On the Efficacy of Frequency Hopping in Coping with Jamming Attacks in 802.11 Networks

Konstantinos Pelechrinis, Christos Koufogiannakis, and Srikanth V. Krishnamurthy, *Member, IEEE*

***Abstract*—Frequency hopping (FH) has been the most popularly considered approach for alleviating the effects of jamming attacks. We re-examine, the efficacy of FH based on both experimentation and analysis. Briefly, the limitations of FH are: (a) the energy spill over between adjacent channels that are considered to be orthogonal, and (b) the small number of available orthogonal bands. In a nutshell, 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 employing FH. Our model accounts for the above limiting factors and provides bounds on the performance of proactive FH in coping with jamming. (b) Extensive experimentation to quantify the impact of a jammer on 802.11a/g/n networks. Interestingly, we find that 802.11n devices can be more vulnerable to jamming as compared with legacy devices. We carefully analyze the reasons behind this observation. (c) Application of our framework to quantify the efficacy of proactive FH and validation of our analytical bounds across various 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 FH seems to be inadequate in coping with jamming attacks in current 802.11 networks.**

***Index Terms*—Measurements, analysis, performance, security, IEEE 802.11, frequency hopping, game theory, jamming.**

I. Introduction

THE availability of commercial jamming devices makes 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[1]. 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[2] that is orthogonal to that being used by a jammer - i.e., there is no overlap associated with the spectral masks - 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 understand the interactions between a jammer and a communication link and 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 legacy 802.11 (802.11a and 802.11g) as well as its current 4G extension,***

Manuscript received January 26, 2010; revised April 7, 2010 and June 3, 2010; accepted August 1, 2010. The associate editor coordinating the review of this paper and approving it for publication was W. Lou.

K. Pelechrinis is with the School of Information Sciences, University of Pittsburgh (e-mail: kpele@pitt.edu).

C. Koufogiannakis and S. V. Krishnamurthy are with the Department of Computer Science and Engineering, University of California, Riverside (e-mail: {ckou, krish}@cs.ucr.edu).

This work is supported in part by the U.S. Army Research Office under the Multi-University Research Initiative (MURI) grants W911NF-07-1-0318 and the NSF NeTS:WN / Cyber trust grant 0721941. C. Koufogiannakis is partially supported by the Greek State Scholarships Foundation (IKY).

Digital Object Identifier 10.1109/TWC.2010.09.100113

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

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

1536-1276/10$25.00 © 2010 IEEE

802.11n: We perform extensive experiments on our 802.11 indoor testbed in order to quantify the impact of a jammer that 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 of legacy 802.11 significantly. The throughput achieved by the legitimate pair can be reduced to just 10% of the throughput possible under benign conditions. This effect significantly limits the effectiveness of frequency hopping in 802.11 networks.

In addition, our experiments with 802.11n reveal additional vulnerabilities. 802.11n utilizes channel bonding as a way to increase the transmission rate [10]. In a nutshell with channel bonding, two or more adjacent channels are used in conjunction to form a new *wider* channel. Our measurements indicate that this property (in conjunction with the CSMA/CA policy inherited from legacy 802.11) can make 802.11n links more susceptible to jamming attacks. We provide a detailed discussion on why this is the case.

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 5 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 measurement-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 validate our analytical results on our testbed. The performance of an 802.11n MIMO link under the presence of a jammer is considered in section VII. Section VIII discusses the applicability of our framework across a variety of jamming models, while our conclusions form section IX.

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.

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 [11]; however, in professional implementations these penalties are likely to be extremely small.

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. Hu *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 [12]. 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 [13] [14]. 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 [15] 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 highest power is the winner of the game. In contrast, our work is measurement driven and is validated via experimentation; it provides a comprehensive look at the performance of proactive frequency hopping in coping with jamming attacks. In [16], 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 [17] 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* [18] 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.

Finally, in some more recent efforts, emulation attacks in cognitive communication systems are being cast as game theoretic problems. In particular, Li *et al* [19] study a primary user emulation jamming attack in a cognitive radio network utilizing game theoretic notions. The authors provide numerical solutions for different variations of the attack model and show that the performance of a secondary user is improved when the number of available channels is increased. Thomas *et al* [20] model the interactions between a selfish radio and a well behaved radio, as well as between two selfish radios, using the Bayesian game framework. They show that both types of interactions result in games with imperfect knowledge which can lead to Bayesian Nash Equilibrium (BNE) with both pure and mixed strategies. The also show that under different system parameters different BNEs arise.

D. Prior work on energy spill over between 802.11 channels

The authors in [21] try to exploit partially overlapped channels to improve the end-to-end application throughput. The efforts in [22] [23] and [24] 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 with the jammer 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 a zero-sum, two person game. This means that **an equilibrium always exists** [25][3]. 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.

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 $v = x^T Ay$. 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, each of the constraints, (2) and (6), are used to describe the $|\Sigma|$ inequalities in a compact way. In particular, $A^T x$ and Ay are $|\Sigma| \times 1$ vectors, and each element of these vectors should satisfy the corresponding inequality with respect to v. Furthermore, $|x|$ 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[4].

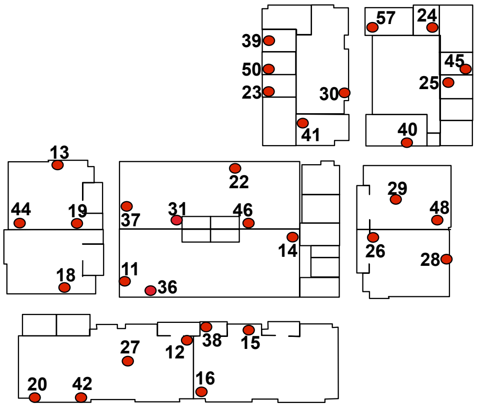

Fig. 1. Deployment of our wireless testbed.

