A Review of Multi Agent Perimeter Defense Games*

Daigo Shishika[1 [0000−0002−8514−9634]] and Vijay Kumar[2 [0000−0002−3902−9391]]

[1] George Mason University, Fairfax VA 22030, USA. dshishik@gmu.edu
[2] University of Pennsylvania, Philadelphia PA 19104, USA. kumar@seas.upenn.edu

Abstract. This paper reviews a series of works done on the multi-agent perimeter defense scenario, in which a team of intruders try to score by reaching the target region while a team of defenders try to minimize the score by intercepting those intruders. We describe how the small-scale differential games are solved and are leveraged to design team strategies in the large-scale swarm versus swarm scenarios. Three different approaches to analyze the large-scale games (MM, MIS, and LGR) are introduced with comments on their relative strengths and weaknesses. As a unique contribution of this paper, we discuss how the existing results can be extended into more general problem formulations. Furthermore, we point out the limitations of the current work and suggest potential directions for future research.

Keywords: Pursuit-evasion games · Multi-agent systems · Cooperative control · Reach-avoid games · Perimeter defense.

1 Introduction

Multi-robot systems (MRS) have gained significant attention in the past few decades. Since MRS are naturally robust due to their redundancy, and are also able to distribute into large areas, there are various application spaces that are anticipated. Of particular relevance to this paper is a class of scenarios related to security and defense applications. Specifically, we discuss a scenario in which a group of defenders is tasked to protect a region from a group of intruders.

Since we are concerned with two parties with conflicting objectives, the study of effective strategies naturally fits into game-theoretic analyses. The problem which we call the *perimeter defense game* considers a scenario where intruders try to reach the perimeter of a target region without being intercepted by the defenders, whereas a team of defenders seek to intercept or capture those intruders before they reach the perimeter. This is a variant of pursuit-evasion games for which various versions of one pursuer vs. one evader scenarios have been considered, and the differential-game approach has been applied successfully to derive equilibrium strategies [1, 16].

In the past decade, there have been increasing efforts in solving pursuit-evasion games that involve multiple pursuers or multiple evaders. When the game is set up so that a group of agents faces a single opponent, the optimal solution leads to coordination strategies for the group. There are works that look at the problem from the perspective

* Supported by ARL grant ARL DCIST CRA W911NF-17-2-0181.

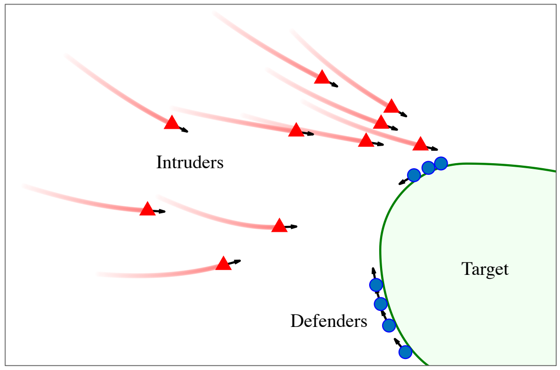

Fig. 1: Illustration of a perimeter defense game. Multiple intruders are approaching the target, while the defenders are tasked to intercept them.

of the pursuers [29, 30, 35], and also those that consider the problem from the evaders' side [6, 12, 24].

The derivation of the optimal strategies becomes more challenging when the game is played between teams of agents: i.e., multiple pursuers and multiple evaders [11]. The main challenge is in the dimensionality of the state space that prohibits us from naively applying the differential-game techniques. Similar to other multi-agent problems, the Voronoi tessellation has been widely employed to reduce the large-scale problems into a local area-minimization problem, or to assign pursuers to evaders [15, 18, 20, 39].

When there is a target region to be protected, the problem becomes a variant of the *target guarding problem*, originally introduced in [16]. A version of this problem is called the target-attacker-defender game [9, 17, 22], which has relevance to missile guidance applications [9, 21]. A similar scenario is also called the reach-avoid game [5, 37, 38], and it has been studied in many different variants [2, 3, 14, 34–36], including coast-line guarding or boarder defense [7, 10, 13, 31], and three-dimensional environments [8, 33].

This paper is concerned with the perimeter defense game, which is a variant of the target guarding problem where defenders are constrained to move on a convex target region [25–28]. This additional constraint leads to convenient closed-form solutions, as well as their interesting geometric interpretations. Moreover, the problem has high relevance to scenarios involving ground vehicles defending a building, or aerial vehicles patrolling around a no-fly zone.

