2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 25-29, 2020, Las Vegas, NV, USA (Virtual)

Adaptive Partitioning for Coordinated Multi-agent Perimeter Defense

Douglas G. Macharet, Austin K. Chen, Daigo Shishika, George J. Pappas, and Vijay Kumar

***Abstract*— Multi-Robot Systems have been recently employed in different applications and have advantages over single-robot systems, such as increased robustness and task performance efficiency. We consider such assemblies specifically in the scenario of perimeter defense, where the task is to defend a circular perimeter by intercepting radially approaching targets. Possible intruders appear randomly at a fixed distance from the perimeter and with azimuthal location determined by some unknown probability density. Coordination among multiple defenders is a complex combinatorial optimization problem. In this work, we focus on the following two aspects: (i) estimating the probability density that describes the direction from which the next intruders are going to arrive, and (ii) partitioning of the space so that the defenders focus on capturing a disjoint subset of intruders. Results show that the proposed strategy increases the number of captures over a naive baseline strategy, especially in scenarios with non-uniform spatial distributions of intruder arrival. The proposed approach is also efficient and able to quickly adapt to time-varying intruder distributions.**

I. Introduction

The use of autonomous robots in a vast range of tasks such as exploration, environmental monitoring, and surveillance has significantly increased in recent years. In this context, Multi-Robot Systems (MRS) have several advantages over single-robot systems, e.g. increased robustness and task performance efficiency [1]. Here we consider the use of such systems in the problem of defending a perimeter from approaching intruders.

This *perimeter defense* scenario has been previously formulated as a pursuit-evasion game [2], [3], [4], where both the defender and intruder team strategies were considered. Under the assumption that the defenders move along the perimeter of an enclosed compound, the intruders can ensure that one defender cannot make multiple captures [2].

In this paper, we consider the case where intruders move radially towards the perimeter. This is no longer a game because the intruder strategy is fixed. Furthermore, intruders are intermittently inserted in the environment accordingly to an unknown random distribution, which demands consecutive captures by the defenders. Finding the optimal defender team strategy is challenging because the coupled problems of task allocation (which intruder should be captured by whom) and the routing (in what order intruders should be captured) must be solved.

The authors are with the GRASP Lab at the University of Pennsylvania, Philadelphia, USA. Emails: {macharet, akchen, shishika, pappasg, kumar}@seas.upenn.edu.

Douglas G. Macharet is also with the Computer Vision and Robotics Laboratory (VeRLab), Department of Computer Science, Universidade Federal de Minas Gerais, Brazil. E-mail: doug@dcc.ufmg.br.

We gratefully acknowledge the support of ARL grant ARL DCIST CRA W911NF-17-2-0181. This work was also supported in part by CNPq/Brazil (204974/2018-6).

We tackle these problems using task allocation for team coordination and sequencing to capture the maximum number of intruders. Given a restricted number of defenders, the coordination problem is directly related to coverage control, i.e., defenders should place themselves in a position that best enables the capture of upcoming intruders. The routing component emerges due to their limited speed, so even if they are aware of several intruders, it is necessary to select a subset of intruders they are actually able to capture.

The proposed methodology consists of an intruder arrival estimator based on past incursion information combined with partition-based coverage control. Next, accordingly to the ratio between the intruders' and defenders' velocities, we select intruders that are *reachable* and define a sequence that maximizes the number of captures.

Figure 1 shows an example of the considered scenario. The circular region in the center represents the perimeter, with associated set defenders (blue dots) and approaching intruders (red dots). Each defender has an associated sector (colored segmented annulus) that determines which intruders it is responsible for. The orange line represents the probability of new intruder arrivals given by the estimator and the blue lines describe the sequence of captures.

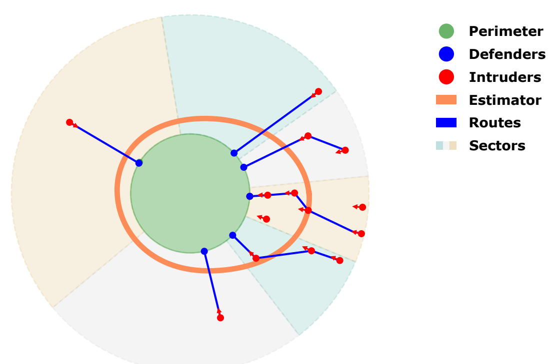

Fig. 1: Example scenario considering a non-uniform intruder spatial distribution. Sectors are adapted according to the current arrival estimation, and defenders are indirectly guided to cover the estimator's high probability regions. Blue lines illustrate the sequence of captures only, as defenders are constrained to move on the boundary of the perimeter.

Simulation results show that the proposed approach achieves high capture rates in different scenarios when compared with a naive uniform partitioning strategy, especially in cases with non-uniform intruder distributions. Our method is also capable of quickly adapting to time-varying intruder distributions.

II. Related Work

Related problems such as area coverage and monitoring [5], [6], pursuit-evasion [7], and perimeter patrol and defense [8], [9] can directly benefit from the application of MRS. However, their use also raises many challenges regarding coordination, control, and planning.

