Power Napping with Loud Neighbors: Optimal Energy-Constrained Jamming and Anti-Jamming

Bruce DeBruhl, Christian Kroer, Anupam Datta, Tuomas Sandholm, and Patrick Tague
Carnegie Mellon University
{debruhl@, ckroer@cs., danupam@, sandholm@cs., tague@}cmu.edu

ABSTRACT

The openness of wireless communication and the recent development of software-defined radio technology, respectively, provide a low barrier and a wide range of capabilities for misbehavior, attacks, and defenses against attacks. In this work we present *finite-energy jamming games*, a game model that allows a jammer and sender to choose (1) whether to transmit or sleep, (2) a power level to transmit with, and (3) what channel to transmit on. We also allow the jammer to choose on how many channels it simultaneously attacks. A major addition in finite-energy jamming games is that the jammer and sender both have a limited amount of energy which is drained according to the actions a player takes.

We develop a model of our system as a zero-sum finite-horizon stochastic game with deterministic transitions. We leverage the zero-sum and finite-horizon properties of our model to design a simple polynomial-time algorithm to compute optimal randomized strategies for both players. The utility function of our game model can be decoupled into a recursive equation. Our algorithm exploits this fact to use dynamic programming to construct solutions in a bottom-up fashion. For each state of energy levels, a linear program is solved to find Nash equilibrium strategies for the subgame. With these techniques, our algorithm has only a linear dependence on the number of states, and quadratic dependence on the number of actions, allowing us to solve very large instances.

By computing Nash equilibria for our game models, we explore what kind of performance guarantees can be achieved both for the sender and jammer, when playing against an optimal opponent. We also use the optimal strategies to simulate finite-energy jamming games and provide insights into robust communication among reconfigurable, yet energy-limited, radio systems. To test the performance of the optimal strategies we compare their performance with a random and adaptive strategy. Matching our intuition, the aggressiveness of an attacker is related to how much of a discount is placed on data delay. This results in the defender often choosing to sleep despite the latency implication, because the threat of jamming is high. We also present several other findings from simulations where we vary the strategies for one or both of the players.

Categories and Subject Descriptors

[**Security and privacy**]: Mobile and wireless networks; [**Computing methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence, Multi-agent systems*

Keywords

Jamming; Game Theory; Wireless Networks; Equilibrium Computation; Energy-Constrained Jamming

1. INTRODUCTION

The flexibility of wireless communication enables untethered mobility, agile deployment, and on-the-fly reorganization of connected devices. However, the openness of wireless communication and the recent development of software-defined radio (SDR) technology, respectively, provide a low barrier and a wide range of capabilities for misbehavior and attacks. One class of attacks which has benefited significantly from SDR technology is jamming, or injection of intentionally interfering signals into the wireless medium. While jamming has been a topic of research for several decades [37], partially due to the devastating potential and difficulty of defense, the SDR revolution has sparked continued innovation on jamming and anti-jamming techniques.

Much of the early work in developing jamming models and technologies focused on attackers with unlimited energy resources, postulating that a jamming attacker would use a generator or be connected to the power grid [32]. Such assumptions have led to overly wasteful or boisterous attackers that make no attempt to conserve energy or to hide their attack activity. Likewise, anti-jamming mechanisms are often designed assuming that the jamming attack is trivial to detect, so many techniques reduce to either advanced signal processing [33] or localizing the attack source to take further action [40].

Our increasingly battery-operated mobile world has recently inspired exploration of attackers with limited energy resources [29]. Constrained attackers, however, are not necessarily less effective, as they can leverage the advanced technologies of SDR, software-defined networking, agile and reconfigurable protocols, sensing, and machine learning. Such capabilities can also provide increases in attack stealth, allowing attackers to avoid detection or localization [9]. Ex-

WiSec'14, July 23–25, 2014, Oxford, UK.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2972-9/14/07 ...$15.00.
http://dx.doi.org/10.1145/2627393.2627406 .

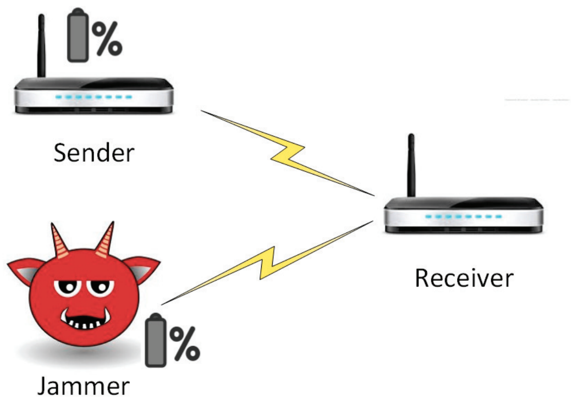

Figure 1: The players involved in our Finite-Energy Jamming Game are the attacker and the defender, made up of the sender-receiver combination. All game decisions on the defender side are actually made by the sender in our model.

amples of recent energy-conscious attacks include periodic jamming [10] and random jamming [40] that alternate between jamming and sleeping to save energy; control channel jamming [8] and similar attacks that leverage protocol structure for efficiency; and reactive jamming [39], adaptive jamming [9, 31], and mesh jamming [22] which respond to observed activity instead of attacking statically.

Fortunately, the same innovative technologies that enable energy-efficient and stealthy attacks can also enable more robust and agile anti-jamming techniques. The agility provided by SDRs allows defenders and attackers alike to have more fine-grained control of protocols and parameters, enabling the ability to adapt on the fly [25]. However, this mutual agility increases system complexity and presents a significant challenge to our understanding of various performance, security, and reliability metrics required for effective system design. Understanding how mutually agile opponents interact in a resource-constrained scenario remains an active research field. In this work we explore a battery-operated jammer and battery-operated sender, where the sender's goal is to successfully transmit and the jammer's goal is to prevent that.

To increase our understanding of mutually agile, resource-constrained players in the context of wireless communications, we look to game theory for tools to analyze optimal strategies for jamming and anti-jamming. We make the following contributions in this work.

- We design a new model for energy-constrained jammer-defender interaction, allowing players to transmit or sleep during any round. This provides for the exploration of a realistic scenario where opponents have similar energy levels and freedom to reconfigure.
- We model this interaction as a zero-sum finite-horizon stochastic game with deterministic transitions to find optimal player strategies. We leverage the properties of our game to design a simple polynomial-time dynamic programming algorithm that solves a series of small linear programs to compute optimal strategies (a Nash equilibrium).
- We implement a simulation of three scenarios to gain insights on the performance of energy-constrained Direct Sequence Spread Spectrum (DSSS) and Frequency Hopping Spread Spectrum (FHSS) systems.

Our first contribution is considering attacking and defending opponents that can choose (1) whether to transmit or sleep, (2) what power level to transmit with, and (3) what channel (or how many channels) to transmit on, with the understanding that each choice has a different energy usage and that the outcome also depends on the action chosen by their opponent. Due to the energy constraints and the fact that every choice has a non-zero cost (even sleeping is subject to non-trivial energy leakage), both players can only participate for a finite amount of time. Moreover, since the value of data may significantly decrease with latency, we allow the sender's utility to decay with time. Since our game incorporates aspects of power control and sleeping for throughput and latency management, we refer to our game formulation as a *finite-energy jamming game*. We are the first to mathematically model and analyze accumulative energy-constrained jammer-sender strategic interaction. In related energy-constrained work, the focus has been on average energy consumption [5] or over-heating [26].

In particular, our finite-energy jamming game formulation imposes maximum energy expenditure on both the jammer and defender, while allowing both players to adjust their transmit power levels. The sender and receiver, collectively comprising the defender in our scenario, communicate either using single-channel DSSS or multi-channel FHSS. We model DSSS and FHSS because many modern communication systems use variants of these techniques. Likewise, the attacker can jam on one or many channels depending on which technique is employed by the defender. If the sender selects a sufficiently high power level compared to the jammer's selected power level (or if the jammer chooses a different channel), then the transmission is successful. Regardless, both player's expend an amount of energy that corresponds to their chosen power level (or sleeping).

