [9] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.
- [10] A. Doucet, N. De Freitas, N. J. Gordon, *et al.*, *Sequential Monte Carlo methods in practice*. Springer, 2001, vol. 1, no. 2.
- [11] S. S. Blackman, "Multiple hypothesis tracking for multiple target tracking," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 5–18, 2004.
- [12] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE journal of Oceanic Engineering*, vol. 8, no. 3, pp. 173–184, 1983.
- [13] R. P. Mahler, "Multitarget bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic systems*, vol. 39, no. 4, pp. 1152–1178, 2003.
- [14] L. D. Stone, R. L. Streit, T. L. Corwin, and K. L. Bell, *Bayesian multiple target tracking*. Artech House, 2013.
- [15] B. Ristic, D. Clark, and B.-N. Vo, "Improved smc implementation of the phd filter," in *2010 13th International Conference on Information Fusion*. IEEE, 2010, pp. 1–8.
- [16] P. Dames, P. Tokekar, and V. Kumar, "Detecting, localizing, and tracking an unknown number of moving targets using a team of mobile robots," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1540–1553, 2017.
- [17] P. M. Dames, "Distributed multi-target search and tracking using the phd filter," *Autonomous robots*, vol. 44, no. 3-4, pp. 673–689, 2020.
- [18] J. Chen and P. Dames, "Distributed multi-target search and tracking using a coordinated team of ground and aerial robots," in *Robotics science and systems*, 2020.
- [19] H. Van Nguyen, B.-N. Vo, B.-T. Vo, H. Rezaatofghi, and D. C. Ranasinghe, "Multi-objective multi-agent planning for discovering and tracking multiple mobile objects," *arXiv preprint arXiv:2203.04551*, 2022.
- [20] P. Xin and P. Dames, "Comparing stochastic optimization methods for multi-robot, multi-target tracking," in *International Symposium on Distributed and Autonomous Systems*, 2022.
- [21] H. Jeong, H. Hassani, M. Morari, D. D. Lee, and G. J. Pappas, "Deep reinforcement learning for active target tracking," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1825–1831.
- [22] L. Zhou, V. D. Sharma, Q. Li, A. Prorok, A. Ribeiro, P. Tokekar, and V. Kumar, "Graph neural networks for decentralized multi-robot target tracking," in *2022 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2022, pp. 195–202.
- [23] M. Tzes, N. Bousias, E. Chatzipantazis, and G. J. Pappas, "Graph neural networks for multi-robot active information acquisition," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3497–3503.
- [24] R. P. Mahler, *Advances in statistical multisource-multitarget information fusion*. Artech House, 2014.
- [25] R. Ghods, A. Banerjee, and J. Schneider, "Decentralized multi-agent active search for sparse signals," in *Uncertainty in Artificial Intelligence*. PMLR, 2021, pp. 696–706.
- [26] N. A. Bakshi, T. Gupta, R. Ghods, and J. Schneider, "Guts: Generalized uncertainty-aware thompson sampling for multi-agent active search," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 7735–7741.
- [27] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, Z. Wen, *et al.*, "A tutorial on thompson sampling," *Foundations and Trends® in Machine Learning*, vol. 11, no. 1, pp. 1–96, 2018.
- [28] Z. Zhou, B. Hajek, N. Choi, and A. Walid, "Regenerative particle thompson sampling," *arXiv preprint arXiv:2203.08082*, 2022.
- [29] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.
- [30] B. Ristic, B.-N. Vo, and D. Clark, "A note on the reward function for phd filters with sensor control," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 2, pp. 1521–1529, 2011.