This paper serves as an overview of various results and developments presented in our previous publications [25–28]. After formulating the perimeter defense problem in Section 2, we review solution methods and the associated results in Section 3. The potential extensions and generalizations of those results are considered in Section 4. Finally, Section 5 discusses the limitations in the current work and suggests directions for future research.

2 Problem Statement

We consider planer dynamics of point particles representing N_A intruders and N_D defenders. The symbols $_A$ and $_D$ are reserved to denote the intruders and the defenders respectively. The agents all have simple dynamics (i.e., first-order integrator), implying that they can change their velocity instantaneously, and this velocity vector is the control variable for each agent. Without the loss of generality, the intruders and the defenders have bounds on there speed, ν and 1, respectively. The parameter ν has a constraint $\nu \leq 1$, implying that the intruders are at most as fast as the defenders.

We make the following assumptions: (A1) the defenders move along the perimeter of the target region; (A2) the target region is convex; (A3) each agent has access to full-state information (i.e., positions of all agents); and (A4) intercept or capture occurs when the distance between the intruder and the defender becomes zero.

Viewing the game from an individual intruder's perspective, either one of the following three happens: (i) it reaches the perimeter without being intercepted by any of the defenders, (ii) the distance with the defender becomes zero when it reaches the perimeter, or (iii) it does not reach the perimeter in finite time. The third case is caused by the defender's maneuver to place itself between the target and the intruder. Therefore, we define both (ii) and (iii) as *capture*. Any intruder that achieves the outcome (i) will score a point.

Viewing the game in the team vs. team level, the objective function Q is the number of intruders that score, i.e., reach the perimeter without being captured by the defenders (outcome (i) above). The intruder team seeks to maximize Q while the defender team seeks to minimize Q. The problem is to find the equilibrium strategies that give

$$\min_{\boldsymbol{\Gamma}_D} \max_{\boldsymbol{\Gamma}_A} Q = \max_{\boldsymbol{\Gamma}_A} \min_{\boldsymbol{\Gamma}_D} Q,$$

where $\boldsymbol{\Gamma}_D$ and $\boldsymbol{\Gamma}_A$ denote the defender and intruder team strategies respectively.

3 Solution Method

A common approach in the variants of the multi-player target guarding problem is to leverage the results of the games played between a small number of players. The low-level velocity control strategies obtained in this small-scale game are combined with the high-level assignment policies to design the team-level strategies.

3.1 Agent-level Control Policy

One vs. One The smallest problem is the game played between one defender and one intruder. The key result that we need from this small-scale problem is the solution to the *game of kind*, which consists of the *barrier* surface and the corresponding control strategies. The barrier surface divides the state space into two regions: the intruder-winning region and the defender-winning region. If the initial configuration of the game falls in the intruder-winning region, then the intruder has a strategy to score, whereas if the game starts in the defender-winning region, then the defender has a strategy to guarantee capture.

To derive the barrier surface, one can consider a *game of degree* in which the payoff J is defined as the distance between the defender and the intruder at the time of breaching, i.e., when the intruder reaches the perimeter. The intruder tries to maximize this distance while the defender tries to minimize it. If an equilibrium exists, the Value function is defined as

$$V = \min_{\omega_D} \max_{\mathbf{v}_A} J = \max_{\mathbf{v}_A} \min_{\omega_D} J, \tag{1}$$

where ω_D denotes the defender's control on a one-dimensional space, and $\mathbf{v}_A$ denotes intruder's velocity. Once the equilibrium strategies ω_D^* and $\mathbf{v}_A^*$ are derived, the Value V is a function of the player positions, and the barrier surface can be identified as the level set $V = 0$.

The barrier for a circular perimeter was first derived in [25] using geometric approach, and later verified with differential game approach in [32]. The equilibrium defender strategy is to simply move in the direction of the intruder. Using the polar angle shown in Fig. 2, the equilibrium defender strategy is

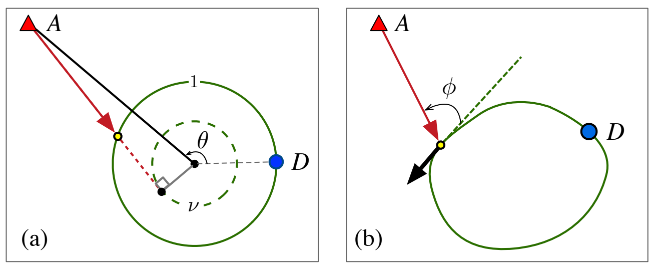

Fig. 2: Construction of the intruder strategies. (a) For circular perimeters, circles scaled by ν are used. (b) For non-circular perimeters, the approach angle ϕ is used.

$$\omega_D^* = \text{sign}(\theta).$$

The intruder's equilibrium strategy is to move towards the tangent point on a circle scaled by the speed ratio ν [25, 32].

