

On the Performance of IEEE 802.11 under Jamming

E. Bayraktaroglu*, C. King†, X. Liu*, G. Noubir*, R. Rajaraman*, B. Thapa*

*College of Computer & Information Science, Northeastern University, Boston MA 02115

†Department of Mathematics, Northeastern University, Boston MA 02115

Abstract—In this paper, we study the performance of the IEEE 802.11 MAC protocol under a range of jammers that covers both channel-oblivious and channel-aware jamming. We study two channel-oblivious jammers: a *periodic* jammer that jams deterministically at a specified rate, and a *memoryless* jammer whose signals arrive according to a Poisson process. We also develop new models for channel-aware jamming, including a *reactive* jammer that only jams non-colliding transmissions and an *omniscient* jammer that optimally adjusts its strategy according to current states of the participating nodes.

Our study comprises of a theoretical analysis of the saturation throughput of 802.11 under jamming, an extensive simulation study, and a testbed to conduct real world experimentation of jamming IEEE 802.11 using GNU Radio and USRP platform. In our theoretical analysis, we use a discrete-time Markov chain analysis to derive formulae for the saturation throughput of IEEE 802.11 under memoryless, reactive and omniscient jamming. One of our key results is a characterization of optimal omniscient jamming that establishes a lower bound on the saturation throughput of 802.11 under arbitrary jammer attacks. We validate the theoretical analysis by means of Qualnet simulations. Finally, we measure the real-world performance of periodic and memoryless jammers using our GNU radio jammer prototype.

I. INTRODUCTION

The IEEE802.11 CSMA/CA MAC protocol is widely used and operates over many physical layers such as DSSS/FHSS/IR, CCKFHSS (IEEE802.11b), OFDM (IEEE802.11a), and MIMO (IEEE802.11n) [1]. It is reasonably efficient for controlling medium access and delivers a throughput significantly higher than other non-explicit reservation MAC protocols such as Aloha, and variants of CSMA [2]. However, efficiency is achieved through a relatively sophisticated control mechanism, and by making assumptions on the behavior of competing nodes and the characteristics of the channel. Such control mechanisms are usually the target of choice for malicious attackers.

A natural objective of adversaries is to drastically reduce the throughput of the communicating nodes while using as little energy as possible. This can be achieved by carefully jamming critical packets or bits at the right moment, frequency, and location. Such a strategy enables an adversary to devise sophisticated attacks including the partitioning of a network, or redirecting traffic through areas under the control of the adversary. Conserving energy increases the lifetime of jammer nodes (also called cybermines), which then remain a threat for a longer period of time. Building such smart jammers is within the reach of the public at large, due to the availability of low-cost fully controllable Software Defined Radio platforms such as USRP/GNU-Radio [3], [4] and many other partially controllable sensor network platforms operating over the 2.4GHz ISM band [5]. Since IEEE802.11 MAC is widely used and common

to many physical layers, it is important to understand its limits in terms of resiliency to smart jammers.

A. Our Contributions

In this paper, we study the performance of IEEE802.11 MAC in the presence of various types of jammers through a systematic theoretical analysis, extensive simulations, and a prototype implementation.

- Building on the discrete Markov model of [2], we analyze the saturation throughput of 802.11 (basic mode) under both channel-oblivious and channel-aware jammer models. Our theoretical analysis framework is general and can be used to analyze the resilience of other MAC protocols to jamming.
- We introduce the notion of a channel-aware omniscient jammer and derive key properties of an optimal omniscient jammer. In addition to identifying damaging jamming techniques, our analysis of an optimal jammer provides a lower bound on the throughput achievable by 802.11 under arbitrary adversarial jamming.
- We validate our theoretical analysis through an extensive simulation study using Qualnet. We also develop a GNU Radio prototype jammer testbed for implementing memoryless and periodic jammers and compare the prototype results with theory and simulations.
- Our results indicate that while a periodic channel-oblivious jammer is fairly damaging for large packet sizes and large saturated networks, it is significantly less effective than channel-aware jamming, allowing orders of magnitude more throughput for small jamming rates. Furthermore, an optimal omniscient jammer is even 20-30% more effective than other natural channel-aware jammers, and is especially efficient against networks with a small number of active sessions.

B. Related Work

Wireless networks are highly sensitive to denial of service attacks. The wireless communication medium is a broadcast channel, exposing the physical layer of wireless communication to jamming originating at arbitrary locations [6], [7]. There has also been considerable research on attacks on the control mechanisms at higher layers as well as cross-layer attacks (e.g., [8], [9]). The focus of this paper is on the MAC layer, which is sensitive to attacks targeting the control channels and mechanisms owing to the limited sensing capabilities in the wireless medium [10]–[12]. The work [11] analyzes the throughput of CSMA/CA under adversarial jamming, assuming the Poisson arrival of packets. The recent work of [5] classifies jammer attack models and presents jamming detection techniques.

The IEEE 802.11 MAC protocol is widely used and has been extensively analyzed with respect to various performance issues, including throughput, power control, fairness, as well as hidden terminal jamming problems [2], [13], [14]. With the increased ease of building low-cost jammers and increased interest in studying DoS attacks, researchers have started studying the effect of adversarial jamming on 802.11 [12], [15], [16]. A recent series of studies analyzes the energy-efficiency of several jamming techniques against 802.11 [16], [17]; they demonstrate through extensive simulations that intelligent jamming by concentrating jamming signals on control packets (e.g., CTS or ACK) is significantly more energy-efficient than jammers that are oblivious to the channel. In our work, we have analyzed a wider range of jammers through theoretical analysis, simulations, as well as a GNU radio prototype testbed. Another difference between [16], [17] and our work is that while their performance measure of interest is the jammer energy needed to completely shut down the channel, our study considers the entire throughput range and analyzes how 802.11 throughput varies as a function of jammer rate (and, hence, energy). Another recent work studies the impact of periodic jammers on an 802.11 LAN supporting simultaneous Voice over IP (VoIP) connections through simulations [14], while [18] and [19] propose channel hopping and protocol hopping techniques to increase the robustness of 802.11.