In most related work on modeling jamming games, energy constraints have not been considered, so single-shot or repeated game approaches have been adopted. In contrast to that work, the energy constraint means that our game has state, and thus needs more advanced modeling techniques. The two papers in the literature that are closest to our work are the following. First, Altman et. al. [5], consider jamming in a stochastic game setting. Whereas we assume that actions and energy levels are fully observed, their work goes the opposite direction and requires that actions are completely unobserved. The truth lies, of course, somewhere in the middle, but both our and their work can shed light on the possibilities and limitations for the general problem. Second, Mallik et. al. [26] consider a dynamic game where temporal energy constraints exist, in the form of over-heating. This means that energy usage only impacts the immediate rounds afterwards, as opposed to expending energy from a finite supply. Like us, they assume that actions are fully observed. They propose a dynamic game, almost akin to a repeated game, with slight variations in the available actions.

Our second contribution is to develop algorithms for computing optimal strategies for our system, formulating it as a zero-sum finite-horizon stochastic game with deterministic transitions. We use Nash equilibria as our framework for optimal strategies. Nash equilibria are a compelling solution concept especially for zero-sum settings such as ours, as they guarantee the highest utility against optimal opponents and sub-optimal opponents only increase our utility. As such, Nash equilibria and their associated expected utility represent the best guarantee on utility that one can hope for, when faced with potentially optimal adversaries.

We leverage the zero-sum and finite-horizon properties to design a simple polynomial-time dynamic programming algorithm that solves a series of small linear programs to compute a Nash equilibrium. The dynamic programming aspect is similar to the work of Mallik et. al. [26], who also use the finite-horizon aspect to obtain a dynamic programming description. However, they further use specific properties of their setting to derive analytical solutions, whereas our work relies on algorithms for computing strategies. Their consideration of temporal constraints could easily be incorporated into our more general framework and algorithms, along with our finite-resource energy constraints.

Our third contribution is a series of simulations of finite-energy jamming games, which provide insights into robust communication among reconfigurable yet energy-limited radio systems. To further understand the benefit of our game-theoretic models, we compare the rational player using the finite-energy jamming game model with a random player and an adaptive player, demonstrating several cases where the game-theoretic strategies provided by finite-energy jamming game provides significant gains over other strategies. The game theoretic strategies also provides interesting insights about the tradeoffs of energy-constrained jamming-defender interaction. Of particular interest and matching our intuition, we observe that the jammer's optimal strategy is extremely aggressive when the sender highly values low-latency communication, resulting in an attack strategy using high-power jamming in the beginning. This forces the sender to transmit with low probability in the beginning of the game, even when highly valuing low latency. In addition to these observations, we evaluate a number of different attack and defense scenarios, and identify a number of interesting trends and tradeoffs in the realm of finite-energy jamming games. In order to mimic a realistic scenario, we set the jammer and defender's initial energies to be within an order of magnitude of each other for our simulations.

The remainder of this work is organized as follows. In Section 2, we explore related work in jamming and game theory. We introduce our system model and assumptions in Section 3, and we present finite-energy jamming games in Section 4. In Section 5, we present our simulation and evaluation setup, and we discuss our simulation results in Section 6. Lastly, in Section 7 we briefly discuss limitations and future research directions.

2. RELATED WORK

Due to the potential risk of jamming, a large body of work has recently focused on how to effectively avoid and mitigate the effects of jamming attacks. Much of the work on basic and advanced jamming techniques through the last decade has been summarized in a 2010 survey [29]. Efficient jamming and anti-jamming techniques can be classified into two categories: static and adaptive. Static jamming and anti-jamming techniques rely on specification of protocols, parameters, and strategies in advance, while adaptive techniques rely on context, measurements, and observations to choose protocols, parameters, and strategies on the fly to improve performance.

Traditional jamming mitigation techniques have focused on static strategies and shared secrets to perform spread spectrum techniques such as direct sequence spread spectrum (DSSS), frequency hopping spread spectrum (FHSS), code division multiplexing (CDMA), and orthogonal frequency division multiplexing (OFDM) [28]. Efficient static strategies include random [40], periodic [10], and deceptive jamming [40]. Both random and periodic jamming alternate between attacking and sleeping in an attempt to attack in an efficient manner. Deceptive jamming on the other hand sends legitimate packets in an attempt to stealthily interfere with communications, making its effect very similar to greedy MAC misbehavior techniques [30]. More recent strategies have explored adaptation of protocols and parameters at multiple layers either randomly or in response to observations and measurements. The SPREAD system uses multi-layer adaptation as an extension of spread spectrum [25], providing a more robust communication system but still depending on the same secret-sharing fundamentals. Adaptive jamming strategies using observation-based agility [9] and offline optimization using long-term measurement data [36]. Moreover, adaptive anti-jamming techniques have included the use of advanced signal processing and filtering at the receiver [33], jamming-aware traffic management [35], and adaptive beamforming [6].

Game theory has provided a potent tool to investigate and analyze jamming and anti-jamming [1, 12, 27] as well as other security problems. In the domain of jamming, game theory has provided a framework to select parameters and strategies for both static and adaptive jamming and anti-jamming scenarios. We briefly discuss three types of related games: power management games, jammer-versus-defender games, and friendly jamming games.

Power management games study the choice of transmission power levels among nodes in a network to achieve sufficient signal quality while limiting interference with neighbors [4]. Power management games are useful in maximizing the signal-to-interference-and-noise ratio (SINR) of wireless communication in the network. The authors derive a Nash equilibrium for transmission power selection to maximize SINR over the network in both a selfish and cooperative setting.

A second class of relevant games involves explicit competition between jamming and defending players. Previous work has studied the equilibrium behavior of a rate-adaptive defender versus a power-limited jammer [13], choosing jamming power to avoid detection [23], choosing jamming and communication transmission power to balance over-heating concerns [26], choosing jamming strategies considering impact and per-round energy drain [3], and team-versus-team jamming where each team maximizes their own throughput while minimizing the opposing team's throughput [7, 17, 18].

Friendly jamming games aim to use jamming to enforce communication secrecy or privacy against eavesdroppers. In this scenario, utility is defined by the ability to relay data to an intended receiver while preventing eavesdropping by an unintended receiver [16, 19]. Variations on the game include using a coexisting network of active jamming attackers that can also prevent the intended nodes from receiving the data [41].

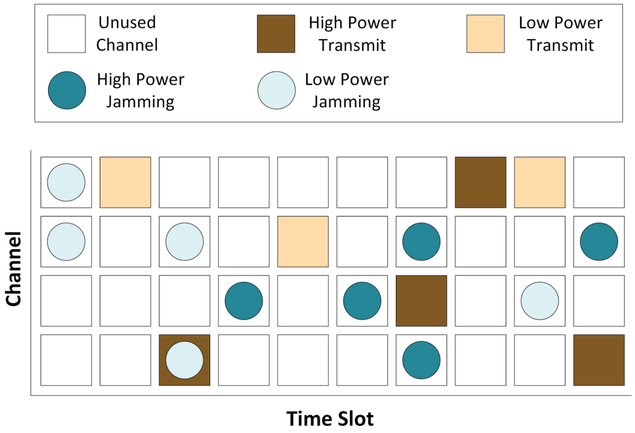

Figure 2: We illustrate our system and show the Finite-Energy Jamming game. The jammer and sender both are able to choose to power nap, transmit at a low power, or transmit at a high power.

Our two-player game with energy-constrained players has similarities with many of these related works, but we include the additional consideration of multi-round optimization with a fixed energy budget for the entire game. The closest of the related works in this regard is the optimal jamming and anti-jamming work of Li et al. [23], but that work differs in that the goal of the attacker is to avoid detection, while in our work the attacker aims to maximally ruin the sender's throughput.