Note here that a probabilistic analysis could be used to model the interactions between a jammer and the communication link. However, as the dimensionality of the problem increases and/or the components of Σ change (e.g., different frequency allocations across large wireless networks), such an analysis is likely to increase in compexity or become intractable. Our game theoretic model on the other hand, is easily applicable in such contexts.

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 802.11a/g wireless testbed consists of 32 Soekris net4826 nodes [26]. 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 the Atheros chipset and an *Intel-2915*. The layout of our testbed is depicted in Fig. 1.

With our *EMP-8602 6G* cards, we use the MadWifi driver [27]. In addition, we use a proprietary version of the *ipw2200* AP and client driver/firmware with the *Intel-2915* cards. 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.

The architecture of our 802.11n testbed is similar to the one described above. However, the nodes are utilizing 15 Soekris net5501 boxes[5], which are equipped with an *RT2860 mini-PCI* card that supports 802.11n communications.

B. Experimental Methodology

Our measurements are on a large set of individual links on our testbed. We perform experiments by varying the

[3] 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.

[4] We will also show how we can easily extend our framework to account for cases with more than one jammer.

[5] These boxes have higher processing capabilities - as compared to net4826 -and can realize the MIMO benefits in terms of achievable throughput [28].

transmission powers of both the jammer(s) and the legitimate transceivers. We perform experiments with all modes, namely, 802.11a/g/n. Our experiments with 802.11g/n are conducted late at night in order to avoid interference from other co-located WLANs that operate at the same frequency band (note that RT2860 operate only in the 2.4GHz band in 802.11n mode). 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 $2.4GHz$ band (i.e., 802.11g), while there are 12 orthogonal channels in the $5GHz$ 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 sending 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 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 Legacy 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 Fig. 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,

[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.

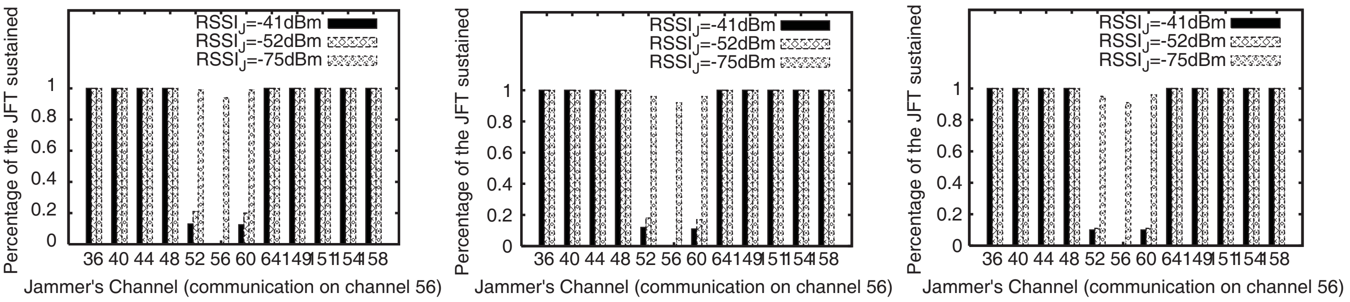

Fig. 2. Percentage of the jamming free throughput (JFT) achieved when the jammer is on various channels, and for various $RSSI_J$, for the case of 802.11a. In the three figures we have $RSSI_l = -37dBm$, $RSSI_l = -47dBm$ and $RSSI_l = -66dBm$, respectively.

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 [13]. 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 would like here to emphasize the fact that the above observations do not hold for channels 64 and 149. These channels are more than $400\ MHz$ apart and as our measurements indicate are completely isolated.

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 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 Fig. 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.

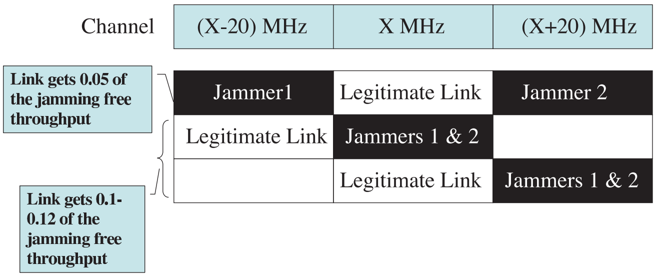

Channel	(X-20) MHz	X MHz	(X+20) MHz
	Jammer1	Legitimate Link	Jammer 2
	Legitimate Link	Jammers 1 & 2	
		Legitimate Link	Jammers 1 & 2

Link gets 0.05 of the jamming free throughput

Link gets 0.1-0.12 of the jamming free throughput

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

VI. Applying and Validating our Framework in Legacy 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[8]. For ease of presentation, we label the channels: 1, 2, . . ., 12. The central frequencies of the channels are 20MHz apart, with the exception of the eighth and ninth channel pair that are 425MHz apart. 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 (with the exception of the eight and ninth channel pair), we assume that the link can achieve only 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. Again, the eighth and ninth channels are very far apart and so, if the link resides on one of those channels and the jammer is on the other one, then the link 's performance is not deteriorated. Note here that, if the link were to operate on any of the channels 1, 8, 9 or 12, the jammer would only impact the link if it resides on the same channel or the immediate 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:

[8] 802.11a channels are 36, 40, 44, 48, 52, 56, 60, 64, 149, 153, 157, 161 in North America.

