Exposing Transmitters in Mobile Multi-Agent Games

Mai Ben-Adar Bessos
Bar-Ilan University, Israel
mai.bessos@gmail.com

Simon Birnbach
University of Oxford, United Kingdom
simon.birnbach@cs.ox.ac.uk

Amir Herzberg
Bar-Ilan University, Israel
amir.herzberg@gmail.com

Ivan Martinovic
University of Oxford, United Kingdom
ivan.martinovic@cs.ox.ac.uk

ABSTRACT

We study the trade-off between the benefits obtained by communication, vs. the risks due to exposure of the location of the transmitter. To study this problem, we introduce a game between two teams of mobile agents, the **P-bots** team and the **E-bots** team. The E-bots attempt to eavesdrop and collect information, while evading the P-bots; the P-bots attempt to prevent this by performing patrol and pursuit. The game models a typical use-case of micro-robots, i.e., their use for (industrial) espionage. We evaluate strategies for both teams, using analysis and simulations.

1. INTRODUCTION

Wireless communication is essential for effective cooperation among (bot or human) agents. Yet, communication may expose the location of the transmitter. Techniques to detect and locate the transmitter are well known, e.g., had a major role in the electronic warfare during WWII; this remains an active area in research and practice, esp. relevant in scenarios where intelligence is collected in hostile environments (see related works, Section 7). A good example is industrial espionage, which is a major concern for companies dealing with valuable confidential information. While these corporate secrets are usually well protected against outsiders, they are discussed inside the company on a day-to-day basis, and communicated in the clear on-premise. Eavesdropping on meetings, conversations between employees and internal communication may expose sensitive information.

Recent advances in (micro) robotics facilitate the development of dedicated *eavesdropping microbots (E-bots)*, exploiting their tiny size, mobility and communication capabilities to allow eavesdropping in highly guarded, sensitive areas. Examples of microbot research include several crawling agents [5, 27, 1] and few flying agents [25]. The Russian SpyRoach [1] is deliberately designed to resemble a living cockroach, explicitly motivated by the desire to avoid detection for intelligence-gathering operations, i.e., to be used as E-bot. Due to their small size, cf. to persons, E-bots may be especially attractive for use in scenarios where the goal is to collect information from hostile environments.

We envision a future where the use of E-bots is ubiquitous. These devices could penetrate a sensitive organization, e.g., through air-vents, or they could be planted by an insider or even a visitor that gets access to the building, exploiting their tiny size to avoid detection and to penetrate through physical barriers. Once inside, they can eavesdrop on internal computer and human communication and activity. Their mobility allows them to get close to meeting rooms, offices or communal areas where discussions usually take place. Their size allows them to stay undetected while eavesdropping, communicating and moving, e.g., within the grid (suspended) ceiling typical in many offices. When they have gathered enough data, they can either crawl to a safe place, waiting to be collected, or transmit the data to a sink that is nearby but off property. Transmitting the data to an off-site receiver (sink) may often be the only or best option to extract the data.

Due to the omnipresence and stealthiness of these devices, detecting and locating them is challenging; physically screening entire premises upon detecting suspect transmissions is hardly a solution. Hence, to minimize the threat posed by these intruders, companies will have to deploy advanced defense mechanisms. In particular, we believe that a main defense mechanism would be the use of *protecting robots (P-bots)*, used to prevent eavesdropping by E-bots, by patrolling sensitive areas and pursuing and disabling E-bots. The P-bots could take advantage of communication by the E-bots to detect and locate them.

We introduce the *Game of Eves*, which models the problem as a game between two teams of bots: the *E-bots* vs. the *P-bots*. The E-bots represent tiny, stealthy devices, used for *eavesdropping* on information from some target locations, and transporting it to some sink locations, while *evading* the P-bots. The P-bots have the opposite goal, i.e., they try to *protect* the information, e.g., by *patrol* and *pursuit* of E-bots. The P-bots do not necessarily have to be tiny or stealthy; we believe that a more realistic model for them is of rapidly-moving, larger, and visible agents.

Our analysis and simulations of the Game of Eves show that the use of emitter-locating mechanisms can have significant impact on the strategies of both teams and on the outcome of the game (amount of information leakage). The

CPS-SPC '16, October 28, 2016, Vienna, Austria.
© 2016 ACM. ISBN 978-1-4503-4568-2/16/10. . . $15.00
DOI: http://dx.doi.org/10.1145/2994487.2994495

resulting strategies are interesting and, while non-trivial, rather elegant, with multiple challenges, questions and insights related to intra-team communication and cooperation in adversarial scenarios.

Eavesdroppers should collect information from specific target locations, and transfer it to *sinks* (collection-points), by transmitting it or by 'crawling' and dropping it on a sink. Transmitting is significantly quicker, but exposes the location of the E-bot to the P-bots. One round of the Game of Eves, with one move for E-bots and one for P-bots, is illustrated in Fig. 1.

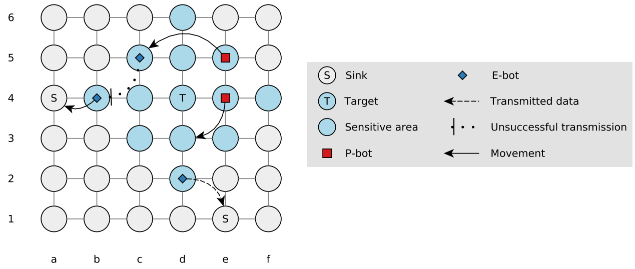

Figure 1: Example of one round of 'Game of Eves'. P-bot *e5* detected the broadcast by E-bot *c5* and captures it, interrupting its transmission; E-bot *d2* transmits to sink *e1*; E-bot *b4* has accumulated data and now moves to sink *a4*; P-bot *e4* is patrolling.

The fact that communication may expose the E-bots location (further discussed in Section 2), introduces interesting challenges and questions; often agents refrain from communicating, in order to hide their location, and sometimes, they take extra risks to communicate together, to reduce the effectiveness of pursuit.

Game of Eves is related to several well-studied problems (see Section 7 for a more thorough survey), yet we are not aware of any current research involving coordination between mobile agents which suggests methods for coping with the risk of exposure due to the transmission itself. A few studies focus on minimizing the use of communication operations in a general-purpose framework (*DEC-POMDP*) [26, 11], which requires unfeasible computation time for most environments, including Game of Eves. E-bots are reminiscent of pervasive, mobile devices trying to ensure *location privacy* [6], however, in most current works on location privacy, the location is exposed via the content and identifiers of messages sent. The Game of Eves shares similarities with other games that involve mobile agents in adversarial environments, in particular, *patrolling* and *pursuit-evasion.* Few works on these games discuss the aspect of locating adversaries using emitter-locating mechanisms [33], but moderating the use of communication as part of the strategy was not considered.

The Game of Eves further extends the patrolling problems, by considering a long-term game (the game does not end after an E-bot breaches a target or is captured). This motivates more complex strategies such as coordinated distractions. That is, E-bots may deliberately expose their location (e.g. as preparation before simultaneous, coordinated intrusions). We identify several 'sub-tasks', most of which are challenging by themselves, and provide corresponding basic solutions, analysis and experiments. These sub-tasks may be useful in many other similar problems involving microbots in adversarial environments.

Contributions.

- We show that emitter-locating mechanisms can have significant impact on operating strategies of both agent-teams.
- We develop and analyze different P-bot strategies, prove corresponding bounds on amount of information leakage, and show experimentally that these bounds are reasonably tight. In particular, among the tested P-bots strategies, we show that best results require combining pursue and patrol activities, and furthermore, the use of both perimeter patrol and area patrol.
- We evaluate two strategies for E-bots (1) crawl (no transmissions), (2) transmit (no crawling). We show that if E-bots may use both strategies, their performance increases significantly.
- We show that E-bots may cleverly adapt their strategy, based on the P-bots strategy. One counter-intuitive adaptation we found experimentally, is that against 'aggressive patrol' policies, the best E-bots strategy is often to simultaneously transmit using several E-bots.

2. SYSTEM & THREAT MODEL

E-bots are microbots deployed in an office environment where sensitive information is communicated. Their goal is to transfer the information to a *data sink*, by transmitting it and/or by 'crawling' to it.

The collected data may lose relevancy over time. For example, if an attacker needs to know what the company is doing at a certain moment, the data has to be transmitted in real time. These transmissions can be detected and help to locate the transmitting E-bot.

We assume that these E-bots cannot communicate freely without revealing their location. Hence, E-bots only coordinate outside the protected area or when they transmit data. Due to their small size E-bots are only capable of slow movement, but are hard to spot at a distance. They can sense P-bots, however, only in their immediate proximity.

