Proceedings of ASME 2017 Dynamic Systems and Control Conference
DSCC 2017
October 11-13, 2017, Tysons, Verginia, USA

DSCC2017-5133

COMPETING SWARMS OF AUTONOMOUS VEHICLES: INTRUDERS VERSUS GUARDIANS

Daigo Shishika, Katarina Sherman, and Derek A. Paley
Department of Aerospace Engineering and
the Institute for Systems Research
University of Maryland
College Park, Maryland 20742
Email: daigo.shishika@gmail.com, dpaley@umd.edu

ABSTRACT

We consider a competition between two swarms of aerial vehicles, where multiple intruder vehicles try to approach and then leave an area that multiple guardian vehicles are protecting. Pre-existing swarming strategies for the guardians to maximize the probability of capturing a single intruder are summarized. This work considers the case where multiple intruders approach the protected area sequentially with varied time intervals, to study the impact of intrusion frequency on the probability of capture. In addition, we formulate a payoff function treating the competition as a zero-sum game, and use this function to design strategies for the intruders, i.e., how to optimize the time interval between intrusions. We propose an intrusion strategy and demonstrate its performance with numerical simulations.

INTRODUCTION

Pursuit-evasion games involve many different variations depending on the assumptions made on the pursuer, evader (or target), and the environment. Taxonomy and surveys of the research in the field have been presented, for example, in [1] and [2]. In the past decade, pursuit-evasion problems that involve multiple players are gaining increasing research interest. For example, strategies to encircle a target with a team of pursuers are proposed in [3, 4]. From the evasion perspective, the optimal strategies for multiple interacting agents have been studied too [5].

Although the aforementioned works consider strategies for a group of agents against a single agent, there are also studies of the competitions between two groups of agents. For example, the multiplayer reach-avoid game deals with a scenario where one group seeks to reach an area quickly while the other group tries to delay or prevent it [6–8].

One example of a reach-avoid game is the so-called Capture-the-Flag (CTF) [9]. In this game, each team owns a flag, and their goal is to capture the opponent's flag and safely return. This game has a pursuit-evasion aspect, since an agent can be intercepted by its opponent in the opponent's territory. However, the capture-the-flag problem is very complex because it also involves an attacking-defending aspect and, possibly, switching of the agent roles, as well as the assignment problem to decide which agent should pursue which opposing agent. Some works have tackled this complex problem by combining the tools from differential games and graph theory [6, 9], whereas some other works made some simplifying assumptions on the attackers' dynamics and used optimization techniques [7, 8].

Agent dynamics are typically modeled as first order (i.e., control input is velocity) and target capture is often defined by interception (i.e., distance becomes sufficiently small) [6–9]. In contrast, we previously considered a related problem from a dynamics perspective, inspired by the swarming behavior of mosquitoes [10, 11]. Male mosquitoes aggregate and form mating swarms to attract female mosquitoes that fly faster than the males. When the female enters the swarm, a male's pursuit behavior is triggered when the distance to the female becomes small, which we call a *close encounter*. After the pursuit phase, the male and female exhibit coupling flight during which they fly

in approximately the same direction [12]. For a male to achieve coupling flight, simply intercepting a female is insufficient; he also has to align his velocity with the female. Since the female flies faster than a swarming male, a male has to accelerate after the close-encounter in order to successfully track the female. Therefore, the mosquito pursuit is modeled by agents with variable speed.

The combination of limited perceptual range, target tracking instead of interception, and the dynamical model of the agent raises the importance of quick response, i.e., when a male detects a female, it has to speed up and match the velocity of a fast female before the female escapes from the male's perceptual range. Inspired by these characteristics, we previously considered the scenario in which multiple guardians with limited perceptual range and bounded acceleration are deployed to protect an area from an intruder [13, 14]. We focused on the guardian's side and proposed various swarming strategies to increase the probability of capturing the intruder.

This work considers the scenario of a zero-sum game between a guardian team and an intruder team. We pit intruder strategies against the guarding strategies considered previously. The contributions of this work are (1) formulation of a novel pursuit-evasion game involving multiple intruders and guardians; (2) theoretical predictions of the outcome of the game; and (3) the design of intrusion strategy and its demonstration with computer simulation. The problem studied in this work can be applied to a situation where multiple vehicles are deployed to enforce a no-fly zone, for the application to drone countermeasures, or for convoy protection. This work may also provide a guideline in selecting the capabilities of the vehicles for such applications, and provide a methodology to fully utilize those capabilities.

PROBLEM FORMULATION

Consider a planar system of point particles with unit mass representing N_P guardians and N_T intruders (we use the subscripts $_T$ and $_P$ to denote the intruder/target and guardian/pursuer, respectively). The intruders seek to pass through (approach and then safely leave) a region that is protected by the guardians. Figure 1 illustrates the case with only one intruder. The timing and the direction of the intruder trajectories are unknown to the guardians. Once the intruder enters the perceptual range of a guardian, the roles of the agents change—the intruder becomes a target and the guardian becomes a pursuer. The goal of the pursuer is to capture the target (i.e., approach the target and stay close to it).

Agent Dynamics and Sensing Capabilities

Consider the case where the protected region is sufficiently small to be approximated as a point O. Let O to be the origin of the inertial frame; $\mathbf{r}_i$, $\mathbf{v}_i$, and $\mathbf{a}_i$ denote the position, velocity, and acceleration of particle i in the inertial frame. The guardians have second-order dynamics, i.e., $\dot{\mathbf{r}}_i = \mathbf{v}_i$ and $\dot{\mathbf{v}}_i = \mathbf{a}_i$. In addition, assume the following capabilities of the guardians:

(A1) the magnitude of the guardian's acceleration is bounded according to $\|\mathbf{a}_P\| \leq u_{\max}$; and
(A2) each guardian perceives the position and velocity of all other agents within the range ρ_a.

Based on the perceptual range ρ_a, consider target capture defined as follows.

Definition 1. *Let $\mathbf{r}_{T/P} = \mathbf{r}_T - \mathbf{r}_P$ denote the relative position of the target with respect to the pursuer. Target capture is successful if there exists t_{cap} such that $\|\mathbf{r}_{T/P}\| < \rho_a$, for all $t > t_{\text{cap}}$.*

This definition requires the guardian to track the target, in contrast to target intercept where distance condition has to be satisfied only instantaneously [6–9].

We also introduce another perceptual range that determines when the pursuit behavior is triggered:

(A3) Each guardian becomes a pursuer once the distance to an intruder becomes less than ρ_p, which we call the *close-encounter distance*.