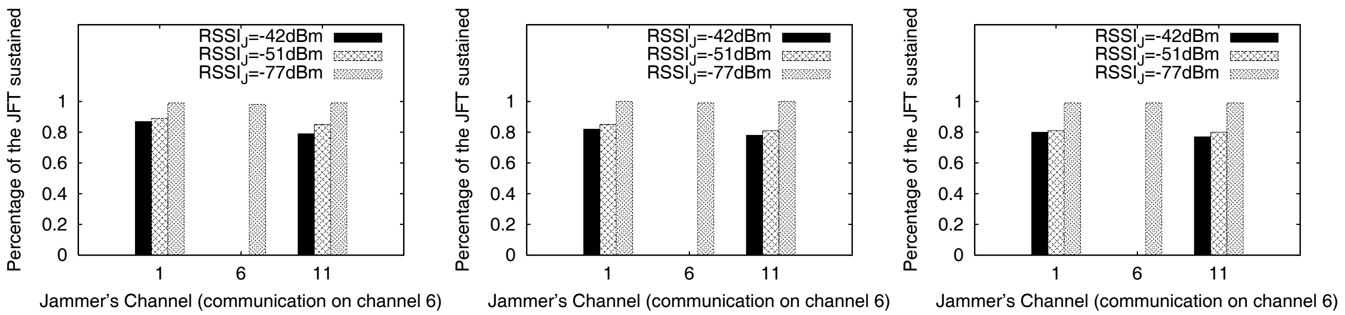

Fig. 4. Percentage of the jamming free throughput (JFT) achieved when the jammer is on various channels, and for various $RSSI_J$, for the case of 802.11g. In the three figures we have $RSSI_l = -39dBm$, $RSSI_l = -45dBm$ and $RSSI_l = -68dBm$, respectively.

TABLE I
MIXED STRATEGY FOR ONE JAMMING DEVICE IN 802.11A

channel j	1	2	3	4	5	6	7	8	9	10	11	12
y_j^*	.0894	.1155	0	.1016	.1016	0	.1155	.0894	.1728	.0207	.0207	.1728

TABLE II
MIXED STRATEGY FOR THE COMMUNICATION LINK IN 802.11A

channel i	1	2	3	4	5	6	7	8	9	10	11	12
x_i^*	.1910	0	.026	.0894	.0894	.0260	0	.191	.1728	.0207	.0207	.1728

$$A^{1,a}_{i,j} = \begin{cases} 0 & \text{if } i = j, \\ 1 & \text{elseif } (i = 8 \text{ and } j = 9) \text{ or } (i = 9 \text{ and } j = 8), \\ 0.12 & \text{elseif } |i - j| = 1, \\ 1 & \text{otherwise.} \end{cases}$$

We can now use the linear programs (1)-(4) and (5)-(8) in order to compute equilibrium 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 shown in Tables I and II.

The strategy y^* gives the probability distribution as per which the jammer should choose the next channel to hop. We show that the equilibrium strategy for the jammer is unique. For the link, x^* is one possible equilibrium probability distribution according to which the next channel can be chosen; however, it is not unique. If the players play as per these equilibrium strategies, the value of the game is $v = 0.809$. This implies that the expected throughput on the link is about 81% of its jamming-free throughput.

Uniqueness: The following corollaries formally prove that (i) the jammer's equilibrium strategy is unique and, (ii) the link should not use channels 2 and 7.

Corollary 1: The linear program (5)-(8), with $A = A^{1,a}$, has just one optimal solution $y = y^*$, where y^* is given in Table I.

Proof: We prove the corollary by contradiction. Let there be another 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:

$$\begin{aligned} \text{maximize} \quad & |\hat{y} - y^*| && (9) \\ \text{subject to} \quad & A\hat{y} \leq 0.809 && (10) \\ & |\hat{y}| = 1 && (11) \\ & \hat{y} \geq 0 && (12) \end{aligned}$$

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 $i = 1, \ldots, 12$,

$$\begin{aligned} \text{maximize} \quad & \hat{y}_i - y_i^* && (13) \\ \text{subject to} \quad & A\hat{y} \leq 0.809 && (14) \\ & |\hat{y}| = 1 && (15) \\ & \hat{y} \geq 0 && (16) \end{aligned}$$

$$\begin{aligned} \text{maximize} \quad & y_i^* - \hat{y}_i && (17) \\ \text{subject to} \quad & A\hat{y} \leq 0.809 && (18) \\ & |\hat{y}| = 1 && (19) \\ & \hat{y} \geq 0 && (20) \end{aligned}$$

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

Corollary 2: Any equilibrium strategy x^* for the maximizing player (the link) has $x_2 = x_7 = 0$.

Proof: To prove that in any optimal solution, $x_2 = x_7 = 0$, we formulate the following linear program.

$$\begin{aligned} \text{maximize} \quad & x_2 + x_7 && (21) \\ \text{subject to} \quad & A^T x \geq 0.809 && (22) \\ & |x| = 1 && (23) \\ & x \geq 0 && (24) \end{aligned}$$

The linear program tries to find the maximum value for the sum $x_2 + x_7$ under the constraint that the achieved payoff is at least 0.809 (this is the maximum achievable payoff). The

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

# jammers	1	2	3	4	5
v	80.9%	61.8%	42.7%	23.6%	4.5%

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

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

solution to the above linear program yields an objective value of zero. In other words, there cannot be any optimal solution with either $x_2 \neq 0$ or $x_7 \neq 0$. ■

Corollary 3: If the jammer plays the strategy of corollary 1, then the link player can set $x_1 + x_3 + x_4 + x_5 + x_6 + x_8 + x_9 + x_{10} + x_{11} + x_{12}$ to any non-negative value, as long as their sum is 1.

Proof: The value of the game is $x^T Ay$. Substituting $A = A^{1,a}$ and $y = y^*$ we have:
$v = x^T Ay = 0.809(x_1 + x_3 + x_4 + x_5 + x_6 + x_8 + x_9 + x_{10} + x_{11} + x_{12}) + 0.8059(x_2 + x_7)$
In order to maximize v we should set $x_2 = x_7 = 0$, and then set the remaining variables to any non-negative values such that $x_1 + x_3 + x_4 + x_5 + x_6 + x_8 + x_9 + x_{10} + x_{11} + x_{12} = 1$. ■

Recall that the solution x^*, provides the best response strategy of the communication link to the strategy y^* of the jammer (and vice versa). The set of channels available can be separated into two disjoint sets in terms of interference, that is, channels 1-8 and 9-12. In the first subset, the jammer picks channels 2 and 7 with the highest probability, since it can then block a set of 3 channels that cannot be simultaneously blocked otherwise. As a result, the link should avoid these channels (i.e., $x_2 = x_7 = 0$) and place its device with high probability on the *edge* channels (i.e., 1, 8, 9 and 12). In the second subset, the jammer picks the edge channels with higher probability, since it then can effectively block channels 9-12. Note here that, if we were to compare the probabilities with which the edge channels are occupied by the link, we have $x_9 = x_{12} < x_1 = x_8$, because $y_9 = y_{12} > y_1 = y_8$.