P-bots engage in coordinated patrols and pursuit. In order to determine the position of transmitting E-bots, P-bots can rely on a pre-established *emitter-location* system, that detects and locates radio transmissions within the perimeter of their facility. This system is able to distinguish E-bots from legitimate senders, as we assume that the use of wireless transmitters is restricted to authenticated devices by company policy. We further assume that if the localization is successful, it is precise enough to get in visual range of a stationary transmission source (modeled as being 'in the same cell'). Once the P-bots get into the same cell as an E-bot, they can disable it.

The Game of Eves has parameters $(G, t, V_S, \eta, \psi, r_p, r_e, p_d, R)$, known to both teams:

- $G = (V, E)$ is an unweighted, undirected 4-connected grid map, representing the environment. Agents move along the edges, from node to node.
- η is the number of E-bots (eavesdropping microbots).
- ψ is the number of P-bots (patrol robots).

- $r_p \in \mathbb{N}$: distance that P-bots may go per round.
- $r_e \in \mathbb{N}$: eavesdropping distance of an E-bot; we refer to the area within r_e around the target as the *sensitive area.*
- $p_d \in [0, 1]$: probability of detecting a transmission of an E-bot, by P-bots.
- $t \in V$ is a node representing a target (meeting room, office, communal area)
- $V_S \subset V$: set of nodes representing sinks (collection point). We assume sink points are in distance $r_e + 1$ from the target (i.e. adjacent to the sensitive area).
- $R : \mathbb{N} \to (0, 1]$: a nonincreasing function - the reward given to E-bots for a data item which reaches the sink x rounds after it was eavesdropped. We used $R(x) = \delta^x$ where $0 < delta \leq 1$ is the *discount factor* (typically, with $\delta = 0.9$).

The game executes as a sequence of *rounds*, each consists of two *turns* (moves): first for the P-bots, then for the E-bots. In the P-bots turn, each P-bot can move by up to $r_p \geq 1$ hops. In contrast, E-bots can only move by one hop each turn, but may also transmit. The amount of transmitted data is limited to the amount of eavesdropped data in a single round. Whenever a turn results in a P-bot and an E-bot occupying the same node, except for sink nodes, then the E-bot is *captured*, i.e., disabled for the rest of the game; any message it sent during this round is lost. For example, P-bot *e5* in Fig. 1 captures E-bot *c5*. The relevancy of each collected data item depends on the number of rounds x since it got eavesdropped. When E-bots successfully transmit it to a sink (or flush it by physically reaching the sink point), E-bot reward increases by $R(x)$.

At the beginning of the game, the ψ P-bots can be placed anywhere; each of the η E-bots can be placed in every round at any sink node, until all are placed. The game ends after all of the E-bots are captured, and E-bots are given the accumulated reward of all previously collected data.

In the scope of this work we discuss scenarios with a single target, but the model may be extended to include any subset of target nodes in V in order to examine scenarios with different sensitive area types. Additionally, even though we assume that sink points are adjacent to the sensitive area, the model may be extended to describe other scenarios as well. That is, if the sinks are further away, other E-bots may stay close to the area and collect transmitted data instead of the sink, then forward it. Table 1 summarizes the capabilities of each agent type. For a of summary of all game parameters and base values, see Table 2 in Section 6.

Table 1: Capabilities of agents

E-bots	**P-bots**
Movement range: 1 hop	Movement range: r_p
Eavesdrop range: r_e	Transmission detection (in G)
Transmit data to sinks	Locate transmitting E-bots
Sense adjacent P-bots	Capture E-bots on same node

3. E-BOT STRATEGIES

In this section, we introduce two E-bot strategies: *stealthy E-bot* and *transmitting E-bot*. In the *stealthy E-bot* strategy, each E-bot avoids transmissions and instead flushes eavesdropped data by physically reaching the sink point. In the *transmitting E-bot* strategy, each E-bot transmits the data to speed up its delivery to the sink, at risk of detection. These strategies, while basic, allow for a simple but informative analysis. In particular, Figs. 4, 5 illustrate the improved E-bot performance when facing an optimized P-bots strategy (discussed in the following sections), if E-bots are given the option of choosing between the two strategies.

3.1 Stealthy, Crawling E-bot

When the E-bots are dropped at a place where they can either be safely collected or transmit without fear of being detected, they do not necessarily have to take the additional risk of transmitting. Instead, they can use a more stealthy approach and eavesdrop silently, before returning to this safe place to deliver the data.

This strategy involves a single agent which enters the sensitive area through a random point and remains within the area for some rounds to eavesdrop. The E-bot then escapes and transmits the accumulated data to the sink once it is safe. It repeats this process until it gets captured. Fig. 2 illustrates the strategy with an example.

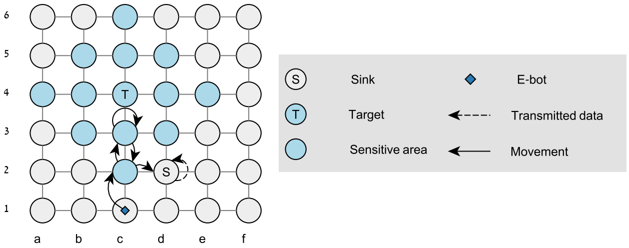

Figure 2: Stealthy E-bot *c1* intrudes the sensitive area, stays for a couple of rounds and escapes; Once it reaches the sink *d2*, where it is safe from the P-bots, it transmits the accumulated data.

3.2 Transmitting E-bot

In cases where the E-bots are not able to reach a safe place or when the data needs to be extracted in a short period of time to be useful to the attacker, E-bots cannot rely only on the more stealthy, crawling approach. Instead they have to risk exposure and transmit data directly to a sink.

The simplest way for implementing this is by having a single E-bot enter the sensitive area, then immediately transmit every eavesdropped data unit it eavesdropped, instead of leaving the sensitive area with the collected data as before. Between transmissions the E-bot will change its location randomly (but remains in the sensitive area), since the transmission may have disclosed the previous location.

Surprisingly, this method may be significantly improved (as we later show) by letting several E-bots enter at once. The E-bots will continue to collect data until they have as many unique data units as there are E-bots eavesdropping, then each of them will transmit a unique data unit in the same round as the others. Simultaneous transmission could be achieved by using non-overlapping channels. As with a single transmitting E-bot, all E-bots change their location after each transmission. If an E-bot gets caught, a new E-bot enters the sensitive area to replace it (if available). The entire process gets repeated until all E-bots are captured. We assume for this strategy that there are sinks in transmission range of the E-bots. An example round is displayed in Fig. 3. Note that when we refer to the transmitting E-bots strategy, any amount of simultaneous transmissions may be used (including only one).

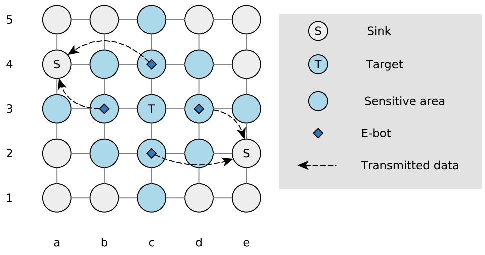

Figure 3: The E-bots *b3,c2,c4,d3* are within eavesdropping range, and transmit eavesdropped-data simultaneously to the sinks at *a4,e2*. In order to have enough unique data units, they have to wait for 4 rounds before transmitting.

4. P-BOT STRATEGIES

In this section, we discuss two P-bot strategies: *Area patrol* and *Patrol and pursuit*. The strategies rely on three algorithms: one pursuit algorithm, and two patrol algorithms (area patrol and circumference patrol). See Appendix A for a high-level description of the algorithms and their performance.

Notations. Let $dist_G(v, w)$ denote the length of the shortest path between $v, w \in V$; since G is an undirected 4-connected grid, then $dist_G$ is simply Manhattan distance. $Ring_G(p, K) = \{v \in G | dist_G(p, v) = k, k \in K\}$.

4.1 Area Patrol Strategy

Consider an E-bot that waits outside the circumference of the sensitive area. If it is able to anticipate that one of its adjacent points (on the circumference) will necessarily not be visited in the following round, it may enter and eavesdrop to at least one data unit, without risk. Similarly, if any point in the sensitive area is reachable and never gets visited, it will be a preferred destination for E-bots. In both cases, the amount of leaked data may be unbounded.

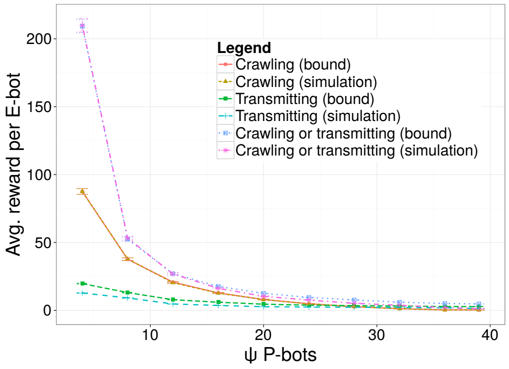

Figure 4: Compare E-bot strategies, for different number of P-bots. Even when crawling is preferable to transmitting, if both are possible, this affects P-bots strategy.