Note that the parameter ρ_p permits two interpretations. First, it can be interpreted as the limitation of the guardians to distinguish between a friendly guardian vehicle and the intruder, i.e., if $\rho_p < \|\mathbf{r}_{j/i}\| < \rho_a$, guardian i does not know whether an agent j (in its perceptual range) is an intruder or not. Second, ρ_p may be a control parameter that the guardian can choose: i.e., the guardian will ignore the intruder unless it gets closer than the distance ρ_p. In either case, the value of ρ_p will not exceed ρ_a.

For the intruder's motion, assume the following:

(A4) Each intruder approaches O at their maximum speed $\|\mathbf{v}_T\| = v_T$ on a straight path.

In other words, their trajectories are straight lines that pass through O. Assume that the intruders do not react to the guardians by changing its direction of motion. Although it seems that the intruders' behavior is highly restricted, the next section shows that they still have certain degrees of freedom in planning how they intrude O, by modifying the timing and the direction of their arrival.

Strategy Sets

The intrusion-capture problem can be separated into two parts. The first is the *swarming phase* in which the guardian does not know where the intruder is. Once the intruder enters the circle with radius ρ_p around a guardian, the *pursuit phase* starts. The continuous strategy set for each guardian is the commanded acceleration $\mathbf{a}_i$ during the above two phases, which were the main focus of previous work [13, 14].

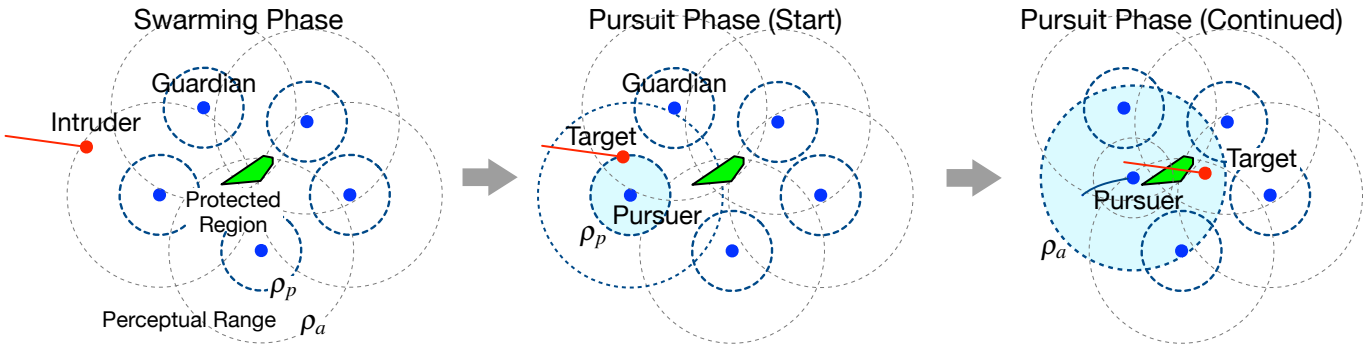

FIGURE 1. Illustration of the intrusion and pursuit scenario. In the swarming phase, an intruder (red) is approaching the protected region (green). The guardians (with static formation here for clarity) are deployed to wait for the intruder. Once the intruder enters the perceptual range, the guardian turns into a pursuer and the intruder becomes the target.

The success of target capture depends on how quickly a guardian can respond (i.e., close the distance and match the velocity) to the intruder once it is in perceptual range ρ_p. If the response is too slow, then the target will escape from the range ρ_a. Previous work [13] showed the importance of the swarming phase for guardians' success, and proposed strategies for how they should prepare for the intruder to maximize the probability of target capture.

This work focuses on the intruders' side, and considers their strategies as follows. Let $t_j^{\text{int}} \in [t_0, \infty)$ and $\psi_j^{\text{int}} \in [0, 2\pi)$ denote the time and azimuthal direction that the jth intruder arrives at O (assuming it is not captured), and let positive number $T_j^{\text{int}} = t_{j+1}^{\text{int}} - t_j^{\text{int}}$ denote the time interval between two successive intruders. The strategy sets for the intruders are $\Psi_{\text{int}} \triangleq \{\psi_j\} \in \mathbb{R}^{N_T}$ and $\mathbf{t}_{\text{int}} \triangleq \{t_j^{\text{int}}\} \in \mathbb{R}^{N_T}$, i.e., they can manipulate the timing and the direction that they approach O. We assume that the time of arrival $\mathbf{t}_{\text{int}}$ is constrained as follows:

(A5) The jth intruder vehicle becomes available for deployment only after a certain time, which allows it to reach O no earlier than s_j^{int} (i.e., $t_j^{\text{int}} \geq s_j^{\text{int}}$ as illustrated in Fig. 2).

The set $\mathbf{s}_{\text{int}} \triangleq \{s_j^{\text{int}}\}$ is introduced to reflect various time constraints, which may exist in practice, for the intruder vehicles to be ready for deployment, e.g., refueling, repair, or going to a set position to achieve a given ψ_j^{int}. If there is no such constraints, one can simply use $s_j^{\text{int}} = t_0$, $j = 1, 2, ..., N_T$.

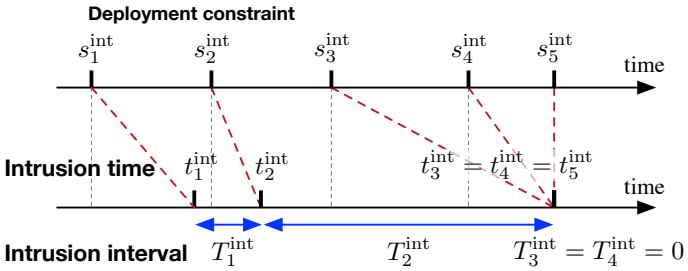

FIGURE 2. Illustration of how the deployment constraint $\mathbf{s}_{\text{int}}$ affects the choice of intrusion time $\mathbf{t}_{\text{int}}$, and how the intrusion interval $\mathbf{T}_{\text{int}}$ is computed from the intrusion time.

Payoff Function

We consider the intrusion-capture competition as a zero-sum game between intruders and guardians by considering the payoff function J_I and $J_G = -J_I$ associated to the team of intruders and guardians respectively. The payoff for the intruder team is the sum of the individual payoff for each intruder, i.e.,

$$J_I = \sum_i^{N_T} J_i, \tag{1}$$

where J_i is the payoff function associate to ith intruder. (Note that we occasionally omit the subscript that indicates J_i is for an intruder. Also note that the payoff for individual guardian is not defined, because the performance is considered in terms of the entire guardian swarm.)

