Cascading Attacks on Wi-Fi Networks: Theory and Experiments

Liangxiao Xin, *Member, IEEE,* David Starobinski, *Senior Member, IEEE,* and Guevara Noubir, *Senior Member, IEEE*

***Abstract*—We unveil the existence of a vulnerability in Wi-Fi (802.11) networks, 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, theoretical analysis, and ns-3 simulations. The experiment shows that an attacker can cause the throughput of a node outside its communication range to vanish. To gain insight into the root-causes of the attack, we model the network as a dynamical system and analyze its limiting behavior and stability. The model predicts that a phase transition (and hence a cascading attack) is possible in linear networks when the retry limit parameter of Wi-Fi is greater or equal to 7, and also 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 governments, 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 CSMA cannot prevent collisions caused by hidden nodes (cf. Section III for more details about the hidden node problem). As a consequence, different networks do not always evolve independently, even if they are located far away.

To understand the consequence of such interactions, suppose that some node A_0 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 a node, say node B_1, concurrently receiving packets from another node, say node A_1. Node A_1 may be unable to hear transmissions by node A_0 due to the hidden node problem. As a result, node A_1 keeps retransmitting packets which collide with the packets sent by node A_0. These retransmissions by node A_1 may in turn affect the ability of other nodes in the network to successfully communicate, thus causing this phenomenon to propagate. We note that the total number of packet retransmissions (including the original transmission) cannot exceed the so-called *retry limit*, after which a packet must be dropped. We will show in the sequel that the retry limit plays a major role in sustaining the propagation effect.

An optional mechanism, called request-to-send and clear-to-send (or RTS/CTS), has been designed to combat the hidden node problem. However, this mechanism 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 as explicitly stated in [6]–[9]. Therefore, most Wi-Fi systems today operate without RTS/CTS.

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 introduce a new dynamic system model that sheds light into the network behavior under attack. The model shows 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.

The theoretical model shows that the sequence of node utilizations always converges to a fixed point (the utilization of a node is defined as the fraction of time during which the

L. Xin, and D. Starobinski are with the Division of Systems Engineering, Boston University, Boston, MA 02215 USA (e-mail: xlx@bu.edu; staro@bu.edu).

G. Noubir is with the Khoury College of Computer and Information Science, Northeastern University, Boston, MA 02115 USA (e-mail: noubir@ccs.neu.edu).

node transmits). We characterize the different types of fixed points (stable and unstable) and show that a phase transition is associated with the existence of an unstable fixed point. The model explicitly predicts for which values of the retry limit a phase transition (and hence a cascading attack) can occur. In particular, we show that a phase transition can occur for the default value of the retry limit in Wi-Fi, which is 7.

Finally, we concretely demonstrate the attack through experiments on a testbed composed of nodes equipped with real Wi-Fi cards and provide simulation results obtained with the ns-3 simulator that corroborate the theoretical results in various network topologies.

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 the network model and attack scenario. Section IV presents our theoretical analysis. We present experimental and simulation results that verify the findings in Section V. Section VI concludes the paper.

An earlier version of this paper appeared in the proceedings of the IEEE Conference on Communications and Network Security (CNS 2016) [10]. This journal version significantly expands the theoretical analysis, including detailed proofs and new results on stability analysis and heterogeneous traffic load, which can be found in Section IV. Moreover, new simulation results for networks based on a realistic indoor building model and ring networks are presented in Sections V.

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. 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 [14], [15]. However, to the best of our knowledge, none of these rate adaptation algorithms are used in practice. Our work is based on Minstrel [16], 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 [17], [18]. 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, [19], [20] 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 generation of traffic 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 [21] shows that interference coupling can affect the stability of multi-hop networks. In the case of a greedy source, a three-hop network is stable while a four-hop network becomes unstable. In contrast, in our work, the path of each packet consists of a single-hop. Thus, network instability is not due to multi-hop communication in our case.

The work in [22] shows that local coupling due to interferences can have global effects on wireless networks. Thus, it 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.

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. 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

A. IEEE 802.11 Back-off Mechanism

The IEEE 802.11 standard uses the CSMA/CA mechanism to control access to the transmission medium and avoid collisions. After a packet is sent, a node waits for a short interframe slots (SIFS) period to receive an acknolwedgment (ACK). Whenever the channel becomes idle, the node waits for a distributed interframe space (DIFS $>$ SIFS) period and a random backoff before contending for the channel. The random backoff consists of a random number of backoff slots, which depends on the so-called contention window. Specifically, at the r 1 retransmission attempt (retry count), the contention window CW_r is given by

$$CW_r = \begin{cases} 2^{r\ 1}(CW_1 + 1) \quad 1 & CW_r < CW_{max}, \\ CW_{max} & \text{otherwise.} \end{cases} \tag{1}$$

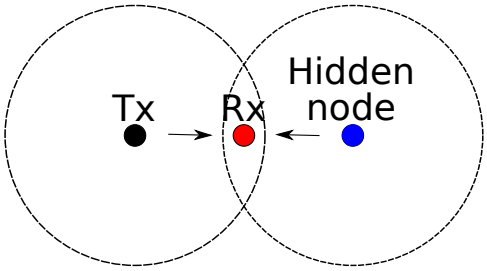

Fig. 1. Classical hidden node problem. The transmitter and the hidden node cannot sense each other. The collision happens when they transmit simultaneously.

The number of backoff slots is chosen uniformly at random in the interval $[0, CW_r]$. For IEEE 802.11b, the initial contention window size is $CW_1 = 31$, the maximum contention window size is $CW_{max} = 1023$, and the duration of a backoff slot is $20\ \mu s$. Note that the case $r = 1$ corresponds to the initial packet transmission attempt.

B. The Hidden Node Problem

A typical instance of the hidden node problem is illustrated in Figure 1. 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 r of a packet exceeds the retry limit, the packet is dropped, the retry count is reset to $r = 1$, and a new packet transmission can start. 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 2. This configuration could arise over different time and space in more complex network topologies. We consider $N + 1$ 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.

In general, the linear topology considered here represents a propagation path used by an attack. It is possible for an attack to be launched in a more general network as long as such a propagation path exists. We show a concrete example in our simulations in Section V-C.

