

A Plug-n-Play Game Theoretic Framework For Defending Against Radio Window Attacks

Pruthuvi Maheshakya
Wijewardena
University of Utah
pmaheshakya4@gmail.com

Aditya Bhaskara
University of Utah
bhaskaraaditya@gmail.com

Sneha Kumar Kasera
University of Utah
kasera@cs.utah.edu

Syed Ayaz Mahmud
University of Utah
ayaz.mahmud@utah.edu

Neal Patwari
Washington University in St. Louis
npatwari@wustl.edu

ABSTRACT

The large scale deployment of multi-antenna wireless networks in homes and office buildings introduces new privacy concerns for people residing in these spaces. By measuring the signal strength using receivers placed outside the premises, an attacker can track the movement of people inside. One way to defend against such an attack is to have the signal strengths of the transmitters vary (sometimes reducing to zero) according to some randomized schedule. We show that the question of finding the schedule that minimizes the worst-case “privacy loss” can be formulated as a constant-sum Stackelberg game between an attacker, whose goal is to place receivers in order to learn the movement of users, and a defender who tries to prevent the attacker while maintaining the connectivity and QoS requirements of the network. We introduce a flexible framework that enables us to capture the constraints of the attacker and the defender. The framework allows us to capture features of modern wireless systems such as directional antennas and also allows us to plug in different path-loss models with minimal changes to the setup. We then formulate the problem of finding the optimal defender strategy as a linear program and show that it can be solved efficiently. We also perform numerical evaluations on how the payoffs are affected as the requirements of the defender and the resources the attacker can afford to exhaust change.

CCS CONCEPTS

- Security and privacy → Mobile and wireless security.

KEYWORDS

game theory, wireless security, privacy, radio window attacks

ACM Reference Format:

Pruthuvi Maheshakya Wijewardena, Aditya Bhaskara, Sneha Kumar Kasera, Syed Ayaz Mahmud, and Neal Patwari. 2020. A Plug-n-Play Game Theoretic Framework For Defending Against Radio Window Attacks. In *13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '20, July 8–10, 2020, Linz (Virtual Event), Austria*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WiSec '20, July 8–10, 2020, Linz (Virtual Event), Austria

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8006-5/20/07...\$15.00

<https://doi.org/10.1145/3395351.3399368>

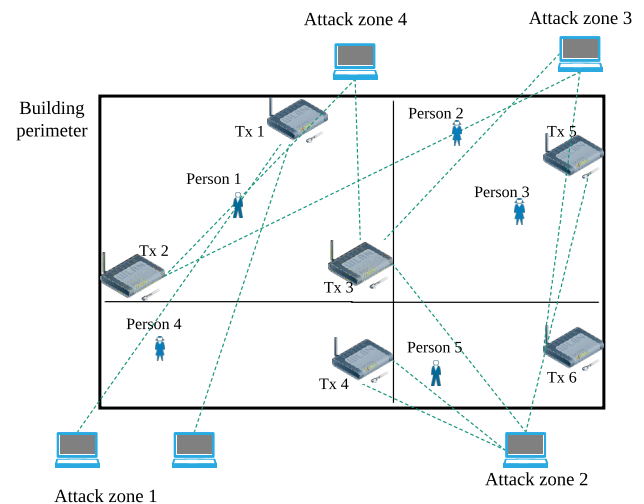


Figure 1: Illustration of a radio window attack, where an attacker deploys receivers outside a premise to detect movement of users inside.

'20), July 8–10, 2020, Linz (Virtual Event), Austria. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3395351.3399368>

1 INTRODUCTION

New wireless technologies that use multiple antennas and high-frequency directional beams have the potential to boost bandwidth and reliability for consumers at home, in office buildings, and in other closed premises. However, these technologies come with new vulnerabilities that can lead to a breach of privacy for the users of the network. An illustrative example is the class of so-called radio-window attacks (Figure 1). In such attacks, an adversary can deploy receivers and measure the *received signal strength (RSS)* from transmitters in a wireless network from *outside* the premise. Recent works of [4, 18, 19] showed that by using these values, the adversary can determine to a high accuracy the location of moving objects inside the premise. Roughly speaking, these attacks are based on the idea that when an object crosses the area between an RSS measuring receiver and a transmitter, the RSS values fluctuate appropriately, and this can be used for detecting motion.

An important property of such attacks is that the receivers placed by the adversary are *passive* (they do not generate signals of their own) [1, 2, 5, 16, 21]. Therefore, the *defender* (i.e. the administrator of the wireless network system) is unlikely to even know about the presence of adversaries trying to breach the location privacy of the users. Work of [5] demonstrates how movement can be tracked through solid walls using receivers from outside even when artificial noise is added to the signals in the network. The main questions are thus, how can the defender guard the privacy of the users? Can we quantify the privacy loss? Can we devise strategies that allow us to *guarantee* a small privacy loss?

The recent work of [9] suggests a promising approach: by turning off certain transmitters (or reducing their signal strength significantly) at random time intervals, an adversary will find it difficult to detect movement using the signals of the transmitters. I.e., if the signal quality is poor, the adversary will be unable to make confident inferences about movement based on the RSS values, leading to the protection of user privacy. However, note that in this process, the *utility* of the users of the wireless system could be reduced: users might now need to repeatedly switch between access points, or worse still, might lose connectivity altogether. We thus have an example of the well-known tussle between utility and privacy. The setting of radio window attacks is one example of a more general security scenario, in which we have a defender whose goal is to protect certain assets (in our case, the movement of users). There is a price to pay in terms of loss in utility to the users, or alternately constraints to be met in terms of QoS on the part of the defender. The goal is then to minimize the risk to the assets *irrespective* of the attacker's strategy. A powerful way to view such security problems is as a two-player game between the attacker and the defender, where the *objective* of the attacker is to maximally compromise the number of assets. The defender's goal is to minimize this objective. Depending on the setting, we have different constraints on the attacker and defender's strategies, as well as different ways in which the objective is defined as a function of the strategies.

The game theoretic view has been applied successfully in applications such as the setting up of security checkpoints in airports [15] and the scheduling of air marshals [17]. These works view the problem as a Stackelberg game [7, 8, 11–14] where the defender, *leader*, commits to a strategy first and the attacker, *follower*, observes the leader's strategy and picks an optimal response strategy afterwards. [9] first applied a similar approach to our problem of radio window attacks, and gave a greedy algorithm for finding the optimal *schedule* for turning off the transmitters. The main drawback with these works is that the algorithms proposed are not guaranteed to be efficient, involving the solution of mixed integer programs [10, 13, 14]. In some cases (such as the prior work on radio window attacks [9]), the solution produced is not guaranteed to be optimal. Furthermore, in real settings, we can have additional constraints on the attacker and the defender (as we will see, the defender in our wireless network setting would wish to maintain QoS requirements for the users, and the attacker might need to consider the costs of placing receivers and the chance of being "discovered").

