Cascading Denial of Service Attacks on Wi-Fi Networks

Liangxiao Xin
Division of Systems Engineering
Boston University
Boston, MA 02215
Email: xlx@bu.edu

David Starobinski
Division of Systems Engineering
Boston University
Boston, MA 02215
Email: staro@bu.edu

Guevara Noubir
College of Computer and Information Science
Northeastern University
Boston, MA 02115
Email: noubir@ccs.neu.edu

***Abstract*—We unveil the existence of a vulnerability in Wi-Fi, which allows an adversary to remotely launch a Denial-of-Service (DoS) attack that propagates both in time and space. This vulnerability stems from a coupling effect induced by hidden nodes. Cascading DoS attacks can congest an entire network and do not require the adversary to violate any protocol. We demonstrate the feasibility of such attacks through experiments with real Wi-Fi cards, extensive ns-3 simulations, and theoretical analysis. The simulations show that the attack is effective both in networks operating under fixed and varying bit rates, as well as ad hoc and infrastructure modes. To gain insight into the root-causes of the attack, we model the network as a dynamical system and analyze its limiting behavior. The model predicts that a phase transition (and hence a cascading attack) is possible when the retry limit parameter of Wi-Fi is greater or equal to 7, and characterizes the phase transition region in terms of the system parameters.**

I. Introduction

Wi-Fi (IEEE 802.11) is a technology widely used to access the Internet. Wi-Fi connectivity is provided by a variety of organizations operating over a shared RF spectrum. These include schools, libraries, companies, towns and government, as well as ISP hotspots and residential wireless routers. Wi-Fi traffic is also rapidly rising due to increased offloading by cellular operators [1]. The importance of Wi-Fi networks and the need to strengthen their resilience to intentional and non-intentional interference have been recognized by companies, such as Cisco [2].

Wi-Fi networks rely on simple, distributed mechanisms to arbitrate access to the shared spectrum and optimize performance. Such mechanisms include carrier sensing multiple access (CSMA), exponential back-offs, and bit rate adaptation. The behavior of these mechanisms in isolated single-hop networks has been extensively studied and is generally well-understood (see, e.g., [3]). However, due to interference coupling, these mechanisms result in complex interactions in multi-hop settings. As a consequence, different networks do not always evolve independently, even if they are located far away.

Figure 1 serves to illustrate this phenomenon at a high level. Suppose that an attacker increases the rate at which it generates packets, and transmits these packets in accordance with the IEEE 802.11 protocol. These transmissions may cause packet collisions at nodes concurrently receiving packets from other sources. Due to the infamous hidden node problem, which is hard to avoid in wireless networks, transmitters may be unable to hear transmission by other nodes, even when using CSMA, and hence keep retransmitting packets until they reach the so-called retry limit of the back-off procedure. These retransmissions affect other neighbours and may propagate.

While an optional mechanism, called RTS/CTS, has been designed to combat the hidden node problem, it increases overhead and latency especially at high bit rates. Since the cost of the RTS/CTS exchange usually does not justify its benefits, it is commonly disabled [4], [5]. Indeed, most manufacturers of Wi-Fi cards disable RTS/CTS by default and discourage changing this setting [6]–[9].

The coupling phenomenon induced by interferences creates multi-hop dependencies, which an adversary can take advantage of to launch a widespread network attack from a single location. We refer to such an attack as a *cascading Denial-of-Service (DoS) attack*. Cascading DoS attacks are especially dangerous because they affect the entire network and do not require the adversary to violate any protocol (i.e., the attacks are protocol-compliant).

The contributions of this paper are as follows. First, we unveil the existence of a vulnerability in the IEEE 802.11 standard, which allows an attacker to launch protocol-compliant cascading DoS attacks. In contrast to existing jamming attacks, the attacker does not need to be in the vicinity of the victims.

Second, we provide a concrete attack that exploits this vulnerability in certain network scenarios. We demonstrate the attack through experiments on a testbed composed of nodes equipped with real Wi-Fi cards, and through extensive ns-3 simulations.

Third, we show the existence of a phase transition. When the packet generation rate of the attacker is lower than the phase transition point, it has vanishing effect on the rest of the network. However, once the packet generation rate exceeds the phase transition point, the network becomes entirely congested.

Finally, we develop a new analytical model that sheds light into the phase transition observed in the simulations and experiments. The analysis predicts for which values of the retry limit a phase transition (and hence a cascading attack) can occur, and explicitly characterizes the phase transition region in terms of the system parameters. In particular, we show that

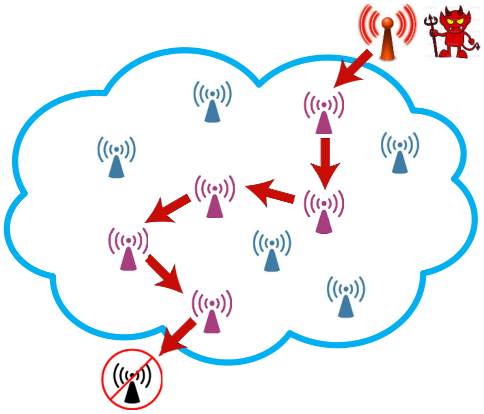

Fig. 1: Illustration of a cascading denial of service attack. Transmissions by an attacker impact nodes located far away, due to interference coupling caused by hidden nodes.

a phase transition can occur for the default value of the retry limit in Wi-Fi, which is 7.

The rest of the paper is organized as follows. In Section II, we discuss related work. In Section III, we provide brief background on Wi-Fi and hidden nodes, and introduce our network model. We present and discuss experimental and simulation results in Section IV. In Section V, we present an analytical model that explains the behavior of the network and the impact of various parameters, and compare the analytical and simulation results. In Section VI, we conclude the paper and discuss possible mitigation methods. Due to space limitation, most proofs are omitted. They can be found in [10].

II. Related Work

In general, the main goal of a DoS attack is to make communication impossible for legitimate users. Within the context of wireless networks, a simple and popular means to launch a DoS attack is to jam the network with high power transmissions of random bits, hence creating interferences and congestion. Jamming at the physical layer, together with *anti-jamming* countermeasures, have been extensively studied (cf. [11] for a monograph on this subject).