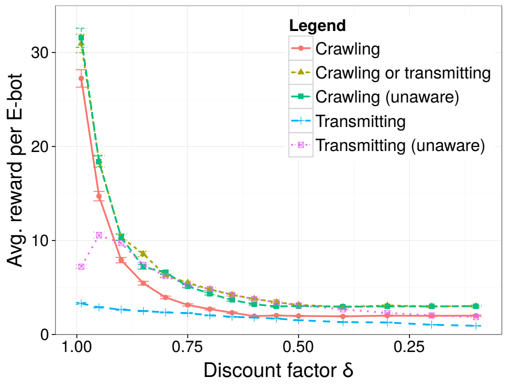

Figure 5: Compare E-bot strategies, for different δ values and where P-bots are either aware of the limitations on E-bots strategy, or oblivious to them (and assume no limitations). In all cases, reward for crawling decreases drastically with δ, since E-bots lose the incentive to accumulate many data units.

The Area patrol strategy avoids this situation by guarding the entire sensitive area around the target (using area patrol algorithms). Its goal is to ensure that every point in the sensitive area is being visited at *each round* with probability of at least $p_a > 0$. That is, every point must be reachable by at least one P-bot at each round.

4.2 Patrol and Pursuit Strategy

E-bots must cross the circumference to reach the sensitive area, where they can eavesdrop. Hence, P-bots are motivated to intensify the patrol in the circumference, in addition or instead of area patrol. However, if no pursuit after transmiting E-bots is done, E-bots could transmit eavesdropped data directly to a sink from within the area, thus limiting the advantage of focusing on the circumference.

The pursuit algorithm allocates P-bots that remain within the sensitive area, and pursue E-bots if their location was disclosed. For simplicity, we assume these P-bots do not take part in any of the patrol patterns. The strategy separates P-bots into the following three distinct groups:

- *Circumference Patrol:* P-bots dedicated to patrolling the circumference, i.e., ensuring that each point circumference $Ring_G(t, \{r_e\})$ is visited with probability at least p_c, at each round.
- *Area Patrol:* P-bots dedicated to patrolling the sensitive area, i.e., ensuring that each point in the sensitive area (circumference excluded), is visited with probability at least $p_a < p_c$, at each round.
- *Pursuit:* P-bots dedicated to the pursuit of E-bots that transmitted from within the sensitive area. The pursuit lasts for only a single turn of the P-bots. If multiple transmissions were detected simultaneously, the P-bots will target only one of them, chosen at random with uniform probability. The probability for capturing a transmitting E-bot is $0 \leq p_p \leq 1$.

Section 5 provides a method for bounding E-bots performance, given p_a, p_c, p_p. The patrol and pursuit strategy considers all possible allocations of the ψ P-bots to the three groups, and selects the one that minimizes E-bots' performance - assuming E-bots will attack using the best strategy available to them.

Note that the strict limitations on the group dedicated for pursuit were set in order to simplify the analysis below. After a transmission, the transmitting E-bot must be in either the same cell or one of the four adjacent cells. If pursuing P-bots are distributed such that exactly $i \leq 5$ of these five points will be covered by a P-bot, this results in a capture probability of $p_p = \frac{i}{5} \in [0, 1]$. This allows the P-bots to initiate a new pursuit in the following round. That is, in all three P-bots groups, the effect in each round is independent from the state of P-bots in the previous round.

5. BOUNDING EFFECTIVE LEAKAGE

In this section, we present upper bounds on the reward for any E-bot strategy, given that P-bots are limited to Area Patrol or Patrol and Pursuit strategies as described in Section 4. We assume that only the number of P-bots allocated for each of the roles *Area Patrol* , *Circumference Patrol* and *Pursuit* may vary (the capture probabilities p_a, p_c, p_p are derived respectively), and P-bots focus on protecting the target point t. Additionally, we strictly assume all data is flushed immediately to the sink if an E-bot escapes from the sensitive area. As mentioned in Section 2, η is the amount of E-bots, p_d is the transmission-detection probability and $R(x)$ is the reward given for an x rounds old data unit.

5.1 Area Patrol Strategy

Theorem 1. *If P-bots use the Area Patrol strategy, the expected reward gained by E-bots is bounded by* $\eta \cdot (\frac{1}{p_a} - 1) \cdot R(1)$.

Proof. At the end of the P-bots turn, only uncaptured E-bots that are within the sensitive area may accumulate data. Consider any specific E-bot and let X represent the total number of rounds during which it is inside the sensitive area (and hence able to collect data). Since at each such round the E-bot is in some cell in the sensitive area, and this cell is visited with probability of at least p_a, it follows that $E(X) \leq \frac{1}{p_a}$ (the expectancy of geometric distribution with probability p_a for success). In total, the E-bot is expected to accumulate data for at most $\frac{1}{p_a} - 1$ rounds before it is captured. For each data unit it receives reward, and the claim follows since this holds for each of the η E-bots. □

5.2 Patrol and Pursuit Strategy

We use $R_n = \sum_{i=1}^{n} R(i)$ to denote the reward given for the n latest consecutively-eavesdropped units. For an E-bot that uses the strategy described in Section 3.1 (repeatedly enters and escapes the sensitive area) and spends l rounds within the sensitive area each time it enters, we denote with $u(l)$ the expected reward it gains before being captured. $C(l)$ denotes the probability that it will be captured before exiting the sensitive area. Additionally, we use $l_{escape} = \arg\max_{l \in \mathbb{N}} u(l)$. Let x be the reward given for data that reached the sink exclusively by flush. $E_{escape}(x, p_a, p_c, R)$, abbreviated to $E_{escape}(x)$, is the lower bound on the expected number of captured E-bots before they received the reward i.e. for crawling E-bot strategy, at least $E_{escape}(1)$ E-bots are expected to be captured for every collected data unit.

E-bots that follow the stealthy strategy in Section 3.1 may have an incentive to stay longer in the sensitive area due to the increased risk of staying only in the circumference area. Since each E-bot gains reward each time it successfully exits the sensitive area, it holds that: $u(l) = (\frac{1}{C(l)} - 1) \cdot R_l$. This is similar to Theorem 5.1, where an E-bot repeatedly enters and exits the area until it is captured. The reward is given for the data accumulated in the span l rounds.

Lemma 1. *1. $C(x)$ may be bounded as follows:*

$$C(x) \geq \begin{cases} p_c & x = 1 \\ (1 - (1 - p_c)^2 (1 - p_a)^{x-2}) & o.w. \end{cases}$$

2. $u(l)$ has a single extremum point in $[3, \infty)$

3. $E_{escape}(x, p_a, p_c, R) \geq \frac{x}{R_{l_{escape}}} \cdot \frac{C(l_{escape})}{1 - C(l_{escape})}$

A proof for Lemma 1 is given in Appendix B. This lemma allows us to bound E_{escape} from below, if l_{escape} is given. For the constant reward function: $R(x) = c$, we may compute l_{escape} by strictly assuming the above bound on $C(x)$ is tight, then solving:

$$\frac{\Delta}{\Delta l_{ex}} u(l_{ex}) = \frac{\Delta}{\Delta l_{ex}} (\frac{1}{C(l_{ex})} - 1) = 0 \longleftrightarrow l_{ex} =$$

$$- c \frac{W(-\frac{p_c^2 - 2p_c + 1}{e \cdot (-1 + p_a)^2}) + 1}{ln(1 - p_a)}, l_{escape} = \arg\max_{l \in \{1, 3, \lceil l_{ex} \rceil, \lfloor l_{ex} \rfloor\}} u(l)$$

(where W is Lambert's W function). In this case E_{escape} has at most one extreme point in the range $[3, \infty)$, and this similarly holds for other nonincreasing reward functions. Specifically, given $\delta \in [0, 1]$ and a reward function: $R(x) = \delta^{x-1}$, we may compute l_{escape} by utilizing simple numerical methods. In the following section, reward functions of this form are used to illustrate the leakage bound in different scenarios.

Let x be the reward given for leaked data units that reached the sink exclusively from inside the sensitive area (by transmission). $E_{stay}(x, p_a, p_p, p_d, R)$, abbreviated to $E_{stay}(x)$, denotes the lower bound on the expected number of captured E-bots by the time they received the reward. If in all the transmissions the E-bots transmitted n units simultaneously, $E^n_{stay}(x, p_a, p_p, p_d, R)$, abbreviated to $E^n_{stay}(x)$, denotes the same.