Beyond the papers mentioned in Section 1, much work in the stochastic games literature has been focused around iterative algorithms that eventually converge to a Nash equilibrium or approximate Nash equilibrium, posing additional constraints on the game for convergence, such as existence of global optima, or saddle points [20,24,38]. For the finite-horizon case, polynomial-time algorithms have been developed, with a running time that is quadratic in the size of the state space [21]. Iterative convergence approaches have also been combined with optimal solving of stage games [14,15], but without runtime guarantees. Our context allows us to develop a significantly more efficient algorithm that only requires traversing the state space once.

3. SYSTEM MODEL AND ASSUMPTIONS

In this work, we explore a three-node scenario consisting of a sender, a receiver, and a jamming attacker over a time interval $\mathcal{T}$, as illustrated in Figure 1. The sender and receiver collectively comprise the defender, able to use single-channel DSSS or N-channel FHSS, while all of the defender's decisions in our scenario are made by the sender. We assume that both the attacker and defender are energy constrained, starting with initial energy $E_{a,0}$ and $E_{d,0}$, respectively, so they are forced to balance between maximum performance and minimum energy expenditure.

We assume that the time interval $\mathcal{T}$ is divided into distinct sub-intervals referred to as *rounds*. In each round, the defender chooses a transmission power p_d from a discrete set of power levels $\mathcal{P}_d \subseteq \{0, 1, \ldots, p_{d,\max}\}$, where $p_{d,\max}$ is the defender's maximum transmission power. When the defender transmits with power p_d in a round, it incurs an energy cost $\epsilon_d(p_d)$, and we assume two fundamental properties of this cost function: monotonicity and strict positivity. Monotonicity of the cost function simply means that higher transmission power incurs higher energy cost, while strict positivity means that all actions incur an energy cost, even a play of $p_d = 0$, in which case the defender pays a leakage cost while sleeping. In the case of FHSS, the defender also chooses which of the N channels it will use for communication. We assume that the underlying synchronization, configuration, and channel switching costs are negligible, and thus we treat them as free. We illustrate this system in Figure 2.

The attacker's energy model is similar. In each round, the attacker chooses a jamming power p_a from a discrete set of power levels $\mathcal{P}_a \subseteq \{0, 1, \ldots, p_{a,\max}\}$, where $p_{a,\max}$ is the attacker's maximum jamming power. The energy cost of the attacker's action in the round is also dictated by a function $\epsilon_a(p_a)$, which is both monotonic and strictly positive as with the defender, but with one significant difference. In the FHSS case, the attacker is allowed to reconfigure the radio front-end to jam any k out of the N channels, where $1 \leq k \leq N$, using power p_a per channel. The cost for the attacker for this round is then given by $k\epsilon_a(p_a)$ proportional to the number of channels jammed.

To measure the defender's performance in each round, we consider the throughput $T(p_d, p_a)$ achieved in the round when the defender transmits with power p_d and the attacker transmits on the same channel with power p_a. Moreover, since the value of the sender's data to the receiver may decrease with time, we introduce a discount factor $\delta \in (0, 1]$ for every round in which the data does not reach the receiver. This discount function is representative of a system where the sender has all data at the beginning of the time interval and desires rapid transmission. To compensate for the latency induced by the jamming attack in the defender's utility function, we multiply the throughput $T(p_d, p_a)$ by δ for each round of delay, so any throughput attained during round i is valued according to the *latency-adjusted throughput* $\delta^i T(p_d, p_a)$. When the attacker jams on a channel different from that used by the defender, we use the equivalent throughput $T(p_d, 0)$, since the attack has no effect.

We assume a perfect knowledge scenario starting at the beginning of the time interval, so the players know the initial energy of both players. In addition, we assume that each player can observe their opponent's actions in that round by the end of the round, so each player always has complete knowledge of their opponent's residual energy at the beginning of the next round when they have to decide what to do in that round.

4. FINITE-ENERGY JAMMING GAMES

We model a finite-energy jamming game in the described wireless system as a zero-sum finite-horizon stochastic game with deterministic transitions. In much of the related literature, jamming scenarios have been modeled as single-shot or repeated games [2,34]. Both of these approaches are sensible when no state is present, for example, when energy constraints do not apply, as the same strategy remains optimal throughout the game. Since our game has state in the form of residual energy, neither single-shot nor repeated game models can adequately capture our setting. Instead, we turn to stochastic games, where the game state transitions at every time step. Since residual energy is monotonically decreasing, we can model our game as a finite-horizon game. We allow all transitions between states to be deterministic, since the energy cost of different actions is assumed to be fixed and known.

In two-player zero-sum games, the solution concept of Nash equilibria is particularly compelling. In general-sum games, there can be many Nash equilibria with different expected utilities for the players, and playing a Nash equilibrium strategy says nothing about the expected utility for a player, if the opponent does not play a best response. This is not so for zero-sum games, where playing a Nash equilibrium strategy guarantees at least a certain level of utility in expectation. That guaranteed utility is called the value of the game, and it is achieved when the opponent responds optimally to the player's optimal (Nash equilibrium) strategy. The zero-sum property guarantees that the player can only benefit (and get more than the value of the game) if the opponent does not play optimally. We will formally define a Nash equilibrium in Section 4.1.

We first introduce our game framework in the context of the single-channel communication system. We then extend our study to include FHSS with the jammer transmitting on a fixed number of channels. After this, we explore a further extension using FHSS where the attacker can vary the number of jammed channels at each time step. For the FHSS settings, we are not assuming that a single channel is chosen at each time step. Rather, we assume that the frequency hopping is so effective that the best the jammer can do is jam a random subset of channels in the hopes of disrupting communication. Finally, we show how to compute Nash equilibria for these games.

4.1 Single-Channel Game

The first game we explore is the single-channel finite-energy jamming game. This game uses a single DSSS channel and has the attacker and defender select power levels from a discrete set. The parameters to the game are the discount factor δ and the initial energies $E_{d,0}, E_{a,0}$ for the players. Based on the defender's residual energy E_d at the start of a round, the defender's action set $A_d(E_d)$ for that round is defined as $A_d(E_d) = \{p \in \mathcal{P}_d : \epsilon_d(p) \leq E_d\}$. The attacker's action set $A_a(E_a)$ is similarly defined. When the defender and attacker choose respective actions $p_d \in A_d(E_d)$ and $p_a \in A_a(E_a)$, the immediate utility to the defender is $u_d(p_d, p_a) = T(p_d, p_a)$, which is later discounted by δ^i during round i to compensate for latency. The attacker's immediate utility is $u_a(p_d, p_a) = -u_d(p_d, p_a)$. The defender chooses its action based on an energy-dependent *strategy* $\sigma_d^{E_d,E_a}$ that specifies a probability distribution over actions in $A_d(E_d)$. For example, $\sigma_d^{E_d,E_a}(p)$ is the probability the defender will transmit at power level $p \in \mathcal{P}_d$. We analogously define the attacker's strategy $\sigma_a^{E_d,E_a}$. Once the players choose their actions in a round i, with the defender and attacker respectively transmitting at power levels p_d and p_a, the game transitions to round $i+1$, where the players have residual energy $E_d - \epsilon_d(p_d)$ and $E_a - \epsilon_a(p_a)$. The game continues in this way until the defender's residual energy is such that $A_d(E_d) \subseteq \{0\}$, after which $u_d = u_a = 0$.