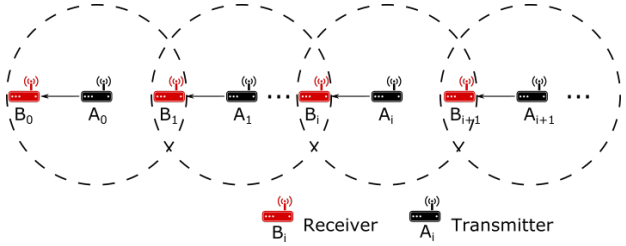

Fig. 2. Topology of the network. Node A_i transmits packets to node B_i. Node A_i is a hidden node with respect to A_{i+1}.

We assume that all the nodes communicate over the same channel. Note that there are only three non-overlapping channels in the 2.4GHz band. Hence, it is common that several nodes use the same channel over time and space in crowded areas. For instance, in a dense Wi-Fi network, each cell has multiple neighboring cells. Since there are only three non-overlapping channels, some neighboring cells will necessarily share the same channel (i.e., there could be other pairs of nodes using different channels which are not shown in Figure 2).

D. Attack Scenario

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.

E. Impact of exponential back-off

When a hidden node retransmits its packets, it must back off after each retransmission, which leaves the channel idle for a certain period of time. The duration of the backoff period is generally too short to allow for a successful transmission. Indeed, a packet transmission is successful only if

1) The size of the contention window of the hidden node is longer than the packet transmission time.
2) The transmitter starts and ends its transmission entirely during the backoff period of the hidden node.

At 1 Mb/s, the transmission time of an 1500 bytes packet lasts 12 ms. This is longer than the contention window as long as $CW_r < CW_{max} = 1023$. Hence, by Eq. (1), a transmission cannot be successful during the backoff period preceding the $r < 6$ retransmission attempt by a hidden node. Note that in the attack scenarios considered in this paper, each transmitter is a hidden node (i.e., it does not hear the transmissions of other nodes). Hence, the backoff counter keeps counting down and never freezes.

At the r 6 retransmission attempt by a hidden node A_i, $CW_r = CW_{max} = 1023$. Node A_i back-offs for n slots, where n is an integer between 0 and 1023 that is picked uniformly at random (i.e., with probability $1/1024$). Since the length of a backoff slot is 20 μs, the backoff delay is $0.02n$ ms. Without loss of generality, assume that node A_i starts backing off at time $t = 0$ and ends its backoff at time $t = 0.02n$ (all the time units are in milliseconds). Node A_i then starts a packet transmission, which ends at time $t = 0.02n + 0.12$.

Node A_{i+1} can transmit a packet successfully only if it starts its transmission during the time interval $[0, 0.02n \quad 12]$. This requires $n > 600$. Assuming that the starting time of the packet transmission by node A_{i+1} is uniformly distributed in the time interval $[0, 0.02n + 12]$, the probability that a packet is successfully transmitted by node A_{i+1} is

$$\sum_{n=600}^{1023} \frac{1}{1024} \quad \frac{0.02n \quad 12}{0.02n + 12} = 0.059.$$

Thus, the likelihood of a successful packet transmission is low, a result validated by the experimental and simulation results of Section V.

IV. Analysis

In this section, we develop an analytical model that provides insight into the network behavior under attack. Specifically, our goals are to explain why and under what conditions an attacker can congest a remote node and cause its throughput to vanish, and to 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 2. Packet generations 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 $1 \quad \overline{r}_i \quad R$ represent the mean retry count at node A_i. Note that the initial packet transmission is included in that count. 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 spacing are ignored. Therefore the upper limit of the utilization equals 1 in our analysis.

We denote the utilization of node A_i by $0 \quad u_i \quad 1$, 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 $i \quad 0$ is

$$u_i = \min\{\overline{r}_i \lambda_i T, 1\}. \tag{2}$$

Note that there is no retransmission at node A_0 and $\overline{r}_0 = 1$.

Our model represents a special case of interacting queues, which are notoriously difficult to analyze [23]. 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.

Our model is similar to the "random-look" model used by Kleinrock and Tobagi in their analysis of (single hop) random access networks [24] (see also Ch. 4 of [25]). We stress that beside these assumptions, the rest of our analysis is exact. Note that the experiments and simulations shown in Section IV do not incorporate the simplifications used to make the analysis tractable, yet they produce the same effects.

B. Iterative analysis of the utilization

Our goal is to find the utilization at each node $i \quad 0$ and in the limit as $i \; ! \; 1$. We consider the same scenario as in our simulations, whereby node A_0 (the attacker) varies its traffic load

$$\rho_0 \; , \; \lambda_0 T, \tag{3}$$

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

$$\rho \; , \; \lambda_i T, \tag{4}$$

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 function of the parameters ρ and R.

First, we get the utilization at node A_0:

$$u_0 = \min\{\rho_0, 1\}. \tag{5}$$

We next develop an iterative procedure to derive u_{i+1} from u_i. From (2) and (4),

$$u_{i+1} = \min\{\overline{r}_{i+1}\rho, 1\}. \tag{6}$$

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 $1 \quad r \quad R$ attempts is $(1 \quad p_{i+1})(p_{i+1})^{r \quad 1}$ 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

$$\begin{aligned} \overline{r}_{i+1} &= \sum_{r=1}^{R} r \quad (1 \quad p_{i+1}) \quad (p_{i+1})^{r \quad 1} + R \quad (p_{i+1})^R \\ &= \sum_{r=1}^{R} (p_{i+1})^{r \quad 1}. \end{aligned} \tag{7}$$

We next relate p_{i+1} to u_i. First, suppose $u_i < 1$ (i.e., node A_i is uncongested). Assume that node A_{i+1} starts a packet transmission (or retransmission) at some arbitrary time $t = t^0$. 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 $t = t^0$ 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 with probability 1. If node A_i does not transmit at time t^0, which occurs with probability $1 \quad u_i$, then a collision

occurs only if node A_i starts a transmission during the interval $[t^0, t^0 + T]$. Since the packet inter-arrival time on the channel is exponentially distributed with mean $\overline{r}_i T$, such an event occurs with probability

$$(1 \quad e^{\overline{r}_i \ i T}) = (1 \quad e^{\ u_i}), \tag{8}$$

based on Assumption 1. Therefore, the unconditional probability that a packet transmitted by node A_{i+1} collides is

$$\begin{aligned} p_{i+1} &= 1 \ u_i + (1 \quad e^{\ u_i}) \ (1 \quad u_i) \\ &= 1 \quad e^{\ u_i}(1 \quad u_i). \end{aligned} \tag{9}$$

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 (9) still provides the correct result.

