2018

Cascading attacks in Wi-Fi networks: demonstration and counter-measures

BOSTON UNIVERSITY

COLLEGE OF ENGINEERING

Dissertation

CASCADING ATTACKS IN WI-FI NETWORKS: DEMONSTRATION AND COUNTER-MEASURES

by

LIANGXIAO XIN

B.Eng., Zhejiang University, 2012
M.Eng., Boston University, 2014

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2018

Approved by

First Reader

David Starobinski, Ph.D.
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

Second Reader

Ari Trachtenberg, Ph.D.
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

Third Reader

Manuel Egele, Ph.D.
Assistant Professor of Electrical and Computer Engineering

Fourth Reader

Guevara Noubir, Ph.D.
Professor of the College of Computer and Information Sciences
Northeastern University

To Father, Mother, and Yue

Acknowledgments

First, I would like to give my truehearted veneration and gratitude to my advisor Prof. David Starobinski. Without his academic guidance and support, I am unable to achieve my research in the past 5 years. Hence, I really appreciate the opportunity to work with him and learn from him not only on handling experiments and analyzing data, but also on writing a coherent paper and keeping patience and perseverance on whatever I am conducting.

Secondly, I am deeply thankful to Prof. Guevara Noubir, who gave me tremendous supports during the first year when I was a truly novice in this field. He helped to arrange a lab with all sound devices and gave me any support, such as providing me with assistance from more experienced lab mates, that I may need in Northeastern University.

I also want to thank Prof. Ari Trachtenberg and Prof. Manuel Egele for being my committee members as their insightful comments on my thesis are helpful. I also deeply appreciate Prof. Wenchao Li who is willing to sacrifice his time to be my defense chair when I cannot find anyone else to chair the defense.

I intended to list the names of my lab colleagues and friends, but I then realized that there are so many names should be shown on this page. I appreciate their academic help and personal support, namely listen to my complaint, in these years.

I am fully indebted to my family. Thanks to my father Zhiping Xin and my mother Suyue Hu who have raised my up and always hold faith on me unconditionally. I also would like to express my thanks and love to my girlfriend Yue Guo for always cheering me up during my Ph.D. life. I dedicate this dissertation to them.

Liangxiao Xin

Ph.D.

Division of Systems Engineering

CASCADING ATTACKS IN WI-FI NETWORKS: DEMONSTRATION AND COUNTER-MEASURES

LIANGXIAO XIN

Boston University, College of Engineering, 2018

Major Professor: David Starobinski, Ph.D.
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

ABSTRACT

Wi-Fi (IEEE 802.11) is currently one of the primary media to access the Internet. Guaranteeing the availability of Wi-Fi networks is essential to numerous online activities, such as e-commerce, video streaming, and IoT services. Attacks on availability are generally referred to as Denial-of-Service (DoS) attacks. While there exists significant literature on DoS attacks against Wi-Fi networks, most of the existing attacks are localized in nature, i.e., the attacker must be in the vicinity of the victim. The purpose of this dissertation is to investigate the feasibility of mounting global DoS attacks on Wi-Fi networks and develop effective counter-measures.

First, the dissertation unveils the existence of a vulnerability at the MAC layer of 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. The dissertation demonstrates 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, the network is modeled as a dynamical system and its limiting behavior is analyzed. 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.

Next, the dissertation identifies a vulnerability at the physical layer of Wi-Fi that allows an adversary to launch cascading attacks with weak interferers. This vulnerability is induced by the state machine's logic used for processing incoming packets. In contrast to the previous attack, this attack is effective even when interference caused by hidden nodes do not corrupt every packet transmission. The attack forces Wi-Fi rate adaptation algorithms to operate at a low bit rate and significantly degrades network performance, such as communication reliability and throughput.

Finally, the dissertation proposes, analyzes, and simulates a method to prevent such attacks from occurring. The key idea is to optimize the duration of packet transmissions. To achieve this goal, it is essential to properly model the impact of MAC overhead, and in particular MAC timing parameters. A new theoretical model is thus proposed, which relates the utilization of neighboring pairs of nodes using a sequence of iterative equations and uses fixed point techniques to study the limiting behavior of the sequence. The analysis shows how to optimally set the packet duration so that, on the one hand, cascading DoS attacks are avoided and, on the other hand, throughput is maximized. The analytical results are validated by extensive ns-3 simulations. A key insight obtained from the analysis and simulations is that IEEE 802.11 networks with relatively large MAC overhead are less susceptible to cascading DoS attacks than networks with smaller MAC overhead.

Contents

List of Tables

List of Figures

List of Abbreviations

ACK	.....	Acknowledgement
CSMA	.....	Carrier Sense Multiple Access
DCF	.....	Distributed Coordination Function
DIFS	.....	DCF Interframe Space
DoS	.....	Denial of Service
SIFS	.....	Short Interframe Space

Chapter 1

Introduction

Wi-Fi (IEEE 802.11) is a popular wireless technology that allows users to communicate over the unlicensed 2.4 GHz and 5 GHz channel bands. Wi-Fi is widely used for ubiquitous Internet access by a variety of organizations including schools, libraries, companies, towns and governments, as well as ISP hotspots and residential wireless routers. According to the statistics in Cisco Visual Networking Index 2017 (Cisco Systems, Inc., 2017), 42% of the Internet traffic was transmitted through Wi-Fi in 2015 and this ratio is expected to increase to 49% by 2020. The total number of Wi-Fi hotspots should grow six-fold from 2016 to 2021, from 94.0 million in 2016 to 541.6 million by 2021.

The widespread deployment of Wi-Fi networks increases the variety of services, which include cellular offload, video/voice streaming, Wi-Fi phone, online banking and so on. Indeed, Wi-Fi is the *de-facto* backbone communication technology for the Internet-of-Things (IoT). The state-of-the-art IoT products, such as Apple HomePod, Amazon Echo, and Google Home, rely on Wi-Fi for the communication. Therefore, it is critical to ensure the security of Wi-Fi networks. The importance of Wi-Fi networks and the need to strengthen their security have been recognized by companies, such as Cisco (Cisco Systems, Inc., 2016).

Security guarantees are typically divided into three categories: confidentiality, integrity, and availability. For the first two categories, IEEE 802.11 standard has ratified the 802.11i standard (i.e., WPA2 or WPA3) (IEEE 802.11 Working Group and

others, 2004) which has two versions: WPA-Personal for home networks and WPA-Enterprise for enterprise networks. WPA-Personal uses a shared secret key (password) between the access point and its clients for network authentication. WPA-Enterprise employs RADIUS servers to handle client authentication instead of pre-shared key, which effectively scales better to a large organization and makes key management more secure.

Wi-Fi is vulnerable to availability attacks due to the shared-medium nature of radio propagation (Zou et al., 2016). Violation of availability is generally referred to as a *denial of service* (DoS) attack. A DoS attack makes a Wi-Fi network inaccessible to its legitimate users, and may be launched by exploiting vulnerabilities at the physical layer, MAC layer, or even higher layers of Wi-Fi networks.

The surveys (Bicakci and Tavli, 2009; Zou et al., 2016) list a variety of DoS attacks on Wi-Fi networks and their counter-measures. For instance, attackers can exploit vulnerabilities at the physical layer to launch attacks, such as preamble attack, SFD attack (Gummadi et al., 2007), symbol attack (Lin and Noubir, 2005) and so on. Their counter-measures are rapid frequency hopping, spatial retreat, and multi-hop forwarding. Attacks at the MAC layer include deauthentication/disassociation attack, duration inflation attack, and attacks against sleeping nodes (Bellardo and Savage, 2003). Counter-measured against those attacks include cryptographic protection.

So far, most existing DoS attacks localized, which means that the attacker and the victim are within communication range. To our knowledge, only few works, such as (Haenggi et al., 2009; Kong and Yeh,), study DoS attacks from a global perspective. On the other hand, radio attacks, such as jamming, are non-compliant with Wi-Fi protocols. In this dissertation, we investigate DoS attacks which are protocol-compliant and have impact on a global (multi-hop) Wi-Fi network.

The main goal of this dissertation is to evaluate the feasibility of launching global

DoS attacks and propose counter-measures. The attacks proceed by exploiting coupling mechanisms that are intrinsic to wireless networks, in general, and the Wi-Fi protocol, in particular. Specifically, through experiments with real Wi-Fi cards, realistic ns-3 simulation, and analysis, the attacks unveiled in this dissertation exploit vulnerabilities at the MAC layer and the physical layer of the IEEE 802.11 standard.

We first unveil the existence of a vulnerability at the MAC layer of the IEEE 802.11 standard, which allows an attacker to launch protocol-compliant attacks that congest a global Wi-Fi network. Indeed, Wi-Fi networks rely on a simple, distributed mechanism, known as carrier sensing multiple access (CSMA), to arbitrate access to the shared medium and optimize performance. The behavior of this mechanism in isolated single-hop networks has been extensively studied and is generally well-understood (see, e.g., (Bianchi, 2000)). However, due to interference coupling, this mechanism results 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 1·1 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 neighbors and may propagate.

The coupling phenomenon induced by interference creates multi-hop dependencies, which an adversary can take advantage of the interference coupling to create

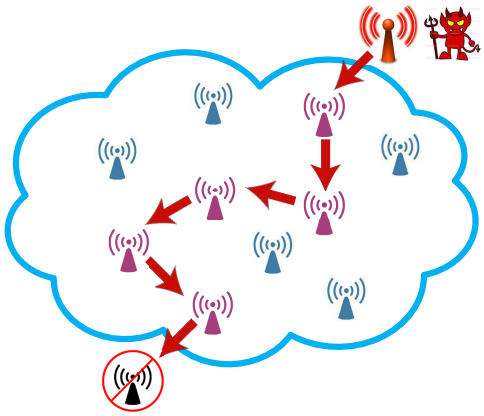

Figure 1·1: Illustration of a cascading denial of service attack. Transmissions by an attacker impact nodes located far away, due to interference coupling caused by hidden nodes.

multi-hop dependencies between different cells and launch a widespread network Denial-of-Service (DoS) 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 main contributions of the first part of this dissertation are as follows:

1. 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.

2. 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.

3. 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.

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

The attack described in the first part assumes that, at the receiver's end, the power of interference caused by a hidden node is stronger than the power of the signal of the sending station (in other words, the signal-to-interference ratio is -10dB) such that any overlap between transmissions of the station and the hidden node causes a loss of the packet transmitted by the station irrespective of the bit rate. However, such a disparity in the power strengths of the received signals may be rare in practice: in effect, it means that either the hidden node transmits at a much higher power than the station or that it is located much closer to the receiver. However, the attacker controls only the first hidden node in the chain, but not other hidden nodes in the network that help propagate the coupling effect.

In the second part of this dissertation, we investigate cases where interference caused by hidden nodes is on the same order or weaker than the signals of sending stations. Our main objective is to find out whether cascading attacks are still feasible in those situations. Through extensive ns-3 simulations and mathematical analysis, we provide an affirmative answer to this question. The attack leverages two phenomena, which we describe in detail as follows.

The first phenomenon is a PHY-layer effect known as *receiver capture* (Jiang and Liew, 2007). Accordingly, if the PHY header of the packet transmitted by the hidden

node is decoded first, the packet sent by the station is lost (assuming the two packet transmissions overlap). Thus, even though not all packets transmitted by the station are lost, a large fraction still is.

The second phenomenon relates to bit rate adaptation. Specifically, rate adaptation algorithms vary the bit rate used for packet transmissions based on the observed quality of the channel. While different algorithms have been proposed in the literature (Biaz and Wu, 2008), most gradually lower the bit rate upon experiencing packet losses. Since packet losses are still possible due the receiver capture effect, rate adaption algorithms may end up significantly lowering the bit rate of Wi-Fi stations, sometimes down to the base rate of 1 Mb/s. As a result, the capacity of the shared channel is drastically reduced (since each packet transmission uses the shared channel for a longer amount of time) leading to traffic congestion.

The main contributions of the second part of this dissertation can thus be summarized as follows:

1. We identify and document a coupling effect between neighboring cells, due to hidden nodes and receiver capture.

2. We analyze and provide simulations of the packet loss probability with and without receiver capture. We show that with receiver capture, a packet sent by a station is lost irrespective of the signal-to-interference ratio (SIR) and the bit rate.

3. Leveraging the above coupling effect, we demonstrate the feasibility of launching cascading attacks on Wi-Fi networks using weak hidden nodes (i.e., hidden nodes producing weak interference). Through extensive ns-3 simulations, including for an indoor building model, we show that the coupling effect may propagate, thus reducing the channel capacity across an entire chain of Wi-Fi cells. These results apply to several rate adaptation algorithms.

4. We provide an analysis of the limiting behavior of the channel utilization in a chain of Wi-Fi cells, assuming nodes implement the Auto Rate Fallback (ARF) rate adaptation algorithm (Kamerman and Monteban, 1997). In particular, we show how the average bit rate experiences a sharp drop as the channel utilization gets higher, which provides insight into how the attack is able to propagate throughout the network.

Given the serious consequences of cascading DoS attacks, it is important to find methods to mitigate them. 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 (Forouzan Behrouz, 2004; Gast, 2005). Indeed, most manufacturers of Wi-Fi cards disable RTS/CTS by default and discourage changing this setting as explicitly stated in (Netgear, 2016; TP-Link, 2016; Linksys, 2016; D-link, 2016). Therefore, most Wi-Fi systems today operate without RTS/CTS. Therefore, it is necessary to design mitigation other than RTS/CTS.

In the third part of this dissertation, we focus on the mitigation of cascading DoS attacks in Wi-Fi networks. Our key idea is to optimize the duration of packet transmissions (or, equivalently, the packet length divided by the bit rate) in order to ensure that interference coupling does not propagate and amplify. To achieve this goal, we show that it is essential to properly model the impact of MAC overhead, and in particular MAC timing parameters. We propose a refined theoretical model where we relate the utilization of nodes in neighboring cells using iterative equations. We then perform a fixed point analysis to characterize the limiting behavior of the sequence of node utilization and the feasibility of launching a cascading DoS network against a Wi-Fi network.

Our contributions to the mitigation of cascading DoS attacks are listed as follows:

1. We show how to set the packet duration in order to avoid a cascading DoS attack, namely to prevent the initial value of the sequence of node utilization (which can be set by the attacker) to affect the limit of the sequence.

2. We show that it is possible to simultaneously optimize the packet duration in order to achieve maximum throughput.

3. We validate the analytical results using ns-3 simulations, including for an office building model. A key insight obtained from our analysis and simulation is that IEEE 802.11 networks with relatively large MAC overhead (e.g., IEEE 802.11b) are less susceptible to cascading DoS attacks than networks with smaller overhead (e.g., IEEE 802.11g and IEEE 802.11n). We also show that our method achieves higher throughput performance than the RTS/CTS method, especially at high bit rates.

To summarize, in this dissertation,

- We unveil the existence of a vulnerability at the MAC layer of the IEEE 802.11 standard, which an adversary can take advantage of to launch a cascading DoS attack. The vulnerability stems from a coupling effect induced by hidden nodes. We first consider the scenario with strongly interfering hidden nodes, i.e., when any collision between transmissions of the station and the hidden node causes a loss of the packet transmitted by the station. 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 and stability.

- We investigate the feasibility of cascading DoS attacks with weakly interfering hidden nodes. Through extensive ns-3 simulations, including for an indoor building model, we show that cascading DoS attacks are still feasible. The attacks leverage two PHY-layer phenomena: receiver capture and bit rate adaptation. We provide supporting analysis for these phenomena and also investigate the limiting behavior of the channel utilization in a chain of Wi-Fi cells, assuming nodes implement the Auto Rate Fallback (ARF) rate adaptation algorithm. Our analysis sheds light into a coupling effect, whereby the average bit rate of a transmission pair drops sharply as the channel utilization of a neighboring pair gets higher. This coupling effect facilitates the propagation of the attack throughout the network.

- We provide a counter-measure against the occurrence of cascading DoS attacks. Our key idea is to optimize the duration of packet transmissions. To achieve this goal, we show that it is essential to properly model the impact of MAC overhead, and in particular MAC timing parameters. We propose a new theoretical model where we relate the utilization of neighboring pairs of nodes using a sequence of iterative equations and use fixed point techniques to study the limiting behavior of the sequence.

