Mitigating Malicious Interference via Subcarrier-Level Radio Agility in Wireless Networks

Ahmed Osama Fathy Atya*, Azeem Aqil*, Shailendra Singh*, Ioannis Broustis*,
Karthikeyan Sundaresan† and Srikanth V. Krishnamurthy*
*University of California, Riverside, †NEC Labs, USA
{*afath001, aaqil001, singhs, broustis, krish*}*@cs.ucr.edu, karthiks@nec-labs.com*

***Abstract*—Malicious interference injection or jamming is one of the simplest ways to disrupt wireless communications. Prior approaches can alleviate jamming interference to a limited extent; they are especially vulnerable to a reactive jammer i.e., a jammer that injects noise upon sensing a legitimate transmission or wideband jamming. In this paper, we leverage the inherent features of OFDM (Orthogonal Frequency Division Multiplexing) to cope with such attacks. Specifically, via extensive experiments, we observe that the jamming signal experiences differing levels of fading across the composite sub-carriers in its transmission bandwidth. Thus, if the legitimate transmitter were to somehow exploit the relatively unaffected sub-carriers to transmit data to the receiver, it could achieve reasonable throughputs, even in the presence of the active jammer. We design and implement *JIMS*, a Jamming Interference Mitigation Scheme that exploits the above characteristic by overcoming key practical challenges. Via extensive testbed experiments and simulations we show that JIMS achieves a throughput restoration of up to 75% in the presence of an active jammer.**

I. Introduction

Wireless communications can be easily disrupted by malicious injection of interference, aka jamming. Given the commercial availability of jamming devices today [1], [2], [3], mounting Denial-of-Service (DoS) attacks using jamming is an easy task.

How easy is it to combat jamming? Previous efforts have tried to mitigate jamming by tuning several physical layer knobs. Examples include adaptive power and rate control, or the use of lower modulation rates in order to reduce the packet error rates (PER) [4], [5], [6] in the presence of jamming interference. Frequency hopping has also been considered in cases where there is significant additional available bandwidth for use [7], [8]. All of these prior studies conclude that in general, it is very difficult to overcome the impact of active jamming, especially when jammers account for the inherent properties of MAC layer protocols [9]. Our extensive testbed measurements using legacy WiFi devices as well as programmable wireless boards [10] support such an argument.

Our measurements however, also reveal a new, promising dimension for malicious interference avoidance in OFDM (Orthogonal Frequency Division Multiplexing) settings [11]. Specifically, we identify a feature that can be exploited with OFDM to mitigate jamming; more importantly, this can be applied in conjunction with most previously proposed anti-jamming schemes.

Exploiting an intrinsic aspect of OFDM signal propagation: OFDM is currently a widely adopted transmission scheme in many different wireless network technologies (e.g., LTE [12], WiMAX [13] and 802.11 [14]). In traditional OFDM implementations, the transmission power is uniformly distributed across a predefined set of frequency subcarriers; the number and width of these subcarriers dictates the available channel bandwidth [11]. Due to physical obstructions and interference, signal power (even that of a jammer) undergoes different levels of fading across the different subcarriers. As a result, on some of the subcarriers the received jamming signal strength can be high, while on other subcarriers it is likely to be low[1] [15].

Employing subcarrier-level radio agility: Our testbed measurements also indicate that jamming signals are likely to experience varying levels of fading on different OFDM subcarriers. As a result, some subcarriers may not be "significantly affected" by the malicious power emission; such "cleaner" portions of the available spectrum could be temporarily used for legitimate packet transmissions, as long as a transceiver pair is made aware of which those subcarriers are.

Thus, we design and implement a framework that allows a transceiver pair to exchange information that reveals the "clean" subcarriers in the available spectrum, where the jamming signal experiences significant fading. Once such subcarriers are identified, we pool power onto them (to the extent allowed), and utilize them for packet transmissions to increase the probability of successful packet delivery and thereby the long-term throughput (while being actively jammed). More specifically, our contributions in this paper are the following:

1) Experimental characterization of jamming interference: We perform a large set of testbed measurements using WARP reference boards in order to observe the impact of fading on jamming transmissions, across different OFDM subcarriers in a spectral band. We experiment on different network topologies and with various jamming patterns. We validate our hypothesis that there may be portions of the spectrum where the jamming signal experiences deep fading.

2) Design of our subcarrier-level radio-agile anti-jamming

This work was supported by the US Army Research Office under the Multi-University Research Initiative (MURI) grant W911NF-07-1-0318, as well as the NSF NeTS grants 1017012 and 1016748. 978-1-4799-1270-4/13/$31.00
©2013 IEEE.

[1]As discussed in section III, we focus on the case of OFDM jammers due to the difficulty in detecting their presence [5].

framework: We design a framework that enables a pair of legitimate transceiver pair, say Alice and Bob, to exchange information regarding which subcarriers should be used for packet transmissions in each link direction (note that the fading patterns for the jamming signal will differ at Alice and at Bob). For this, we leverage *raptor* codes [16] to securely and efficiently exchange information on the "clean" sub-carriers. Subsequently, we design an algorithm (executed at each transceiver) that considers the subcarrier-specific SINR and the expected number of packet retransmissions, in order to make subcarrier-level transmission decisions, such that the long-term user throughput is maximized. These components constitute our jamming interference mitigation scheme, JIMS.

3) Implementation and evaluation of our framework: We implement JIMS on the WARP platform [10] and evaluate its efficiency via extensive experiments in the 2.4 GHz ISM band. We involve legacy WiFI nodes in many of our experiments. We also implement our framework on the NS3 simulator in order to observe its efficacy in large-scale wireless network settings. Both our experiments and simulations incorporate diverse jamming patterns that reflect a large number of settings. We observe that the application of JIMS can achieve a user throughput restoration of up to 75% and network-wide throughput improvements that range between 30% and 75%.

II. Background and Related Work

In this section, we first provide the relevant background on jamming attacks. Subsequently we discuss previous related studies on anti-jamming and differentiate our work.

Malicious interference injection: As discussed in section I, various types of jamming devices are readily available in the market today [1], [2], [3]. Although initial models were very simple in their operation (i.e., they were simply emitting energy all the time), newer devices have incorporated intelligent power emission patterns, in order to conserve battery power and avoid detection. More specifically, jammers can emit power continuously or intermittently. Intermittent jammers are further categorized based on the duration of the active and inactive time intervals; for example, periodic jammers use fixed durations for these intervals. Moreover, *reactive* jammers act more intelligently, by emitting power only if they overhear traffic; this makes them more energy-efficient and more difficult to detect [17].

Previous related studies: Most of the previous efforts on alleviating or avoiding malicious interference employ frequency hopping, or power and/or bit rate adaptation techniques. With frequency hopping, legitimate users decide on a hopping pattern across the set of available channels in an effort to avoid the jammer [7], [18], [8]. However, frequency hopping techniques cannot avoid jammers that can distribute their power across multiple bands simultaneously [19]. In [5] the authors use power control and bit rate adaptation towards mitigating jamming interference; however, the proposed approaches can exacerbate interference due to increased power levels or inappropriately set carrier-sense thresholds. The proposed bit rate adaptation is only useful if the jammer is silent intermittently. Similar techniques are proposed in [4] and [6]. A survey on jamming attacks and mitigation solutions can be found in [20]. Our approach, unlike the prior approaches leverages OFDM diversity; most importantly, *it can be used in conjunction with many of such techniques*.