Our theoretical contributions build on the framework of [2] for analyzing the saturation throughput of 802.11. There have been several subsequent studies that refine the model of [2] or consider different traffic models, channel conditions, or performance measures (e.g., [20], [21]). To the best of our knowledge, our work is the first theoretical analysis of the IEEE 802.11 MAC under adversarial jamming.

There has also been considerable interest recently on jamming attacks against sensor networks; [22] gives a taxonomy of attacks, [23] formulates the jammer-network interaction as an optimization problem, while [24] studies the resiliency of several sensor MAC protocols.

II. MODELS OF COMMUNICATION AND JAMMING

Medium Access Control Model: IEEE802.11. Our focus is on the IEEE802.11 Distributed Coordination Function (DCF) [1]. DCF is a distributed MAC protocol based on CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) for networks with or without infrastructure. It has two modes: a basic mode which has a DATA/ACK exchange and an extension with RTS (Request To Send)/CTS (Clear To Send) handshake prior to DATA/ACK. The RTS/CTS exchange was designed to reserve the channel in advance and minimize the impact of collisions but obviously does not help against jamming [16], [17]. In this paper, we only consider the basic mode.

IEEE802.11 defines four types of IFSs (Inter Frame Space): SIFS, DIFS, PIFS, and EIFS [1]. SIFS (Short Interframe Space) is the shortest IFS and is used between RTS, CTS, DATA and ACK frames. The PIFS (PCF Interframe Space) is used under PCF (Point Coordination Function) but not in the DCF mode. DCF requires the wireless nodes to defer the transmission until the medium has been idle without interruption for a period of

DIFS (DCF Interframe Space) or EIFS (Extended Interframe Space). If the last frame reception is successful, DIFS is applied. If the last frame reception does not result in a correct frame check sequence, EIFS must be applied. In our previous work, we have devised an efficient attack against IEEE802.11 using periodic pulses with a period of EIFS. We have also shown how one can protect against this type of attack [19]. In this paper, we consider DCF without the EIFS functionality. IEEE802.11 uses an exponential backoff scheme for contention avoidance, whose details we defer to Section III-A, where we present the Markov chain model for our analysis.

Jammer Models for MAC-Layers. We classify jammers of the MAC layer into four abstract categories according to their capability of sensing and reacting to the medium state (*Channel-Oblivious vs. Aware*), and maintaining a state that dictates their future actions (*Memoryless vs. Stateful*):

- **Channel-Oblivious & memoryless** jammers make jamming decisions without sensing the channel, and independently from their past actions. There are only two types of channel-oblivious & memoryless jammers: (a) in continuous time, jamming pulses arrive according to a Poisson distribution; (b) in discrete time, the jammer has a fixed probability of transmitting a pulse every timeslot.
- **Channel-Oblivious & stateful** jammers do not have access to the channel state; however, their actions may be dependent on their past behavior. The simplest example is a *periodic jammer*. A more sophisticated jammer of this type may send a burst of pulses and then stop for a long period of time before repeating. Such a jammer could attempt to drive the nodes into a long backoff period where they do not attempt to send packets even though no jamming is occurring.
- **Channel-Aware & memoryless** jammers have basically one jamming rate for each possible state of the channel (e.g., busy, idle). In a continuous-time model, the pulses are generated according to a Poisson process with different rates for the two states.
- **Channel-Aware & stateful** jammers are the most sophisticated jammers. One such jammer is a *reactive jammer*, which senses the medium and transmits a jamming pulse with a specified probability whenever it detects a non-colliding transmission. The strongest channel-aware and stateful jammer is an *omniscient jammer*, which senses the medium and can identify the number of retransmissions that a packet went through. Whenever such a jammer detects a non-colliding transmission, it transmits a jamming pulse with a probability that may depend on the the backoff stage of the transmitter.

Our paper focuses on four classes of jammers: channel-oblivious & memoryless jammers in continuous time (henceforth abbreviated as *memoryless* jammers), periodic jammers which are a special case of channel-oblivious & stateful jammers, and two channel-aware & stateful jammers: reactive jammers and omniscient jammers.

III. THEORETICAL ANALYSIS

Consider a wireless network with n pairs of 802.11 nodes and a jammer that jams the channel at a specified rate. Through-

out this section, we make the following assumptions for our analysis: (i) *ideal channel conditions*, that is, any transmission can be heard by every node in the network; thus, there are no hidden terminals or exposed terminals [2]; (ii) *saturation conditions*, that is, every node always has packets to send; and (iii) *ideal jamming conditions*, that is, a jamming signal destroys an 802.11 packet once their transmissions overlap.

Under the above assumptions, we derive the throughput of an 802.11 LAN under three probabilistic jamming models: memoryless, reactive, and omniscient. We derive formulae for the throughput under the three models, and establish key properties of an optimal omniscient jammer. Our characterization of an optimal jammer is, perhaps, the most significant theoretical contribution of this paper. These analyses are developed in Sections III-B through III-D. First, we present an analysis framework that is common to all these jammers.

A. An Analysis Framework

Following [2], we model the exponential backoff mechanism of 802.11 MAC protocol using a bidirectional discrete-time Markov Chain. Unlike [2], we adopt the protocol standard of a finite retransmission limit (this refinement of [2]'s model has been studied in [21]).

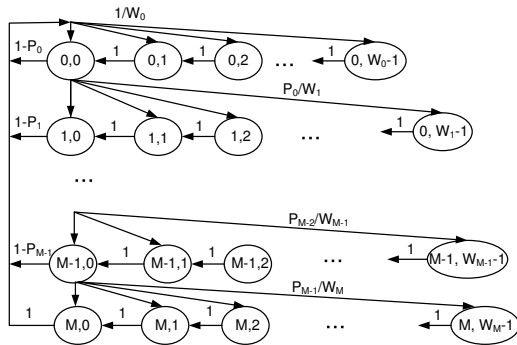


Fig. 1. Markov Chain model for 802.11 node under probabilistic jamming

Let $W_0 = CW_{min}$ denote the minimum contention window (CW), W_i be the CW of the i th backoff stage and M be the maximum retransmission limit. CW doubles when a transmission fails, i.e. $W_i = 2^i W_0$, until it reaches the maximum contention window CW_{max} . If the number of retransmissions exceeds M , the sender discards the current packet in the queue and resets CW to W_0 . Note that if $M > \log_2 \frac{CW_{max}}{CW_{min}}$, the last several backoff stages stay constant at CW_{max} . For simplicity, we will assume for our analysis that $M = \log_2 \frac{CW_{max}}{CW_{min}}$.