This dissertation has six chapters. The present chapter is Chapter 1. In Chapter 2, we present background and work related to this dissertation. In Chapter 3, we demonstrate the feasibility of cascading DoS attacks on Wi-Fi networks with strong interferers and investigate sufficient conditions to launch the attack. In Chapter 4, we investigate feasibility of the attacks in cases where interference caused by hidden nodes is on the same order or weaker than the signals of sending stations. In Chapter 5, we introduce, analyze, and experiment with a method for cascading DoS attacks in Wi-Fi networks. We conclude the dissertation in Chapter 6.

Chapter 2

Background and Related Work

In this chapter, we provide background and discuss work related to this dissertation. We first provide background on key aspects of IEEE 802.11 standard. In particular, we review the details of carrier sensing multiple access (CSMA), physical layer reception, cyclic redundancy check, and rate adaptation, whose vulnerabilities are exploited to launch the cascading attacks.

We then review the *hidden node* problem in Wi-Fi networks. Hidden nodes induce an interference coupling phenomenon that the cascading attacks leverage. We next describe the *receiver capture effect* which represents another vulnerability of Wi-Fi networks that can be exploited for the purpose of launching a cascading attack.

Next, we review prior work on DoS attacks in Wi-Fi networks. We review literature work on interference coupling, phase transitions, capture phenomena, and the impact of MAC timings, all of which play an important role in determining the feasibility of cascading DoS attacks.

Finally, we summarize the limitations of prior work and highlight on our new contributions.

2.1 IEEE 802.11 protocol

Wi-Fi is a wireless local area network (WLAN) technology, which mainly runs on 2.4 GHz ISM bands and 5 GHz bands (Gast, 2005). The IEEE 802.11 standard is a series of specifications, such as the media access control (MAC) and physical

layer (PHY) interfaces. The first 802.11 standard that gained widespread success is 802.11b. It runs on the 2.4 GHz band and supports 4 bit rates (1, 2, 5.5, 11 Mb/s). The subsequent standards (e.g., 802.11a, g, n, and ac) support higher 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.

2.1.1 Carrier-sense multiple access

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 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 $r \geq 1$ retransmission attempt (retry count), the contention window CW_r is given by

$$CW_r = \begin{cases} 2^{r-1}(CW_1 + 1) - 1 & CW_r < CW_{max}, \\ CW_{max} & \text{otherwise.} \end{cases} \tag{2.1}$$

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 μs. Note that the case $r = 1$ corresponds to the initial packet transmission attempt.

The number of backoff slots is an element of the set $\{0, 1, \ldots, CW_r\}$ chosen uniformly at random. We denote the duration of a backoff slot by T_{slot}. The average

backoff delay at the rth retransmission attempt is

$$\overline{T}_{\text{backoff},r} = \frac{1}{2} CW_r \cdot T_{\text{slot}}. \tag{2.2}$$

After sending a packet, a node waits for a short interframe space (SIFS) period before expecting to receive an ACK. If the ACK is received (i.e., the transmission is successful), then the average duration of the MAC overhead at the rth retransmission attempt is

$$d_r^{(s)} = T_{\text{DIFS}} + \overline{T}_{\text{backoff},r} + T_{\text{SIFS}} + T_{\text{ACK}}, \tag{2.3}$$

where T_{DIFS} and T_{SIFS} represent respectively the durations of the DIFS and SIFS intervals and T_{ACK} represents the duration of an ACK transmission.

If a node does not receive an ACK within an *ACK timeout* period (e.g., due to a collision caused by a hidden node), then it increments r and repeats the procedure. Thus, if a transmission fails, the average duration of the MAC overhead at the rth retransmission attempt is

$$d_r^{(f)} = T_{\text{DIFS}} + \overline{T}_{\text{backoff},r} + T_{\text{ACK\_timeout}}, \tag{2.4}$$

where $T_{\text{ACK\_timeout}}$ is the duration of the ACK timeout interval. This process continues as long as the number of retransmissions r does not exceed the (short) retry limit R. Once this limit is exceeded, the packet is dropped, r is reset to 1, and the transmission of a new packet can start. In all our analysis and simulations, we use the default value of the retry limit, namely $R = 7$ (Riley and Henderson, 2010).

The IEEE 802.11 standard has several variants, which differ in their physical and MAC layer specifications (Intel, 2017). These variants support transmissions at different bit rates going up to 11 Mb/s for IEEE 802.11b, 54 Mb/s for IEEE 802.11g, and 600 Mb/s (theoretically) for IEEE 802.11n. In practice, IEEE 802.11n networks often operate with bit rates going up to 54 Mb/s (Intel, 2017).

Table 2.1: IEEE 802.11 parameters (Gast, 2005)

	802.11b	802.11g/n
CW_1	31	15
$CW_{\max}$	1023	1023
T_{DIFS} (μs)	50	28
T_{SIFS} (μs)	10	10
T_{slot} (μs)	20	9 or 20

Table 2.1 shows settings of the timing parameters of IEEE 802.1b and IEEE 802.11g/n that are relevant to this dissertation. Note that IEEE 802.11g/n networks can use either a long slot time (i.e., $T_{\text{slot}} = 20$ μs) or a short slot time (i.e., $T_{\text{slot}} = 9$ μs) (ns-3, 2018). The long slot time is typically used in a mixed environment composed of both 802.11b and 802.11g/n nodes.

2.1.2 Physical layer reception

IEEE 802.11 uses the *Physical Layer Convergence Procedure (PLCP)* to implement physical layer functionalities (IEEE Standards Association and others, 2012). The format of a packet at the physical layer is shown in Figure 2·1, and consists of a PLCP preamble, a PLCP header and data (payload). The PLCP preamble consists of a synchronization pattern (SYNC) and a start frame delimiter (SFD) to indicate the start of a packet. The PLCP header, which follows the PLCP preamble, contains radio information about the packet, such as the bit rate (RATE) and the packet length (LENGTH). The last part of the packet is the data (payload) which corresponds to the MAC frame.

When a receiver processes packets at the physical layer, it transits between two PLCP states, the Carrier Sense/Clear Channel Assessment (CS/CCA) state and the Receive (Rx) state[1]. The transitions between those states are illustrated in Figure 2·2. In the CS/CCA state, the receiver monitors the state of the channel. Once it detects a

[1]Each state has its own internal state machine, see (IEEE Standards Association and others, 2012) for details.

Figure 2·1: Packet format at the physical layer of IEEE 802.11.

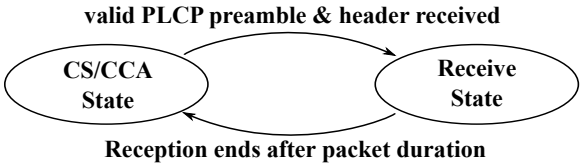

Figure 2·2: Transitions of the PLCP state-machine upon receiving a packet.

valid PLCP preamble and header, it moves to the Rx state and processes the payload using information provided in the PLCP header. In particular, using packet length information, the receiver stays in the Rx state until the last bit of the packet is presumably received. Regardless of whether the payload is correctly received or not, the reception ends at that point and the receiver moves back to the CS/CCA state.

2.1.3 Cyclic redundancy check

A cyclic redundancy check (CRC) is a coding algorithm that detects errors in a packet. In the IEEE 802.11 standard, the CRC field corresponds to the last 32 bits of the MAC frame. When a PHY packet is received, the physical layer forwards the MAC frame to the MAC layer. The MAC layer checks the integrity of the MAC frame using the 32-bit CRC code. If the MAC frame is damaged, then the frame is discarded.

2.1.4 Rate adaptation

The IEEE standard leaves it to the implementer to determines the bit rate for each packet. Accordingly, many rate adaptation algorithms have been proposed in the literature. Several of them such as ARF (Kamerman and Monteban, 1997), Onoe (Project, 2018), and AMRR (Lacage et al., 2004) react to the occurrence of packet losses or lack thereof. If packet losses occur, the algorithms lower the bit rate, while if no

packet loss occurs, they raise the bit rate.

ARF

We next explain the operations of ARF in more detail, since we theoretically analyze its performance in the sequel. In ARF, packets can be transmitted at N different bit rates, denoted by $b_1, b_2, \ldots, b_N$ (e.g., $N = 12$ for IEEE 802.11b/g). The bit rates are sorted from the lowest to the highest. That is, $b_j < b_{j+1}$. ARF changes the bit rate when a certain number of consecutive packets are either transmitted successfully or lost. Specifically, if s consecutive packets are transmitted successfully and the current bit rate b_j is lower than b_N, then ARF raises its bit rate from b_j to b_{j+1}. If f consecutive packets are lost and b_j is higher than b_1, then ARF lowers the bit rate from b_j to b_{j-1}. By default, $s = 10$ and $f = 2$ (Lacage et al., 2004).

Minstrel

Minstrel is a practical, state-of-the-art rate adaptation algorithm that has been implemented within the MadWiFi project and Linux mac80211 driver framework (Berg, 2016). It chooses the bit rate of a transmission based on the throughput measured over past transmissions at different rates. Technically, it selects a bit rate following a retry chain, as shown in Table 2.2.

In Minstrel, 90% of the packets are transmitted at a "normal rate" (fourth column in Table 2.2). The remaining 10% are transmitted at a "lookaround rate" (second and third columns in Table 2.2). Each packet is transmitted at a rate following a retry chain (rows in Table 2.2). For example, consider a packet being transmitted at "lookaround rate". If a random rate is lower than the rate with "best throughput", the packet is first transmitted at the "best throughput" rate, then at the "random rate", then at the "best probability" rate, and finally at the "lowest baserate". The packet is dropped if the transmission fails at the "lowest baserate". The retry chain

Table 2.2: Minstrel Retry Chain (Berg, 2016)

Try	Lookaround rate		Normal rate
	random < best[2]	random > best	
1	Best throughput	Random rate	Best throughput
2	Random rate	Best throughput	2nd best throughput
3	Best probability[3]	Best probability	Best probability
4	Lowest baserate	Lowest baserate	Lowest baserate

table is updated 10 times every second based on performance statistics.

Therefore, a large amount of packet loss does not necessarily cause Minstrel to switch to a low bit rate. Another advantage of Minstrel is that it probes the throughput of different bit rates randomly. This makes the rate adaptation more robust in complicated environment and against some adversaries.

2.2 Hidden node problem

A typical instance of the hidden node problem is illustrated in Figure 2·3. 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 congested if it is forced to retransmit packets.

The hidden node problem can in principle be avoided by enabling the RTS/CTS exchange, which is implemented in Wi-Fi networks. However, the RTS/CTS ex-

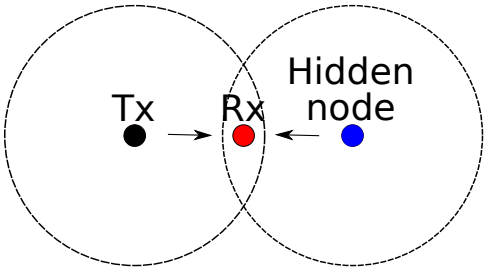

Figure 2·3: Classical hidden node problem. The transmitter and the hidden node cannot sense each other. The collision happens when they transmit simultaneously.

change has not only high overhead, but also does not always fully prevent packet collisions (Ray et al., 2005a) and may lead to deadlocks in multi-hop configurations (Ray and Starobinski, 2007). Generally, it is either turned off (Bellardo and Savage, 2003) or only used for packets whose length exceeds the so-called RTS threshold. Most manufacturers of Wi-Fi cards, including Netgear (Netgear, 2016), TP-LINK (TP-Link, 2016), Linksys (Linksys, 2016) and D-Link (D-link, 2016), disable RTS/CTS altogether by setting the RTS threshold to a sufficiently high default value (e.g., 2346 bytes, which corresponds to the maximum length of an IEEE 802.11 frame). They furthermore recommend to not change the default setting.

2.3 Receiver capture effect

Since the PLCP header only contains radio information, the address of the intended receiver of a packet is unknown until it is checked by the MAC layer. Thus, at the physical layer, a node processes any packet heard from the air. Once the node detects a valid PLCP preamble and header, it moves from the CS/CCA to the Rx state. Therefore, the node cannot detect the PLCP preamble and header of other packets (including those destined to itself) until the current packet reception ends. This phenomenon is known as the *receiver capture effect* (Jiang and Liew, 2007).

2.4 Related work

2.4.1 Denial of service attacks

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. (Poisel, 2011) 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 (Pelechrinis et al., 2011). For instance, (Lin and Noubir, 2005) shows that the energy consumption of a smart jamming attack can be four orders of magnitude lower than continuous jamming. The works in (Noubir et al., 2011; Orakcal and Starobinski, 2014) 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 thesis 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 (Chen et al., 2007; Rayanchu et al.,). However, to the best of our knowledge, none of these rate adaptation algorithms are used in practice.

2.4.2 Interference coupling

Interference coupling caused by hidden nodes is studied by (Chen et al., 2007; Ray et al., 2005b; Broustis et al., 2007), though none of these works consider security ramifications. The work in (Chen et al., 2007) shows that coupling causes nodes to transmit at low bit rates, thus aggravating packet losses. The work in (Ray et al., 2005b) conducts a queuing-theoretic analysis of a chain of neighboring cells with hidden nodes. The analysis reveals that the impact of hidden nodes propagates through the network, causing some nodes to congest at load as low as 15% of the capacity.

The work in (Broustis et al., 2007) perform measurements of a multi-cell IEEE 802.11 network in an indoor testbed. The experiments clearly shows the existence of hidden nodes and the effects of interference coupling in a real world setting. The experimental results also show that hidden nodes cause fairness issues. These fairness issues as well as throughput performance of the network get even worse when RTS/CTS is enabled. Other drawbacks of the RTS/CTS procedure are discussed in (Ray et al., 2003; Xu et al., 2003).

2.4.3 Phase transition

The attacks that we are investigating bear similarity to cascading failures in power transmission systems (Kinney et al., 2005; Soltan et al., 2014). 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, (Haenggi et al., 2009; Kong and Yeh,) 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 dissertation 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 (Aziz et al., 2009; Aziz et al., 2011) show 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 (Ray et al., 2005b; Saligrama and Starobinski, 2006) show that local coupling due to interferences can have global effects on wireless networks. Thus, (Ray et al., 2005b) 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 (Saligrama and Starobinski, 2006) 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.

2.4.4 Physical capture phenomenon

We next explain the differences between the well-known *capture* effect and the lesser known *receiver's capture* effect, which is the focus of our work. The *capture* effect pertains to the fact that two overlapping transmissions may not necessarily result in a packet loss. Specifically, if the power of a detected packet exceeds the combined power of interfering signals beyond a certain threshold, then that packet can still be decoded successfully. In Wi-Fi networks, this effect occurs only when the packet with the strongest power is received before others. That is, the packet with the

highest power is transmitted first. The works in (Durvy et al., 2007; Hadzi-Velkov and Spasenovski, 2003; Li and Zeng, 2006; Nyandoro et al., 2007; Daneshgaran et al., 2008) provide models of Wi-Fi networks integrating the capture effect and show that the packet loss probability can be significantly lower than in models that ignore the capture effect, e.g., Bianchi's Markov model (Bianchi, 2000).

Under the *receiver capture* effect (Jiang and Liew, 2007), a receiver aligns its state machine with information provided by the PHY header of the first transmission, before the second packet arrives. We stress that the receiver does not need be the intended recipient of the first transmission (because there is no destination address in the PHY header). Under the receiver capture effect, the second transmission (which may be a packet destined to the receiver) cannot be properly decoded and this packet is lost. Note that such a scenario occurs in network configurations where hidden nodes are present, but not in configurations where nodes can all hear each other (Jiang and Liew, 2007). Modern simulators, such as ns-3, take the receiver capture effect into account in their physical layer models. In this work, we show how an adversary can exploit the receiver capture effect to launch a cascading attacks on a Wi-Fi network.