These are the two main motivations of our work. Firstly, we wish to give a general framework that allows us to capture a variety of security problems. The framework should allow us to compute, in an efficient manner, the optimal strategies for the defender. Moreover,

the framework must be able to incorporate problem-specific constraints. Second, we focus on the specifics of radio window attacks, and wish to give a complete, deployable solution for this setting. Modern wireless systems allow us to either use omni-directional (relatively low-frequency) transmitters, or high frequency directional ones. This raises questions such as: how does this choice affect the potential privacy of the users? How does the precise trade-off between the utility (measured via the probability of loss in connectivity) and the privacy look like? We address these questions by using known models for the path loss in signal strength for directional and omni-directional transmitters.

Formulation as min-max optimization

Our setting consists of an indoor wireless network with multiple transmitters (possibly directional). There are devices (i.e. legitimate receivers) which connect to these transmitters to obtain WiFi service. These devices have some quality-of-service (QoS) expectations/requirements. The attacker places receivers outside the premises in order to measure RSS (received signal strength) from transmitters, with the goal of tracking the movements of persons inside using the variation of RSS values. An attacker tracks people based on the drop in RSS values due to the obstacle (user's physical presence) along the path from the transmitter to the attack receiver.

The network administrator (whom we will refer to as the defender) needs to ensure that the attacker's ability to detect movement in the different parts of the building are minimized. Our main idea is that if a transmitter is turned *off* or has its power lowered by a certain level, the attacker cannot detect movement using the signals from that transmitter. Further, if the region of interest (where movement is being detected) is not roughly along the line joining a transmitter and an attacker's receiver, the chances of detection are minimal. Since there are multiple transmitters, the defender can hope to provide service to all the regions while still maintaining privacy. For a moment, consider Figure 1 and suppose that we *know* where the attacker places the receiver at attack zone 1 (and nowhere else). Now, to protect the privacy of Person 1 in the figure, the defender should use transmitter T_2 (written as Tx2 in the figure) and switch off T_1 ; this works because Person 1 is not anywhere close to the line joining T_2 to the attack zone. In general, a defender's strategy consists of turning off a subset of the transmitters. This can be done in a probabilistic way. Further, we may even assume that the defender changes the choice of transmitters to turn off at regular intervals of time. This lets us view the defender's strategy as consisting of a sequence of probability values: $c_{j,t}$ is the probability of turning off transmitter T_j at time t .

Now, in the example above (Fig. 1), we assumed that the defender knew the location of the attacker's receiver, which led to a simple solution (turn off T_1 and use T_2). Even in this case, the strategy protects the privacy of Person 1 but not necessarily Person 4. In general, there is no knowledge of the attacker's exact location but defender knows the set of all the possible locations at which an attacker may place receivers, and the goal is to optimize *for any possible choice* of the defender. Suppose we have a function $U(\sigma_a, \sigma_d)$ that takes a defender strategy σ_d , and attacker strategy (information about the receivers) σ_a and computes the "expected privacy loss". Then, the defender's objective is to find σ_d that minimizes $\max_{\sigma_a} U(\sigma_a, \sigma_d)$.

If the set of all *valid* attacker and defender strategies are denoted by \mathcal{A} and \mathcal{D} respectively, this is equivalent to solving the min-max optimization problem:

$$Z^* = \min_{\sigma_d \in \mathcal{D}} \max_{\sigma_a \in \mathcal{A}} U(\sigma_a, \sigma_d).$$

The value of Z^* is a *guaranteed* bound on the amount of privacy loss (no matter how the attacker plays). The computation of Z^* becomes equivalent to finding the value of a Stackelberg game. We can think of the defender as first *committing* to a strategy σ_d (this could be a “mixed” strategy as we explain below), and based on this, the attacker picks σ_a . If we view the attacker’s goal as maximizing the privacy loss $U(\sigma_a, \sigma_d)$, then we obtain a zero-sum game.

Mixed strategies and randomization. In our setting, the defender can choose a randomized or mixed strategy to turn off transmitters. For instance, in the example earlier from Figure 1, T_1 and T_2 could be turned off with probabilities $2/3$ and $1/2$ respectively (they need not add up to 1 as turning off one transmitter is independent of the others). Further, by dividing up the total duration into different time slots, the defender can turn different transmitters off at different time steps. In this setting, we are interested in the *expected value* of the privacy loss. Randomization gives added power to the defender because even if the attacker were to know the probability of a receiver being turned off, they cannot place the receivers to take advantage of this situation with certainty.

A formulation as above is also possible for other security games (whenever we can reasonably define a utility function as above and identify the set of strategies \mathcal{A} and \mathcal{D}). **Our main contributions can thus be summarized as follows:**

- We show that the min-max optimization formulation above can be solved via a linear programming approach for a range of utility functions U , yielding an optimal solution for the defender’s strategy. The framework is quite general (we explain the setup shortly), and we can plug-in different utility functions, as well as a variety of problem setups (see Section 4).
- For radio-window attacks, we provide a concrete instantiation of the framework, giving expressions for the utility for the case of omni-directional as well as directional transmitters. This utilizes relevant path loss models, taking into account the distance of the attacker’s receivers, the number of obstacles/walls, etc. We show numerically the power of directional transmission in reducing the privacy loss.

1.1 Related Work

Stackelberg security games. Many security games have been formulated as general-sum Stackelberg games (e.g., [13, 14]). Unlike the more standard setting of zero-sum games, the two players in a Stackelberg game do not play simultaneously. Instead, the leader (often the defender in security contexts) plays first, and the follower (typically the attacker) plays next, fully aware of the strategy of the leader, and thus has an advantage. In case of randomized strategies, the follower is aware of the distribution but not the outcomes of the random choice. Typically in applications, the leader and follower have different objectives. But in our formulation, we are only interested in the objective value of the defender, i.e., the total privacy

loss of the users. Thus we incorporate the other aspects of the problem (such as QoS requirements for the users and constraints on the number of receivers an attacker can place) into the formulation as additional constraints on the set of feasible strategies \mathcal{A} and \mathcal{D} . This lets us avoid dealing with multiple objectives.

Linear programming for Stackelberg security games. The focus on a single objective enables us to express the solution to the Stackelberg game (the min-max optimization) as a linear program, which can be solved efficiently. In works such as [14], the attacker strategy is represented using a vector a of length $|A|$ where the i th entry represents the probability that attacker chooses the pure strategy i . They note that (as is standard in min-max theory), for a fixed defender strategy, the attacker’s utility is maximized by a pure strategy. This enables them to focus on binary vectors a , and using these variables, the authors formulate a mixed integer program where attack vector and defender’s vector are variables. In our approach, since we have only a single objective function, we can embed the set of attacker strategies A in constraints of the optimization problem (see Section 3). This allows us to avoid having integer variables in our optimization problem, thus making it possible to find the optimum solution efficiently.