More recently, several works have developed and demonstrated *smart jamming* attacks. These attacks exploit protocol vulnerabilities across various layers in the stack to achieve high jamming gain and energy efficiency, and a low probability of detection [12]. For instance, [13] shows that the energy consumption of a smart jamming attack can be four orders of magnitude lower than continuous jamming. The works in [14], [15] show that several Wi-Fi bit rate adaptation algorithms, such as SampleRate, ONOE, AMRR, and RARF, are vulnerable to smart jamming. However, both conventional and smart jamming attacks are usually non-protocol compliant. Moreover, they require physical proximity. These limitations can be used to identify and locate the jammer.

In contrast, in this work we show how a protocol-compliant DoS attack can be remotely launched by exploiting coupling due to hidden nodes in Wi-Fi. Rate adaptation algorithms further amplify this attack due to their inability to distinguish between collisions, interferences, and poor channels. One potential mitigation is to design a rate adaptation algorithm whose behaviour is based on the observed interference patterns [16], [17]. However, to the best of our knowledge, none of these rate adaptation algorithms are used in practice. Our work is based on Minstrel [18], which is the most recent, popular, and robust rate adaptation algorithm for Linux systems.

The attacks that we are investigating bear similarity to cascading failures in power transmission systems [19], [20]. When one of the nodes in the system fails, it shifts its load to adjacent nodes. These nodes in turn can be overloaded and shift their load further. This phenomenon has also been studied in wireless networks. For instance, [21], [22] model wireless networks as a random geometric graph topology generated by a Poisson point process. They use percolation theory to show that the redistribution of load induces a phase transition in the network connectivity. However, the cascading phenomenon that we investigate in this paper is different from cascading failure studied in those works. In our work, the exogenous traffic arriving at each node is independent. That is, a node will not shift its load to other nodes. The amount of traffic measured on the channel increases due to packet retransmissions caused by packet collisions, rather than due to traffic redistribution.

The work in [23], [24] show that local coupling due to interferences can have global effects on wireless networks. Thus, [23] proposes a queuing-theoretic analysis and approximation to predict the probability of a packet collision in a multi-hop network with hidden nodes. It shows that the sequence of the packet collision probabilities in a linear network converges to a fixed point. The work in [24] evaluates the impact of rate adaption and finds out that traffic increase at a single node can congest an entire network, and points out the existence of a phase transition.

Our paper differs in several aspects. First, it considers an adversarial context, and shows how interference-induced coupling can be exploited to cause denial of service. Second, to our knowledge, it is the first work to demonstrate the existence of such coupling on real commodity hardware. Third, our simulations are based on a high-fidelity wireless simulator (ns-3), capable of capturing the effects of rate adaptation algorithms and accurately modeling infrastructure networks. Finally, our analytical model is original and captures the impact of the retry limit and traffic parameters. A key result is that a cascading attack can be launched for the default value of the retry limit in Wi-Fi, a result validated by the experiments and simulations.

III. Background and Model

We first review key aspects of IEEE 802.11 and then describe the network model under consideration.

A. Wi-Fi Summary

Wi-Fi is a wireless local area network (WLAN) technology, which mainly runs on 2.4 GHz ISM bands and 5 GHz bands [5]. The IEEE 802.11 standard is a series of specifications, such as the media access control (MAC) and physical layer (PHY)

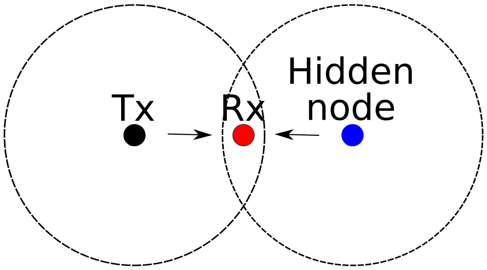

Fig. 2: Classical hidden node problem. The transmitter and the hidden node cannot sense each other. The collision happens when they transmit simultaneously.

interfaces. The first 802.11 standard that gained widespread success is 802.11b. It runs on 2.4 GHz bands and has up to 11 Mb/s bit rate. The subsequent standards (e.g., 802.11a, g, n, and ac) increased the bit rates using higher order modulation along with coding, OFDM, MIMO, and wider bands. It is noteworthy that 802.11b is the only mode that supports communication at 1 Mb/s. Hence, when the bit rate reduces to 1 Mb/s, Wi-Fi network reverts to the 802.11b mode. Generally, this lower bit rate has higher resistance to interference during transmission and is able to operate over lower SNR channels.

The IEEE 802.11 standard uses a CSMA/CA mechanism to control access to the transmission medium and avoid collisions. After a packet is sent, the node waits for a short interframe slots (SIFS) period to receive an ACK. Whenever the channel becomes idle, nodes wait for a distributed interframe space (DIFS $>$ SIFS) period to start contending for the channel.

B. Hidden Node Problem

A typical instance of the hidden node problem is illustrated in Figure 2. The figure shows three nodes: a transmitter, a receiver and a hidden node. The dashed circle represents the transmission range of the node. Since the transmitter and the hidden node cannot sense each other, a collision happens when both of them transmit packets at the same time.

A packet collision triggers a retransmission. In IEEE 802.11, there is an upper limit on the number of retransmissions that a packet can incur, called *retry limit* and denoted by R (the default value is $R = 7$). If the retry count of a packet exceeds the retry limit, the packet is dropped. The channel utilization of a node increases with the probability of a packet collision. In the worst case, the utilization can be R times larger than in the absence of packet collisions. Therefore, the access channel of a node can easily be saturated if it is forced to retransmit packets.

C. Network Model

The network model considered in this paper is shown in Figure 3. We consider N pairs of nodes. Each node A_i ($i = 0, 1, 2, \ldots, N$) transmits packets to node B_i. The dashed circle represents the range of transmission. Node B_{i+1} can receive packets from both node A_i and node A_{i+1}. However, node A_i and node A_{i+1} cannot hear each other. That is, node A_i is a hidden node with respect to node A_{i+1} (and vice-versa). A packet collision happens at node B_{i+1} when packet transmissions by node A_i and A_{i+1} overlap.

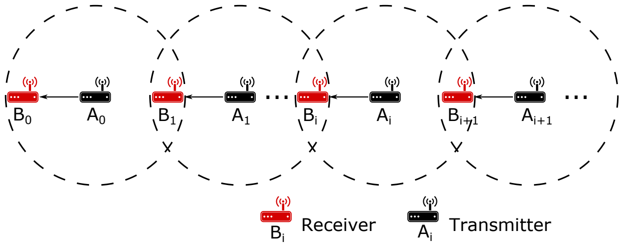

Fig. 3: Topology of the network. Node A_i transmits packets to node B_i. Node A_i is a hidden node with respect to A_{i+1}.