1) Multiple jammers: We consider the scenario where the jammer can employ more than one jamming device, that is, it can block more than one channel. This case of multiple jammers can still be modeled as a zero-sum two-player game and described by a matrix A_{ij}. Here i is the channel on which the link resides and j represents the channels where the jamming devices reside. In order to reduce the dimension space due to the multiple jamming devices, we use a row/column major order representation. As an example, let us consider the case of two jamming devices on channels j_1 and j_2. There are $12^2 = 144$ possible placements of these devices on the frequency spectrum. Each placement can be encoded by a single value j. It is easy to see that by setting $j = 12(j_1 - 1) + j_2$, every combination of j_1 and j_2 is encoded into a unique value.

Table III summarizes the expected percentage of the jamming-free throughput in equilibrium for the case of one, two, three, four and five jamming devices.

It is straightforward to extend Corollary 2 and Corollary 3 for the multiple jammer cases. Thus, x^* given in Table II is an equilibrium strategy for any of the cases. The jammer's equilibrium strategy is no longer unique but still y_3 and y_6 are 0. Moreover, it makes no sense to put multiple jamming devices on the same channels.

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).

Note here that, our framework can be extended (in a similar way) to account for multiple communication links (i.e., *maximizing* players). Again, a major row/column representation will be used for the second dimension (i.e., the one representing the links' strategies) of the payoff matrix. Note that the number of rows will increase as well.

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:

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

For two jamming devices the payoff matrix is given by

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

Note here that interestingly, our measurements indicate that for a link that is being (partially) jammed by a jammer residing on an adjacent channel, adding one more jammer on the other adjacent orthogonal channel does not further impact the link's throughput (as it does in the case of 802.11a). This can be attributed to the relatively large 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 IV.

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

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

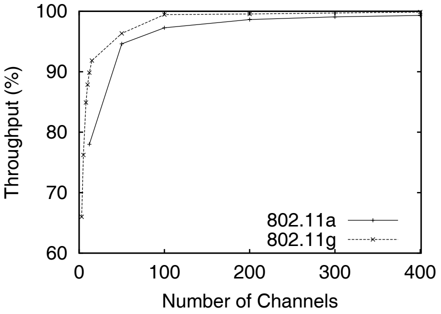

Fig. 5. Increasing the spectrum availability, significantly increases FH's robustness against jamming.

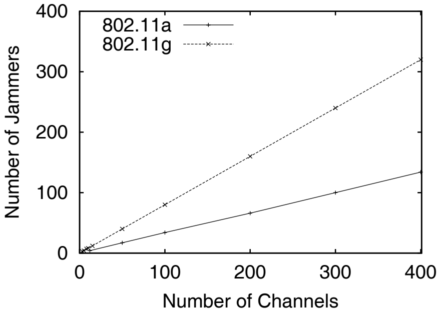

Fig. 6. Number of jammers needed to drop throughput below 20% of the jamming free performance enjoyed.

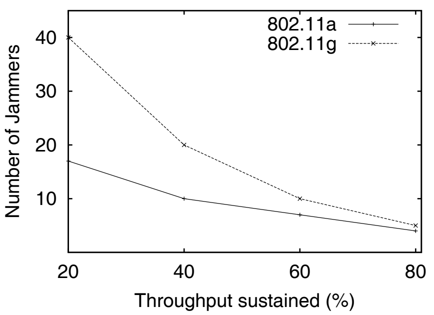

Fig. 7. Number of jammers needed to drop the throughput at a specific percentage (50 channels).

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

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

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

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

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

channel i	1	2	3
x_i^*	0.3333	0.3333	0.3333

among the three channels, uniformly at random. The strategies are shown in tables VI and VII.

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.

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

Proof: We prove the corollary 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 V. An optimal solution $y = y^*$ given by Table V makes all the constraints tight i.e.,

$$0.88y_2 + y_3 = v \tag{25}$$

$$0.88y_1 + 0.88y_3 = v \tag{26}$$

$$y_1 + 0.88y_3 = v \tag{27}$$

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 V 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[9] 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 Fig. 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

[9] For ease of presentation, here we assume that the central frequencies of all the channels are 20MHz apart. Although this might not be true (i.e. the cases of channels 64 and 149 with 802.11a that are 425MHz apart) it only affects the results by a negligible factor if the number of such pairs is small compared to the total number of channels.

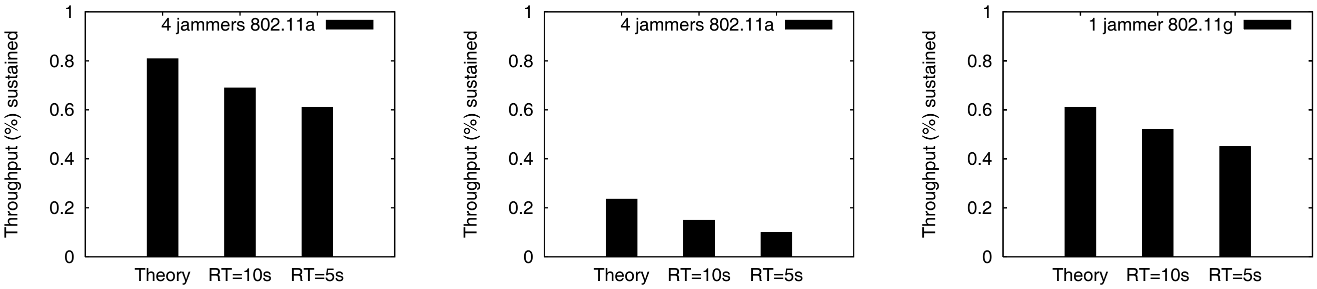

(a) 1 jammer with 802.11a.