The above defender strategy for the circular perimeter is almost trivial since any intruder position that gives $\theta = 0$ clearly defines the "front" of the defender's position. When considering a non-circular perimeter (e.g., polygon), this surface that divides the left and the right side from the defender's position is not immediately obvious. In addition, the intruder strategy can no longer be parameterized by the polar angle.

In [26], these issues are addressed by parameterizing the intruder strategies based on the *approach angle*, which is defined by the angle between the intruder's direction of motion and the tangent vector of the curve at the breach point. The optimal approach angle was derived to be

$$\phi = \cos^{-1} \nu.$$

This result tells us that the intruder should approach the tangent point on the target when $\nu = 1$, whereas it should approach the closest point on the target when $\nu \to 0$. It is easy to verify that this result matches with the special case where the target is circular.

The barrier is also obtained for the case when the perimeter is some arbitrary convex shape [26]. A convenient way to visualize the barrier surface is to look at the slice of the state space at a particular defender position. Then the barrier for that particular defender position is shown as a closed curve that completely surrounds the target region (Fig. 3a).

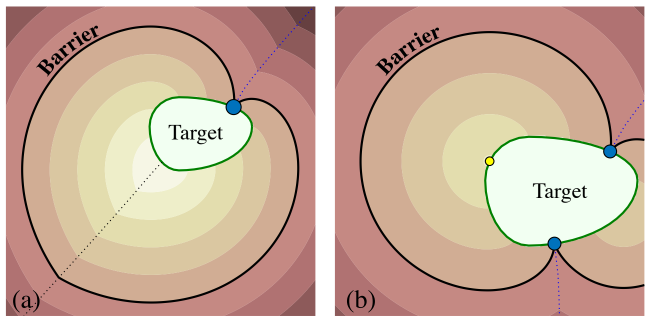

Fig. 3: Contour of the Value function and the barrier curve for a non-circular perimeter. The region enclosed by the barrier is the intruder-winning region, and the region outside is the defender-winning region. (a) One vs. one game. (b) Two vs. one game.

The results obtained in the above accommodate any value of the speed ratio $\nu \in (0,1]$, however, the result for the case with $\nu = 1$ lends itself to a nice geometric interpretation. Firstly, the intruder's equilibrium strategy is to move towards the tangent of the target region (for both circular and non-circular perimeters). Secondly, the barrier curve is constructed by pieces of geometry called the *involute*. An involute of a convex shape is the locus of a point on a piece of taut string as the string is either unwrapped from or wrapped around the shape. In our case, the shape of interest is the target region.

Two vs. One To consider an explicit form of cooperation among the defenders, we extended the one vs. one results into a game played between two defenders and one intruder [25, 26, 32]. Since the intruder must avoid both defenders, the optimal breach point against one defender may be suboptimal against anoteher defender. Depending on the configuration, the safest breach point now becomes either the optimal breach point against the closer defender or the mid-point between the two defenders. The defender pair's optimal strategy is to approach the intruder from both sides, which we call the *pincer movement*. Importantly, the defender-winning region provided by the pincer movement is larger than the union of the winning regions provided by the individual defenders [25, 28, 32]. The barrier derived for the two vs. one game has a greater implication beyond the additional consideration in the assignment strategy. The intruder-winning regions for the two vs. one game gives us a way to directly analyze an arbitrary size of the game played between n_D defenders and n_A intruders, which will be discussed at the end of Sec. 3.2.

3.2 Team-level Coordination Policies

The results from small scale games are used to develop team strategies when there are N_D defenders and N_A intruders. Let Q denote the number of intruders that reach the perimeter without being captured. Let the team strategies $\boldsymbol{\Gamma}_D$ and $\boldsymbol{\Gamma}_A$ denote the mappings from the current states (positions of all the agents) to the control actions $\boldsymbol{\omega}_D = [\omega_{D_1}..,\omega_{D_{N_D}}]$ and $\mathbf{v}_A = [\mathbf{v}_{A_1}..,\mathbf{v}_{A_{N_A}}]$ respectively. The large-scale game uses this score Q as the payoff function, which the defender team minimizes and the intruder team maximizes. If an equilibrium exists, the Value function is

$$Q^* = \min_{\boldsymbol{\Gamma}_D} \max_{\boldsymbol{\Gamma}_A} Q = \max_{\boldsymbol{\Gamma}_A} \min_{\boldsymbol{\Gamma}_D} Q. \tag{2}$$

The goal is to find the equilibrium team strategies $(\boldsymbol{\Gamma}_D^*, \boldsymbol{\Gamma}_A^*)$ and the Value Q^*.