Our goal is to investigate how node A_0 can trigger a cascading DoS attack, resulting in a congestion collapse over the entire network. We start by increasing the packet generation rate at node A_0. Node A_0 transmits packets over its channel, in compliance with the IEEE 802.11 standard. The transmissions by node A_0 cause packet collisions at node B_1. These collisions require node A_1 to retransmit packets. The increased amount of packet transmissions and retransmissions by node A_1 impact node A_2 and so forth. If this effect keeps propagating and amplifying, then the result is a network-wide denial of service, which we refer to as a cascading Denial of Service (DoS) attack. Because this attack is protocol-compliant, it is difficult to detect or trace back to the initiator. We note that the topology studied here could arise over different time and space in more complex network configurations.

IV. Experimental and Simulation Results

In this section, we demonstrate the feasibility of launching cascading DoS attacks both through experiments and simulations. We first show results on an experimental testbed using real Wi-Fi cards. We then use ns-3 simulations to investigate how this attack can be performed in significantly larger scale networks, and under different settings (ad hoc, infrastructure, fixed bit rate, and adaptive bit rate).

A. Experiments

We set up an experimental testbed composed of six nodes. The testbed configuration is shown in Figure 4. We establish an IEEE 802.11n ad hoc network consisting of three pairs of nodes. Each node consists of a PC and a TP-LINK TL-MN722N Wireless USB Adapter. We use RF cables and splitters to link the nodes, isolate them from external traffic, and obtain reproducible results.

We place 70 dB attenuators on links between node A_i and B_i ($i \in 0, 1, 2$), and 60 dB attenuators on links between nodes A_i and B_{i+1}. The difference in the signal attenuation of different links insures that a packet loss occurs if a hidden node transmits. In practice, such a situation may occur if nodes A_i and B_{i+1} communicate without obstacles, while node A_i and B_i are separated by an office wall [25]. The transmission power of each node is set to 0 dBm. We use iPerf [26] to generate

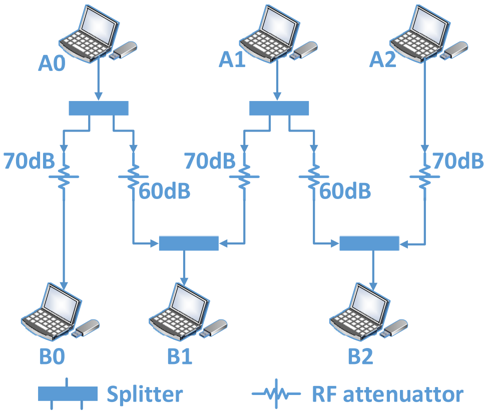

Fig. 4: Experimental testbed.

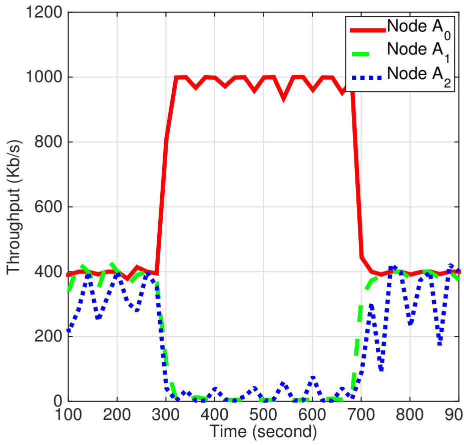

Fig. 5: Throughput performance measurements in testbed. When node A_0 starts increasing its packet generation rate, the throughput of nodes A_1 and A_2 vanishes.

UDP data streams and to measure the throughput achieved on each node. The length of a packet is the default IP packet size of 1500 bytes.

Figure 5 demonstrates the cascading DoS attack on the experimental testbed. At first, the packet generation rates of nodes A_0, A_1 and A_2 are set to 400 Kb/s. We observe that the throughput of all the nodes remains in the vicinity of 400 Kb/s during the first 300 seconds. After 300 seconds, A_0 starts transmitting packets at 1 Mb/s. As a result, the throughput of nodes A_1 and A_2 suddenly vanishes. Once node A_0 resumes transmitting at 400 Kb/s, the throughput of node A_1 and node A_2 recovers.

B. *Simulations*

In the previous section, we demonstrated the feasibility of launching a cascading DoS attack on an experimental testbed. This testbed relies on commercial cards that are black boxes for all purposes. For instance, the driver of the Wi-Fi card and the rate adaptation algorithm are closed-source. There are also substantial usage restrictions, such as parameter settings.

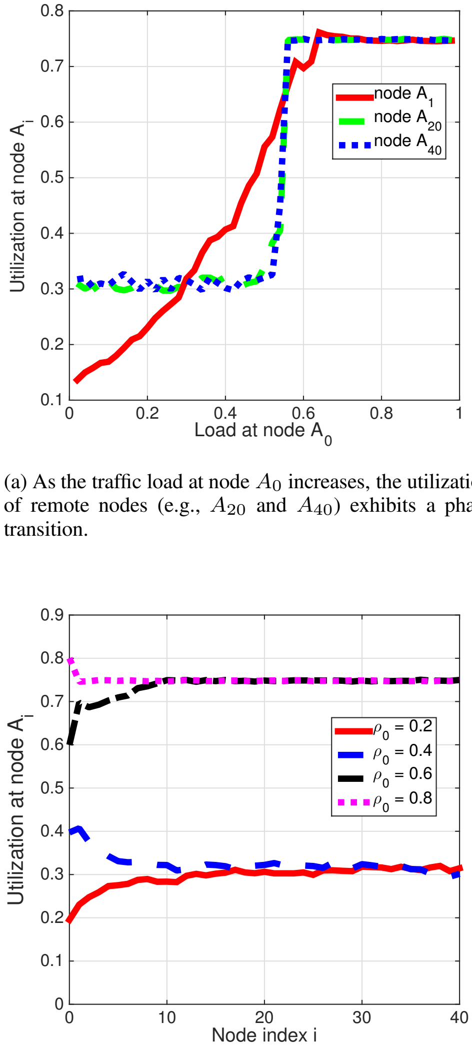

(a) As the traffic load at node A_0 increases, the utilization of remote nodes (e.g., A_{20} and A_{40}) exhibits a phase transition.

(b) Utilization of nodes A_i $(i \geq 1)$ for different traffic loads at node A_0. The utilization converges as i gets large. When the load at node A_0 changes from 0.4 to 0.6, the sequence of utilization converge to different limits, illustrating the phase transition.