Considering the entire game over multiple rounds, a *strategy profile* σ is a pair of strategies $\sigma = \{\sigma_d, \sigma_a\}$ that fully specifies the game. Using the strategy profile σ, we can then compute the defender's total expected utility $u^\sigma(E_{d,0}, E_{a,0})$ using a recursive definition over diminishing energy levels as

$$u^\sigma(E_d, E_a) = \sum_{p_d \in A_d(E_d)} \sum_{p_a \in A_a(E_a)} \sigma_d^{E_d,E_a}(p_d)\sigma_a^{E_d,E_a}(p_a) \times \Big(u_d(p_d,p_a) + \delta u^\sigma(E_d - \epsilon_d(p_d), E_a - \epsilon_a(p_a))\Big) \tag{1}$$

where $u_d(p_d, p_a) = T(p_d, p_a)$ for the single-channel game. This can be viewed as a series of normal-form games, where the payoff matrix for each game depends on the values of the subgames induced by the various choices of actions.

A Nash equilibrium is a strategy profile $\sigma^* = \{\sigma_d^*, \sigma_a^*\}$ that satisfies

$$\sigma_d^* = \arg\max_{\sigma_d} u^{\{\sigma_d, \sigma_a^*\}}(E_{d,0}, E_{a,0})$$
$$\sigma_a^* = \arg\max_{\sigma_a} u^{\{\sigma_d^*, \sigma_a\}}(E_{d,0}, E_{a,0})$$

In other words, in a Nash equilibrium, each player maximizes their own utility, given the strategy of the other player.

4.2 Multi-Channel Game with FHSS

We next consider a finite-energy jamming game in which the defender spreads its transmissions over N orthogonal channels by choosing a different channel randomly in each round of the game. In our first FHSS-based game, the attacker chooses k channels to jam every round, where k is constant for the duration of the game. In each round, the attacker has a probability of k/N of interfering with the defender's transmission, so the immediate utility for the defender in this case is given by

$$u_d(p_d, p_a) = \frac{k}{N}T(p_d, p_a) + \frac{N-k}{N}T(p_d, 0) \tag{2}$$

The defender's total expected utility is given by substituting (2) into (1). In this game, the energy expenditure of the attacker is increased by a factor of k, meaning that an attack action with power p_a incurs a cost $k\epsilon_a(p_a)$.

4.3 Multi-Channel Game with FHSS and Selection of Number of Channels to Jam

Similar to our second game, we consider a generalization of the previous FHSS game in an N-channel communication system. In our second FHSS-based game, the attacker is free to choose any value of $k \in \{1, \ldots, N\}$ in each round as part of its attack strategy. Given the additional game parameter, the attacker's action set $A_a(E_a)$ in each round is extended to

$$A_a(E_a) = \{(p, k) \in \mathcal{P}_a \times \{1, \ldots, N\} : k\epsilon_a(p) \leq E_a\}$$

and the utility function $u_d(p_d, p_a)$ is extended to $u_d(p_d, p_a, k)$, using the same form as (2). In contrast to the previous game with fixed k, treating k as a variable game parameter allows the attacker to effectively balance the tradeoff between higher utility and greater energy expenditure of jamming more channels. In addition, since the attacker's action set $A_a(E_a)$ has increased in dimensionality compared to the fixed-k case, the complexity of solving the game increases linearly in N.

4.4 Computing a Nash Equilibrium

For each of the three game models described above, we can use the same basic approach for computing a Nash equilibrium. Each of those three games can be viewed as a series of normal-form games, each of which depend on the values of subgames to fill out their payoff matrix. We use this subgame property, along with the well-known fact that zero-sum normal-form games can be solved in polynomial time using linear programming, to solve our problem. Using dynamic programming, solutions are constructed bottom up through successively solving linear programs that compute Nash equilibria of subgames. The pseudocode is presented as Algorithm 1.

```
**Input**: Energy levels $E_d, E_a$, discount factor $\delta$
**Output**: Nash equilibrium strategy profile $\sigma$
$U \leftarrow [\,]$        // dynamic programming table
**for** $E'_d \in \{0, \ldots, E_d\}$ **do**
    **for** $E'_a \in \{0, \ldots, E_a\}$ **do**
        $M \leftarrow [\,]$        // payoff matrix
        **for** $p_d \in A_d(E'_d), p_a \in A_a(E'_a)$ **do**
            $M[p_d, p_a] =$
            $u(p_d, p_a) + \delta \cdot U[E'_d - \epsilon_d(p_d), E'_a - \epsilon_a(p_a)]$
        **end**
        $U[E'_d, E'_a] = \textsc{GameValue}(M)$
        $\sigma^{E'_d, E'_a} = \textsc{StrategyProfile}(M)$
    **end**
**end**
```

Algorithm 1: Bottom-up dynamic program for computing Nash equilibria in finite-energy jamming games.

The dynamic program iterates over all possible energy levels for the two players, starting from the smallest levels possible. For each pair of energy levels, a payoff matrix M is computed. Line 1 implements the recursive equation for utility given in (1) or (2) depending on the game played. That is, it sets the payoff to the immediate payoff achieved from the actions taken plus the value of the subgame reached by the power loss, weighted by the discount factor δ. Lines 1 and 1 extract the value of the game and a strategy profile that achieves a Nash equilibrium.

Our dynamic program crucially relies on the fact that every set of energies E_d, E_a induces a subgame, where the path traveled to get to these energy levels does not matter. Depending on the round where the energy levels are reached, the discount factor might be different. However, in terms of computing a strategy for E_d, E_a, we can assume without loss of generality that we are at round 0, since for any other round i, every entry in M will be scaled by the same discount factor δ^i, and so the optimal strategies will be the same.

For the function calls GameValue and StrategyProfile in Lines 1 and 1, a solver for computing a Nash equilibrium of M is needed. Since M is a standard payoff matrix for a normal-form game (entries are constants, because values for the subgames have already been computed), we can adopt the standard linear programming approach for computing a Nash equilibrium strategy. We will show how to compute a Nash equilibrium strategy for the defender, with the case for the attacker being completely analogous.

The linear program is shown in Figure 3. The variable v denotes the utility for the defender, which is to be maximized. The first two constraints ensure that the defender's strategy at the subgame forms a probability distribution. The last constraint ensures that no matter which action the attacker selects, the defender is guaranteed value v. For any optimal solution, the value of v will be the value of the game, and the computed strategy will be a Nash equilibrium strategy.

$$\max v \tag{3}$$

$$\sum_{p_d \in A_d(E_d)} \sigma_d^{E_d,E_a}(p_d) = 1 \tag{4}$$

$$\sigma_d^{E_d,E_a}(p_d) \geq 0 \qquad \forall p_d \in A_d(E_d) \tag{5}$$

$$\sum_{p_d \in A_d(E_d)} \sigma_d^{E_d,E_a}(p_d) \cdot M[p_d, p_a] \geq v \quad \forall p_a \in A_a(E_a) \tag{6}$$

Figure 3: The linear program used in computing a Nash equilibrium strategy for the defender. $A_d(E_d)$ and $A_a(E_a)$ are the sets of actions available for the defender and attacker respectively, given their current energy levels.

The number of linear programs can be upper-bounded by the number of possible energy levels in subgames. Given initial energy levels $E_{d,0}$ and $E_{a,0}$, the number of linear programs solved is $O(\frac{E_{d,0}}{\epsilon_d(0)} \cdot \frac{E_{a,0}}{\epsilon_a(0)})$, since the power cost of sleeping divides all other power costs. Each linear program has size $O(|\mathcal{P}_d| \cdot |\mathcal{P}_a|)$.

Technically, our algorithm computes a subgame perfect equilibrium, a refinement of Nash equilibria. A subgame perfect equilibrium is a Nash equilibrium such that for any subgame, even those reached with probability zero, the players are playing Nash equilibrium strategies for the subgame. This provides an extra level of robustness over Nash equilibria, as we are not only guaranteed the value of the game, but also guaranteed to play optimally if the opponent chooses a sub-optimal action, assuming the game is played optimally onwards from there. This is not the same as optimally responding to any strategy of the opponent. Rather, it means that we optimally respond to any current game state, assuming that the opponent will play optimally from then on, even with mistakes in the past.