2.4.5 Effect of MAC timings

The effect of MAC timing parameters on the performance of IEEE 802.11 networks has been extensively studied in the literature (Bianchi, 2000; Calì et al., 2000; Magistretti et al., 2011; Sun and Dai, 2015; Kumar et al., 2007; Dai and Sun, 2013; Foh and Tantra, 2005; Duda et al., 2008). In particular, an analysis carried out in (Duda et al., 2008) shows that in the absence of contention between nodes, MAC overhead significantly affects throughput, especially at high bit rates. In contrast to those works, the focus of our work is to assess the impact of the MAC overhead on the feasibility of launching a cascading DoS attack. Interestingly, we show that a larger MAC overhead can help prevent such attacks (by mitigating the impact of hidden

nodes).

2.5 Summary

After reviewing the related work, we summarize our findings as follows:

- A DoS attack aims to make communication impossible for legitimate users. It can be launched by creating high power interference or by exploiting protocol vulnerabilities across various layers in the stack to achieve high gain and energy efficiency. However, most of those attacks require physical proximity and are usually non-protocol complaint. These limitations can be used to identify and locate the attacker. In contrast, in this dissertation we show how a protocol-compliant DoS attack can be remotely launched by exploiting coupling due to hidden nodes.

- Wi-Fi supports the RTS/CTS mechanism to solve the hidden node problem. However, RTS/CTS causes fairness issues and degrades throughput performance in many scenarios. Therefore, most Wi-Fi systems today operate without RTS/CTS and hidden nodes are still present.

- There exist two capture phenomena at the physical layer of Wi-Fi. One is the well-known *capture* effect, which pertains to the fact that two overlapping transmissions may not necessarily result in a packet loss. The other one is the so-called *receiver capture* effect, under which a receiver aligns its state machine with information provided by the PHY header of the first transmission, before the second packet arrives.

- The MAC timing parameter settings in IEEE 802.11 networks have a high impact on throughput performance, especially at high bit rates.

This dissertation explore the security issues associated with these cross-layer phenomena in Wi-Fi and demonstrate a new type of DoS attack, which is protocol-compliant and can be remotely launched. Our work differs in several aspects from previous work in the literature. 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 models are original and capture 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.

Chapter 3

Cascading Attacks with Strong Interferers

3.1 Motivation

In this chapter, we demonstrate cascading DoS attacks on Wi-Fi networks and investigate sufficient conditions to launch such attacks. Toward that end, we consider a network with a strong coupling effect between neighboring Wi-Fi cells. As described in Chapter 1, the attack is due to an interference coupling effect caused by hidden nodes. We consider *strong hidden nodes*, whereby interference due to a hidden node destroys all the ongoing transmissions within its communication range, hence creating strong coupling between neighboring Wi-Fi network cells.

The contributions of this chapter 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. Thus, under a phase

transition, the utilization of a remote node experiences no change until it is suddenly forced to congestion (Saligrama and Starobinski, 2006).

Finally, we introduce a new analytical model that sheds light into the phase transition observed in the simulations and experiments. We apply fixed point theorems to this model. 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 a phase transition can occur for the default value of the retry limit in Wi-Fi, which is 7. We carry out a stability analysis and demonstrate that in the phase transition region the system must have multiple fixed points, one of which being unstable.

The rest of the chapter is organized as follows. In Section 3.2, we describe the attack scenario to investigate the feasibility of the attack. We present and discuss experimental and simulation results in Section 3.3. In Section 3.4, we present an analytical model that explains the behaviour of the network and the impact of various parameters, and compare the analytical and simulation results. In Section 3.5, we summarize the chapter and discuss possible mitigation methods.

3.2 Attack Scenario

We first explain how a cascading DoS attack can unfold. We consider a network configuration consisting of a chain of N pairs of nodes. Figure 3·1 depicts the configuration. The ith pair is denoted (A_i, B_i), where $i \geq 0$. Each node A_i transmits packets to node B_i (one-hop communication). Furthermore, each node A_i is a *hidden node* with respect to node A_{i+1}, which means that node A_i cannot sense a transmission by node A_{i+1}. Here, we assume that the hidden node is always strong, i.e., if a transmission by node A_i overlaps with a transmission by node A_{i+1}, a packet collision occurs at node B_{i+1}. This collision forces node A_{i+1} to retransmit its packet using

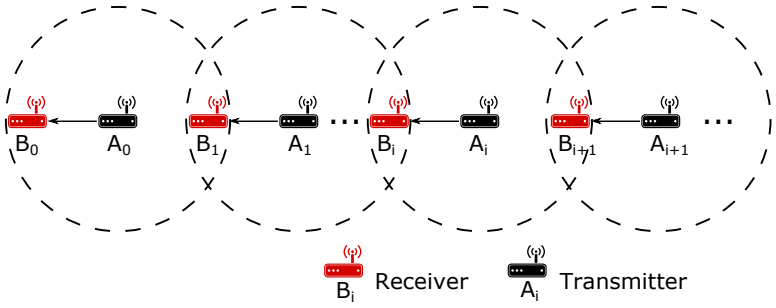

Figure 3·1: Network configuration. The dotted circles represent the communication range of nodes A_i. Nodes A_i transmit packets to nodes B_i $(i = 0, 1, \ldots)$. Each transmission pair (A_i, B_i) belongs to a different cell. Nodes A_i are hidden nodes with respect to nodes A_{i+1}.

the procedure described in Section 2.1.1.

In this configuration, suppose node A_0 (the attacker) starts increasing the rate at which it generates packets and transmits them over the channel (in compliance with the IEEE 802.11 standard). These transmissions will cause collisions at node B_1, which forces node A_1 to increase the rate at which it attempts to transmit packets over the channel (due to retransmissions). The increased rate of transmission attempts by A_1 will in turn impact pair (A_2, B_2) and so forth. Under certain conditions, this effect may amplify along the chain and cause a large fraction of transmission attempts to fail and result in unstable queues (i.e., the rate at which nodes can successfully transmit packets over the channel is lower than the rate at which packets are generated).

3.3 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-3.22 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).

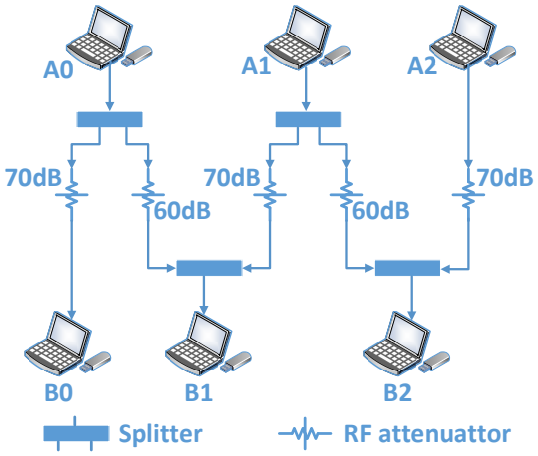

Figure 3·2: Experimental testbed.

3.3.1 Experiments

We set up an experimental testbed composed of six nodes. The testbed configuration is shown in Figure 3·2. 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 ($i \in 0, 1, 2$), and 60 dB attenuators on links between nodes A_i and B_{i+1}. The difference in the signal attenuation of different links 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 (Stein, 1998). The transmission power of each node is set to 0 dBm. We use iPerf 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 3·3 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

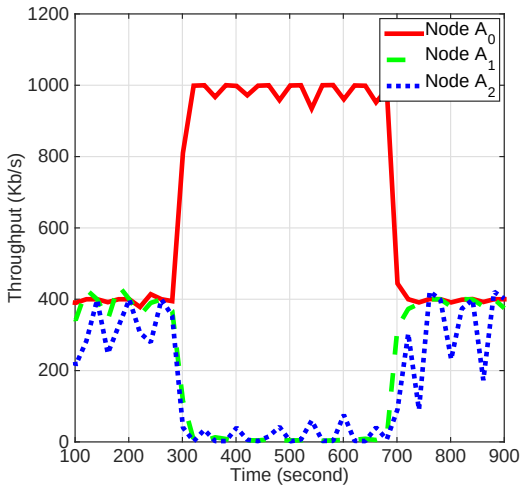

Figure 3·3: Throughput performance measurements in testbed. When node A_0 starts increasing its packet generation rate, the throughput of nodes A_1 and A_2 vanishes.

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. This result is similar to the simulation result in Section 3.3.2. This experimentally proves that the cascading DoS attack can be achieved in practice.

3.3.2 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.

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;
4. In ring topology networks;
5. In an indoor scenario;

and the countering of cascading DoS attacks

6. In networks with RTS/CTS enabled.

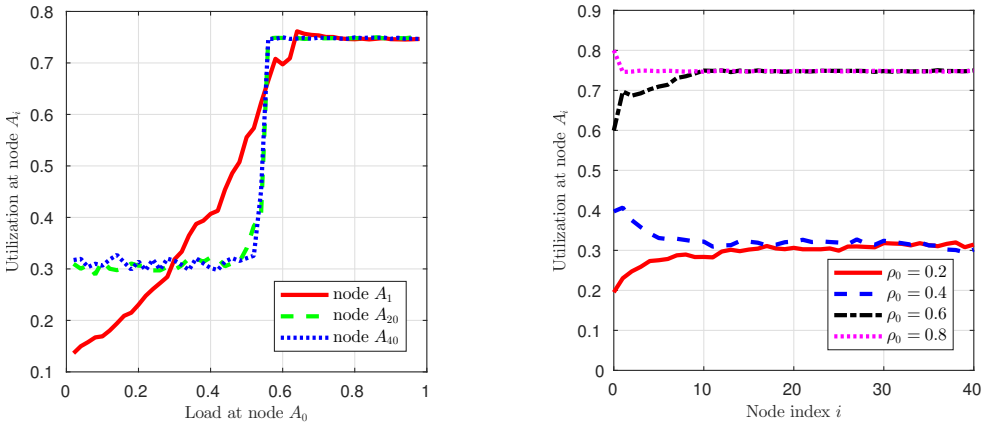

(a) As the traffic load at node A_0 increases, the utilization of remote nodes (e.g., A_{20} and A_{40}) exhibits a phase transition.

(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.

Figure 3·4: Occurrence of cascading DoS attacks in ad hoc networks with fixed bit rate.

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 41 pairs of nodes (i.e. a sequence of 41 hidden nodes), as shown in Figure 3·1. Each packet is transmitted over a single-hop path (similar to Wi-Fi Direct). We fix the bit rate to 1 Mb/s and the retry limit to $R = 7$.

We set up a Wi-Fi network using the standard IEEE 802.11 library in ns-3. At each node A_i, $i \geq 1$, the generation rate of UDP packets is $\lambda_i = 8.125$ pkts/s. The generation rate of UDP packets at node A_0, λ_0, varies from 1.25 to 61.25 pkts/s. Packets at each node are generated according to a Poisson process, hence different nodes start transmitting at different times. 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 3·4(a) 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.

Figure 3·4(b) illustrates the phase transition in a different way. The figure depicts

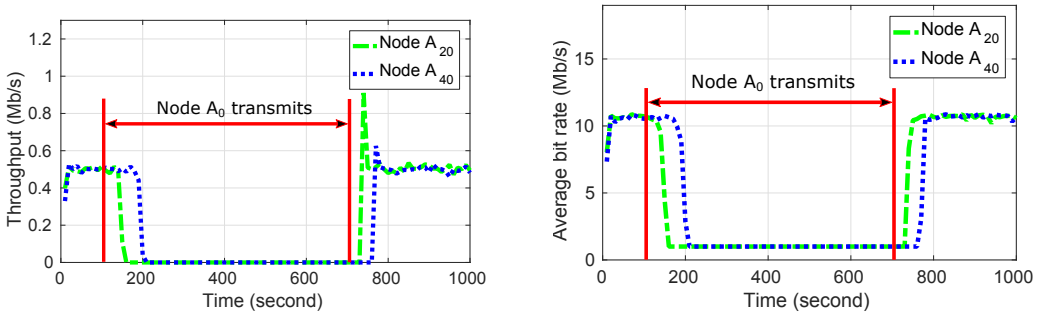

(a) Throughput **(b)** Bit rate

Figure 3·5: 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).

the utilization of each node A_i for $i \geq 1$, as i increases. Again, we observe that different values of ρ_0 lead to two completely distinct behaviour for the sequence of utilizations $(u_i)_{i=0}^{40}$ (i.e., $u_{40} \simeq 0.3$ when $\rho_0 = 0.2$ and $\rho_0 = 0.4$, while $u_{40} \simeq 0.75$ when $\rho_0 = 0.6$ and $\rho_0 = 0.8$). Note that the upper limit of the utilization does not reach 1, due to inter-frame spacing requirements and (random) backoff delays mandated by IEEE 802.11.

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 3.3.2. In our simulations, the parameter *EWMA* of Minstrel is set to 0.25 (Xia et al.,).

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 3·5, 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.

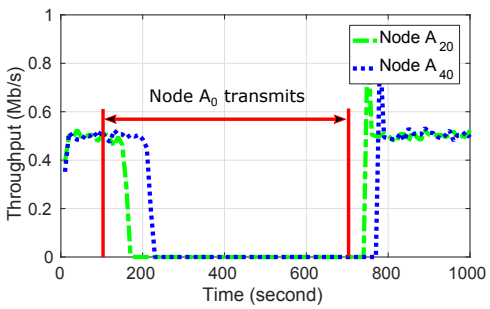

(a) When node A_0 generates packets at 5 Mb/s and transmits, the throughput of nodes A_{20} and A_{40} vanishes.

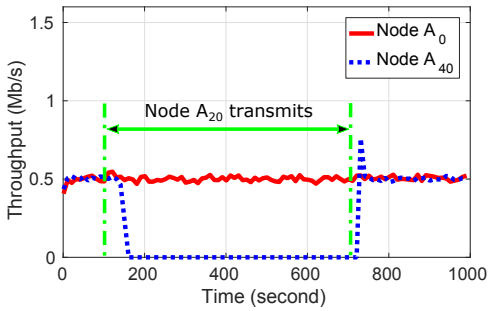

(b) When node A_{20} generates packets at 5 Mb/s and transmits, the throughput of node A_{40} vanishes while the throughput of node A_0 does not.

Figure 3·6: Simulation results under AP mode without reassociation. Nodes A_i are stations and nodes B_i are access points, for $i \in \{0, 1, 2, \dots\}$.

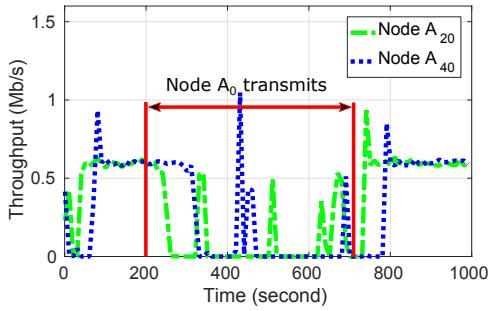

Figure 3·7: Simulation results under AP mode with reassociation. When node A_0 generates packets at 5 Mb/s and transmits, the throughput of node A_{20} and A_{40} significantly decreases.

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 3.3.2 except that we set nodes B_i as access points, and nodes A_i as stations. The initial beacon starting time at each AP is a random variable that is uniformly distributed between 0 and 102.4 ms.

We first investigate the cases where stations do not restart association when beacons are missing. Toward this end, we set the number of consecutive beacons that must be missed before restarting association, i.e. the attribute `MaxMissBeacons` in ns-3, to a large value. Otherwise, we use the default settings of ns-3 for the APs and the stations (ns-3, 2018). Figure 3·6 shows similar results as in Section 3.3.2, namely when a cascading DoS attack is launched by node A_0, as shown in Figure 3·6(a), the remote nodes A_{20} and A_{40} in the sequence exhibit a phase transition. If the attacker is node A_{20}, the simulation result in Figure 3·6(b) shows that the throughput of node A_{40} vanishes but the throughput of node A_0 does not. This result shows that an attack can be launched from any node A_i in the topology and the following nodes in the sequence (i.e., $A_{i+1}, A_{i+2}, \ldots$) will experience congestion.