The value of J_i is determined as described in the payoff matrix in Fig. 3. Successful intrusion, corresponding to the first row of the matrix, is the case where intruder i reaches O without being captured by any of the guardians. (Note, this definition is independent from the capture that might occur *after* the intrusion.) Once the intruder successfully reaches O, it scores $V_{int}(t)$, which is a decreasing function of time defined as follows:

$$V_{int}(t) = e^{-(t-t_0)/\tau}, \tag{2}$$

Intruder \ Guardian swarm	No capture	Capture
Intrusion	$V_{int}(\mathrm{t})$	$V_{int}(\mathrm{t})-\eta$
No intrusion	0	$-\eta$

FIGURE 3. Value of individual intruder payoff J_i depending on the performance of ith intruder and the guardian swarm.

where t_0 is the start time of the game and τ is the *intruder-payoff time constant*. If the intruder i misses O, or if it is captured by any of the guardians before reaching O, it does not score $V_{int}(t)$, corresponding to the second row of the matrix.

The parameter $\tau > 0$ models the risk of waiting incurred by the intruder's team. Consider, for example, the Capture-the-Flag problem [9] where two teams (A and B) divide their members into offensive and defensive players. The problem studied in our work can be considered as the competition between team A's offensive players (intruders) and team B's defensive players (guardians). In this case, there is another identical game played simultaneously where team A is guarding their flag against team B's offensive players. The decaying payoff $V_{int}(t)$ with the time constant τ quantifies the risk that A's flag is attacked by B within the duration of time $t - t_0$.

The use of an exponential function in (2) is inspired by the Poisson distribution (see Appendix), which is often used to model the probability of an event occurring in a fixed interval of time (e.g., an earthquake, or a customer entering a store). In our problem, we are modeling the attack from B's offensive players as a Poisson process, i.e., the probability that A's flag is *not* attacked by B decays exponentially. Furthermore, with this Poission-process assumption, the time constant τ corresponds to the expected time interval between successive attacks from B. Therefore, higher risk on team A (intruder's side) can be modeled by a smaller value of τ, i.e., faster decay in $V_{int}(t)$. Note, if the game is one sided and A does not have to defend their flag, then the constant can be set to $\tau = \infty$ so that $V_{int}(t) \equiv 1$.

Next, consider the performance of the guardian swarm. If any of the guardian captures the intruder i, then the guardian swarm scores $\eta > 0$ (equivalently, the ith intruder scores $-\eta$), which corresponds to the second column in Fig. 3. Note that capture may be before or after the intrusion. The parameter η describes the penalty on the intruder swarm to lose its vehicles. Depending on the value of η, the scenario can be categorized into the following three cases:

- $\eta = 0$: Pure guarding scenario. Guardians cannot win the game, but they lose unless they capture every target before intrusion.
- $0 < \eta < 1$: The payoff $J_i = V_{int}(t) - \eta$, which corresponds to capture after intrusion, changes its sign from positive to negative during the game.
- $\eta > 1$: The payoff $J_i = V_{int}(t) - \eta$ is already negative at the beginning of the game. There is no benefit for intruders to approach O unless they can escape without capture.

Various intrusion-capture scenario can be modeled by the proper choices of parameter η and τ. As one example, consider a scenario where guardian vehicles are protecting an area against bombers or missiles. What happens after the intrusion is not so important (guardians have to capture the intruders *before* they attack the area), so we use $\eta \ll 1$ for this case. Another example is the scenario where intruders are manned aircraft for a reconnaissance mission. The penalty for losing a vehicle will be particularly high if it is manned, so we use $\eta > 1$ for this case. Consequently, intruders will only attempt to approach if probability of capture is sufficiently low. The other parameter τ can be chosen to model how sensitive the intruder's mission is to time.

GUARDIAN STRATEGY

This section briefly summarizes our previous results on the analyses of the intrusion-capture problem, and guardian's swarming strategies [13, 14]. First, we introduce nondimensional parameters that describe which swarm (guardians or intruders) has the advantage in terms of their capabilities (i.e., dynamics and sensing). Second, we introduce guardian's swarming strategies, which we use as the baseline to test various intrusion strategies in the following section. Readers are also referred to [*https://youtu.be/Cnz75WZ88rI*] for an illustration of prior experimental results.

Nondimensional Parameters

Considering the individual encounter of a guardian and an intruder, the difficulty of target capture is given in terms of the intruder's speed v_T and the guardian's capability $u_{\max}$, ρ_a, and ρ_p. To explore this parameter space efficiently, we introduce the following two nondimensional parameters [13, 14]:

$$\alpha = \frac{\rho_p}{\rho_a} \quad \text{and} \quad \Gamma = \frac{2u_{\max}(\rho_a + \rho_p)}{v_T^2}. \tag{3}$$

The first parameter $\alpha \in (0, 1]$ is the *pursuit activation distance*, which describes the ratio between the two perceptual ranges defined in assumptions (A2) and (A3). The second parameter Γ is the nondimensionalized *guardian acceleration*, which describes the ratio between the guardian's capability and the intruder's speed.

A large value of Γ corresponds to the case where guardians have an advantage, i.e., they can accelerate quickly and/or sense the intruder from far away. On the other hand, a small value of

Γ corresponds to the case where the intruder moves so fast that it is hard to capture. We showed in [13] that $\Gamma > 1$ is a necessary condition for a static guardian to capture the target.

There exist sufficient conditions for target capture when the guardian's acceleration during pursuit phase is given by a force resembling a damped spring attached to the target, i.e., $\mathbf{a}_P = \mathbf{F}_P^{(\text{pursuit})} = c\mathbf{r}_{T/P} + b\mathbf{v}_{T/P}$, where c and b are control parameters. We previously showed that the pursuit is guaranteed if the relative velocity between the pursuer and the target, $\mathbf{v}_{T/P} = \mathbf{v}_T - \mathbf{v}_P$, at the time of close encounter satisfy the following condition:

$$\|\mathbf{v}_{T/P}\| \le v_0, \quad \text{where} \quad v_0 = v_T\sqrt{\frac{\Gamma(1-\alpha)}{2}}. \tag{4}$$

(See Proposition 2 in [13] for details.) The condition (4) states that the guardian's velocity $\mathbf{v}_P$ at the time of close encounter has to be sufficiently aligned with the intruder's velocity $\mathbf{v}_T$ to guarantee target capture. This condition was also used as a practical definition for target capture:

Definition 2. *Target j is captured if the condition $\|\mathbf{v}_{j/i}(t)\| \le v_0$, where v_0 is defined in* (4)*, is satisfied at any point in time with any guardian $i \in \{1,2,...,N_P\}$.*

Definition 2 gives an instantaneous representation of target capture, in contrast to Definition 1, which requires the observation of the system for all t. Swarming strategies to achieve the velocity alignment even when $\Gamma < 1$ was the focus of our previous work, which we summarize in the next.