Fig. 6: The occurrence of cascading DoS attacks in ad hoc networks with fixed bit rate.

In order to gain a better insight into the attack in large-scale networks, we resort to ns-3 simulations, a state-of-the-art simulator which includes high-fidelity wireless libraries. We show the occurrence of cascading DoS attacks:

1) In ad hoc networks with fixed bit rate.
2) In ad hoc networks under Minstrel rate adaptation.
3) In infrastructure networks.

1) Fixed bit rate: We first describe the occurrence of a cascading DoS attack in an ad hoc network with fixed bit rate. We consider a linear topology consisting of 40 pairs of nodes (i.e. a sequence of 40 hidden nodes), as shown in Figure 3. We fix the bit rate to 1 Mb/s and the retry limit to $R = 7$.

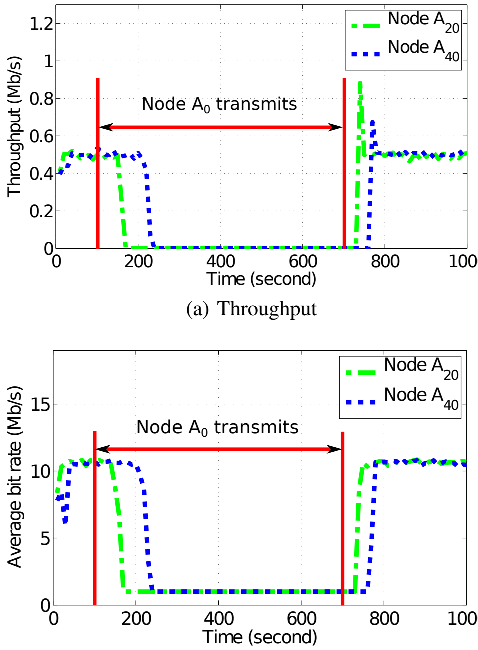

(a) Throughput

(b) Bit rate

Fig. 7: Simulation results with Minstrel rate adaptation. When node A_0 generates packets at 5 Mb/s and transmits, the throughput of nodes A_{20} and A_{40} vanishes. The average bit rates of nodes A_{20} and A_{40} also reduce to 1 Mb/s. This result indicates that nodes A_{20} and A_{40} are transmitting packets at the lowest bit rate, however with no throughput (all their packets collide).

We set up a Wi-Fi network using the standard IEEE 802.11 library in ns-3. At each node A_i, $i \geq 1$, UDP packets arrive at rate $\lambda_i = 8.125$ pkts/s. The packet generation rate at node A_0, λ_0, varies from 1.25 to 61.25 pkts/s. The size of each packet is 2000 bytes. Each node has the same transmission power (40 mW). We set the propagation loss between node A_i and B_i to 80 dB and the propagation loss between node A_i and B_{i+1} to 70 dB. We run each simulation five times for 1,000 seconds, and average out the results.

The *(exogenous) load* at each node A_i is denoted $\rho_i = \lambda_i T$, where T represents the duration of each packet transmission attempt (0.016 second in our case). The *utilization* of a node A_i, denoted u_i, is defined as the fraction of time the node is busy transmitting bits on the channel.

Figure 6 depicts the utilization u_1, u_{20}, and u_{40} as a function of ρ_0, the load at node A_0. The utilization of node A_1, u_1, increases smoothly until it reaches its upper limit. However, the utilizations of nodes A_{20} and A_{40} remain low until u_0 reaches a certain threshold around $\rho_0 = 0.5$, at which point u_{20} and u_{40} suddenly jump to a high value. This sudden jump corresponds to a phase transition, and the critical threshold represents the phase transition point.

2) Rate Adaptation: We next consider the same network setting as in the previous section, but this time we assume that nodes can transmit at different bit rates. We specifically assume that nodes implement the Minstrel rate adaptation algorithm. In this case, the attack works by coercing the rate adaptation algorithm to reduce the bit rate to 1 Mb/s at each node, thus leading to similar results to those shown in Section IV-B1. In our simulations, the parameter *EWMA* of Minstrel is set to 0.25. [27]

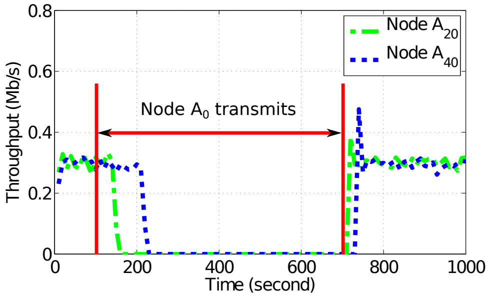

Fig. 8: Simulation results under AP mode. Nodes A_i are stations and nodes B_i are access points, for $i \in \{0, 1, 2, \dots\}$. When node A_0 generates packets at 5 Mb/s and transmits, the throughput of nodes A_{20} and A_{40} vanishes.

We set $\lambda_0 = 312.5$ pkts/s and $\lambda_i = 31.25$ pkts/s ($i \geq 1$) for the packet generation rates. As shown in Figure 7, packet transmissions at node A_0 start after $t = 100$ s. During the first 100 seconds, the throughput of nodes A_{20} and A_{40} remain around 0.5 Mb/s, which implies that all the packets are received. Once node A_0 starts transmitting packets, the throughput of nodes A_{20} and A_{40} is brought down to close to zero. We also observe that the bit rates at node A_{20} and A_{40} go down to 1 Mb/s, due to the repeated packet collisions. Once node A_0 stops transmitting at $t = 700$ s, nodes A_{20} and A_{40} recover.

3) Infrastructure networks: We next show that cascading DoS attacks are also feasible in infrastructure networks. Since the infrastructure mode is more widely used than ad hoc in practice, the feasibility of the cascading DoS attack in infrastructure networks increases its severity and potential impact. We repeat the simulations of Section IV-B2 except that we set nodes B_i as access points, and nodes A_i as stations.

Figure 8 shows similar results as in Section IV-B2. When the cascading DoS attack starts, the remote nodes in the sequence exhibit a phase transition. However, we observe a larger lag in the propagation of the DoS attack compared to ad hoc networks.

V. Analysis

In this section, we develop an analytical model that provides insight into the network behaviour observed in the simulations and experiments. Specifically, our goal is to explain why and under what conditions the phase transition occurs, and shed light into the roles played by the retry limit R and the traffic load at the different nodes.