We next consider the case where stations restart association when beacons are missing. We set `MaxMissBeacons` $= 10$, which is the default value in ns-3 (ns-3, 2018). The simulation results are shown in Figure 3·7. When Node A_0 starts to transmit packets, we observe a significant throughput degradation at nodes A_{20} and A_{40}, but the throughput does not vanish completely. The reason is that if A_i disassociates from its AP B_i over a certain period then node A_{i+1} is not affected by interference coupling during that period. This result indicates that reassociations help mitigate cascading DoS attacks, though throughput performance is still significantly impaired.

Ring topology

We investigate cascading DoS attacks in a ring topology with 41 pairs of nodes, as shown in Figure 3·8. 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. Node A_i $(i = 0, 1, \dots)$ generate packets at rate 0.5 Mb/s, following a Poisson process. At time $t = 300$ s, node A_0 increases its packet generation rate to 11 Mb/s and the throughput of all the 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.

This result is illustrated in Figure 3·9. During the first 100 seconds, all the nodes A_i $(i = 0, 1, \dots)$ generate packets at 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.

Building model

In this section, we use the ns-3 `HybridBuildingsPropagationLossModel` library (ns-3, 2018) to demonstrate the feasibility of cascading DoS attacks in an 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, 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,

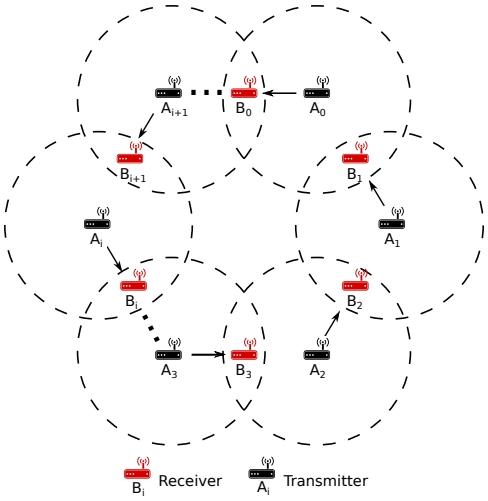

Figure 3·8: Ring topology under cascading DoS attack. The dash circle represents the transmission range of the transmitter.

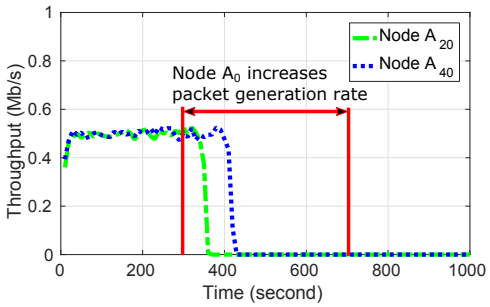

Figure 3·9: Simulation results under a ring topology. When the packet generation rate of node A_0 increases, the throughput of nodes A_{20} and A_{40} vanishes. This effect continues even when the packet generation rate of node A_0 decreases.

office, and commercial).

In our simulations, we consider a 20-floor office building with six rooms in each floor, as shown in Figure 3·10. 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, $i \geq 1$, 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 3·11. 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.

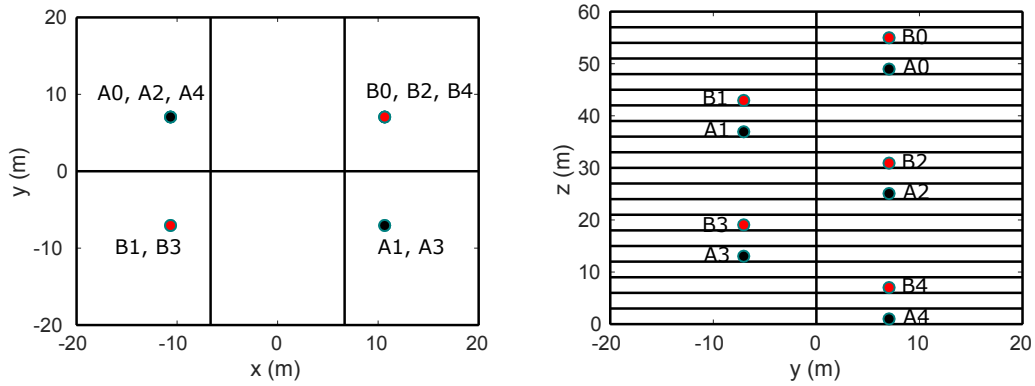

(a) Top view. **(b)** Side view.

Figure 3·10: Office building model. The building has 20 floors (z-axis) and 6 rooms in each floor (x and y axes).

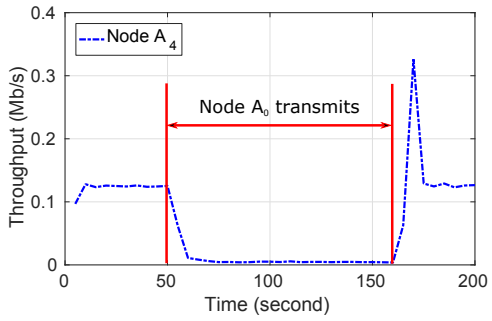

Figure 3·11: Simulation results using ns-3 building model. When node A_0 transmits, the throughput of remote node A_4 collapses.

RTS/CTS

We next evaluate the impact of enabling RTS/CTS in the topology under consideration. Specifically, we repeat the simulations of Section 3.3.2, but with RTS/CTS enabled. Figure 3·12 shows that transmissions by node A_0, which start after 100 s,

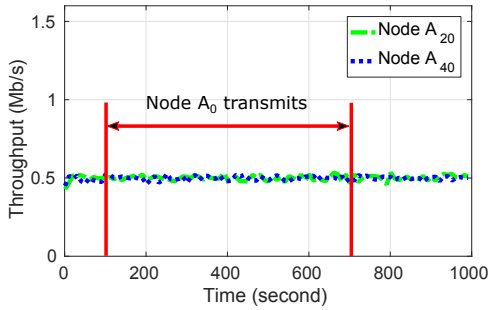

Figure 3·12: Simulation results when enable RTS/CTS. The increase of the packet generation rate of node A_0 does not affect the throughput of nodes A_{20} and A_{40}.

have no effect on the throughput of remote nodes A_{20} and A_{40}. This shows that RTS/CTS is an effective solution against cascading DoS attacks in this scenario.

3.4 Analysis

In this section, we develop a stylized, analytical model that provides qualitative insight into the network behavior observed in the simulations and experiments for the linear topology. 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.

3.4.1 Model

We consider the linear topology shown in Figure 3·1. 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 $1 \leq \overline{r}_i \leq R$ 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 $0 \leq u_i \leq 1$, 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

$$u_i = \min\{\overline{r}_i \lambda_i T, 1\}. \tag{3.1}$$

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 (Rong and Ephremides, 2009). 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.

Though the assumption of Poisson retransmissions is not fully consistent with the Wi-Fi protocol, it is similar to the "random-look" model used by Kleinrock and Tobagi in their analysis of (single hop) random access networks (Kleinrock et al., 1975) (see also (Bertsekas and Gallager, 1992)[Ch. 4]). The simulations do not incorporate the simplifications used to make the analysis tractable, yet lead to the same effects. We stress that beside these assumptions, the rest of our analysis is exact.

3.4.2 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

$$\rho_0 \triangleq \lambda_0 T, \tag{3.2}$$

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

$$\rho \triangleq \lambda_i T, \tag{3.3}$$

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 function of the parameters ρ and R.

First, we get the utilization at node A_0:

$$u_0 = \min\{\rho_0, 1\}. \tag{3.4}$$

We next develop an iterative procedure to derive u_{i+1} from u_i. From (3.1) and (3.3),

$$u_{i+1} = \min\{\overline{r}_{i+1}\rho, 1\}. \tag{3.5}$$

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

$$\begin{aligned}\overline{r}_{i+1} &= \sum_{r=1}^{R} r \cdot (1-p_{i+1}) \cdot (p_{i+1})^{r-1} + R \cdot (p_{i+1})^{R} \\ &= \sum_{r=1}^{R} (p_{i+1})^{r-1}. \end{aligned} \tag{3.6}$$

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 $t = t'$. We compute p_{i+1} by conditioning on whether or not node A_i is transmitting at time t'. Note that due the Poisson Arrivals See Time Averages (PASTA) property, the transmission state of node A_i at time $t = t'$ is the same as at any random point of time.

If node A_i transmits at time t', 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', which occurs with probability $1 - u_i$, then a collision occurs only if node A_i starts a transmission during the interval $[t', t' + T]$. Since the packet inter-arrival time on the channel is exponentially distributed with mean $\overline{r}_i T$, such an event occurs with probability

$$(1 - \mathrm{e}^{-\overline{r}_i \lambda_i T}) = (1 - \mathrm{e}^{-u_i}), \tag{3.7}$$

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

$$\begin{aligned} p_{i+1} &= 1 \cdot u_i + (1 - \mathrm{e}^{-u_i}) \cdot (1 - u_i) \\ &= 1 - \mathrm{e}^{-u_i}(1 - u_i). \end{aligned} \tag{3.8}$$

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 (3.8) still provides the correct result.

Putting (3.5), (3.6), and (3.8) together, we obtain

$$u_{i+1} = \min\left\{\rho \sum_{r=1}^{R} \left(1 - \mathrm{e}^{-u_i}(1-u_i)\right)^{r-1}, 1\right\}. \tag{3.9}$$

3.4.3 Limiting behavior of the utilization

We next analyze the limiting behavior of the iteration given by (3.9). The sequence $(u_i)_{i=0}^{\infty}$ corresponds to a discrete non-linear dynamical system (Lynch, 2004). 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

$$f(u_i) \triangleq \rho \sum_{r=1}^{R} \left(1 - \mathrm{e}^{-u_i}(1-u_i)\right)^{r-1}. \tag{3.10}$$

We then rewrite (3.9) as follows:

$$u_{i+1} = \min\left\{f(u_i), 1\right\}. \tag{3.11}$$

We say that $\omega \in [0,1]$ is a *fixed point* of (3.11) if

$$\omega = \min\left\{f(\omega), 1\right\}. \tag{3.12}$$

Suppose (3.12) has K different fixed points (Theorem 4 in the sequel will show that $K \geq 1$). We denote by Ω the ordered set of all the fixed points of (3.12). That is,

$$\Omega \triangleq \{\omega_1, \ldots, \omega_k, \ldots, \omega_K\}, \tag{3.13}$$

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

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

Lemma 1. *Let $u, u' \in (\omega_k, \omega_{k+1})$, where $k \in \{1, \ldots, K-1\}$. If $f(u) > u$, then $f(u') > u'$. If $f(u) < u$, then $f(u') < u'$.*

Proof. The proof goes by contradiction. Let $u, u' \in (\omega_k, \omega_{k+1})$. Suppose $f(u) > u$ and $f(u') < u'$. Since f is continuous in (ω_k, ω_{k+1}), then by the intermediate-value theorem there exists a point u'' between u and u' such that $f(u'') = u''$. Thus, u'' is a fixed point of (3.12). 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 2.

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.*

Proof.

1. Let $\omega_k < u_0 < \omega_{k+1}$, where $k \in \{1, \ldots, K-1\}$. Since $p_i \in (0,1)$. Therefore, the function f is continuous and monotonically increasing, $f(\omega_k) < f(u_0) < f(\omega_{k+1})$. Hence, according to (3.11) and (3.12), we get

$$\omega_k \leq u_1 \leq \omega_{k+1}. \tag{3.14}$$

   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 (3.14), we have $u_2 = f(u_1) > u_1$. Applying the same argument inductively, either there exists some value $M \geq 2$ such that $u_i = \omega_{k+1}$ for all $i \geq M$, or the sequence $(u_i)_{i=0}^{\infty}$ 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 $(u_i)_{i=0}^{\infty}$ 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 $u \in [0, \omega_1)$ such that $f(u) > u$, then $f(u') > u'$ for all $u' \in [0, \omega_1)$. Since $f(0) = \rho > 0$, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_1.

3. This is handled similarly to case 2.

□

3.4.4 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 (3.12). First, we investigate the case $\omega = 1$.

Lemma 3. *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 (3.12) 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 2.
3. This is handled similarly to Part 2.

□

Lemma 3 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, (3.12) can be rewritten

$$f(\omega) = \omega. \tag{3.15}$$

We look for solutions of (3.15) that belong to the interval $[0, 1]$. Each such solution is an element of Ω.

Equation (3.15) is difficult to work with because it contains two unknown variables, ρ and R. To circumvent this difficulty, we introduce the function

$$h_R(\omega) \triangleq \frac{\rho\omega}{f(\omega)} = \frac{\omega}{\sum_{r=1}^{R} \left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)^{r-1}}. \tag{3.16}$$

For each value of ρ, the solutions of (3.15) must satisfy

$$h_R(\omega) = \rho. \tag{3.17}$$

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

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

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

Theorem 4.

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, (3.17) 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, (3.17) 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 3. Thus, there are $K \ge 2$ fixed points.

   By Part 3 of Lemma 3, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_K for all $u_0 \in (\omega_{K-1}, \omega_K]$. Moreover, by Theorem 2, the limit of the sequence $(u_i)_{i=0}^{\infty}$ is no larger than ω_{K-1} for all $u_0 \le \omega_{K-1}$. Hence, a phase transition exists at ω_{K-1}.

3. If $\rho > h_R^{max}$, then (3.15) has no solution. Moreover, since $\rho > h_R^{max} \ge h_R(1) = 1/R$, we get $\rho > 1/R$. By Parts 1 and 2 of Lemma 3, 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 4 for different values of R, using Figure 3·13. First, consider $R = 4$ as shown in Figure 3·13(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 3·13(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 4, 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 3·13(c). Furthermore, the size of this region is larger since $(1/R, h_R^{max}) = (0.1, 0.162)$.

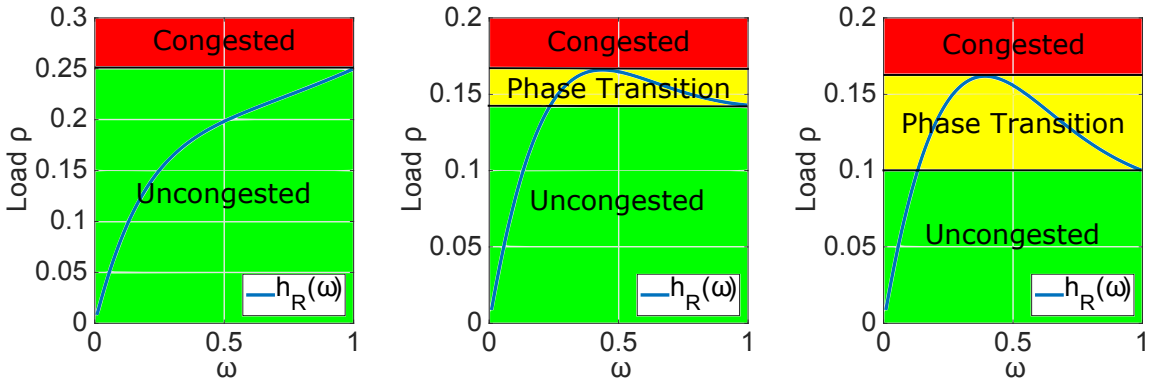

(a) $R = 4$ **(b)** $R = 7$ **(c)** $R = 10$

Figure 3·13: Illustration of the different network regimes for different values of R. For each value of 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 $h_R(1) = 1/R$.

3.4.5 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 5. *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 $R > R^*$.*

Proof. Let $\omega \in [0,1]$. Since

$$\left(h_R^{-1}(\omega)\right)' = -\frac{h'_R(\omega)}{h_R(\omega)^2}, \tag{3.18}$$

the sign of $h'_R(\omega)$ is opposite to $\left(h_R^{-1}(\omega)\right)'$. Hence, we investigate the sign of

$$\left(h_R^{-1}(\omega)\right)' = \sum_{r=1}^{R} \Psi'_r(\omega), \tag{3.19}$$

where

$$\Psi_r(\omega) \triangleq \frac{\left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)^{r-1}}{\omega}. \tag{3.20}$$

We check the sign of each term $\Psi'_r(\omega)$ in (3.19), for $r \in \{1, 2, \ldots, R\}$. For $r = 1$, we have

$$\Psi'_1(\omega) = \frac{d}{d\omega}\left(\frac{1}{\omega}\right) = -\frac{1}{\omega^2} < 0.$$

For $r \geq 2$, we have

$$\Psi'_r(\omega) = -\frac{e^{-\omega}\left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)^{r-2}\Phi_r(\omega)}{\omega^2}, \tag{3.21}$$

where

$$\Phi_r(\omega) \triangleq -1 + e^{\omega} + (3-2r)\omega + (r-1)\omega^2.$$

Clearly, the terms $e^{-\omega}$, $\left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)^{r-2}$ and ω^2 in (3.21) are all positive. Thus, the signs of $\Phi_r(\omega)$ and $\Psi'_r(\omega)$ are opposite.