Figure 1 depicts the state transition of one 802.11 node in a discrete-time Markov chain. The communication is divided into *timeslots*. At the beginning of each timeslot, the backoff counter decreases by one, as shown in the figure, and the node transits from state (i, j) to state $(i, j - 1)$. When the backoff counter reaches zero, the node initiates a transmission. If the transmission succeeds, the node resets its backoff stage and enters $(0, j)$, where j is chosen uniformly at random from $[0, W_0 - 1]$. Otherwise, the node doubles its CW and enters state $(i + 1, j)$, where j is chosen uniformly at random

from $(0, W_{i+1} - 1)$. We note that the amount of time that a node spends in a state (the length of a timeslot), is variable, depending on whether the channel is busy (owing to an ongoing transmission, or even a jamming signal) or the channel is idle (in which case it equals the 802.11 physical slot parameter σ).

There are two reasons for the failure of a packet transmission by a node (which happens in a state of the form $(i, 0)$): the packet collides with a packet transmission initiated by another node, or the packet (or the associated ACK) is jammed by the jammer. We follow one fundamental assumption in Bianchi's model that in steady state, the probability that a packet transmission collides with a packet transmitted by another 802.11 node is independent of the current state of the transmitting node [2]. This assumption is justified, especially for a sufficiently large number of nodes and sufficiently large contention window size. Let P_c denote this collision probability.

Each of the three jammers we analyze in this section are probabilistic jammers, and can be captured by the probability with which they jam the channel at a given time. For the memoryless jammer, this probability is constant, independent of the state of the nodes. A reactive jammer, on the other hand, jams only when a transmission is ongoing and there are no collisions, but the jamming probability is independent of the backoff stage. Finally, the jamming probability of an omniscient jammer may depend on the backoff stage of the transmitting node. We define the *jamming probability* q_i of a jammer to be the probability of jamming an ongoing transmission in state $(i, 0)$ conditioned on the event that there is no collision with another 802.11 transmission. We now obtain that P_i , the probability that a transmission in backoff stage i fails, is given by $P_i = P_c + (1 - P_c)q_i$.

We are now ready to derive the state occupancy probabilities, which follow using standard Markov chain techniques. Let $b_{i,j}$ denote the probability for a node to be in the backoff stage i with backoff counter equals j in a steady state. We formulate the state transitions by following set of equations.

$$b_{i,j} = \begin{cases} b_{i,j+1} + P_i b_{i-1,0}/W_i & i > 0, j < W_i - 1 \\ P_i b_{i-1,0}/W_i & i > 0, j = W_i - 1 \neq 0 \\ b_{0,j+1} + b_{M,0}/W_0 & i = 0, j < W_0 - 1 \\ b_{M,0}/W_0 & i = 0, j = W_0 - 1 \end{cases}$$

Given values for the failure probabilities P_i , the above equations, together with the normalization condition that the $b_{i,j}$'s sum to one, can be solved to obtain the $b_{i,j}$ values. Since each node transmits only when its backoff counter reaches zero, the steady state transmission probability τ is given by $\sum_{i=0}^M b_{i,0}$. Given τ , P_i , and the protocol-related parameters packet length L , header size H , acknowledgment length ACK , we compute the throughput by determining the channel time-wise utilization for successful payload transmissions. The normalized throughput Γ is expressed by (1) [2].

$$\Gamma = \frac{E[\text{Payload transmitted in a timeslot}]}{E[\text{length of a timeslot}]} \quad (1)$$

The numerator of the above equation equals $P_s L$, where P_s is the probability that there is a successful transmission in a given timeslot (this depends on τ and the P_i 's, and the particular

jammer model), and L is the duration of the payload of a packet. The denominator of Equation 1 is given by

$$E[\text{length of a timeslot}] = P_{tr}T_{tr} + (1 - P_{tr})T_{id},$$

where $P_{tr} = 1 - (1 - \tau)^n$ is the probability that at least one node transmits in a given timeslot, T_{tr} is the time taken by a timeslot during which a transmission occurs, and T_{id} is the time taken by an idle timeslot (when no 802.11 node transmits). Specifying τ , P_s , T_{tr} and T_{id} then yields the throughput using (1). Finally, we define the rate of a jammer to be simply the fraction of time it jams the channel; it lies in $[0, 1]$. For example, a periodic jammer that emits pulses of width w μ s every ms has a rate of $w/1000$, and a continuous jammer has rate 1.

B. Memoryless Jammers

A memoryless jammer generates jamming signals such that the idle time between successive signals is drawn from an exponential distribution specified by the jamming rate R . The probability that a jamming signal is generated during a time interval t_0 is $(1 - e^{-Rt_0})$; this is, indeed, the jamming probability q_i for all i . Since q_i is independent of i , the failure probability P_i is also independent of i ; let p denote this common failure probability. We obtain that

$$p = P_c + (1 - P_c)(1 - e^{-R(\text{DATA}+\text{ACK})}),$$

where DATA and ACK refer to the duration of a data and ACK packet, respectively. The DATA term includes both payload length L as well as any headers.

Following our framework of Section III-A, we now derive the throughput by specifying τ , P_s , T_{tr} , and T_{id} . Plugging in the failure probability into the b_{ij} equations, we get:

$$b_{0,0} = \frac{2(1 - 2p)(1 - p)}{(1 - p)(1 - (2p)^{M+1})W + (1 - 2p)(1 - p^{M+1})} \quad (2)$$

The steady state transmission probability τ is given by

$$\tau = \frac{2(1 - 2p)(1 - p^{M+1})}{(1 - p)(1 - (2p)^{M+1})W + (1 - 2p)(1 - p^{M+1})} \quad (3)$$

Solving (2), (3), and the equation $P_c = 1 - (1 - \tau)^{n-1}$ over the three unknowns τ , p , and P_c yields τ .