A. Model

We consider the linear topology shown in Figure 3. Packet arrivals at each node A_i form a Poisson process with rate λ_i. The packet size is fixed and the duration of each packet transmission attempt is T (we assume a fixed bit rate). A transmission by node A_{i+1} is successful only if does not overlap with any transmission by (hidden) node A_i.

If a packet collides, it is retransmitted until either it is successfully received or the retry count reaches the limit R. Let $\overline{r}_i$ represent the mean retry count at node A_i. Then, the mean service time of a packet at node A_i is $\overline{r}_i T$. To keep the analysis tractable, timing details of Wi-Fi, such as DIFS, SIFS, and back-off inter-frame spacings are ignored.

We denote the utilization of node A_i by $0 \leq u_i \leq 1$, where u_i represents the fraction of time node A_i transmits. If $u_i = 1$, node A_i is congested and transmits continuously. Otherwise, node A_i is uncongested and transmits packets at rate $\overline{r}_i \lambda$. Therefore, the utilization of node A_i for all $i \geq 0$ is

$$u_i = \min\{\overline{r}_i \lambda_i T, 1\}. \tag{1}$$

Note that $\overline{r}_0 = 1$.

Our model represents a special case of interacting queues, which are notoriously difficult to analyze [28]. To make the analysis tractable, we *assume* that

1) Packet transmissions and retransmissions at each uncongested node A_i form a Poisson process with rate $\overline{r}_i \lambda$.
2) The probability that a packet transmitted by node A_i collides is independent of previous attempts. This probability is denoted p_i.

The above assumptions are similar to the "random-look" model used by Kleinrock and Tobagi in their analysis of (single hop) random access networks [29]. We stress that beside these assumptions, the rest of our analysis is exact.

B. Iterative analysis of the utilization

Our goal is to find the utilization at each node $i \geq 0$ and in the limit as $i \to \infty$. We consider the same scenario as in our simulations, whereby node A_0 (the attacker) varies its traffic load

$$\rho_0 \triangleq \lambda_0 T, \tag{2}$$

while all other nodes A_i $(i \geq 1)$ have the same traffic load

$$\rho \triangleq \lambda_i T, \tag{3}$$

where $0 < \rho < 1$. We aim to understand if and how changes in the value of ρ_0 affect the utilization of nodes that are located far away as a function of the parameters ρ and R.

First, we get the utilization at node A_0:

$$u_0 = \min\{\rho_0, 1\}. \tag{4}$$

We next develop an iterative procedure to derive u_{i+1} from u_i. From (1) and (3),

$$u_{i+1} = \min\{\overline{r}_{i+1} \rho, 1\}. \tag{5}$$

We first relate $\overline{r}_{i+1}$ to p_{i+1}, the probability that a packet transmitted by node A_{i+1} collides. Based on Assumption 2, the probability that a packet is successfully received after $1 \leq r \leq R$ attempts is $(1 - p_{i+1})(p_{i+1})^{r-1}$ while the probability that a packet fails to be received after R attempts is $(p_{i+1})^R$. Hence, the mean retry count at node A_{i+1} is

$$\begin{aligned} \overline{r}_{i+1} &= \sum_{r=1}^{R} r \cdot (1 - p_{i+1}) \cdot (p_{i+1})^{r-1} + R \cdot (p_{i+1})^R \\ &= \sum_{r=1}^{R} (p_{i+1})^{r-1}. \end{aligned} \tag{6}$$

We next relate p_{i+1} to u_i. First, suppose $u_i < 1$ (i.e., node A_i is uncongested). WLOG, assume that node A_{i+1} starts a packet transmission (or retransmission) at time $t = 0$. We compute p_{i+1} by conditioning on whether or not node A_i is transmitting at time $t = 0$. Note that due the Poisson Arrivals See Time Averages (PASTA) property, the transmission state of node A_i at time $t = 0$ is the same as at any random point of time.

If node A_i transmits at time $t = 0$, which occurs with probability u_i, then the packet transmitted by node A_{i+1} collides. If node A_i does not transmit at time 0, then a collision occurs only if node i starts a transmission during the interval $[0, T]$. Such an event occurs with probability $(1 - \mathrm{e}^{-\overline{r}_i \lambda_i T}) = (1 - \mathrm{e}^{-u_i})$, according to Assumption 1. We therefore obtain

$$\begin{aligned} p_{i+1} &= 1 \cdot u_i + (1 - \mathrm{e}^{-u_i}) \cdot (1 - u_i) \\ &= 1 - \mathrm{e}^{-u_i}(1 - u_i). \end{aligned} \tag{7}$$

Next, suppose $u_i = 1$ (i.e., node A_i is congested). In that case, all the transmissions by node A_{i+1} collide and $p_{i+1} = 1$. We note that (7) still provides the correct result.

Putting (5), (6), and (7) together, we obtain

$$u_{i+1} = \min\left\{\rho \sum_{r=1}^{R} \left(1 - \mathrm{e}^{-u_i}(1 - u_i)\right)^{r-1}, 1\right\}. \tag{8}$$

C. Limiting behavior of the utilization

We next analyze the limiting behaviour of the iteration given by (8). The sequence $(u_i)_{i=0}^{\infty}$ corresponds to a discrete non-linear dynamical system [30]. Such systems are generally complex as they may converge to a point, to a cycle (i.e., they exhibit periodic behavior), or not converge at all (i.e., they exhibit chaotic behavior).

The main result of this section is to show that the sequence $(u_i)_{i=0}^{\infty}$ always converges to a point. However, the limit depends on the initial utilization u_0.

To simplify notation, we define the function

$$f(u_i) \triangleq \rho \sum_{r=1}^{R} \left(1 - \mathrm{e}^{-u_i}(1 - u_i)\right)^{r-1}. \tag{9}$$

We then rewrite (8) as follows:

$$u_{i+1} = \min\{f(u_i), 1\}. \tag{10}$$

We say that $\omega \in [0, 1]$ is a *fixed point* of (10) if

$$\omega = \min\{f(\omega), 1\}. \tag{11}$$

Suppose (11) has K different fixed points (Theorem 2 in the sequel will show that $K \geq 1$). We denote by Ω the ordered set of all the fixed points of (11). That is,

$$\Omega \triangleq \{\omega_1, \ldots, \omega_k, \ldots, \omega_K\}, \tag{12}$$

where $\omega_1 < \ldots < \omega_k < \ldots < \omega_K$.

The following theorem states that the limit of the sequence $(u_i)_{i=0}^{\infty}$ is always one of the elements in Ω.