5. SIMULATION

To show the benefits of the finite-energy jamming game we simulate three different games. To use realistic parameters in the simulation we base our parameters on measurement data taken from communication nodes and a jammer implemented with GNUradio on USRP2 software-defined radio. We consider an attacker that is able to adapt their power level and also the number of channels they jam on. The defender is able to choose a channel to transmit on and also choose a power level to transmit at. For our measurements the sender and jammer are connected to the receiver via wire with equal attenuation. This mimics the location of the senders being equidistant from the receiver. In this section, we discuss the parameters we use for our simulation, the optimization results, and the game play simulation we use.

5.1 Game parameters

We take RF power measurements at the connection port and find that the power expended for a low-power attack as $1.16\mu w$ and for a high-power attack as $3.22\mu w$. We also

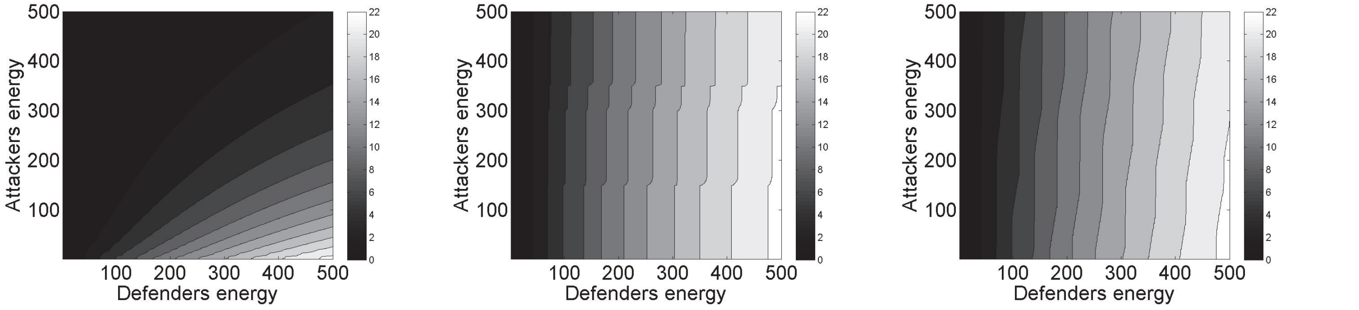

(a) Single channel

(b) Fixed FHSS

(c) Optimal FHSS

Figure 4: To demonstrate the optimization, we show the expected utility for 3 different games with varying initial energy levels. For all the games a discount factor of .975 is used and in the frequency hopping game the defender uses 50 channels. The color scale shows the utility of the game.

assume a continuous energy drain per round that we estimate as $.5\mu w$. This constant drain controls for calculation, battery leakage, and other constant sources of drainage. We normalize the cost of energy usage per round and define $\{1, 3, 7\} \in \mathcal{P}_a$ as the values for sleeping, low power, and high power attacking, respectively. The jammer is able to simultaneously jam on multiple channels during any round. We assume a linear cost increase per channel for the low- or high-power attacks. Sleeping does not use channels so we assume it has no increased costs.

Likewise for the defender we find power at the port for a low-power transmission as $6.5\mu w$ and a high-power transmission as $7.83\mu w$. We again assume a constant energy drain of $.5\mu w$. Normalizing and approximating the cost per round of each play we find costs of $\{1, 14, 16\} \in \mathcal{P}_d$ for sleeping, low power transmissions, and high power transmissions, respectively. We assume that synchronization and key-sharing is done beforehand and that there is no extra cost for the sender to use frequency hopping.

If the defender is transmitting we assume a constant rate so the normalized throughput per round is approximated by packet delivery ratio (PDR). Because of this we use packet delivery ratio in lieu of throughput when the defender is transmitting and assume zero throughput when the defender is sleeping. We measure packet delivery ratio in our single channel 802.15.4 system as

$$\mathrm{pdr}(\mathcal{P}_a, \mathcal{P}_d) = \begin{pmatrix} 0 & .96 & 1 \\ 0 & .58 & .92 \\ 0 & 0 & 0 \end{pmatrix} \qquad (7)$$

where the attacker is the row player and the defender is the column player. We assume that there is no cross-channel interference so if the jammer is not attacking a particular channel there is no added interference.

In order to mimic both players using the same class of devices, we constrain the attacker and defender to have similar initial energy resources. We define similar initial energy resources as both players having an initial energy that is within one order of magnitude of the other.

5.2 Optimization

We use Algorithm 1 to arrive at an optimal strategy and expected utility. In Figure 4(a) we show the values for the single channel game with a .975 discount factor. In Fig-

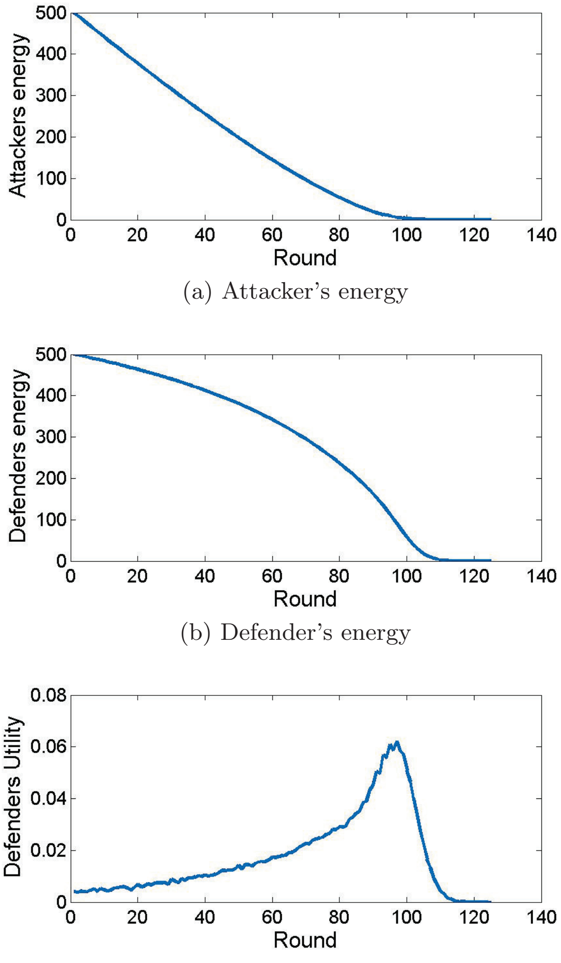

(a) Attacker's energy

(b) Defender's energy

(c) Defender's utility

Figure 5: The average over 10,000 runs of a simulation of the single channel game with two rational players and a .975 discount factor.

ure 4(b) we show the expected defender utility for a defender with 50 channels and an attacker with 50 channels and .975 discount. In Figure 4(c) we show the defender's utility when the defender has 50 channels and the attacker optimizes power and number of channels. The optimization provides confirmation of what is intuitively expected. The single channel game heavily favors the attacker while either of the frequency hopping games with 50 channels favors the defender.

5.3 Game play

We designed a simulator to explore the performance of our computed strategies and compare them to other strategies. Other strategies we use for comparison include a constant strategy, a uniform random strategy, and a weighted average algorithm [11]. The random strategy that we consider uniformly samples from all possible strategies. The weighted average algorithm was designed for a similar power game. It works by keeping a weighted vector of the likelihood of their opponents strategy as well as a matrix of the expected utility for given combinations of plays. The player then uses these to compute their strategy.

We designed a simulator for each of the games introduced in Section 4. For the single channel, input parameters include both players' strategies as well as the initial energy of both players, and the discount factor for the players. In the frequency hopping spread spectrum case with a constant number of attacker channels the simulator also takes the number of channels used by the defender N and attacker k. The simulator also accepts the precomputed optimal strategy for both players for the given game and the discount factor.