Yao et al. in [21] propose a DSSS (Direct Sequence Spread Spectrum) anti-jamming method for broadcast transmissions. The method relies on spreading codes to encode a bit stream of data. However, due to its nature, this work is not applicable in wireless communication systems that are based on OFDM.

III. Sub-carrier Radio Agility aids Anti-jamming

In this section, we describe our testbed experiments on assessing the behavior of malicious interference from the perspective of OFDM sub-carrier level propagation. Our measurements offer insights on how the jamming power is distributed across the subcarriers of the available spectrum. These insights motivate and form the foundation of our radio-agile anti-jamming framework design, which we discuss in section IV.

In a nutshell, our measurement-based, key findings are the following:

- The Received Jamming Signal to Noise Ratio (or **RJSNR**) experienced by legitimate users (transceivers), can often be quite low on some OFDM subcarriers.
- Due to the asymmetry in the perceived RJSNR per subcarrier, a transceiver pair needs to exchange information regarding the subcarriers with respect to which the RJSNR is low, at each end (of the link).
- Due to variations in RJSNR over time, nodes need to periodically send updated channel feedback. A low-overhead feedback frequency of the order of once every 1000 msec suffices in relatively static settings.

In what follows, we describe our threat model and experimental configuration; subsequently we present our observations.

The threat model: We consider a jammer (Eve) that transmits OFDM signals with the same transmission power budget as legitimate users, thereby imitating a typical legitimate device to avoid detection. Other than this, we do not require any other constraint on the jammer.

We also assume that a pair of legitimate transceivers use a shared symmetric key to encrypt the channel state information that they need to exchange. The key may be either preloaded, or derived via an authentication and key agreement protocol (e.g. [22], [23]).

We also do not address the problem of jammer detection; we restrict ourselves to the mitigation of active jamming. We assume that schemes such as those proposed in [24] can be used to distinguish between benign and malicious interference.

Experimental setup: Although our study is generally applicable with any wireless OFDM system, throughout the rest of the paper we particularly focus on measurements in the ISM 2.4 GHz band. We consider a 20 MHz channel in the ISM band (channel 6, centered at 2.437 GHz); the channel consists of 64 subcarriers, 48 of which are used for data transmissions. We perform our experiments late at night in a campus building,

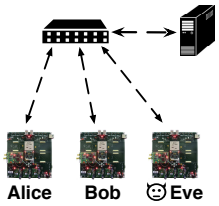

Fig. 1: Alice, Bob and Eve are all placed on a straight line.

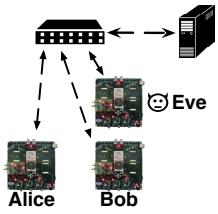

Fig. 2: Eve is placed at 90 degrees to Alice and Bob.

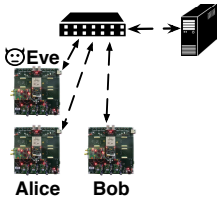

Fig. 3: Eve is placed at 90 degrees to Alice and Bob.

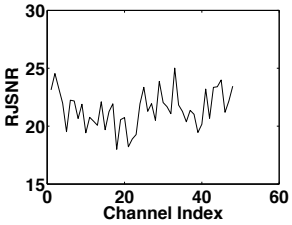

Fig. 4: Alice's perceived RJSNR.

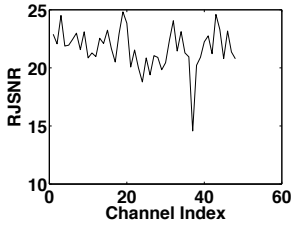

Fig. 5: Bob's perceived RJSNR.

and we verify that this channel is not used by any collocated WLAN networks. Our experimental assessment on the effects of fading on jamming signals involves a pair of legitimate devices (Alice and Bob), and a custom-made jammer (Eve). Alice, Bob and Eve are all stationary nodes that use fixed power budgets. Note that although all nodes operate in the ISM band, for this set of experiments they do not follow the IEEE 802.11 CSMA-CA MAC protocol; instead, Alice transmits packets to Bob as soon as they arrive at her output queue. All three devices are based on WARP programmable boards [10], which are connected to a management server (Fig. 1-3). Each reference board is equipped with a Xilinx Virtex-II Pro FPGA and 4 daughter-boards operating in the ISM band. The OFDM implementation that we use (WARPLab v6) supports BPSK, QPSK and 16 QAM modulation rates, and a 40 MHz sampling rate. Legitimate packets carying CSI information have a length of 240 bytes, while data packets have a length of 1500 bytes; each experiment lasts for 5 minutes and is repeated 20 times.

Experimental insights: Next, we elaborate on specific network configurations and discuss our observations.

i. Jamming signals often experience deep fading on some OFDM subcarriers: We configure the jammer to constantly emit electromagnetic energy on channel 6 (2.437 GHz), and we capture the observed RJSNR at legitimate nodes[2]. A sample of our measurements is depicted in Fig. 4 and Fig. 5, which shows the RJSNR as perceived by Alice and Bob on each of the 48 data subcarriers. We observe that on quite a few of the subcarriers, the RJSNR can be quite low.

This promising observation serves as the main motivation in designing JIMS: *If legitimate transmitters could somehow estimate the subcarrier-level RJSNR values at the receivers, they could simply use only those subcarriers where the RJSNR is low, for packet transmissions*.

ii. The jammer's fading profile differs at each receiver: This is evident from our RJSNR measurements depicted in Fig. 4 and Fig. 5: the perceived RJSNR per subcarrier differs at each legitimate node. This observation is in line with previous studies [12]. Our measurements suggest that *in order for Alice to utilize low-RJSNR subcarriers when transmitting packets to Bob in the presence of Eve, Bob should send reliable channel state information (CSI) feedback to Alice to indicate the subcarriers relatively unaffected by Eve*.

[2]As we discussed earlier, channel 6 was not used by other wireless networks in the neighborhood.

iii. The RJSNR value changes over time on each subcarrier: We perform experiments to observe the variation of RJSNR over time. Similarly as above, in this set of experiments Eve continuously emits electromagnetic energy, while Alice and Bob measure the corresponding RJSNR for her signal. As intuitively expected, due to fading, scattering and power decay, Alice and Bob observe different RJSNR values over time on each subcarrier of the available spectrum. While we have performed measurements with approximately 120 different intervals for sending CSI feedback, in Figures 6a and 6b we plot how the RJSNR values for three specific subcarriers, for two different feedback intervals, i.e., for 2300-msec and 1000-msec, respectively.