Lemma 2. $E^n_{stay}(x) \geq (n - \frac{(-1+p_a)((1-p_a)^n - 1)}{p_a} + p_p p_d^n) \cdot \frac{x}{R_n}$.

A proof for Lemma 2 is given in Appendix B. Ideally for E-bots, no agent is captured and the optimal number of simultaneous transmissions is repeatedly used. Therefore, it holds that: $E_{stay}(x) \geq \min_{n \in \{1, 2 \ldots, \eta\}} (E^n_{stay}(x))$. Lemma 2 allows calculating a lower bound for E_{stay}, given the value: $\arg\max_{n \in [1, \ldots, \eta]} E^n_{stay}$. This value may be found using numerical methods (or exhaustive search, since η is finite and discrete). An approximation of this bound is illustrated in the following section.

Theorem 2. *Denote* $k = \arg\max_{n \in [1, \ldots, \eta]} E^n_{stay}$, $E_e = \frac{1}{R_{l_{escape}}} \cdot \frac{C(l_{escape})}{1 - C(l_{escape})}$ *and* $E_s = (k - \frac{(-1+p_a)((1-p_a)^k - 1)}{p_a} + p_p p_d^k) \cdot \frac{1}{R_k}$ *(E_e and E_s bound $E_{escape}(1, p_a, p_c, R)$ for $E^k_{stay}(1, p_a, p_p, p_d, R)$). The expected reward of the E-bots is bounded from above by:* $\frac{\eta}{\min(E_e, E_s)}$.

Proof. P-bots use a stateless strategy (i.e. the probability of capturing E-bots in each round does not depend on the state of P-bots in the previous round) and by definition, in every round only a single new unique data unit is generated. Consider any data unit that was eventually collected by E-bot e. Before the data unit is accumulated by e , it has a probability of at least $p = p_a$ or $p = p_c$ for being captured (or $p + (1 - p) \cdot \frac{p_p}{k}$ if it transmitted data at the previous round). This minimal risk is independent of the method used to eventually collect this data. The only exception to this, is for E-bots that escape the sensitive area while transmitting data and flushing simultaneously. If additional E-bots transmitted simultaneously (without escaping), the risk associated with pursuit may be decreased since pursuing P-bots might target an escaped E-bot. However, this method necessarily decreases E-bots performance (with respect to the bound). Consider an escaping E-bot in this scenario. If it was not in the sensitive area long enough to accumulate all the data that is transmitted in that round, then letting it enter sooner in order to accumulate additional data would be preferrable (accumulation of each unit is conditioned with p_a, which is the minimal possible risk). But otherwise, if the E-bot escapes with all the accumulated data, then no additional reward is given for the transmissions. That is, the expected risk for all E-bots in the sensitive area in every round is at least $E_{escape}(1)$ or at least $E_{stay}(1)$, and no additional data units may be accumulated in that round. □

5.3 Improving the Pursuit Algorithm

As specified in Section 4.2, the patrol and pursuit strategy requires that the pursuit algorithm will last for a single turn, the pursuing agents will not be used for patrolling tasks and it is assumed that at most one E-bot is targeted each round. Integration of more general pursuit algorithms may increase the deterrence against transmissions, but may also complicate the strategies of both sides, which makes their contribution harder to evaluate. For example, if the pursuing agents continue the search after transmitting E-bots for several rounds, it may increase the expected amount of captured E-bots. However, E-bots may be able to exploit this behavior by transmitting random noise, then transmitting a large amount of actual data from unguarded areas. Similarly, if the same P-bots are used for both pursuit and patrol, then E-bots may find a way to increase the probability of intrusion and escape. Adaptation of existing pursuit algorithms is left for future work. In particular, more work is needed for deciding under which conditions each P-bot should take part in a pursuit, for how many rounds the pursuit should last and in which cases the targeted E-bot(s) should be switched.

6. EVALUATION AND RESULTS

In this section, we compare the theoretical bounds of the previous section to results obtained by simulations. A confidence interval of 99% is used, and displayed with each value in the figures. The base values used are as specified in Table 2.

Table 2: Summary of game parameters

Symbol	Typical Value	Description
η	20	E-bots #
ψ	20	P-bots #
r_e	20	Sensitive area radius
r_p	20	P-bots speed
p_d	1.0	Transmission detection prob.
δ	0.9	Discount factor
$R(x)$	δ^x	E-bots reward for x-old data

Our results compare the different E-bot and P-bot strategies and limitations. Specifically, we compared results when E-bots use only crawling, only transmitting, or when E-bots use the better strategy (crawling or transmitting) against P-bots that use the patrol and pursuit strategy. In the results, we assume P-bots are aware of the limitations and the methods available for E-bots, and allocate agents to the different roles in order to minimize the E-bots' performance (then the E-bots evaluate and choose the best available response). Unless stated otherwise, both sides are free to use any of the presented strategies. In most cases, providing P-bots with knowledge about the E-bots' limitations only makes a small difference, as demonstrated in Fig. 5. The figure examines the case where P-bots may not assume that E-bots are limited to any specific strategy, even though they are. Since P-bots are optimized for the worst case, the E-bots gain only a slight advantage. Fig. 4 also compares different E-bots strategies. In this figure, we show both the analytical bounds (Theorem 2) and the simulation results; as can be seen here, they are very close. Hence, to reduce clutter, we display only simulation results in most of our graphs. The match between analysis and simulations increases our confidence in the simulation results and in the quality of the analytical bounds. The results in Figs. 4,5 show that the ability to transmit is important for E-bots, in spite of the increased risk of detection. On the other hand, detection *is* important for P-bots; as shown in Fig. 6, some P-bots are always used for pursuit.

Recall, that we modeled decaying reward for information, i.e., reward upon delivery of data at sink is given by $R(x) = \delta^x$, where x is the number of rounds since eavesdropping, and δ is the *discount factor*. In several experiments, we compared the impact of different δ values. For example, Fig. 5 shows that a lower discount factor δ increases the advantage of transmissions. Similarly, Fig. 6b shows that with lower δ, there is reduced need in area patrol, and most or all P-bots are allocated to circumference patrol and pursuit. Fig. 5 shows, however, that the combined approach yields a significantly higher reward than both uncombined strategies. This is due to the fact, that the P-bots have to allocate resources for both the pursuit and the circumference patrol if they can't be sure which strategy the E-bots will use. That is, even if E-bots will not utilize their ability to transmit data (or to crawl back), providing them with the choice improves their performance. Not surprisingly, the results of Fig. 5 show that the crawling-only strategy is not that bad when the freshness of the responses is not critical, i.e., as $\delta \to 1$.

Figs. 6a, 6b show the most effective allocations of P-bots for the different sub-tasks, by number of pursuers and by the discount factor. Notice that the pursuit sub-task is very important, motivating the use of transmitter-locating facilities by P-bots side. The next pair of figures, Figs. 7a and 7b, present the resulting impact on the E-bot reward as a function of the number of P-bots or discount factor, respectively. These figures are esp. informative, when considered together

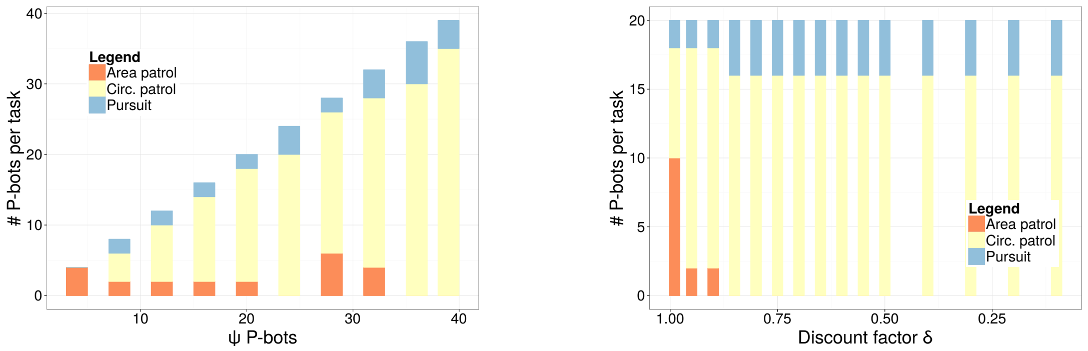

(a) The relative number of P-bots in each role is mostly indifferent to ψ.

(b) At some point, P-bots increase the effort on reducing the effectiveness of transmitting.

Figure 6: Optimal allocation of P-bots to tasks, as function of number of P-bots (6a) and discount factor (6b). E-bots use crawling or transmitting, and P-bots are allocated in a way that minimizes the E-bots reward, regardless of their strategy.