We next investigate the signs of the first and second derivatives of the function $\Phi(\omega)$. We have

$$\Phi_r'(\omega) = e^{\omega} + 3 - 2r + 2(r-1)\omega, \tag{3.22}$$

$$\Phi_r''(\omega) = e^{\omega} + 2(r-1) > 0, \tag{3.23}$$

for all $\omega \in [0,1]$ and $r \geq 2$. From (3.23), we find that $\Phi_r'(\omega)$ is monotonically increasing with ω.

For any $r \geq 2$, we obtain from (3.22) that

$$\Phi_r'(0) = 4 - 2r, \tag{3.24}$$

$$\Phi_r'(1) = e + 1. \tag{3.25}$$

We distinguish between three possible cases regarding the sign of $\Phi_r(\omega)$:

1. For $r = 2$, $\Phi_2'(0) = 0$. Hence, $\Phi_2'(\omega) > 0$. The function $\Phi_2(\omega)$ is monotonically increasing with ω. Since $\Phi_2(0) = e - 1 > 0$, $\Phi_2(\omega)$ is always positive.

2. For $r = 3$, $\Phi_3'(0) < 0$. The function $\Phi_3(\omega)$ first decreases then increases as ω increases from 0 to 1. Since $\Phi_3(0) = 0$ and $\Phi_3(1) > 0$, the sign of the function $\Phi_3(\omega)$ turns from negative to positive as ω increases from 0 to 1.

3. For $r > 3$, $\Phi_r'(0) < 0$. The function $\Phi_r(\omega)$ first decreases then increases as ω increases from 0 to 1. Since $\Phi_r(0) = 0$ and $\Phi_r(1) < 0$, the sign of the function $\Phi_r(\omega)$ is always negative.

Therefore, by (3.19), for any given $\omega \in [0,1]$, the sign of the function $\Phi_r(\omega)$ turns from being positive to being negative as r increases. Equivalently, the sign of the function $\Psi_r'(\omega)$ turns from being negative to being positive as r increases.

Thus, by (3.19), if $\big(h_R^{-1}(\omega)\big)'$ is nonnegative for $R = R^*$, then it is also nonnegative for all $R \geq R^*$. Equivalently, by (3.18), if $\big(h_R^{-1}(\omega)\big)'$ is nonpositive for $R = R^*$, then it is also nonpositive for all $R \geq R^*$, which completes the proof. □

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

$$\begin{aligned} h_\infty(\omega) &= (1 - \big(1 - \mathrm{e}^{-\omega}(1-\omega)\big))\omega \\ &= e^{-\omega}(1-\omega)\omega, \end{aligned} \tag{3.26}$$

and its derivative

$$h'_{\infty}(\omega) = e^{-\omega}(1 - 3\omega + \omega^2). \tag{3.27}$$

The next corollary is the logical transposition of Lemma 5.

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 $[0, \alpha)$, where

$$\alpha \triangleq \frac{3 - \sqrt{5}}{2}. \tag{3.28}$$

Lemma 6. *Let $0 \leq \omega < \alpha$. Then, $h'_R(\omega) > 0$, for all $R \geq 1$.*

Proof. Let the function $h_{\infty}(\omega)$ and its derivative $h'_{\infty}(\omega)$ be defined as in (3.26) and (3.27), respectively. Since $e^{-\omega}$ is always positive, $h'_{\infty}(\omega)$ 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 $\bar{w}$ as defined in (3.28).

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

□

The consequence of Lemma 6 is that for all $R \geq 1$,

$$h_R^{max} \geq h_R(\alpha). \tag{3.29}$$

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.

Theorem 7. *Let α be defined as in (3.28) and suppose $h_R(\alpha) > 1/R$. Then, a phase transition is guaranteed to exist for any $\rho \in (1/R, h_R(\alpha))$.*

Proof. From Theorem 4, we know that a phase transition exists if $1/R < \rho < h_R^{max}$. By (3.29) and the assumption that $h_R(\alpha) > 1/R$, the proof follows. □

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

Theorem 8. *Let $h_{\infty}(\omega)$ and α be defined as in (3.26) and (3.28), respectively. Then, $h_R^{max} \geq h_{\infty}(\alpha) \simeq 0.161$.*

Proof. By (3.16),

$$\begin{aligned} h_R(\alpha) &= \frac{\omega}{\sum_{r=1}^{R}(1-e^{-\omega}(1-\omega))^{r-1}} \\ &> \frac{\omega}{\sum_{r=1}^{\infty}(1-e^{-\omega}(1-\omega))^{r-1}} = h_\infty(\alpha). \end{aligned} \tag{3.30}$$

Thus, by (3.29) and (3.30), $h_R^{max} > h_\infty(\alpha) \simeq 0.161$. Note that this bound is asymptotically tight as $R \to \infty$ since $h_\infty^{max} = h_\infty(\alpha)$. □

From Theorems 4 and 8, 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 (3.16), since for any $\omega \in [0,1]$ the denominator increases as R gets larger).

3.4.6 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 Ω and diverges from *unstable* fixed points of Ω. We will show that the stability of the fixed points of (3.15) 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 $\omega \in \Omega$ 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 3, a special fixed point of (3.12) 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 (3.15). To establish the stability of those fixed points, we will employ the following proposition.

Proposition 1 ((Lynch, 2004)). *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 $\omega \in \Omega$ with respect to the function $h_R(\omega)$.

Theorem 9. *Consider a fixed point $\omega \in \Omega$, where $\omega < 1$. Then ω is stable if $h'_R(\omega) > 0$ and unstable if $h'_R(\omega) < 0$.*

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

$$h'_R(\omega) = \frac{1}{\Gamma(\omega)} - \frac{\omega}{(\Gamma(\omega))^2} \cdot \Gamma'(\omega) > 0, \tag{3.31}$$

where

$$\Gamma(\omega) \triangleq \sum_{r=1}^{R} \left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)^{r-1} = \frac{f(\omega)}{\rho}. \tag{3.32}$$

If one can show that (3.31) implies $|f'(\omega)| < 1$, then according to Proposition 1, the fixed point ω is stable. We multiply both sides of (3.31) by $(\Gamma(\omega))^2$ and obtain

$$\Gamma(\omega) - \omega\Gamma'(\omega) > 0. \tag{3.33}$$

Using (3.32) and (3.15), we can rearrange (3.33) as follows:

$$\Gamma'(\omega) < \frac{\Gamma(\omega)}{\omega} = \frac{f(\omega)}{\rho\omega} = \frac{1}{\rho}. \tag{3.34}$$

From (3.32) and (3.34), we get

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

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.

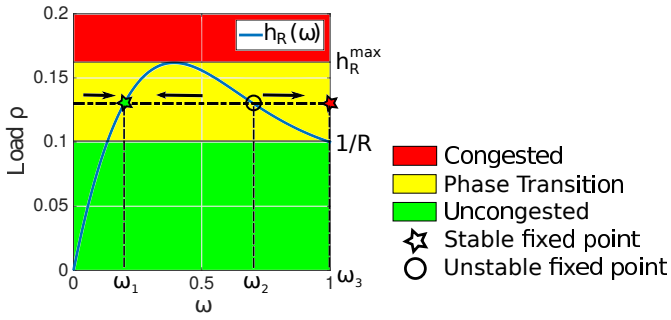

Figure 3·14: Stability of fixed points with $R = 10$. Given a load $\rho = 0.13$ (dash line), Ω contains three fixed points: $\omega_1 = 0.2$, $\omega_2 = 0.7$ and $\omega_3 = 1$. The fixed point ω_1 is stable because $h'_R(\omega_1) > 0$ and ω_2 is unstable because $h'_R(\omega_2) < 0$. The fixed point $\omega_3 = 1$ exists and is stable because $\rho > 1/R$. Therefore, the sequence $(u_i)_{i=0}^{\infty}$ converges to ω_1 if $u_0 < \omega_2$, and to ω_3 if $u_0 > \omega_2$.

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

□

We next show how the stability analysis of the fixed points helps to determine the limit of the sequence $(u_i)_{i=0}^{\infty}$. Consider, for instance, the example shown in Figure 3·14 with parameters $R = 10$ and $\rho = 0.13$. Under these parameters, $\Omega = \{\omega_1, \omega_2, \omega_3\} = \{0.2, 0.7, 1\}$.

The fixed points ω_1 and ω_2 are the solutions of $h_R(\omega) = \rho$. According to Theorem 9, 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 4, ω_2 is a phase transition point. Hence, the sequence $(u_i)_{i=0}^{\infty}$ 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.

3.4.7 Heterogeneous traffic load

In previous subsections, we assumed that node A_0 varies its traffic load ρ_0, but all other nodes A_i $(i \geq 1)$ have the same traffic load ρ. We now relax this assumption

and assume that nodes A_i ($i \geq 1$) 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 10. *Suppose $h_R^{max} > 1/R$. If $\rho_i \in (1/R, h_R^{max})$ for all $i \geq 1$, then a phase transition occurs.*

Proof. Let $\rho_{max} = \max_{i \geq 1} \rho_i$ and $\rho_{min} = \min_{i \geq 1} \rho_i$. According to Theorem 4, 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 $i \geq 1$. □

3.4.8 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 41 pairs of nodes, as described in Section 3.3.2.

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 3·15 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 our simulations for any $R \leq 6$.

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 anal-

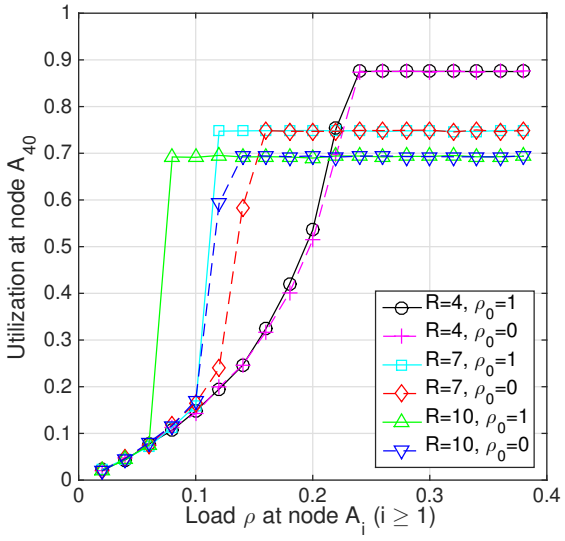

Figure 3·15: Simulation of the limiting behaviour of the node utilization in a network of 41 pairs of nodes. For $R = 4$, the limit is the same when $\rho_0 = 0$ and $\rho_0 = 1$, hence no phase transition is observed. However, for $R = 7$ and $R = 10$, the limits are different, hence showing the existence of a region of load ρ in which a phase transition occurs.

ysis 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.

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 3.4.7. 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 3·16 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

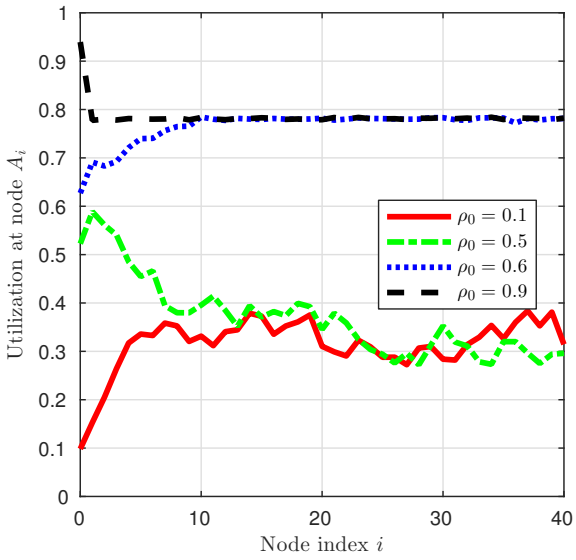

Figure 3·16: Simulation with heterogeneous traffic load in a network with 41 pairs of nodes. The traffic load of nodes A_i ($i \geq 1$) are uniformly distributed between 0.11 and 0.15. For $R = 7$, when the load ρ_0 changes from 0.5 to 0.6, the limiting behavior of the sequence of node utilizations differs, thus indicating the occurrence of phase transition.

exceeds 0.6, the sequence of node utilizations converges to its upper limit, implying that the network is congested.

3.5 Summary

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 behaviour. 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 $R \geq 7$ and provides a simple and explicit estimate of traffic load at each node under which a phase transition occurs (i.e., $\rho_i \in (1/R, 0.161)$ for all $i \geq 1$). 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. Although the analysis is based on some simplifying assumptions, the estimate is not far from the values observed in the simulations.

Chapter 4

Cascading Attacks with Weak Interferers

4.1 Motivation

In Chapter 3, we investigated cascading DoS attacks with strong hidden nodes. However, since the effect of the attack propagates through a chain of nodes in the network, we cannot control the strength of the interference except for the first node (i.e., the attacker). In a more general scenario, interference coupling between two neighboring networks can be weak, i.e., the interference caused by hidden node does not always disrupt an ongoing transmissions. Is the cascading DoS attack still feasible in such a scenario?

In this chapter, we investigate cases where interference caused by hidden nodes is on the same order or weaker than the signals of sending stations. Our main objective is to find out whether cascading attacks are still feasible in those situations. Through extensive ns-3 simulations and mathematical analysis, we provide a positive answer to this question. The attack leverages two phenomena (Xin and Starobinski, 2018), which we describe in detail in the chapter.

The first phenomenon is a PHY-layer effect known as *receiver capture* (Jiang and Liew, 2007). Accordingly, if the PHY header of the packet transmitted by the hidden node is decoded first, the packet sent by the station is lost (assuming the two packet transmissions overlap). Thus, even though not all packets transmitted by the station are lost, a large fraction still is.

The second phenomenon relates to bit rate adaptation. Specifically, rate adapta-

tion algorithms vary the bit rate used for packet transmissions based on the observed quality of the channel. While different algorithms have been proposed in the literature (Biaz and Wu, 2008), most gradually lower the bit rate upon experiencing packet losses. Since packet losses are still possible due to the receiver capture effect, rate adaption algorithms may end up with significantly lowering the bit rate of Wi-Fi stations, sometimes down to the base rate of 1 Mb/s. As a result, the capacity of the shared channel is drastically reduced (since each packet transmission uses the shared channel for a longer amount of time) while leading to a traffic congestion.

The main contributions of this chapter can thus be summarized as follows. We first identify and document a coupling effect between neighboring cells, due to hidden nodes and receiver capture. We then analyze and provide simulations of the packet loss probability with and without receiver capture. We show that with receiver capture, a packet sent by a station is lost irrespective of the signal-to-interference ratio (SIR) and the bit rate. Leveraging the above coupling effect, we demonstrate the feasibility of launching cascading attacks on Wi-Fi networks using weak hidden nodes (i.e., hidden nodes producing weak interference). Through extensive ns-3 simulations, including for an indoor building model, we show that the coupling effect may propagate, thus reducing the channel capacity across an entire chain of Wi-Fi cells. These results apply to several rate adaptation algorithms. We next provide an analysis of the limiting behavior of the channel utilization in a chain of Wi-Fi cells, assuming nodes implement the Auto Rate Fallback (ARF) rate adaptation algorithm (Kamerman and Monteban, 1997). In particular, we show how the average bit rate experiences a sharp drop as the channel utilization gets higher, which provides insight into how the attack is able to propagate throughout the network.