Incomplete information games. The systems discussed in [13, 14] also consider a scenario where there are multiple attackers with different abilities. In our context, this can correspond to the usage of different movement detection algorithms on the part of the attacker, or receivers with different strengths. In the works above, Bayesian Stackelberg games are used to handle this setting. Here, one assumes that the defender is aware of prior probabilities of each attacker type, and the goal is to minimize the expected privacy loss. Our framework can be easily extended to this setting as long as the number of attackers is small and the product of their strategy spaces is without having to know the prior probabilities. If there are k attackers with each having A_1, \dots, A_k different set of strategies, the linear program we construct will have more than $|A_1| \times |A_2| \times \dots \times |A_k|$ constraints. Finally, we show how to extend our framework if we know the prior probabilities of the attackers by allowing the defender to minimize the expected payoff. Here we solve a constant-sum Bayesian Stackelberg game.

Prior work on radio-window attacks. Game theoretic approaches for defending against radio-window attacks have been first studied in [9]. Here, the problem is formulated as a general-sum Stackelberg game and the corresponding optimization problem is solved using a greedy heuristic algorithm. The approach has many limitations: (a) the defender strategy that is produced is not guaranteed to be optimal, (b) constraints about the quality of service for the users is captured in a coarse manner, by limiting the total number of transmitters that can be turned off, (c) the framework is restricted to omni-directional signals, and models for high-frequency, directional transmitters were not considered. Our work allows us to overcome all of these issues, as we will see. We also remove technical restrictions in the work of [9], so we end up making transmitters and target regions “decoupled” from one another (hence making the framework more general).

2 PROBLEM FORMULATION, ADVERSARY ASSUMPTIONS, & FRAMEWORK

Suppose that there are N targets (formally referred to as *target regions*) $\{R_1, \dots, R_N\}$ that defender tries to defend. There are M transmitters T_1, \dots, T_M defender use located in different areas of the premise. There are V time slots, and the defender is allowed to turn off/on different sets of transmitters at different time slots. Legitimate users or devices in each target region R_i connect to one of the transmitters that is within range.

We model the attacker as having S “attack zones”, and the attacker is allowed to place a certain number of receivers (denoted r_a) at one of the zones (denoted s_a).

Now we define the strategies of the two players.

- Defender strategy σ_d is the matrix C of size $M \times V$ where $c_{j,t} \in [0, 1]$ is the probability of turning the j th transmitter off in at the time slot t .
- Attacker strategy σ_a is (r_a, s_a) : setting up r_a attack receivers at the attack zone s_a .

Note that we assume that the attacker cannot change the receiver positions with time. This is because the time slots of interest are relatively short (few seconds or less); in this case, it is unrealistic for an adversary to be able to change the locations of the receivers. (However, as we will see in Section 4.1, a slight modification of our solution allows us to also capture moving or adapting adversaries.) In practical scenarios, it is reasonable to also take into account the overheads due to a defender strategy that “changes too quickly” (users have to keep connecting to new access points). Such an evaluation is left to future work.

2.1 Objective functions of attacker and defender

As discussed earlier, the attacker’s payoff depends on the number of target regions they can successfully “monitor”, and factors such as the cost of placement of the attack receivers at the different attack zones. On the other hand, the defender’s payoff depends on the privacy loss as well as factors that capture the loss of *utility* (e.g., the probability that a user in a certain region loses connectivity).

Let us focus on the privacy loss, or the number of targets that the adversary can successfully monitor. As we will discuss, this will be the “main” objective value in our formulation (it is the key component in both the attacker and defender’s objective), and the other terms will be viewed as constraints.

To formalize this, we need a model that captures the probability that an attacker detects movement in a given target region (conditioned on movement in that region) by placing receivers at a certain attack zone. **Our optimization framework allows us to plug in any model that defines the parameter $AD(r_a, s_a, j, k)$ described below**, and in this sense, our framework is plug-n-play. But in the context of radio window attacks, we give a concrete proposal (as it also illustrates the flexibility in the framework). First, note that the attacker could use RSS values from any of the transmitters in order to detect movement in a particular target region (and the overall probability of detection is that of the union of these events). Let us consider the attacker strategy (r_a, s_a) , and denote

the ability of the attacker to detect movement in target region j using the signal from transmitter k as $AD(r_a, s_a, j, k)$. Our framework allows to plug any definition of $AD()$ depending on the problem scenario. Here we propose a simple approach to derive a form for $AD(r_a, s_a, j, k)$.

The considerations here are as follows. First, for omnidirectional transmitters, the attacker is assumed to be able to measure the RSS from all the transmitters, while for *directional* transmission, the attacker can only measure the RSS from transmitters within a certain angle around the transmission direction (see Figure 2). Next, the probability of detecting movement depends on the RSS values (weak signals lead to low confidence in detection), and also the number of receivers placed (r_a). Finally, we assume that detecting movement in a region R_j is possible only if R_j is approximately along the line joining the transmitter and the attack zone. To capture these parameters, we first introduce a term that we call the “detection coefficient”.

Let $CF(s_a, R_j, T_k)$ be the detection coefficient corresponding to target region j , attack zone s_a , using signals from the k th transmitter. Again, note that what we provide below is one choice of CF; it can be replaced with other models. We adopt a path loss model similar to the models discussed in [18] along with Friis equations to approximate the received signal strength at the attack zones and target regions. Let y'_{s_a, T_k} be the signal strength from transmitter T_k at the attack zone s_a . We compute this quantity using the power at the transmitter minus the path loss. I.e., if P_k is the transmitted power at transmitter k in dB, then

$$y'_{s_a, T_k} = P_k - f_{s_a, T_k}, \quad (1)$$

where f_{s_a, T_k} path loss at s_a due to the distance between T_k and s_a and obstacles such as walls along the path. The detection coefficient will be normalized such that the maximum value is 1. This is done by defining

$$y_{s_a, T_k} = \left(y'_{s_a, T_k} - \max_{s'_a, T'_k} y'_{s'_a, T'_k} \right) \times \delta \quad (2)$$

where $0 < \delta < 1$ is a scaling factor which can be used to control the magnitude of the $|y_{s_a, T_k}|$. (If $|y_{s_a, T_k}|$ is too large, $CF(s_a, R_j, T_k)$ may vanish when taking the exponent of y_{s_a, T_k} as we explain next).

Let D be an array of dimensions $M \times N \times S$, where D_{T_k, R_j, s_a} is 1 if attack zone s_a and target region R_j are in the same direction from transmitter T_k , otherwise 0 (these values can also be chosen to be fractions; if the line joining T_k to s_a very slightly intersects R_j , the fraction can be set to a smaller value than in the case when the line has a significant overlap with R_j ; to keep things simple, our model uses binary values). Now we compute the detection coefficient as:

$$CF(s_a, R_j, T_k) = 10^{y_{s_a, T_k}} \times D_{T_k, R_j, s_a}. \quad (3)$$

To summarize, this captures the fact that an attack receiver can detect movement in a target region only if the signal is strong enough and if the region is on the path from the transmitter to the attack zone. Let us now see how to compute CF values for two types of protocols.