We now determine P_s , T_{tr} , and T_{id} .

$$\begin{aligned} P_s &= n\tau(1 - \tau)^{n-1}e^{-R(\text{DATA}+\text{ACK})} \\ \text{EDIFS} &= \frac{1 - e^{-R \cdot \text{DIFS}}}{R \cdot e^{-R \cdot \text{DIFS}}} \\ T_{tr} &= \text{DIFS} + \text{SIFS} + \text{DATA} + \text{ACK} \\ T_{id} &= (1 - R)\sigma + R(\text{EDIFS} + \sigma) \end{aligned}$$

The above equations in conjunction with the equations of Section III-A give us the throughput of the system.

C. Reactive Jammers

We specify a reactive jammer by its jamming probability q , which is the probability that the jammer jams an ongoing packet transmission that has not undergone a collision.

Since the jamming probability is independent of the backoff stage, the failure probability is also constant for all backoff stages. Let this probability be p . We obtain:

$$p = P_c + (1 - P_c)q \quad (4)$$

The steady state transmission probability τ is given by the same equation (3). Solving (4), (3), and $P_c = 1 - (1 - \tau)^{n-1}$ yields τ . The probability of success of a given transmission, P_s , is given by $P_s = n\tau(1 - \tau)^{n-1}(1 - q)$, while T_{tr} and T_{id} are DIFS + SIFS + DATA + ACK and σ , respectively.

The above equations in conjunction with Equations of Section III-A give us the throughput of the system. The rate of a reactive jammer with jamming probability q is given by

$$R = \frac{qn\tau(1 - \tau)^{n-1}w}{E[\text{length of a timeslot}]},$$

where w is the length of a jamming pulse.

D. Omniscient Jammers

In this section, we analyze an omniscient jammer that is aware of the current state of each 802.11 node and adopts a jamming strategy that minimizes system throughput subject to constraints on the jamming rate. While a completely omniscient jammer may not be realizable in practice, effective approximations can be implemented (see Section VI for a brief discussion). Furthermore, an accurate analysis of omniscient jammers provides a useful lower bound on the system throughput of 802.11 against all jammers, and also provides a measure for MAC resiliency.

We first make several observations about an optimal omniscient jammer: (a) An optimal omniscient jammer only jams the channel when a transmission of an ACK occurs. (b) An optimal omniscient jammer jams an ongoing transmission only if it incurs no collision. (c) When a transmission is ongoing, the probability with which an optimal omniscient jammer jams the transmission is independent of the particular nodes involved in the transmission. We omit a formal proof of the above three claims owing to space constraints.

1) *Throughput calculation:* We model an omniscient jammer by a *jamming vector* $\vec{q} = (q_0, q_1, q_2, \dots, q_m)$, where q_i is the probability that the jammer jams an ongoing transmission of a node in the i th backoff stage, conditioned on the fact that there is no collision. Given the jamming vector \vec{q} , the throughput of the system and the rate of the jammer can be calculated using the framework of Section III-A.

The failure probability P_i is given by $P_c + (1 - P_c)q_i$ and the product P_s is given by

$$P_i = n \sum_{i=0} M b_{i,0} (1 - P_c) (1 - q_i) \quad (5)$$

The times T_{tr} and T_{id} are DIFS + SIFS + DATA + ACK and σ , respectively. Since the expected length of a timeslot equals $(1 - (1 - \tau)^n)T_{tr} + (1 - \tau)^n\sigma$, the normalized throughput of the system equals

$$\Gamma = \frac{nL \sum_{i=0} M b_{i,0} (1 - P_c) (1 - q_i)}{(1 - (1 - \tau)^n)T_{tr} + (1 - \tau)^n\sigma}$$

The rate of an omniscient jammer with jamming vector \vec{q} is

$$R = \frac{nw \sum_{i=0}^M b_{i,0}(1-P_c)q_i}{(1-(1-\tau)^n)T_{tr} + (1-\tau)^n\sigma},$$

where w is the length of a jamming pulse. The above two equations can be combined to yield

$$\Gamma = \frac{nL(1-P_c)\tau}{(1-(1-\tau)^n)T_{tr} + (1-\tau)^n\sigma} - \frac{LR}{w} \quad (6)$$

In the remainder of this section, we analyze *optimal* rate-constrained omniscient jammers. For convenience, we represent all times as a multiple of σ , and replace T_{tr} by T and σ by 1.

2) *Properties of an optimal omniscient jammer:* Let R denote the rate at which an optimal jammer is jamming the channel. The optimal jammer, constrained by jamming rate R , aims to minimize the total throughput, and is specified by the solution to the following optimization problem.

$$\begin{aligned} & \text{minimize} && \frac{Ln(1-P_c)\tau}{(1-(1-\tau)^n)T_{tr} + (1-\tau)^n\sigma} - \frac{LR}{w} \\ & \text{subject to} && \sum_{i=0}^M \frac{nw b_{i,0}(1-P_c)q_i}{(1-(1-\tau)^n)T_{tr} + (1-\tau)^n\sigma} = R, \end{aligned}$$

The above optimization problem is a complex non-linear program and does not appear to admit a closed-form solution. Nevertheless, we are able to establish a number of interesting properties of an optimal solution for the omniscient jammer.

For the purposes of analysis, we focus our attention on the effect of the jammer on a single node N . Towards this end, we separate out the transmission probability of N as τ_0 , letting τ be the common transmission probability of other nodes.

Lemma 1: For a fixed jammer rate R , the throughput Γ is a monotonic function of τ_0 ; i.e., the sign of the partial derivative $\partial\Gamma/\partial\tau_0$ is independent of τ_0 .