Putting (6), (7), and (9) together, we obtain

$$u_{i+1} = \min\left\{\rho \sum_{r=1}^{R} \left(1 \quad e^{\ u_i}(1 \quad u_i)\right)^{r \ 1}, 1\right\}. \tag{10}$$

C. Limiting behaviour of the utilization

We next analyze the limiting behaviour of the iteration given by (10). The sequence $(u_i)_{i=0}^{1}$ corresponds to a discrete non-linear dynamical system [26]. Such systems are generally complex as they may converge to a point, to a cycle (i.e., they exhibit periodic behaviour), or not converge at all (i.e., they exhibit chaotic behaviour).

The main result of this section is to show that the sequence $(u_i)_{i=0}^{1}$ always converges to a point. However, the limit depends on the initial utilization u_0.

To simplify notation, we define the function

$$f(u_i) \ , \ \rho \sum_{r=1}^{R} \left(1 \quad e^{\ u_i}(1 \quad u_i)\right)^{r \ 1}. \tag{11}$$

We then rewrite (10) as follows:

$$u_{i+1} = \min f f(u_i), 1g. \tag{12}$$

We say that $\omega \ 2 \ [0, 1]$ is a *fixed point* of (12) if

$$\omega = \min f f(\omega), 1g. \tag{13}$$

Suppose (13) has K different fixed points (Theorem 2 in the sequel will show that $K \quad 1$). We denote by the ordered set of all the fixed points of (13). That is,

$$, \ f\omega_1, \ldots, \omega_k, \ldots, \omega_K g, \tag{14}$$

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

We are next going to show that for any $u_0 \ 2 \ [0, 1]$, the limit of the sequence $(u_i)_{i=0}^{1}$ is one of the elements in . To prove this result, we will use the following lemma.

Lemma 1: Let $u, u^0 \ 2 \ (\omega_k, \omega_{k+1})$, where $k \ 2 \ f1, \ldots, K \quad 1g$. If $f(u) > u$, then $f(u^0) > u^0$. If $f(u) < u$, then $f(u^0) < u^0$.

Proof: The proof goes by contradiction. Let $u, u^0 \ 2 \ (\omega_k, \omega_{k+1})$. Suppose $f(u) > u$ and $f(u^0) < u^0$. Since f is continuous in (ω_k, ω_{k+1}), then by the intermediate-value theorem there exists a point u^{00} between u and u^0 such that $f(u^{00}) = u^{00}$. Thus, u^{00} is a fixed point of (13). This contradicts the fact that no fixed point exists between ω_k and ω_{k+1}. ■

We now present the main result of this section.

Theorem 1:

1) Let $u_0 \ 2 \ (\omega_k, \omega_{k+1})$, where $k \ 2 \ f1, \ldots, K \quad 1g$. If $f(u_0) > u_0$, the sequence $(u_i)_{i=0}^{1}$ converges to ω_{k+1}. If $f(u_0) < u_0$, the sequence $(u_i)_{i=0}^{1}$ converges to ω_k.
2) If $u_0 \ 2 \ [0, \omega_1)$, the sequence $(u_i)_{i=0}^{1}$ converges to ω_1.
3) If $\omega_K < 1$ and $u_0 \ 2 \ (\omega_K, 1]$, the sequence $(u_i)_{i=0}^{1}$ converges to ω_K.

Proof:

1) Let $\omega_k < u_0 < \omega_{k+1}$, where $k \ 2 \ f1, \ldots, K \quad 1g$. Since $p_i \ 2 \ (0, 1)$. Therefore, the function f is continuous and monotonically increasing, $f(\omega_k) < f(u_0) < f(\omega_{k+1})$. Hence, according to (12) and (13), we get

$$\omega_k \quad u_1 \quad \omega_{k+1}. \tag{15}$$

Now, suppose $u_1 = f(u_0) > u_0$. If $u_1 = \omega_{k+1}$, then the result is proven. If $u_1 < \omega_{k+1}$, then by Lemma 1 and Equation (15), we have $u_2 = f(u_1) > u_1$. Applying the same argument inductively, either there exists some value $M \quad 2$ such that $u_i = \omega_{k+1}$ for all $i \quad M$, or the sequence $(u_i)_{i=0}^{1}$ is monotonically increasing and upper bounded by ω_{k+1}. According to the monotone convergence theorem, the sequence converges. Since there is no other fixed point between u_0 and ω_{k+1} and f is continuous, the sequence $(u_i)_{i=0}^{1}$ must converge to ω_{k+1}. The case $u_1 = f(u_0) < u_0$ is handled similarly.
2) Similar to Lemma 1, one can show that if there exists $u \ 2 \ [0, \omega_1)$ such that $f(u) > u$, then $f(u^0) > u^0$ for all $u^0 \ 2 \ [0, \omega_1)$. Since $f(0) = \rho > 0$, the sequence $(u_i)_{i=0}^{1}$ converges to ω_1.
3) This is handled similarly to case 2.

■

In summary, the existence of fixed points is determined by the utilization of all the nodes except the attacking node. The fixed points can be computed by solving (13). Once the fixed points are known, Theorem 1 provides the ranges of utilization of the attacking node u_0 for which the sequence converges to each fixed point.

D. Phase transition analysis

In the previous section, we showed that the limit of the sequence of node utilizations $(u_i)_{i=0}^{1}$ must be one of the fixed points in the set . 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 $(u_i)_{i=0}^{1}$ 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 $\omega \ 2$, 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 (13). First, we investigate the case $\omega = 1$.

Lemma 2: 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.