The perimeter defense task can be seen as a particular case in problems involving multi-agent systems, since we have the presence of *adversarial* agents, i.e., agents that will be acting against the success of the entire mission.

A fundamental problem when leveraging MRS is how to efficiently make use of available resources. In this context, task allocation is defined as an assignment problem, which consists in finding the best distribution for a set of tasks among the robots [10]. This combinatorial optimization problem is usually solved with graph-based models [11] or market-based approaches [12].

If we take into account the distribution of the tasks in the environment, a simple strategy is to make the assignment based on the spatial relation between tasks and robots. Region decomposition or partitioning is commonly used in motion planning [13], and is applied for segmenting an area into a finite set of sub-regions that should be traversed. In the case of task allocation, the assignment can be based on these sub-regions, where a robot is responsible for all tasks within its allocated portion of the environment [14].

The distribution of robots in the environment to better attend to current or new demands/tasks can be achieved through a coverage control strategy [15]. This approach works as a control law that drives the robots to certain positions in order to optimize a pre-defined cost function.

Perimeter defense can also be seen as a Dynamic Vehicle Routing (DVR) problem [16]. In this context, intruders act as new demands that must be attended within a certain time, and the objective is to minimize the expected time of service. Approaches are typically built upon solutions to instances of the Traveling Salesman Problem (TSP).

Routing-based approaches often assume the use of a single defender. In [17] the task is to defend a linear boundary, upon which the defender's position is constrained. The authors' proposed policy determines a route through the maximum number of reachable intruders. More recently, and based on the same concept, [18] considered the case of radially incoming targets where the defender is able to move freely in the environment. A slightly different problem is presented in [19], where the defender must capture targets before they escape from a circular environment.

Nonetheless, the most common approach is to consider the use of multiple defenders. In [20] the authors propose a simple coordination strategy based on a threshold denoted *shame-level*. The work in [21] uses a particle-based method to jointly minimize a cost function using Bayes risk. In [22], the environment is partitioned and each vehicle executes a policy within its designated region. A graph-based approach is presented in [2], [3], where a onetime maximum matching problem between defenders and intruders is solved.

In this work, we estimate a probability density for the arrival direction of new intruders. This density function is used as input to a coverage control policy, which promotes region partitioning. Finally, an independent routing policy is executed by each defender which considers intruders within the defender's assigned region.

III. Problem Formulation

Given an obstacle-free planar environment $\mathcal{E}$, consider a circular region, $\mathcal{P}$. This region must be guarded from radially incoming intruders by a set of defender robots.

Definition 1 (Perimeter). *Without loss of generality, we assume $\mathcal{P}$ to be a unit circle ($r_p = 1$) centered at the origin ($\mathbf{n} = \langle 0, 0\rangle$) and we aim to defend its perimeter (boundary) $\partial\mathcal{P}$ which can be formally defined as*

$$\partial\mathcal{P} = \{\mathbf{p} \mid \mathbf{p} \in \mathbb{R}^2, \|\mathbf{p} - \mathbf{n}\| = r_p\}.$$

Definition 2 (Defenders). *A homogeneous team of defenders $\boldsymbol{\mathcal{D}} = \{\mathcal{D}_1, ..., \mathcal{D}_N\}$ is distributed over $\partial\mathcal{P}$. The defender's position is represented in polar coordinates: $\mathbf{p}_i^d = [r_i^d,\ \theta_i^d]^T$. Defenders are constrained to move on $\partial\mathcal{P}$, i.e., $r_i^d = r_p$, and with limited velocity ν_d, i.e., $|\dot{\theta}_i^d| \leq \nu_d$.*

Definition 3 (Intruders). *An intruder $\mathcal{A}_j$ is an agent approaching $\mathcal{P}$, and its position is represented by a polar coordinate $\mathbf{p}_j^a = [r_j^a,\ \theta_j^a]^T$. Intruders move linearly towards the center of the perimeter, i.e., $\nu_a = -\dot{r}_j^a$.*

The insertion of intruders in the environment occurs sequentially over time accordingly to a Poisson process with rate λ. The azimuthal location (i.e. the angle θ^a) of new intruders is randomly chosen accordingly to some unknown distribution (ϕ^*), but at a fixed distance η from the perimeter (See Fig. 1).

We assume that as soon as a new intruder is inserted it starts moving towards the center of the region in a straight line with constant linear velocity ν_a. Hereafter we assume that ν denotes the ratio $\nu = \nu_a/\nu_d$.

An intruder j is *captured* (intercepted) by a defender i when $||\mathbf{p}_j^a - \mathbf{p}_i^d|| = 0$, otherwise, it *escaped* if it was able to reach the perimeter, i.e., $||\mathbf{p}_j^a - \mathbf{n}|| \leq r_p$, without being captured.