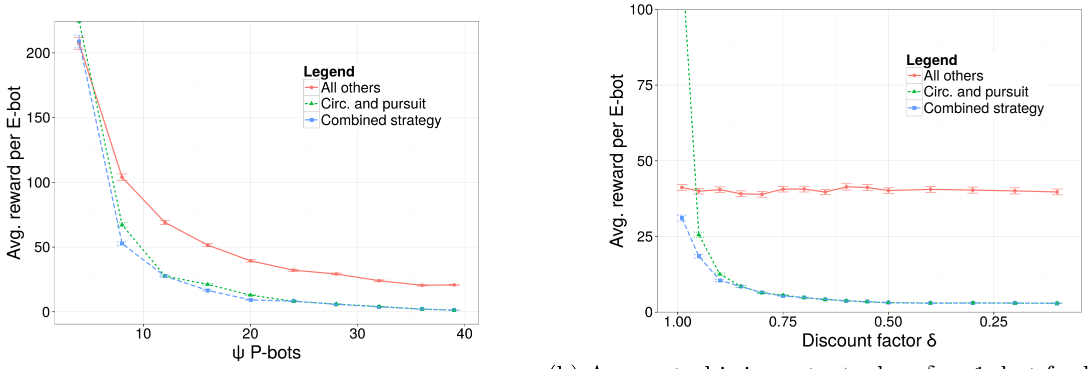

(a) If enough P-bots are available, circumference patrol and pursuit are particularly significant.

(b) Area patrol is important when $\delta \to 1$, but for higher values, *both* circumference patrol and pursuit are essential.

Figure 7: Reward as a function of the number of P-bots (7a) and discount factor (7b), where P-bots may not allocate agents for certain roles. We see that all of the roles are necessary for minimizing the reward given to E-bots.

with the allocation of P-bots to the different roles, presented and discussed above.

Again, we see that locating transmitters is an important ability of the P-bots. As can be seen from these figures, if transmitting E-bots face no risk of pursuit, they can use this to their advantage and increase their reward significantly. Hence, transmitting would be even more effective if P-bots would not apply transmitter-localization at all, or if the effectiveness of transmitter-localization would be reduced due to hardware or other limitations, or due to the use of evasive techniques such as spread spectrum communication by the E-bots. The use of such evasive techniques may also have disadvantages such as increased costs, energy consumption and delay; further research is needed.

We also see that circumference patrol is an essential part of P-bot strategies. If the circumference is not well guarded, then repeatedly entering and escaping the sensitive area to reduce long term risks becomes a viable option for the E-bots. E-bots are especially inclined to reduce the duration of their stay in the sensitive area, if the value of the eavesdropped data gets reduced by the discount factor otherwise. In this case, the circumference patrol is even more critical. On the other hand, the results show that as $\delta \to 1$, the area patrol also becomes important.

Fig. 8 illustrates the effectiveness of the E-bot technique of simultaneous transmissions, which may be counter-intuitive. As P-bots can only pursue a limited amount of exposed transmitters, additional simultaneous transmission are possible without the added risk of pursuit. But to enable simultaneous transmissions, several E-bots have to wait in the sensitive area facing the risk of getting caught by the area patrol. Additionally, they have to trade-off their waiting time and thus the amount of simultaneous transmissions against the diminishing reward gained for the data.

Fig. 9 shows that even an inaccurate transmitter-locating mechanism has a significant impact. For lower p_d values, E-bots will only use a single transmitting agent. At some point, using simultaneous transmissions is preferable, and this prevents further reduction in reward.

7. RELATED WORK

Detecting Transmissions vs. Low-Probability-of-Interception (LPI) Techniques: A transmitter that wants to remain undetected may use LPI techniques, such as spread spectrum technology, to decrease the likelihood of detection. These methods have been used in military applications, e.g. military GPS, but they are also widely used in commercial protocols, although mainly to be less susceptible to interference [29].

Spread spectrum technologies use codes to spread data over a large bandwidth. Different methods exist, with differences in their detectability. In particular, in Frequency Hopping Spread Spectrum (FHSS) communication, senders and recipients use different frequencies, selected using pseudorandom sequence; as the signal is still transmitted with full power, it is detectable [2]. Direct-Sequence Spread Spectrum (DSSS), on the other hand, spreads the signal over the whole bandwidth and can, in the ideal case, even manage to get the signal below the thermal noise threshold. Detectability depends on the available bandwidth and the length of the codes used to spread the signal. Commercial devices are more restricted in terms of available bandwidth than military systems and thus may be detectable even if they use DSSS. Muntwyler et al. [23] studied how commercial spread spectrum systems like the 802.15.4 standard can be used to avoid detection by substituting the public spreading codes with random ones. However, the obfuscation gain achieved by their adaptation is not sufficient for senders facing opponents with superior radio equipment, like in our case.

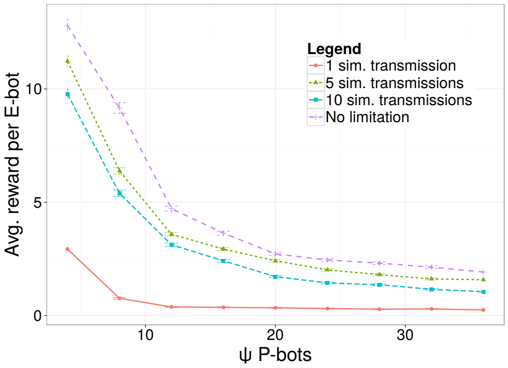

Figure 8: Compare different specific amounts of simultaneously-transmitted data (where E-bots may not use the crawling strategy). Despite the added risk of keeping several E-bots within the sensitive area at once, an increased amount may be effective for E-bots.

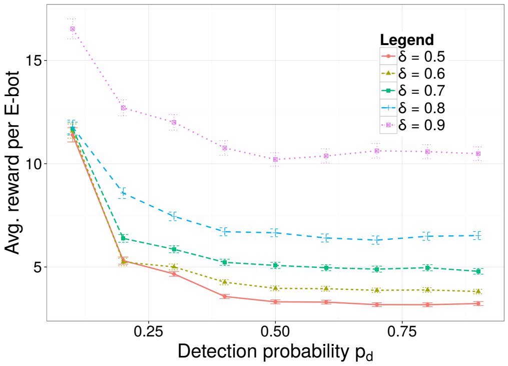

Figure 9: Impact of different detection probabilities. The rapid change in reward illustrates the point where E-bots drastically increase the amount of simultaneously-transmitted data, since using a single E-bot is no longer effective.

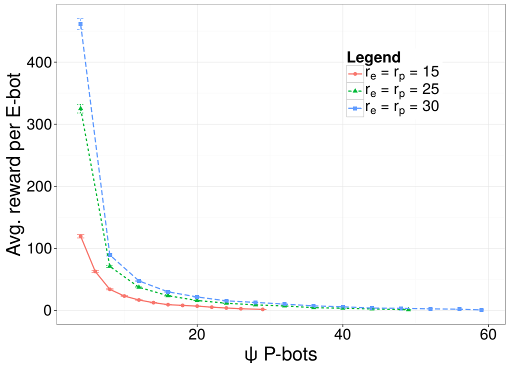

Figure 10: Impact of increasing r_e, r_p together. After ψ allows P-bots to allocate enough agents for circumference patrol and for pursuit, the difference in reward diminishes.

Directed antennas can be used to limit the area of exposure of the transmission. If fewer nodes can receive the signal, localization gets harder [28]. However, this can be counteracted by using a denser deployment of receivers, and may not be applicable to our scenario, since it complicates the transceiver design.

Emitter Location: A general introduction to basic emitter localization techniques can be found in [2]. Localization based on properties of wireless transmissions has received a lot of attention in recent years [16, 15]. For localization of non-cooperative transmitters, properties such as the Angle of Arrival (AoA) [20] of the transmission or the differing signal arrival times at different receivers (TDoA) [10] can be used. The previously mentioned methods can in general achieve a high accuracy when there is line-of-sight between sender and receivers.

Due to the multipath properties of indoor environments, most of the work on indoor positioning systems has focused, instead, on methods using the Received Signal Strength (RSS) at several receivers [21, 14]. Hybrid systems [12, 17, 35] are able to leverage the advantages of different approaches. For a thorough discussion of wireless positioning systems please refer to [36].