- (1) *WiFi (802.11a, 802.11b, 802.11c)*: Here, the signals are low frequency signals (2.4 GHz, 5 GHz) and they typically travel long distances due to their low path loss. They also pass through thick walls without suffering huge losses in signal




Figure 2: Fraction of target regions covered by an attack zone with respect to transmitters.

strength therefore a receiver quite far from a transmitter (around 75 m) can observe reasonable signal strengths with or without direct line of sight (due to the presence of walls).

- (2) *Millimeter wave (802.11ad)*: Here, the signals are high frequency (typically 30 – 70 GHz) and the path loss is very high, therefore these signals do not travel far. But the receivers within the range of the antennas enjoy high data rates. These signals suffer huge losses in signal strength as they travel through walls therefore without direct line of sight, a receiver at a large distance from the transmitter cannot expect to observe measurable signal strengths.

In the example in figure 2, $CF(s_a, R_2, T_2)$ is much higher than $CF(s_a, R_2, T_1)$ as transmitter T_1 is much further away from the attack zone (in addition to the loss in signal strength due to multiple walls). To consider another example, $CF(s_a, R_3, T_1) \neq 0$ while $CF(s_a, R_3, T_2) = 0$ because R_3 has no overlap with the line joining T_2 and s_a . We discuss more details on how the precise detection coefficients are derived in our experiments (§5). We note that this framework is not limited to the given example protocols, but it can be extended to handle different types of antennas (directional, omni-directional as we will see in Section 5) and different protocols.

Finally, we allow the attacker to place multiple receivers at an attack zone (denoted r_a in our definition of the attacker’s strategy). Intuitively, placing more signals allows a receiver to boost the detection probability. We denote this boost factor by $\Lambda(r_a, s_a)$. Again, our framework does not place any restriction on how to define this. For simplicity, we follow the model from [9], where the factor increases linearly until a certain value λ and then flattens. Formally,

$$\Lambda(r_a, s_a) = \begin{cases} \frac{r_a}{r_{\max}(s_a)} \times \lambda & \text{if } r_a \leq r_{\max}(s_a), \\ \lambda & \text{otherwise} \end{cases} \quad (4)$$

Here the parameters are $r_{\max}(s_a)$ and λ .

Using the detection coefficient and the boost parameter, we compute $AD(r_a, s_a, j, k)$

$$AD(r_a, s_a, j, k) = CF(s_a, j, k) \times \Lambda(r_a, s_a) \quad (5)$$

Lastly, to write down the objective values of the players, we need to consider the *actual* movement inside the premise. As discussed in the introduction, we assume that the defender knows the likelihood of movement in every target region in each time slot. We denote the probability of movement in target region j at time slot t by $pm_{j,t}$. The attacker’s goal is to successfully detect movement in the different target regions. As a simple way to capture this, we assume a gain of α (for the attacker) for every successful detection, for some constant α .

We can now write an expression for the *expected payoff* or objective value of the attacker. For strategies σ_a and σ_d of the attacker and defender respectively, this will be denoted as $U(\sigma_a, \sigma_d)$, where $\sigma_a = (r_a, s_a)$ and $\sigma_d = C$. The expression for this is

$$U(\sigma_a, \sigma_d) = \sum_{t=1}^V \sum_{j=1}^N \alpha \cdot pm_{j,t} \sum_{k=1}^M (1 - c_{k,t}) AD(r_a, s_a, j, k). \quad (6)$$

At each time step, the expected payoff is the expected number of j in which movement is successfully detected. The term $\sum_k (1 - c_{k,t}) AD(r_a, s_a, j, k)$ is the probability that movement in R_j is detectable using one of the k transmitters. With this probability, we get a reward of α , but only conditioned on movement in region j at that time, which happens with probability $pm_{j,t}$. The expected payoff after setting α to 1 quantifies the expected total movement detected by the attacker. This objective also directly measures the “total privacy loss” of the defender. Thus the defender’s goal is to minimize $U(\sigma_a, \sigma_d)$. One limitation of the framework is that even if we consider slightly different utility functions, e.g., number of locations where the attacker “reliably” measures movement, the objective not linear, and thus the LP framework does not apply.

Attacker’s risk. An important quantity that was not considered above is the *risk* involved for the attacker – placing too many receivers can be expensive (and can involve a risk of being detected). We model this via constraints in our formulation (as we see below). One simple setting is to have a cost β for placing each receiver, and having a risk associated with each attack zone. We denote the former by $\text{cost}(r_a) := \beta \cdot r_a$ and the latter by $\text{risk}(r_a, s_a) = \gamma_{s_a} \cdot r_a$ where γ_{s_a} is a constant that depends on the attack zone s_a . We assume that the maximum cost + risk the attacker can afford to withstand is E . I.e., the attacker only chooses strategies such that

$$\text{cost}(r_a) + \text{risk}(r_a, s_a) \leq E. \quad (7)$$

3 OUR APPROACH

3.1 Min-max optimization problem

As we discussed earlier, the defender’s goal is to choose a strategy s_d that minimizes the *worst possible damage* that the attacker can do. Let us define \mathcal{D} and \mathcal{A} to be the set of feasible strategies for the defender and attacker respectively. As discussed before, \mathcal{D} will consist of matrices with probability values (probabilities of turning off transmitters at different times), and \mathcal{A} consists of pairs (r_a, s_a) . Additionally, we will have constraints on these sets, as we will discuss. The defender’s goal is to find:

$$\min_{\sigma_d \in \mathcal{D}} \max_{\sigma_a \in \mathcal{A}} U(\sigma_a, \sigma_d).$$

From the discussion above on the attacker's maximum risk,

$$\mathcal{A} = \{(r_a, s_a) : \text{cost}(r_a) + \text{risk}(r_a, s_a) \leq E, r_a \in [1, r_{\max}], s_a \in S\}$$

Clearly, $|\mathcal{A}| \leq |S| \cdot r_{\max}$.

3.2 A Linear Program

The nice property of the min-max formulation above is that because of the form of $U()$, we can convert it into an equivalent linear program (LP). $U()$ is a linear function of the defender strategy matrix C . We let the variables of the LP be the entries of the matrix C . I.e., for each transmitter j and time step t , we have a variable $c_{j,t}$. We have an additional variable x which is "intended" to take the value $\max_{\sigma_a \in \mathcal{A}} U(\sigma_a, \sigma_d)$. To enforce this, we impose the constraint:

$$x \geq U(\sigma_a, \sigma_d), \text{ for all } \sigma_a \in \mathcal{A}.$$

Given any σ_a , the constraint above is linear in the variables $c_{j,t}$. Thus consider the following linear program:

$$\min_{\sigma_d \in \mathcal{D}} x \quad \text{subject to} \quad (8)$$

$$x \geq U(\sigma_a, \sigma_d), \text{ for all } \sigma_a \in \mathcal{A}, \quad (9)$$

$$0 \leq c_{j,t} \leq 1, \text{ for all } j \in [M], t \in [V] \quad (10)$$

Now, for any choice of the $c_{j,t}$ variables, the best choice of x is $\max_{\sigma_a \in \mathcal{A}} U(\sigma_a, \sigma_d)$ (because we are minimizing x subject to the conditions (9)). Thus the optimal solution finds $c_{j,t}$ to minimize the max value, which is precisely the solution to the min-max problem.