Theorem 1:

1) Let $u_0 \in (\omega_k, \omega_{k+1})$, where $k \in \{1, \ldots, K-1\}$. If $f(u_0) > u_0$, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_{k+1}. If $f(u_0) < u_0$, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_k.
2) If $u_0 \in [0, \omega_1)$, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_1.
3) If $\omega_K < 1$ and $u_0 \in (\omega_K, 1]$, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_K.

D. Phase transition analysis

In the previous section, we showed that the limit of the sequence of node utilizations $(u_i)_{i=0}^{\infty}$ must be one of the fixed points in the set Ω. A phase transition represents a situation where a small change of u_0 leads to an abrupt change of the limit. Specifically, we focus on the case when the limit jumps to 1. Formally:

Definition 1 (Network congestion): A network is said to be *congested* if $(u_i)_{i=0}^{\infty}$ converges to 1. Else, the network is said to be *uncongested*.

Definition 2 (Phase transition): A network experiences a phase transition if there exists a fixed point $\omega \in \Omega$, such that if $u_0 < \omega$ the network is uncongested, and if $u_0 > \omega$ the network is congested. We refer to ω as the phase transition point.

We note that a phase transition can possibly occur only if $\omega_K = 1$, since otherwise the network is never congested, irrespective of u_0.

A network must fall in one of the following three regimes:

1) The network is uncongested for all $u_0 \in [0, 1]$.
2) The network is congested for all $u_0 \in [0, 1]$.
3) A phase transition occurs.

Our goal in the following is to determine what regime prevails under different network parameters.

For this purpose, we investigate the existence and properties of solutions of (11). First, we investigate the case $\omega = 1$.

Lemma 1: If $\rho > 1/R$, then

1) $\omega_K = 1$.
2) If $K = 1$, then for all $u_0 \in [0, \omega_K]$ the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_K.
3) If $K \geq 2$, then for all $u_0 \in (\omega_{K-1}, \omega_K]$ the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_K.

Lemma 1 indicates that the sequence $(u_i)_{i=0}^{\infty}$ can converge to 1 (depending on u_0), if $\rho > 1/R$. Besides this special case, (11) can be rewritten

$$f(\omega) = \omega. \tag{13}$$

We look for solutions of (13) that belong to the interval $[0, 1]$. Each such solution is an element of Ω.

Equation (13) is difficult to work with because it contains two unknown variables, ρ and R. To circumvent this difficulty, we introduce the function

$$h_R(\omega) \triangleq \frac{\rho\omega}{f(\omega)} = \frac{\omega}{\sum_{r=1}^{R} \left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)^{r-1}}. \tag{14}$$

For each value of ρ, the solutions of (13) must satisfy

$$h_R(\omega) = \rho. \tag{15}$$

We denote the maximum of $h_R(\omega)$ by

$$h_R^{max} \triangleq \max_{0 \leq \omega \leq 1} h_R(\omega).$$

The following theorem establishes the prevailing network regimes for different parameters.

Theorem 2:

1) If $\rho < 1/R$, then the network is uncongested for all $u_0 \in [0, 1]$.
2) If $h_R^{max} > 1/R$ and $1/R < \rho < h_R^{max}$, then a phase transition occurs and the phase transition point is ω_{K-1}.
3) If $\rho > h_R^{max}$, then the network is congested for all $u_0 \in [0, 1]$.

Proof:

1) If $\rho < 1/R$, then $R\rho < 1$ and the utilization of each node is always less than 1. Hence, for any $u_0 \in [0, 1]$, the network is always uncongested. Note that since $h_R(0) = 0$, $h_R(1) = 1/R$, and h_R is continuous, (15) must have at least one solution (i.e., at least one fixed point exists).
2) Let $\rho \in (1/R, h_R^{max})$. We know that $h_R(0) = 0$ and $h_R(1) = 1/R$. Since the function h_R is continuous, (15) must have at least one solution (i.e, at least one fixed point strictly smaller than 1 exists). Also, because $\rho > 1/R$, a fixed point point at $\omega = 1$ exists (i.e., $\omega_K = 1$), by Part 1 of Lemma 1. Thus, there are $K \geq 2$ fixed points.
   By Part 3 of Lemma 1, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_K for all $u_0 \in (\omega_{K-1}, \omega_K]$. Moreover, by Theorem 1, the limit of the sequence $(u_i)_{i=0}^{\infty}$ is no larger than ω_{K-1} for all $u_0 \leq \omega_{K-1}$. Hence, a phase transition exists at ω_{K-1}.
3) If $\rho > h_R^{max}$, then (13) has no solution. Moreover, since $\rho > h_R^{max} \geq h_R(1) = 1/R$, we get $\rho > 1/R$. By Parts 1 and 2 of Lemma 1, the sequence $(u_i)_{i=0}^{\infty}$ converges to 1 for any $u_0 \in [0, 1]$, and the network is always congested.

■

We next illustrate Theorem 2 for different values of R, using Figure 9. First, consider $R = 4$ as shown in Figure 9(a). Since $h_R^{max} = 1/R = 0.25$, there exists no traffic load ρ for which a phase transition exists. Either the network is always uncongested (for $\rho < 1/R$), or it is always congested (for $\rho > 1/R$).

Next, consider $R = 7$ as shown in Figure 9(b). There, $h_R^{max} = 0.166 > 1/R = 0.143$. Hence, a phase transition occurs if $\rho \in (0.143, 0.166)$. For instance, consider the case $\rho = 0.15$. Then, the equation $h_R(\omega) = \rho$ has two solutions. Including the fixed point $\omega = 1$ (since $\rho > 1/R$), the set Ω has $K = 3$ fixed points: $\{\omega_1 = 0.265, \omega_2 = 0.777, \omega_3 = 1\}$. Hence, by Theorem 2, the network is uncongested if $u_0 < 0.777$, and congested if $u_0 > 0.777$.

The case $R = 10$ also has a phase transition region, as shown in Figure 9(c). Furthermore, the size of this region is larger since $(1/R, h_R^{max}) = (0.1, 0.162)$.

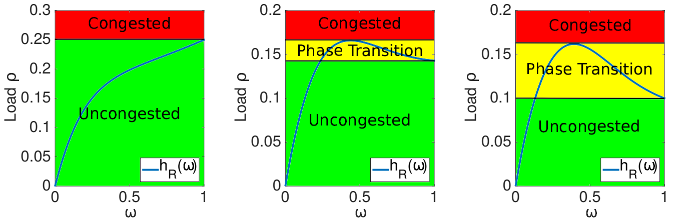

(a) $R = 4$ (b) $R = 7$ (c) $R = 10$

Fig. 9: Illustration of the different network regimes for different values of the retry limit R. For each value of the load ρ, the fixed points are the solutions of $h_R(\omega) = \rho$. In addition, the fixed point $\omega = 1$ always exists when $\rho > 1/R$. A phase transition region exists if the maximum of $h_R(\omega)$, h_R^{max}, is strictly greater than $h_R(1) = 1/R$.

