

Multi-Agent Perimeter Monitoring for Uncertainty Reduction

Brian Reily¹, Bradley Woosley¹, Christopher Reardon², and John G. Rogers¹

Abstract—Perimeter monitoring is a valuable capability for a multi-agent system, with multiple defense- and security-focused applications. In this work in progress, we consider the problem of perimeter monitoring for uncertainty reduction, in which a team of defender agents must position and orient themselves to minimize the uncertainty about attackers approaching the perimeter. We envision perimeter monitoring as a crossover of perimeter defense and area coverage, and formulate it as a distributed optimization problem, where defenders take actions to maximize their individual sensing of attackers. We describe both single agent and multi-agent scenarios, and propose multiple evaluation setups to analyze our approach.

I. INTRODUCTION

Multi-agent systems present many advantages over single agents, primarily in their robustness and their ability to accomplish multiple tasks simultaneously across an environment. Multi-agent systems can be composed of simpler agents, for example with restrictions on operational range or sensor complement, and still operate more effectively than expensive, well-equipped individual agents.

A key capability of defense and security systems is perimeter monitoring, and multi-agent systems are particularly well-suited for this task, as a real-world perimeter, such as around an installation or military unit, would be too large for an individual agent to patrol. While often in physical environments we consider robots to be the agents, multi-agent systems could also consist of fixed sensors, mobile non-robot sensors, or even humans acting as teammates.

We specifically consider the problem of *perimeter monitoring*, and consider it a combination of, but separate from, the problems of perimeter defense and area coverage. In perimeter monitoring, agents position themselves along a fixed perimeter, but with the aim of maximizing coverage of an environment or attacker agents outside the perimeter. An illustration of the problem is seen in Figure 1. Perimeter defense, however, is the version of the problem where defender agents must intercept attackers before or while they breach the perimeter (despite the differences between the problems, we do adopt the nomenclature of *defenders* and *attackers*). Work in this area has largely utilized game theoretic formulations, treating these as pursuit-evasion games [1], Blotto games [2], or polymatrix games [3]. Similarly, area coverage is a related but distinct research area. In this problem, agents distribute themselves throughout an




Fig. 1. An overview of the perimeter monitoring problem: Blue circles represent available defender locations, with the dashed lines emerging from two locations indicating possible fields of view. Red elements denote probability distributions representing uncertain observations of attackers. The goal of our approach is to position defenders at available locations and orientations to maximize their observations of these attacker distributions with the aim of reducing the uncertainty.

environment to maximize sensor coverage of the entire area [4] or of events occurring in the environment [5]. Despite active research in both these areas, they do not address the problem of perimeter monitoring - how can a multi-agent system optimally position its agents on a fixed perimeter to maximally reduce uncertainty about attacker agents?

In this work in progress, we define the value of sensed points based on their distance to defender locations, while considering sensor overlap. We formulate the entire problem as finding an optimal set of defender positions and orientations that maximizes the sensed points. As this work is ongoing, we have not yet evaluated this formulation. However, we propose multiple evaluation setups to compare to possible baseline approaches (e.g., task allocation approaches), and to learn the effects of assumptions within the problem (e.g., field of view, numbers of defenders, or heterogeneous capabilities).

II. RELATED WORK

The problem we define as perimeter monitoring is a combination of two active research areas in multi-robot and multi-agent systems. First, perimeter defense aims to optimally position defender agents in order to intercept attacker agents that attempt to breach a perimeter. Second, area coverage aims to optimally position sensors in order to maximize the coverage of an area or events occurring in an area.

Perimeter defense has been most commonly addressed through the lens of game theory [6]. Games are typically defined as variants of pursuit-evasion games, where attackers

¹Brian Reily, Brad Woosley, and John G. Rogers are with the DEVCOM Army Research Laboratory in Adelphi, MD, 20783, USA. email: {brian.j.reily, bradley.r.woosley, john.g.rogers59}.civ@army.mil

²Christopher Reardon is with the Department of Computer Science at the University of Denver in Denver, CO, 80210, USA. email: christopher.reardon@du.edu

take on the role of evaders attempting to move to a point within the perimeter, while defenders play as pursuers [7]. Because of the high-dimensional state space, local games are often played by individual defenders [1]. Other proposed game theory-based solutions utilize Blotto games [2] or polymatrix games [3]. Research has also examined various team compositions [8], coordination strategies [9], and the probability of perimeter breach based on different environment compositions [10]. Perimeter defense has also been addressed from fixed position defenders, such as floodlights that can rotate but not change their position [11].

Area or sensor coverage has seen extensive research into multiple variations of its central problem of observing as much of an area as possible. Many homogeneous sensor coverage formulations focus on evenly distributing sensors within an environment [12], with approaches dividing an area into Voronoi cells [13], estimating density functions [14], or partitioning a graph [15].

