Gaming the Jammer: Is Frequency Hopping Effective ?

Konstantinos Pelechrinis, Christos Koufogiannakis, Srikanth V. Krishnamurthy
Department of Computer Science and Engineering
University of California, Riverside
{*kpele, ckou, krish*}*@cs.ucr.edu*

Abstract*—Frequency hopping has been the most popularly considered approach for alleviating the effects of jamming attacks. In this paper, we provide a novel, measurement-driven, game theoretic framework that captures the interactions between a communication link and an adversarial jammer, possibly with multiple jamming devices, in a wireless network employing frequency hopping (FH). The framework can be used to quantify the efficacy of FH as a jamming countermeasure. Our model accounts for two important factors that affect the aforementioned interactions: (a) the number of orthogonal channels available for use and (b) the frequency separation between these orthogonal bands. If the latter is small, then the energy spill over between two adjacent channels (considered orthogonal) is high; as a result a jammer on an *orthogonal* band that is adjacent to that used by a legitimate communication, can be extremely effective. We account for both these factors and using our framework we provide bounds on the performance of proactive frequency hopping in alleviating the impact of a jammer. The main contributions of our work are: (a) Construction of a measurement driven game theoretic framework which models the interactions between a jammer and a communication link that employ FH. (b) Extensive experimentation on our indoor testbed in order to quantify the impact of a jammer in a 802.11a/g network. (c) Application of our framework to quantify the efficacy of proactive FH across a variety of 802.11 network configurations. (d) Formal derivation of the optimal strategies for both the link and the jammer in 802.11 networks. Our results demonstrate that frequency hopping is largely inadequate in coping with jamming attacks in current 802.11 networks. In particular, we show that if current systems were to support hundreds of additional channels, FH would form a robust jamming countermeasure[1].

Index Terms – Measurements, Analysis, Performance, Security, IEEE 802.11, Frequency hopping, Game theory, Jamming

I. Introduction

The availability of commercial jamming devices make it easy for malicious attackers to disrupt operations of a wireless network [1] [2]. Numerous jamming attacks have been reported in the recent past [3] [4] [5]; this makes the defense against such attacks very critical. A jammer continually emits electromagnetic signals on the medium in order to prevent legitimate data exchanges. In particular a jammer achieves its goal in a CSMA/CA network (e.g. 802.11, sensor networks) by exploiting two transceiver functionalities: **(a)** the MAC protocol requires a transmitter to sense the medium to be idle prior to transmitting its packet; thus, in the presence of illegitimate jamming packets on the medium, a node will defer its transmissions, and **(b)** the packets from the jammer collide with legitimate packets at the receiver. Both of the above effects cause a drastic degradation in the achieved throughput.

Traditionally, frequency hopping has been considered to be a solution that can help alleviate the effects of jamming; both proactive and reactive frequency hopping strategies have been proposed in the literature [6] [7] [8] [9]. The ease of implementation has made proactive frequency hopping more popular; reactive frequency hopping has associated synchronization challenges between the transmitter and the receiver (to be discussed). In this paper, we construct a measurement-driven, analytical framework for quantifying the efficacy of proactive frequency hopping[2]. Our framework accounts for two factors that affect such a strategy. First, the number of available orthogonal channels dictates the effectiveness of frequency hopping. Second, depending on the separation between adjacent orthogonal channels on the available spectrum, there might be an energy spill over between the bands. All prior efforts on frequency hopping assume that operating on a channel[3] that is orthogonal to that being used by a jammer, automatically protects a link. However if the aforementioned separation between bands is small, then a jammer (on a specific channel) can significantly hurt a legitimate communication that is on an adjacent orthogonal channel.

Our objective in this work is to quantify the efficacy of frequency hopping in coping with jamming attacks. In a nutshell, our contributions in this paper are as follows:

1. *Construction of a measurement-based game theoretic framework to capture the interactions between a link and a jammer employing proactive FH :* We model the interactions between a legitimate link and the jammer as a two-player, zero-sum game. The strategies followed by each player and the payoff matrix account for the factors mentioned above. Our framework assumes that the jammer and the network, iteratively and selfishly try to adapt their strategies to stimulate the *best response* to the strategy of the opponent. Thus, the framework yields bounds on the performance of proactive frequency hopping. We extend our framework to cases with more than one jammer.

2. *Quantifying the impact of a jammer via experiments on an indoor wireless testbed with both 802.11a and 802.11g:* We perform extensive experiments on our 802.11 indoor testbed in order to quantify the impact of a jammer that

[1] This work was done partially with support from the US Army Research Office under the Multi-University Research Initiative (MURI) grants W911NF-07-1-0318 and the NSF NeTS:WN / Cyber trust grant 0721941.

[2] We consider proactive frequency hopping since a practically viable reactive strategy is yet to emerge.

[3] We use the terms band and channel interchangeably.

resides on channels that are orthogonal to the one used by a pair of legitimate transceivers. The results of our experiments show that the presence of a jammer on an adjacent, albeit orthogonal channel to that of the legitimate pair, can still degrade the performance significantly. The throughput achieved by the legitimate pair can be reduced to up to just 10% of the throughput possible under benign conditions. This effect significantly limits the effectiveness of frequency hopping in 802.11 networks.

3. ***Applying our framework to quantify the efficacy of proactive frequency hopping in 802.11 networks:*** The measurements from our indoor testbed are then used to drive our framework, applying which we obtain bounds on the anti-jamming performance of a frequency hopping scheme in 802.11 networks. Our result indicate that proactive frequency hopping provides very limited protection to an 802.11 network, from jamming attacks. Our results show that with just 4 jammers one can basically block all the possible channels with 802.11a; this result is in stark contrast with previous efforts as per which, as many as 12 jammers are required to produce this effect.

4. ***Formal derivation of the optimal strategies for both the link and the jammer in 802.11 networks:*** We formally prove that the jammer has a *unique* optimal FH strategy when only a single jamming device is being employed. We extend the result for cases where multiple devices are used. We also prove certain key properties that have to be fulfilled by an optimal FH strategy, followed by a communication link.

Scope of our work: The main application of our framework is the evaluation of FH as a jamming countermeasure. We wish to point out however that our model captures the interactions between communication links and jammers when FH is used by all entities in the wireless network. As such, it can be used from both perspectives (the communication link's and the jammer's) and provide useful insights based on each player's objective.

The rest of the paper is organized as follows. In section II we discuss related work in brief. Section III describes our measurement-driven, game theoretic framework. We describe our wireless testbed and the experimental methodology in section IV. In section V, we present the experimental results that serve as measurements-inputs for our framework for an 802.11a/g network. Section VI describes the application of our framework and the computation of performance bounds of a generic, proactive, frequency hopping scheme for the case of 802.11 networks; the optimal strategies are derived for both the legitimate communication pair and the jammer. We further examine the impact of having additional channels in current 802.11 systems on the effectiveness of FH. Our conclusions form section VII.