Random Swarming Strategy

The two objectives of guardian's swarming motion are to (i) maintain high density around O where the intruder passes through; and (ii) maintain high speed. The first objective increases the probability of encountering a target, and the second objective increases the probability of satisfying condition (4).

The control law (strategy) for the guardian is described by the combination of artificial forces that generates the desired acceleration of the agent. The overall forcing on guardian i is given by $\mathbf{F}_i = (1-\lambda_i^P)\mathbf{F}_i^{(\text{swarm})} + \lambda_i^P\mathbf{F}_i^{(\text{pursuit})}$, where the switching parameter $\lambda_i^P \in \{0,1\}$ takes the value $\lambda_i^P = 0$ (resp. 1) in the swarming (resp. pursuit) phase.

The swarming algorithm $\mathbf{F}_i^{(\text{swarm})}$ consists of three forces; central, spacing, and random force, i.e., $\mathbf{F}_i^{(\text{swarm})} = \mathbf{F}_i^{(\text{cent})} + \mathbf{F}_i^{(\text{spac})} + \mathbf{F}_i^{(\text{rand})}$. The central force $\mathbf{F}_i^{(\text{cent})}$ resembling a damped spring attached to O maintains the cohesiveness of the swarm: $\mathbf{F}_i^{(\text{cent})} = -k_c\mathbf{r}_i - b_c\mathbf{v}_i$, where positive constants k_c and b_c are the control gains. The spacing force $\mathbf{F}_i^{(\text{spac})}$, which also resembles a damped spring, generates attraction, repulsion, and alignment behavior between the agents:

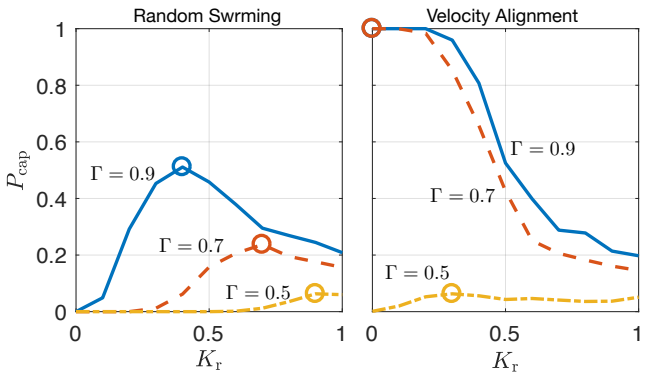

FIGURE 4. Probability of target capture as a function of K_r with $N_P = 10$. Left (resp. right) figure shows the results from random-swarming (resp. velocity-alignment) strategy [13].

$$\mathbf{F}_i^{(\text{space})} = -k_s \sum_{j \in S_i^{(\rho_a)}} \left(1 - \frac{x_0}{\|\mathbf{r}_{i/j}\|}\right)\mathbf{r}_{i/j} - b_s\mathbf{v}_{i/j}. \tag{5}$$

The positive parameter x_0 denotes the rest length of the spring, and the set $S_i^{(\rho_a)} = \{\, j \mid \|\mathbf{r}_{i/j}\| \le \rho_a\}$ consists of all the agents within the range ρ_a from agent i. The random force $\mathbf{F}_i^{(\text{rand})}$ has a constant magnitude, $\|\mathbf{F}_i^{(\text{rand})}\| = K_r u_{\max}$, in a random direction θ_i, i.e.,

$$\mathbf{F}_i^{(\text{rand})} = K_r u_{\max}[\cos\theta_i,\ \sin\theta_i]^T, \tag{6}$$

where $K_r \in [0,1)$. The random variable θ_i is generated by the random process $\dot{\theta} = Ww_i$, where w_i denotes the unit-intensity white noise, and $W > 0$ is a parameter describing the intensity. The intensity W determines how much (on average) the force $\mathbf{F}_i^{(\text{rand})}$ changes its direction in each time step. Finally, the magnitude of $\mathbf{F}_i^{(\text{swarm})}$ can exceed the limit $u_{\max}$, in which case the control is saturated while preserving its direction, i.e., $\mathbf{F}_i = u_{\max}\mathbf{F}_i^{(\text{swarm})}/\|\mathbf{F}_i^{(\text{swarm})}\|$.

In [13], we focused on the control gain K_r to study how the random forcing affects the probability of target capture. First we showed the existence of a trade off between the two objectives (swarm density and agent speed), i.e., with higher randomness, guardians acquire higher velocities, but they also spread out and the swarm density decreases. Then we found the optimal randomness of the swarm for a given set of system parameters Γ and α. The optimality was considered in terms of probability of target capture P_{cap} (see Fig. 4). Note that we did not use the payoff function J_G in our previous work.

Velocity-Alignment Strategy

The swarming algorithm introduced in the previous section focused on the individual motion of the guardians. The strategy we introduce in this section, which previously appeared in [13], has cooperation among the guardians to improve the target-capture capability. In particular, we consider a collaboration that is generated from a velocity-alignment behavior.

Consider a one-digit binary signal (i.e., communication of "Yes" or "No," instead of a serial communication like "010010...") that each vehicle can broadcast to other vehicles within the range ρ_a. The signal from vehicle i tells other vehicles whether it is in a *regular* swarming state or in an *alerted* state, which is the union of pursuit phase and velocity-alignment phase. In practice, the signal can be based on vision or acoustic sensing received by cameras or microphones, for example.

The algorithm for the velocity-alignment behavior is as follows. Let $S^{(\text{alert})}$ be the set of guardians that are either in pursuit phase or velocity-alignment phase. A guardian i in the swarming phase switches to velocity-alignment phase if it sees any guardian in the set $S^{(\text{alert})}$, i.e., if the following set is nonempty: $S_i^{(\text{align})} = \{\ j \mid \|\mathbf{r}_{i/j}\| \le \rho_a,\ j \in S^{(\text{alert})}\}$. The velocity-alignment phase will terminate in one of the following two ways: (i) guardian i switches back to the swarming phase when $S_i^{(\text{align})} = \emptyset$; or (ii) it switches to the pursuit phase when it encounters the target, i.e., $\|\mathbf{r}_{T/i}\| < \rho_p$.

Additional forcing for guardian i in the velocity-alignment phase is $\mathbf{F}_i^{(\text{align})} = b_{\text{a}} \sum_{j \in S_i^{(\text{align})}} \mathbf{v}_{j/i}$, which is equivalent to changing the damping constant b_{s} in the spacing term $\mathbf{F}^{(\text{spac})}$ to b_{a}, only for those guardians in the set $S_i^{(\text{align})}$. The constant b_{a} $(> b_{\text{s}})$ is sufficiently large that it dominates the other control terms during the the velocity-alignment phase.