(b) 4 jammers with 802.11a.

(c) 1 jammer with 802.11g.

Fig. 8. Experimenting with our prototype proactive FH. Our framework, indeed, bounds the performance of FH as jamming countermeasure.

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 Fig. 7 we present the number of jamming devices needed in order to drop the throughput of the link 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).

In summary, as one might expect, our results suggest that if current systems could support a larger number of orthogonal bands, frequency hopping has the potential of being a robust anti-jamming technique.

From a different point of view, we are interested in examining the effect of one or multiple jammers in a scenario where two adjacent orthogonal channels are completely isolated (i.e., $A_{i,j} = 0$). Such scenarios can exist if we were able to (i) reallocate the available bandwidth in such a way that adjacent orthogonal channels are isolated (which would result in fewer channels as compared to current systems), or (ii) use additional resources/bandwidth and assure that the frequency bands used do not interfere with each other. These results, can provide useful guidelines for future frequency allocation policies that are resilient to jamming attacks.

Figure 9(a) depicts the sustainable throughput for different number of jamming devices versus the number of isolated frequency bands. As one might expect, increasing the number of isolated frequency bands, causes frequency hopping to be more robust to jamming attacks. As an example, with 100 isolated channels, even under the presence of 10 jammers the sustainable, jamming-free, throughput is as high as 90%. Furthermore, it is interesting to notice, that if we were to reallocate/reassign the 5 GHz band in such a way that there is a 20MHz spacing between the channels (which results in 6 orthogonal bands), the sustained throughput with one jammer is 83%, with 2 jammers is 66% and with 4 jammers is 33%. All these values are higher than the corresponding values with the current 12 channel allocation (i.e., 80.9%, 61.8%, 23.6%).

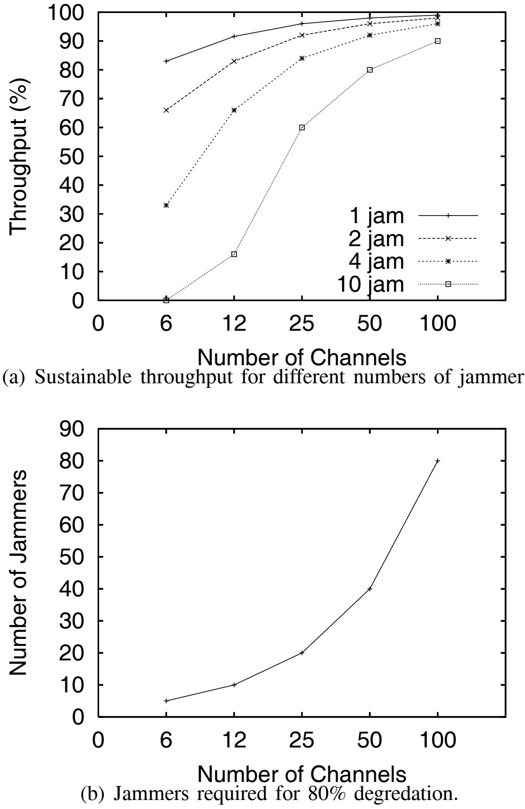

(a) Sustainable throughput for different numbers of jammers.

(b) Jammers required for 80% degredation.

Fig. 9. Isolated orthogonal channels.

Finally, Fig. 9(b) presents the number of jamming devices required for a sustained throughput of at most 20% as compared with the jamming free environment. As with current channel allocations, the number of jammers required increases as the number of available, isolated bands increases.

D. Validation Of Our Framework

In this section we build a proof of concept prototype of a proactive frequency hopping scheme. Note that our goal is to validate the performance bounds that were theoretically computed in the previous section and *not* the implementation of a *full fledged* distributed implementation of a frequency hopping technique.

1) System design and implementation: Our system implements a simple, generic proactive frequency hopping scheme. The scheme is based on the game described in the previous section. In particular, the network nodes switch between the available frequency bands, once every k seconds. The hopping sequence is known by all network nodes, but not by the jammer. This is achieved by an offline computation of the hopping sequence by using the linear programs from the previous section and a priori loading of the computed sequence on all the nodes in the network. A similar procedure is followed for the jammer's hopping sequence. An offline emulation of the sampled frequencies demonstrated that the system converges after approximately 70 frequency hops. Accordingly, we create various sequences of 100 frequencies each and experiment with them.

An important design parameter is the *residence time* of a node on the channel (denoted RT from now on). RT is defined to be the time that a node spends on a channel prior to hopping to a different channel. In the first set of experiments described in this section, we use *fixed* RT values of 5 and 10 seconds for both the jammer(s) and the link. Optimizing the RT is beyond the scope of this paper. However, we experimentally study a plurality of scenarios where the jammer and the link use different RT values and discuss the implications thereof later in this section.

All nodes are synchronized using the Network Time Protocol (NTP) [29] through our testbed server. Thus, all nodes share the same clock and hop between the channels simultaneously. The hopping is implemented using the `ioctl()` [30] interface. The delay that ioctl() interface imposes is of the order of μsec [31] [32], and as a result the overall performance is not affected. The reader should also recall, that implementing a *professional*, proactive frequency hopping scheme is beyond the scope of this work, as mentioned in the beginning of this section.