With regards to localizing LPI signals, Uysal et al. [32]

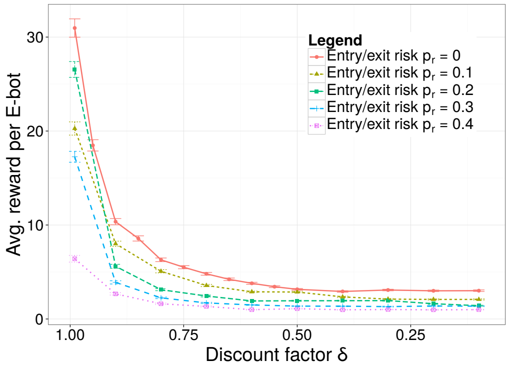

Figure 11: Impact of an added risk when going through the circumference (in both directions), e.g., due to physical barrier. The risk significantly affects the effectiveness of crawling, and P-bots utilize this.

recently proposed a method to localize a non-cooperative transmitter even when it uses DSSS techniques to avoid detection. More background on localizing LPI transmissions is provided by [3].

Eavesdropping: Eavesdropping games have been subject to study before. In [22] an overview is given over game theoretic approaches applied to eavesdropping and other network security problems. According to this survey the main concept used in related work is the secrecy capacity that can be achieved, i.e. the maximum rate at which two nodes can communicate without an eavesdropper being able to decode the information. In contrast, we study cooperating eavesdroppers that can be detected due to their transmissions and focus on the consequences of this assumption.

Industrial Espionage: Industrial espionage is a real and current threat to innovative companies. A comprehensive work on traditional methods and defences can be found in [9]. In recent years, cases of corporate espionage have increased significantly; e.g., see [19]. Nowadays, attackers go beyond installing bugs, and also use social engineering [34] and tailored malware [31] to get the desired information. In our work, however, we discuss the emerging threat of small, mobile eavesdropping devices, which is increasingly becoming a threat, due to advances in microbot research and development.

Pursuit-Evasion and Security Games: Game of Eves is related to security games [24] (adversarial patrolling) and to pursuit-evasion [13, 4]. Most studies on security games attempt to maximize the probability that a patrolling agent will visit one of several targets while it is attacked by an intruder. Only little work focuses on the ability of several intruders to cooperate. For example, [30] studies a two-intruder team that may perform a diversion, which motivates the defending player to consider several game steps ahead. Game of Eves involves intrusion into a physically protected area, but unfortunately none of the existing patrolling algorithms is compatible with our needs, and adaptation of such algorithms remained outside the scope of this work.

Pursuit-evasion games involve pursuing and evading mobile agents, where pursuers attempt to capture evaders by minimizing the distance (e.g. occupying the same node in a graph) or by surrounding them. Most studies on pursuit-evasion games are concerned with the analysis of distinct variations of the game such as specific graph classes or specific limitations imposed on the agents. A lot of work was done on limited visibility; [18] provides a patrolling heuristic that helps in pursuing an evader which is visible only if a line of sight exists between the pursuer and the evader and [8] discusses scenarios where all agents have a limited sensory range. It is often assumed that evaders' sole incentive is to avoid capture and will not disclose their location willingly (e.g. in order to transfer data). Therefore, integrating existing pursuit algorithms will not necessarily contribute to reducing the amount of data leaked by E-bots.

8. CONCLUSIONS AND FUTURE WORK

The emergence of microbots will have significant implications and applications for security, in particular, eavesdropping. We introduced the Game of Eves, a game between eavesdropping E-bots and patrolling P-bots, and where E-bots may communicate - but at risk of exposing their location. Our results indicate that utilizing transmitter-locating mechanisms can have significant impact on operating strategies of both sides. We consider this work as the first step in exploring adversarial scenarios involving mobile agents with an incentive to transmit accumulated data or to communicate with teammates.

There is a wide room for improving and innovating, not only of our algorithms, but also in study of related theoretical and applied problems. Introducing new patrol and pursuit strategies may improve the performance of P-bots. The game model can also be extended in order to study realistic scenarios more accurately, as the current model neglects energy considerations for E-bots and does not allow representation of physical obstacles nor location-dependent detection probability, all of which may have implications on the strategies. The presented analysis does not include methods for E-bots to covertly route data through a long distance (in case that sink points are not adjacent to the sensitive area). The need for routing eavesdropped data gives the P-bots an additional opportunity to capture E-bots, and requires E-bots to use decentralized algorithms which may include control communication. In a realistic system, deploying such communication would be challenging, esp. considering the potential exposure of location. In order to realize the results in this line of study, additional hardware experiments are needed for determining the feasibility and costs of locating hidden emitters with varying bandwidths and physical environments. Another aspect which should be addressed in follow-up work is the security of this communication mechanism, in particular, assuming that the patrolling team captures some of the devices, and furthermore, considering energy and other constraints. This may require adoption of low-energy cryptographic primitives resilient to exposure of some devices, e.g., key pre-distribution schemes such as [7], often proposed for similar goals, e.g., in sensor networks.

9. ACKNOWLEDGMENTS

We thank Hezi Moriel, Kris Pister, Noa Agmon and Sarit Kraus for their valuable comments. This work is supported by the Check Point Institute for Information Security, and the Israeli Ministry of Science and Technology.

10. REFERENCES

[1] Russian Scientists Create Cockroach Spy Robot. https://thestack.com/iot/2015/09/25/russian-scientists-create-cockroach-spy-robot. Accessed: 2015-10-07.

[2] D. L. Adamy. *EW 101: a first course in electronic warfare*. Artech House, 2000.

[3] D. L. Adamy. *EW 103: Tactical battlefield communications electronic warfare*. Artech House, 2008.

[4] F. Amigoni and N. Basilico. A game theoretical approach to finding optimal strategies for pursuit evasion in grid environments. In *ICRA*, pages 2155–2162. IEEE, 2012.

[5] B. G. D. I. R. J. W. Andrew T. Baisch, Onur Ozcan. High Speed Locomotion for a Quadrupedal Microrobot. *The International Journal of Robotics Research*, 2014.

[6] A. R. Beresford and F. Stajano. Location Privacy in

Pervasive Computing. *IEEE Pervasive computing*, (1):46–55, 2003.

[7] C. Blundo, A. D. Santis, A. Herzberg, S. Kutten, U. Vaccaro, and M. Yung. Perfectly Secure Key Distribution for Dynamic Conferences. *Inf. Comput.*, 146(1):1–23, 1998.

[8] S. D. Bopardikar, F. Bullo, and J. P. Hespanha. On Discrete-Time Pursuit-Evasion Games with Sensing Limitations. *Robotics, IEEE Transactions on*, 24(6):1429–1439, 2008.

[9] N. R. Bottom and R. R. Gallati. *Industrial espionage: Intelligence techniques and countermeasures.* Butterworth, 1984.

[10] J. J. Caffery and G. L. Stüber. Radio location in urban CDMA microcells. In *Personal, Indoor and Mobile Radio Communications, 1995. PIMRC'95. Wireless: Merging onto the Information Superhighway., Sixth IEEE International Symposium on*, volume 2, pages 858–862. IEEE, 1995.

[11] A. Carlin and S. Zilberstein. Value of Communication in Decentralized POMDPs. In *Proc. of the AAMAS Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (MSDM)*, pages 16–21, 2009.

[12] C. Cheng, W. Hu, and W. P. Tay. Localization of a moving non-cooperative RF target in NLOS environment using RSS and AOA measurements. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 3581–3585. IEEE, 2015.

[13] T. H. Chung, G. A. Hollinger, and V. Isler. Search and Pursuit-Evasion in Mobile Robotics. *Autonomous robots*, 31(4):299–316, 2011.

[14] D. Denkovski, M. Angjelicinoski, V. Atanasovski, and L. Gavrilovska. Practical assessment of RSS-based localization in indoor environments. In *Military Communications Conference, MILCOM 2012 IEEE*, pages 1–6. IEEE, 2012.

[15] Z. Farid, R. Nordin, and M. Ismail. Recent advances in wireless indoor localization techniques and system. *Journal of Computer Networks and Communications*, 2013, 2013.

[16] F. Gustafsson and F. Gunnarsson. Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements. *Signal Processing Magazine, IEEE*, 22(4):41–53, 2005.

[17] M. Hedley and Q. Zhai. Wireless sensor network using hybrid TDOA/RSS tracking of uncooperative targets. In *Wireless Personal Multimedia Communications (WPMC), 2014 International Symposium on*, pages 385–390. IEEE, 2014.

[18] A. Kehagias, G. Hollinger, and S. Singh. A Graph Search Algorithm for Indoor Pursuit/Evasion. *Mathematical and Computer Modelling*, 50(9):1305–1317, 2009.

[19] C.-M. Lee. Industrial espionage and police investigation. *International Journal of Security and Its Applications*, 7(1):155–162, 2013.

[20] M. Li and Y. Lu. Angle-of-arrival estimation for localization and communication in wireless networks. In *EUSIPCO*, pages 1–5. IEEE, 2008.

[21] H. Liu, H. Darabi, P. Banerjee, and J. Liu. Survey of wireless indoor positioning techniques and systems. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 37(6):1067–1080, 2007.