To demonstrate the operation of our simulator we show the average run of 10,000 trials of the single channel game with a .975 discount factor and two rational players in Figure 5. The initial energy for assigned to both player is 500 units in this experiment, and the power levels and corresponding energy usage are given in Section 5.1. Figures 5(a) and 5(b) show the average remaining energy for the attacker and defender, respectively, at the given time. Figure 5(c) on the other hand shows the average instantaneous utility for the defender.

To simulate the frequency hopping game the defender selects one channel $n \in [1, N]$ at the beginning of every round. Similarly, the attacker selects k of N channels to interfere with. If n is one of the k channels selected then the attacker is successful and the throughput is calculated using (7). Otherwise the throughput is calculated using

$$\text{pdr}(\mathcal{P}_a, \mathcal{P}_d) = \begin{pmatrix} 0 & .96 & 1 \\ 0 & .96 & 1 \\ 0 & .96 & 1 \end{pmatrix} \qquad (8)$$

6. SIMULATION RESULTS

In this section, we discuss simulated scenarios using the setup and parameters presented in Section 5. We explore all three games from Section 4 and describe insights gained from various experiments.

6.1 Single-Channel Game

For the single channel game we compare the performance of the rational, random, and adaptive weighted average algorithm for both players. In Table 1 the utility averaged over 100,000 runs for various attacker and defender strategy pairs is presented. Both players start with 200 units of energy and choose optimal strategies for the given discount factor. The results for the rational player always outperform the random and weighted average player's performance, but sometimes this is less pronounced. Although the gain from rationality is marginal with the .9 discount factor, a second factor to consider is that rationality decreases deviation of results. In Table 2 we show that either player playing rationally greatly decreases the standard deviation in utility. This decrease in variance can be a significant benefit for designing secure systems in that it is able to provide performance guarantees and less uncertainty.

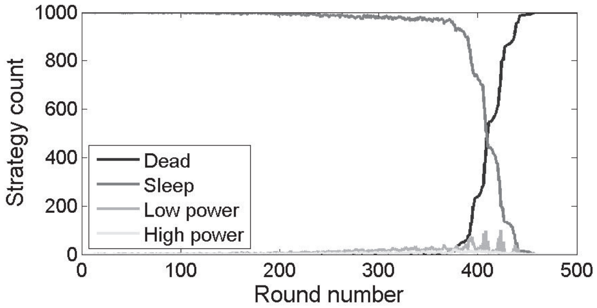

Figure 6: Counts of how many times out of a thousand a defender chooses a strategy against a constantly sleeping attacker with a .9 discount factor. We define a dead node as a node that has expended all of its energy.

The single channel game also provides an interesting insight on the effect of rationality on the defender's utility. In Table 3 we see that rational play increases the defender's overall throughput. The smaller the discount factor, the greater the gain in defender's throughput from rationality against a rational attacker.

Another interesting result is highlighted in Figure 6. In this figure the attacker always chooses to power nap while the defender is rational. This results in an attacker that has a slow but constant energy fade. The rational defender plays as if the attacker was also rational, and therefore he transmits with very low probability in the beginning of the game. This is highly counterintuitive from a throughput perspective, since the sender could transmit freely, and gain much higher utility. This is an example of how inoptimal opponents are not exploited optimally by a Nash equilibrium strategy, since the sender has to assume that the jammer might start playing optimally at each round, in order to guarantee attaining the value of the game.

We also explored the effect of a difference in energy between the two players. In Figure 7 we show the defender's utility for various advantages in the attacker's energy. The curve here, while qualitatively intuitive, can be instructive in how much extra energy a defender must have to perform well in the presence of an attacker.

6.2 Multi-Channel Game with FHSS

The second set of experiments we conduct considers a defender frequency hopping over a set of N channels and a jammer blocking a set of K channels per round. In Figure 8 we show the defender's utility for various sets of attacker

		Defense								
		.9 Discount Factor			.95 Discount Factor			.99 Discount Factor		
		Rational	Random	Weighted	Rational	Random	Weighted	Rational	Random	Weighted
Attack	Rational	0.2575	0.2573	0.2564	1.2424	1.2351	1.2369	6.0236	6.0031	6.0119
	Random	0.2583	3.2958	3.1432	1.2452	4.7213	4.5393	6.08	7.0388	6.8867
	Weighted	0.2577	2.5974	2.5286	1.25	3.545	3.4438	6.425	5.0245	4.9568

Table 1: Mean defender's utility for the single channel game.

		Defense								
		.9 Discount Factor			.95 Discount Factor			.99 Discount Factor		
		Rational	Random	Weighted	Rational	Random	Weighted	Rational	Random	Weighted
Attack	Rational	0.0956	0.3124	0.2942	0.3682	0.7693	0.7255	1.0194	1.3730	1.3590
	Random	0.2651	0.9328	0.9878	0.5104	1.1131	1.1843	1.0983	1.4895	1.5331
	Weighted	0.2554	1.3741	1.4083	0.5049	1.9931	2.0147	1.1843	3.1574	3.1580

Table 2: Standard deviation of the defender's utility for the single channel game.

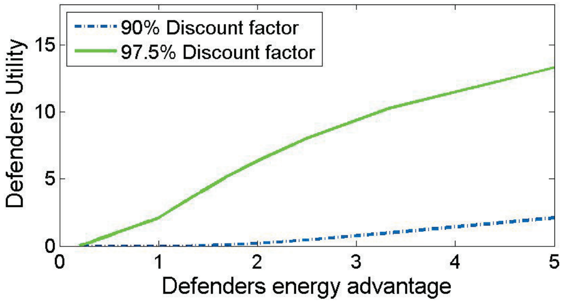

Figure 7: Advantage gained by an attacker or defender having a energy advantage with varying discount factors. The advantage shown is the multiplicative advantage such that defender's advantage $= \frac{E_d}{E_a}$.

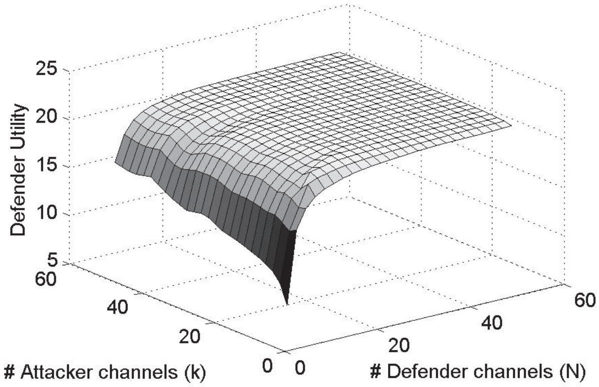

Figure 8: Defender's utility for the set number of attacker channel FHSS game. The attacker and defender both choose their power levels optimally for the number of channels they are using.

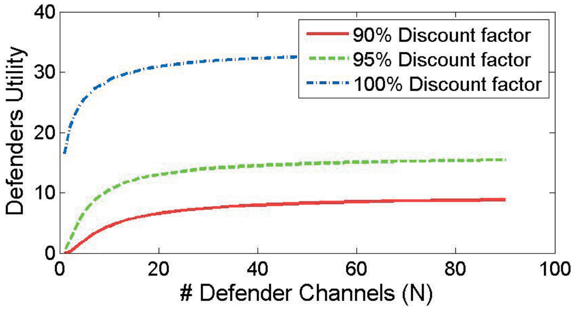

Figure 9: In this figure, we show the mean defender's utility for varying numbers of defending channels and an optimal attacker.

and defenders channel numbers with a .975 discount factor when both players are rational. This leads to two conclusions when the defender and attacker have similar initial energy. First, a defender with 20 or more channels effectively mitigates the jamming threat. Second, an attacker jamming fewer channels in this case can be beneficial to the attacker. One likely explanation for this would be sensitivity to power cost, since only being able to jam a large number of channels (as opposed to being able to vary this) expends a large portion of the energy.