2) Experiments with 802.11a: We perform experiments on several 802.11a links with jammers in their vicinities. Both the link being considered and the jammer, hop frequencies as per the equilibrium schedule (as discussed earlier). In particular, we conduct experiments with: **(a)** 40 different links on our testbed and, **(b)** 30 different equilibrium hopping sequences. Each of these hopping sequences consist of 100 sequential frequency hops for both the link under consideration and the associated jammer. The hopping sequences are samples generated with the probabilistic distributions from the output of our game theoretic framework, **(c)** 1, 2, 3 and 4 jammers active at a time, **(d)** $RT = 5sec$ and $RT = 10sec$. Note that in all our experiments we have used the Sample rate algorithm [33] (the default settings).

The results from our experiments with one active jammer are shown in figure 8(a). We observe that in practice, the throughput achieved in the presence of a jammer with a proactive frequency hopping strategy is lower than what is theoretically expected. This is because the model used in section III assumes *zero dwell times* between the channel hops, and perfect synchronization. Neither of these assumptions holds in a real deployment. Furthermore, note that the throughput is lower due to a higher switching[10] and synchronization overhead if $RT = 5sec$ as compared to the case where $RT = 10sec$. In practice there is never perfect synchronization, even with NTP.

We experimented with 2, 3 and 4 jammers with similar results. In figure 8(b) we present the results for 4 jammers. We notice again that in practice the performance is poorer as compared to what is theoretically expected. In particular, with 4 jammers the throughput achieved is *only 8-10%* of the jamming free throughput.

3) Experiments with 802.11g: We report experiments with only one jammer with 802.11g. Our experiments suggest (as one might expect from our analysis) that the performance degrades significantly with 2 jammers and with 3 jammers the entire spectrum is blocked. As with 802.11a, we compute the equilibrium hopping sequences for both the link and the jammer, and experiment with two different values of RT. The hopping sequences were again of length 100. As previously, it was verified offline that 100 hops were enough for the game to converge to its optimal value. The results are shown in figure 8(c). As with 802.11a, we observe that the performance in practice is lower than what is theoretically expected (due to the same reasons as before).

4) The sensitivity to the choice of RT: Our framework provides long term performance bounds and as a result, by itself does not yield insights on the right choice of the value of RT (for either the link or the jammer). Computing the optimal value for RT is beyond the scope of this work. However in our experiments we provide results when the link and the jammer have different values for this parameter (RT_L and RT_J, respectively).

In figures 10(a) and 10(b) we present the results of our experiments with 802.11a for the case of a single jammer. First, in figure 10(a), we hold the RT for the link fixed at $20sec$. The RT of the jammer is varied. Reducing the RT value of the jammer can have two conflicting effects. On the one hand, the jammer can hit multiple channels during the $20sec$ RT period of the link; this can increase its effectiveness. On the other hand, it might incur a switching penalty each time it switches channels. We observe that when the RT of the jammer is reduced from $20sec$ to $15sec$, the first factor has a higher impact; however, further reducing the value of RT causes the second factor to be dominant. A similar behavior is observed when we keep $RT_J = 20sec$ and we vary RT_L. The sweet spot again appears when $RT_L = 15sec.$

We wish to point out that irrespective of the choice of RT, the practical schemes cannot do better than what is theoretically predicted by our framework in the long term. Our framework is independent of the RT of each player and the potential switching penalty (note that our analysis implicitly assumes zero switching penalty). Thus, although the performance of a frequency hopping strategy might be improved by tuning the frequency with which the link switches between channels, it is still limited and cannot provide better performance in the long run, than what is predicted by our framework.

VII. Experimenting with 802.11n.

The use of antenna arrays or MIMO (multi-input multi-output) technology promises higher reliability; the 802.11n standard supports transmissions on MIMO links. In this sec-

[10] Note that with appropriate driver/firmware modifications - specific to the hardware in use - one can make this penalty extremely small.

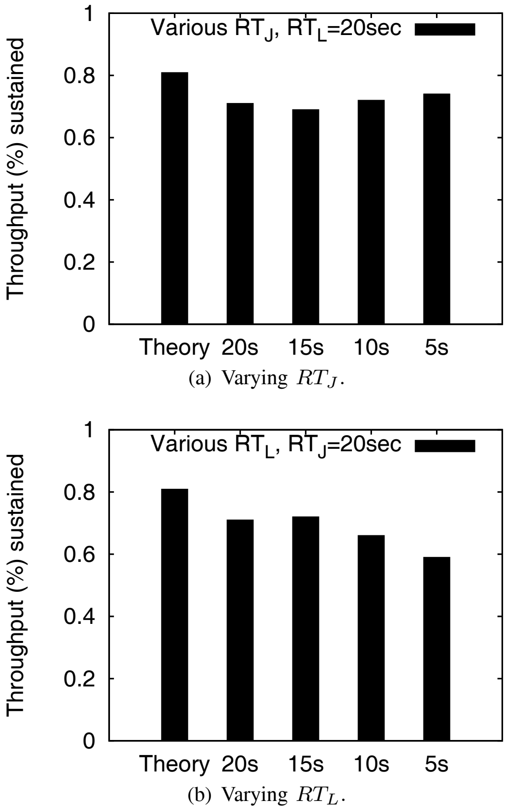

(a) Varying RT_J.

(b) Varying RT_L.

Fig. 10. Validation of our framework.

tion, our objective is to evaluate the efficacy of frequency hopping in 802.11n networks against jamming attacks.

A technique that is exploited to allow transmissions at higher rates with 802.11n is channel bonding. In a nutshell, we find that channel bonding makes frequency hopping less effective with regards to jamming attacks. We begin this section with a brief overview of channel bonding; subsequently we apply our game theoretic framework and evaluate the performance of 802.11n in the presence of a jamming attack.