For heterogeneous systems, the problem is often formulated as multiple different event types are occurring in the environment, and agents with the appropriate sensors must position themselves in order to maximize sensing of these events [5]. This has been done through Voronoi cells as well [16], but most commonly is done by identifying a distribution of robots based on learned cost functions [17].

III. OUR PROPOSED APPROACH

A. Problem Statement

We model the perimeter to monitor as an undirected graph, with a finite number of defensive locations and a finite number of movement possibilities between these locations. Formally, we consider a graph $G = \mathcal{F}V; E\mathcal{g}$, where $V = [v_1; \dots; v_m]$ is the set of vertices, representing the m available defensive locations. We treat this as a finite set, as opposed to a continuous perimeter due to real-world practicalities; consider the application of deploying agents onto a ridge to protect a military installation. The ridge may be continuous, but only a finite number of spots are feasible to move equipment to, offer cover and concealment, and provide a vantage point for overwatch. We then treat the edge set E as the possible transitions between perimeter locations. Again, due to real-world constraints we don't assume this graph to be fully connected. An example perimeter can be seen in Figure 1, denoted with the blue circles and connecting edges, where $m = 6$.

To monitor this perimeter, we consider n defenders to occupy the m possible defensive locations. Each defender has a location (a spot on the perimeter) and an orientation (the direction in which to point its possible limited field of view). We model defender locations with an occupancy matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, where x_{ij} denotes whether the i -th location is occupied by the j -th defender. A column $\mathbf{x}_j \in \mathbb{R}^m$ denotes the occupancy for an individual agent. We restrict \mathbf{X} to be binary, so that either $x_{ij} = 0$ or $x_{ij} = 1$ (i.e., occupancy is not probabilistic). We record defender orientations in a vector $\theta \in \mathbb{R}^n$, where θ_j denotes the orientation of the j -th defender. While some sensors (e.g., many commercial

LiDAR systems) can provide 360° field of view, making orientation largely irrelevant, many sensors are limited to smaller angles.

Next, we model *attackers* as probability distributions. Often, observations about unknown or possibly hostile entities are uncertain. While this uncertainty can extend to many aspects of the attacker, for example, size, number of attackers, etc., in our formulation we generalize it to uncertainty solely about location. We consider p attackers for a set of distributions $P = \{P_1; \dots; P_p\}$, where $P_i = \langle \mu_i; \Sigma_i \rangle$ represents the i -th attacker, with μ_i representing the mean of its location and Σ_i representing the covariance of its location.

Finally, we re-frame the problem of uncertainty reduction as that of maximizing the sensor coverage of the attackers. To do this, we tie the perimeter, defenders, and attackers together by quantifying the coverage of attackers each defender is able to provide.

First, we must consider three aspects beyond the defender and attacker locations: distance, attacker occlusions, and observational overlap.

- *Distance*: Almost universally, the closer a sensor is to what is being sensed, the more accurate it is. From this, we consider observations from a smaller distance (i.e., a closer location on the perimeter) to be more useful for uncertainty reduction. To account for distance, we propose a linear attenuation:

$$y_{ij}^a = \frac{1}{d_a} p_{ij} \quad (1)$$

Here, we consider the *possible coverage* y_{ij}^a of a point p_{ij} within P based on its distance d_a to perimeter location v_a .

- *Attacker Occlusions*: Although in diagrams there may appear to be occlusions (such as in the field of view in Figure 2(c)), in reality attackers most likely occupy only small points of the available field of view. We do not account for attacker occlusions - that is, from the Figure, observations of P_2 are only less useful than observations of P_1 because they are further away, not because they are occluded.
- *Observational Overlap*: While it may be the case that multiple sensors observing a single entity can increase the overall information about that entity, in practice this is difficult to quantify, difficult to capture in simulation, and often eschewed in area coverage research. Based on this, we do not double count observational overlaps - for example, in Figure 3(a), the defenders positioned at locations v_4 and v_5 both observe P_2 , yet the overlap between the fields of view should only be counted once, by the nearer defender. In the case of overlapping observations of a point p_{ij} , we quantify this by

$$y_{ij}^a = \begin{cases} \frac{1}{d_a} p_{ij} & \text{if } d_a \leq d_b \\ 0 & \text{if } d_a > d_b \end{cases} \quad (2)$$

which extends Eq. (1) to handle multiple observations of a point. In the case of an equal sensing distance from