If the guardian in pursuit phase aligns its velocity to the target, and if the velocity-alignment interaction propagates through the swarm, guardians that are far from the target can start moving in the direction that matches the velocity of the target. This mechanism allows the guardians to effectively increase their perceptual range ρ_p to the size of the swarm in order to gain favorable initial conditions for pursuit. Simulation results (see Fig. 4 right) showed that the velocity-alignment significantly improves P_{cap} if Γ is sufficiently large ($\Gamma > 0.5$ for this case), and also that it works best with crystallized swarm generated with $K_r = 0$.

The previous work only considered the case with a single intruder $N_T = 1$. In addition, the capture that occurs before and after the intrusion were not distinguished. This work addresses these two aspects with the payoff function J_I and the strategies for intruder swarm. The simulation section shows that the random-swarming strategy outperforms the velocity-alignment strategy when the intruders use a certain intrusion strategy.

ESTIMATED PAYOFF

This section analyzes the estimated outcome of the game by making some simplifying assumptions. Let N_T^e, N_T^a, and N_T^b be the number of intruders that escape after intrusion, captured after intrusion, and captured before intrusion, respectively. Intruders in each of the above groups score $J_i = V_{int}$, $J_i = V_{int} - \eta$, and $J_i = -\eta$, respectively. Note that we do not consider the case where an intruder misses O and is not captured (the left bottom entry in Fig. 3). Hence, we have the relation $N_T = N_T^e + N_T^a + N_T^b$.

Recalling that P_{cap} is the overall probability that the intruder is captured by any of the guardians, we have the relation

$$P_{cap} = \frac{N_T^a + N_T^b}{N_T} = \frac{N_T^a}{N_T} + \frac{N_T^b}{N_T} \triangleq P^a + P^b, \tag{7}$$

where we define $P^a \triangleq N_T^a/N_T$ and $P^b \triangleq N_T^b/N_T$ to be the probability of capture after intrusion and before intrusion, respectively. The probability of escaping (after intrusion) is

$$P_{esc} \triangleq 1 - P_{cap} = 1 - P^a - P^b. \tag{8}$$

We predict the expected outcome of the game assuming the following, though these assumptions are not enforced as part of the game:

- (B1) The expected time interval between two intruders, which we denote by $\bar{T}_{int}$, is constant throughout the game, i.e., $\bar{T}_{int} = \frac{1}{N_T - 1} \sum_{j=1}^{N_T - 1} T_j^{\text{int}}$.
- (B2) The number of intruders N_T is sufficiently large that the game may continue indefinitely.
- (B3) The guardians return to the swarm $T_{\text{cap}} < \infty$ seconds after the condition (4) is satisfied, so that the number of guardians protecting O is approximately constant throughout the game.
- (B4) The probabilities P^a and P^b are constant throughout the game.

Under assumptions (B1)–(B4), the expected number of intruders that arrived at O by time t is $n_T(t) = (t - t_0)/\bar{T}_{int}$, among which $n_{cap} = P_{cap} n_T$ are captured by the guardians. Therefore, the estimated payoff associated to the loss of intruder vehicles are

$$\hat{J}_\eta(t) = -\eta n_{cap} = -\eta P_{cap} \frac{t - t_0}{\bar{T}_{int}}. \tag{9}$$

(The accent $\hat{}$ is used to denote that it is an estimate.) On the other hand, the estimated score that the intruders gain by successfully intruding O is

$$\hat{J}_V(t) = \int_{t_0}^{t} (1 - P^b) \omega_{int}(s) V_{int}(s) ds, \tag{10}$$

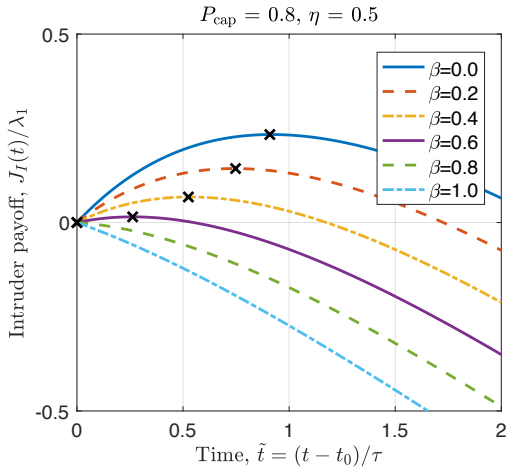

FIGURE 5. Expected time history of intruder payoff $\hat{J}_I(\tilde{t})$ normalized by λ_1. The crosses highlight the critical time $\tilde{t}^*$ and the value $\hat{J}_I^* = \hat{J}_I(\tilde{t}^*)$, for each value of β.

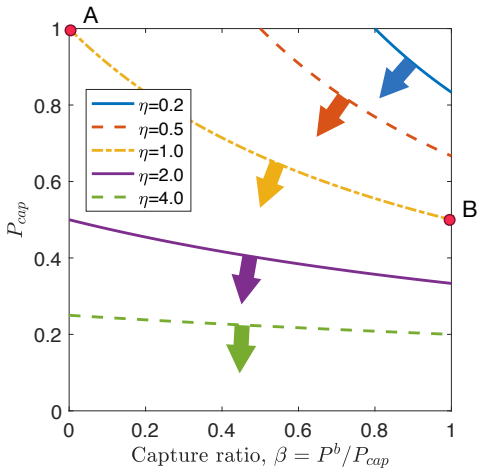

FIGURE 6. Necessary condition for intruder's win. (Sufficient condition for guardian's win)

where ω_{int} describes the expected number of intruders that arrive at O at time s, irrespective of whether they are captured or not. By assumption (B1), we have $\omega_{int} = 1/\bar{T}_{int}$. Multiplying $(1-P^b)$ excludes those intruders who are captured before arriving at O, i.e., only count the intruders who score $V_{int}(s)$. We compute $\hat{J}_V(t)$ explicitly as follows:

$$\hat{J}_V(t) = \left(1-P^b\right)\int_{t_0}^{t} \frac{e^{-(s-t_0)/\tau}}{\bar{T}_{int}} ds \tag{11}$$

$$= \left(1-P^b\right)\frac{\tau}{\bar{T}_{int}}\left(1-e^{-(t-t_0)/\tau}\right). \tag{12}$$

Let $\tilde{t} \triangleq (t-t_0)/\tau$ denote the nondimensionalized time, $\lambda_1 = \tau/\bar{T}_{int}$ denote the ratio between two time constants, and $\beta = P^b/P_{cap}$ denote the conditional probability that the capture occurs before intrusion. Now, the expected value of the payoff function J_I at time $\tilde{t}$ is described as

$$\hat{J}_I = \hat{J}_V + \hat{J}_\eta = \lambda_1\left((1-\beta P_{cap})\left(1-e^{-\tilde{t}}\right) - \eta P_{cap}\tilde{t}\right). \tag{13}$$

Figure 5 shows the expected time history of the payoff $\hat{J}_I(\tilde{t})$ for various values of β, when $P_{cap} = 0.8$ and $\eta = 0.5$. Since the penalty η is kept constant while the benefit V_{int} diminishes with time, the slope of J_I converges to $-\eta P_{cap}$ as time goes to infinity. (The validity of this prediction is shown in the simulation section.)