Problem 1 (Cooperative Perimeter Defense). *Given a region $\mathcal{P}$ and a group of N defender agents placed on its boundary $\partial\mathcal{P}$, define a cooperative defense strategy that minimizes the number of intruders that reach the perimeter of the region before being intercepted by a defender.*

The overall performance of the proposed methodology is assessed by the fraction of intruders that the team of defenders manage to capture [17], [18], i.e.,

$$\mathcal{C} = \lim_{t\to\infty} \mathrm{E}\left[\frac{n_{\mathrm{cap}}(t)}{\lambda t}\right], \qquad (1)$$

considering the arrival of new intruders is governed by a Poisson process with rate λ, and with $n_{cap}(t)$ denoting total captures at time t.

IV. Intruder Arrival Estimation

A. Motivation

In the simplest case of perimeter defense with a single defender and unlimited sensing, we are able to see all intruders at all times. Then, the problem of maximizing the number of captures is reduced to that of finding the longest path/max cardinality (i.e. the path which visits the most intruders) within the set of reachable intruders [17], [18]. However, intruders are continually inserted into the environment, and with limited sensing their positions cannot be known before they are present in the environment. We may continue to use a longest path strategy over current intruders, but the performance may be suboptimal as the resultant behavior does not account for the yet unseen intruders.

Besides this issue, the optimal control policy for teams of defenders is more complex than that for a single defender. Defenders may no longer blindly use the longest path strategy, as they must coordinate their actions, avoiding overlapping assignments or else achieve poor performance scaling. For example, the defenders could all eventually end up on one side of the perimeter, meaning any intruders on the opposite side would be unreachable. Partitioning is one suggested method, but introduces the question of how the partitions should be selected. Poor partitioning (e.g. uniform partitions for a non-uniform intruder distribution) can lead to heavily degraded performance, as we show in Section VII.

To address the above two difficulties, we employ a distribution estimator which records the arrival locations of intruders. The estimator is then integrated into the adaptive partitioning method discussed in Sec. V. Intuitively, we may expect temporal locality and non-uniformity in many scenarios; for instance, intruders in defense scenarios may arrive in groups, and may prefer certain routes of attack (e.g. when there are roads). Distribution estimation exploits both temporal locality and non-uniformity to bias defender actions towards defending regions that are more likely to be attacked. In addition, by using a distribution estimation, we may build upon existing work in coverage control to optimally partition the perimeter amongst the defenders, thereby obtaining a coordinated control strategy.

B. Formulation and Distribution Updates

The estimated distribution ϕ exists over the polar angles θ of our perimeter. That is, $\phi(\theta)$ is the relative likelihood of an intruder appearing at angle θ. Note that

$$\int_\theta \phi(\theta) = 1,$$

or in the discrete case

$$\sum_\theta \phi(\theta) = 1.$$

The information from this distribution can then be used to inform the adaptive partitioning strategy in Section V. Updates are made when new intruders appear, which causes an update for the adaptive partitioning strategy.

Note, however, that it may be desirable to not update the distribution, e.g. if it is more critical to defend a certain side of the perimeter. In these cases, it is possible to construct a distribution estimation to induce a certain partition, as opposed to changing based on the appearance of new intruders. For instance, we may construct a wrapped normal (von Mises) distribution centered around the angle where we expect the most intruders to appear, causing our defenders to bias their positions toward the relevant angle. If our prior belief is accurate, we can capture more intruders than without a prior. We demonstrate this with our omniscient adaptive partitioning tests in Section VII.

When a new intruder appears, we may update our estimated distribution ϕ. First, we consider a distribution ϕ_j^a for an individual intruder a_j. ϕ_j^a attempts to capture the locality that the presence of a new intruder presents with a wrapped normal distribution that has mean $\mu = \theta_j^a$ (the polar angle of the intruder) and standard deviation $\sigma = .5$. Then, we may either take a sliding window of the most recent n intruders or all the seen intruders to calculate an average distribution among all the selected intruders. Call this selected set of intruders $\boldsymbol{\mathcal{A}}^*$. We take our updated ϕ as

$$\phi(\theta) = \frac{1}{M_{\boldsymbol{\mathcal{A}}^*}} \sum_{a_j \in \boldsymbol{\mathcal{A}}^*} \phi^a(\theta) \tag{2}$$

where

$$M_{\boldsymbol{\mathcal{A}}^*} = \int_\theta \sum_{a_j \in \boldsymbol{\mathcal{A}}^*} \phi^a(\theta) \tag{3}$$

or in the discrete case

$$M_{\boldsymbol{\mathcal{A}}^*} = \sum_\theta \sum_{a_j \in \boldsymbol{\mathcal{A}}^*} \phi^a(\theta) \tag{4}$$

where Equation 2 averages over all selected distributions. $M_{\boldsymbol{\mathcal{A}}^*}$ is a normalizing factor that ensures the resulting $\phi(\theta)$ is a valid probability distribution.

Though we only consider the wrapped normal distribution (the analogue to the Gaussian distribution for a circular perimeter), other distribution families may also be considered for ϕ_j^a. These could include a uniform distribution close to θ_j^a and zero everywhere else, a linearly decreasing distribution from θ_j^a, or other probability distribution that incorporates prior knowledge of possible intruder distributions.