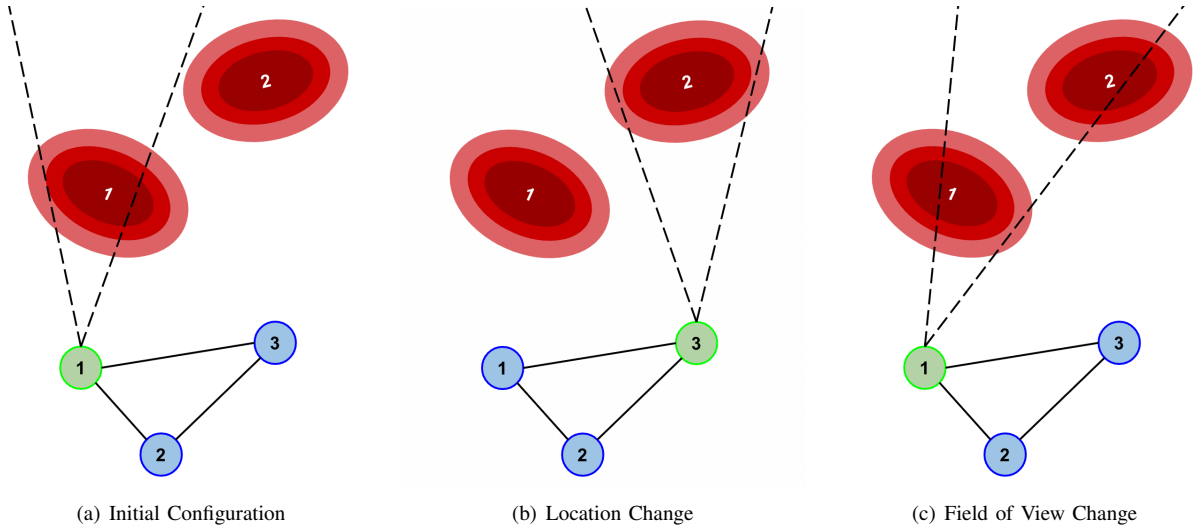


Fig. 2. The single agent scenario: In order for a single agent to maximize its observations of multiple attackers, it can change its location or its field of view. Figure 2(a) shows an agent observing only a single attacker. By changing its location in Figure 2(b), it is still limited to observing just one attacker. However, by changing its field of view in Figure 2(c), it is able to observe both attackers.

multiple vantage points, the coverage is credited to the earlier perimeter location (i.e., v_4 instead of v_5).

We then formulate the overall problem of perimeter monitoring as a maximization of the coverage of each point p_{ij} in the environment, based on the defender occupancy matrix \mathbf{X} , defender orientations θ_i , and the factors introduced by Eqs. (1) and (2):

$$\max_{\mathbf{X}} \sum_i f(\mathbf{x}_i; \theta_i; P) \quad (3)$$

where $f(\mathbf{x}_i; \theta_i; P)$ returns the total coverage value of each point y_{ij} that is within the field of view defined by the perimeter occupancy in \mathbf{x}_i and the defender orientation θ_i .

To solve this proposed formulation, we allow agents only two actions. First, they can change locations, such that $(v_i; v_j) \in E$ and they are currently located at v_i . Second, they can adjust their orientation θ_i .

B. Agent Behavior

In the case of a single agent, behavior is somewhat simple due to the limited actions available. Figure 2 illustrates the actions available to a single agent, in a limited version of the perimeter described in Figure 1.

Figure 2(a) presents a possible initial configuration for the system, in which a defender's field of view is directed towards attacker P_1 . From this location and orientation, the defender can only observe the single attacker.

By changing its location, as in Figure 2(b), it can observe P_2 , but still remains only able to observe a single attacker (and, at a further distance, this makes its observations less valuable as per Eq. (1)). However, by remaining in its original position and changing its orientation, as in Figure 2(c), it can now observe the majority of both P_1 and P_2 .

In the case of a multi-agent system, behaviors become more complicated but are based in the same atomic actions as

the single agent case. Each individual agent can still change its location (to those locations accessible from its current location) or its orientation. We do not restrict multiple agents from occupying the same location on the perimeter; however, this is discouraged from happening due to the observational overlap accounted for in Eq. (2).

Because of these limited behaviors and the limited information available to the agents, the formulation in Eq. (3) cannot be directly solved. Instead, a behavior policy must be learned to make iterative moves in the direction of maximizing the covered area. Such a policy may even make investigative moves, such as rotating the field of view to gain more information about the environment.

IV. PROPOSED EVALUATION

There are a number of possible evaluation strategies for our proposed approach. We divide these into two groupings: first, we discuss strategies to evaluate assumptions within our formulation; and second, we discuss existing methods to compare to.

Within our formulation, we believe it will be important to evaluate the effect of the field of view restriction on the time it takes for the defending team to converge and find the local maxima of Eq. (3). As the field of view grows towards 360 deg, is the decrease in convergence time linear? How does this change as the number of defenders and attackers change? How is this effected by a heterogeneous team of defenders, where perhaps agents have varying field of view capabilities?