Proof:

1) Let $\rho \geq 1/R$. We compute the RHS of (13) at $\omega = 1$ and obtain $\min\{f(1), 1\} = \min\{R\rho, 1\} = 1$, which proves that a fixed point indeed exists at $\omega = 1$.
2) If $\rho > 1/R$, then $f(1) = R\rho > 1$. Since $f(1) > 1$, then for all $u_0 \in (0, \omega_K)$, we have $f(u_0) > u_0$, based on an argument similar to Lemma 1, and the sequence $(u_i)_{i=0}^{\infty}$ converges to 1, following an argument similar to Theorem 1.
3) This is handled similarly to Part 2.

■

Lemma 2 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, (13) can be rewritten

$$f(\omega) = \omega. \tag{16}$$

We look for solutions of (16) that belong to the interval $[0,1]$. Each such solution is an element of .

Equation (16) is difficult to work with because it contains two unknown variables, ρ and R. To circumvent this difficulty, we introduce the function

$$h_R(\omega) \triangleq \frac{\rho\omega}{f(\omega)} = \frac{\omega}{\sum_{r=1}^{R} (1 - e^{-\gamma(1-\omega)})^{r-1}}. \tag{17}$$

For each value of ρ, the solutions of (16) must satisfy

$$h_R(\omega) = \rho. \tag{18}$$

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, (18) 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, (18) 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 2. Thus, there are $K \geq 2$ fixed points.
   By Part 3 of Lemma 2, 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 (16) 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 2, the sequence $(u_i)_{i=0}^{\infty}$ converges to 1 for any $u_0 \in [0,1]$, and the network is always congested.

■

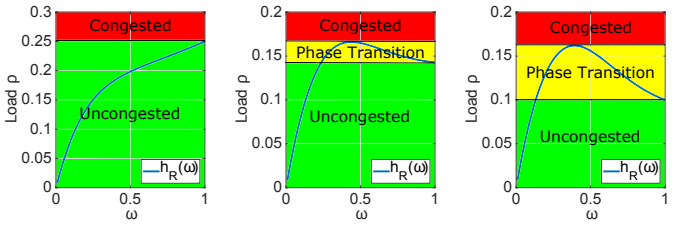

(a) $R = 4$ (b) $R = 7$ (c) $R = 10$

Fig. 3. Illustration of the different network regimes for different values of R. For each value of , the fixed points are the solutions of $h_R(\omega) =$. In addition, the fixed point $\omega = 1$ always exists when $> 1/R$. A phase transition region exists if the maximum of $h_R(\omega)$, h_R^{max}, is strictly greater than $h_R(1) = 1/R$.

Theorem 2 establishes whether the network is always uncongested, is susceptible to a phase transition, or is always congested, depending on the network parameters. We illustrate this theorem for different values of R and ρ, using Figure 3. First, consider $R = 4$ as shown in Figure 3(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 3(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 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 3(c). Furthermore, the size of this region is larger since $(1/R, h_R^{max}) = (0.1, 0.162)$.

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. The proof of the following lemma can be found in [27].

Lemma 3: For $\omega \in [0,1]$, if there exists R such that $h'_{R^*}(\omega) \leq 0$, then $h'_R(\omega) \leq 0$ for all $R > R^*$.

Next, consider the function $h_R(\omega)$ as $R \to 1$:

$$\begin{aligned} h_1(\omega) &= (1 - (1 - e^{-1}(1 - \omega)))\omega \\ &= e^{-1}(1 - \omega)\omega, \end{aligned} \tag{19}$$

and its derivative

$$h'_1(\omega) = e^{-1}(1 - 3\omega + \omega^2). \tag{20}$$

The next corollary is the logical transposition of Lemma 3.

Corollary 1: If $h'_1(\omega) \leq 0$, then $h'_R(\omega) \leq 0$ for all $R \geq 1$.

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

$$\overline{\omega} \triangleq \frac{3 - \sqrt{5}}{2}. \tag{21}$$

Lemma 4: Let $0 \leq \omega < \overline{\omega}$. Then, $h'_R(\omega) > 0$, for all $R \geq 1$.

Proof: Let the function $h_1(\omega)$ and its derivative $h'_1(\omega)$ be defined as in (19) and (20), respectively. Since e^{-1} is always positive, $h'_1(\omega)$ has the same sign as $(1 - 3\omega + \omega^2)$. The unique root of $(1 - 3\omega + \omega^2) = 0$ for $\omega \in [0,1]$ is w as defined in (21).

Thus, $(1 - 3\omega + \omega^2)$ is positive when $0 \leq \omega < \overline{\omega}$, and so is $h'_1(\omega)$. By Corollary 1, $h'_R(\omega) > 0$ for $0 \leq \omega < \overline{\omega}$ and for all $R \geq 1$. ∎

The consequence of Lemma 4 is that for all $R \geq 1$,

$$h_R^{max} \geq h_R(\overline{\omega}). \tag{22}$$

This equation provide a lower bound on h_R^{max} that can easily be computed. We then obtain the following sufficient condition for the existence of phase transition.

Lemma 5: Let $\overline{\omega}$ be defined as in (21) and suppose $h_R(\overline{\omega}) > 1/R$. Then, a phase transition is guaranteed to exist for any $\rho \in (1/R, h_R(\overline{\omega}))$.

Proof: From Theorem 2, we know that a phase transition exists if $1/R < \rho < h_R^{max}$. By (22) and the assumption that $h_R(\overline{\omega}) > 1/R$, the proof follows. ∎

The next theorem establishes an even more explicit lower bound on h_R^{max}.

Theorem 3: Let $h_1(\omega)$ and $\overline{\omega}$ be defined as in (19) and (21), respectively. Then, $h_R^{max} \geq h_1(\overline{\omega}) \approx 0.161$.

Proof: By (17),

$$\begin{aligned} h_R(\overline{\omega}) &= \frac{\omega}{\sum_{r=1}^{R}(1 - e^{-1}(1 - \omega))^{r-1}} \\ &> \frac{\omega}{\sum_{r=1}^{1}(1 - e^{-1}(1 - \omega))^{r-1}} = h_1(\overline{\omega}). \end{aligned} \tag{23}$$

Thus, by (22) and (23), $h_R^{max} > h_1(\overline{\omega}) \approx 0.161$. Note that this bound is asymptotically tight as $R \to 1$ since $h_1^{max} = h_1(\overline{\omega})$. ∎

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 $\rho \in [1/R, 0.161]$.