The rest of the this chapter is organized as follows. In Section 4.2, we introduce the network model used for studying cascading attacks. In Section 4.3, we present

numerical simulations illustrating receiver capture and cascading attacks in various scenarios. In Section 4.4, we propose analytical models to shed light on and provide further validation of the results observed in the simulations. We conclude the chapter in Section 4.5.

4.2 Attack Scenario

We now describe the attack scenario studied in this chapter. Our goal is to assess the feasibility of a chain reaction. Hence, we consider the network topology shown in Figure 4·1 which is the same as in Section 3.2. This topology consists of M single-hop pairs of Wi-Fi nodes. Each pair belongs to a different cell $i \in (1, 2, 3, \ldots, M)$. In each cell i, node A_i transmits packets to node B_i. The dash circle around each node A_i represents its communication range. We assume that node B_i can sense transmissions from both nodes A_i and A_{i-1}, but node A_i and node A_{i-1} cannot sense each other (in Section 4.3.2, we present a realistic network configuration where such a scenario can occur). Thus, node A_{i-1} is a hidden node with respect to node A_i.

In this chapter, we consider the practical scenario where at node B_i the signal received from the transmitter A_i is stronger than that received from the hidden node A_{i-1}. This situation occurs, for instance, when the transmission powers of nodes A_i and A_{i-1} are the same, but the path loss between A_i and B_i is lower than between A_{i-1} and B_i (Giustiniano et al., 2007). This scenario is likely when nodes belonging to the same cell are located closer than nodes belonging to different cells. In this scenario, the hidden node causes weak interference, but a packet loss is still possible due to the receiver capture effect.

If a packet loss occurs at node B_i, node A_i keeps retransmitting that packet as long as the number of retransmissions (i.e, the *retry count*) does not exceed a certain upper limit (i.e., the *retry limit*). We also assume that each node A_i in the topology

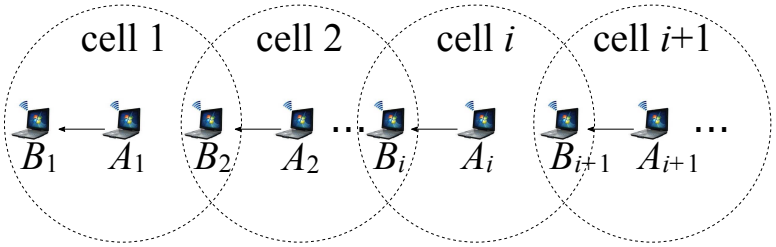

Figure 4·1: Attack scenario. Each node A_i transmits packets to node B_i. Node A_{i-1} is a hidden node with respect to A_i. Node A_1 is the attacker (first hidden node in the chain).

runs a rate adaptation algorithm, such as ARF, Onoe or AMRR.

The attack proceeds as follows. An adversary that controls node A_1 increases its packet generation rate. Though the interference generated by node A_1 is weak, it stills leads to packet losses on the link between A_2 and B_2 due to receiver capture. Hence, node A_2 retransmits packets and gradually lowers its bit rate. The lower bit rate prolongs the duration of packet transmissions. Hence, transmissions by node A_2 add interference on the link between A_3 and B_3 which increases the rate of packet losses between A_3 and B_3 and reduces the bit rate of node A_3, and so on. In the next section, we investigate the feasibility and impact of this attack under different network settings.

4.3 Simulations

In this section, we run ns-3 simulations of IEEE 802.11g/n networks for the scenario depicted in Figure 4·1. We define the *utilization* of node A_i as the average fraction of time during which node A_i transmits. In addition, we denote by *SIR* the ratio of the signal strength of node A_i (signal) to the signal strength of node A_{i-1} (interference) observed at node B_i.

4.3.1 Hidden node

Our first simulation evaluates the impact of a hidden node on the packet loss probability for different SIR and bit rates. We consider a network comprising $M = 2$ cells. In each cell $i \in \{1, 2\}$, node A_i sends 1500-byte UDP packets to node B_i, $i \in \{1, 2\}$. Thus, node A_1 is a hidden node with respect to node A_2. We repeat each simulation 10 times and average the results (we also compute 95% confidential intervals, but since they are very narrow, they cannot be seen on the graphs).

We consider two cases:

1. Node A_1 starts transmitting a packet right after node A_2 finishes transmitting the PLCP preamble and header of a packet.

2. Node A_2 starts transmitting a packet right after node A_1 finishes transmitting the PLCP preamble and header of a packet.

Figure 4·2 shows the results for different SIR and bit rates. Case (1) (the curve in blue) corresponds to the classical *capture* effect. If the SIR is above a certain threshold, the packet transmitted by node A_2 is received with high probability by node B_2, otherwise it is lost. We note that the threshold is pretty sharp and depends on the bit rate. The higher the bit rate, the higher the SIR needed to successfully receive a packet. For instance, Figure 4·2 (a) shows that the threshold for a successful reception by B_2 at 1 Mb/s is about 0 dB. Figures 4·2(b), (c), and (d) show that for bit rates of 12 Mb/s, 24 Mb/s, and 54 Mb/s, the SIR threshold for successful reception is about 10 dB, 15 dB, and 25 dB, respectively.

On the other hand, case (2) (the curve in red) corresponds to the *receiver capture* effect. We observe that no matter how high the SIR is, a packet transmitted by node A_2 is always lost. We conclude that the order of packet arrivals plays a critical role and that a weak hidden node can still induce significant packet losses, even at

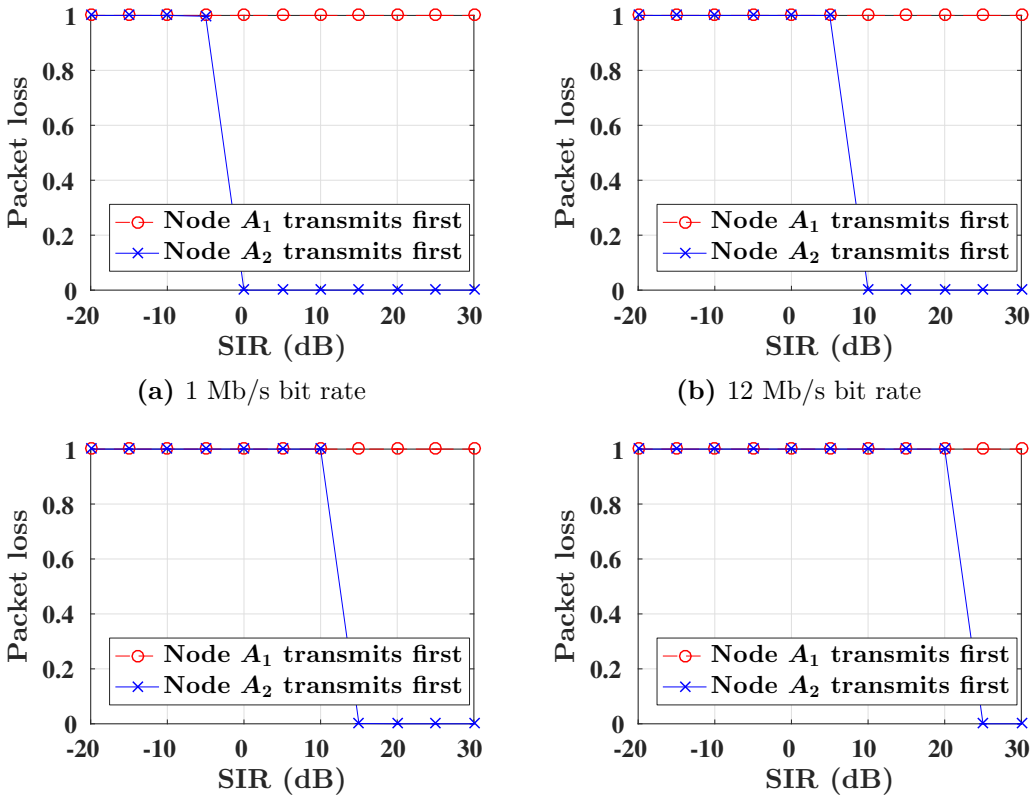

(a) 1 Mb/s bit rate

(b) 12 Mb/s bit rate

(c) 24 Mb/s bit rate

(d) 54 Mb/s bit rate

Figure 4·2: Packet loss probability due to a hidden node in a two-cell network. The performance depends on the order of packet arrivals at the receiver.

high SIR. In Section 4.4.1, we provide further theoretical explanations of the results depicted in Fig. 4·2.

4.3.2 Cascading attack in an office building

We next demonstrate the impact of a cascading attack in an IEEE 802.11g/n network comprising $M = 3$ transmission pairs. The network is deployed in an office floor of a building containing 11 rooms, as shown in Figure 4·3. The physical parameters of the wireless channel are based on the *building model* (ns-3, 2018) of ns-3. The propagation losses between nodes are estimated according to the *hybrid buildings propagation loss model* (ns-3, 2018) of ns-3. The parameter settings of these two models are listed in

Table 4.1.

Each node A_i transmits 1500 byte UDP packets to nodes B_i, where $i \in \{1, 2, 3\}$. The retry limit is set to 7. Nodes A_2 and A_3 generate packets according to Poisson processes with mean rate 0.7 Mb/s. Simulations are run for 600 seconds. Node A_1 (the attacker) starts its transmissions after 200 seconds and ends after 400 seconds. When node A_1 is active, it generates packets at rate 54 Mb/s, which implies that its transmission queue is never empty.

Figure 4·4(a) depicts the average bit rate at node A_3 for different rate adaptation algorithms, namely ARF, Onoe, and AMRR. In all cases, during the first 200 seconds, node A_1 does not transmit. Hence, the bit rate at nodes A_3 climbs from its initial value of 1 Mb/s to 54 Mb/s. We note that the bit rate under ARF and Onoe converges faster to 54 Mb/s than under the more conservative AMRR algorithm. Once node A_1 starts transmitting, though, the bit rate of A_3 drops quickly to low values between 1 and 2 Mb/s for all three algorithms. The reason for the drop are packet retransmissions by node A_2 which cause packet loss at node B_3 due to the receiver capture effect. These packet losses induce the rate adaptation algorithm at node A_3 to lower the bit rate. We stress here that node B_3 is outside the interference range of node A_1. Hence, this result demonstrates the cascading nature of the attack launched by node A_1.

The change of the bit rate at node A_3 impact its utilization. Figure 4·4(b) illustrates this fact. During the first 200 seconds, node A_1 does not transmit. We observe that the utilization of node A_3 decreases as its bit rate increases. This is because the high bit rate shortens the duration of packet transmissions, which in effect means that the channel capacity is higher from the perspective of node A_3. Similarly, when node A_1 is transmitting, the utilization of node A_3 increases reaching values at or above 0.8. This implies that the channel gets congested.

As a consequence of channel congestion, B_3 does not receive some UDP packets

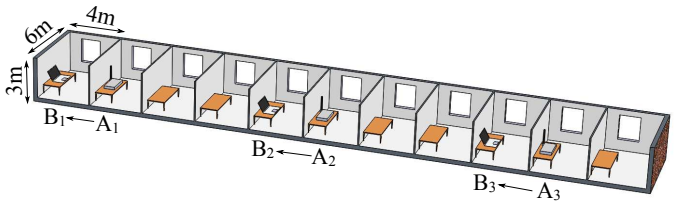

Figure 4·3: Cascading attack in an office building, with three transmission pairs (A_i, B_i), where $i \in \{1, 2, 3\}$. Note that node B_3 is outside the interference range of node A_1 and therefore packet losses at node B_3 are caused by transmissions of node A_2.

Table 4.1: Parameter settings of ns-3 simulation in office building scenario

Building model			
Parameter	Values	Parameter	Values
Building type	Office	External Wall type	Concrete with windows
# of floors	1	Height of floor	3 m
# of rooms at each floor	11	Size of each room	$6 \times 4 \times 3$ m
Hybrid buildings propagation loss model			
Frequency	2.4 GHz	Shadow sigma indoor	8
Internal wall loss	12 dB	Shadow sigma external walls	5

transmitted by node A_3, as shown in Table 4.2. Specifically, when node A_1 transmit, B_3 misses about 3% of the UDP packets for Onoe and AMRR, and about 8% of the packets for ARF. Yet, when node A_1 is not transmitting, all the UDP packets transmitted by node A_3 are received by node B_3. We note that the loss of a UDP packet implies that all transmissions of that packet at the MAC layer fail (i.e., 7 consecutive packet losses when the retry limit is set to 7).

Table 4.2: Fraction of UDP packets not received in the office building scenario

	UDP packets not received by node B_3	
	Node A_1 transmits	Node A_1 does not transmit
ARF	7.67%	0
Onoe	2.51%	0
AMRR	3.14%	0

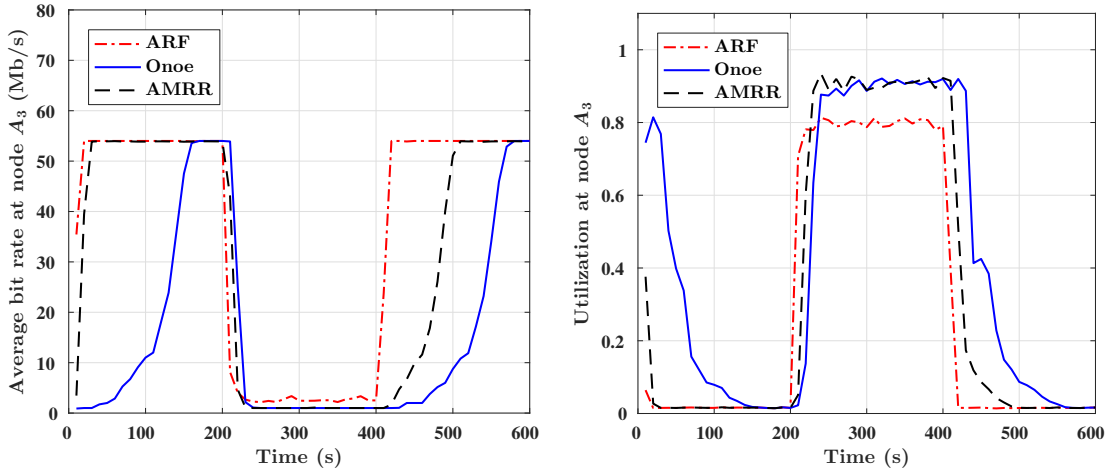

(a) Bit rate **(b)** Utilization

Figure 4·4: Cascading attack in an office building scenario. Node A_1 (attacker) transmits between 200 seconds and 400 seconds. The bit rate of node A_3 drops and its utilization increases significantly during the attack.

4.3.3 Cascading attack in a large network

We now consider a large Wi-Fi network consisting of $M = 20$ cells. Again, we assume each node A_i transmits 1500 byte UDP packets to node B_i, where $i \in \{1, 2, \ldots, 20\}$. Each node uses the ARF rate adaptation algorithm. The retry limit is set 7. To ensure that the packet loss is only due to the receiver capture effect, we set SIR = 30 dB (i.e., the interference caused by each hidden node is very weak). The running time of each simulation is 500 seconds. The packet generation rate of node A_1 (the attacker) varies from 1.08 to 10.8 Mb/s. The packet generation rates of the other nodes range from 0.054 to 2.16 Mb/s. At each node, the UDP packet generation follows a Poisson process.

Figure 4·5(a) depicts the utilization of node A_{20} versus the utilization of node A_{19}, obtained for different packet generation rates of node A_1. The results indicate that the utilizations of node A_{19} and A_{20} are essentially identical, which indicates that the utilization in a large network converges to a limit.