Proof: Expressed as a function of τ_0 , the throughput Γ of the system is given by

$$\frac{L(1-P_c)(\tau_0 + (n-1)\tau)}{(1-(1-\tau)^{n-1}(1-\tau_0))T_{tr} + (1-\tau)^{n-1}(1-\tau_0)\sigma} - \frac{LR}{w}$$

where τ is the transmission probability of any node and R is the jammer rate. Since Γ is of the form $(A\tau_0 + B)/(C\tau_0 + D) + E$ for some terms A , B , C , D , and E , independent of τ_0 , we obtain that $\partial\Gamma/\partial\tau_0$ equals $(BC - AD)/(C\tau_0 + D)^2$, whose sign is independent of τ_0 , completing our proof. ■

We next present the main theorem of this section, which provides a key characterization of an optimal jamming vector.

Theorem 1: For any achievable rate R , there exists an optimal omniscient jammer with rate R which satisfies the following condition: there exists at most one i , $0 \leq i \leq m$, such that q_i lies in the open interval $(0, 1)$.

Proof: Consider an optimal jammer's actions against node N , while keeping the jammer's actions against other nodes fixed. Suppose the jammer is defined by a vector \vec{q} in which q_i and q_j are both in $(0, 1)$. We will analyze the impact of the jammer changing the jamming probabilities q_i and q_j for N while maintaining all other jamming probabilities the same as in \vec{q} ; i.e., the jamming probabilities remain the same for all levels against other nodes, and for all levels $\neq i, j$ against node N . We will prove that q_i and q_j can be changed continuously

(for node N) without increasing the throughput of the system and without changing the total jamming rate R , eventually ending up with two new values, one of which is either 0 or 1. Repeating this argument for all other fractional pairs, and for all nodes, will imply that the optimal strategy can be achieved with values where at most one of the q_i is in $(0, 1)$.

Suppose that q_i and q_j are fractional, with $i < j$, and define for convenience

$$x = p_i = P_c + (1 - P_c)q_i, \quad y = p_j = P_c + (1 - P_c)q_j \quad (7)$$

The pair (x, y) lies in the square $[P_c, 1] \times [P_c, 1]$, since $0 \leq q_i, q_j \leq 1$. We will write $\tau_0(x, y)$ for the transmission probability of N , ignoring the dependence of τ_0 on the other jamming rates q_k which are held constant throughout. In order to keep the total jamming rate R constant we cannot vary x and y independently. The constraint that R is fixed implies a relation between x and y , which we will determine shortly (18). The relation (18) can be solved to give y as a function of x in the interval $[P_c, 1]$, and so the constrained transmission probability is $\tau(x, y(x))$.

By Lemma 1, the throughput is a monotonic function of τ_0 . We note that the transmission probabilities of all other nodes remain fixed. So the relevant question is to determine how τ_0 varies as a function of the jamming probabilities at each level. We will compute the derivative

$$\frac{d}{dx} \tau_0(x, y(x)) \quad (8)$$

and show that either it is identically zero, or else is never zero.

In the first case where (8) is zero, it follows that τ_0 is constant along the curve $(x, y(x))$. The graph $(x, y(x))$ intersects the boundary of $[P_c, 1] \times [P_c, 1]$ at two points. At these points one or both of q_i, q_j is 0 or 1. Therefore by choosing these values in place of the original ones we can reduce the number of fractional values among the jamming probabilities without changing R or τ_0 , and hence Γ , as claimed.

In the second case where (8) is never zero, it follows that τ_0 is strictly monotone along the curve $(x, y(x))$. Since the sign of $\partial\Gamma/\partial\tau_0$ is independent of τ_0 , Γ is smaller at one of the points where this curve intersects the boundary of the square. By choosing the values of q_i, q_j at this point we again reduce the number of fractional values without increasing Γ or R .

It remains to derive the relation (18) and compute the derivative (8). To simplify notation define

$$\gamma_0 = 1, \quad \gamma_k = \prod_{l=0}^{k-1} P_l \quad k = 1, \dots, M \quad (9)$$

where P_l is the probability of a failed transmission at level l . We then have $b_{k,0} = b_{0,0}\gamma_k$ for $i = k, \dots, M$, and from the normalization condition we deduce

$$b_{0,0}^{-1} = \sum_{k=0}^M \gamma_k W_k \quad (10)$$

where $W_k = (2^k W + 1)/2$. Note that γ_k does not depend on

x or y for $k \leq i$, so the right side of (10) can be written

$$b_{0,0}^{-1} = A + \sum_{k=i+1}^M \gamma_k W_k \quad (11)$$

where A is a constant. Now γ_k is a linear function of $x = p_i$ for all $k \geq i + 1$, so we can write (11) as

$$b_{0,0}^{-1} = A + xB + \sum_{k=j+1}^M \gamma_k W_k \quad (12)$$

with another constant B . Finally γ_k is also a linear function of $y = p_j$ for all $k \geq j + 1$, so we end up with the expression

$$b_{0,0}^{-1} = A + xB + xyC \quad (13)$$

with a third constant C . We now consider τ_0 :

$$\tau_0 = \sum_{k=0}^M b_{k,0} = b_{0,0} \sum_{k=0}^M \gamma_k \quad (14)$$

By applying the same reasoning we get the expression

$$\tau_0 = b_{0,0}(H + xJ + xyK) \quad (15)$$

Separating out the jamming component against node N from that against other nodes, we can write the jamming rate R as

$$\frac{(1 - P_c)w \sum_{k=0}^M b_{k,0} q_k + D}{(1 - (1 - \tau)^{n-1}(1 - \tau_0))T_{tr} + (1 - \tau)^{n-1}(1 - \tau_0)\sigma}, \quad (16)$$

where D is a constant (since the jamming probabilities against nodes other than N remain the same). Since $(1 - P_c)q_k = p_k - P_c$, we can apply similar reasoning on the right side of (16) to end up with the expression for R

$$\frac{b_{0,0}(E + xF + xyG)}{(1 - (1 - \tau)^{n-1}(1 - \tau_0))T_{tr} + (1 - \tau)^{n-1}(1 - \tau_0)\sigma} \quad (17)$$

where E, F, G are again constants. The denominator of (17) is of the form $\alpha\tau_0 + \beta$ for constants α and β . Combining (13), (15), and (17) yields the desired relation between x and y :

$$axy = bx + c \quad (18)$$

where a, b, c are constants. Assuming that $a \neq 0$ the solution $y(x)$ of (18) is defined for all x in $[P_c, 1]$, as claimed (we will consider the case where $a = 0$ at the end). Combining (15) with (13) and using (18) to remove the terms with xy we get

$$\tau_0 = \frac{c_1 + c_2x}{c_3 + c_4x} \quad (19)$$

for some constants c_i . The derivative with respect to x is

$$\frac{d}{dx}\tau_0 = \frac{c_2c_3 - c_1c_4}{(c_3 + c_4x)^2} \quad (20)$$

Therefore the derivative is either identically zero (if $c_2c_3 = c_1c_4$) or else is never zero, as claimed.