A. Channel Bonding

802.11n devices can operate on channels that span either $20MHz\ or\ 40MHz$ bandwidth. In the latter case, channel bonding is used [10]. With channel bonding, two or more adjacent channels are used in conjunction to form a new *wider* channel. The expansion helps achieve higher data rates (practically doubles the possible rate). The thesis is that, the increased reliability possible on MIMO links (due to diversity and the use of space time codes) [34] can support transmissions at higher rates[11]. To elucidate the concept of channel bonding, consider channel 6 (as specified with 802.11g). Without channel bonding, the 802.11n signal utilizes the spectrum between $2.427MHz$ and $2.447MHz$. However, with channel bonding the spectrum that is used spans the frequencies between $2.417MHz$ and $2.457MHz$.

At this point we should note that 802.11n systems employ carrier sensing for medium access. This makes them susceptible to interference due to collocated links operating at the same or overlapping frequencies [28]. Consequently, 802.11n systems cannot take full advantage of the benefits of the underlying PHY layer technology (e.g. interference cancellation, support of simultaneous multiple transmissions etc).

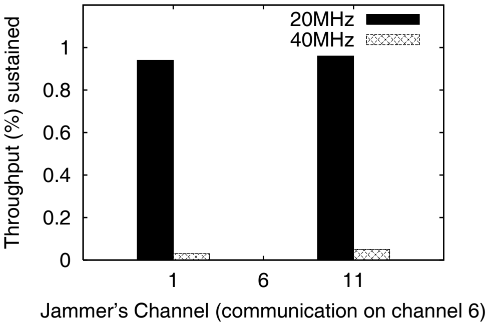

Fig. 11. Channel bonding can degrade MIMO performance under jamming.

B. MIMO Performance Under Jamming

As mentioned, MIMO links with Space-Time Block Codes (STBC) are expected to provide robustness to signal variations. Thus, the required SINR for achieving a target bit error rate is expected to be lower than the corresponding requirement with SISO (Single-Input Single-Output) links[12].

For our experiments we use Ralink's RT2860 chipset, which supports 802.11n communications [35]. These cards operate in the $2.4GHz$ band. We used channels 1, 6 and 11 for our experiments; these are essentially, the only orthogonal channels in this band. We experimented with 40 MIMO STBC links on our testbed, each of which was under the influence of a jammer. Our experiments include both the cases of 20 and $40MHz$ bandwidth for the link, while the jammer uses a bandwidth of $22MHz$ (802.11g mode). Figure 11 depicts the results from our experiments. We only present the case where the communication is taking place on channel 6; other cases yielded very similar results.

From figure 11 we observe that the performance of 802.11n in the case where a bandwidth of $20MHz$ is used, is almost the same as that with 802.11g. 802.11n seems to offer a slightly better performance than 802.11g; an adjacent orthogonal jammer degrades the performance by only 5%. This can be due to two reasons: (a) MIMO links and STBC offer marginally better robustness to the jammer than SISO and (b) the secure zone with 802.11n is $5MHz$ as compared to $3MHz$ with 802.11g (with 802.11n the bandwidth is $20MHz$ while with 802.11g, it is $22MHz$).

The results with channel bonding show that the effectiveness of the jammer is dramatically increased in this case. The reason for this is that channel bonding practically eliminates orthogonality. Even if a jammer operates on a band of $22MHz$ (as in our experiments) and is active on the furthest channels (i.e., channels 1 or 11) from that of the link (channel 6) there is an overlap; in other words, the jammer's signals interfere

[11] With SISO, the higher the transmission rate, the lower the reliability.

[12] Due to CSMA/CA though, such benefits might become obsolete if the transmitter can sense the jamming signals.

TABLE VIII
EXPECTED LINK THROUGHPUT FOR 802.11N WITH $20MHz$ BW.

# jammers	1	2	3
v	64.26%	32.12%	0%

TABLE IX
EXPECTED LINK THROUGHPUT FOR 802.11N WITH $40MHz$ BW.

# jammers	1	2	3
v	5.06%	1%	0%

with the link. The link is safe only when it operates on channel 1 and the jammer occupies channel 11 and vice versa.

Based on the above measurements, we use the framework presented in section III to quantify the performance of a proactive frequency hopping strategy with 802.11n. Applying the model to the case where a bandwidth of $20MHz$ is used yields table VIII, while for the case where a $40MHz$ bandwidth is used, we get table IX.

The results suggest that while immensely useful in terms of increasing the data rates under benign conditions, channel bonding can increase the vulnerability of a frequency hopping technique to jamming. More importantly, we observe that the limitations of frequency hopping as a jamming mitigation technique carry over to 802.11n networks.

VIII. Discussion

Our game theoretic framework can be applied with other variants of a jamming attack. As an example Xu *et al* [36] introduce the random and the reactive jamming model. With the former, the jammer transits between active and idle periods. Each of these periods follows a random distribution. A reactive jammer, senses the medium for ongoing communications, and whenever there is a legitimate packet on the air it jams the medium. The model presented in this paper can be applied to account for these jamming strategies as well. In the following, we will present its application for the case of a random jammer.

Let us assume that the jammer picks its active periods T_a from a uniform distribution U$[a, b]$ secs and its idle period T_i from the uniform distribution U$[c, d]$ secs. Thus, the average active and idle times are $E[T_a] = \frac{a+b}{2}$ and $E[T_i] = \frac{c+d}{2}$ respectively. Consequently, the effectiveness of a random jammer is reduced by a factor α as compared to the case of the constant jammer, where:

$$\alpha = \frac{E[T_a]}{E[T_a] + E[T_i]} \qquad (28)$$

Incorporating this factor, the corresponding payoff matrix for a single jamming device is now given by the equation at the top of the next page. Solving the game using the above payoff matrix will provide us with the solutions that correspond to the random jamming model.