Since directly solving this large-scale game is very challenging, various approximation methods have been proposed. We describe three approaches in analyzing the multi-player games represented by: maximum matching (MM), maximum independent set (MIS), and the local game regions (LGR).

Man-to-man Defense The results of the one-defender vs. one-intruder game immediately leads to a naive coordination strategy for the defender team. We can assign each defender to an intruder that it can capture, i.e., one in the corresponding defender-winning region. What we must avoid in these assignments is any overlap in the defenders or the intruders. More specifically, we must not assign a single defender to multiple intruders, since capture is only guaranteed against one intruder. In addition, we must not assign multiple defenders to a single intruder, since such redundant assignments will reduce the overall number of capture, leading to a higher score.

The optimal assignment of defenders to intruders in the man-to-man defense framework is provided by *matching* in graph theory. Considering a bipartite graph, where one set of nodes represents the defenders and the other represents the intruders, we draw edges from each defender to all the intruders that it can capture. By performing *maximum-cardinality matching* (MM) on this bipartite graph, we obtain the optimal man-to-man defense that gives us an upper-bound on the score

$$Q^* \leq Q_{\mathrm{MM}} = N_A - N_{\mathrm{MM}}, \tag{3}$$

where N_{MM} denotes the number of matches.

This general approach was originally proposed in [3, 4], and it has been applied to other variants of the problems [25, 36]. Due to its simplicity, there is a potential for application to many other scenarios, which will be discussed further in Sec. 4.1. The downside of this approach is that this naive coordination strategy does not account for a tighter cooperation that can happen between the defenders or the intruders. The cooperation between the defenders was incorporated by leveraging the result of two-defender vs. one-intruder game.

Two-on-one Defense A naive extension of the maximum-cardinality matching approach is to incorporate the results of the two vs. one game by assigning pairs of defenders to intruders. We are tempted to simply augment the bipartite graph with nodes representing the pairs of defenders, and adding edges from each defender pair to all the intruders that it can capture. However, we cannot perform maximum-cardinality matching on this bipartite graph since a node representing a single defender and one that represents a pair may share the same defender, which leads to an overlapping assignment.

To circumvent this problem, we can construct a new graph in which each node represents an assignment in the original graph, and the edges represent any conflict between the assignments: i.e., whenever two assignments share any defender or intruder, they are connected. By solving the Maximum Independent Set (MIS) problem on this transformed graph, we obtain the set of assignments without any overlapping defenders or intruders. This assignment method guarantees that the number of capture is greater than the one provided by the maximum matching, i.e., the provided upper bound is tighter:

$$Q^* \leq Q_{\text{MIS}} = N_A - N_{\text{MIS}} \leq Q_{\text{MM}}, \tag{4}$$

where N_{MIS} denotes the cardinality of the maximum independent set. However, the down side of this method is its computational complexity. Since MIS problem is NP-hard, it does not scale well with large number of agents. An improvement over the MM and MIS approaches is presented next.

Local Game Decomposition As was mentioned in Sec. 3.1, the barrier curve for the two vs. one game leads to a stronger coordination policy beyond two-on-one assignments. Consider a specific pair of defenders and the intruder-winning region that it defines (Fig. 3b or Fig. 4). We call it the *Local Game Region* (LGR) as it leads to a region-based decomposition method [28].

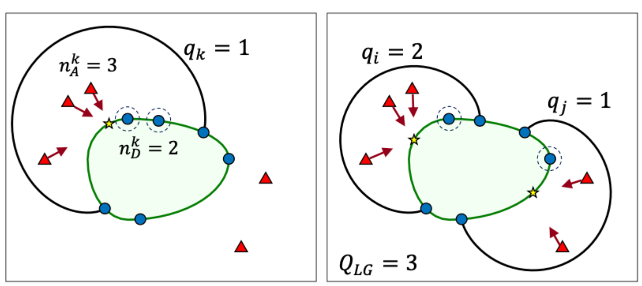

Fig. 4: Local game regions (LGRs) and the associated subteams. (a) An LGR that contains three intruders and two defenders. (b) The intruder team's optimal selection of the LGRs, $\mathbf{G}^*$, leading to the overall score of $Q = 3$.