Complexity. The number of variables in this linear program is $MV + 1$ and the number of constraints is $\leq |S| \cdot r_{\max} + 2MV$, and thus the LP can be solved efficiently (polynomial time in theory; and in practice, we can solve it with tens of thousands of variables and constraints). The limitation of our approach is that it performs poorly when case when the number of attacker strategies is large (e.g., when an attacker is allowed select l attack zones as we see below). Developing other methods for such cases is an interesting direction for future work.

3.3 Defender QoS constraints

Even though the defender's main goal is to protect the privacy of the users, it cannot entirely compromise the connectivity of users within the premise (note that the trivial solution of switching all the transmitters off, i.e., $c_{j,t} = 1$ for all j, t , perfectly preserves privacy; the objective value is 0, but it also nullifies the utility of the system). Ideally, if a certain transmitter turns off, the defender needs to ensure that all the legitimate receivers of target regions that receive signals from this transmitter can connect to different transmitters within the network. We describe this requirement using following notation.

- For each target region R_j , find the set of transmitters that devices can connect (within the range of transmitter signals).
- Consider the $N \times M$ binary matrix Q , where $Q_{j,k}$ indicates whether transmitter k is reachable from the target region j . If reachable $Q_{j,k} = 1$ otherwise 0.

Now we can write the following set of constraints for parameter p_{\min} which indicates minimum connectivity "level" for each user at every time step. We impose the constraint:

$$\sum_{k=1}^M Q_{j,k}(1 - c_{k,t}) \geq p_{\min}, \quad \forall j \in [N] \text{ and } \forall t \in [V].$$

Having p_{\min} large enough (close to 1) ensures that the probability of turning off *all* reachable transmitters (setting all of these $c_{j,t}$ values to 1) is unlikely.

3.3.1 Stronger attacker strategies. So far, we considered that the attacker can place receivers at precisely one attack zone s_a . What if the attacker can place receivers at multiple zones? Our framework can be extended to handle this case as well, except that now, the number of distinct attacker strategies is significantly higher; thus the number of constraints in the LP formulation increases. We now show how to model this scenario.

- Suppose the attacker can choose l attack zones out of $|S|$.
- Let us define $\sigma_a^i = (s_a^i, r_a^i)$ where $i \in [l]$. This indicates the strategy of the attacker for the i th chosen location.
- Let \mathcal{A} be the set of all attacker strategies. These are determined by $\sum_{i \in [l]} \text{cost}(r_a^i) + \text{risk}(r_a^i, s_a^i) \leq E$. We thus have

$$|\mathcal{A}| \leq \binom{|S|}{l} \times r_{\max}.$$

- Now an overall strategy of the attacker is defined by the tuple: $\sigma_a = [\sigma_a^1, \dots, \sigma_a^l]$.
- Let $u(\sigma_a^i, \sigma_d)$ be the payoff for i th attack zone using eq. (6). Then the total utility of the attacker is

$$U(\sigma_a, \sigma_d) = \sum_{i=1}^l u(\sigma_a^i, \sigma_d)$$

As before, the min-max optimization is captured via the following linear program.

$$\min_{\sigma_d \in \mathcal{D}} x \quad \text{subject to} \quad (11)$$

$$x \geq U(\sigma_a, \sigma_d), \text{ for all } \sigma_a \in \mathcal{A} \quad (11)$$

$$0 \leq c_{j,t} \leq 1, \text{ for all } j \in [M], t \in [V] \quad (12)$$

$$\sum_{k=1}^M Q_{j,k}(1 - c_{k,t}) \geq p_{\min}, \text{ for all } j \in [N], t \in [V] \quad (13)$$

Time complexity. The variables are once again the probabilities $c_{k,t}$ and x . Therefore, number of variables is $MV + 1$. The first set of constraints has $\binom{|S|}{l} \times r_{\max}$ constraints. The second set has $2MV$ constraints. The third has NV constraints. Therefore, the total number of constraints is $\leq \binom{|S|}{l} \times r_{\max} + V(2M + N)$.

4 APPLICATIONS OF THE FRAMEWORK

We now give examples of more general settings to which the framework in Section 2 applies with little modifications. The first two will be extensions of radio window attacks, while the next one will be unrelated security games that have been considered in prior works.

4.1 Adaptive attackers

We consider an attacker who is capable of changing their strategy at each time step. This captures a situation where the attacker potentially learns some information about the future movements of the users and then decides to change their strategy. In this case, we effectively have a different σ_a for every time step. While this leads to an explosion in the number of attacker “strategies”, we observe that in this case, the defender strategies at the different time steps are all independent of one another, and hence the defender can solve a separate optimization problem for each time step.

In the optimization problem for the time step t , when computing the payoff function, we consider probability of movement $pm_{j,t}$ in target regions and the defender strategy vector $c_{k,t}$. With this setting we obtain the following payoff function for the time step t .

$$U_t(\sigma_a, \sigma_d) = \sum_{j=1}^N pm_{j,t} \sum_{k=1}^M (1 - c_{k,t}) AD(r_a, s_a, j, k) \quad (14)$$

Now we solve an independent optimization problem for each time step t . For $t = 1$ to $t = V$, we do:

$$\begin{aligned} \min_{\sigma_d \in D} \quad & x \quad \text{subject to} \\ & x \geq U_t(\sigma_a, \sigma_d), \text{ for all } \sigma_a \in \mathcal{A} \\ & 0 \leq C_{j,t} \leq 1, \forall j \in [M], \forall t \in [V] \\ & \sum_{k=1}^M Q_{j,k}(1 - c_{k,t}) \geq p_{\min}, \forall j \in [N] \end{aligned}$$

The resultant matrix C gives the optimal strategy for the defender when the attacker is adaptive.

4.2 Incomplete information

We consider the situation where there are multiple attackers with different kinds of resources. For example, the types of receivers they use can affect the antenna gains and therefore utilities can vary. One way to model such scenarios was proposed in [13], where the defender is assumed to have a prior distribution that gives the probability of each attacker type. Such games are known as Bayesian Stackelberg games, or incomplete information games.

Let there be H different attackers and let A_1, \dots, A_H be the sets of attacker strategies. Let the payoff functions of attackers be $U(\sigma_a^1, \sigma_d), \dots, U(\sigma_a^H, \sigma_d)$. Let the prior probabilities of the attackers be p_1, \dots, p_H . Here the defender tries to minimize the **expected value** of the maximum payoffs of attackers. We solve the following optimization problem in order to compute the optimal defender strategy for this setting.

$$\begin{aligned} \min_{\sigma_d \in D} \quad & \sum_{h=1}^H p_h x_h \quad \text{subject to} \\ & x_h \geq U(\sigma_a^h, \sigma_d), \forall \sigma_a^h \in A_h, \forall h \in [H] \\ & 0 \leq C_{j,t} \leq 1, \forall j \in [M], \forall t \in [V] \\ & \sum_{k=1}^M Q_{j,k}(1 - c_{k,t}) \geq p_{\min}, \forall j \in [N] \text{ and } \forall t \in [V] \end{aligned}$$

Note that we introduced a separate variable for each attacker type that takes a value equal to the maximum payoff over the strategies for that attack type. The defender then minimizes the expectation of these maximum payoffs.