Finally if $a = 0$ in (18) then y can be freely varied in (15) without changing R , and its derivative with respect to y is either identically zero or never zero. Therefore the value of τ_0 is minimized at either $q_j = 0$ or $q_j = 1$. ■

Intuitively, this suggests that an optimal jammer assigns a certain priority order to the backoff stages of the node N ,

completely jamming transmissions made at a certain backoff stage before jamming transmissions made at a different stage. Our experiments indicate that this is indeed the case. It turns out, however, that this priority order among the backoff stages may depend in subtle ways on the number of nodes n , L (and, thus, T_{tr}), and P_c . We have been able to establish tight characterizations for two important subcases, when $n = 1$, and when the packet sizes are small.

Theorem 2: For $n = 1$, an optimal jamming vector is $(q, 1, 1, \dots, 1, 0)$ or of the form $(1, 1, \dots, 1, q)$.

Proof: When $n = 1$, $P_c = 0$. So, for any positive jamming rate R , $q_0 > 0$. Furthermore, if q_i is zero for some i , we can assume without loss of generality that q_j , for $j > i$, are all zero since the node will never reach backoff stage $i + 1$ or higher.

We first show that $q_M \leq q_i$, for all $i < M$. The proof is by contradiction. Suppose there exists an optimal jammer for which there is an i such that $q_i < q_M$. We can keep the jammer rate fixed by increasing q_i and decreasing q_M appropriately. It is easy to see that τ is unaffected by a change to q_M ; on the other hand, it decreases when q_i increases. By (6), when $n = 1$, Γ is a monotonically increasing function of τ , so we obtain that for an optimal jammer, $q_i \geq q_M$ for all $i < M$.

We next show that if $q_0 = q \in (0, 1)$, then $q_i = 1$ for $0 < i < M$. If the claim is not true, then there exists an $i \in [1, M - 1]$ such that $q_i \neq 1$. Let k be the smallest such i . By Theorem 1, this implies that $q_i = 0$. We now consider an alternative jamming vector that is identical to \vec{q} except that the jamming probability is q^* (to be specified shortly) at level 0 and 1 at level i . We set q^* such that the expected time to return to backoff stage 0 is the same under both vectors.

$$\frac{q^*}{q} = \frac{\sum_{i=1}^k (W_i + 1)/2 + k(T - 1)}{\sum_{i=1}^{k+1} (W_i + 1)/2 + (k + 1)(T - 1)} \quad (21)$$

(Here we use T to denote T_{tr}/σ .) In each case, every packet is eventually successfully transmitted, so the throughput is identical under both jamming vectors. On the other hand, the ratio of the rate of \vec{q}^* and that of \vec{q} is $(k + 1)q^*/(kq)$, which by (21) and the fact that $W_i = 2^i W_0$, is at most one, completing the proof of this case.

The remaining case is when $q_0 = 1$. In this case, we prove that if there exists $i < M$ such that $q_{i-1} = 1$ and $q_i = q \in (0, 1)$, we can set q_i to 1 and q_{i-1} to a suitably chosen q^* without increasing either the throughput or the jammer rate. The argument is similar to the above. Again, we choose q^* such that the expected time to return to backoff stage 0 is identical for both jamming vectors. In particular,

$$q^* = \frac{(W_i + 1)/2 + q(W_{i+1}/2 + 1/2)}{W_i/2 + W_{i+1}/2 + 1}.$$

The difference between the jamming rate of the new vector and that of the old one is proportional to

$$2q^* - q - 1 = (W_{i+1} - W_i)(q - 1)/2 \leq 0,$$

establishing the desired claim and completing the proof. ■

Second, we have studied the case where the packet size is very small, so we can assume that T_{tr} is close to the slot length σ . We omit the proof due to space constraints.

Theorem 3: If $T_{tr} = 1$ and $P_c \leq 0.5$, the jamming vector of an optimal omniscient jammer satisfies the following conditions: $q_i \leq q_{i+1}$, for $0 \leq i < M - 1$, and $q_M \leq q_0$.

IV. SIMULATION EXPERIMENTS

We validate our theoretical analysis for memoryless, reactive and omniscient jamming models by an extensive simulation study. We also investigate and compare the performance of these jammers, as well as the periodic jammer, in terms of their efficiency for different network configurations.

We run our experiments on the Qualnet 3.9.5 simulator [25]. We set up a 1Mbps 802.11 network that satisfies the ideal channel and jamming conditions, and the saturation scenario, discussed in Section III. Towards this end, we locate n sender-receiver pairs in a $300 \times 300 m^2$ area, for varying n , and set their transmission powers to 10.0dBm so that they can hear one another. Each sender has an unbounded queue of packets, so that the network is in a saturated state. We run 802.11 DCF in the basic model with EIFS disabled.

We implement the memoryless and periodic jammers by attaching an exponential and periodic traffic generator, respectively, with an independent jammer node. We emulate the reactive and omniscient jammers by dropping the packets according to the jamming probability of the associated jammer, which is conditioned on the event of no collision. In all our simulations, the jammer is located next to the receiver. For the experiments discussed in the section, we set the jamming pulse width to be $2\mu s$, and the transmission power of the jammer to be sufficient to destroy the packet reception at the receiver while not disturbing the sender.