II. Background and Related Work

In this section we provide a brief overview on previously proposed frequency hopping schemes; we also discuss the practical limitations of these strategies.

A. Frequency Hopping Strategies

Frequency hopping strategies can be divided into two main categories.

1) Proactive frequency hopping: In a proactive frequency hopping scheme the pair of transceivers that form a link switch channels once every k seconds, irrespective of whether or not there is a jammer on the current channel. Gummadi *et al* [8] propose a rapid proactive frequency hopping scheme to alleviate the impact of specific patterns of narrow-band interference. Navda *et al* [6] implement a proactive frequency hopping protocol with pseudo-random channel switching for coping with a jammer. They compute the optimal residence time on a channel, assuming that the jammer is aware of the hopping protocol. However, they do not account for the energy spill over between adjacent orthogonal channels. A proactive strategy has the advantage of obviating the need for a jamming detection module. We wish to point out here that depending on the implementation, hopping between channels can also potentially incur a performance penalty due to the loss of throughput during the periods used for switching between frequencies [10]; however, in professional implementations these penalties are likely to be extremely small.

2) Reactive frequency hopping: In a reactive frequency hopping scheme, a node switches to a new channel only if and when it detects the presence of a jammer. With such a scheme, when one member of a communicating node pair switches to a new channel, the other member will have to somehow detect the event and change its band as well. Xu *et al* [7] [9] propose a reactive channel hopping strategy. The key idea is that when a node is jammed it switches to a new but predetermined channel. The other node of the communicating pair switches to the same channel upon not hearing from its partner for a prolonged period of time. The authors point out the challenges in the implementation of such a strategy but do not provide solutions. In particular, there are issues related to synchronization, scalability, loss of packets and latency.

Given the ease of implementation, proactive frequency hopping strategies have been more popularly considered for coping with jamming. An effective reactive frequency hopping strategy is yet to emerge. Given this, we primarily consider a proactive approach in this work.

B. Practical Limitations of Frequency Hopping

Channel surfing (switching between channels) tries to avoid the jammer by switching between multiple orthogonal narrow spectral bands. The method can be effective in the presence of a narrow band jammer. In the presence of a wide band jammer that can simultaneously jam multiple bands (and in the extreme case, all possible bands) frequency hopping will not offer any benefits [11]. Given this, we only examine frequency hopping from the perspective of its effectiveness in coping with narrow band jammers.

The performance of frequency hopping will be limited by the extent to which an interferer on an adjacent (considered orthogonal) channel affects a considered channel [12] [13].

In [7] the authors take it for granted that 802.11a supports 12 *perfectly* orthogonal channels; this would imply that the presence of a jammer on one specific channel does not affect the other channels. In [8] the authors measure the throughput that is achieved when there is an interferer on a frequency band that is $15MHz$ apart from the one being used by a legitimate communication. Given that the channel bandwidth with 802.11a is $20MHz$ ($22MHz$ with 802.11g), this scenario reflects the case of partially overlapped channels. The authors show that under these conditions, the overall throughput reduces to $2-3$ Mbps from the base rate of 6 Mbps; they conclude that 50% of the interference-free throughput is achievable if the interferer is present on a partially overlapped channel. We observe that the presence of a jammer on even *an adjacent orthogonal* channel ($20MHz$ apart from the channel of the legitimate communication) causes the throughput to drop to $3-4Mbps$. This is discussed in detail with our 802.11 measurements in section V. We observe that the jamming-free throughput that is achievable on these links is around 27 Mbps (the links inherently support data rates that are much higher than the 6Mbps considered in [8]) and thus, the jammer degrades the throughput to about just $10-15\%$ of what is achievable. In summary, the presence of a jammer on an adjacent orthogonal channel can significantly hurt the performance of a legitimate communication; this in turn limits the effectiveness of frequency hopping strategies.

C. Game theoretic formulations of attacks

In the literature, game theoretic approaches have been used to model various wireless network problems. The work in [14] studies the problem of a legitimate node and a jammer transmitting to a common receiver and models it as a dynamic game. However, this work is theoretical; it suggests that the player that transmits with the higher power is the winner of the game. In contrast, our work is measurement driven and captures the interactions observed on a real network; it provides a comprehensive look at the performance of proactive frequency hopping in coping with jamming attacks. In [15], the authors examine the interactions between a single channel sensor network and a jammer. They are concerned with the detection of the jammer and more specifically, they try to minimize the detection time. They formulate and solve non-linear optimization problems to compute best responses of the attacker and the network to the worst-case strategy of the other. The authors of [16] use linear programming to model a specific class of attacks on network flows. Their work however, differs substantially from ours; it is not based on experimentation and does not consider channel surfing. Liu *et al* [17] propose a novel approach SPREAD, to address the problem of cross layer DoS attacks in wireless data networks. They use a game theoretic approach to describe the interactions between a smart jammer that takes into account protocol specific parameters and the possible decisions of SPREAD. However, their work is neither based on experimentation nor does it examine the performance of frequency hopping.

D. Prior work on energy spill over between 802.11 channels

The authors in [18] try to exploit partially overlapped channels to improve the end-to-end application throughput. The efforts in [19] [20] and [21] try to understand the impact of the use of adjacent channels on a multi-radio, multi-hop 802.11 mesh network. Their findings indicate that multi-hop performance in mesh networks is affected by the adjacent channel interference that one NIC (Network Interface Card) imposes on the other NIC of the same node. However, none of the above efforts consider the presence of a malicious node, which injects packets on the medium to launch an attack.

To the best of our knowledge, our work is the first attempt to construct a measurement based analytical framework which quantifies the performance of a generic proactive frequency hopping strategy in coping with jamming attacks in any given wireless network.

III. Our Framework: The Generic Model of the Game

In this section we present our game which models the interactions between the legitimate communication link and the jammer. Both entities employ frequency hopping in order to achieve their objectives. On the one hand the link switches between bands in order to avoid the jammer; on the other hand the jammer hops across bands in order to find the communication link and hurt its performance. We model this interaction as a game. A game in normal form can be represented by a triplet $< N, (\Sigma_i), A >$. In this representation, N is the finite set of players, Σ_i is the set of possible strategies for player i and A is the payoff matrix of the game.