4.3 Other security games

We show how our approach of designing a linear program can be used to solve ORIGAM-MILP in [10]. In ORIGAM-MILP, the defender solves the following mixed-integer program. Let T be the set of targets. Let c, a be length $|T|$ vectors that denote the defender’s coverage vector and attacker’s attack vector respectively. Let $U(t, c)$ be the payoff of the attacker for when attacker chooses target t and defender plays coverage vector c . Here both t and c are parameters to U . Let Z be a very large number.

$$\begin{aligned} \min_c \quad & k \quad \text{subject to} \\ & a_t \in \{0, 1\} \quad \forall t \in T \\ & c_t \in [0, 1] \quad \forall t \in T \\ & \sum_{t \in T} c_t < m \\ & U(t, c) \leq k \quad \forall t \in T \\ & k - U(t, c) \leq (1 - a_t)Z \quad \forall t \in T \\ & c_t \leq a_t \quad \forall t \in T \end{aligned} \quad (15)$$

Constraint 4 makes sure that the defender minimizes the maximum payoff of the attacker and the constraint 5 ensures that payoff of $U(t, c)$ is maximized for the target t chosen by the attacker. Now consider a payoff function $U_t(c)$ for target t where only c is a parameter. We eliminate the attack vector a (which is a binary vector) and write the following linear program.

$$\begin{aligned} \min_c \quad & k \quad \text{subject to} \\ & c_t \in [0, 1] \quad \forall t \in T \\ & \sum_{t \in T} c_t < m \\ & U_t(c) \leq k \quad \forall t \in T \end{aligned} \quad (16)$$

Now there are $|T|$ payoff functions for each target and the defender minimizes the maximum of these payoff functions. Here we would not know which target is being attacked. But the optimum values of the optimization problems (15) and (16) are the same. The last constraint in (15) sets the coverage probabilities of all targets that were not attacked to 0. When we solve (16), a linear program solver may have assigned non-zero values for these targets. An important assumption in ORIGAM is that payoff of the attacker does not depend on the targets that are not attacked. Therefore, changing the values of un-attacked targets in the coverage vector does not affect the solution. Thus, after solving our linear program, we can check which target gave the optimum for the solution coverage vector and set all other coverage probabilities to 0. In this approach we solve a linear program with half the variables and same number of constraints as in ORIGAM-MILP. In [10], the authors show that the solution to program 15 is also a solution to ERASER-MILP [13]. Thus, our solution above can be used as a solution for ERASER-MILP as well.




Figure 3: Experiment layout of the building with transmitters, target regions, and attack zones.

5 EXPERIMENTAL RESULTS

In this section, we evaluate our framework, the goal is to observe the following: (a) how the schedules obtained by solving our optimization problem perform compared to simple baselines, such as random scheduling, (b) how the payoff the attacker varies depending on changing QoS requirements of the users (captured by the minimum probability the defender allows for target regions to lose connectivity) as well as the amount of resources the attacker has, (c) how directional antennas and omni-directional antennas compare with one another in terms of the defender’s utility. In these experiments, we consider the optimization problem we designed for the adaptive attacker where each time step is independent.

We consider two types of wireless antennas in these experiments. a) WiFi signals, b) millimeter wave signals. We model how the signals propagate from transmitters till they reach attack receivers. Using this, we derive the values of $CF(.)$ (eq (3)) for attack zones.

The layout of the building we consider with attack zones, transmitters, and target regions is shown in figure 3. Here there 6 transmitters T_1, \dots, T_6 , 5 attack zones S_1, \dots, S_5 , and 12 target regions R_1, \dots, R_{12} . Distances are measured in meters(m).

The table 1 shows which transmitters are reachable from different target regions. The matrix form of this table will be used to ensure that QoS requirements of the devices in the target regions are maintained in our experiments.

5.1 Computing the detection coefficients $CF(.)$

In this section we discuss how we compute $CF(.)$ values for each attack zone, target region, and transmitter. Here we consider two types of protocols.

- (1) Standard WiFi (802.11a, 802.11b, 802.11c, etc.)

	T_1	T_2	T_3	T_4	T_5	T_6
R_1	1	1	0	1	1	0
R_2	1	1	0	1	0	0
R_3	1	1	1	1	0	0
R_4	0	1	1	1	0	0
R_5	1	1	0	1	1	0
R_6	1	1	1	1	1	1
R_7	1	1	1	1	0	0
R_8	0	1	1	1	0	0
R_9	1	0	0	1	1	1
R_{10}	1	0	1	1	1	1
R_{11}	0	0	0	1	1	1
R_{12}	0	0	0	1	1	1

Table 1: Transmitters reachable from the different target regions - matrix Q

- (2) Millimeter wave (802.11ad)

We derive $CF(.)$ values for the building layout in figure 3 when each of these protocol types used in transmitter antennas separately. We assume that the gains of the transmitters and the attack receivers are 10dB.

5.1.1 Standard WiFi signals. Let d be the distance between receiver and the transmitter. Here we consider transmitters of frequency 2.4GHz. We assume the power at transmitter is $-3dB$. Let P_r be the received signal strength at the receiver. Let η be the path loss exponent. Now using Friis equations,

$$P_r = P_0 - 10\eta \log_{10}\left(\frac{d}{d_0}\right)$$

We use Friis equations to calculate P_0 . Here $d_0 = 1m$. Let P_t be the transmitted power at the transmitter. Let $G_t, G_r = 10dB$ be gains at the transmitter and the receiver. Then

$$P_0 = P_t + G_t + G_r - 20 \log_{10}\left(\frac{4\pi d_0}{\lambda}\right)$$

For the values we defined we get $P_0 = -23dB$. Assuming $\eta = 2$, we compute P_r at the receiver as

$$P_r = -23 - 20 \log_{10}(d)$$

5.1.2 Millimeter wave protocols. Here we assume the frequency of the signals is 60GHz. We use Friis equations to compute P_0 as the standard WiFi signal propagation model. From that we get $P_0 = -51dB$. The obstacles such as wall on the path affect high frequency waves adversely unlike low frequency WiFi waves. Therefore, when computing P_r at a receiver there is an extra penetration loss term (as discussed in path loss models in [3]). For dry walls in indoor offices this is typically 6dB. In our building layout, for each wall encountered on the path from a transmitter to a receiver, we add a penetration loss component to the overall path loss. Thus,

$$P_r = -51 - 20 \log_{10}(d) - 6w$$

where w is the number of walls on the path from transmitter to receiver.