For evaluation against existing solution methods, we believe it is important to evaluate against the global maximum of Eq. (3), i.e., the exact \mathbf{X} and θ that provides the most coverage of the attackers. Solution methods that could be compared are policies learned through reinforcement learning (both with full and limited knowledge) and task assignment methods such as the Hungarian algorithm.

V. CONCLUSION

In conclusion, perimeter monitoring for uncertainty reduction is an important capability for a multi-agent security system. Although related to the research areas of sensor coverage and perimeter defense, it is a distinct and challenging problem.

REFERENCES

- [1] D. Shishika and V. Kumar, "Local-game decomposition for multiplayer perimeter-defense problem," in *2018 IEEE conference on decision and control (CDC)*, pp. 2093–2100, IEEE, 2018.
- [2] A. K. Chen, B. L. Ferguson, D. Shishika, M. Dorothy, J. R. Marden, G. J. Pappas, and V. Kumar, "Path defense in dynamic defender-attacker blotto games (ddab) with limited information," *arXiv preprint arXiv:2204.04176*, 2022.
- [3] A. Langley, V. Dhiman, and H. Christensen, "Heterogeneous multi-robot adversarial patrolling using polymatrix games," in *International Symposium on Automation, Mechanical and Design Engineering*, pp. 13–27, Springer, 2023.
- [4] M. Corah and N. Michael, "Efficient online multi-robot exploration via distributed sequential greedy assignment.," in *Robotics: Science and Systems*, vol. 13, 2017.
- [5] B. Reily, T. Mott, and H. Zhang, "Adaptation to team composition changes for heterogeneous multi-robot sensor coverage," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9051–9057, IEEE, 2021.
- [6] D. Shishika and V. Kumar, "A review of multi agent perimeter defense games," in *International Conference on Decision and Game Theory for Security*, pp. 472–485, Springer, 2020.
- [7] 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.
- [8] D. Shishika, J. Paulos, M. R. Dorothy, M. A. Hsieh, and V. Kumar, "Team composition for perimeter defense with patrollers and defenders," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 7325–7332, IEEE, 2019.
- [9] D. G. Macharet, A. K. Chen, D. Shishika, G. J. Pappas, and V. Kumar, "Adaptive partitioning for coordinated multi-agent perimeter defense," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7971–7977, IEEE, 2020.
- [10] N. Agmon, S. Kraus, and G. A. Kaminka, "Multi-robot perimeter patrol in adversarial settings," in *2008 IEEE International Conference on Robotics and Automation*, pp. 2339–2345, IEEE, 2008.
- [11] S. Bereg, J. M. Díaz-Báñez, M. Fort, M. A. López, P. Pérez-Lantero, and J. Urrutia, "Continuous surveillance of points by rotating floodlights," *International Journal of Computational Geometry & Applications*, vol. 24, no. 03, pp. 183–196, 2014.
- [12] I. Rekleitis, V. Lee-Shue, A. P. New, and H. Choset, "Limited communication, multi-robot team based coverage," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, vol. 4, pp. 3462–3468, IEEE, 2004.
- [13] K. Guruprasad and D. Ghose, "Performance of a class of multi-robot deploy and search strategies based on centroidal voronoi configurations," *International Journal of Systems Science*, vol. 44, no. 4, pp. 680–699, 2013.
- [14] S. G. Lee, Y. Diaz-Mercado, and M. Egerstedt, "Multirobot control using time-varying density functions," *IEEE Transactions on robotics*, vol. 31, no. 2, pp. 489–493, 2015.
- [15] S.-k. Yun and D. Rus, "Distributed coverage with mobile robots on a graph: locational optimization and equal-mass partitioning," *Robotica*, vol. 32, no. 2, pp. 257–277, 2014.
- [16] O. Arslan and D. E. Koditschek, "Voronoi-based coverage control of heterogeneous disk-shaped robots," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4259–4266, IEEE, 2016.
- [17] M. Santos and M. Egerstedt, "Coverage control for multi-robot teams with heterogeneous sensing capabilities using limited communications," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5313–5319, IEEE, 2018.

(a) Initial Configuration

(b) Field of View Change

(c) Location and Field of View Change

Fig. 3. The multi-agent scenario: when multiple agents monitor the perimeter, they each make individual changes to improve the overall monitoring of the system. Figure 3(a) shows a possible initial configuration, in which two agents are making overlapping observations. By adjusting field of view in Figure 3(b), all attackers can be observed. However, by multiple agents making adjustments in Figure 3(c), all can attackers can be observed from a closer distance.