The sliding window n is taken as a tuning parameter. With a low n, the estimator has very high variance and is unable to accurately estimate the underlying intruder distribution. On the other hand, if n is too high, the estimator is unable to adapt to time-varying distributions.

V. Adaptive Partitioning

A. Centroidal Voronoi

For coordinated defense, we consider a strategy based on partitioning. The space is divided into N disjoint regions, where each region is assigned to a single defender. A natural way to partition the space in our scenario is by the polar angles. The sector (connected component on the perimeter),

S_i assigned to defender $\mathcal{D}_i$ is defined by the starting angle θ_{is} and ending angle θ_{ie}: i.e., $S_i = \{\theta \mid \theta \in [\theta_{is}, \theta_{ie}]\}$.

A naive approach is to uniformly partition the space using $[\theta_{is}, \theta_{ie}] = [\frac{2\pi(i-1)}{N}, \frac{2\pi i}{N}]$. However, this approach does not perform well when the intruder's arrival probability is non-uniform. To accommodate scenarios where the intruder's arrival probability is non-uniform and possibly time varying, we propose an adaptive partitioning strategy based on the idea of coverage control [15]. In coverage control, robots are distributed in the space based on some density function that quantifies the *importance* of each point in the space.

An efficient control policy based on Lloyd's algorithm has been proposed [23]. First, the area is partitioned into Voronoi cells. In our case, each robot is responsible for the set of points on the circle closest to itself:

$$V_i = \{\theta \mid |\theta - \theta_i^d| \leq |\theta - \theta_j^d|, \forall j \neq i\}. \tag{5}$$

The set of regions $\{V_1, ..., V_N\}$ is called the Voronoi diagram for the generators $\{\theta_1^d, ..., \theta_N^d\}$.

Next, each robot computes the centroid of the Voronoi cell weighted by the known density function ϕ:

$$C_{V_i} = \frac{1}{M_{V_i}} \int_{V_i} \theta \phi(\theta) d\theta \tag{6}$$

where

$$M_{V_i} = \int_{V_i} \phi(\theta) d\theta. \tag{7}$$

In our case, the density function $\phi(\theta)$ describes the estimated probability for intruder arrivals.

Let us consider first order dynamics for the defender:

$$\dot{\theta}_i^d = \omega_i, \tag{8}$$

then Lloyd's algorithm commands each robot to move towards the centroid according to:

$$\omega_i = -k_{\text{prop}}(\theta_i^d - C_{V_i}), \tag{9}$$

where k_{prop} is a positive gain. This algorithm is known to stabilize the system and achieve $\theta_i^d = C_{V_i}, \ \forall \ i$. Such a configuration minimizes the cost function

$$\mathcal{H} = \sum_i \int_{W_i} \|\theta - \theta_i^d\|^2 \phi(\theta) d\theta, \tag{10}$$

where W_i is the *dominance region* that $\mathcal{D}_i$ is responsible for.

B. Partitioning Strategy

Noting that the defenders must be pursuing the intruders and not the Voronoi center, we consider a virtual agent that generates the partition. Every defender keeps track of a virtual agent $\hat{\theta}_i^d$, which evolves according to the dynamics in (9) and satisfies $\hat{\theta}_i^d(t_0) = \theta_i^d(t_0)$. To compute the evolution of the virtual agent, each defender must share the location of its own virtual agent with its neighboring defenders.

At every time instant, the sector is defined by

$$S_i = V_i \tag{11}$$

where the Voronoi cell is computed using the virtual agents. Each defender plans its behavior based on the intruders in its own sector.

VI. Independent Routing Policy

In this section, we detail a simple distributed and non-cooperative routing strategy based on [17]. The approach is non-cooperative in the sense that there is no explicit coordination among defenders, and the planning step occurs separately within each individual partition.

An intruder a_j is said to be *reachable* by a defender d_i if its time to reach the boundary is greater than the time the defender takes to go from its current position to an intercepting point. More formally, the intruder is reachable by the defender if and only if

$$r_{ij} \geq \nu|\theta_{ij}|, \tag{12}$$

with $r_{ij} = r_j^a - r_i^d$ and $\theta_{ij} = \theta_j^a - \theta_i^d$.