In our case the set N contains only two players; the jammer and the legitimate link. Both these players have the same set of strategies; $\Sigma = \{set\ of\ available\ orthogonal\ bands\}$. The payoff matrix should represent the objectives of each player. In our case the objective of the legitimate link is to increase its throughput by hopping channels - i.e. changing its strategy - while the objective for the jammer is to reduce this throughput. As a result, an appropriate definition for the payoff matrix is the following: $A_{i,j}$ is the percentage of the jamming-free throughput that the legitimate link enjoys when it resides on channel i and the jammer is residing on channel j. With this definition of the payoff matrix, the value (or the payoff) v of the game is defined to be the percentage of the jamming-free throughput that is achieved on the link. On the one hand, the link is trying to maximize its payoff; on the other hand the jammer is trying to minimize the same payoff. As a result our game is zero-sum, two person game. This means that an **an equilibrium always exists** [22][4]. Our analysis yields the probabilities with which the legitimate link and the jammer ought to occupy the various channels in order to achieve the equilibrium performance.

[4] We wish to stress that our goal is not to provide a system that will compute this equilibrium in real time, but to quantify the performance of a proactive frequency hopping scheme.

The link chooses its channel randomly, using a probability distribution (mixed strategy) x, while the jammer picks its channel as per a probability distribution y. With this, the expected throughput achieved on the link (value of the game) is simply $v = x^T A y$. We can always find the equilibrium strategies x^* and y^*, by solving the above game. The optimal mixed strategy x for the maximizing player (the legitimate link) can be found by solving the following linear program:

$$\text{maximize} \quad v \tag{1}$$

$$\text{subject to} \quad A^T x \geq v \tag{2}$$

$$|x| = 1 \tag{3}$$

$$x \geq 0 \tag{4}$$

and the optimal strategy y for the minimizing player (the jammer) is found as the solution to the dual linear program:

$$\text{minimize} \quad v \tag{5}$$

$$\text{subject to} \quad Ay \leq v \tag{6}$$

$$|y| = 1 \tag{7}$$

$$y \geq 0 \tag{8}$$

In the above formulation, $|x|$ is the 1-norm of vector x, i.e., the sum of all its coordinates. If both players play the game according to their equilibrium mixed strategies x^* and y^*, (computed by solving the above linear programs) the game would be in an equilibrium state. *At equilibrium, no player would benefit from changing the probability distribution with which they choose their channels.*

From the above formulation one can see that our framework accounts for both (i) the number of available orthogonal channels of the wireless technology under consideration and (ii) the effectiveness of a jammer which resides in a different orthogonal band. In the following sections we will show how we can apply our framework to an 802.11 network[5].

IV. Experimental Setup

Prior to applying our framework to various 802.11 configurations, we describe our wireless testbed and the methodology followed in our experiments.

A. Testbed Description

Our wireless testbed consists of 32 Soekris net4826 nodes [23]. Each node mounts a Debian Linux distribution with kernel v2.6.16.19 over NFS. The nodes are synchronized with an NTP server. The Soekris boxes have 2 miniPCI slots. These nodes are equipped with two miniPCI 802.11a/g WiFi cards; in particular, they have an *EMP-8602 6G* with Atheros chipset and an *Intel-2915*. The layout of our testbed is depicted in Figure 1.

With our *EMP-8602 6G* cards, we use the MadWifi driver [24]. In addition, we use a proprietary version of the *ipw2200* AP and client driver/firmware with the *Intel-2915* cards. This was provided to us by Intel Research. With this version we are able to tune the CCA (**C**lear **C**hannel **A**ssessment) threshold parameter; note that this functionality has been implemented in the prototype firmware. The ability to tune the CCA threshold helps us implement a jammer as discussed later in this section.

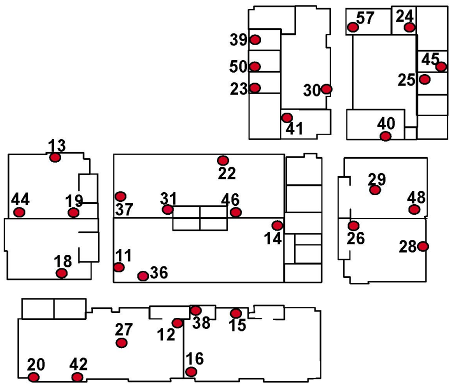

Fig. 1. Deployment of our wireless testbed.

B. Experimental Methodology

Our measurements are on a large set of individual links on our testbed. We perform experiments by varying the transmission powers of both the jammer(s) and the legitimate transceivers. We perform experiments with 802.11a and 802.11g. Our experiments with 802.11g are conducted late at night in order to avoid interference from other co-located WLANs that operate at the same frequency band. In our experiments, we have used all the orthogonal channels that are available with all modes of operation. There are only 3 orthogonal channels in the $2.4 GHz$ band (i.e., 802.11g), while there are 12 orthogonal channels in the $5 GHz$ band (i.e., 802.11a).

C. Implementing a Jammer

To facilitate our experiments, we implement our own jamming utility. The implementation of a jammer with an 802.11 legacy device has to ensure that: **(a)** other packets on the medium do not prevent the jammer from transmitting its packets, and **(b)** when active, the jammer should be able to send its malicious packets at the maximum possible rate in order to cause high impact on legitimate connections. The former requires the tuning of the CCA threshold, while the latter calls for the use of specific types of packets.

We implement our jammer on an 802.11 legacy device by setting the CCA threshold to a very high value (0 dBm). This ensures that the device ignores the traffic in transit over the wireless medium. We observe that packets always arrive at the jammer's circuitry with power less than 0 dBm even if the distances between the jammer and the legitimate transceivers are very small.

In order to ensure that the jammer continuously transmits packets on the medium, we have developed a user-space software utility. With this, the jammer continuously *broadcasts* UDP packets. Given that the backoff functionality is by default

[5] We will also show how we can easily extend our framework to account for the case of more than one jammer.

disabled in 802.11 for broadcast traffic, our software utility can ensure that packets are sent as fast as possible. With such transmissions the jammer does not wait for any ACK packets[6]. Our utility employs *raw sockets*, which allow the construction of a UDP packet from scratch and the forwarding of the packet directly down to the hardware, for transmission. Note here that such an operation requires administrative privileges. To summarize, our jammer utility consists of a specific NIC configuration that sets CCA=0 and a software utility for continuously generating and transmitting broadcast packets. The former feature is possible with our *Intel-2915* cards, since we have access to the firmware.

For our experiments we also utilized the *iperf* measurement tool to generate data traffic with packets of size 1500 bytes, on a legitimate link. Note that, we use the terms *the communication link*, *the link* and *legitimate link* interchangeably. We initiate traffic between the nodes and immediately after, we turn on the jammer(s). In the following section we present the results of our experiments.

V. Measuring the Impact of a Jammer in 802.11 Networks

In this section we present the measurements that will drive the payoff matrix of our game in the context of 802.11 networks. The measurements quantify the impact of a jammer that resides on a channel that is orthogonal to that of the communication link; we observe how this affects the performance of the legitimate link and incorporate these observations into our framework. We describe our experiments with both 802.11a and 802.11g.