The results of the two vs. one game tell us that the n_A^k intruders contained in this LGR can win against the defender pair, by approaching near the mid-point. Additionally, this result implies that any defender outside of this LGR cannot capture any of those intruders. Therefore, only those n_D^k defenders contained in the LGR can possibly contribute to capture. If the intruders have a local numerical advantage, i.e.,

$q_k = n_A^k - n_D^k > 0$, then the intruders can achieve a score of q_k. The subscript k denotes the indexing of the defender pairs, or the LGRs.

Now, the intruder team can maximize the guaranteed score by intelligently selecting the set of LGRs in which they play the local game [28]. Let $\mathbf{G}$ denote the set of disjoint LGRs, then the intruder team can consider the following optimization problem:

$$Q_{\mathrm{LG}} = \max_{\mathbf{G}} \sum_{k \in \mathbf{G}} q'_k, \tag{5}$$

where $q'_k = \max\{0, q_k\}$.

Theorem 1 (From [28]). *Let $\boldsymbol{\Gamma}_A^*$ denote the intruder strategy corresponding to the teaming into $\mathbf{G}^*$ and then playing the two vs. one game against the defender pair associated to each LGR. The strategy $\boldsymbol{\Gamma}_A^*$ guarantees*

$$Q^* \geq Q_{LG} \tag{6}$$

for all permissible defender strategies. ■

Importantly, this is the first result in the team-level coordination policy that gives an intruder team strategy and a corresponding lower bound on the score. In [28] we also propose a defender team strategy $\boldsymbol{\Gamma}_D^*$ based on the local game decomposition. The strategy first removes any uncapturable intruders based on the result of Theorem 6. Then the two-on-one defense are assigned to intruders that are near the boundary of the LGRs with the score $q_k = 0$. Intuitively, this assignment is used to ensure that no additional intruder enters the LGRs that already has an equal number of intruders and defenders, which we call the *occupied* LGRs [28]. Finally, one-on-one defense is assigned to the remaining agents. In [28], we explain how this strategy accounts for the sequence of captures that occur in time, in contrast to the MM and MIS approaches that are ignorant of such dynamical aspects. Finally, [28] discusses the class of initial configurations from which the following relationship is guaranteed

$$Q(\boldsymbol{\Gamma}_D^*, \boldsymbol{\Gamma}_A) \leq Q(\boldsymbol{\Gamma}_D^*, \boldsymbol{\Gamma}_A^*) \leq Q(\boldsymbol{\Gamma}_D, \boldsymbol{\Gamma}_A^*), \tag{7}$$

which implies that the strategies are at an equilibrium.

4 Extensions and Generalizations

This section discusses potential extensions and generalizations of the results described in the previous section.

4.1 Assignment-based Defense Policies

The solution method involving the one vs. one results and the maximum-cardinality matching (MM) gives us a generic framework to solve various swarm vs. swarm problems. An immediate extension of the perimeter defense problem is the engagement between intruders and defenders that can move in three-dimensional space. A similar problem has been considered for agents that can freely move in three-dimensional space, with target region defined as a plane [33].

A direct extension of the perimeter defense scenario is the *hemisphere guarding problem*, in which defenders that move on a hemisphere seek to intercept aerial intruders. An initial step to solve this problem is presented in [23], where one vs. one engagement between an aerial defender and a ground intruder is considered. This is an intermediate step towards intruders that can freely move in three-dimensional space. Once we have the barrier surface from this small-scale problem, we can determine whether the intruder is capturable or not for every defender-intruder pair. This information allows us to convert the design of strategies in continuous velocity space into a matching problem, as was discussed in Sec. 3.2.

Heterogeneous Speed The assignment approach also extends to teams with heterogeneous capabilities. For example, defenders may have various speed limits. In this case, the barrier surface will be different for every defender. However, this variation does not affect the assignment method, since all we need in the high-level matching algorithm is the pair-wise information of whether a given defender can capture a given intruder or not. Similarly, the intruders may have different speeds without affecting the overall solution method. It is worth noting that the closed-form solution for the one vs. one game provided in our work becomes useful in the presence of these speed variations, since the required win/loss information can be obtained computationally efficiently. When the numerical approach with HJI PDE is used (as in [3,4]), the extension becomes more challenging because the barrier surfaces for all possible values of ν have to be computed a priori and be stored to be used as a look-up table.

Intruder Weights Another immediate extension of the assignment based approach is the case when different intruders carry different weights in terms of the damage they incur when the perimeter is breached. For example, the threat level of each intruder may be described by some weighting factor. In this case, the maximum cardinality matching will be modified to a Linear Sum Assignment Problem that can be solved by the well-known Hungarian Algorithm.