Definition 4 (Reachable set)**.** *The set of all points moving at constant speed ν_j that are reachable by a vehicle at position $\mathbf{p}_i$ and capable of move at maximum speed ν_i is given by*

$$R(\mathbf{p}_i; \nu) = \{(r_j, \theta_j) \mid r_{ij} \geq \nu|\theta_{ij}|\}. \tag{13}$$

The set $\boldsymbol{\mathcal{A}}^t$ is composed of all the intruders currently in the environment, at a time instant t, that have neither been captured nor escaped, i.e.:

$$\boldsymbol{\mathcal{A}}^t = \{\mathcal{A}_j \mid r_j^a \leq \eta \wedge r_j^a > r_p\}. \tag{14}$$

The subset $\boldsymbol{\mathcal{A}}_i^t \subset \boldsymbol{\mathcal{A}}^t$ represents the intruders that are within the bounds of a certain sector S_i (associated to defender $\mathcal{D}_i$), defined as

$$\boldsymbol{\mathcal{A}}_i^t = \boldsymbol{\mathcal{A}}^t \setminus \{\boldsymbol{\mathcal{A}} j \mid \theta_j^a \notin [\theta_{is}, \theta_{ie}]\}. \tag{15}$$

A Directed Acyclic Graph (DAG) $G = (V, E)$ is a finite directed graph with a set of vertices V and a set of directed edges $E \subseteq V \times V$ with no directed cycles. A simple path in G is a sequence of non-repeated vertices with a direct edge $e \in E$ connecting it to the next one in the sequence.

Definition 5 (Reachability graph)**.** *A reachability graph is represented as a DAG with vertex set V and edge set E, where for $v_i, v_j \in V$ the edge $(v_i, v_j) \in E$ if and only if $v_i \neq v_j$, $\langle \mathbf{p}_j, \nu_j \rangle \in R(\mathbf{p}_i, \nu_i)$, and $\mathbf{p}_j \in \boldsymbol{\mathcal{A}}_i^t$.*

The main idea of the strategy is that at time t, each defender individually determines which intruders within its partition are reachable, and then computes a route that maximizes the capture number by finding the longest path (i.e., the one that traverses the maximum number of vertices) on the reachability graph [17] (Algorithm 1).

Algorithm 1 IndependentRoutingPolicy($\boldsymbol{\mathcal{D}}$, t)

for each $\mathcal{D}_i \in \boldsymbol{\mathcal{D}}$ **do**
  $\boldsymbol{\mathcal{A}}_i^t \leftarrow$ `UpdateIntruderSet`$(\mathcal{D}_i)$
  $G_t \leftarrow$ `CreateReachabilityGraph`$(\mathcal{D}_i, \boldsymbol{\mathcal{A}}_i^t)$
  $P_t \leftarrow$ `DetermineLongestPath`(G_t)
end for
`MoveTowardsFirstNode`(P_t)

Fig. 2 shows the limits of a defender's (blue dot) reachability set (red lines), as well as its associated sector (shaded wedge). Intruders (red dots) may be in different situations: (a) reachable by the defender, but outside its sector; (b) inside the sector, but not reachable; and (c) within the sector and reachable. Considering only the last case for potential valid captures, we define a route that maximizes the number of captures (blue line).

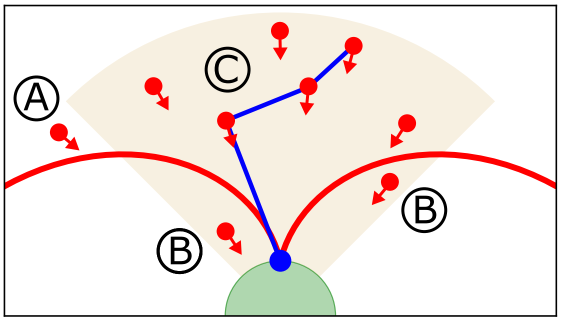

Fig. 2: Example showing the different cases for the intruders (red dots) considering a defender sector (shaded wedge) and reachability set (area above red lines). Only intruders in case (C) are valid potential captures for the defender. The defender (blue dot) selects a route that maximizes the captures (blue line).

VII. Experiments

A. Illustrative examples

We first show an example of the proposed estimator and adaptive partitioning strategy[1]. As an example, we generate intruders according to a Poisson process with arrival rate $\lambda = 4$ over a period of $T = 30$ seconds. We first consider a uniform angle distribution (i.e. intruders are equally likely to appear from any direction). Our estimator is initialized to a uniform distribution and has a sliding window of $n = 25$, which is to say it forms its estimate from the information of the most recent 25 intruders. Various phases are shown in Figure 3. During the run, virtual agents (denoted by the small green circles) are free to move according to the adaptive partitioning strategy, independent of their corresponding defenders (blue circles). Recall that the virtual agents determine the partitions for which their corresponding defenders are responsible for. As intruders appear, the estimator changes its estimated distribution and the partitions assigned to each defender shift accordingly. At the end of the run in Figure 3d, the estimator is no longer completely uniform, but does not weight one side much more greatly than others.

Contrast this result with that generated when the intruders' incoming angle is distributed according to a nonuniform distribution. As an example of a non-uniform distribution, we use a von Mises (circular normal) distribution with mean μ and variance $1/\kappa$:

$$\phi^*(\theta) = \frac{1}{M_v} \exp\left(\kappa \cos(\theta - \mu)\right), \tag{16}$$

where M_v is a normalization factor.