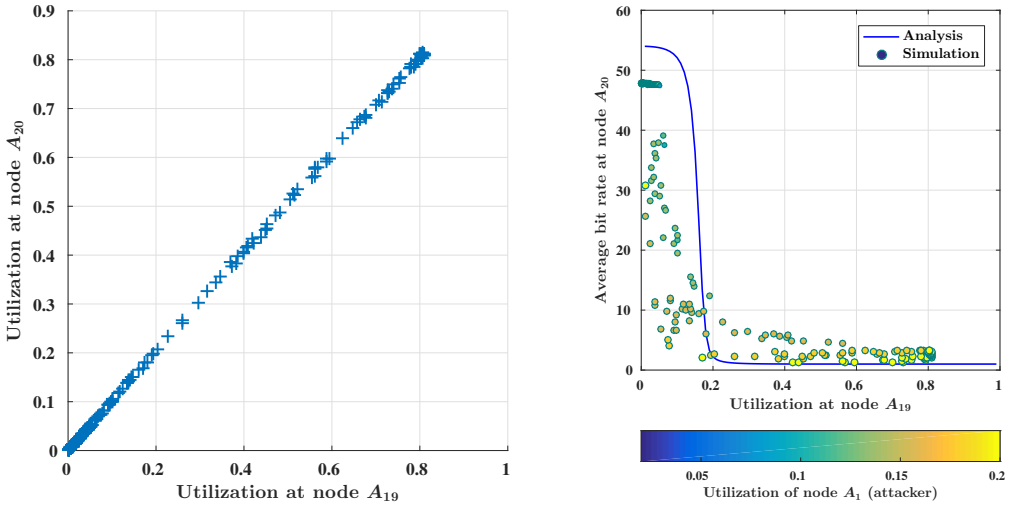

(a) Utilization **(b)** Bit rate

Figure 4·5: Utilization and bit rate in a large network. (a) The utilization of nodes A_{19} and A_{20} is about the same, which implies the existence of a limit. (b) Relationship between bit rate and utilization at the limit. As the utilization of the attacker A_1 increases, the utilization limit jumps and the bit rate limit drops.

Figure 4·5(b) depicts the average bit rate of node A_{20} with respect to the utilization of node A_{19}. When the utilization of node A_{19} is low (say below 0.05), the bit rate of node A_{20} can be as high as 48 Mb/s. Yet, when the utilization of node A_{19} is high (say above 0.4), the bit rate of node A_{20} drops to the base rate of 1 Mb/s. This trend matches the behavior predicted by the analysis, which we discuss in detail in Section 4.4.2.

Figure 4·5(b) also shows the impact of node A_1 on the other nodes. When the utilization of node A_1 is below 0.1 (blue and green dots), the average bit rate of node A_{20} is near 48 Mb/s. However, when the utilization of node A_1 exceeds 0.15 (orange and yellow dots), the average bit rate of node A_{20} drops sharply. This results implies that node A_1 can launch a cascading attack by increasing the rate of packet transmissions over its channel.

4.4 Analysis

In this section, we propose analytical models to shed light into the simulation results shown in Section 4.3. We first analyze the packet loss probability for the two cases of Section 4.3.1. Next we consider cascading attacks in a large network, based on the model of Section 4.3.3. We first estimate the limit of the packet loss due to the receiver capture effect as the number of cells $M \to \infty$. We then model the limiting behavior of the ARF rate adaptation algorithm.

4.4.1 Hidden node

We analyze the packet loss probability in the two cases presented in Section 4.3.1. We show that a packet loss in the first case is caused by a CRC error, which depends on the SIR and the bit rate. On the hand, the packet loss in the second case is caused by the receiver capture effect, which does not depend on the SIR and the bit rate.

Case (1): A_2 transmits before A_1

When A_2 transmits before A_1, packet loss at B_2 is due to a CRC error. Specifically, a CRC error occurs if one or more bits of the received packet are wrong.

The probability of a CRC error can be estimated as follows. Suppose a packet has n bits. Denote by β_j the error rate (probability) of bit j in the packet and assume that it only depends on the signal to interference ratio during the transmission of that bit, which is denoted SIR_j. The probability of a packet loss, denoted by L, is then

$$L = 1 - \prod_{j=1}^{n}(1 - \beta_j), \tag{4.1}$$

In IEEE 802.11, each bit rate has different resilience to the interference due to the adoption of different signal modulations and coding schemes. Thus, the bit error model at each bit rate is different. Generally, a lower bit rate has better resilience

to interference. For example, as shown in Figure 4·2, a bit rate of 1 Mb/s requires a SIR of 0 dB or higher to ensure successful reception (capture) of a packet with high probability. At a bit rate of 54 Mb/s, on the other hand, a packet is successfully received if the SIR is above 25 dB.

Bit errors models for different modulation and coding schemes in Wi-Fi can be found in (Pei and Henderson, 2009; Pei and Henderson, 2010). We illustrate here the computation for the 1 Mb/s bit rate. Denote the channel bandwidth by B and the channel bit rate by f_b. In that case, the error rate of bit j is

$$\beta_j = \frac{1}{2} e^{-10^{\mathrm{SIR}_j/10} * B/f_b}. \tag{4.2}$$

We now explain the result for the packet loss probability at 1 Mb/s associated, shown in Figure 4·2(a). At 1 Mb/s bit rate, $B = 22$ MHz, and $f_b = 1$ Mb/s. According to (4.1) and (4.2), when $\mathrm{SIR}_j = 0$ dB for all the bits in a packet, the loss probability of a 1500-byte packet is 1.67×10^{-6}. That is, the transmission is virtually always successful when the power of interference is equal to the power of the signal. However, when $\mathrm{SIR}_j = -5$ dB, the packet loss probability of a 1500-byte packet is close to 1. In summary, for case (1), a packet sent by node A_2 can be captured if the SIR is high enough. The SIR thresholds for a successful packet capture depends on the bit rate.

Case (2): A_1 transmits before A_2

We next explain why in the second case considered in Section 4.3.1, packet losses always occur irrespectively of the SIR or bit rate. Since node A_1 transmits a packet before node A_2, node B_2 detects the PLCP preamble and header of the packet sent by node A_1 and transits into the Rx state. Therefore, it cannot detect the PLCP preamble and header of the packet sent by node A_2. Therefore, a packet loss occurs

regardless of the SIR. Moreover, in the IEEE 802.11 protocol, the PLCP preamble and header are always transmitted at the lowest supported rate of the band (e.g., 1 Mb/s for 2.4 GHz). Hence, the packet loss caused by the receiver capture effect does not relate to the bit rates of the packets sent by nodes A_1 and A_2.

4.4.2 Asymptotic analysis of cascading attacks for large Wi-Fi networks

In this section, we conduct an asymptotic analysis for a large Wi-Fi network model similar to that introduced in Section 4.3.3, where each node implements the ARF algorithm. This analysis sheds light into the relationships between the limits of the packet loss probability, utilization, and average bit rate for each transmission pair $A_i - B_i$, when i is very large.

Our analysis is based on a number of simplifying assumptions required to keep the analysis tractable. These assumptions are:

- The limit of the utilization (or shortly, the *utilization limit*) exists and is unique (this fact was empirically verified with Fig. 4·5(a)). We denote the utilization limit by ω.
- MAC layer packet transmissions and retransmissions at each node are independent and form a Poisson process.
- If a hidden node transmits first, the receiver can always decode the PLCP preamble and header of the packet sent by that node.
- Packet loss is only due to the receiver capture effect (i.e., the SIR is above 25 dB).
- MAC overhead timing parameters (such as SIFS, DIFS, etc.) are ignored.

Beside these assumptions, the rest of the analysis presented in this section is exact.

In particular, it captures peculiarities of the ARF protocol and interactions between neighboring Wi-Fi cells.

Estimation of the limit of the packet loss probability

We first estimate the limit of the packet loss caused by receiver capture in a large network. We denote this limit $L(\omega)$.

Proposition 2. *Let the utilization limit be ω. Then, the limit of the packet loss probability is*

$$L(\omega) = \omega(1 - \frac{\omega}{2}). \tag{4.3}$$

Proof. Assume i is very large, such that the utilization of A_i and A_{i+1} is ω. Remember that A_i is a hidden node with respect to A_{i+1}. We define three states of the channel which are idle, captured by node A_i, and captured by node A_{i+1}. The packet loss occurs if node A_{i+1} starts to transmit when the channel is captured by node A_i. When there is no packet transmissions, the channel is idle with probability $(1-\omega)^2$. From our assumption, the PLCP preambles and headers of node A_i can always be decoded if the channel is not captured first by node A_{i+1}. Hence, by symmetry, nodes A_i and A_{i+1} have the same probability to capture the channel. The probability that the channel is captured by node A_i (or node A_{i+1}) is $(1-(1-\omega)^2)/2 = \omega - \omega^2/2$. Due to the Poisson Arrivals See Time Averages (PASTA) property, the probability that a node transmits at any random point of time is the same. Without loss of generality, assume node A_{i+1} starts transmitting a packet at time t'. The probability that the channel is captured by node A_i at time t' is $\omega - \omega^2/2$. □

We note that when the channel is fully utilized (i.e., $\omega = 1$), the probability of a packet loss due to receiver capture is $1/2$. Figure 4·6 depicts the packet loss probability at different utilizations using our analysis and simulations. The simulation results match the analysis results well until the node utilization reaches its upper limit. Due to MAC overhead (which we ignore in the analysis), at a bit rate of 1 Mb/s, the upper limit of the node utilization is about 0.9. This value is much smaller when nodes communicate at 54 Mb/s bit rate. Nevertheless, even though it is based on a number of simplifications, our analysis captures the correct trend.

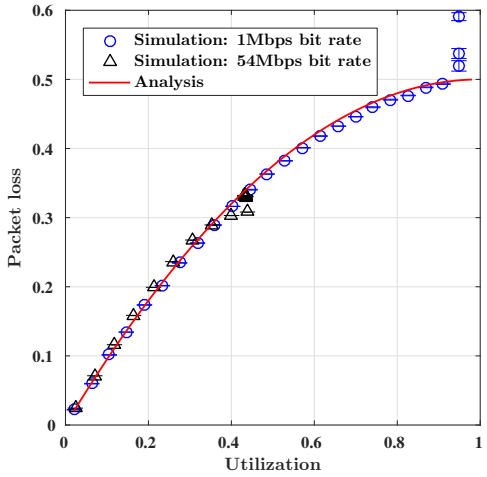

Figure 4·6: Packet loss probability vs. utilization under receiver capture: asymptotic analysis (Proposition 1) versus simulations.

Limiting behavior of ARF

We next investigate the limiting behavior of ARF under a cascading attack. Given the utilization limit ω and the limit of the packet loss probability $L(\omega)$, our goal is to characterize the limit of the average bit rate $\overline{b}(\omega)$. In the sequel, we derive the proportion of packets transmitted at each bit rate b_j of ARF, which we denote by $\alpha_j(\omega)$. Given $\alpha_j(\omega)$, we can readily compute $\overline{b}(\omega)$:

Theorem 11. *Given the utilization limit ω, the limit of the average bit rate is*

$$\overline{b}(\omega) = \Big(\sum_{j=1}^{N} \frac{\alpha_j(\omega)}{b_j}\Big)^{-1}. \tag{4.4}$$

Proof. Suppose that a node has large amount of bits X to transmit. The amount of data that is transmitted at bit rate b_j is $X\alpha_j(\omega)$. Therefore, the average time spent transmitting at bit rate b_j is $X\alpha_j(\omega)/b_j$ and the average bit rate is $\frac{X}{\sum_{j=1}^{N} \frac{X\alpha_j(\omega)}{b_j}} = \Big(\sum_{j=1}^{N} \frac{\alpha_j(\omega)}{b_j}\Big)^{-1}$. □

Figure 4·5(b), introduced earlier, compares the theoretical result of Theorem 11 and simulations. As one can see from the figure, the theorem captures well the relationship between the bit rate versus the utilization, and in particular the sharp

drop of the bit rate of a transmission pair as the channel utilization of its neighboring pair increases. This phenomenon is a consequence of the receiver capture effect.

We next provide details on how to obtain an analytical expression for α_j. We analyze the behavior of ARF using the theory of semi-Markov processes (Medhi, 2012). The analytical model is based on the work in (Singh and Starobinski, 2007). However, (Singh and Starobinski, 2007) only considers the case when each node always has packets available to transmit, i.e., each node is congested. Unlike the analysis in (Singh and Starobinski, 2007), our analysis also covers the case when nodes are uncongested. We calculate the fraction of packets that are transmitted at each bit rate.

Let denote by j the state during which ARF operates at bit rate b_j, that is, when entering state j ARF transmits at bit rate b_j until it exits that state. The number of packets transmitted in that state and the transition probability to the next possible states (i.e., $j-1$ or $j+1$) only depend on the current state j (i.e., these quantities are independent of the previous states). Thus, the behavior of ARF rate can be modeled with a semi-Markov process, whose embedded Markov chain is shown in Figure 4·7. We denote by $\overline{X}_j(\omega)$ the average number of packets transmitted during each visit to state j and by $\pi_j(\omega)$ the steady-state proportion of transitions into state j. Since the process is irreducible and ergodic, according to Theorem 7.2 in (Medhi, 2012), we obtain that the fraction of packets transmitted by node A_{i+1} at bit rate b_j is

$$\alpha_j(\omega) = \frac{\pi_j(\omega)\overline{X}_j(\omega)}{\sum_{n=1}^{N} \pi_n(\omega)\overline{X}_n(\omega)}. \tag{4.5}$$

Based on Equation (4.5), we next provide expressions for $\pi_j(\omega)$ and $\overline{X}_j(\omega)$. We first consider $\pi_j(\omega)$. The embedded Markov chain of the transitions of the bit rates is depicted in Figure 4·7. We denote $T_j^+(\omega)$ and $T_j^-(\omega)$ to be the transition probabilities from state j to states $j+1$ and $j-1$, respectively.

Proposition 3. *The steady-state proportion of transitions into state j is*

$$\pi_j(\omega) = \begin{cases} \frac{1}{1+\sum_{b=2}^{N}\prod_{n=1}^{b-1}\frac{T_n^+(\omega)}{T_{n+1}^-(\omega)}} & \text{if } j = 1, \\ \pi_1(\omega)\prod_{n=1}^{b-1}\frac{T_n^+(\omega)}{T_{n+1}^-(\omega)} & \text{if } j > 1. \end{cases} \tag{4.6}$$

We next derive the transition probabilities $T_j^+(\omega)$ and $T_j^-(\omega)$. The operation of ARF at bit rate b_j $(1 < j < N)$ is shown in Figure 4·8. States S_k^j and S_{-k}^j $(k > 0)$ represent the states where k consecutive packets are transmitted successfully and unsuccessfully at bit rate b_j, respectively. The initial state is S_0^j. Thus, the transition probability $T_j^+(\omega)$ is the probability that that the process reaches state S_s^j before state S_{-f}^j, starting from state S_0^j.

Proposition 4. *Given $L(\omega)$, the transition probability from state j to state $j+1$ is*

$$T_j^+(\omega) = \begin{cases} 1 & \text{if } j = 1, \\ \frac{(1-L(\omega))^s\sum_{k=0}^{f-1}L(\omega)^k}{1-\sum_{k=1}^{s-1}(1-L(\omega))^k\sum_{k=1}^{f-1}L(\omega)^k} & \text{if } 1 < j < N, \end{cases} \tag{4.7}$$

and the transition probability from state j to state $j-1$ is

$$T_j^-(\omega) = \begin{cases} 1 - T_j^+(\omega) & \text{if } 1 < j < N, \\ 1 & \text{if } j = N. \end{cases} \tag{4.8}$$

Based on Propositions 2, 3 and 4, the steady-state proportion of transitions into state j, namely $\pi_j(\omega)$, can be expressed as a function of the utilization limit ω. The following proposition states that the average number of consecutive packet attempts transmitted at bit rate b_j, namely $\overline{X}_j(\omega)$, can also be expressed as a function of ω.

Proposition 5. *Given the packet loss $L(\omega)$, the average number of consecutive packets transmitted at bit rate b_j is*

$$\overline{X}_j(\omega) = \begin{cases} \frac{\sum_{k=0}^{s-1}(1-L(\omega))^k}{(1-L(\omega))^s} & \text{if } j = 1, \\ \frac{\sum_{k=0}^{s-1}(1-L(\omega))^k\sum_{k=0}^{f-1}L(\omega)^k}{1-\sum_{k=1}^{s-1}(1-L(\omega))^k\sum_{k=1}^{f-1}L(\omega)^k} & \text{if } 1 < j < N, \\ \frac{\sum_{k=0}^{f-1}L(\omega)^k}{L(\omega)^f} & \text{if } j = N. \end{cases} \tag{4.9}$$

Plugging the results of Propositions 2, 3, 4, and 5 into Equation (4.5), we finally

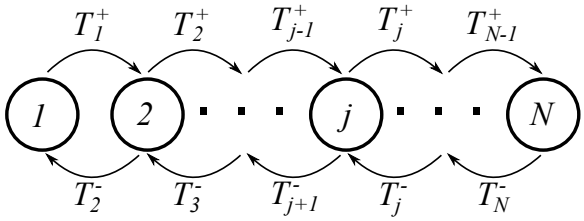

Figure 4·7: Embedded Markov chain of the semi-Markov process which represents the transitions of the bit rates operated by ARF. State j corresponds to the state where ARF operates at bit rate b_j.