We use $\mathbf{RSSI_J = max(RSSI_{JT}, RSSI_{JR})}$ to denote the maximum RSSI (**R**eceived **S**ignal **S**trength **I**ndicator) value that is observed on a link with regards to the signal from the jammer[7]. $RSSI_{JT}$ is the RSSI due to the signal from the jammer at the transmitter, while $RSSI_{JR}$ is the corresponding RSSI as observed at the receiver. As mentioned earlier, the jammer can affect both the transmitting and receiving functions of a node; in particular, it can cause interference at the receiver while it can cause the transmitter to defer its transmissions. By choosing the maximum value, we capture the case wherein the jammer has the maximum impact on the considered link. $\mathbf{RSSI_l = min(RSSI_{TR}, RSSI_{RT})}$ denotes the minimum RSSI value between the end points of the communication link. $RSSI_{TR}$ is the RSSI of the signal from the transmitter at the receiver, while $RSSI_{RT}$ is the RSSI in the reverse direction. $RSSI_l$ represents the worst case RSSI for the link in the realistic scenario where the link is not symmetric.

A. Impact of Jamming in 802.11a

The 802.11a standard supports 12 orthogonal bands or channels. Each of these channels is of $20MHz$ bandwidth. The spacing between the central frequencies of these bands is $20MHz$ as well. In general, when two links communicate on orthogonal bands it is assumed that one does not interfere with the other. This observation drives all the frequency hopping schemes proposed thus far. These schemes assume that via a transition to a channel that is orthogonal to that of the jammer, a communication link can be completely protected. However, this assumption does not hold with two adjacent orthogonal channels. We first present our experimental results to demonstrate this and later, discuss the reasons for this effect.

In our experiments a legitimate connection is initiated on one of the 12 orthogonal channels of 802.11a. Subsequently, the jammer is turned on. The jammer sequentially sweeps the 12 orthogonal channels, one channel at the time. We measure the throughput of our legitimate connection in each case. We repeat the experiments for various $RSSI_J$ and $RSSI_L$ values, in order to account for various topologies. In Figure 2 we present the results for the case where the communication channel was channel 56. The results were similar when the legitimate connection was established on any other different channel.

Our main observation is that ***a jammer which transmits signals on an orthogonal band that is adjacent to that of the legitimate communication, can significantly degrade the throughput performance. Specifically, the throughput of the connection drops to approximately 10 to 15 % of the jamming-free throughput.*** The exact degradation depends on the distance between the jammer and the link and the corresponding channel characteristics. However, our measurements indicate that when $RSSI_J \gg CCA$ for a co-channel user, that user gets at most 15% of the jamming-free throughput if it were to use the adjacent orthogonal bands. The reason for this may be attributed to the fact that RF filters typically do not provide sharp cut-offs at the specified boundaries of the channels [12]. As a result, the spectral power from the signal in one channel (that of the jammer) may spill over to an adjacent channel (that of the legitimate communication), even if in theory they are considered orthogonal. In order to completely avoid the effects of jamming, the legitimate connection will have to be at least 2 orthogonal channels apart from the channel on which the jammer is present.

Next, we conducted experiments with two jammers. We considered all possible placements of the jammers on the 12 orthogonal channels. Our main observations are summarized in figure 3. When the two jammers reside on the two orthogonal channels adjacent to that of the communication link, the degradation in the link throughput can be as high as 95%.

We use these measurements as inputs to our game-theoretic framework in section VI.

B. Impact of a Jammer With 802.11g

In contrast with 802.11a, 802.11g has only 3 orthogonal channels, each of which is of $22MHz$ bandwidth. The central frequencies of these bands are however, $25MHz$ apart. This

[6] This configuration allows the deferral of back-to-back transmissions for the minimum possible time (i.e. $DIFS + min_{BackOff}$).

[7] This is measured when both the jammer and the communication link are on the same channel.

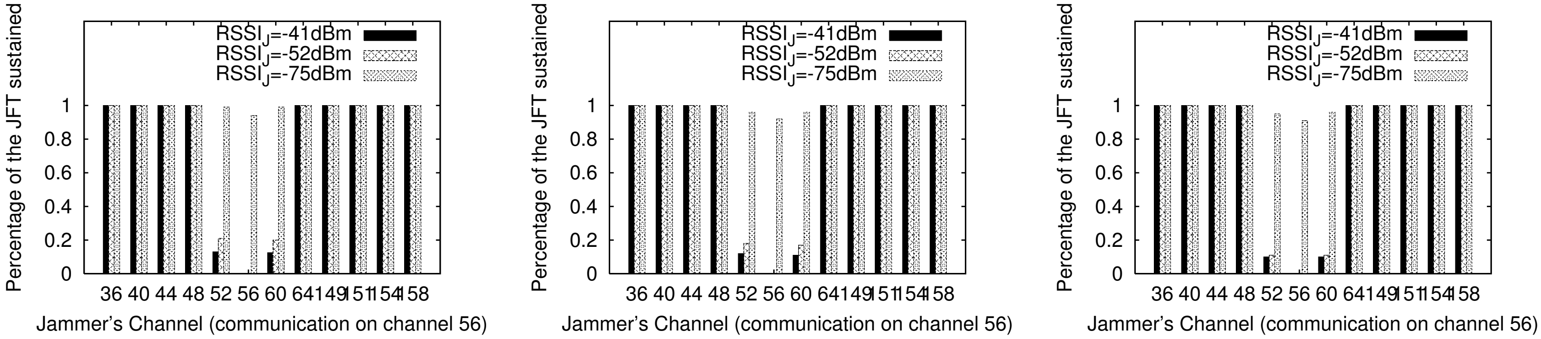

Fig. 2. Percentage of the jamming free throughput (JFT) achieved when the jammer is on various channels, and for various $RSSI_J$, for the case of 802.11a. In the three figures we have $RSSI_l = -37dBm$, $RSSI_l = -47dBm$ and $RSSI_l = -66dBm$, respectively.

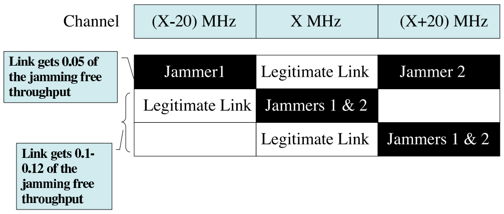

Fig. 3. The case of 2 jamming nodes on adjacent communication channels.

implies that there is a *secure zone* of $3MHz$ between the adjacent orthogonal channels. Conducting the same experiments as before, we obtain the results in Figure 4.

As with 802.11a, we observe that in the presence of a jammer on an orthogonal, adjacent channel, the performance of a legitimate connection is still degraded. However, with 802.11g the degradation is significantly lower. This can be primarily attributed to the *larger* channel separation between adjacent orthogonal channels; this results in a reduced seepage of the spectral power of the jammer into the adjacent channel being used by the legitimate connection. However, since there are only 3 orthogonal bands in 802.11g, frequency hopping is not expected to be very effective.