MIS approach The solutions to the two-defender vs. one-intruder game were used to improve the defender team's performance by considering explicit cooperation among pairs of defenders. A similar approach will apply for other problems if any n-defender vs. m-intruder subproblem is identified and solved. In the transformed graph, each node represents an assignment of a team of n defenders to a group of m intruders, and each edge represents a conflict. Each node (assignment) also carries a weight based on the number of intruders or their weighted sum according to the relative threat level as was discussed with intruder weights. Each edge represents a conflict whenever two different assignments contain the same defender or the intruder. By finding the Maximum-weight Independent Set on this graph, we can optimize the defender to intruder assignments that account for explicit cooperation.

4.2 Cooperative Intruder Strategies

The tasks given to the defender team and the intruder team are fundamentally different, and it is easy to see that the assignment-based approaches do not work well for the intruder team. More specifically, each intruder cannot simply select a single defender and play a one vs. one game against it. Instead, an intruder has to avoid all defenders in order to reach the target and score a point. This asymmetry makes it difficult for us to design a cooperative team strategy for the intruders.

The decomposition method using the local game regions (LGRs) presented in [28] was most useful in the sense that it led to a cooperative intruder strategy, which was not provided by either the MM or the MIS formulations. The essence of this approach is in the analysis based on local *numerical advantages*. In the perimeter defense game, a local overmatch situation was created by the intruder subteams that simultaneously attack a single point on the perimeter.

A similar strategy will likely to appear in the multi-player reach avoid game if it is assumed that the defender is consumed by the intruder whenever a capture occurs, i.e., a single defender cannot capture multiple intruders. The main challenge in extending the LGR decomposition to the reach avoid game or any other variants of the target guarding problem is in the definition of the regions that leads to an efficient combinatorial optimization. Each region must have two elements: (a) an associated intruder subteam that, and (b) a set of defenders that every intruder in the subteam can win against. The latter gives us the complement of the defender subteam associated to the region, which leads to the sufficient condition for the intruder subteam to score: $n_A^k > n_D^k$.

5 Limitations and Future directions

In this section we discuss the limitations in the existing works and consider directions for future developments.

Defender Dynamics We assumed that the defender's motion is constrained on the perimeter. The most significant consequence of this constraint is the fact that the intruders can ensure no defender captures more than one intruder, by approaching the perimeter simultaneously. This property makes the perimeter defense game to be interesting only when $N_A < N_D$, since the intruder team can always guarantee a score otherwise.

Sequential Capture If we allow the defenders to leave the perimeter and pursue the intruder, there is a possibility of multiple captures achieved by a single defender. Such scenario leads to several interesting avenues for future work. First, the defender team strategy will now contain a *vehicle routing problem*, which is a complex combinatorial optimization problem even in the case where the locations to be visited by the agents are stationary. Secondly, the intruder behavior will now have a more meaningful distinction between an evasive maneuver and an offensive maneuver. More specifically, some of the intruders may try to lure the defenders away from the perimeter so that other intruders can successfully reach the target. Such cooperation schemes have been considered for small-scale games, e.g., [29]. The design of defender's strategy will also become complex because there is a potential conflict between (i) pursuing an intruder, and (ii) staying close to the target for subsequent captures. The design of coordination strategies for large teams is therefore an interesting problem from the perspective of both the defenders and the intruders.

Fast Intruders The existing results are provided for any speed ratio $\nu \in (0, 1]$, i.e., the intruders cannot be faster than the defenders. It is easy to see that this assumption is necessary for one vs. one game since if the intruder has a speed advantage, then it can come arbitrarily close to the perimeter and outrun the defender to always guarantee intrusion. This is still the case even when there are multiple defenders. The degeneracy arises due to the combination of assumptions (A1) and (A4) in Sec. 2. To consider a faster intruder scenario, we must relax either or both of those two assumptions (e.g., consider nonzero capture radius) and design strategies for the defenders to coordinate their motion.

Partial Information Finally, one of the strongest assumptions in the existing variants of the target guarding problems is the assumption that each agent knows the location of all other agents (i.e., full state information). In a realistic setting the detection of the intruders is an important consideration. As a step towards a more holistic solution, [27] considers the design of a patroller team that ensures the detection of any agent that approaches a certain distance from the perimeter. Additionally, for large-scale problems, the defenders must be able to act on locally collected information. The work in [19] developed a training methodology to learn communication strategies from a centralized expert policy. The proposed method was used to design a decentralized version of the MM assignment strategy without any centralized coordination mechanism.