We use these equations to calculate the received signal strengths at the attack zones. Then we use these values to compute y_{s_a, R_j, T_k} and the corresponding $CF(s_a, R_j, T_k)$ values. When computing y_{s_a, R_j, T_k}

	T_1	T_2	T_3	T_4	T_5	T_6
WiFi	0.503	0.005	0.503	0.992	0.294	0.214
Millimeter wave	0.355	0.314	0.355	0.830	0.215	0.454

Table 2: Average probabilities of transmitters turning off

values, we fix $\delta = 0.1$ (Eq. 2). Note that one can replace these functions with different path loss models or different received signal strength calculation methods based on the requirement.

5.2 Numerical evaluations

In this section we evaluate the effectiveness of the transmitter turn-off schedule we obtain from our framework. In these experiments we set the constant for receiver strength $\lambda = 2$ (Eq. (4)), constant for the gain of the attacker by observing movement at a target region $\alpha = 10$, constant for cost of the attacker $\beta = 0.005$, and constants for risk in all attack zones $\gamma_{s_a} = 0.005$ (Eq. (7)). Note that these constants can be changed based on the priority of each component in the payoff function in this framework. Here we set the number of time steps $V = 100$. The duration of a time slot can be varied depending on the requirement. We assume the maximum number of receivers that the attacker can deploy at an attack zone is 30. We consider the stronger version of the attacker where they are able to choose 4 attack zones in the building layout in Fig. 3. For the purpose of the experiments we generate the movement probability matrix pm according to $Unif(0, 1)$. We set the required QoS threshold $p_{min} = 1.5$. In these experiments, unless mentioned otherwise we assume that the sum of total cost and the total risk the attacker withstand is less than 0.2. Attacker does not choose a strategy if the cost + risk (E) exceeds this value (Eq. (7)).

After solving our optimization problem with the aforementioned settings, the average probabilities of turning transmitters off over 100 time slots are shown in the table 2.

5.2.1 Payoff of the attacker with p_{min} . In section 3 we discussed how the defender can maintain the undisturbed connectivity for the devices in the target regions by controlling the the minimum threshold probability required for the transmitters reachable from a single target region to keep turned on. Intuitively the ability of the attacker to detect movement at a given target region decreases as we decrease this probability threshold p_{min} . But smaller p_{min} affects the connectivity of legitimate devices in the network adversely. Here we evaluate how increasing p_{min} affects the overall attacker payoff with respect to two protocols we consider. We compute the average payoff over 100 time steps for each p_{min} value.

In figures 4 and 5 we notice that increasing p_{min} threshold indeed increases the attacker payoff since keeping the transmitters on increases the ability to detect movement, but keeping p_{min} low comes with the cost of losing connectivity of legitimate devices. The payoff increases in a slower rate for smaller values of p_{min} and in a faster rate after larger p_{min} values. This pattern is visible clearly in the WiFi signals. This helps in deciding what would be the optimal value a defender can set for p_{min} to minimize the attacker payoff while maintaining high QoS standards.

5.2.2 Payoff of the attacker with attacker resources. Here we intend to demonstrate how an attacker exhausting more resources in order

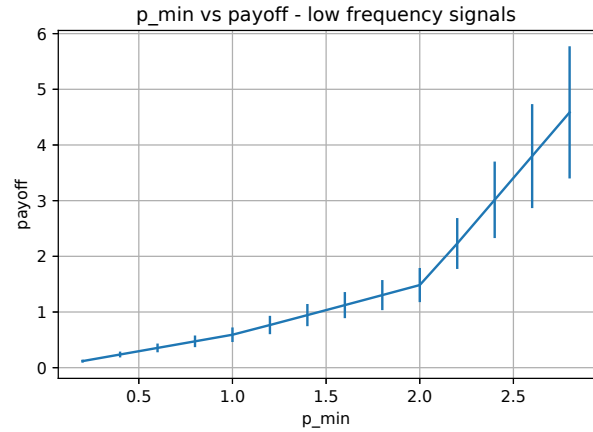


Figure 4: Attacker payoff rise with p_{min} for WiFi signals.




Figure 5: Attacker payoff rise with p_{min} for millimeter wave protocol

to determine movement in target areas affects the overall payoff of the attacker. In the section 2, we discuss the constant for cost β and the constant for risk at attack zone $s_a: \gamma_{s_a}$ and these can be used to control amount of resources the attacker can exhaust. Let $E = \sum \text{cost}(r_a) + \sum \text{risk}(r_a, s_a)$ (eq 7) be the amount of resources the attacker is allowed to use. We demonstrate how the payoff of the attacker changes as E is increased. For the WiFi signals, we set $\beta = 0.005$ and $\gamma_{s_a} = 0.005 \forall s_a \in S$. For the millimeter wave signals, we set $\beta = 0.005$ and $\gamma_{s_a} = 0.005 \forall s_a \in S$. Figures 6 and 7 show how the average payoffs over 100 time steps behave as we increase E . We observe that the payoff of the attacker increases as E increases up to a certain E then it saturates.

We note that this behavior depends on the priority we set for the cost and the risks of the attacker by changing β and γ_{s_a} . If the detectability of movement in target region is of the highest priority compared to costs and risks, then by increasing E , attacker can continuously expect increased payoffs.

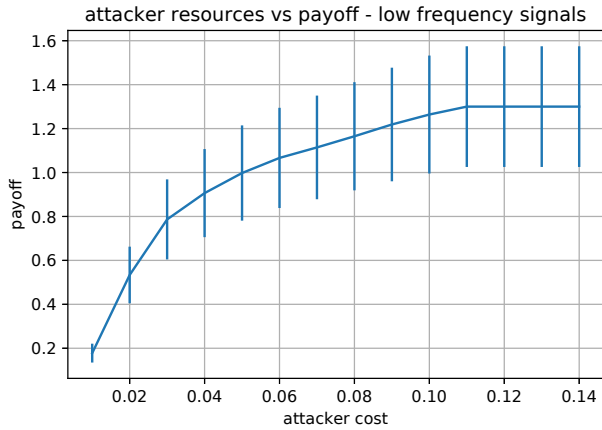


Figure 6: Payoff of the attacker with cost+risk for WiFi signals.

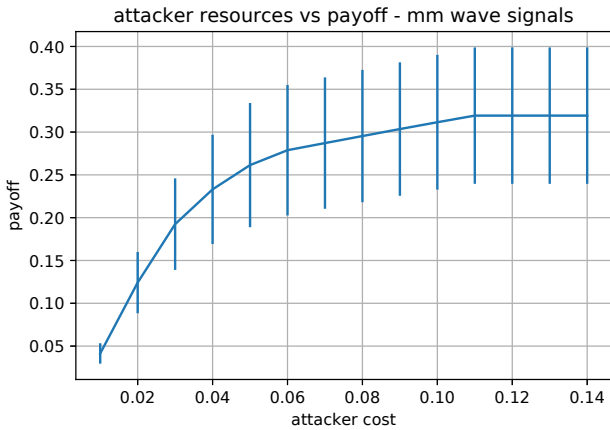


Figure 7: Payoff of the attacker with cost+risk for millimeter wave protocol

5.2.3 *Performance of the framework compared to a random transmitter scheduling.* In this experiment we fix $l = 4$ and $p_{min} = 1.5$. First we compute the optimal payoff given by solving the optimization problem for $V = 100$ time steps. With the same set of parameters, we also compute a random transmitter schedule $C'_{.,t}$ that satisfies all constraints in the optimization problem. We use this $C'_{.,t}$ to compute the maximum payoff for this schedule. Then we take the average of these maximum payoffs over 100 time steps.