[22] M. H. Manshaei, Q. Zhu, T. Alpcan, T. Bacşar, and J.-P. Hubaux. Game theory meets network security and privacy. *ACM Computing Surveys (CSUR)*, 45(3):25, 2013.

[23] B. Muntwyler, V. Lenders, F. Legendre, and B. Plattner. Obfuscating IEEE 802.15. 4 communication using secret spreading codes. In *Wireless On-demand Network Systems and Services (WONS), 2012 9th Annual Conference on*, pages 1–8. IEEE, 2012.

[24] T. H. Nguyen, D. Kar, M. Brown, A. Sinha, M. Tambe, and A. X. Jiang. Towards a Science of Security Games. In *New Frontiers of Multi-Disciplinary Research in STEAM-H*, 2016.

[25] N. O. Pérez-Arancibia, P.-E. J. Duhamel, K. Y. Ma, and R. J. Wood. Model-Free Control of a Flapping-Wing Flying Microrobot. In *Advanced Robotics (ICAR), 2013 16th International Conference on*, pages 1–8. IEEE, 2013.

[26] M. Roth, R. Simmons, and M. Veloso. What to Communicate? Execution-Time Decision in Multi-Agent POMDPs. In *Distributed Autonomous Robotic Systems 7*, pages 177–186. Springer, 2006.

[27] M. Rubenstein, C. Ahler, and R. Nagpal. Kilobot: A Low Cost Scalable Robot System for Collective Behaviors. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3293–3298. IEEE, 2012.

[28] C. Santivanez and J. Redi. On the use of directional antennas for sensor networks. In *Military Communications Conference, MILCOM 2003 IEEE*, volume 1, pages 670–675. IEEE, 2003.

[29] M. K. Simon, J. K. Omura, S. R. A., and B. K. Levitt. *Spread Spectrum Communications Handbook.* McGraw-Hill, New York, NY, USA, 1994.

[30] E. Sless, N. Agmon, and S. Kraus. Multi-Robot Adversarial Patrolling: Facing Coordinated Attacks. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1093–1100. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[31] O. Thonnard, L. Bilge, G. O'Gorman, S. Kiernan, and M. Lee. Industrial espionage and targeted attacks: Understanding the characteristics of an escalating threat. In *Research in attacks, intrusions, and defenses*, pages 64–85. Springer, 2012.

[32] C. Uysal and T. Filik. A joint detection and localization method for non-cooperative DS-SS signals. In *Military Communications Conference, MILCOM 2015 IEEE*, pages 523–528. IEEE, 2015.

[33] M. Wei, G. Chen, J. B. Cruz, L. Haynes, K. Pham, and E. Blasch. Multi-pursuer multi-evader pursuit-evasion games with jamming confrontation. *Journal of Aerospace Computing, Information, and Communication*, 4(3):693–706, 2007.

[34] I. S. Winkler. Case study of industrial espionage through social engineering. In *Proceedings of the 19 th*

Information Systems Security Conference, pages 1–7. Citeseer, 1996.

[35] W. Xu, F. Quitin, M. Leng, W. P. Tay, and S. G. Razul. Distributed localization of a non-cooperative RF target in NLOS environments. In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8. IEEE, 2014.

[36] R. Zekavat and R. M. Buehrer. *Handbook of position location: Theory, practice and advances*, volume 27. John Wiley & Sons, 2011.

APPENDIX

A. PATROLLING ALGORITHMS

The implementation of *Area Patrol* , *Circumference Patrol* and *Pursuit* algorithms (presented below) implicitly rely on the following proposition.

Proposition 1. *Given distance $r > 1$, a 4-connected grid graph $G = (V, E)$ and any two points $v_1, v_2 \in Ring_G(v_c \in V, \{i|0 \le i \le r\})$, it holds that* **1)** *$dist_G(v_1, v_2) \le 2r$ and* **2)** *if the numerated points do not intersect the edges of the grid graph, $|Ring_G(v_c, r)| = 4r$ and $|Ring_G(v_c, \{i|0 \le i \le r\})| = 2r(r + 1) + 1$.*

Proof. **1)** By definition, $dist_G(v_1, v_c) \le r, dist_G(v_2, v_c) \le r$. By concatenating the paths that correspond to the distances we create a path of length $\le 2r$. **2)** By induction: For $r = 2$, trivial. We assume for $r > 2$. Let $v_c = V[x_0, y_0]$, and $v_t = V[x_0, y_0 + r + 1] \in Ring_G(v_c, r + 1), v_b = V[x_0, y_0 - r - 1] \in Ring_G(v_c, r + 1)$. Except for v_t, v_b, for each $V[x, y] \in Ring_G(v_c, r + 1)$: if $x \le x_0$ then $V[x + 1, y] \in Ring_G(v_c, r + 1)$, and if $x \ge x_0$ then $V[x - 1, y] \in Ring_G(v, r + 1)$ since the vertex is closer by 1 edge to v. Only $V[x_0, y_0 + r], V[x_0, y_0 - r]$ are matched in both conditions, and we get $|Ring_G(v, r)| + 2 + 2$ distinct vertices in $Ring_G(v, r + 1)$. By summation, $\sum_{0 \le i \le r} 4i = 4\frac{r(r+1)}{2}$. Including the graph center gives $2r(r + 1) + 1$. □

A.1 Area Patrol

The amount of available P-bots and their velocity affects the possible patrolling patterns. For example, for $r_p \ge 2r_e$, a single P-bot may move between any two points in the area, each round. Otherwise, the area must be divided into smaller distinct areas, where each area is reachable by different P-bots. Given distance $d \in \mathbb{N}$ and a point $t \in G$, the following algorithm guarantees that each of the points in the area $A = Ring_G(t, \{i|0 \le i \le d\})$ has a certain minimal probability p_a of being visited each round.

- Each two P-bots p_0, p_1 are designated an area $A_{p0,p1} = Ring_G(v, \{i|0 \le i \le r_p\})$ for some center point v. The center points are spread evenly within area A.
- Each round, for each two P-bots p_0, p_1 of the same area whose center is c, one of the following occurs:
  - If no P-bot is in c, then (with probability p_a) one of them moves to c, and the other randomly moves to any other unoccupied point in the area (that is reachable for it). Note that a P-bot that moves to c may reach any point in the area, in the following round.
  - One of the P-bots $p_i \in \{p_0, p_1\}$, which is currently not in the center, randomly moves to another unoccupied point in the area $v_i \in A_{p0,p1}$. The P-bot p_{1-i} moves to a point v_j for which $dist_G(v_i, c) = dist_G(v_j, c)$ and $dist_G(v_i, v_j)$ is maximized (this ensures each point in the area will be reachable by at least one of the two P-bots.

It is possible to designate a larger group of P-bots to each area instead of two (the group's size must be even, to ensure uniform visitation probability). This algorithm allows utilization of about a half of each P-bot's reachable area.
Performance: For an area $Ring_G(t, \{i|0 \le i \le r_p\})$, denoted by A_{t,r_p}, a P-bot in point $(x_f, y_f) \in A_{t,r_p}$, if $x \le y$ ($x \ge y$), every point $(x_d, y_d) \in A_{t,r_p}, x_d \le y_d$ (respectively $x_d \ge y_d$) is reachable. Therefore, the 2 P-bots used by the area patrol algorithm are enough to ensure that an area with radius r_p is reachable by at least one of them, i.e. two P-bots are designated to an area of size $\left|A_{t,r_p}\right| = 2r_p(r_p + 1) + 1$.

In some cases, an alternative algorithm which assigns the area $A_{t,\frac{r_p}{2}}$ to a single P-bot is preferable (although it was not used in our analysis). Since for every two points $p_1, p_2 \in A_{t,\frac{r_p}{2}}$ it holds that $dist_G(p_1, p_2) \le r_p$, every point in the area will be reachable. Even though more P-bots are required in this method for covering the same area, the method may be preferable if the area is small.

A.2 Circumference Patrol:

A simple circumference patrol algorithm may be implemented as follows. Given distance $d \in \mathbb{N}$ and a point $t \in G$, this algorithm spreads the P-bots in $Ring_G(t, d)$. It is assumed that $2 \mid r_p$, and only the minimal amount of P-bots, for which $p_c > 0$, is used.

- Position P-bots in $Ring_G(t, d)$, with distance of at most $2r_p$ from each other.
- At each round, $-\frac{r_p}{2} \le x \ne 0 \le \frac{r_p}{2}$ is chosen randomly. Each P-bot then moves clockwise (counter-clockwise if $x < 0$) from point p to $q \in Ring_G(t, d), dist_G(p, q) = |2x|$.