Importantly, neither of the above works directly addressed the issue of unseen intruders. The consideration of such partial information will make any pursuit-evasion game much more difficult to solve. When an equilibrium is infeasible to obtain, a design of reasonable team strategies will be of great importance to the field.

6 Conclusion

This paper presented an overview of the works done on the multi-agent perimeter defense game, in which a team of intruders seek to reach a target region while a team of defenders try to intercept those intruders. The multi-player game can be solved by first deriving the strategies for the games played between a small number of agents (one vs. one and two vs. one) and leveraging those results in the team-level coordination strategies. Three different analysis methods (maximum matching, maximum independent set, and local-game decomposition) have varying levels of cooperation. While the local-game decomposition generates the most sophisticated team behavior, the maximum matching has the most straightforward ways of being applied to other problems. There are several limitations in the current problem formulation that lead us to future developments in the aspects related to the agent dynamics, payoff function, and the information structure.

References

1. Basar, T., Olsder, G.J.: Dynamic Noncooperative Game Theory, 2nd Edition. Society for Industrial and Applied Mathematics (2011)
2. Chen, M., Zhou, Z., Tomlin, C.J.: A path defense approach to the multiplayer reach-avoid game. IEEE Conf. Decis. Control (CDC) pp. 2420–2426 (2014)
3. Chen, M., Zhou, Z., Tomlin, C.J.: Multiplayer reach-avoid games via low dimensional solutions and maximum matching. Proc. Amer. Control Conf. (ACC) pp. 1444–1449 (2014)
4. Chen, M., Zhou, Z., Tomlin, C.J.: Multiplayer reach-avoid games via pairwise outcomes. IEEE Trans. Autom. Control **62**(3), 1451–1457 (mar 2017)
5. Fisac, J.F., Chen, M., Tomlin, C.J., Sastry, S.S.: Reach-avoid problems with time-varying dynamics, targets and constraints. Proc. 18th Int. Conf. Hybrid Sys. Comp. Control (ACM) pp. 11–20 (2015)
6. Fuchs, Z.E., Khargonekar, P.P., Evers, J.: Cooperative defense within a single-pursuer, two-evader pursuit evasion differential game. IEEE Conf. Decis. Control (CDC) pp. 3091–3097 (2010)
7. Garcia, E., Casbeer, D.W., Pachter, M.: The barrier surface in the cooperative football differential game. arXiv preprint arXiv:2006.03682 (2020)
8. Garcia, E., Casbeer, D.W., Pachter, M.: Optimal strategies for a class of multi-player reach-avoid differential games in 3d space. IEEE Rob. Autom. Lett. **5**(3), 4257–4264 (2020)
9. Garcia, E., Casbeer, D.W., Pham, K., Pachter, M.: Cooperative aircraft defense from an attacking missile. J. Guid. Control Dyn **38**(8), 1510–1520 (2015)
10. Garcia, E., Casbeer, D.W., Von Moll, A., Pachter, M.: Cooperative two-pursuer one-evader blocking differential game. In: Proc. Amer. Control Conf. (ACC). pp. 2702–2709. IEEE (2019)
11. Garcia, E., Casbeer, D.W., Von Moll, A., Pachter, M.: Multiple pursuer multiple evader differential games. IEEE Trans. Autom. Control (2020)
12. Garcia, E., Von Moll, A., Casbeer, D.W., Pachter, M.: Strategies for defending a coastline against multiple attackers. In: IEEE Conf. Decis. Control (CDC). pp. 7319–7324. IEEE (2019)
13. Garcia, E., Von Moll, A., Casbeer, D.W., Pachter, M.: Strategies for defending a coastline against multiple attackers. In: IEEE Conf. Decis. Control (CDC). pp. 7319–7324 (2019)
14. Huang, H., Ding, J., Zhang, W., Tomlin, C.J.: A differential game approach to planning in adversarial scenarios: A case study on capture-the-flag. IEEE Int. Conf. Rob. Autom. (ICRA) pp. 1451–1456 (2011)
15. Huang, H., Zhang, W., Ding, J., Stipanović, D.M., Tomlin, C.J.: Guaranteed decentralized pursuit-evasion in the plane with multiple pursuers. IEEE Conf. Decis. Control (CDC) pp. 4835–4840 (2011)
16. Isaacs, R.: Differential games: A mathematical theory with applications to warfare and pursuit, control and optimization. Courier Corporation (1999)
17. Liang, L., Deng, F., Peng, Z., Li, X., Zha, W.: A differential game for cooperative target defense. Automatica **102**, 58–71 (2019)
18. Makkapati, V.R., Tsiotras, P.: Optimal Evading Strategies and Task Allocation in Multi-player Pursuit–Evasion Problems. Dynamic Games and Applications pp. 1–20 (2019)
19. Paulos, J., Chen, S.W., Shishika, D., Kumar, V.: Decentralization of multiagent policies by learning what to communicate. In: IEEE Int. Conf. Rob. Autom. (ICRA). pp. 7990–7996. IEEE (2019)
20. Pierson, A., Wang, Z., Schwager, M.: Intercepting rogue robots: An algorithm for capturing multiple evaders with multiple pursuers. IEEE Rob. Autom. Lett. **2**(2), 530–537 (2017)