E. Sufficient condition for phase transition

In the previous section, we showed that a phase transition exists in the region $1/R < \rho < h_R^{max}$, if $h_R^{max} > 1/R$. In this section, we derive an explicit lower bound on h_R^{max}, which provides a simple condition for the existence of a phase transition. First, we establish a relationship between the derivatives of $h_R(\omega)$ for different values of R, but a given value of ω.

Lemma 2: For $\omega \in [0, 1]$, if there exists $R^* \geq 1$ such that $h'_{R^*}(\omega) \leq 0$, then $h'_R(\omega) \leq 0$ for all $R > R^*$.

Consider the function $h_R(\omega)$ as $R \to \infty$:

$$\begin{aligned} h_\infty(\omega) &= \left(1 - \left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)\right)\omega \\ &= e^{-\omega}(1-\omega)\omega, \end{aligned} \tag{16}$$

and its derivative

$$h'_\infty(\omega) = e^{-\omega}(1 - 3\omega + \omega^2). \tag{17}$$

The next corollary is the logical transposition of Lemma 2.

Corollary 1: If $h'_\infty(\omega) \geq 0$, then $h'_R(\omega) \geq 0$ for all $R \geq 1$.

The following lemma establishes that the function $h_R(\omega)$ is always strictly increasing in the interval $[0, \varpi)$, where

$$\varpi \triangleq \frac{3 - \sqrt{5}}{2}. \tag{18}$$

Lemma 3: Let $0 \leq \omega < \varpi$. Then, $h'_R(\omega) > 0$, for all $R \geq 1$.

The consequence of Lemma 3 is that for all $R \geq 1$,

$$h_R^{max} \geq h_R(\varpi). \tag{19}$$

This equation provides a lower bound on h_R^{max} that can easily be computed. The next theorem establishes an even more explicit lower bound on h_R^{max}.

Theorem 3: Let $h_\infty(\omega)$ and ϖ be defined as in (16) and (18), respectively. Then, $h_R^{max} \geq h_\infty(\varpi) \simeq 0.161$.

From Theorems 2 and 3, it follows that a phase transition exists if $1/R < 0.161$. Hence:

Corollary 2: A phase transition is guaranteed to exist for $R \geq 7$ and $\rho \in [1/R, 0.161]$.

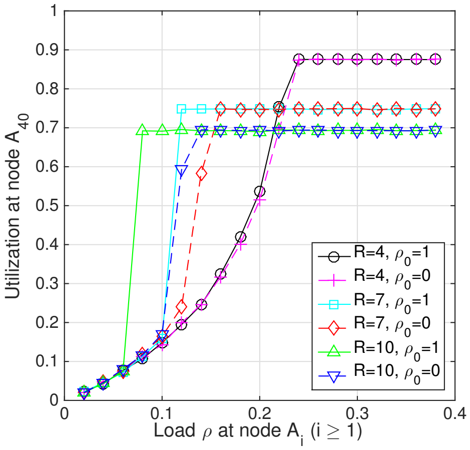

Fig. 10: Simulation of the limiting behavior of the node utilization in a network of 40 nodes. For $R = 4$, the limit is the same when $\rho_0 = 0$ and $\rho_0 = 1$, hence no phase transition is observed. However, for $R = 7$ and $R = 10$, the limits are different, hence showing the existence of a region of load ρ in which a phase transition occurs.

We note that the lower bound on h_R^{max} is quite tight. For instance, $h_7^{max} = 0.166$. Moreover, h_R^{max} decreases with R (this follows from (14), since for any $\omega \in [0, 1]$ the denominator increases as R gets larger).

F. Comparison with simulation results

We compare the results of our analysis with ns-3 simulations, for different settings of the retry limit R and load ρ. For the simulations, we consider an ad hoc network composed of 40 pairs of nodes, as described in Section IV-B1. To check whether a phase transition exists for a given R, we run simulations both for $\rho_0 = 0$ and $\rho_0 = 1$. If the node utilizations in the limit (i.e., for node A_{40}) is the same in both cases, then we assume that there is no phase transition. If the limits are different, then a phase transition exists.

Figure 10 indicates that the existence of a phase transition is related to the retry limit, as predicted by our analysis. For the case $R = 4$, there is no phase transition, while a phase transition occurs in the cases $R = 7$ and $R = 10$.

The analysis also reasonably approximates the phase transition region. For $R = 7$, the simulations show that a phase transition exists if $\rho \in (0.12, 0.16)$, while the analysis predicts $\rho \in (0.14, 0.17)$. For $R = 10$, the simulation results are $\rho \in (0.08, 0.14)$ while the analysis predicts $\rho \in (0.10, 0.16)$. We note that the size of the phase transition region increases with R, as predicted by the analysis.

VI. Conclusion

We describe a new type of DoS attacks against Wi-Fi networks, called cascading DoS attacks. The attack exploits a coupling vulnerability due to hidden nodes. The attack propagates beyond the starting location, lasts for long periods of time, and forces the network to operate at its lowest bit rate. The attack can be started remotely and without violating the IEEE 802.11 standard, making it difficult to trace back.

We demonstrate the feasibility of such attacks, both through experiments on a testbed and extensive ns-3 simulations. The simulations show that the attack is effective in networks operating under fixed and varying bit rates, as well as ad hoc and infrastructure modes. We show that a small change in the traffic load of the attacker can lead to a phase transition of the entire network, from uncongested state to congested state.

We develop an iterative analysis to characterize the sequence of node utilizations, and study its limiting behavior. We show that the sequence always converges to fixed points. Based on the system parameters, we identify when the system remains always uncongested, congested, or experiences a phase transition caused by a DoS cascading attack. The analysis predicts that a phase transition occurs for $R \geq 7$ and provides a simple and explicit estimate of traffic load under which a phase transition occurs (i.e., $\rho \in (1/R, 0.161)$). Although the analysis is based on some simplifying assumptions, the estimate is not far from the values observed in the simulations.