Figure 4 shows a result with a von Mises distribution ($\mu = 0$, $\kappa = 2$). When the distribution is nonuniform, our strategy places larger numbers of defenders where more intruders have appeared in the past. Also, smaller partitions are assigned to the defenders covering highly weighted areas. These characteristics are desirable, as they help to ensure that intruders are more likely to be within the reachability set of the defenders and that fewer intruders are not captured because a defender is occupied with other captures.

B. Numerical analysis

Since the proposed system has two components (the estimator and adaptive partitioning strategy), we conduct three sets of tests:

- Estimator: We compare the estimated distribution with the true intruder distribution, and measure its performance through Kullback–Leibler (KL) divergence.
- Adaptive partitioning: We quantify the performance of the adaptive partitioning strategy when it knows the true intruder distribution. Performance is measured by fraction of intruders captured.
- Entire system: We evaluate the estimator and adaptive partitioning strategy, and measure performance through fraction of intruders captured.

For these tests, we consider a team of $N = 5$ defenders, an intruder arrival schedule over $T = 30s$, intruder arrival rate of $\lambda = 4$, and $\nu = 1$, where $\nu_d = \nu_a = 0.5$m/s. The bimodal distribution is generated by combining two von Mises distributions with $\mu = 0$ and $\mu = 2.5$, and the time varying distribution is generated by changing the mean of the von Mises distribution every 10 seconds to $\mu = \{0, \pi, \frac{3\pi}{2}\}$.

1) Estimator: First, we consider the performance of the estimator. We calculate the average KL divergence between the true distribution and the estimated distribution at each of the simulator's discrete time steps, and average over all time steps. Finally, we run the simulation 100 times, and report the average KL divergence. Note that the metric used is reverse KL divergence to avoid the problem of zero avoidance associated with forward KL divergence.

Estimator results are shown in Figure 5. As we increase the size of the estimator's window (i.e. the number of intruders considered for the estimate), the KL divergence between the estimated distribution and true distribution trends towards zero for the static distributions, which is expected. On the other hand, increasing the window size too much results in poor performance for the time varying case, as the estimator is unable to quickly adapt.

2) Adaptive partitioning and System: For our second and third sets of tests, we consider two additional defender strategies: a naive strategy ($Naive$), which always partitions the perimeter equally between the defenders, and the omniscient adaptive partitioning (AP) strategy, which uses

[1] Video of the execution: `https://youtu.be/b6L86WKPO_w`

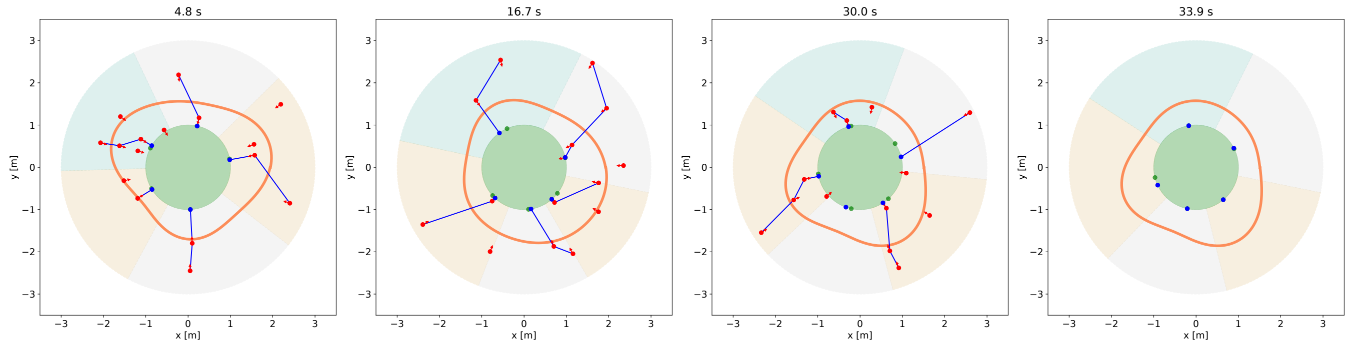

(a) (b) (c) (d)

Fig. 3: Illustrative example of the estimator and adaptive partitioning strategy considering an *uniform* intruder arrival distribution. The estimator does not heavily weight a side, which is expected for the uniform intruder distribution. Defenders are spaced roughly evenly around the perimeter in accordance with the adaptive partitioning strategy.