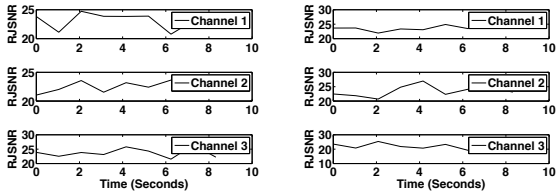

(a) Per subcarrier RJSNR with feedback every 2.3 sec. (b) Per subcarrier RJSNR with feedback every 1 sec.

Fig. 6: Impact of varying the CSI feedback interval

We observe that if the RJSNR information is fed back once per 1000 msec, the intermediate RJSNR variations are captured much more accurately than with the 2300-sec interval. Clearly, the smaller the feedback interval, the higher the probability that a significant variation in RJSNR is captured. In other words, more frequent CSI feedback increases the accuracy in Alice's determination of the subcarriers where the jammer's signal strength is low. On the other hand, as the frequency of feedback messages increases, so does the network overhead. We carefully examine our measurements with different feedback intervals (where each feedback message contains a vector of pointers to subcarriers that should be used by the transmitter, as we discuss in the following section). We conclude that *while the jammer is active, a Channel State Information (CSI) feedback periodicity that limits the network overhead to acceptable extents, while providing an accurate view of the channel is on the order of once every 1000 msec.*

iv. Effect of power allocation on the SINR value: In an interference dominated setting (as in the presence of a jammer), the SINR which is $\frac{P_R}{N+I}$ can be approximated to be $\frac{P_R}{I}$; in

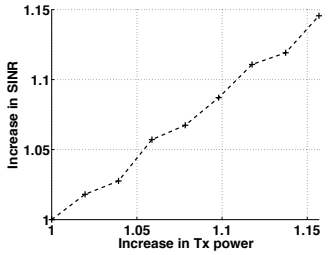

Fig. 7: The effect of increase in transmission power on the SINR in the presence of a jammer.

these expressions, P_R is the received power, I the interference power and N the ambient noise power. P_R is proportional to the transmit power P_T and the channel attenuation, say ζ. In static settings (slow fading), ζ can be expected to be fairly stable over time. If the interference from the jamming signal I is also stable over time (jammer does not move much[3], which is the scenario we consider here), one can expect the SINR to change in proportion to P_T. In other words, if we change P_T to, say, αP_T, the SINR can also be expected to scale by the factor α. We conducted testbed measurements in order to validate this hypothesis. Specifically, Fig. 7 depicts the increase in the SINR, relative to an increase in P_T, in the presence of the jammer. The results demonstrate that our hypothesis holds. We use this observation to determine the gains in SINR with power reallocation with JIMS.

IV. Our Subcarrier-Level Radio-Agile Design

In this section, we describe the design of our jamming interference mitigation scheme (JIMS), which is based on the key observations made in Section III. The scheme consists of three major steps. **First**, the legitimate pair of transceivers independently determine the OFDM subcarriers that are relatively unaffected by the jamming signal. **Second,** by means of using Raptor codes [16], they exchange the information they have determined (CSI) in the first step. **Third**, each transceiver uses this information, to transmit symbols on only an appropriately chosen set of subcarriers (that are relatively unaffected at the receiver).

To maximize the likelihood of correct reception, and facilitate higher transmission rates on the relatively unaffected subcarriers, we further consider an extended version of JIMS, which involves pooling power from the subcarriers that remain unused (to the extent allowed by regulations) to those subcarriers on which, symbols are actively transmitted. We call this extended version of JIMS as JIMS-PA (for Power Allocation).

A. Determining the subcarriers affected by the jamming signal

We consider two ways for detecting the subcarriers that are affected by the jammer. For ease of discussion, let us assume that Alice is executing this step. She simply measures the signal from the jammer when there are no other transmissions in the vicinity. Towards this, we first assume that somehow Alice knows that a jammer is in operation using one of the techniques proposed in [24]. Next, we assume that Alice can simply listen and detect the jamming signal. If the jammer emits energy continuously or without regards to whether or not Alice and Bob are transmitting, this can be done easily. If the jammer is reactive i.e., only transmits upon sensing a transmission from Alice, Alice can send a short pilot to trigger the jammer and subsequently go silent (assuming half-duplex mode of operations as is common with legacy systems); the jamming signal that spills beyond Alice's prompt can then be captured to determine the jammer's profile. The signal can be then decomposed to determine the SNR on each of the subcarriers in the operational band. In other words, the subcarrier level RJSNR can be determined. We call this approach the *explicit approach* of determining the affected subcarriers.

Second, let us assume again that using an appropriate technique from those reported in [24], the presence of the jammer is detected. Bob then sends a pilot signal to Alice. Alice then determines the SINR on each of the subcarriers in that pilot signal. In a nutshell, if either the signal quality is low and/or the jamming signal is high, on a specific subcarrier, that subcarrier is deemed unfit for communication. Other subcarriers where neither of the above scenarios hold true, are appropriate for transmission. We call this approach the *implicit approach* of determining the affected subcarriers.

B. Subcarrier selection

Using either the explicit or implicit approach, Alice is able to determine the quality of communications on each of her subcarriers. Now, she has to determine the appropriate set of subcarriers for use by Bob, for him to communicate with her. The process of selecting this set is different with the explicit and implicit approaches described above.

With the explicit approach, the good subcarriers (to be used for communication by Bob) are chosen based on simple RJSNR threshold. Specifically, if the RJSNR is lower than a certain threshold on a subcarrier it is deemed a *good* subcarrier.

A simple way to choose the RJSNR threshold is to determine average RJSNR from that observed on all subcarriers, and use those that have RJSNRs lower than the average. Specifically, the threshold η is computed to be:

$$\eta = \frac{1}{n}\sum_{i=1}^{n} RJSNR_{c_i} \tag{1}$$

where, $RJSNR_{c_i}$ is the measured RJSNR on subcarrier, i and n is the total number of data subcarriers. Upon computing η, Alice classifies those subcarriers, c_i with $RJSNR_{c_i}$ greater than η to be unsuitable for reception. Later, in Section V, we examine the performance of JIMS with other possible thresholds as well.

With our implicit approach, Alice measures the SINR on each of the subcarriers on a pilot transmitted by Bob. Thus, with this approach, she first computes the average SINR, ξ, by considering all the subcarriers as follows.

$$\xi = \frac{1}{n}\sum_{i=1}^{n} SINR_{c_i} \tag{2}$$

[3]Typically jammers are strategically placed in areas where they can disrupt ongoing communications for a prolonged period of time; this suggests that they do not move frequently.

where, $SINR_{c_i}$ is the SINR with respect to subcarrier, i and n is the total number of data subcarriers. Alice then only chooses those subcarriers, c_i with $RINR_{c_i}$ higher than ξ as the good subcarriers (for Bob to communicate with her).

Choice of the right threshold: One of the challenges that arises with both the explicit and the implicit schemes is "How do we choose the right threshold (be it RJSNR or SINR depending on whether the explicit or implicit approach is used)?" For simplicity, let us just consider the implicit approach; instead of choosing ξ as above, let us assume that we choose a different static threshold ξ'. If we are liberal, and chooose ξ' to be low, we include a large set of subcarriers; however, the SINRs on some of these subcarriers will be unacceptably low. If instead, we are conservative and choose a high value for ξ', we may end up excluding a large number of subcarriers (on a few of which, communications may in fact be possible), and thus, end up achieving a lower throughput than what is possible. We find via experiments that choosing the average value (as discussed above) to be the threshold, provides a good compromise between the two extreme cases, in most scenarios. We evaluate this choice, by comparing the performance with other cases where a static threshold is chosen, in Section V.