We first verify the performance of the memoryless jammer for various packet sizes. Figure 2 shows the throughput of one session with packet size varying from 100 bytes to 1500 bytes, under three jamming rates. The simulation results match the theory well. The figure also indicates that there exists a tradeoff between packet size and throughput, for a given jamming rate. Large packet sizes incur less overhead and yield higher throughput in the absence of jamming. However, larger packets are more susceptible to jamming, so when the jamming rate is high, small packet sizes yield higher throughput.

Second, we validate the analysis of memoryless jamming with respect to network size. Figure 3 compares simulation with theory for network sizes from 1 to 50 under 5 different jamming rates. Similarly, we verify the analysis of reactive jamming and omniscient jamming in Figures 4 and 5.

We now compare the effectiveness of different jamming models. Figures 6 and 7 present the performance of the four jammers for two extreme network sizes, 1 session and 50 sessions, respectively, exchanging 500 byte packets. The group of curves which has a higher throughput under no jamming corresponds to the single session case, and the other group to the 50 sessions case. It is easy to see that for a given jamming rate, 802.11 achieves least throughput under omniscient jamming, followed by reactive jamming, then periodic jamming, with memoryless jamming being least effective, requiring at least an order of magnitude higher energy to bring the throughput to zero compared to other jammers. As a general trend, the

gap between the other three jammers decreases with increasing network size. Nevertheless, we observe a significant difference among them, if we analyze the data carefully. Figures 8 and 9 plot the reduction in throughput, as a function of the jamming rate, for periodic, reactive, and omniscient jammers. For a large fraction of the jamming rates, the omniscient jammer reduces the throughput 20-30% more than a reactive jammer and 20-50% more than a periodic jammer.

V. PROTOTYPE EXPERIMENTS

We build a prototype implementation of two channel-oblivious jammers, memoryless and periodic, to compare our theoretical analysis under such jamming with the respective real-world experimentation results. Our prototype testbed uses GNU Radio [4] and USRP [3] to implement such jammers. We observe that the throughput of 802.11 under different jamming parameters and different payload sizes qualitatively match our theoretical analysis. This prototype will serve as a basis for future development of channel-aware jammers.

Prototype: Our experiment setup consists of three nodes: a sender, a receiver and a jammer. We establish a 1 Mbps 802.11b adhoc network between the sender and the receiver and generate a UDP traffic using the *udpclient* and the *udpserver* programs. We select UDP as our transport protocol so that we can focus our attention on MAC layer and avoid issues related to TCP congestion control. Both nodes use Hawking Wireless-B USB network adapters which operate on a ZyDAS chipset.

The jammer is implemented using GNU Radio Software v.3.0.3 on the USRP v.4 (Universal Software Radio Peripheral) platform. Our USRP motherboard uses a 2.4 - 2.5 GHz transceiver/receiver RFX2400 daughterboard along with a vertical (ideal) directional antenna.

We run our experiments with two different placements, based on the position of the jammer relative to the sender and the receiver. Due to space constraints, we only present the results in the case where jammer is close to the receiver. For the jamming pulse width, we ran several experiments with a periodic jammer with different pulse widths. We set the length of the jamming signal to be $22 \mu s$, which is the smallest value at which a pulse is guaranteed to corrupt a packet it hits. (Note that this width is 11 times more than the width used in the plots of Section IV.)

We run our experiments for memoryless and periodic jammers indoors. We scan the network over a period of time and pick a channel with least interference. We also ensure that the jamming signal is strong enough to corrupt every packet it hits.

Parameters and Analysis: We carry out an extensive experimental analysis of the saturation throughput of 802.11 under channel-oblivious jamming. The three major variables that constitute our experiments are: (i) memoryless or periodic jamming; (ii) packet sizes, ranging from 100 bytes to 1400 bytes; (iii) jamming rates, ranging between 0.0018 and 0.022. For a given jamming model and jamming rate, the sender constantly sends packets with different sizes for a duration of 5 seconds. We repeat each this experiment 10 times and report the overall throughput as the average over the 10 runs.

Results: Figures 10 and 11 show the overall experiment results for memoryless and periodic jammers, respectively. Fig-




Fig. 2. Throughput of one IEEE 802.11 session under memoryless jamming with different jamming rates.




Fig. 3. Throughput of multiple IEEE802.11 sessions under memoryless jamming with different jamming rates.




Fig. 4. Throughput of multiple IEEE802.11 sessions under reactive jamming with different jamming rates.




Fig. 5. Throughput of multiple sessions under omniscient jamming with different rates.




Fig. 6. Comparison of the four jammers. Packet size 500 bytes, 1 802.11 session.




Fig. 7. Comparison of four types jammers. Packet size is 500 bytes, 50 sessions.




Fig. 8. Jammer efficiency comparison of omniscient, reactive, and periodic jammers. The network size is 1 and the packet size is 500 bytes.




Fig. 9. Jammer efficiency comparison of omniscient, reactive, and periodic jammers. The network size is 50 and packet size 500 bytes.




Fig. 10. Experimental throughput for different packet sizes under memoryless jammer with various mean jamming rates.

ure 12 compares our experimental results with our theoretical analysis for memoryless jammer, with mean jamming rates 0.0073 and 0.0022. Although the experiments and theory follow a similar trend in both graphs, there is a higher discrepancy in the case where mean jamming rate is 0.0073. Our preliminary investigation suggests that the wireless chipset we use (ZyDAS) does not fully abide by 802.11 specifications in terms of some parameters such as initial backoff window, which lets us get a higher throughput given a busy medium. This issue of differences between commercial adapters and the IEEE 802.11 standard is discussed at length in [26]. Finally, Figure 13 compares our experimental results with simulation results for periodic jammer, with jamming rates 0.0073 and 0.0022. Experimental results here have a better match with theory

in both cases, partly owing to the fact that we use periodic jammers as our reference for selecting jammer parameters such as pulse width.

VI. DISCUSSION AND CONCLUSION

The IEEE802.11 MAC protocol is widely used with support for many physical layers. Given the recent availability of many SDR and sensor networking platforms that make smart jamming relatively easy to build, it is important to understand the limits of IEEE802.11 in the presence of jammers. We have analyzed the saturation throughput performance of IEEE802.11 MAC against several jammers and studied the impact of the jamming rate, packet size, and network size, using mathematical analyses, simulations, as well as a prototype




Fig. 11. Experimental throughput for different packet sizes under periodic jammer with various mean jamming rates.