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 (17), since for any $\omega \in [0,1]$ the denominator increases as R gets larger).

F. Stability of fixed points

In this subsection, we use stability theory to shed further light into the limiting behaviour of the sequence $(u_i)_{i=0}^{\infty}$. Specifically, the sequence $(u_i)_{i=0}^{\infty}$ converges to *stable* fixed points of and diverges from *unstable* fixed points of . We will show that the stability of the fixed points of (16) are determined by the sign of $h'_R(\omega)$ at those points.

Informally, a fixed point ω is stable (or an *attractor*), if there exists a domain containing ω, such that if u_0 belongs to that domain, then $(u_i)_{i=0}^{\infty}$ converges to ω.

Definition 3 (Stability of a fixed point): Let $u_0 \in [0,1]$. A fixed point $\omega \in$ is *stable* if there exists $\epsilon > 0$ such that if $|u_0 - \omega| < \epsilon$, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω. It is *unstable* if for all $u_0 \neq \omega$ the sequence $(u_i)_{i=0}^{\infty}$ does not converge to ω.

Recall that according to Lemma 2, a special fixed point of (13) exists at $\omega = 1$, if $\rho > 1/R$. According to Definition 3, this fixed point is stable. Besides this special case, the rest of the fixed points satisfy Equation (16). To establish the stability of those fixed points, we will employ the following proposition.

Proposition 1 ([26]): Suppose that a continuously differentiable function f has a fixed point ω. Then, ω is stable if $|f'(\omega)| < 1$ and unstable if $|f'(\omega)| > 1$.

The next theorem provides a criterion to establish the stability of a fixed point $\omega \in$ with respect to the function $h_R(\omega)$.

Theorem 4: Consider a fixed point $\omega \in$, where $\omega < 1$. Then ω is stable if $h'_R(\omega) > 0$ and unstable if $h'_R(\omega) < 0$.

Proof: Let $\omega \in$. The derivative of $h_R(\omega)$ with respect to ω is

$$h'_R(\omega) = \frac{1}{\Phi(\omega)} - \frac{\omega}{(\Phi(\omega))^2}\Phi'(\omega) > 0, \tag{24}$$

where

$$\Phi(\omega) \triangleq \sum_{r=1}^{R}\left(1 - e^{-1}(1 - \omega)\right)^{r-1} = \frac{f(\omega)}{\rho}. \tag{25}$$

If one can show that (24) implies $|f'(\omega)| < 1$, then according to Proposition 1, the fixed point ω is stable. We multiply both sides of (24) by $(\Phi(\omega))^2$ and obtain

$$\Phi(\omega) - \omega\Phi'(\omega) > 0. \tag{26}$$

Using (25) and (16), we can rearrange (26) as follows:

$$\Phi'(\omega) < \frac{\Phi(\omega)}{\omega} = \frac{f(\omega)}{\rho\omega} = \frac{1}{\rho}. \tag{27}$$

From (25) and (27), we get

$$f'(\omega) = \rho\,\Phi'(\omega) < 1.$$

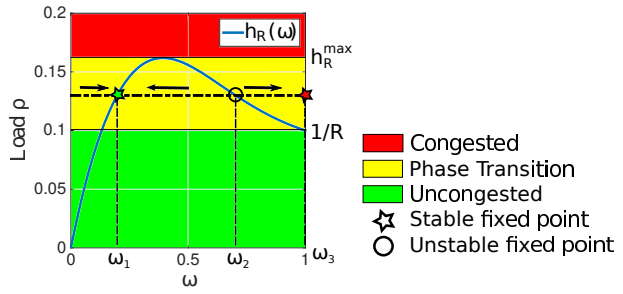

Fig. 4. Stability of fixed points with $R = 10$. Given a load = 0:13 (dash line), contains three fixed points: $!_1 = 0{:}2$, $!_2 = 0{:}7$ and $!_3 = 1$. The fixed point $!_1$ is stable because $h'_R(!_1) > 0$ and $!_2$ is unstable because $h'_R(!_2) < 0$. The fixed point $!_3 = 1$ exists and is stable because $> 1/R$. Therefore, the sequence $(u_i)_{i=0}^{\infty}$ converges to $!_1$ if $u_0 < !_2$, and to $!_3$ if $u_0 > !_2$.

Since $f(\omega)$ is monotonically increasing with ω, for $\omega \in [0,1]$, we conclude

$$0 < f'(\omega) < 1.$$

Hence, by Proposition 1, ω is a stable fixed point.

Similarly, $h_R^0(\omega) < 0$ implies $f'(\omega) > 1$, which means that ω is unstable. ■

Theorem 4 provides a stability analysis of the fixed points and helps determine the limit of the sequence $(u_i)_{i=0}^{1}$. Consider, for instance, the example shown in Figure 4 with parameters $R = 10$ and $\rho = 0.13$. Under these parameters, $= f\omega_1, \omega_2, \omega_3 g = f0.2, 0.7, 1g$.

The fixed points ω_1 and ω_2 are the solutions of $h_R(\omega) = \rho$. According to Theorem 4, ω_1 is stable and ω_2 is unstable. The fixed point $\omega_3 = 1$ exists and is stable, since $\rho > 1/R$.

According to Theorem 2, ω_2 is a phase transition point. Hence, the sequence $(u_i)_{i=0}^{1}$ converges to ω_1 if $u_0 < \omega_2$ and the network is uncongested. If $u_0 > \omega_2$, the sequence converges to ω_3 and the network is congested.

G. Heterogeneous traffic load

In previous subsections, we assumed that node A_0 varies its traffic load ρ_0, but all other nodes A_i (i 1) have the same traffic load ρ. We now relax this assumption and assume that nodes A_i (i 1) have different traffic loads $\rho_i = \lambda_i T$. We next prove that a phase transition still occurs, as long as all the traffic loads fall in the appropriate range.

Theorem 5: Suppose $h_R^{max} > 1/R$. If $\rho_i \in (1/R, h_R^{max})$ for all i 1, then a phase transition occurs.

Proof: Let $\rho_{max} = \max_{i \ 1} \rho_i$ and $\rho_{min} = \min_{i \ 1} \rho_i$. According to Theorem 2, the network is uncongested when $\rho_0 = 0$ and the load at each node A_i is $\rho_{max} < h_R^{max}$. Hence, the network must remain uncongested when the load at each node A_i is smaller than ρ_{max}.