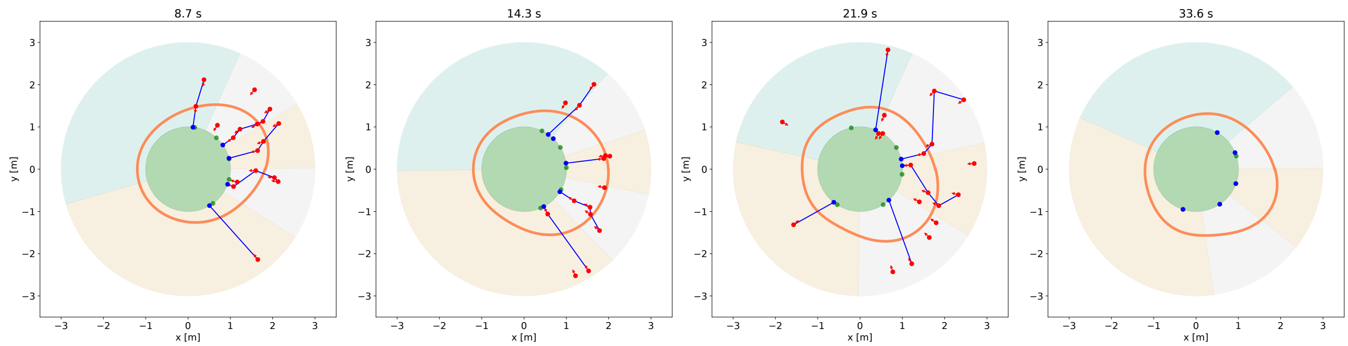

(a) (b) (c) (d)

Fig. 4: Illustrative example of the estimator and adaptive partitioning strategy considering a *nonuniform* intruder arrival distribution. As seen in the figure, the estimator assigns a higher weight to the right side of the perimeter, which corresponds with $\mu = 0$ in the von Mises intruder distribution. Defenders are then placed towards the right side in accordance with the adaptive partitioning strategy.

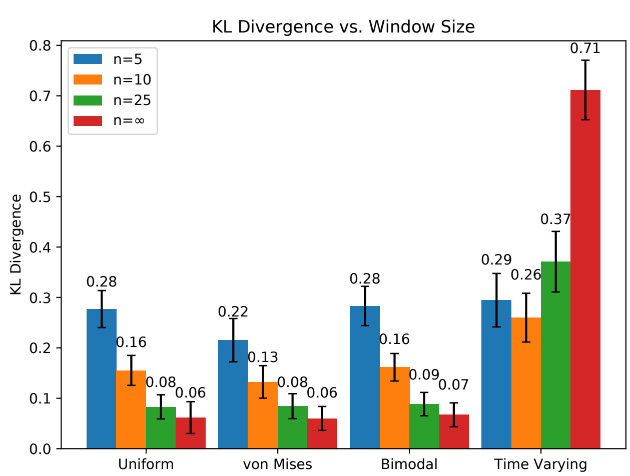

Fig. 5: Average KL divergence and standard deviation vs. window size n over 100 runs for the estimator. For static intruder distributions, bigger windows result in more accurate estimates. As expected for the time varying distribution, the infinite size window performs poorly.

the true intruder distribution for its adaptive partitioning updates. Note that both the naive strategy and omniscient adaptive partitioning strategy are static for a static true intruder distribution.

Finally, we test our estimated adaptive partitioning (EAP_n), where n denotes the window size. Results for both the uniform true distribution and von Mises true distribution are shown in Table I.

Strategy	**Uniform**	**von Mises**	**Bimodal**	**Time Varying**
$Naive$	64.9 ± 3.2	58.8 ± 3.6	64.8 ± 3.2	58.2 ± 3.1
AP	64.9 ± 3.2	72.4 ± 3.1	67.8 ± 2.9	65.6 ± 4.1
EAP_5	64.1 ± 3.7	68.0 ± 4.0	64.8 ± 4.1	62.6 ± 4.4
EAP_{10}	65.6 ± 3.7	70.1 ± 4.7	67.1 ± 4.3	64.0 ± 3.8
EAP_{25}	66.2 ± 3.2	70.2 ± 3.9	67.9 ± 3.9	64.7 ± 3.7
EAP_∞	66.0 ± 3.7	69.6 ± 4.0	68.1 ± 3.5	60.6 ± 3.6

TABLE I: Average percentage and standard deviation of captured intruders over 100 runs for different defender strategies. Note that $(Naive)$ and (AP) are equivalent for the uniform true distribution, so the results are identical.

In the uniform case, where the naive and omniscient adaptive partitioning are equivalent, the proposed strategy matches the performance of the other two strategies. In the von Mises case, the proposed strategy outperforms the naive strategy by around a 10% increase in capture rate, while also coming very close to the omniscient adaptive partitioning strategy. These results are similar for the bimodal distribution, where (EAP) outperforms $(Naive)$ and matches (AP). Through these three static distributions, we see that a window size of $n = 10, 25$ or ∞ perform roughly the same and come close to (AP). EAP_n exhibits degraded performance, which suggests that a window size of $n = 5$ is too small.

For the time varying intruder distribution, both (EAP_{10}) and (EAP_{25}) perform similarly to (AP). In this case, the infinite window size does not allow the estimator to adapt quickly enough, which causes degraded performance. Thus, to best accommodate both static and time varying distributions, we would select a window of $n = 10$ or $n = 25$ for this test setup. With these window sizes, the proposed system is able to match or significantly exceed the performance of our baseline $(Naive)$, and mostly match the performance of the omniscient (AP).

VIII. Conclusion and Future Work

In this paper, we tackle the perimeter defense problem against radially incoming intruders and consider a collection of defender agents. We propose a coordination strategy based on intruder arrival estimation and coverage control. Defenders then apply a routing approach in order to maximize the fraction of captured intruders within their partitions.