Table 3 shows that the optimal payoff given by our framework is less than the payoff given by a random schedule. We also note that this pattern holds for every time step as well. Therefore, the defender is always better off playing the optimal strategy given by this framework with respect to this payoff function.

5.2.4 *Attacker payoff of omni-directional and directional transmission of millimeter wave signals.* Here we consider two cases where the defender transmitters transmitting signals in a single direction

	Average optimum	Average maximum payoff for random C
WiFi	1.034	5.372
Millimeter wave	0.226	0.591

Table 3: Average optimal payoff computed by the framework and average maximum attacker payoff with random schedule over 100 time steps.

	Average payoff
Single directional transmission	0.226
Omni-directional transmission	1.616

Table 4: Average optimal payoffs of directional and omni-directional transmission for millimeter wave signals over 100 time steps.

(targeted at a legitimate receiver) and the transmitters transmitting signals in an omni-directional manner. In the omni-directional setting while the received signal strength at the target regions that are located in the same direction as the attack zone remain the same, we assume that gain at the target regions that are located within an angle of 90 degrees from the direction of the attack zone is reduced by 5dB [6, 20]. Then we set $D_{k,j,s_a} = 1$ for all target regions R_j that are within an angle of 90 degrees from the line between transmitter T_k and the attack zone s_a and calculate $CF(\cdot)$ values. We solve the optimization problem for these two cases.

Table 4 shows the average payoff of the attacker over 100 independent time steps. It is seen that the defender is better off using directional transmitters as opposed to the omni-directional transmitters with respect to this payoff model.

6 CONCLUSIONS AND FUTURE DIRECTIONS

We developed a novel, flexible, plug-n-play framework based on a constant-sum Stackelberg game for defending against radio window attacks. Our framework enables us (i) to comprehensively capture the constraints of the attacker and the defender, (ii) to capture features of modern wireless systems such as directional antennas, and (iii) allows us to plug in different path-loss models with minimal changes to the setup. We formulated the problem of finding the optimal defender strategy as a linear program and showed that it can be solved efficiently. We also performed numerical evaluations to demonstrate the applicability of our framework. In the future, we will collect signal strength data in various experimental settings and plug into our framework for data-driven comparisons. We will investigate the practical concerns that arise such as controlling and coordination of access points and determining the size of time slots when implementing the strategies obtained by this framework.

ACKNOWLEDGMENTS

This material is based upon work supported by the Army Research Office under Grant No. W911NF-17-1-0457.

REFERENCES

- [1] Fadel Adib, Zach Kabelac, Dina Katabi, and Robert C Miller. 2014. 3D tracking via body radio reflections. In *11th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 14)*. 317–329.
- [2] Fadel Adib and Dina Katabi. 2013. See through walls with WiFi!. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. 75–86.
- [3] Christopher R Anderson and Theodore S Rappaport. 2004. In-building wideband partition loss measurements at 2.5 and 60 GHz. *IEEE transactions on wireless communications* 3, 3 (2004), 922–928.
- [4] Paramvir Bahl and Venkata N Padmanabhan. 2000. RADAR: An in-building RF-based user location and tracking system. In *Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No. 00CH37064)*, Vol. 2. Ieee, 775–784.
- [5] Arijit Banerjee, Dustin Maas, Maurizio Bocca, Neal Patwari, and Sneha Kasera. 2014. Violating privacy through walls by passive monitoring of radio windows. In *Proceedings of the 2014 ACM conference on Security and privacy in wireless & mobile networks*. 69–80.
- [6] Guillermo Bielsa, Adrian Loch, Irene Tejado, Thomas Nitsche, and Joerg Widmer. 2018. 60 GHz networking: Mobility, beamforming, and frame level operation from theory to practice. *IEEE Transactions on Mobile Computing* 18, 10 (2018), 2217–2230.
- [7] Matthew Brown, Bo An, Christopher Kiekintveld, Fernando Ordóñez, and Milind Tambe. 2014. An extended study on multi-objective security games. *Autonomous agents and multi-agent systems* 28, 1 (2014), 31–71.
- [8] Debarun Kar, Thanh H Nguyen, Fei Fang, Matthew Brown, Arunesh Sinha, Milind Tambe, and Albert Xin Jiang. 2017. Trends and applications in Stackelberg security games. *Handbook of Dynamic Game Theory* (2017), 1–47.
- [9] Mojgan Khaledi, Mehrad Khaledi, Sneha Kumar Kasera, and Neal Patwari. 2016. Preserving Location Privacy in Radio Networks Using a Stackelberg Game Framework. In *Proceedings of the 12th ACM Symposium on QoS and Security for Wireless and Mobile Networks*. ACM, 29–37.
- [10] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. 689–696.
- [11] Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. 2010. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- [12] Dmytro Korzhyk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. 2011. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research* 41 (2011), 297–327.
- [13] Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. 2008. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 895–902.
- [14] Praveen Paruchuri, Jonathan P Pearce, Milind Tambe, Fernando Ordonez, and Sarit Kraus. 2007. An efficient heuristic approach for security against multiple adversaries. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*. 1–8.
- [15] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. 2008. Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport.. In *AAMAS (Industry Track)*. 125–132.
- [16] Qifan Pu, Sidhant Gupta, Shyamath Gollakota, and Shwetak Patel. 2013. Whole-home gesture recognition using wireless signals. In *Proceedings of the 19th annual international conference on Mobile computing & networking*. 27–38.
- [17] Jason Tsai, Shyamsunder Rathi, Christopher Kiekintveld, Fernando Ordonez, and Milind Tambe. 2009. IRIS-a tool for strategic security allocation in transportation networks. *AAMAS (Industry Track)* (2009), 37–44.
- [18] Joey Wilson and Neal Patwari. 2010. Radio tomographic imaging with wireless networks. *IEEE Transactions on Mobile Computing* 9, 5 (2010), 621–632.
- [19] Moustafa Youssef, Matthew Mah, and Ashok Agrawala. 2007. Challenges: device-free passive localization for wireless environments. In *Proceedings of the 13th annual ACM international conference on Mobile computing and networking*. 222–229.
- [20] Yanzi Zhu, Ying Ju, Bolun Wang, Jenna Cryan, Ben Y Zhao, and Haitao Zheng. 2018. Wireless side-lobe eavesdropping attacks. *arXiv preprint arXiv:1810.10157* (2018).
- [21] Yanzi Zhu, Zhujun Xiao, Yuxin Chen, Zhijing Li, Max Liu, Ben Y Zhao, and Haitao Zheng. 2018. Et Tu Alexa? When Commodity WiFi Devices Turn into Adversarial Motion Sensors. *arXiv preprint arXiv:1810.10109* (2018).