Similarly, the network is congested when $\rho_0 = 1$ and the load at each node A_i is $\rho_{min} > 1/R$. Hence, it must remain congested when the load at each node A_i is larger than ρ_{min}. Thus, a phase transition occurs when $1/R < \rho_i < h_R^{max}$ for all i 1. ■

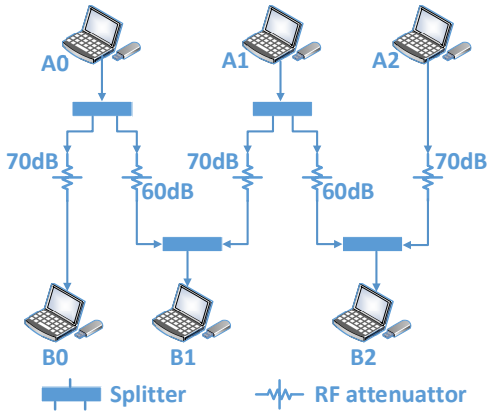

Fig. 5. Experimental testbed.

This result shows that phase transitions are also possible in linear networks with heterogeneous traffic loads.

V. Experiments and Simulations

A. Experiments

We demonstrate the practical feasibility of launching cascading DoS attacks through experiments on a testbed composed of six nodes. The testbed configuration is shown in Figure 5. 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-WN722N 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 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 ensures 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 [28].

The transmission power of each node is set to 0 dBm. We use iPerf [29] to generate 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 6 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.

Note that if the values of the attenuators are set equal, some packets transmitted at the lowest bit rate (i.e., 1 Mb/s) may be successfully received, even if the packets overlap. The analysis of this scenario is more complicated but simulations show that even in this case, cascading attacks are feasible [27, Ch. 4].

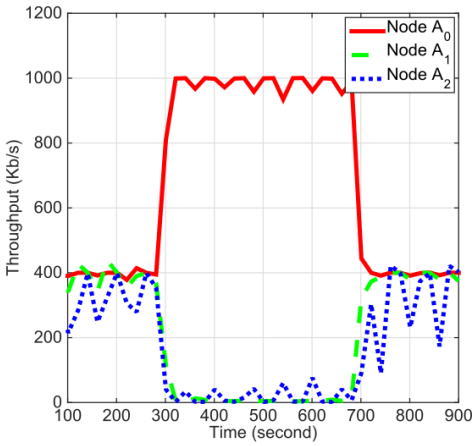

Fig. 6. Throughput performance measurements in testbed. When node A_0 starts increasing its packet generation rate, the throughput of nodes A_1 and A_2 vanishes.

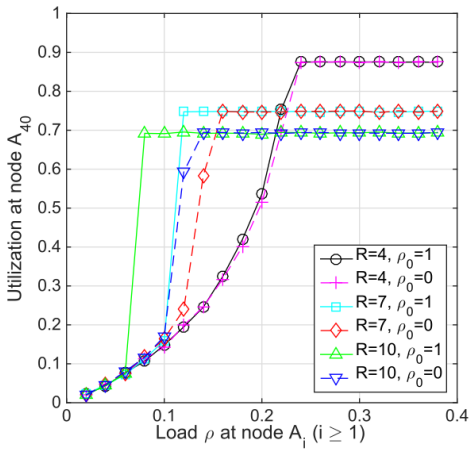

Fig. 7. Simulation of the limiting behaviour of the node utilization in a network of 41 pairs of nodes. For $R = 4$, the limit is the same when $\rho_0 = 0$ and $\rho_0 = 1$, hence no phase transition is observed. However, for $R = 7$ and $R = 10$, the limits are different, hence showing the existence of a region of load ρ in which a phase transition occurs.

B. Simulation results for linear topologies

We next compare the results of the analysis of Section IV 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 41 pairs of nodes.

1) Region of phase transition: 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 7 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 $R = 10$. In fact, we observed no phase transition in our simulations for any $R \geq 6$.

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.

2) Heterogeneous traffic load: We next show the feasibility of a cascading DoS attack in a network where the traffic load at different node is heterogeneous, in line with the analysis of Section IV-G. Specifically, the traffic load ρ_i at each node A_i ($i \geq 1$) is a continuous random variable that is uniformly distributed between 0.11 and 0.15.

Figure 8 shows the simulation results for retry limit $R = 7$. When ρ_0, the load of node A_0, is below 0.5, the network is uncongested and the utilizations of nodes A_i oscillate around 0.35 as i gets large. Note that the sequence does not converge to a fixed value due to the different traffic loads at the different nodes. However, when ρ_0 exceeds 0.6, the sequence of node utilizations converges to its upper limit, implying that the network is congested. We note that the convergence to steady-state is pretty fast, i.e., it is reached after about 10 nodes.

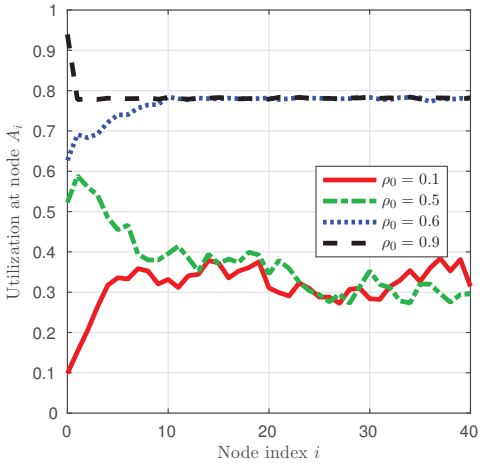

Fig. 8. Simulation with heterogeneous traffic load in a network with 41 pairs of nodes. The traffic load of nodes A_i ($i \geq 1$) are uniformly distributed between 0.11 and 0.15. For $R = 7$, when the load ρ_0 changes from 0.5 to 0.6, the limiting behavior of the sequence of node utilizations differs, thus indicating the occurrence of phase transition.

C. Simulation results for other topologies

We next next investigate cascading attacks in other topologies, specifically a realistic three-dimensional indoor topology and a ring topolgy.