The figure also highlights the critical points where the intruders achieve their maximum payoff. Solving the equation $\frac{d}{d\tilde{t}}(\hat{J}_I) = 0$ for the critical time $\tilde{t}$, we obtain

$$\tilde{t}^* = -\log\left(\frac{\eta P_{cap}}{1-\beta P_{cap}}\right). \tag{14}$$

If the quantity inside the logarithm is greater than 1, there does not exist a critical point $\tilde{t}^* > 0$ that yields a positive value for $\hat{J}_I$. In this case, the intruders cannot win the game (see $\beta = 0.8$ and 1.0 in Fig. 5), and the optimal strategy for the intruders is to terminate the game immediately at $\tilde{t} = 0$ (i.e., no attempt of intrusion). Therefore, the necessary condition for the intruders to win the game is

$$\frac{\eta P_{cap}}{1-\beta P_{cap}} < 1, \tag{15}$$

which is depicted in Fig. 6. The top right corner of the figure corresponds to the case where all intruders are captured before intrusion, whereas the bottom left corresponds to the case where all intruders successfully escape after intrusion. To further understand Fig. 6, look at the case with $\eta = 1.0$. If the overall capture probability is $P_{cap} = 1.0$ (highlighted as point A), then all of the captures have to occur after the intrusion (i.e., $\beta = 0$) for the intruders to win. If the overall capture probability is $P_{cap} = 0.5$ (highlighted as point B), then the intruders can win even if all of the capture occurs before the intrusion (i.e., $\beta = 1.0$).

When a positive $\tilde{t}^*$ exists, the critical value of the payoff function is

$$\hat{J}_I^* = \hat{J}_I(\tilde{t}^*) = \lambda_1\lambda_2\left(\frac{\lambda_3}{\lambda_2} - \log\frac{\lambda_3}{\lambda_2} - 1\right), \tag{16}$$

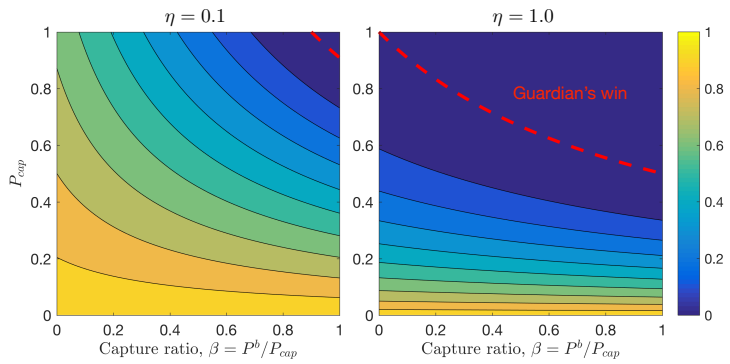

FIGURE 7. Estimated optimal intruder payoff $\hat{J}_I(\tilde{t}^*)$ normalized by λ_1.

where $\lambda_2 = \eta P_{cap}$ and $\lambda_3 = 1 - \beta P_{cap}$. Note that $\lambda_2 < \lambda_3$ from condition (6). Also note that λ_1 scales the entire $\hat{J}_I^*$. By differentiating $\hat{J}_I^*$ with respect to β, we obtain

$$\frac{d\hat{J}_I^*}{d\beta} = P_{cap}\lambda_1\lambda_2\left(\frac{1}{\lambda_3} - \frac{1}{\lambda_2}\right), \tag{17}$$

which is negative for $\lambda_2 < \lambda_3$. This result suggests that intruders should minimize β in order to maximize $\hat{J}_I^*$. Similarly, differentiation of $\hat{J}_I^*$ with respect to P_{cap} gives a negative value, indicating that P_{cap} should be minimized as well. The effect of β and P_{cap} is illustrated in Fig. 7, which shows $\hat{J}_I^*/\lambda_1$ for two values of η.

The intruder team needs to find strategies that minimize both the overall probability of capture and the conditional probability that captures occur before intrusion. The relative importance between the two probabilities depends on the vehicle cost η (i.e., the overall probability P_{cap} is more important than the conditional probability β when η is high). The intruders also have to estimate when they will achieve their optimal payoff and terminate the game before the payoff starts decreasing.

SIMULATION RESULTS

This section presents simulation results. First, we show how the probability of capture varies when intruders arrive at different frequencies. Based on the results, we propose a grouping strategy that improves the payoff of the intruders.

Effect of Intrusion Frequency

We first show the validity of the estimate $\hat{J}_I$ (see (13)) derived in the previous section. Figure 8 shows the time history of J_I from a single run. For this simulation, we used a fixed intrusion interval $\bar{T}_{\text{int}} = 3$ (i.e., $T_j^{\text{int}} = 3$ for all j). The approaching direction Ψ_{int} is a random variable uniformly distributed in $[0, 2\pi)$. The figure shows that J_V asymptotes to a constant value

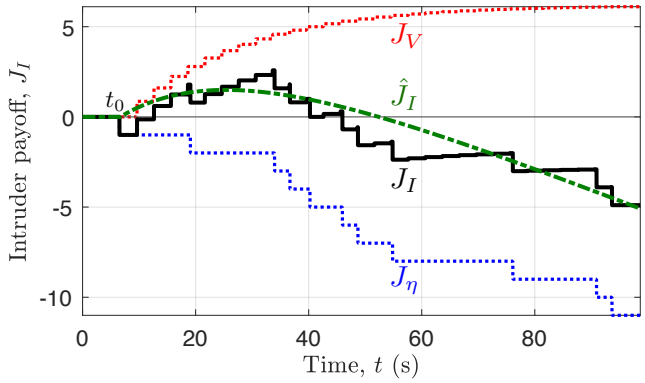

FIGURE 8. Time history of payoff J_I with constant intrusion interval $T_j^{\text{int}} = 3$ (s).