VI. Applying our Framework in 802.11 Networks

In this section we will apply our game-theoretic framework based on the measurements presented in the previous section.

A. Model for 802.11a

An 802.11a wireless network can support twelve orthogonal channels. For ease of presentation, we label the channels: 1, 2, …, 12. Based on the measurement results obtained in the previous section, if the jammer is on a channel that is adjacent to that of the link, we assume that the link can achieve 12% of its jamming-free throughput; if the jammer is on the same channel as that of the link, no throughput is achieved. If two jamming devices reside on the two adjacent channels of the link, the throughput achieved on the link is just 5% of the jamming-free throughput. Note here that, if the link were to operate either on channel 1 or 12, the jammer could only impact the link via one adjacent channel; for the other cases, there are two such possible adjacent channels.

First, we consider the case where the communication link is on channel i and we have a single jamming device on channel j. The payoff matrix is then given by:

$$A^{1,a}_{i,j} = \begin{cases} 0 & \text{if } i = j, \\ 0.12 & \text{if } |i-j| = 1, \\ 1 & \text{otherwise.} \end{cases}$$

If the link is on channel i and the jammer uses two jamming devices, one on channel j_1 and the second on j_2 (where $j_1 \leq j_2$ without loss of generality) the payoff matrix is given by:

$$A^{2,a}_{i,j_1j_2} = \begin{cases} 0 & \text{if } i = j_1 \text{ or } i = j_2, \\ 0.05 & \text{elseif } i = j_1 + 1 \text{ and } i = j_2 - 1, \\ 0.12 & \text{elseif } |i - j_1| = 1 \text{ or } |i - j_2| = 1, \\ 1 & \text{otherwise.} \end{cases}$$

Similarly, the payoff matrix when we have three jamming devices on channels $j_1 \leq j_2 \leq j_3$ is:

$$A^{3,a}_{i,j_1j_2j_3} = \begin{cases} 0 & \text{if } i = j_1 \text{ or } i = j_2 \text{ or } i = j_3, \\ 0.05 & \text{elseif } (i = j_1 + 1 \text{ and } i = j_2 - 1) \\ & \text{or } (i = j_2 + 1 \text{ and } i = j_3 - 1), \\ 0.12 & \text{elseif } |i - j_1| = 1 \text{ or } |i - j_2| = 1 \\ & \text{or } |i - j_3| = 1, \\ 1 & \text{otherwise.} \end{cases}$$

Finally, if there are four jamming devices on channels $j_1 \leq$

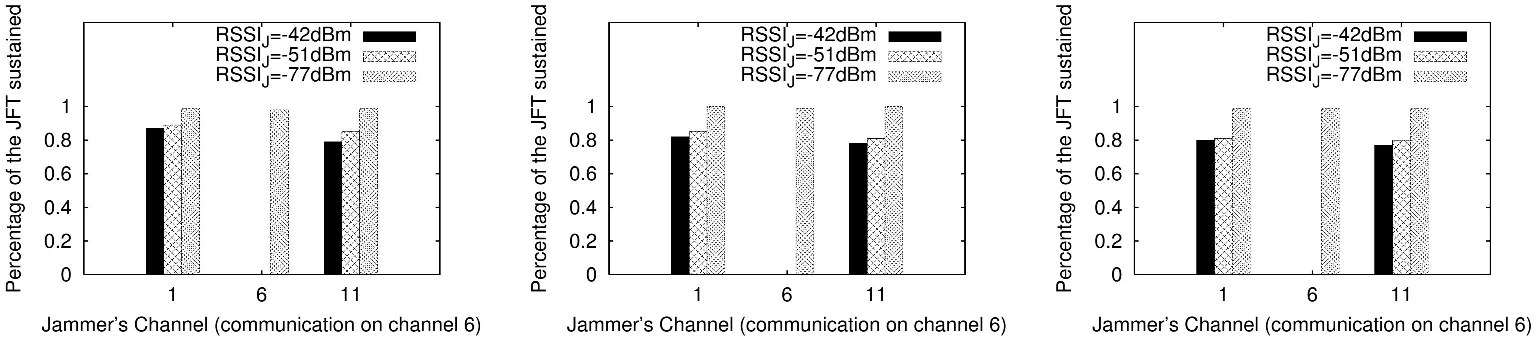

Fig. 4. Percentage of the jamming free throughput (JFT) achieved when the jammer is on various channels, and for various $RSSI_J$, for the case of 802.11g. In the three figures we have $RSSI_l = -39dBm$, $RSSI_l = -45dBm$ and $RSSI_l = -68dBm$, respectively.

$j_2 \leq j_3 \leq j_4$, the the payoff matrix is:

$$A^{4,a}_{i,j_1 j_2 j_3 j_4} = \begin{cases} 0 & \text{if } i = j_1 \text{ or } i = j_2 \text{ or } i = j_3 \text{ or } i = j_4, \\ \\ 0.05 & \text{elseif } (i = j_1 + 1 \text{ and } i = j_2 - 1) \\ & \text{or } (i = j_2 + 1 \text{ and } i = j_3 - 1) \\ & \text{or } (i = j_3 + 1 \text{ and } i = j_4 - 1), \\ \\ 0.12 & \text{elseif } |i - j_1| = 1 \text{ or } |i - j_2| = 1 \\ & \text{or } |i - j_3| = 1 \text{ or } |i - j_4| = 1, \\ \\ 1 & \text{otherwise.} \end{cases}$$

In all cases we use the linear programs (1)-(4) and (5)-(8) in order to compute optimal strategies for the link and the jammer respectively. First, let us consider the scenario where there is just one jamming device. Then, the mixed strategies x^* and y^* are tabulated in I and II.

channel j	2	5	8	11
y_j	0.25	0.25	0.25	0.25

TABLE I
MIXED STRATEGY FOR ONE JAMMING DEVICE IN 802.11A

channel i	1	3	4	6	7	9	10	12
x_i	.184	.066	.1422	.1078	.1078	.1422	.066	.184

TABLE II
MIXED STRATEGY FOR THE COMMUNICATION LINK IN 802.11A

The strategy y^* recommends that the jammer hops uniformly at random between channels 2, 5, 8 and 11. Intuitively, this seems very reasonable since, with the recommended approach, the jammer can harm all the channels to some extent. With the strategy x^*, the link avoids these four channels and hops among the other channels; the distribution x^* is given in table II. If the players play as per these equilibrium strategies, the value v of the game is $v = 0.78$. This implies that the expected throughput on the link is 78% of its jamming-free throughput.