Fig. 12. Experimental and theoretical throughput for different packet sizes under periodic jammer with jamming rates 0.0073 and 0.0022.




Fig. 13. Experimental and simulation throughput for different packet sizes under periodic jammer with jamming rates 0.0073 and 0.0022.

implementation. We note that while we have focused our attention on saturation throughput, our results on reactive and omniscient jammers qualitatively extend to unsaturated scenarios; indeed, the effectiveness of these jammers only increases if communication occurs in infrequent bursts.

The four jammers we study are about four orders of magnitude more efficient than a continuous jammer. Among these, the memoryless jammer is the least efficient, requiring an order of magnitude more energy to bring the throughput down to zero, when compared to the other three jammers. A periodic jammer is easy to implement and is fairly damaging when the network is saturated. It is significantly less effective than the reactive and omniscient jammers for small packet sizes, low number of active sessions, or unsaturated networks. Reactive jammers can dramatically reduce the throughput of IEEE802.11 with only a limited energy cost on the adversary side. Finally, an optimal omniscient jammer is even 20-30% more effective than a reactive jammer in reducing throughput; it is especially efficient against networks with a small number of active sessions (as would be typical in practice). Our theoretical analysis has identified (though not completely resolved) the key characteristics of an optimal jammer. Our numerical calculations and simulation suggest a natural conjecture on the structure of the jammer, which we have confirmed in special cases.

It would be interesting to completely characterize an optimal jammer for various 802.11 protocol parameters. This would help greatly in the design of anti-jamming techniques. We plan to implement variants of reactive and omniscient jammers using the GNU Radio and USRP testbed. The new USRP-2 platform with embedded processing capabilities will allow a jammer node to sense the channel, keep track of retransmissions, and react quickly to any transmissions. Partially controllable sensor nodes also offer a promising platform for designing reactive and omniscient jammers. Finally, we would like to extend our analysis framework to study the resiliency of other MAC protocols.

REFERENCES

- [1] IEEE, "Medium access control (mac) and physical specifications," *IEEE P802.11/D10*, January 1999.
- [2] G. Bianchi, "Performance analysis of the IEEE802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, 2000.
- [3] *Universal Software Radio Peripheral*. <http://gnuradio.org/trac/wiki/USRPs>.
- [4] *GNU Software Defined Radio*. <http://www.gnu.org/software/gnuradio/>.
- [5] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *ACM MOBIHOC*, 2005.
- [6] M. K. Simon, J. K. Omura, R. A. Scholtz, and B. K. Levitt, *Spread Spectrum Communications Handbook*. McGraw-Hill, 2001.
- [7] D. C. Schleher, *Electronic Warfare in the Information Age*, 1999.
- [8] M. Zapata and N. Asokan, "Secure ad hoc on-demand distance vector routing," *ACM Mobile Comp. and Comm. Review*, 2002.
- [9] Y.-C. Hu, A. Perrig, and D. Johnson, "Ariadne: A secure on-demand routing protocol for ad hoc networks," in *ACM MOBICOM*, 2002.
- [10] G. Lin and G. Noubir, "On link layer denial of service in data wireless lans," *Wiley Journal on Wireless Communications and Mobile Computing*, vol. 5, 2004.
- [11] R. Negi and A. Perrig, "Jamming analysis of MAC protocols," Carnegie Mellon University, Tech. Rep., 2003.
- [12] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: Real vulnerabilities and practical solutions," in *USENIX Security Symposium*, 2003.
- [13] V. B. J. Monks and W. Hwu, "A power controlled multiple access protocol for wireless packet networks," in *IEEE INFOCOM*, 2001.
- [14] C. Ware, T. Wysocki, and J. Chicharo, "On the hidden terminal jamming problem in IEEE 802.11 mobile ad hoc networks," in *IEEE ICC*, 2001.
- [15] S.-G. H. Michael Hall, Aki Silvenoinen, "Effect of pulse jamming on IEEE 802.11 wireless LAN performance," in *IEEE MILCOM*, 2005.
- [16] D. Thunte and M. Acharya, "Intelligent jamming in wireless networks with applications to 802.11b and other networks," in *IEEE MILCOM*, 2006.
- [17] M. Acharya, T. Sharma, D. Thunte, and D. Sizemore, "Intelligent jamming in 802.11b wireless networks," in *OPNETWORK-2004 Conference*, 2004.
- [18] V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein, "Using channel hopping to increase 802.11 resilience to jamming attacks," in *IEEE INFOCOM Minisymposium*, 2007.
- [19] X. Liu, G. Noubir, R. Sundaram, and S. Tan, "SPREAD: Foiling smart jammers using multi-layer agility," in *IEEE INFOCOM Minisymposium*, 2007.
- [20] M. M. Carvalho and J. J. Garcia-Luna-Aceves, "Delay analysis of IEEE 802.11 in single-hop networks," in *IEEE ICNP*, 2003.
- [21] H. Wu, S. Cheng, Y. Peng, K. Long, and J. Ma, "IEEE 802.11 distributed coordination function: enhancement and analysis," *Journal of Computer Science and Technology*, vol. 18, no. 5, 2003.
- [22] "Denial of service in sensor networks," *IEEE Computer*, vol. 35, no. 10, 2002.
- [23] M. Li, I. Koutsopoulos, and R. Poovendran, "Optimal jamming attacks and network defense policies in wireless sensor networks," in *IEEE INFOCOM*, 2007.
- [24] Y. W. Law, L. van Hoesel, J. Doumen, P. Hartel, and P. Havinga, "Energy-efficient link-layer jamming attacks against wireless sensor network MAC protocols," in *SASN '05*, 2005.
- [25] *Scalable Network Technologies*. <http://www.scalable-networks.com/>.
- [26] G. Bianchi, A. D. Stefano, c. Giaconia, L. Scalia, G. Terrazzino, and I. Tinnirello, "Experimental assessment of the backoff behavior of commercial IEEE 802.11b network cards," in *IEEE INFOCOM*, 2007.