C. Exchanging CSI

At this point, both Alice and Bob have determined the set of subcarriers on which, they expect to be able to receive symbols from each other, in the presence of the active jammer. Unfortunately, the subcarriers on which Alice can receive information (known only to Alice at this stage) may be different from those on which Bob can receive information (known only to Bob at this stage). Thus, we need a way for Alice to let Bob know "which subcarriers to use" for communicating with her (Bob needs to do likewise).

A low throughput channel using Raptor codes to exchange CSI: Towards, this we leverage Raptor codes to communicate this information (which as previously mentioned, is called the CSI). Raptor codes belong to the class of fountain codes with linear encoding and decoding times. Fountain codes are rateless fault-tolerant codes that can enable reliable communications on erasure channels; examples of fountain codes include Raptor codes [16] and LT-codes [25]. Encoded symbols are generated by the encoder on-the-fly. The decoder recovers the source block by collecting a sufficiently large set of encoding symbols. Hence, Raptor codes facilitate communications in the presence of the jammer (jammed symbols could be considered to be erasures), by utilizing a very low throughput channel (as shown by our experiments later in this paper). Thus, in JIMS we only utilize these for the exchange of CSI information, and later simply utilize the relatively unaffected carriers without applying Raptor codes.

Specifically, Alice uses a bit vector to indicate the subcarriers to be used by Bob, and encodes this using Raptor codes. She transmits the encoded bit vector repeatedly (each time, the vector is encoded differently), until Bob is able to retrieve the source block (the bit vector). Upon this, Bob knows the set of subcarriers to use for correct reception at Alice. He uses only those sub-carriers to send legitimate symbols (from now on), and acknowledges the receipt of the bit vector. He also indicates (again with a bit vector), the subcarriers that are suitable for him, for reception in the presence of the jammer. At this point, both Alice and Bob are aware of the relatively unaffected subcarriers at each other's end.

Encrypting CSI: It is possible for the jammer to sniff the information encoded in the above message exchange. If it is able to retrieve the information with regards to the good subcarriers at Alice or Bob, it can (a) skew its power allocation (when transmitting) on the subcarriers to increase the interference on these specific subcarriers and/or (b) construct and send fake CSI information to Alice and Bob that aids the jammer's goal. To prevent the jammer from gleaning the CSI information, we encrypt the bit vector using a symmetric key that it either pre-provisioned, or securely established via an authenticated key exchange protocol, as discussed earlier.

Summary: In summary, JIMS consists of the three steps described in each of the previous subsections. An algorithmic representation of JIMS is provided in Algorithm 1.

Algorithm 1: JIMS Channel Measurement Procedure

Input: $received\_signal$ is the physical signal received from the antenna.

Output: $channels\_vector$ is the selected channel vector. ;

Initialization: $channels\_vector \leftarrow 0$;

$RJSNR\_vector \leftarrow$ process($received\_signal$) ;
$\eta \leftarrow$ calculate_selection_threshold(SINR_vector) ;
for $i \leftarrow 0$ **to** 48 **do**
    $channels\_vector\ (i)$ = decide(η, $SINR\_vector(i)$) ;
    if $channels\_vector\ (i)$ == *'1'* **then**
        $counter \leftarrow counter$ + 1;

return $channels\_vector$;

D. JIMS with Power Allocation (JIMS-PA)

Thus far, JIMS simply identified those subcarriers that were relatively unaffected by the jamming signal from Eve, and used those subcarriers for the exchange of information between Alice and Bob (in Eve's presence). Since, the information on the other subcarriers, i.e., those that are heavily affected by Eve are relatively unusable, we ask the question, "Can we reallocate some of the power from such subcarriers, to the subcarriers that are being used in order to enhance the throughput?" The answer to this question is that, such a reallocation is possible to some extent. However, one cannot simply reallocate all the power onto the "good" subcarriers for two reasons. First, because of the spectral flatness regulations specified in the 802.11 standard (specifically 802.11n) [14], the difference in the powers allocated to two subcarriers cannot exceed 2 dB. Second, if we blindly assign high powers to the good subcarriers, Eve will notice the anomaly, and can target those subcarriers. Thus, we can only reallocate powers to some extent, and we seek to do so here while adhering to the first constraint.

Specifically, let us assume that Alice has learnt of the SINRs experienced by Bob on each of the composite subcarriers. Based on the received SINR information, Alice now seeks to maximize the throughput, $\tau_a(N_d)$, by finding (a) the appropriate number of "good" subcarriers, N_d, (b) the best bit-rate or modulation for use (M_c), and (c) the optimal power reallocation strategy as discussed above.

Power units: Before formally defining the problem, we define what we call "power units". As mentioned earlier, the spectral flatness constraint requires that we limit the difference in transmission powers between any two subcarriers to 2 dB. This in turn implies that we can only remove at most 1 dB, or add at most 1 dB to a subcarrier. Conservatively, we limit the power removed/reallocated, from/to any subcarrier, to 0.75 dB. It is hard to consider all possible power allocations, by considering the transferred power quantums to be real valued. Therefore, we reduce the search space by quantizing the 0.75 dB budget into discrete power units. We set a power unit to be equivalent to 0.08333 dB (other settings are possible). Power reallocations across subcarriers is always in terms of a "number" of power units (at most 9 units can be transferred with our setting).

Our objective: Now, we formally define the problem to be a throughput maximization problem as follows:

$$\underset{N_d, M_c, \mathcal{X}_\Sigma}{\text{maximize}} \quad \tau_a(N_d) \tag{3}$$

$$\begin{aligned} \text{subject to} \quad & N_d \in \{1,\ 2, \ldots,\ N\} && \text{(I)} \\ & M_c \in \{2,\ 4,\ 16,\ 64\} && \text{(II)} \\ & \mathcal{X}_i < \mathcal{Y} \quad \forall i && \text{(III)} \\ & \mathcal{X}_\Sigma = \sum_{i=1}^{N_d} \mathcal{X}_i \leq \hat{\mathcal{X}}(N_d) && \text{(IV)} \end{aligned}$$

where, N_d is the subset of subcarriers selected for communication and N is the set of all subcarriers. $\mathcal{X}_i$ is the number of power units that allocated to a subcarrier $i \in N_d$, $\mathcal{Y}$ is the maximum number of power units that can be removed or added to a subcarrier. $\hat{\mathcal{X}}(N_d)$ is the maximum total power that can be reallocated (discussed later).

In the above formulation, Alice seeks to maximize the throughput by appropriately selecting a set of subcarriers, N_d, the appropriate modulation M_c on these subcarriers, and the best power allocation strategy. The last two constraints limit the maximum power transferrable to a subcarrier, and the total power that can be transferred.