6.3 Multi-Channel Game with FHSS and Selection of Number of Channels to Jam

Our third set of experiments considers the FHSS game where the attacker can choose power levels and the number of simultaneous channels to attack. In Figure 9 we show the mean defender's utility for various discount factors. This figure suggests that above a certain number of channels, even with an optimal attacker, there is a diminishing return on investment for the defender adding more channels. The attacker's strategy with the lower discount factor causes an attacker to select a very aggressive strategy, often expending all its energy as quickly as possible in hopes of causing some degradation to the transmission.

In Figure 10 we show the defender's utility for various attacker multiplicative energy advantages defined as $\frac{E_a}{E_d}$.

		Defense								
		.9 Discount Factor			.95 Discount Factor			.99 Discount Factor		
		Rational	Random	Weighted	Rational	Random	Weighted	Rational	Random	Weighted
Attack	Rational	7.0159	0.7682	1.0190	6.4976	1.9530	2.1542	6.6333	6.1134	6.1242
	Random	6.8493	7.2073	7.0517	6.6224	7.2075	7.0662	6.6791	7.1996	7.0542
	Weighted	6.2695	5.1466	5.0629	6.4519	5.1496	5.0321	7.0621	5.1268	5.0613

Table 3: Defender's mean throughput for the single channel game.

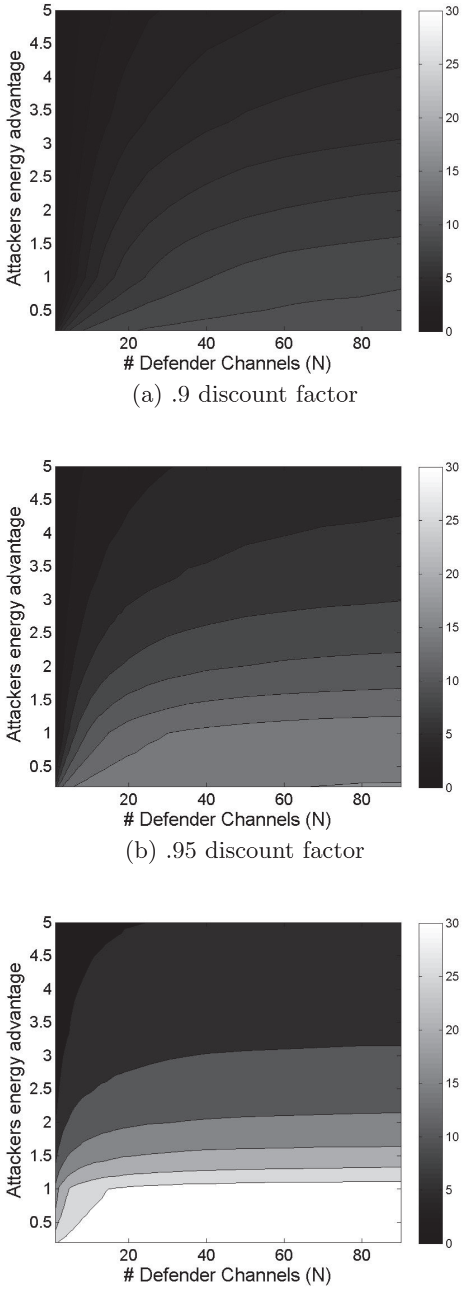

(a) .9 discount factor

(b) .95 discount factor

(c) .99 discount factor

Figure 10: Defender's utility for various channel numbers against an optimal attacker. We define the attacker's energy advantage as $\frac{E_a}{E_d}$.

These curves allow for a decision of how much of an energy advantage an attacker needs to overcome spread spectrum. This also illustrates that the number of channels a defender needs to be protected from a jamming attack varies on the difference in the two players energy.

6.4 Summary of Simulation Results

In this work, we consider three different finite-energy jamming games. The first is a single channel DSSS jamming game, the second is a FHSS game where the attacker jams a constant number of channels, and the third is a power nap game with jammer attacker the optimal number of channels.

In the first game we find that either player playing rational decreases the variance in the game, a beneficial result for designing a secure communication system. We also noted that in this game a rational defender greatly increased the overall throughput of the system. We also showed that rationality can be detrimental to the defender. When the discount factor is small and the attacker chooses a strategy of constantly sleeping the defender is intimidated into not transmitting until most of the energy is drained.

In the second game we show that the a defender that hops over at least 20 channels is effectively able to mitigate the effects of jamming. We also show, counter to intuition, in some scenarios when the attacker jams less simultaneous channels it has a greater impact.

In the third game we confirm the intuition of a diminishing return for the defender past a certain number of channels. We also show the tradeoff in the advantage in the attacker's energy level and the number of channels. These charts provide a basis for choosing the number of channels a defender needs for hopping based on how much extra energy is available for the jammer.

7. CONCLUSION

In this work we introduced *finite-energy jamming games*, a game-theoretic framework to understand energy-constrained jammer-defender interaction. We developed several game models within this framework, where the sender and jammer can vary their power levels and whether to send at all. In our more advanced models, we introduced frequency hopping to the game model, and investigated the effect of allowing the jammer to vary number of channels jammed on. To do this, we modeled our system as a zero-sum finite-horizon stochastic games with deterministic transitions. Leveraging the properties of our game, we designed a simple and fast polynomial-time dynamic programming algorithm for computing a Nash equilibrium. We implemented a simulator to explore the practical properties of our framework across our different game types. Using our simulator, we investigated the possible guarantees that can be achieved under various game settings. We also investigated the practical performance of Nash equilibrium strategies against simpler strategies, such as adaptive or fixed randomized strategies. An interesting result from this analysis was the decrease in variance provided from a rational player, a beneficial property for designing secure systems. Another interesting result provided by this analysis is that an inoptimal opponent that sleeps constanly still leads to a rational sender incurring large performance losses, due to the assumption that the attacker will play optimally.

There are several interesting future research directions for extending our current work. First, to make the problem more practical, it would be interesting to relax the perfect knowledge assumption and replace it with an observation based approach. This would make the game model significantly harder to solve, and so more advanced computational approached would be needed. Second, expanding this work to the setting of multiple jammers and multiple defenders would provide a better understanding of interactions of adversaries in the wild. Third, the scope of both players could be expanded to include multiple layers in the communication stack and cross-layer attacks. Finally, there are several options for expanding the action space of the players, in ways that are easily incorporated in our current model and algorithms. We currently allow the players to select a single power level per round, even when the number of channels jammed on is more than one. There could be cases where the sender and jammer would want to select different power levels for different channels. This would incur an exponential increase in the number of actions available to the players, but our algorithmic results would transfer to such a setting. Related to this, it might be possible to show that certain combinations of power levels over different channels are never optimal, in order to avoid this blowup. Likewise, in this paper we only allowed the sender to transmit on a single channel. In future work, it would be interesting to investigate whether sending on several channels at once is beneficial.

8. ACKNOWLEDGMENTS

DeBruhl and Tague are supported by the NSF under grant CNS-1149582, and DeBruhl is further supported by an NDSEG fellowship. Sandholm and Kroer are supported by the NSF under grants IIS-1320620, CCF-1101668 and IIS-0964579, as well as the CMU Center for Computational Thinking funded by Microsoft Research. Datta is supported in part by the AFOSR MURI on Science of Cybersecurity. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of CMU, NSF, or the U.S. Government or any of its agencies.

9. REFERENCES

[1] T. Alpcan and T. Başar. *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press, 2010.

[2] E. Altman, K. Avrachenkov, and A. Garnaev. Jamming in wireless networks: The case of several jammers. In *IEEE International Conference on Game Theory for Networks, 2009*, pages 585–592.

[3] E. Altman, K. Avrachenkov, and A. Garnaev. A jamming game in wireless networks with transmission cost. In *Network Control and Optimization*, pages 1–12. Springer, 2007.