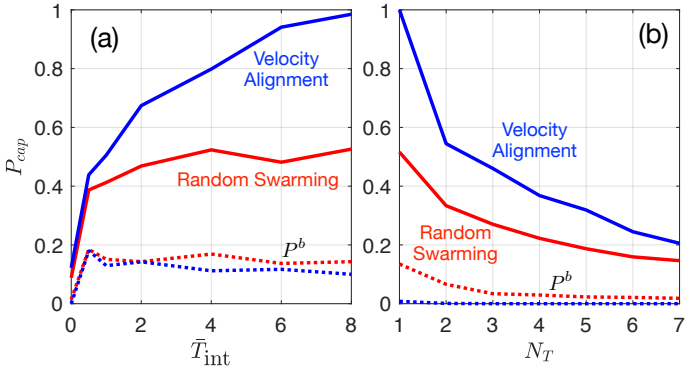

FIGURE 9. Probability of capture in multi-intruder scenario. The dashed lines show the probability of capture before intrusion P^b. (a) The effect of $\bar{T}_{\text{int}}$ tested with fixed intrusion interval. (b) The case with simultaneous intrusion of N_T intruders.

as $V_{\text{int}}(t)$ converges to 0. The estimate $\hat{J}_I$ is computed based on the probabilities P^b and P^a obtained from the simulation. The estimate successfully captures the trend of actual payoff J_I. However, since the estimate $\hat{J}_I$ is based on the knowledge of P^b and P^a, it may be calculated only with some uncertainties in practice.

Figure 9 shows the effect of intrusion frequency by simulating two special cases. The first case is when the intrusion interval is constant $T_j^{\text{int}} = \bar{T}_{\text{int}}, \forall\ j$. Here, we specify $\mathbf{t}_{\text{int}}$ directly assuming that the deployment constraints are satisfied (we will consider $\mathbf{s}_{\text{int}}$ in the next section). The system parameters are chosen as follows: $\Gamma = 0.9$, $\alpha = 0.5$, $N_P = 10$, and $N_T = 20$. Note that η and τ does not directly affect the probability of capture. Various intervals are tested against two guardian strategies (random swarming and velocity alignment). Figure 9-a shows that the probability increases as the intrusion interval $\bar{T}_{\text{int}}$ increases, and converges to the value obtained with $N_T = 1$ (see the opti-

mal P_{cap} with $\Gamma = 0.9$ in Fig. 4). This result demonstrates that the multi-intruder scenario becomes equivalent to the single-intruder scenario as $\bar{T}_{\text{int}}$ goes to infinity.

Figure 9-a also shows that random-swarming strategy reaches its maximum P_{cap} with shorter intrusion interval than the velocity-alignment strategy. This result indicates that the random-swarming strategy is more robust in the multi-intruder scenario. In addition, although the overall P_{cap} is higher with velocity-alignment strategy, the probability of capture before intrusion P^b (dashed lines) is higher with the random-swarming strategy. This trend is also true for the second case shown in Fig. 9-b.

As the second special case, we consider $\bar{T}_{\text{int}} = 0$, i.e., all of the N_T intruders arrive at the same time. The left most datapoint in Fig. 9-b corresponds to the single-intruder scenario, and therefore has the value that matches the right most datapoint in Fig. 9-a. As the number of target increases, the probability of capture decreases. The value at $N_T = 20$ (not shown in the plot) will match the left most datapoint in Fig. 9-a. A significant drop in P^b is noticable for the velocity-alignment strategy. This result is because sufficient number of velocity-matching interactions are necessary for successful target capture, which occurs only after the intruder passed through O, for the system parameters tested in the simulation. If N_P is much larger, for example, the capture will occur before intrusion.

The simulation results shown in this section indicates that the intruders should decrease the intrusion interval, and if possible, cluster the agents into groups and approach O simultaneously. This idea motivates the intruder's strategies considered next.

Intrusion Strategy

This section considers a simple grouping strategy for the intruder team in the presence of deployment constraint $\mathbf{s}_{\text{int}}$. The results in Fig. 9-b showed the benefit of clustering multiple intruders to approach the protected region at the same time. (Note we mean clustering only in terms of arrival time, but not in terms of the direction of intrusion.) An example is shown in Fig. 2—to form a group of three, intruders 3 and 4 wait until intruder 5 becomes ready for deployment.

More generally, we use the parameter N_{group} to define the grouping strategy as follows: intruders wait until they form a group of N_{group} agents, and approach O at the same time. The main interest is in the tradeoff between the penalty in V_{int} due to the waiting, and the benefit from grouping, which we study with numerical simulation.

We use the same systems parameters as before except for $N_T = 10$. For the deployment constraint, we use Poisson distribution to randomly generate $\mathbf{s}_{\text{int}}$. The parameter $\bar{T}_{\text{int}} = 2$ in this case specifies the mean value of the intervals $\{s^{\text{int}}_{j+1} - s^{\text{int}}_{j}\}$. Since the Poisson distribution generates infinitely large time interval with small probability, we saturate sufficiently long intervals with $3\bar{T}_{\text{int}}$ to keep the simulation time reasonable. (Another way to keep the simulation time within a reasonable value is to terminate the simulation at a pre-specified time even if some of the intruders have not been deployed.)

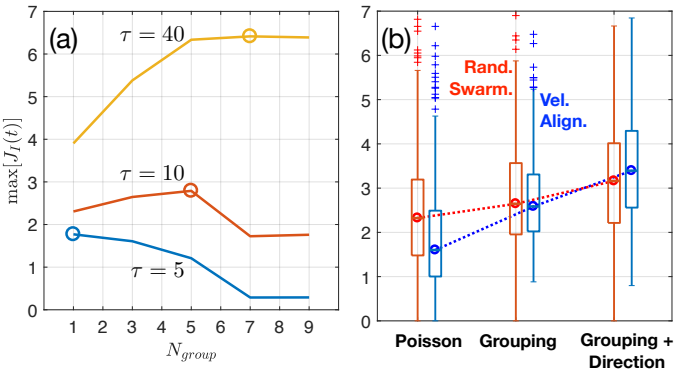

FIGURE 10. Intruder payoff with grouping strategy. (a) Different N_{group} are tested against random-swarming strategy. (b) Three intruder strategies are tested against two intruder strategies (random swarming and velocity alignment).