When the jammer employs two jamming devices, the devices must be used on channels $\{2, 5\}$ with probability 0.5, and on channels $\{8, 11\}$ with probability 0.5. The communication link on the other hand, should avoid these channels. The best strategy for the link, x^* is still as per table II. However, the payoff with this equilibrium pair of strategies is 56%. Note that, as one might expect, this is significantly lower than with just one jamming device.

channels (j_1, j_2)	(2,5)	(8,11)
y_{j_1,j_2}	0.5	0.5

TABLE III
MIXED STRATEGY FOR TWO JAMMING DEVICES IN 802.11A

channel (j_1, j_2, j_3)	(2,5,8)	(2,5,11)	(2,8,11)	(5,8,11)
y_{j_1,j_2,j_3}	0.25	0.25	0.25	0.25

TABLE IV
MIXED STRATEGY FOR THREE JAMMING DEVICES IN 802.11A

If the jammer has three jamming devices, its equilibrium strategy uses the devices on channels $\{2, 5, 8\}$ with probability 0.25, on channels $\{2, 5, 11\}$ with probability 0.25, on channels $\{2, 8, 11\}$ with probability 0.25, and on channels $\{5, 8, 11\}$ with probability 0.25. As before, the link should avoid the channels 2,5,8 and 11; the equilibrium strategy x^* is given by table II. In this case, the link can achieve just 34% of its jamming-free throughput.

channel (j_1, j_2, j_3, j_4)	(2,5,8,11)
y_{j_1,j_2,j_3,j_4}	1

TABLE V
MIXED STRATEGY FOR FOUR JAMMING DEVICES IN 802.11A

Finally if the jammer has four jamming devices, they should be made active on channels 2,5,8 and 11. The link would avoid these channels. The expected payoff is just 12%, which in practice means that the communication is almost completely blocked. Table VI summarizes the expected percentage of the jamming-free throughput for the case of one, two, three and four jamming devices.

# jammers	1	2	3	4
v	78%	56%	34%	12%

TABLE VI
EXPECTED LINK THROUGHPUT FOR 802.11A, USING DIFFERENT NUMBERS OF JAMMERS

Sensitivity to measurements: The results thus far, were based on a premise that if the link was on a channel that was adjacent to that being used by the jammer, only 12% of its jamming-free throughput can be achieved. Note that in practice, the exact degradation experienced varies depending on the locations of the link and the jammer and the environment. Our experiments suggest that only up to 10-15% of the jamming free throughput is achieved. Using any other value in this range for the payoff matrix would not change the results significantly (at most 3% change).

Uniqueness: The following lemmas formally prove that (i) a jammer should not use any jamming device on channels 1, 3, 4, 6, 7, 9, 10, 12 and that it should choose channels 2,5,8,11 with equal probability and, (ii) the link should not use channels 2, 5, 8, 11, since the jammer will *hit* these channels.

Lemma 1: The linear program (5)-(8), with $A = A^1$, has just one optimal solution $y = y^*$, which is $y_1 = y_3 = y_4 = y_6 = y_7 = y_9 = y_{10} = y_{12} = 0$, $y_2 = y_5 = y_8 = y_{11} = 0.25$.

Proof: We prove the lemma by contradiction. Let there be a second optimal solution $\hat{y} \neq y^*$. In other words, if possible, let there be a solution $\hat{y}$ with a non-zero 1-norm distance from y^*. The 1-norm distance is defined as $|\hat{y} - y^*| = \sum_{i=1}^{12} |\hat{y}_i - y_i^*|$. If we cannot find such a solution $\hat{y}$, then the solution y is unique. In other words, we want to check if the following optimization problem has a zero objective value or not. The optimization problem that we want to solve is:

$$\begin{aligned} \text{maximize} \quad & |\hat{y} - y^*| && (9)\\ \text{subject to} \quad & A\hat{y} \leq 0.78 && (10)\\ & |\hat{y}| = 1 && (11)\\ & \hat{y} \geq 0 && (12) \end{aligned}$$

The above formulation is not a linear program (the objective function is non-linear). We reduce the problem into solving $2 \cdot 12 = 24$ linear programs below. For each of the linear programs, our goal is to check if the objective function is zero.

For $i = 1, \ldots, 12$,

$$\begin{aligned} \text{maximize} \quad & \hat{y}_i - y_i^* && (13)\\ \text{subject to} \quad & A\hat{y} \leq 0.78 && (14)\\ & |\hat{y}| = 1 && (15)\\ & \hat{y} \geq 0 && (16) \end{aligned}$$

$$\begin{aligned} \text{maximize} \quad & y_i^* - \hat{y}_i && (17)\\ \text{subject to} \quad & A\hat{y} \leq 0.78 && (18)\\ & |\hat{y}| = 1 && (19)\\ & \hat{y} \geq 0 && (20) \end{aligned}$$

By solving each of the above linear programs, we verify that the objective value is zero. This proves the uniqueness of solution y^*. ■

For any number of jamming devices, the equilibrium strategy for the jammer is the selection of channels 2, 5, 8, 11 with uniform probability. For example, with two jamming devices, an equilibrium strategy selects channels (2,5) with probability 0.5 and channels (8,11) with probability 0.5; note that another equilibrium strategy can be selecting channels (2,11) with probability 0.5 and channels (5,8) with probability 0.5.

Lemma 2: Any equilibrium strategy x^* for the maximizing player (the link) has $x_2 = x_5 = x_8 = x_{11} = 0$.

Proof: To prove that in any optimal solution, $x_2 = x_5 = x_8 = x_{11} = 0$, we formulate the following linear program.

$$\begin{aligned} \text{maximize} \quad & x_2 + x_5 + x_8 + x_{11} && (21)\\ \text{subject to} \quad & A^T x \geq 0.78 && (22)\\ & |x| = 1 && (23)\\ & x \geq 0 && (24) \end{aligned}$$

The linear program tries to find the maximum value for the sum $x_2 + x_5 + x_8 + x_{11}$ under the constraint that the achieved payoff is at least 0.78 (this is the maximum achievable payoff). The solution to the above linear program yields an objective value of zero. In other words, there cannot be any optimal solution with $x_2 = x_5 = x_8 = x_{11} \neq 0$. ■

Lemma 3: If the jammer plays the strategy of lemma 1, then the link player can set $x_1, x_3, x_4, x_6, x_7, x_9, x_{10}, x_{12}$ to any non-negative value, as long as their sum is 1.