21. Prokopov, O., Shima, T.: Linear Quadratic Optimal Cooperative Strategies for Active Aircraft Protection. J. Guid. Control Dyn **36**(3), 753–764 (2013)
22. Rubinsky, S., Gutman, S.: Three-Player Pursuit and Evasion Conflict. J. Guid. Control Dyn **37**(1), 98–110 (2014)
23. S. Lee, E., Shishika, D., Kumar, V.: Perimeter-defense game between aerial defender and ground intruder. In: To appear in IEEE Conf. Decis. Control (CDC). IEEE (2020)
24. Scott, W.L., Leonard, N.E.: Optimal evasive strategies for multiple interacting agents with motion constraints. Automatica **94**, 26–34 (2018)
25. Shishika, D., Kumar, V.: Local-game decomposition for multiplayer perimeter-defense problem. In: IEEE Conf. Decis. Control (CDC). pp. 2093–2100 (2018)
26. Shishika, D., Kumar, V.: Perimeter-defense game on arbitrary convex shapes. arXiv (arXiv:190903989, 2019)
27. Shishika, D., Paulos, J., Dorothy, M.R., Hsieh, M.A., Kumar, V.: Team composition for perimeter defense with patrollers and defenders. In: IEEE Conf. Decis. Control (CDC). pp. 7325–7332 (2019)
28. Shishika, D., Paulos, J., Kumar, V.: Cooperative team strategies for multi-player perimeter-defense games. IEEE Rob. Autom. Lett. **5**(2), 2738–2745 (2020)
29. Von Moll, A., Casbeer, D., Garcia, E., Milutinović, D., Pachter, M.: The multi-pursuer single-evader game. J. Intel. Rob. Syst. **96**(2), 193–207 (2019)
30. Von Moll, A., Casbeer, D.W., Garcia, E., Milutinović, D.: Pursuit-evasion of an evader by multiple pursuers. In: Int. Conf. Unmanned Aircraft Syst. (ICUAS). pp. 133–142. IEEE (2018)
31. Von Moll, A., Garcia, E., Casbeer, D., Suresh, M., Swar, S.C.: Multiple-pursuer, single-evader border defense differential game. J. Aero. Info. Syst. pp. 1–10 (2019)
32. Von Moll1, A., Pachter, M., Shishika, D., Fuchs, Z.: Guarding a circular target by patrolling its perimeter. In: To appear in IEEE Conf. Decis. Control (CDC). IEEE (2020)
33. Yan, R., Duan, X., Shi, Z., Zhong, Y., Bullo, F.: Maximum-matching capture strategies for 3d heterogeneous multiplayer reach-avoid games. arXiv preprint arXiv:1909.11881 (2019)
34. Yan, R., Shi, Z., Zhong, Y.: Escape-avoid games with multiple defenders along a fixed circular orbit. In: 13th IEEE Int. Conf. Control & Autom. (ICCA). pp. 958–963. IEEE (2017)
35. Yan, R., Shi, Z., Zhong, Y.: Reach-avoid games with two defenders and one attacker: An analytical approach. IEEE Trans. Cybern. **49**(3), 1035–1046 (2018)
36. Yan, R., Shi, Z., Zhong, Y.: Task assignment for multiplayer reach–avoid games in convex domains via analytical barriers. IEEE Trans. Rob. **36**(1), 107–124 (2019)
37. Zhou, Z., Ding, J., Huang, H., Takei, R., Tomlin, C.: Efficient path planning algorithms in reach-avoid problems. Automatica **89**, 28–36 (2018)
38. Zhou, Z., Takei, R., Huang, H., Tomlin, C.J.: A general, open-loop formulation for reach-avoid games. In: IEEE Conf. Decis. Control (CDC). pp. 6501–6506 (2012)
39. Zhou, Z., Zhang, W., Ding, J., Huang, H., Stipanović, D.M., Tomlin, C.J.: Cooperative pursuit with Voronoi partitions. Automatica **72**, 64–72 (2016)