In order to increase p_c by including additional P-bots, instead of choosing x each round, a random repetitive pattern is chosen, that tells for each $r_p + 1$ consecutive points in $Ring_G(t, d)$, which ones should be occupied. Since the pattern is repetitive, every randomized pattern is necessarily reachable by the P-bots through either clockwise or counter-clockwise movement.
Performance: $|Ring_G(t, d)| = 4d$, and each P-bot may reach $\frac{r_p}{2}$ points to every direction (and in particular, two directions are in $Ring_G(t, d)$), at least $\lceil\frac{4d}{r_p+1}\rceil$ P-bots are required to cover the area, where each P-bot may reach r_p points other than its current one.

A.3 Single E-bot Pursuit:

Generally, E-bots are hidden from sight of the P-bots and patrolling in the area is necessary. But when an E-bot reveals itself by transmitting data, the P-bots know that it has to be located in one of the five nodes reachable from the transmission source.

- P-bots dedicated to the pursuit are distributed to cover the entire sensitive area. As in the Circular patrol before, every two P-bots p_0, p_1 are designated to an area $A_{p0,p1} = Ring_G(c, \{i|0 \le i \le r_p\})$ for some center point c. But this time, both P-bots wait in the center c until a transmission originates from a point $v_t \in A_{p0,p1}$.

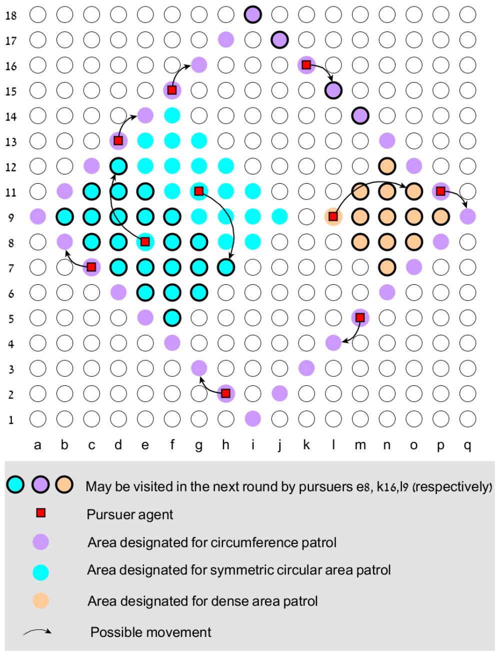
18
17
16
15
14
13
12
11
9
8
7
6
5
4
3
2
1
a b c d e f g h i j k l m n o p q
May be visited in the next round by pursuers e8, k16, l9 (respectively)
Pursuer agent
Area designated for circumference patrol
Area designated for symmetric circular area patrol
Area designated for dense area patrol
Possible movement

Figure 12: Comparison of patrolling algorithms

- One of the P-bots $p_i \in \{p_0, p_1\}$ randomly moves to a point $v_i \in Ring_G(v_t, \{i|0 \leq j \leq 1\})$, which is the area where the transmitting E-bot has to be located. The other P-bot stays in the center to be ready for transmissions in the next round. This way, the P-bots can target at least one transmission originating from any location within the sensitive area each round.

Depending on the number of available P-bots, they can achieve a higher capture probability of $p_p = \frac{k}{5}$ by assigning $2k, k \in [2, 5]$ P-bots to an area instead of just two.

Performance: In a single round, a P-bot may reach any point in distance r_p i.e. $2r_p(r_p + 1)$ different points other than its current one. Since another round is needed before the P-bot returns to the center of that area, at least two P-bots are needed for any area $A_{x,r_p}, x \in G$, and therefore its average designated points per P-bot is similar to that of the area patrol.

B. EFFECTIVE LEAKAGE: PROOFS

In this section, we provide proofs for Lemmas 1 and 2.

Proof of Lemma 1. 1. Consider an E-bot that uses the crawling strategy exclusively. Let l be the length of a particular visit in the sensitive area, which is also the number of data items collected - if the E-bot is not captured. The accumulated data units are necessarily unique, since that E-bot collects data only when no other E-bot is active. Hence, the $\frac{1}{C(l)}$ is the expected number of rounds the E-bot repeats the process until it is captured (geometric distribution), and the expected reward is: $u(l) \equiv \frac{R_l}{C(l)} - R_l$.

- For $l = 1$: the E-bot necessarily visited and immediately escaped a point in $Ring_G(t, r_e)$. Upon escaping, flushing the data does not increase capture probability, and therefore $C(1) = p_c$. Note that if an E-bot remains in $Ring_G(t, r_e)$ for $l = 2$, it risks losing the data it accumulated in the first round, and therefore such a strategy provides no benefit.
- For $l > 2$: the E-bot has the opportunity to occupy points in $Ring_G(t, \{i|0 < i < r_e\})$ (excluding the first and last rounds), thus reducing the capture probability for some of the rounds. Therefore: $C(l) = 1 - (1 - p_c)^2(1 - p_a)^{l-2}$.

2. For any reward function R, the first extremum $x \geq 3$ is the first point for which it holds that: $\frac{R_{x+1}}{(1-(1-\gamma)^2(1-\alpha)^{x-1})} - R_{x+1} > \frac{R_x}{(1-(1-\gamma)^2(1-\alpha)^{x-2})} - R_x \xrightarrow{0<\alpha<\gamma<1} R(x+1) < R_x \frac{(-1+\alpha)\alpha}{((1-\alpha)^x(\gamma^2-2\gamma+1)-\alpha^2+2\alpha-1)}$ (or the opposite, where $R(x+1) >$ from the right-hand side term for the first time). R_x is monotonically increasing since $R(x) > 0$. Additionally, it is multiplied by a monotonic term, since: $\frac{\Delta}{\Delta x}\frac{(-1+\alpha)\alpha}{((1-\alpha)^x\gamma^2-2(1-\alpha)^x\gamma-\alpha^2+(1-\alpha)^x+2\alpha-1)} = 0 \longleftrightarrow (-1+\alpha)\alpha((1-\alpha)^x ln(1-\alpha)\gamma^2 - 2(1-\alpha)^x ln(1-\alpha)\gamma + (1-\alpha)^x ln(1-\alpha)) = 0$ is never satisfied. If the right-hand side is < 0, it will be $< 0 < R(x)$ for any x. If the right-hand side is > 0, then $R(x)$ is nonincreasing and the left-hand is monotonically increasing, and therefore may meet only once.
3. An E-bot that transmits from within the sensitive area does not increase the amount of unique accumulated data, and does not contribute to the amount of data flushed from outside the sensitive area. Additionally, by design of the P-bots in this strategy transmissions may not decrease the probability of the E-bot for being captured. The expected number of transmitted data from outside the sensitive area until an E-bot gets captured $u(l) = R_l(\frac{1}{C(l)} - 1)$ is maximized for $l = l_{escape}$. That is, $E_{escape}(\frac{R_{l_{escape}}}{C(l_{escape})} - R_{l_{escape}}) \geq 1$ holds, and due to the linearity of expected value $E_{escape}(l) \geq l\frac{1}{\frac{R_{l_{escape}}}{C(l_{escape})} - R_{l_{escape}}} = l\frac{C(l_{escape})}{R_{l_{escape}}(1-C(l_{escape})}$ follows.

□

Proof of Lemma 2. Consider an E-bot that exclusively uses the transmitting strategy. Since only one unique data unit is generated in each round, the E-bot that transmitted the oldest data unit had stayed for at least n rounds, at least one other E-bot had stayed for $n - 1$, another for $n - 2$ and so forth. Accordingly, the independent risk each E-bot takes is at least $1 - (1 - p_a)^n$, $1 - (1 - p_a)^{n-1}, \ldots,$ $1-(1-p_a)$, which is summed up to $n-(1-p_a)\frac{(1-p_a)^n-1}{(1-p_a)-1} = n - \frac{(-1+p_a)((1-p_a)^n-1)}{p_a}$. After the transmission of the n units, the pursuit algorithm was invoked and targeted one of the transmitting E-bots that was not yet captured. That is, after any transmission an additional risk of p_p follows for some agent. Therefore, for a reward of R_n, $n - \frac{(-1+p_a)((1-p_a)^n-1)}{p_a}$ E-bots are expected to be captured before the transmissions begin, and additional p_p immediately in the next round. Similarly to the previous lemma, due to the linearity of expected value, $E^n_{stay}(n) = \frac{(n - \frac{(-1+p_a)((1-p_a)^n-1)}{p_a} + p_p)}{R_n}$ and $E^n_{stay}(l) = (n - \frac{(-1+p_a)((1-p_a)^n-1)}{p_a} + p_p)(\frac{l}{nR_n})$

(note that we disregard the option of leaving the sensitive area while transmitting, since this is considered flushing the data). □