Proof: The values of the games with one, two, three or four jamming devices are given by:

$v^1 = x^T Ay = 0.78(x_1 + x_3 + x_4 + x_6 + x_7 + x_9 + x_{10} + x_{12}) + 0.75(x_2 + x_5 + x_8 + x_{11})$

$v^2 = x^T Ay = 0.56(x_1 + x_3 + x_4 + x_6 + x_7 + x_9 + x_{10} + x_{12}) + 0.5(x_2 + x_5 + x_8 + x_{11})$

$v^3 = x^T Ay = 0.34(x_1 + x_3 + x_4 + x_6 + x_7 + x_9 + x_{10} + x_{12}) + 0.25(x_2 + x_5 + x_8 + x_{11})$

$v^4 = x^T Ay = 0.12(x_1 + x_3 + x_4 + x_6 + x_7 + x_9 + x_{10} + x_{12})$

In order to maximize $v^{(i)}$ we should set $x_2 = x_5 = x_8 = x_{11} = 0$, and then set the remaining variables in any non-negative values such that $x_1 + x_3 + x_4 + x_6 + x_7 + x_9 + x_{10} = 1$. ■

B. Model for 802.11g

The model for 802.11g is simpler to solve, given that there are just three orthogonal channels. For one jamming device the payoff matrix is:

$$A_{i,j}^{1,g} = \begin{cases} 0 & \text{if } i = j, \\ 0.88 & \text{if } |i - j| = 1, \\ 1 & \text{otherwise,} \end{cases}$$

For two jamming devices the payoff matrix is given by

$$A_{i,j_1 j_2}^{2,g} = \begin{cases} 0 & \text{if } i = j_1 \text{ or } i = j_2, \\ 0.88 & \text{elseif } |i - j_1| = 1 \text{ or } |i - j_2| = 1, \\ 1 & \text{otherwise,} \end{cases}$$

Note here that interestingly, our measurements indicate that adding one more jamming device on the adjacent orthogonal channel does not further impact the link as compared with the case of one malicious device. This can be attributed to the secure spectral zone with 802.11g; additional energy spillage is negligible. For three jamming devices, all values in the payoff matrix are zero:

$$A^{3,g}_{i,j_1 j_2 j_3} = 0$$

Again, solving the game using linear programming, we get the equilibrium strategies for both players and the expected payoffs (percentage of the link's jamming-free throughput). These payoffs are summarized in table VII.

# jammers	1	2	3
v	61.46%	29.33%	0%

TABLE VII
EXPECTED LINK THROUGHPUT FOR 802.11G, USING DIFFERENT NUMBERS OF JAMMERS

With one jamming device, both players have the same equilibrium strategy; the strategy is tabulated in table VIII.

channel i	1	2	3
y_i	0.3492	0.3016	0.3492
x_i	0.3492	0.3016	0.3492

TABLE VIII
MIXED STRATEGY FOR THE LINK AND ONE JAMMING DEVICE IN 802.11G

If the jammer has two jamming devices, they should be activated in pairs so as to maintain a uniform probability of using each channel. The communication link should also hop among the three channels, uniformly at random. The strategies are shown in tables IX and X.

channels (j_1, j_2)	(1,2)	(1,3)	(2,3)
y_{j_1,j_2}	0.3333	0.3333	0.3333

TABLE IX
MIXED STRATEGY FOR THE TWO JAMMING DEVICES IN 802.11G

channel i	1	2	3
x_i	0.3333	0.3333	0.3333

TABLE X
MIXED STRATEGY FOR THE COMMUNICATION LINK AGAINST TWO JAMMING DEVICES IN 802.11G

With three or more jamming devices, no throughput can be achieved on the link with 802.11g, as one might expect. Next, we prove the uniqueness of the above solutions.

Lemma 4: The solution given in table VIII is the unique optimal solution for the linear programs (1)-(4) and (5)-(8), for $A = A^{1,g}$.

Proof: We prove the lemma for the solution of the dual linear program (5)-(8); a similar proof can be easily constructed for the primal optimal solution x^* in table VIII. The optimal solution for y given by table VIII makes all the constraints tight i.e.,

$$0.88y_2 + y_3 = v \quad (25)$$

$$0.88y_1 + 0.88y_3 = v \quad (26)$$

$$y_1 + 0.88y_3 = v \quad (27)$$

In order to prove this, consider the following:
a) some $\delta > 0$ is subtracted from y_1 and added to y_2 or y_3 or both. Then, the first constraint will yield a value more than v. **b)** some $\delta > 0$ is subtracted from y_2 and added to y_1 or y_3 or both. Then, the second constraint will yield a value more than v. **c)** some $\delta > 0$ is subtracted from y_3 and added to y_1 or y_2 or both. Then, the third constraint will result in a value more than v. **d)** some $\delta_1 > 0$ is subtracted from y_1, some $\delta_2 > 0$ is subtracted from y_2, and $\delta_1 + \delta_2$ added to y_3. Then, the first constraint will yield a value more than v. **e)** some $\delta_1 > 0$ is subtracted from y_2, some $\delta_2 > 0$ is subtracted from y_3, and $\delta_1 + \delta_2$ added to y_1. Then, the third constraint will have value more than v. **f)** some $\delta_1 > 0$ is subtracted from y_1, some $\delta_2 > 0$ is subtracted from y_3, and $\delta_1 + \delta_2$ added to y_2. Then, the sum of the first and the third constraints will be more than $2v$. With this, either the first or the third constraint must result in a value more than v. Thus, there is no way to construct another feasible solution with a value at most v. In other words, the solution in table VIII is unique. ■

C. The Effect of Number of Channels

The number of available channels is a limiting factor on the applicability of frequency hopping in current commodity systems. In this section we want to quantify the efficiency of frequency hopping in coping with jamming with a varying number of orthogonal bands. In other words, we ask the question "what if the commodity systems had higher numbers of orthogonal bands?"; to what extent would it improve the effectiveness of frequency hopping in avoiding a jammer? We solve our game by calibrating a payoff matrix from our measurements but the matrix is appropriately expanded in order to emulate the existence of more channels. In particular, the effect of a jammer residing at an orthogonal band is assumed to be the same as is in current commodity 802.11 systems. We find the solution to our two-player game with new payoff matrices derived from measurements with both 802.11a and g. The results are presented in figure 5. We see that if a fairly large number of channels were available, then frequency hopping would be a very efficient anti-jamming technique. In particular, with a single jammer, the throughput is almost completely restored if the number of channels is close to 100.

In figure 6 we present the number of jamming devices that one would need in order to bring the throughput down to below 20% of the jamming free performance. We notice that the number of devices needed for the model calibrated with measurements using 802.11g are higher than with the model based on 802.11a. This is due to the reduced effect that a jammer residing on an adjacent orthogonal channel has with 802.11g given that the channel spacing is larger. In particular, if 100 channels were available, with the energy spillage between orthogonal channels as with 802.11g, about 80 jammers would be necessary; in the corresponding case, with the energy spillage as with 802.11a, only about 34 jamming devices are sufficient.