The maximum power available for reallocation (referred to as the total power budget) depends on the excess power that can be removed from *usable* subcarriers (this is just a reallocation of power from among the good subcarriers), $\mathcal{X}_j^{excess}$ and the amount of total power units that can be taken from the unused subcarriers. A used subcarrier has an excess power if its SINR value exceeds the minimum required SINR threshold for the current modulation. This threshold is different for different modulation schemes. Similar thresholds has been used in SNR based rate adaptation schemes [26]. We use the mapping provided in section III (specifically Fig. 7) between subcarrier power and SINR in order to calculate the excess or requirement of the power on an subcarrier.

The available power budget for a subset of N_d subcarriers can be expressed as:

$$\hat{\mathcal{X}}(N_d) = (|N| - |N_d|)\mathcal{Y} + \sum_{j=1}^{|N_d|} \mathcal{X}_j^{excess} \tag{4}$$

We reallocate power units to usable subcarriers to ensure that the sender is able to transmit at a higher rate than before and/or be able to convert "bad" subcarriers to "good" subcarriers. Thus, this process increases the overall capacity in the presence of the active jammer.

Considering all possible power reallocations towards finding the maximum possible throughput, results in an exponential number of possibilities. Specifically, there are $O(\mathcal{Y}^{|N_d|})$ combinations as per which, power can be assigned to the $|N_d|$ subcarriers. The number of ways by which power can be *removed* from unused subcarriers is $O(\mathcal{Y}^{|N|-|N_d|})$. There are M_c modulation types. Hence, the number of combinations for the power reallocation to be considered is $O(M_c\mathcal{Y}^{|N_d|}\mathcal{Y}^{|N|-|N_d|})$ or simply $O(\mathcal{Y}^{|N|})$). To reduce computational complexity, we propose a heuristic that runs in polynomial time and is independent of $\mathcal{Y}$.

Algorithm 2: JIMS-PA Algorithm

Input: $\vec{S}$ received SINR vector ;

Output: N_d, M_c, $\vec{\mathcal{X}}$ the selected power strategy.

Initialization: $\vec{S} \leftarrow$ sort_dec($\vec{S}$) ;
$\tau_{avg}^{BPSK}(N_d - \{i\}) \leftarrow 0$;
$N_d \leftarrow \phi$;

for $i \in N$ **do**
  $N_d \leftarrow i$;
  $\tau_{avg}^{BPSK}(N_d) \leftarrow$ calculate_throughput() ;
  if $\tau_{avg}^{BPSK}(N_d) > \tau_{avg}^{BPSK}(N_d - \{i\})$ **then**
    $\vec{\mathcal{X}} \leftarrow$ perform_power_reallocation();
    $M_c \leftarrow$ calculate_throughput_and_MCs() ;
  else
    is_success $\leftarrow$ call Algorithm 3 ;
    if *is_success* $==$ *true* **then**
      $\vec{\mathcal{X}} \leftarrow$ perform_power_reallocation();
      $M_c \leftarrow$ calculate_throughput_and_MCs() ;
    else
      break;

Details of JIMS-PA: Let us assume, for simplicity that Alice is communicating with Bob in the presence of a jammer. Alice transmits a known pilot using all the subcarriers. Upon receiving the pilot, Bob calculates the per subcarrier SINR and sends the computed (SINR) values in an ACK/NACK packet to Alice (Note that here the raw SINR values are sent as opposed to simply a bit vector that indicates the good subcarriers). Then both Alice and Bob apply JIMS-PA as in Algorithm (2).

Initially, JIMS-PA sorts the subcarrier SINR values in descending order. At each step a single subcarrier, i, from the sorted subcarrier list, N, is considered ($i \in N$) as discussed

below. JIMS-PA initializes two subcarrier sets; usable (N_d) and unusable ($N - N_d$). In the beginning, the usable set is empty. In each step, a new subcarrier (specifically the subcarrier that supports the highest SINR) from the unusable set is considered for addition to the usable set. With this new subcarrier, say i, let us assume that the cardinality of the usable set is N_d. JIMS-PA then calculates $\tau_{avg}^{BPSK}(N_d)$, the throughput using BPSK modulation, considering the subcarriers in the usable set (throughput calculation discussed later). It compares this throughput with that achieved without i, i.e., $\tau_{avg}^{BPSK}(N_d - \{i\})$ (this is the throughput with the usable subcarrier set from the previous step). If the throughput degrades by including the new subcarrier for BPSK modulation (less vulnerable to errors) then it will degrade with higher modulations (the packet error probability on this subcarrier would be worse for higher modulations such as QPSK, QAM16, etc.). Thus, at this point there are two possible cases; (1) the value of $\tau_{avg}^{BPSK}(N_d)$ is greater than $\tau_{avg}^{BPSK}(N_d - \{i\})$ or (2) the value of $\tau_{avg}^{BPSK}(N_d)$ is less than $\tau_{avg}^{BPSK}(N_d - \{i\})$. We call these Case 1 and Case 2, and elaborate on them below.

Case 1: In this case, JIMS-PA adds subcarrier i to the usable set. It then considers the reallocation of $\mathcal{Y}(|N| - |N_d|)$ power units from the subcarriers in the unusable set, to those in the usable set. Beginning with the poorest subcarrier (lowest SINR) in the usable set, it incrementally assigns power, one unit at a time. It ensures that it does not violate constraints (III) and (IV) in the maximization formulation 3, when performing the power reallocation. It then does an internal reassignment of powers from among the subcarriers in the usable set, and determines that best applicable modulation scheme (in terms of the achievable throughput) with the resulting SINR values (recall the discussion of how the increase in power can be mapped onto increases in SINR values from Section III). It also computes the maximum achievable throughput with the current set of subcarriers in the usable set.

Algorithm 3: Subcarrier Revival Algorithm

Input: $\tau_{avg}^{BPSK}(N_d - \{i\})$, Subcarrier i ;

Output: $is\_success$ subcarrier can be revived or not ;

Initialization: $is\_success \leftarrow$ false ;

for $j \leftarrow 0$ **to** $\mathcal{Y}$ **do**

  add_power_units_to_i(j) ;

  $\tau_{avg}^{BPSK}(N_d) \leftarrow$ calculate_throughput() ;

  if $\tau_{avg}^{BPSK}(N_d) > \tau_{avg}^{BPSK}(N_d - \{i\})$ **then**

    $is\_success \leftarrow$ true ;

Case 2: In this case, it is clear that simply adding subcarrier i to the usable set will be in fact detrimental to the throughput. However, it may be possible to *revive* or make subcarrier i usable via power reallocation. Thus, JIMS-PA adds the maximum possible power (transferred from subcarriers in the current unusable set) to subcarrier i. If at this point, the throughput with the added subcarrier exceeds that with BPSK computed in the previous step, JIMS-PA proceeds as with Case 1. If not, the process stops. The subcarriers (excluding i) in the usable set are the subcarriers chosen for use. Power reallocation is then applied formally to this set, and communications now take place using this set of subcarriers.