Exploiting the coupling vulnerability in different network configurations represents an interesting area for future work. Experience in the security field indeed teaches that once a vulnerability is identified, more potent attacks are subsequently discovered (consider, for instance, the history of attacks on WEP [31] and MD5 [32]).

Several approaches are possible to mitigate cascading DoS attacks. First, one could enable the RTS/CTS exchange, although this solution has several drawbacks, including major performance degradation under normal network operations, as mentioned in the Introduction. The second approach is to lower the retry limit. However, this could also negatively impact performance. Other approaches include using short packets, collision-aware rate adaptation algorithms, and full-duplex radios. We leave the investigation and comparison of these mitigation techniques as possible areas for future work.

Acknowledgment

This work was supported in part by NSF under grants CNS-1012910, CNS-1409053 and CNS-1409453.

References

[1] K. Lee, J. Lee, Y. Yi, I. Rhee, and S. Chong, “Mobile data offloading: how much can WiFi deliver?” in *Proceedings of the 6th International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*. ACM, 2010, p. 26.

[2] Cisco, “Cisco cleanair technology,” http://www.cisco.com/c/en/us/solutions/enterprise-networks/cleanair-technology/index.html.

[3] G. Bianchi, “Performance analysis of the IEEE 802.11 distributed coordination function,” *IEEE Journal on Selected Areas in Communication*, vol. 18, no. 3, pp. 535–547, 2000.

[4] A. Forouzan Behrouz, *Data Communication and Networking*. 3rd/4th Edition, Tata McGraw, 2004.

[5] M. Gast, *802.11 wireless networks: the definitive guide*. O’Reilly Media, Inc., 2005.

[6] http://documentation.netgear.com/WPN824EXT/enu/202-10310-02/WPN824EXT_UG-4-6.html.

[7] http://www.tp-link.com/resources/document/TD-W8901N_V1_User_Guide_191001.pdf.

[8] http://ui.linksys.com/WAG300N/1.01.01/help/h_AdvWSettings.htm.

[9] http://support.dlink.com/emulators/dir855/Advanced.html.

[10] L. Xin, D. Starobinski, and G. Noubir, “Cascading denial of service attacks on Wi-Fi networks,” http://arxiv.org/abs/1604.05255, 2016.

[11] R. Poisel, *Modern communications jamming principles and techniques*. Artech House Publishers, 2011.

[12] K. Pelechrinis, M. Iliofotou, and S. V. Krishnamurthy, “Denial of service attacks in wireless networks: The case of jammers,” *Communications Surveys & Tutorials, IEEE*, vol. 13, no. 2, pp. 245–257, 2011.

[13] G. Lin and G. Noubir, “On link layer denial of service in data wireless LANs,” *Wireless Communications and Mobile Computing*, vol. 5, no. 3, pp. 273–284, 2005.

[14] G. Noubir, R. Rajaraman, B. Sheng, and B. Thapa, “On the robustness of IEEE 802.11 rate adaptation algorithms against smart jamming,” in *Proceedings of the fourth ACM conference on Wireless network security*. ACM, 2011, pp. 97–108.

[15] C. Orakcal and D. Starobinski, “Jamming-resistant rate adaptation in Wi-Fi networks,” *Performance Evaluation*, vol. 75, pp. 50–68, 2014.

[16] C. Chen, H. Luo, E. Seo, N. H. Vaidya, and X. Wang, “Rate-adaptive framing for interfered wireless networks,” in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*. IEEE, 2007, pp. 1325–1333.

[17] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee, “Diagnosing wireless packet losses in 802.11: Separating collision from weak signal,” in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*. IEEE, 2008.

[18] “Minstrel madwifi documentation,” http://linuxwireless.org/en/developers/Documentation/mac80211/RateControl/minstrel.

[19] R. Kinney, P. Crucitti, R. Albert, and V. Latora, “Modeling cascading failures in the north american power grid,” *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 46, no. 1, pp. 101–107, 2005.

[20] S. Soltan, D. Mazauric, and G. Zussman, “Cascading failures in power grids: analysis and algorithms,” in *Proceedings of the 5th international conference on Future energy systems*. ACM, 2014, pp. 195–206.

[21] M. Haenggi, J. G. Andrews, F. Baccelli, O. Dousse, and M. Franceschetti, “Stochastic geometry and random graphs for the analysis and design of wireless networks,” *Selected Areas in Communications, IEEE Journal on*, vol. 27, no. 7, pp. 1029–1046, 2009.

[22] Z. Kong and E. M. Yeh, “Wireless network resilience to degree-dependent and cascading node failures,” in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2009. WiOPT 2009. 7th International Symposium on*. IEEE, 2009, pp. 1–6.

[23] S. Ray, D. Starobinski, and J. B. Carruthers, “Performance of wireless networks with hidden nodes: a queuing-theoretic analysis,” *Computer Communications*, vol. 28, no. 10, pp. 1179–1192, 2005.

[24] V. Saligrama and D. Starobinski, “On the macroscopic effects of local interactions in multi-hop wireless networks,” in *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, 2006 4th International Symposium on*. IEEE, 2006, pp. 1–8.

[25] J. C. Stein, “Indoor radio WLAN performance part ii: Range performance in a dense office environment,” *Intersil Corporation*, vol. 2401, 1998.

[26] “iperf 2 user documentation,” http://iperf.fr/iperf-doc.php.

[27] D. Xia, J. Hart, and Q. Fu, “Evaluation of the minstrel rate adaptation algorithm in ieee 802.11 g wlans,” in *2013 IEEE International Conference on Communications (ICC)*. IEEE, 2013, pp. 2223–2228.

[28] B. Rong and A. Ephremides, “Protocol-level cooperation in wireless networks: Stable throughput and delay analysis,” in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2009. WiOPT 2009. 7th International Symposium on*, June 2009.

[29] L. Kleinrock, F. Tobagi *et al.*, “Packet switching in radio channels: Part i–carrier sense multiple-access modes and their throughput-delay characteristics,” *Communications, IEEE Transactions on*, vol. 23, no. 12, pp. 1400–1416, 1975.

[30] S. Lynch, *Dynamical systems with applications using MATLAB*. Springer, 2004.

[31] E. Tews, R.-P. Weinmann, and A. Pyshkin, “Breaking 104 bit WEP in less than 60 seconds,” in *Information Security Applications*. Springer, 2007, pp. 188–202.

[32] J. Black, M. Cochran, and T. Highland, “A study of the MD5 attacks: Insights and improvements,” in *Fast Software Encryption*. Springer, 2006, pp. 262–277.