We compare different strategies using the payoff function J_I so that not only P_{cap} but also the ratio $\beta = P^b/P_{cap}$ is taken into consideration. The vehicle cost is set to be $\eta = 1$. Although the timing to terminate the game is an important consideration and the subject of ongoing work, here we simply use the maximum value in the time history $J_I(t)$. Figure 10-a shows how the parameter N_{group} affects the maximum payoff. (Note that $N_{group} = 1$ corresponds to the case where $\mathbf{s}_{\text{int}} = \mathbf{t}_{\text{int}}$, i.e., the intrusion interval is the original Poisson distribution.) The figure shows that the optimal N_{group} that maximizes the payoff is dependent on the time constant τ. For a smaller (resp. larger) τ, penalty due to waiting is higher (resp. lower), so the optimal strategy is to use a smaller (resp. larger) value for N_{group}.

Since the directions of intrusion Ψ_{int} were randomly chosen, some intruders in the same group approach O from similar directions, which is beneficial for guardians using the velocity-alignment strategy. To avoid this situation, we also add a strategy on Ψ_{int} by distributing ψ^{int}_j uniformly in the interval $[0, 2\pi)$ for the intruders in the same group. For example, if intruders $j = 1, 2$ and 3 are in the same group, their directions of intrusion satisfy $\psi^{\text{int}}_2 = \psi^{\text{int}}_1 + \frac{2\pi}{3}$, $\psi^{\text{int}}_3 = \psi^{\text{int}}_2 + \frac{2\pi}{3}$.

Figure 10-b shows the comparison between three intruder strategies: (i) original Poisson intrusion interval with random direction of intrusion, (ii) the grouping strategy ($N_{group} = 5$) with random direction of intrusion, and (iii) the grouping strategy ($N_{group} = 5$) with uniformly distributed direction of intrusion. The results show that the additional strategy on the intrusion direction improves the payoff. The results also show an interest-

ing crossover between the two guardian strategies, i.e., random swarming performs better than velocity-alignment strategy only when the intruder team uses both grouping and uniformly distributed direction of intrusion.

CONCLUSION

This paper formulates a game between a team of intruders and a team of guardians. The game is an extension of the previous study on guardian's strategies against single intruder. The proposed payoff function models various scenarios including missile attacks and reconnaissance missions by the proper choice of the parameter that weights the capture of intruders. Based on the probability of capturing an intruder, analytical expressions are derived to estimate the averaged outcome of the game. Simple grouping strategies for the intruders to simultaneously intrude the protected region are proposed and tested against two guardian strategies (introduced in previous work). Simulation results validated the theoretical results, and demonstrated the performance in muti-intruder scenario.

In ongoing and future work, we are studying a principled way of designing intruder strategies using the knowledge of the system parameters with uncertainties. We also plan to extend the guardian strategies to improve the performance in the multi-intruder case.

REFERENCES

[1] Robin, C., and Lacroix, S., 2016. "Multi-robot target detection and tracking : taxonomy and survey". *Auton. Robot.,* ***40***(4), pp. 729–760.

[2] Chung, T. H., and Hollinger, G. A., 2011. "Search and pursuit-evasion in mobile robotics a survey". *Auton. Robot.,* ***31***(4), pp. 299–316.

[3] Kim, T. H., and Sugie, T., 2007. "Cooperative control for target-capturing task based on a cyclic pursuit strategy". *Automatica,* ***43***(8), pp. 1426–1431.

[4] Bopardikar, S. D., Bullo, F., and Hespanha, J. P., 2009. "A cooperative homicidal chauffeur game". *Automatica,* ***45***(7), pp. 1771–1777.

[5] Scott III, W. L., 2017. "Optimal evasive strategies for groups of interacting agents with motion constraints". PhD thesis, Princeton University.

[6] Huang, H., Ding, J., Zhang, W., and Tomlin, C. J., 2011. "A differential game approach to planning in adversarial scenarios: A case study on capture-the-flag". *IEEE Int. Conf. Robot. Autom.*, pp. 1451–1456.

[7] Chasparis, G. C., and Shamma, J. S., 2005. "Linear-programming-based multi-vehicle path planning with adversaries". *IEEE Amer. Control Conf.*, pp. 1072–1077.

[8] Earl, M. G., and D'Andrea, R., 2007. "A decomposition approach to multi-vehicle cooperative control". *Rob. Auton. Syst.,* ***55***(4), pp. 276–291.

[9] Chen, M., Zhou, Z., and Tomlin, C. J., 2014. "Multiplayer reach-avoid games via low dimensional solutions and maximum matching". *IEEE Amer. Control Conf.*, pp. 1444–1449.

[10] Butail, S., Manoukis, N., and Diallo, M., 2013. "The dance of male *Anopheles gambiae* in wild mating swarms". *J. Med. Entomol.,* ***50***(3), pp. 552–559.

[11] Shishika, D., Manoukis, N. C., Butail, S., and Paley, D. A., 2014. "Male motion coordination in anopheline mating swarms". *Sci. Rep.,* ***4***, pp. 1–7.

[12] Shishika, D., and Paley, D. A., 2015. "Lyapunov stability analysis of a mosquito-inspired swarm model". *54th IEEE Conf. Decision Control*, pp. 482–488.

[13] Shishika, D., and Paley, D. A., 2017. "Mosquito-inspired quadrotor swarming and pursuit for cooperative defense against fast intruders". *Auton. Robot.* Submitted.

[14] Shishika, D., and Paley, D. A., 2017. "Mosquito-inspired swarming algorithm for decentralized pursuit". *IEEE Amer. Control Conf.*, pp. 923–929.

Appendix A: Poisson Distribution

The Poisson distribution for a discrete variable $x = 0, 1, 2, ...$ and real parameter λ is

$$P(x|\lambda) = \exp\left(-\lambda \frac{\lambda^x}{x!}\right). \tag{18}$$

This distribution is often used to describe the probability that an event can occur x times in an interval (unit time), and the average number of events in the unit time is specified by λ.

Let $\lambda = 1/\bar{T}$. Then the parameter $\bar{T}$ describes the average time interval between two events. Using $x = 0$, the quantity

$$P(\bar{T}) = P(x = 0|\lambda = 1/\bar{T}) = e^{-\frac{1}{\bar{T}}} \tag{19}$$

describes the probability that no event occurs over a unit time. Hence, considering the time interval $[t_0,\ t_0 + t]$, the quantity

$$P'(t;\bar{T}) = 1 - P(\bar{T})^t = 1 - e^{-\frac{t}{\bar{T}}} \tag{20}$$

describes the probability that at least one event occurs during this time interval. For $t = 0$, the probability is $P'(0;\bar{T}) = 0$. As t goes to infinity, the probability goes to unity, i.e., $P'(\infty;\bar{T}) = 1$.

In the simulation, we use this distribution to generate the deployment constraint $\mathbf{s}_{\text{int}}$, where $\bar{T}$ specifies the average time interval between two intruders.