There are three stopping conditions for JIMS-PA; *(i)* there is no power budget left for reallocation, *(ii)* a subcarrier cannot be revived, and *(iii)* all the subcarriers are added ($N_d = N$). JIMS-PA declares a solution when one of these three conditions is satisfied. Upon reaching a solution, JIMS-PA returns the set of subcarriers(N_d), the modulation (M_c) to use and the per subcarrier power allocation ($\vec{\mathcal{X}}$).

Computational complexity: It is easy to verify that the run time for JIMS-PA is $O(|N| \times M_c)$, where $|N|$ is the total number of sub-carriers and M_c is the number of available modulations. In brief, JIMS-PA iterates over $|N_d|$ subcarriers and for each sub-carrier, i, it iterates over the available modulation schemes to select the best modulation with power redistribution. Since the subcarriers are a priori sorted, power redistribution only takes $O(N)$ time.

Computing the throughput with a given set of subcarriers: Next, we present the calculation of the throughput based on per subcarrier SINRs. In order to make the analysis generally applicable with different transmission technologies and MAC layer protocols (WiFi, LTE, WiMAX, etc.), we do not consider MAC layer-specific packet retransmissions in our computation and experiments.

The throughput τ_a depends on three factors (1) the maximum number of re-transmissions, (2) the duration of the OFDM symbol, (T_s), and (3) the packet error probability (PER), p_e. The PER for a set of subcarriers is:

$$p_e = 1 - \prod_{i=1}^{|N_d|} (1 - p_b(i))^{L/|N_d|} \tag{5}$$

where p_e is a function of the bit error probabilities (denoted by $p_b(i)$) on each of the subcarriers that are used and L is the packet $p_b(i)$ depend on the modulation in use and the SINR value. We simply use the *erfc* function [27] to calculate $p_b(i)$, given M_c and the SINR.

The average transmission time, T_{avg}, of a packet is a function of the expected number of transmissions, $E(R)$ and the packet transmission time (given the set of subcarriers), $\rho(N_d)$. Specifically,

$$T_{avg}(N_d) = \rho(N_d) \times E(R) \tag{6}$$

If the maximum number of retransmissions possible is R, the expected number of transmission attempts is given by,

$$E(R) = \sum_{r=1}^{R} r p_e^{r-1}(1 - p_e) = \frac{1 - (R+1)p_e^R + R p_e^{R+1}}{(1 - p_e)} \tag{7}$$

The packet transmission time for a given number of subcarriers, $|N_d|$, can be calculated as follows

$$\rho(N_d) = \frac{L}{|N_d| M_I} \times T_s \tag{8}$$

where T_s is the OFDM symbol duration, L is the packet length and M_I is the modulation index ($M_I = \log_2 M_c$). From Equations 6, 7 and 8, the average transmission time $T_{avg}(N_d)$ is given by:

$$T_{avg}(N_d) = \frac{L}{|N_d| M_I} \times T_s \times \frac{1 - (R+1)p_e^R + R p_e^{R+1}}{(1 - p_e)} \tag{9}$$

The average throughput, $\tau_a(N_d)$, is given by

$$\tau_a(N_d) = \frac{1}{T_{avg}(N_d)} \quad (10)$$

V. Implementation and Performance Evaluation

In this section, we evaluate our proposed framework via extensive real testbed as well as simulation experiments.

A. Testbed Implementation

We implement our proposed schemes on top of the WARPLab framework [28], which runs on WARP boards in real time. The proposed framework is implemented as a thin layer between PHY and MAC layers to determine the most appropriate subcarriers. The receiver communicates the obtained information through short messages with the transmitter (as described in Section IV). We modify the transmitter to perform power allocation to the selected subcarriers (described with JIMS-PA) upon receiving CSI feedback, included in ACK/NACK packets from the receiver. The ACK/NACK packets are encoded using Raptor codes. Each ACK/NACK packet contains a subcarrier vector of 48 bits which corresponds to the 48 OFDM subcarriers. A bit with value one means a subcarrier is selected (zero otherwise). In all our experiments CSI feedback rate is set to 1 sec. We implement three jamming models; periodic jamming, random jamming and continuous jamming.

For the periodic jamming, the adversary transmits the jamming signal periodically, thereby effecting the legitimate communication for the period it is ON. The jamming pulse lasts for one second. For the random jamming, the adversary transmits the jamming signal at random time intervals for a random time span between 0.5 & 2 seconds. We use the same experimental setup and parameters described in section III.

Raptor codes: To encode ACK/NACK packets, we implement Raptor based Forward Error Correction (FEC) as specified in the standard [29]. We first divide the CSI packet into a number of source blocks, Z. Each source block consists of k OFDM data symbols. Raptor encoding is applied independently on each symbol. The encoder generates l encoded symbols for the k data symbols in a block that are uniquely defined by a set of constraints [29]. For our experiments we set $l = 2k$; k varies between 10 and 50.

The original data symbols can be recovered from any subset of the encoded symbols of size equal to or slightly larger than the number (k) of original symbols. In addition, the coding rate varies according to the observed average $RJSNR$.

The Advanced Encryption Standard (AES): To prevent an adversary from sniffing the CSI packet, legitimate transceivers employ the AES based encryption for ACK/NACK packets, using a shared 128-bit key. We use the publicly available AES implementation in [30] and integrate it with our framework in WARPLab.

B. Experimental Results

Explicit vs. Implicit jamming approaches for determining the good subcarriers: We first compare the performance of explicit and implicit approaches for determining the good subcarriers. Fig. 8 shows the throughput achieved with the explicit and implicit approaches with different network settings. We create four different network topologies as shown in Fig. 1-3 with a legitimate sender-receiver pair and a continuous jammer as an attacker. We see that both schemes perform equally well.

Choice of SINR threshold: As we discussed earlier, an inherent challenge in subcarrier selection is choosing the appropriate SINR thresholds. Since both the explicit and implicit approaches have similar performance, we only consider the SINR to demonstrate the effectiveness of our approach of using the average (recall Section IV-B) as the threshold for subcarrier selection. For this experiment, we use a legitimate sender-receiver pair, which is communicating in presence of a continuous jammer. We try 6 different static SINR thresholds (from low to high values) and compare the performance with their use (in JIMS) against the average SINR threshold that we advocate. Fig. 9 shows the throughput achieved by the legitimate sender-receiver pair with the different SINR thresholds. We observe that average SINR threshold achieves a higher throughput than any other SINR thresholds; specifically it can achieve up to 5 times more throughput than the poorest static threshold.

CSI exchange: Next we evaluate the effectiveness of the CSI exchange under different jamming attacks. The metric of interest is the number of encoded (re-)transmissions required to deliver CSI packets. Figs. 10 and 11 show the percentage of successfully received packets with respect to the number of transmissions for JIMS and JIMS-PA. With a periodic jamming attack, 80% of CSI packets are delivered with three encoded transmissions, while with the random jammer, four transmissions are required to deliver the same number of CSI packets (the skewed periods cause this effect since sometimes the jamming periods are longer) with both schemes. As one might expect, the continuous jamming attack is the hardest to cope with; we need eight transmissions to deliver 80% of the CSI packets. We observe that in the worst case, the delivery of a CSI packet requires ten encoded transmissions.