[4] E. Altman, K. Avrachenkov, and A. Garnaev. Transmission power control game with SINR as objective function. In *Network Control and Optimization*, pages 112–120. Springer, 2009.

[5] E. Altman, K. Avrachenkov, R. Marquez, and G. Miller. Zero-sum constrained stochastic games with independent state processes. *Mathematical Methods of Operations Research*, 62(3):375–386, 2005.

[6] J. Becker, J. D. Lohn, and D. Linden. An in-situ optimized anti-jamming beamformer for mobile signals. In *IEEE Antennas and Propagation Society International Symposium 2012*, pages 1–2.

[7] S. Bhattacharya, A. Khanafer, and T. Başar. Switching behavior in optimal communication strategies for team jamming games under resource constraints. In *IEEE International Conference on Control Applications 2011*, pages 1232–1237.

[8] A. Chan, X. Liu, G. Noubir, and B. Thapa. Control channel jamming: Resilience and identification of traitors. In *Proc. IEEE International Symposium on Information Theory*, Nice, France, June 2007.

[9] B. DeBruhl, Y. Kim, Z. Weinberg, and P. Tague. STIR-ing the wireless ether with self-tuned, inference-based, real-time jamming. In *Proc. 9th IEEE Conference on Mobile Ad-hoc and Sensor Systems*, Las Vegas, USA, Oct. 2012. IEEE.

[10] B. DeBruhl and P. Tague. How to jam without getting caught: Analysis and empirical study of stealthy periodic jamming. In *IEEE International Conference on Sensing, Communication, and Networking, 2013*.

[11] B. DeBruhl and P. Tague. Keeping up with the jammers: Observe-and-adapt algorithms for studying mutually adaptive opponents. *Pervasive and Mobile Computing*, 2014.

[12] M. Felegyhazi and J. Hubaux. Game theory in wireless networks: A tutorial. Technical report, Technical Report LCA-REPORT-2006-002, EPFL, 2006.

[13] K. Firouzbakht, G. Noubir, and M. Salehi. On the capacity of rate-adaptive packetized wireless communication links under jamming. In *Proceedings of the fifth ACM conference on Security and Privacy in Wireless and Mobile Networks*, pages 3–14. ACM, 2012.

[14] S. Ganzfried and T. Sandholm. Computing an approximate jam/fold equilibrium for 3-player no-limit Texas Hold'em tournaments. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2008.

[15] S. Ganzfried and T. Sandholm. Computing equilibria in multiplayer stochastic games of imperfect information. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 2009.

[16] A. Garnaev and W. Trappe. An eavesdropping game with SINR as an objective function. In *Security and Privacy in Communication Networks*, pages 142–162. Springer, 2009.

[17] A. Gupta, A. Nayyar, C. Langbort, and T. Başar. A dynamic transmitter-jammer game with asymmetric information. In *51st IEEE Annual Conference on Decision and Control, 2012*, pages 6477–6482.

[18] Y. Gwon, S. Dastangoo, C. Fossa, and H. Kung. Competing mobile network game: Embracing antijamming and jamming strategies with reinforcement learning. In *IEEE Conference on Communications and Network Security*, pages 28–36. IEEE, 2013.

[19] Z. Han, N. Marina, M. Debbah, and A. Hjorungnes. Physical layer security game: How to date a girl with her boyfriend on the same table. In *IEEE International Conference on Game Theory for Networks, 2009*, pages 287–294.

[20] J. Hu and M. P. Wellman. Nash Q-learning for general-sum stochastic games. *The Journal of Machine Learning Research*, 4:1039–1069, 2003.

[21] M. Kearns, Y. Mansour, and S. Singh. Fast planning in stochastic games. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, pages 309–316. Morgan Kaufmann Publishers Inc., 2000.

[22] L. Lazos and M. Krunz. Selective jamming dropping insider attacks in wireless mesh networks. *IEEE Network*, 25(1):30–34, 2011.

[23] M. Li, I. Koutsopoulos, and R. Poovendran. Optimal jamming attacks and network defense policies in wireless sensor networks. In *IEEE 26th IEEE International Conference on Computer Communications, 2007*, pages 1307–1315.

[24] M. L. Littman. Friend-or-foe Q-learning in general-sum games. In *ICML*, volume 1, pages 322–328, 2001.

[25] X. Liu, G. Noubir, R. Sundaram, and S. Tan. SPREAD: Foiling smart jammers using multi-layer agility. In *26th IEEE International Conference on Computer Communications*, Anchorage, AK, USA, May 2007.

[26] R. K. Mallik, R. A. Scholtz, and G. P. Papavassilopoulos. Analysis of an on-off jamming situation as a dynamic game. *Communications, IEEE Transactions on*, 48(8):1360–1373, 2000.

[27] M. Manshaei, Q. Zhu, T. Alpcan, T. Başar, and J.-P. Hubaux. Game theory meets network security and privacy. *ACM transaction on Computational Logic*, 5, 2011.

[28] A. Molisch. *Wireless Communications*. John Wiley & Sons, Inc., 2005.

[29] K. Pelechrinis, M. Iliofotou, and S. Krishnamurthy. Denial of service attacks in wireless networks: the case of jammers. *IEEE Comm Surveys and Tutorials*, Apr. 2011.

[30] M. Raya, I. Aad, J.-P. Hubaux, and A. El Fawal. DOMINO: Detecting MAC layer greedy behavior in IEEE 802.11 hotspots. *IEEE Transactions on Mobile Computing*, 5(12):1691–1705, Dec. 2006.

[31] A. Richa, S. Schmid, C. Scheideler, and J. Zhang. A jamming resistant mac protocol for multi-hop wireless networks. In *Proc. of the 24th Int.Symposium on Princ. of Distributed Computing*, 2010.

[32] J. Rodgers. Pulse radar systems, 1962. US Patent 3,029,429.

[33] L. Rosenberg and D. Gray. Anti-jamming techniques for multichannel SAR imaging. *IEE Proceedings-Radar, Sonar and Navigation*, 153(3):234–242, June 2006.

[34] Y. E. Sagduyu, R. Berry, and A. Ephremides. Mac games for distributed wireless network security with incomplete information of selfish and malicious user types. In *IEEE International Conference on Game Theory for Networks, 2009*, pages 130–139.

[35] P. Tague, S. Nabar, J. A. Ritcey, and R. Poovendran. Jamming-aware traffic allocation for multiple-path routing using portfolio selection. *IEEE/ACM Transactions on Networking.*

[36] P. Tague, D. Slater, G. Noubir, and R. Poovendran. Linear programming models for jamming attacks on network traffic flows. In *Proc. 6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, pages 207–216, Berlin, Germany, Apr. 2008.

[37] D. J. Torrieri. *Principles of Secure Communication Systems*. Artech House, Boston, 2nd edition, 1992.

[38] X. Wang and T. Sandholm. Reinforcement learning to play an optimal Nash equilibrium in team Markov games. In *In Proceedings of the Neural Information Processing Systems: Natural and Synthetic (NIPS) conference*, 2002. Extended version at http://www.cs.cmu.edu/ sandholm/oal.ps.

[39] M. Wilhelm, I. Martinovic, J. Schmitt, and V. Lenders. Reactive jamming in wireless networks: How realistic is the threat? In *Proc. 4th ACM Conference on Wireless Network Security*, Hamburg, Germany, June 2011.

[40] W. Xu, K. Ma, W. Trappe, and Y. Zhang. Jamming sensor networks: Attack and defense strategies. *IEEE Network*, 20(3):41–47, May/June 2006.

[41] Q. Zhu, W. Saad, Z. Han, H. V. Poor, and T. Başar. Eavesdropping and jamming in next-generation wireless networks: A game-theoretic approach. In *IEEE Military Communication Conference, 2011*, pages 119–124.