The proposed estimator is able to closely approximate the true static distribution, as measured by reverse KL divergence. The partition-based coverage control tends to allocate more defenders where intruders have appeared in the past. In highly weighted areas partitions will be smaller, making intruders more likely to be within the reachability set of the defenders and therefore more susceptible to capture. Finally, the longest tour approach is an efficient way to select the best order in which to capture intruders.

Results show that for a non-uniform von Mises spatial arrival distribution, the methodology significantly outperforms a naive strategy without sacrificing performance in the uniform intruder distribution case. In addition, when using the estimator's distribution, the adaptive partitioning strategy performs closely to the omniscient adaptive partitioning strategy. These results indicate that our solution is viable in scenarios with both static and time-varying distributions.

Future work includes evaluation with additional environment scenarios, including different intruder distributions, intruder frequencies, and defender numbers. Another possible extension is modification of the estimator's distribution update with something other than a wrapped normal distribution for ϕ^i.

References

[1] Z. Yan, N. Jouandeau, and A. A. Cherif, "A Survey and Analysis of Multi-Robot Coordination," *International Journal of Advanced Robotic Systems*, vol. 10, no. 12, p. 399, 2013.

[2] D. Shishika and V. Kumar, "Local-game Decomposition for Multi-player Perimeter-defense Problem," in *IEEE Conference on Decision and Control (CDC)*, Dec 2018, pp. 2093–2100.

[3] D. Shishika, J. Paulos, M. A. Hsieh, and V. Kumar, "Team Composition for Perimeter Defense with Patrollers and Defenders," in *IEEE Conference on Decision and Control (CDC)*, Dec 2019.

[4] D. Shishika, J. Paulos, and V. Kumar, "Cooperative Team Strategies for Multi-Player Perimeter-Defense Games," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2738–2745, 2020.

[5] 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, June 2012.

[6] T. P. Lambrou and C. G. Panayiotou, "Collaborative path planning for event search and exploration in mixed sensor networks," *The International Journal of Robotics Research*, vol. 32, no. 12, pp. 1424–1437, 2013.

[7] T. H. Chung, G. A. Hollinger, and V. Isler, "Search and pursuit-evasion in mobile robotics," *Autonomous Robots*, vol. 31, no. 4, p. 299, Jul 2011.

[8] N. Agmon, S. Kraus, and G. A. Kaminka, "Multi-robot perimeter patrol in adversarial settings," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2008, pp. 2339–2345.

[9] N. Basilico, N. Gatti, and F. Amigoni, "Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder," *Artificial Intelligence*, vol. 184-185, pp. 78 – 123, 2012.

[10] B. P. Gerkey and M. J. Matarić, "A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.

[11] C. Nam and D. A. Shell, "Assignment algorithms for modeling resource contention and interference in multi-robot task-allocation," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 2158–2163.

[12] M. G. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A. J. Kleywegt, "Simple auctions with performance guarantees for multi-robot task allocation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 1, Sep. 2004, pp. 698–705.

[13] S. M. LaValle, *Planning Algorithms*. New York, NY, USA: Cambridge University Press, 2006.

[14] J. G. M. Fu, T. Bandyopadhyay, and M. H. Ang, "Local voronoi decomposition for multi-agent task allocation," in *IEEE International Conference on Robotics and Automation*, May 2009, pp. 1935–1940.

[15] M. Schwager, D. Rus, and J.-J. Slotine, "Decentralized, adaptive coverage control for networked robots," *The International Journal of Robotics Research*, vol. 28, no. 3, pp. 357–375, 2009.

[16] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. Smith, "Dynamic Vehicle Routing for Robotic Systems," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1482–1504, sept. 2011.

[17] S. L. Smith, S. D. Bopardikar, and F. Bullo, "A dynamic boundary guarding problem with translating targets," in *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, Dec 2009, pp. 8543–8548.

[18] S. Bajaj and S. D. Bopardikar, "Dynamic boundary guarding against radially incoming targets," in *IEEE International Conference on Decision and Control (CDC)*, Dec 2019, pp. 642–649.

[19] P. Agharkar, S. Bopardikar, and F. Bullo, "Vehicle Routing Algorithms for Radially Escaping Targets," *SIAM Journal on Control and Optimization*, vol. 53, no. 5, pp. 2934–2954, 2015.

[20] Y. Zhang and Y. Meng, "A decentralized multi-robot system for intruder detection in security defense," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2010, pp. 5563–5568.

[21] M. J. Bays, A. Shende, and D. J. Stilwell, "An approach to multi-agent area protection using bayes risk," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2012, pp. 642–649.

[22] M. Pavone, E. Frazzoli, and F. Bullo, "Adaptive and distributed algorithms for vehicle routing in a stochastic and dynamic environment," *IEEE Transactions on Automatic Control*, vol. 56, no. 6, pp. 1259–1274, June 2011.

[23] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.