Convergence times of the CSI exchange: JIMS's performance is critically dependent on the exchange of CSI packets. However, as mentioned earlier, these packets are sent using all the subcarriers with Raptor encoding and suffer from interference from the jammer. We examine if this first step in JIMS's design can become a communication bottleneck by measuring the time it takes to successfully exchange CSI packets and begin the usage of the good sub-carriers; we refer to this as the convergence time. We measure the convergence time under 3 attack scenarios and with various network topologies.

Fig. 12 depicts how the convergence time varies for 8 different network topologies. For each topology the results are averages over 20 trials. We see that the observed maximum convergence time is less than 35 ms. This demonstrates that JIMS converges fairly quickly. Fig 12 shows that JIMS-PA requires about 60 msec to converge in the worst case. This is because JIMS-PA has to send more data (SINR information)

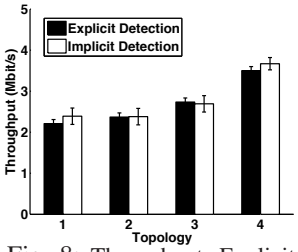

Fig. 8: Throughput: Explicit vs. Implicit detection.

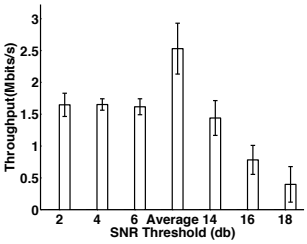

Fig. 9: SINR threshold selection.

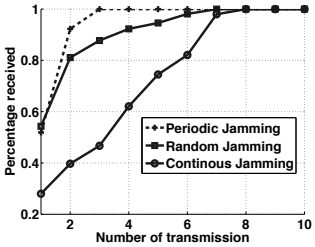

Fig. 10: Number of transmissions with JIMS for encoded CSI packets.

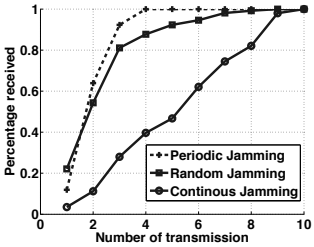

Fig. 11: Number of transmissions with JIMS-PA for encoded CSI packets.

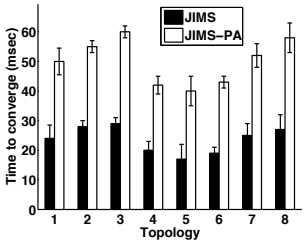

Fig. 12: Convergence Time with JIMS and JIMS-PA.

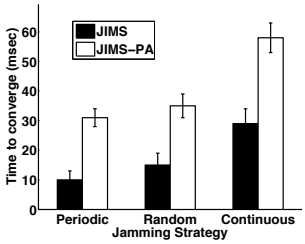

Fig. 13: Convergence times vs different attack strategies.

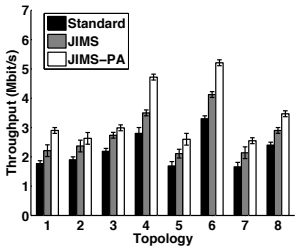

Fig. 14: Throughputs with Standard, JIMS and JIMS-PA.

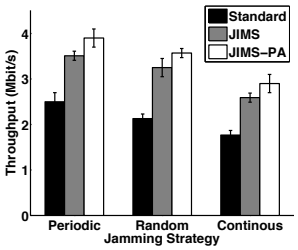

Fig. 15: Throughputs with different attack strategies.

in the ACK/NACK packets as discussed earlier; this larger data transfer is especially difficult due to the jamming signal. Thus, JIMS-PA takes a longer time to converge. If the nodes under consideration (both the legitimate transceivers and the jammer) are relatively static (small amounts of motion), these convergence times are sufficiently small in terms of overhead and can allow sustained use of the "good" subcarriers for relatively, much longer periods.

Fig. 13 depicts the impact of the attack model on the convergence times with JIMS and JIMS-PA. It is immediate that continuous jamming is the most hurtful while the other two attack models have similar (lesser) impact. Even with continuous jamming however, in the topologies considered, the maximum convergence time is 25 ms. Fig. 13 shows that JIMS-PA needs more time to converge compared to JIMS in all cases as expected. With continuous jamming, JIMS-PA has a convergence time of approximately 60 msec. In the best case scenario, it requires about 32 msec to converge.

Throughput performance: Next, we compare the performance of JIMS and JIMS-PA against a standard system which utilizes all the subcarriers for communication. Fig. 14 demonstrates the performance of three schemes in terms of throughput under different network topologies in the presence of a continuous jammer. Results are shown for 8 network topologies, as described in section III. We see that JIMS outperforms standard 802.11 system by up to 65%, while JIMS-PA does so by up to 75%. JIMS-PA provides an additional gain over JIMS (about 10%) in spite of the increased overehad during the CSI exchange process.

In Fig. 15, we depict the performance of the schemes in the presence of three attack strategies viz., periodic, random and continuous jamming. The results with the continuous jamming are as before. With random and periodic jamming, all the three schemes perform better. The JIMS schemes outperform the standard scheme to a slightly lesser degree. For example, JIMS-PA outperforms the standard system by up to 56%.

Impact of mobility: To examine the effect of user mobility on the performance of JIMS, we move the receiver away or towards the transmitter at a constant speed from its original position and measure the SINR at fixed intervals. Fig 16 shows the CDF of the degradation in the average SINR for a mobile receiver as a function of the distance that the receiver has moved from its initial position. The result suggests that the degradation in the average SINR with small extents of mobility (moving a smartphone or walking) is minimal.

We find that when the mobile receiver moves 1 or 2 inches away from its initial position, the throughput gains remain intact. However, if the receiver moves over a distance like 3 inches, the average SINR degrades more than 40%. Thus, the average throughput also decreases. At this point, the CSI values need to be exchanged again and the SINRs recomputed.

C. Simulations

Simulation setup: In order to examine the performance of JIMS and JIMS-PA in larger scale settings, we have implemented both schemes on NS3 version 13, using a detailed PHY layer model called PhySimWifi [31]. We use an on off client server model to generate application level traffic. For the MAC and network layers, we use 802.11a with Friis propagation loss model. The packet size is fixed to 1500 bytes and the transmission power is set to 20 dBm with the background noise varying between -90 and -99 dBm. We have averaged the results over 25 runs where each run lasts for 100 seconds. We consider random and grid topologies, and place the jammer in arbitrarily

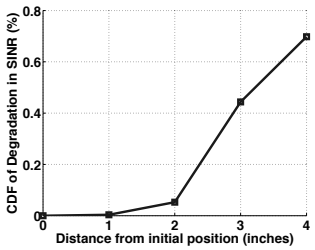

Fig. 16: Degradation in SINR with mobility.

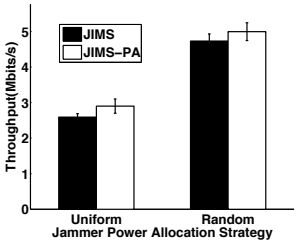

Fig. 17: Throughput comparison for various jamming power allocation strategies.