Finally in figure 7 we present the number of jamming devices needed in order to drop the throughput of the link

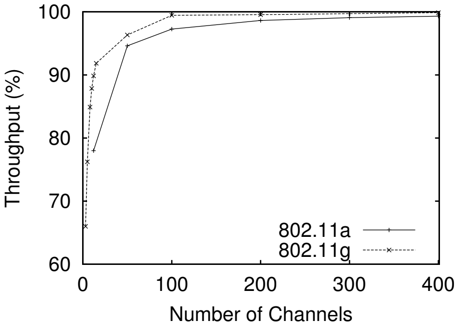

Fig. 5. Increasing the spectrum availability, significantly increases FH's robustness against jamming.

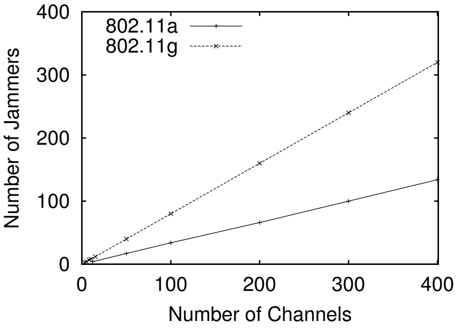

Fig. 6. Number of jammers needed to drop throughput below 20% of the jamming free performance enjoyed.

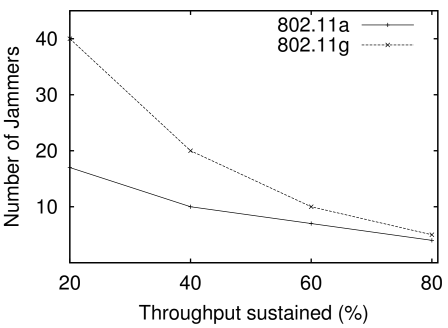

Fig. 7. Number of jammers needed to drop the throughput at a specific percentage (50 channels).

to a specific percentage of the jamming free throughput (x-axis) for a fixed number of channels (50). Again notice, that the jammers will be much more effective if the energy spillage between adjacent channels is higher (as with 802.11a).

VII. Conclusions

In this paper we provide a game theoretic framework in order to capture the interactions between a link and a jammer employing FH. Our framework is measurement driven and accounts for two performance limiting factors; the number of available orthogonal channels as well as the adjacent orthogonal channel, jamming-interference. After formally presenting our framework, we show how we can apply it to 802.11 networks in order to quantify the efficacy of FH as jamming countermeasure. We conduct extensive experiments on our indoor wireless testbed in order to derive the payoff matrix of our game. Our results indicate that frequency hopping is inadequate for protecting 802.11 networks from jamming with current spectrum allocations. We also show that with the same payoff matrix, if the number of orthogonal channels supported was much larger, frequency hopping would be very effective in coping with jamming.

Acknowledgment: We would like to thank Dr. K. Papagiannaki from Intel Research, for providing us with the prototype firmware of *ipw2200* AP.

References

[1] SESP jammers. http://www.sesp.com/.

[2] ISM Wide-band Jammers. http://69.6.206.229/e-commerce-solutions-catalog1.0.4.html.

[3] Jamming attack at hacker conference. http://findarticles.com/p/articles/mi_m0EIN/ is_2005_August_2/ai_n14841565.

[4] Techworld news. http://www.techworld.com/mobility/ news/index.cfm?newsid=10941.

[5] RF Jamming Attack. http://manageengine.adventnet.com/ products/wifi-manager/rfjamming-attack.html.

[6] V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein. Using Channel Hopping to Increase 802.11 Resilience to Jamming Attacks. In *IEEE INFOCOM mini-conference*, 2007.

[7] W. Hu, T. Wood, W. Trappe, and Y. Zhang. Channel Surfing and Spatial Retreats: Defenses Against Wireless Denial of Service. In *ACM Workshop on Wireless Security*, 2004.

[8] R. Gummadi, D. Wetheral, B. Greenstein, and S. Seshan. Understanding and Mitigating the Impact of RF Interference on 802.11 Networks. In *ACM SIGCOMM*, 2007.

[9] W. Hu, K. Ma, W. Trappe, and Y. Zhang. Jamming Sensor Networks: Attacks and Defense Strategies. In *IEEE Network*, May/June 2006.

[10] R. Vedantham, S. Kakumanu, S. Lakshmanan, and R. Sivakumar. Component Based Channel Assignment in Single Radio, Multi-channel Ad Hoc Networks. In *ACM MOBICOM*, 2006.

[11] ISA: Users fear wireless networks for control. http://lists.jammed.com/ISN/2007/05/0122.html.

[12] J. Yee and H. P-Esfahani. Understanding Wireless LAN Performance Tradeoffs. In *http://www.commsdesign.com*, 2002.

[13] P.Li, N.Scalabrino, Y.Fang, E.Gregory, and I.Chlamtac. Channel Interference in IEEE 802.11b. In *Global Telecommunications Conference (GLOBECOM) IEEE*, 2007.

[14] R. Mallik, R. Scholtz, and G. Papavassilopoulos. Analysis of an On-Off Jamming Situation as a Dynamic Game. In *IEEE Trans. Commun., vol. 48, no. 8, pp. 1360-1373*, August 2000.

[15] M.Li, I.Koutsopoulos, and R.Poovendran. Optimal Jamming Attacks and Network Defense Policies in Wireless Sensor Networks. In *IEEE INFOCOM*, 2007.

[16] P. Tague, D Slater, G. Noubir, and R. Poovendran. Linear Programming Models for Jamming Attacks on Network Traffic Flows. In *Network Security Lab (NSL) Technical Report # 002*, 2007.

[17] X.Liu, G.Noubir, R.Sundaram, and S.Tan. SPREAD: Foiling Smart Jammers using Multi-layer Agility. In *IEEE INFOCOM mini-conference*, 2007.

[18] A.Mishra, V.Shrivastava, S.Banerjee, and W.Arbaugh. Partially overlapped channels not considered harmful. In *SIGMETRICS '06/Performance '06: Proceedings of the joint international conference on Measurement and modeling of computer systems*, 2006.

[19] C.M Cheng, P.H Hsiao, H.T Kung, and D Vlah. Adjacent Channel Interference in Dual-radio 802.11a Nodes and Its Impact on Multi-hop Networking. In *Global Telecommunications Conference (GLOBECOM) IEEE*, 2006.

[20] J.Robinson, K.Papagiannaki, C.Diot, X.Guo, and L.Krishnamurthy. Experimenting with a Multi-Radio Mesh Networking Testbed. In *1st workshop on Wireless Network Measurements (WiNMee 2005), Trento, Italy*, 2005.

[21] V.Angelakis, A.Traganitis, and V.Siris. Adjacent channel interference in a multi-radio wireless mesh node with 802.11a/g interfaces. In *IEEE INFOCOM, poster session*, 2007.

[22] Von Neumann J and O.Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press (May 1, 1980) ISBN 0-69-100362-9.

[23] Soekris-net4826. http://www.soekris.com/net4826.htm.

[24] The MAdWiFi driver. http://madwifi.org.