1) 3D indoor building model: In this section, we use the ns-3 `HybridBuildingsPropagationLossModel` library [30] to demonstrate the feasibility of cascading DoS attacks in a 3D indoor scenario. Models in this library realistically characterize the propagation loss across different spectrum bands (i.e., ranging from 200 MHz to 2.6 GHz), different environments (i.e., urban, suburban, open areas), and different node positions with respect to buildings (i.e., indoor,

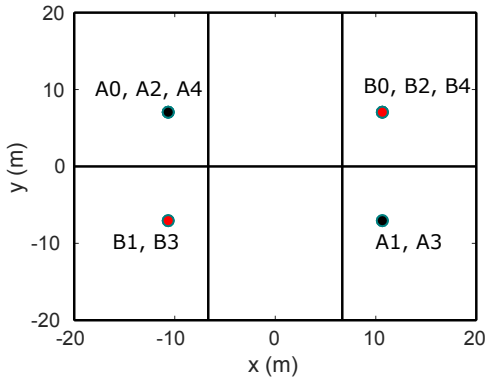

(a) Top view.

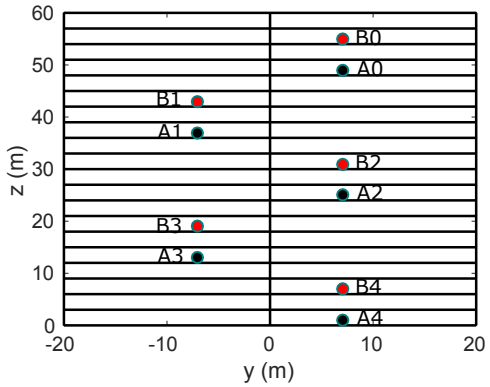

(b) Side view.

Fig. 9. Office building model. The building has 20 floors (z-axis) and 6 rooms in each floor (x and y axes).

outdoor and hybrid). The building models take into account the penetration losses of the walls and floors, based on the type of buildings (i.e., residential, office, and commercial).

In our simulations, we consider a 20-floor office building with six rooms in each floor, as shown in Figure 9. We assume that five pairs of Wi-Fi nodes (A_i, B_i) are active in the building, where node A_i transmits packets to nodes B_i ($i = 0, 1, 2, 3, 4$). The bit rate is set to 1 Mb/s, the retry limit to $R = 7$, and the frequency to 2.4 GHz. The generation rate of UDP packets at nodes A_i, i 1, is $\lambda_i = 8.125$ pkts/s. Packets are 2000 bytes long.

We turn on and off transmissions at node A_0 to observe how it impacts the throughput of other nodes. Simulation results are shown in Figure 10. When node A_0 does not transmit, the throughput of node A_4 is 0.13 Mb/s and it incurs no packet loss. However, when node A_0 starts transmitting, the throughput of node A_4 collapses. The throughput of node A_4 recovers only after node A_0 stops transmitting.

2) Ring topology: We next investigate cascading DoS attacks in a ring topology with 41 pairs of nodes, as shown in Figure 11. In our previous results for linear topologies, the effect of an attack disappears once the attacker reduces its packet generation rate. However, the effect of an attack in a ring topology can last for a long period of time after the attack stops.

This result is illustrated in Figure 12. During the first 100 seconds, all the nodes A_i ($i = 0, 1, \ldots$) generate packets at

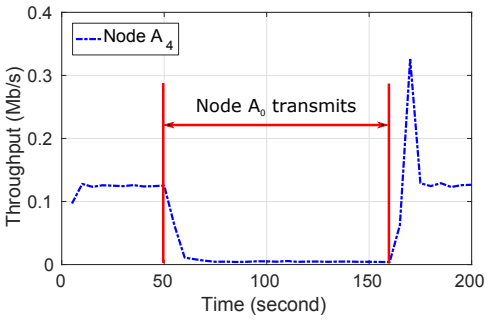

Fig. 10. Simulation results using ns-3 building model. When node A_0 transmits, the throughput of remote node A_4 collapses.

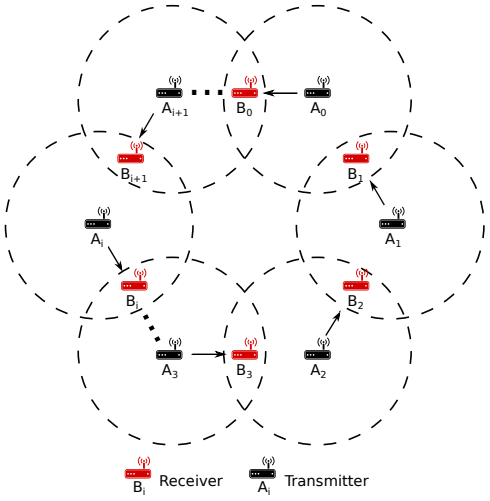

Fig. 11. Ring topology under cascading DoS attack. The dash circle represents the transmission range of the transmitter.

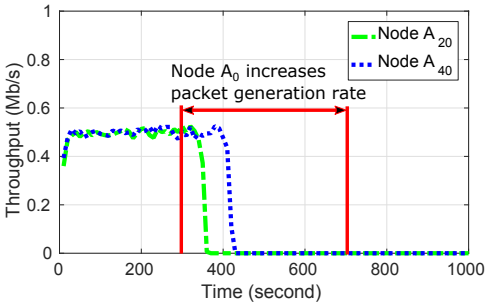

Fig. 12. Simulation results under a ring topology. When the packet generation rate of node A_0 increases, the throughput of nodes A_{20} and A_{40} vanishes. This effect continues even when the packet generation rate of node A_0 decreases.

0.5 Mb/s. At time $t = 300$ s, node A_0 increases its packet generation rate to 11 Mb/s. As a result, the throughput of all nodes vanishes. Yet, unlike results in linear topologies, the throughput of the nodes does not recover after node A_0 reduces its packet generation rate back to 0.5 Mb/s. The cyclic nature of the topology reinforces the attack even after the trigger stops.

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 through experiments on a testbed of nodes equipped with off-the-shelf Wi-Fi cards. The experiments show that change in the traffic load of the attacker can lead to a phase transition of a remote node, from uncongested to congested states.