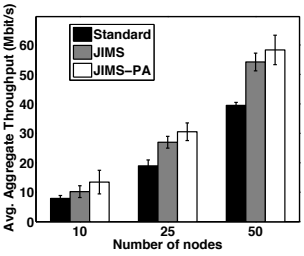

Fig. 18: Throughput with increased node density.

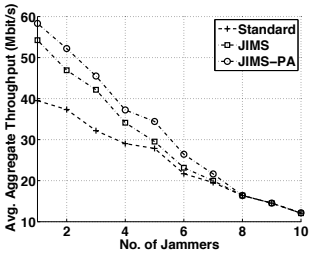

Fig. 19: Effect of multiple jammers on throughput.

chosen positions. To ensure transmissions even when sensing the jamming signal, we disable carrier sensing.

Scalability: We plot the average aggregate throughput with JIMS and JIMS-PA with different numbers of nodes in the presence of a jamming attack. Fig. 18 shows the aggregate throughput with JIMS, JIMS-PA and the standard system with increasing node density. In this scenario we use a single continuous jammer which is trying to disrupt the communication in its neighborhood. With increasing node density, the benign interference levels also go up in addition to the interference from the jammer. As we can see from Fig. 18 that JIMS-PA outperforms the standard system by upto 70%. In addition, JIMS-PA performs better than JIMS by 30% in the best case due to intelligent power reallocation.

Multiple jammers: To show the resiliency of JIMS against multiple jammers, we varied the number of jammers in the network (benign interference still exists) and calculated the aggregate throughput for both JIMS-PA and JIMS. Fig. 19 shows the average aggregate throughput with an increasing number of jammers, in comparison to a case with a single jammer with a 50 node random network topology. We see that JIMS-PA and JIMS both suffer as we increase the number of jammers. However, in a two Jammer case, the throughput loss is about 15%; this still is better than the standard case with a single jammer. If there are more than 7 jammers, JIMS-PA and JIMS cannot restore any throughput since it is likely that all subcarriers are now affected by jamming.

VI. Conclusions

In this paper, we propose JIMS, a jamming interference mitigation scheme, using which, transceivers can identify subcarriers that are relatively unaffected by jamming and utilize them for communications. We show that JIMS restores throughput up to 75%, in the presence of an active jammer via experiments on our WARP testbed. At this time, we rely on prior schemes to detect the jammer, and utilize JIMS only when a jammer is detected. Integrating JIMS with such detection schemes effectively will be considered in future work.

References

[1] PKI 6650 Wideband Jammer. http://bit.ly/10us5My.
[2] Neco Defense Systems. http://www.necodefence.com/rfj.php.
[3] The GSM Jammers. http://bit.ly/TWz15d.
[4] G. Noubir, R. Rajaraman, B. Sheng, and B. Thapa. On the Robustness of IEEE 802.11 Rate Adaptation Algorithms Against Smart Jamming. In *ACM WISEC*, 2011.
[5] K. Pelechrinis, I. Broustis, S. V. Krishnamurthy, and C. Gkantsidis. ARES: An Anti-jamming REinforcement System for 802.11 Networks. In *ACM CoNEXT*, 2009.
[6] C. Orakcal and D. Starobinski. Rate Adaptation in Unlicensed Bands under Smart Jamming Attacks. In *CrownCom*, 2012.
[7] V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein. Using Channel Hopping to Increase 802.11 Resilience to Jamming Attacks. In *IEEE INFOCOM*, 2007.
[8] K. Pelechrinis, C. Koufogiannakis, and S. V. Krishnamurthy. On the Efficacy of Frequency Hopping in Coping with Jamming Attacks in 802.11 Networks. *IEEE Trans. Wireless Commun.*, 2010.
[9] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa. On the Performance of IEEE 802.11 under Jamming. In *IEEE INFOCOM*, 2008.
[10] Rice University WARP Project. http://warp.rice.edu.
[11] H. Sari, Y. Levy, and G. Karam. An Analysis of Orthogonal Frequency-Division Multiple Access. In *IEEE GLOBECOM '97.*
[12] S. Sesia, I. Toufik, and M. Baker. LTE - The UMTS Long Term Evolution: From Theory to Practice. In *Wiley; 2nd Ed.*, Aug. 2011.
[13] M. Roger. IEEE 802.16 WirelessMAN Standard: Myths and Facts. In *Wireless Communications Conference.*, 2006.
[14] IEEE 802.11n Standard. http://bit.ly/126Nw94.
[15] Q.T. Zhang, X. Y. Zhao, Y. X. Zeng, and S. H. Song. Efficient Estimation of Fast Fading OFDM Channels. In *IEEE ICC '06*.
[16] A. Shokrollahi. Raptor codes. *IEEE Trans. Inf. Theory*, 2006.
[17] R.-T. Chinta, T.F. Wong, and J.M. Shea. Energy-Efficient Jamming Attack in IEEE 802.11 MAC. In *IEEE MILCOM*, 2009.
[18] W. Hu, T. Wood, W. Trappe, and Y. Zhang. Surfing and Spatial Retreats: Defenses Against Wireless Denial of Service. In *ACM Workshop on Wireless Security*, 2004.
[19] M. Stahlberg. Radio Jamming Attacks Against Two Popular Mobile Networks. In *Helsinki U. of Tech. Seminar on Network Security*, 2000.
[20] K. Pelechrinis, M. Iliofotou, and S. V. Krishnamurthy. Denial of Service Attacks in Wireless Networks: The case of Jammers. In *IEEE Comm. Surveys and Tutorials*, 2011.
[21] Y. Liu, P. Ning, H. Dai, and A. Liu. Randomized Differential DSSS: Jamming-Resistant Wireless Broadcast Communication. In *IEEE INFOCOM*, 2010.
[22] A. Brusilovsky, I. Faynberg, and Z. Zeltsan. Password-Authenticated Key (PAK) Diffie-Hellman Exchange. RFC 5683, 2010.
[23] V. Cakulev, G. Sundaram, and I. Broustis. IBAKE: Identity-Based Authenticated Key Exchange. RFC 6539, 2012.
[24] W. Xu, W. Trappe, Y. Zhang, and T. Wood. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks. In *MobiHoc*, 2005.
[25] M. Luby. Lt codes. In *IEEE FOCS*, 2002.
[26] Tae-Suk Kim, Hyuk Lim, and Jennifer C. Hou. Improving Spatial Reuse Through Tuning Transmit Power, Carrier Sense Threshold, and Data Rate in Multihop Wireless Networks. In *ACM MobiCom*, 2006.
[27] Harry R. Anderson. *Fixed Broadband Wireless System Design: The Creation of Global Mobile Communications*. 2003.
[28] Wireless Open-Access Research Platform. http://warp.rice.edu/.
[29] Raptor Forward Error Correction (RFC 5053). http://bit.ly/13VJYZq.
[30] AES using Bouncy Castle API. www.itcsolutions.eu.
[31] PhySim-WiFi. http://dsn.tm.kit.edu/english/ns3-physim.php.