In this work, we have mainly focused on *proactive* frequency hopping strategies for both the communication and the jamming. The reactive jamming case is more complicated. The efficacy of a reactive jammer is affected by a number of factors. As examples, the speed with which the medium is sensed, the ability to sense transmissions taking place on an adjacent orthogonal channel etc., affect the performance of the malicious device. In order to apply our framework, all these parameters need to be accurately modeled and measured[13]. However, once their effects have been quantified, our framework can be used as a *black box* to capture the interactions between the reactive jammer and the communication link. For instance, it is clear that if the communication detection time is negligible, the reactive jammer can be very effective (i.e., the link throughput is almost nulled). Reactive jamming strategies are not widely deployed since they require special expertise from the attacker [36]. Nevertheless, their intelligence can further reduce the network throughput. In this sense, our model provides an upper bound on the performance of proactive frequency hopping as anti-jamming technique.

IX. Conclusions

In this paper we seek to examine the effectiveness of FH as anti-jamming technique. 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 further validate our analytical results through experimentation with a prototype proactive FH scheme. 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. Finally, specific features of 802.11n, that is, channel bonding and carrier sensing, make it more susceptible to jamming attacks as compared to legacy systems, reducing further the efficacy of FH.

Acknowledgment

We would like to thank Ralink Technologies for providing us the Linux source driver for the RT2860 AP and Dr. K. Papagiannaki from Intel Research, for providing us with the prototype firmware of *ipw2200* AP. We would also like to thank the editor, Prof. W. Lou, and the anonymous reviewers for their valuable comments that helped us improve the quality of our manuscript.

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.

[13] Note here that, modeling and/or measuring the extent to which, the above factors affect the reactive jammer's performance is out of the scope of our work.

$$A_{i,j}^{1,a,rand} = \begin{cases} (1-\alpha) & \text{if } i = j, \\ 1 & \text{elseif } (i = 8 \text{ and } j = 9) \text{ or } (i = 9 \text{ and } j = 8), \\ (1 - 0.88 \cdot \alpha) & \text{elseif } |i - j| = 1, \\ 1 & \text{otherwise.} \end{cases}$$

[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] B. O'hara and A. Petrick, *IEEE 802.11 Handbook, a Designer's Companion*, IEEE Press, Second Edition, ISBN 0-73-814449-5.

[11] 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.

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

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

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

[15] 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.

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

[17] 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.

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

[19] H. Li and Z. Han, "Dogfight in Spectrum: Jamming and Anti-Jamming in Multichannel Cognitive Radio Systems," in *IEEE GLOBECOM*, 2009.

[20] R. W. Thomas, R. S. Komali, B. J. Borghetti, and P. Mahonen, "A Bayesian Game Analysis of Emulation Attacks in Dynamic Spectrum Access Networks," in *IEEE DySPAN*, 2010.

[21] 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.

[22] C. Cheng, P. Hsiao, H. 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.

[23] 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.

[24] 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.

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

[26] Soekris-net4826, "http://www.soekris.com/net4826.htm."

[27] "The MAdWiFi driver," http://madwifi.org.

[28] K. Pelechrinis, I. Broustis, T. Salonidis, S. V. Krisnamurthy, and P. Mohapatra, "Design and Deployment Considerations for High Performance MIMO Testbeds," in *WICON*, November 2008.

[29] "SNTP, Version 4," http://www.apps.ietf.org/rfc/rfc2030.html.

[30] "ioctl() man page," http://linux.die.net/man/2/ioctl.

[31] V. Raisinghani and S. Iyer, "Architecting Protocol Stack Optimizations on Mobile Devices," in *Cosmoware*, 2006.

[32] V. Navda, A. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "MobiSteer: using steerable beam directional antenna for vehicular network access," in *MobiSys*, 2007.

[33] J. Bicket, "Bit-rate Selection in Wireless Networks," in *MS Thesis, Dept. of Electr. Engin. and Comp. Science, MIT*, 2005.

[34] H. Jafarkhani, *Space-Time Coding, Theory and Practice*. Cambridge University Press, 2005.

[35] "RT2860 driver," http://www.ralinktech.com/ralink/Home/Support/Linux.html.

[36] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks," in *ACM MOBIHOC*, 2005.

Konstantinos Pelechrinis received his Ph.D. from the Computer Science Department of the University of California, Riverside, in 2010. Previously he obtained his M.Sc. degree from the Computer Science Department of the University of California, Riverside in 2008 and the diploma of Electrical and Computer Engineering from the National Technical University of Athens, Greece, in 2006. He is an Assistant Professor at the SIS faculty of the University of Pittsburgh since the Fall of 2010. He has also held research positions at LANL, Thomson Research Labs Paris and MSR Cambridge. He was a visiting researcher at the University of Thessaly during Fall 2008. His research interests include wireless networking, especially security - related issues that span the full protocol stack. He is involved in protocol design, real world experimentation and performance analysis. He is also interested in mathematical foundations of communication networks.

Christos Koufogiannakis received the diploma in Electronics and Computer Engineering from the Technical University of Crete, Chania, Greece in 2004, and the M.Sc. and Ph.D. degrees in Computer Science from the University of California, Riverside in 2007 and 2009, respectively. During the Winter of 2010, he was a Postdoctoral Researcher at the University of California, Riverside. Since summer 2010 he joined KLA-Tencor Corporation as a Software Engineer. His research interests include design and analysis of efficient approximation and distributed algorithms for combinatorial optimization problems.

Srikanth V. Krishnamurthy received his Ph.D. degree in electrical and computer engineering from the University of California at San Diego in 1997. From 1998 to 2000, he was a Research Staff Scientist at the Information Sciences Laboratory, HRL Laboratories, LLC, Malibu, CA. Currently, he is a Professor of Computer Science at University of California, Riverside. His research interests are primarily in wireless networks, network security and Internet technologies. Dr. Krishnamurthy is the recipient of the NSF CAREER Award from ANI in 2003. He has also co-edited the book *Ad Hoc Networks: Technologies and Protocols* published by Springer Verlag in 2005. He served as the editor-in-chief for ACM MC2R between 2007 and 2009 and is a senior member of the IEEE.