To provide insight into this phenomenon, we propose a new dynamical system model to characterize the sequence of node utilizations, and analyze the limiting behavior of this sequence. We show that the sequence always converges to stable fixed points while an unstable fixed point represents a phase transition point. 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 7 in a linear network topology and provides a simple and explicit estimate of traffic load at each node under which a phase transition occurs (i.e., ρ_i 2 $(1/R, 0.161)$ for all i 1). The network is always congested when the traffic load is above the phase transition regime and always uncongested when the traffic load is below the phase transition regime. We also generalize our results to heterogeneous traffic load scenarios.

The theoretical results are corroborated with simulations and experiments. In terms of accuracy, our model is accurate in predicting that the throughput vanishes during cascading attacks (as shown by the real network experiments) as well as predicting the values of the retry parameter R for which cascading attacks are feasible. Notably, cascading attacks are feasible for the default value R = 7 used in IEEE 802.11. The analysis is also accurate in predicting the size of the phase transition region which increases with R. However, the analysis is less accurate in pinpointing the exact boundaries of the phase transition region (which is about 20% off). We defer the refinement of this particular aspect of the analysis to future work, as it would likely require a more complicated model.

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]). In particular, it is possible that interactions between different wireless protocols that use the same spectrum (e.g., Wi-Fi, Bluetooth, and Zigbee [33]) could create a similar security issue.

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. Devising a scheme that triggers RTS/CTS under certain circumstances (e.g., multiple consecutive packet losses) could be an interesting area for future research. The second approach is to lower the retry limit. However, this could also negatively impact performance. Other approaches worth investigating include using shortening packet duration [27, Ch. 5], dynamic channel selection [34], and full-duplex radios [35].

ACKNOWLEDGMENTS

This research was supported in part by the U.S. National Science Foundations under grants CNS-1409453, CNS-1908087, and DGE-1661532.

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.us/support/download-center.

[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," in *Communications and Network Security (CNS), 2016 IEEE Conference*.

[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] 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*.

[15] 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*.

[16] "Minstrel madwifi documentation," http://linuxwireless.org/en/developers/Documentation/mac80211/RateControl/minstrel.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] A. Aziz, D. Starobinski, and P. Thiran, "Understanding and tackling the root causes of instability in wireless mesh networks," *IEEE/ACM Transactions on Networking*, vol. 19, no. 4, pp. 1178–1193, 2011.

[22] 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.

[23] 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*. IEEE, 2009.
[24] 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.
[25] D. Bertsekas and R. Gallager, "Data networks. 1992," *PrenticeHall, Englewood Cliffs, NJ*, 1992.
[26] S. Lynch, *Dynamical systems with applications using MATLAB*. Springer, 2004.
[27] L. Xin, "Cascading attacks in Wi-Fi networks: demonstration and counter-measures," Ph.D. dissertation, Boston University, 2018.
[28] J. C. Stein, "Indoor radio WLAN performance part II: Range performance in a dense office environment," *Intersil Corporation*, vol. 2401, 1998.
[29] "iperf 2 user documentation," http://iperf.fr/iperf-doc.php.
[30] https://www.nsnam.org/doxygen/classns3_1_1_hybrid_buildings_propagation_loss_model.html#details.
[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.
[33] W. Wang, S. He, L. Sun, T. Jiang, and Q. Zhang, "Cross-technology communications for heterogeneous iot devices through artificial doppler shifts," *IEEE Transactions on Wireless Communications*, vol. 18, no. 2, pp. 796–806, 2018.
[34] D. J. Leith and P. Clifford, "A self-managed distributed channel selection algorithm for wlans," in *2006 4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*. IEEE, 2006, pp. 1–9.
[35] D. Bharadia, E. McMilin, and S. Katti, "Full duplex radios," in *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4. ACM, 2013, pp. 375–386.

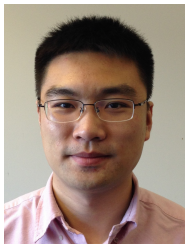

Liangxiao Xin received his B.E. degree in Control Science and Engineering (2012) from Zhejiang University, Zhejiang, China. In 2014 and 2018, he received his M.E. degree and Ph.D degree in Systems Engineering from Boston University, respectively. In 2016, he received best paper award at IEEE CNS 2016 conference. Currently, he is working on IEEE 802.11 standardization.

David Starobinski is a Professor of Electrical and Computer Engineering, Systems Engineering, and Computer Science at Boston University. He received his Ph.D. in Electrical Engineering from the Technion - Israel Institute of Technology, in 1999. He was a visiting post-doctoral researcher in the EECS department at UC Berkeley (1999-2000), an invited Professor at EPFL (2007-2008), and a Faculty Fellow at the U.S. DoT Volpe National Transportation Systems Center (2014-2019). Dr. Starobinski received a CAREER award from the U.S. National Science Foundation (2002), an Early Career Principal Investigator (ECPI) award from the U.S. Department of Energy (2004), the 2010 BU ECE Faculty Teaching Award, and best paper awards at the WiOpt 2010 and IEEE CNS 2016 conferences. He is on the Editorial Board of the IEEE Open Journal of the Communications Society and was on the Editorial Boards of the IEEE Transactions on Information Forensics and Security and the IEEE/ACM Transactions on Networking. His research interests are in cybersecurity, wireless networking, and network economics.

Guevara Noubir is a professor in the College of Computer and Information Science at Northeastern University. His research focuses on privacy and security in networked systems. Professor Noubir received a PhD in computer science from Ecole Polytechnique Fédérale de Lausanne, and an MS degree (Engineering Diploma) from Ecole Nationale d'Informatique et Mathématiques Apliquées de Grenoble. Dr. Noubir received the US NSF CAREER Award in 2005, Google Faculty Research Award on Privacy in 2016, best paper awards at ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec) 2011, and 2018 (and runner-up best paper in 2013), and the IEEE Conference on Communications and Network Security best paper in 2017. He chaired the technical program committee of several security conferences including the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec) in 2015, IEEE Conference on Communications and Network Security 2015. He serve(d) on the editorial boards of ACM Transaction on Privacy and Security, IEEE Transactions on Mobile Computing, Elsevier Journal on Computer Networks, and IEEE Transaction on IEEE Transactions on Information Forensics and Security.