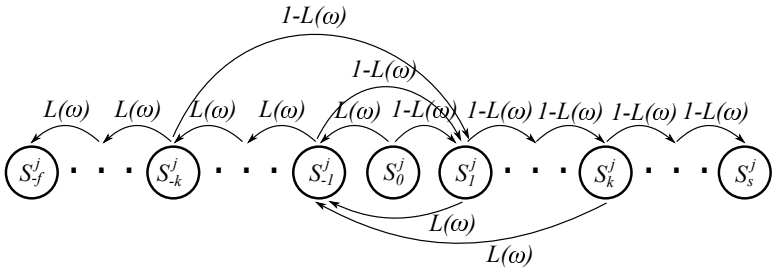

Figure 4·8: Markov model of ARF at an intermediate bit rate b_j $(1 < j < N)$. States S_k^j and S_{-k}^j represent k consecutive successful and failed transmission attempts, respectively.

obtain an expression for the fraction of packets transmitted at each bit rate b_j, that is α_j, as a function of the utilization limit ω and the parameters s, f, and N of the ARF algorithm.

4.5 Summary

In this chapter, we demonstrate the feasibility of launching cascading attacks in Wi-Fi networks with weakly interfering hidden nodes. We observe that weak interference due to hidden nodes can cause significant packet losses due to the receiver capture effect. We validate and contrast the impact of the classical capture effect and that of the receiver capture effect through both analysis and simulation. We show that while the capture effect is sensitive to the SIR and bit rate, the receiver capture effect is insensitive to either of those. Our analysis shows that the receiver capture effect can yield a packet loss rate as high as 50% under high utilization.

Packet losses caused by receiver capture force rate adaptation algorithms to lower their bit rate and induce traffic congestion. Through ns-3 simulations for various rate adaptation algorithms and based on realistic channel models, we show that this coupling between neighboring cells propagates across a network. In fact, a small increase in the traffic generated by the attacker can reduce the bit rate and create congestion across a whole network.

Through simulation and analysis, we also investigate the behavior of a large network implementing the ARF algorithm. We find that the operational bit rate of ARF drops sharply as the packet loss rate increases. In fact, the bit rate is brought down to 1 Mb/s when the packet loss rate is above 40%, which is lower than the highest possible packet loss rate due to the receiver capture effect. We find that if the utilization of nodes in the networks is higher than 20%, the average bit rate is reduced to close to 1 Mb/s. We further show that the utilization and bit rate converge to limits in a large network, and establish a relationship between these two quantities that match well the simulation results.

Chapter 5

Mitigation of Cascading Attacks

5.1 Motivation

In previous chapters, we demonstrated cascading DoS attacks on Wi-Fi networks. This attack exploits an *interference coupling* phenomenon (Xin et al., 2016) between neighboring cells of IEEE 802.11 networks, which is induced by hidden nodes. Using interference coupling, an attacker can locally raise the amount of traffic that it generates and affect its neighboring cells, which in turn affect their own neighboring cells and so on. As a result, the transmitting queue of a distant node can suddenly be brought into instability and get congested. The attack is feasible in both infrastructure and ad-hoc networks, under certain configurations. Moreover, since the attack can be launched remotely and is protocol compliant, it makes it difficult to locate and identify the attacker. Given the serious consequences of cascading DoS attacks, it is important to find methods to mitigate them.

In this chapter, we focus on the mitigation of cascading DoS attacks in Wi-Fi networks. Our key idea is to optimize the durations of packet transmissions (or, equivalently, the packet length divided by the bit rate) in order to ensure that interference coupling does not propagate and amplify. To achieve this goal, we show that it is essential to properly model the impact of MAC overhead, and in particular MAC timing parameters. We propose a new theoretical model where we relate the utilization of nodes in neighboring cells using iterative equations. We then perform a fixed point analysis to characterize the limiting behavior of the sequence of node

utilizations and the feasibility of launching a cascading DoS attack against a Wi-Fi network.

Our main contributions are as follows. We first show how to set the packet duration in order to avoid a cascading DoS attack, namely to prevent the initial value of the sequence of node utilizations (which can be set by the attacker) to affect the limit of the sequence. Second, we show that it is possible to simultaneously optimize the packet duration in order to achieve maximum throughput. Third, we validate the analytical results using ns-3 simulations, including for an office building model with and without the cross traffic. Fourth, we setup an experimental testbed with real Wi-Fi cards to demonstrate the mitigation of the attack. A key insight obtained from our analysis, simulation, and experiment is that IEEE 802.11 networks with relatively large MAC overhead (e.g., IEEE 802.11b) are less susceptible to cascading DoS attacks than networks with smaller overhead (e.g., IEEE 802.11g and IEEE 802.11n). We also show that our method achieves higher throughput performance than the RTS/CTS method, especially at high bit rates.

The rest of this chapter is organized as follows. In Section 5.2, we explain how cascading DoS attacks operate and the impact of the packet length on the feasibility of launching such attacks. In Section 5.3, we introduce our analytical model, derive a sufficient condition for preventing cascading DoS attacks, and show how to optimally set packet durations in order to maximize throughput performance. We present our simulation results in Section 5.4 and investigate the performance of the mitigation in a real Wi-Fi network in Section 5.5. We conclude in Section 5.6.

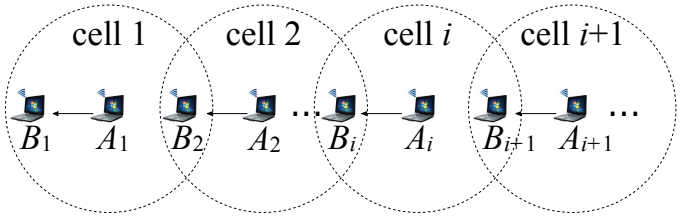

Figure 5·1: Network configuration. The dotted circles represent the communication range of nodes A_i. Nodes A_i transmit packets to nodes B_i $(i = 1, 2, \ldots)$. Each transmission pair (A_i, B_i) belongs to a different cell. Nodes A_i are hidden nodes with respect to nodes A_{i+1}.

5.2 Cascading DoS Attacks

5.2.1 Attack scenario

We next explain how a cascading DoS attack can unfold. We consider a network configuration consisting of a chain of N pairs of nodes. Figure 5·1 depicts the configuration. The ith pair is denoted (A_i, B_i), where $i \geq 1$. Each node A_i transmits packets to node B_i (one-hop communication). Furthermore, each node A_i is a *hidden node* with respect to node A_{i+1}, which means that node A_i cannot sense a transmission by node A_{i+1}. If a transmission by node A_i overlaps with a transmission by node A_{i+1}, a packet collision occurs at node B_{i+1}. This collision forces node A_{i+1} to retransmit its packet using the procedure described in Section 2.1.1.

In this configuration, suppose node A_1 (the attacker) starts increasing the rate at which it generates packets and transmits them over the channel (in compliance with the IEEE 802.11 standard). These transmissions will cause collisions at node B_2, which forces node A_2 to increase the rate at which it attempts to transmit packets over the channel (due to retransmissions). The increased rate of transmission attempts by A_2 will in turn impact pair (A_3, B_3) and so forth. Under certain conditions, this effect may amplify along the chain and cause a large fraction of transmission attempts to fail and result in unstable queues (i.e., the rate at which nodes can successfully transmit packets over the channel is lower than the rate at which packets are generated).

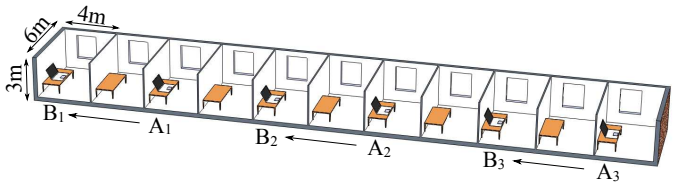

Figure 5·2: Example of an attack in an office building. Three transmission pairs (A_i, B_i), where $i \in \{1, 2, 3\}$, are positioned as shown in the figure.

5.2.2 Example

To help motivate the rest of this chapter, we next present an example to illustrate the occurrence of a cascading DoS attack in a practical scenario, as well as a way to prevent it. Define ρ_i as the *offered load* at node i, that is, the rate at which it generates packets multiplied by the transmission duration of each packet. Further, define the *utilization* of node A_i as the average fraction of time during which node A_i is transmitting, and the *throughput* of node A_i as the average number of bits per second that node A_i successfully transmits to node B_i.

As shown in Figure 5·2, we consider communication within an office building using the ns-3 building model (ns-3, 2018). The external wall of the building is made of concrete with windows. The internal wall loss is 12 dB (Afaqui et al., 2015). All the other parameters are set to default. In the following two examples, we consider an IEEE 802.11g/n network composed of $N = 3$ pairs of nodes and communicating using UDP (examples of realistic applications using UDP include Google Chromecast and Apple TV).

The nodes are located in every other room, as shown in Figure 5·2. Each transmitting node uses a short slot time (i.e., $T_{\text{slot}} = 9$ μs) and a bit rate of 6 Mb/s. The offered load at nodes A_2 and A_3 is set to 0.14 while the attacker A_1 varies its load ρ_1. We run simulations of this configuration using the ns-3 simulator (ns-3, 2018). The running time of each simulation is 200 seconds and the plotted results are averages

computed over three independent runs.

In the first example, we set the packet length to 1500 bytes. Simulation results illustrating the cascading attack are depicted in Figure 5·3. We observe that as node A_1 starts to transmit after 50 s, the utilization of node A_3 suddenly jumps from about 0.25 to 0.65 due to packet collisions and retransmissions. As a result, its throughput drops from about 0.75 Mb/s to 0.5 Mb/s. The utilization and throughput of node A_3 recovers once node A_1 stops transmitting after 150 s.

Now consider the same setting, but with packets of length 200 bytes. The offered load of nodes A_2 and A_3 is maintained the same as in the previous example (by increasing the packet generation rate). In that case, we observe that increased traffic generation by node A_1 has no effect on the utilization and throughput of node A_3. This result holds no matter what packet length is used by the attacker.

Chapter 3 only considers the impact of the traffic load and the retry limit on the feasibility of a cascading attack. Figure 5·3 clearly shows that this is insufficient and that other parameters (e.g., the packet length) need to be taken into account. In the next section, we present and analyze a model that incorporates these other parameters.

5.3 Mitigation of Cascading Attacks: Model and Analysis

We propose an analytical model to find out how to mitigate a cascading DoS attack against an IEEE 802.11 network. The proposed model captures key system parameters, including the offered load, the packet duration (i.e., the packet length divided by the bit rate), and MAC parameters. We consider the network configuration shown in Fig. 5·1, since it is a configuration for which it is known that cascading attacks are feasible. The analysis captures the coupling between the utilizations of neighboring pairs of nodes in the chain through a sequence of iterative equations. We conduct a

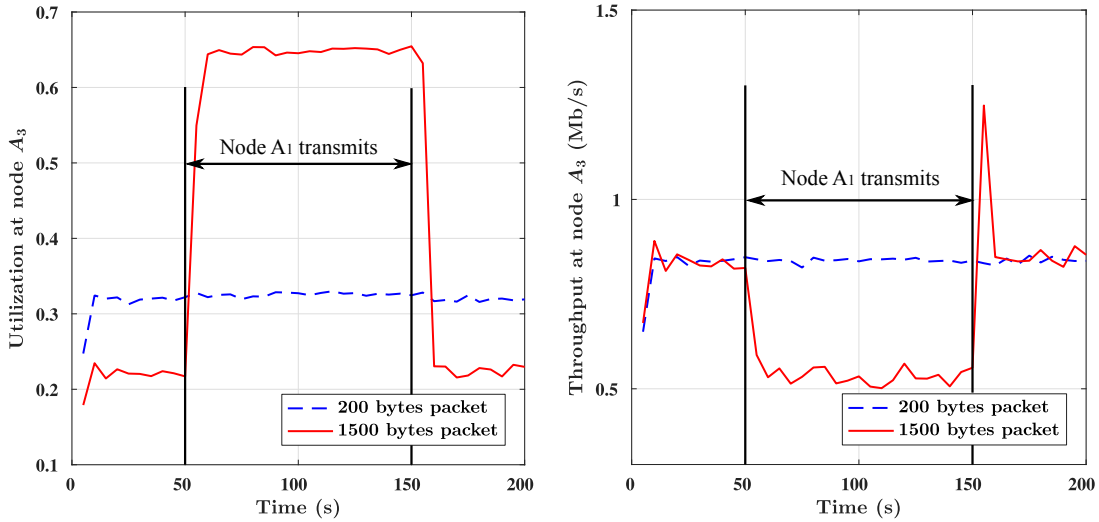

(a) Utilization. **(b)** Throughput.

Figure 5·3: Feasibility of cascading DoS attacks in IEEE 802.11g/n networks of an office building. When nodes in the network use 1500 bytes packets, node A_1 can launch a cascading DoS attack. When node A_1 is transmitting, node A_3 suffers from low throughput and high channel utilization. However, this attack is prevented when nodes use 200 bytes packets.

fixed point analysis to determine the limit of the sequence, as a function of the initial condition (i.e., the utilization of the first node in the chain, which is the attacker). Our goal is to determine when the initial value of the sequence of utilization is guaranteed to have no influence on the limit of the sequence (that is, the utilization of remote nodes) for all possible traffic loads.

5.3.1 Model and assumptions

We now present our model, notation, and assumptions. We denote by λ_1 the packet generation rate at node A_1 (the attacker) and by $\lambda_i = \lambda$ the packet generation rate at all the other nodes A_i ($i \geq 2$). The duration of a packet transmission is T (we assume a fixed bit rate). The offered load at node A_1 is $\rho_1 = \lambda_1 T$ and the offered load at all the other nodes is $\rho = \lambda T$. The average number of transmissions for each

packet at node A_i (i.e., the average retry count) is denoted $\overline{r}_i$. Note that $\overline{r}_1 = 1$. The probability that a packet transmitted by node A_i collides is denoted p_i. Finally, we denote by μ_i the service capacity of the channel, that is the maximum average rate at which packets (both new and retransmissions) can be transmitted over the channel. In the sequel, we derive expressions for $\overline{r}_i$, p_i and μ_i.

The utilization of node A_i (i.e., the fraction of time during which it transmits) is denoted u_i. If $\overline{r}_i\lambda_i < \mu_i$, then the queue of node A_i is stable and by Little's Law (Kleinrock, 1975) its utilization is $\overline{r}_i\lambda_i T$. On the other hand, If $\overline{r}_i\lambda_i > \mu_i$, then the queue of node A_i is unstable and its utilization is $\mu_i T$. We refer to $\mu_i T$ as the *congested utilization*. Hence, the utilization of node A_i ($i \geq 1$) is

$$u_i = \min\{\overline{r}_i\lambda_i T, \mu_i T\}. \tag{5.1}$$

In order to render the analysis of this queueing network tractable, we make use of Kleinrock's random look assumption (Kleinrock et al., 1975), namely:

1. The probability p_i that a packet transmitted by node A_i collides is independent of previous attempts.

2. Packet transmissions and retransmissions at each node A_i form a Poisson process with rate $\min\{\overline{r}_i\lambda_i, \mu_i\}$.

We emphasize that beside these approximations, the rest of the analysis is exact. Note that a key difference between the analysis conducted in this chapter and Chapter 3 is that we develop a method to characterize the congested utilization (see Lemma 14). Because the congested utilization is smaller than 1, the structure of the iterative sequence (see Eq. (5.10)) and the analysis of its limits (see Sections 5.3.3 to 5.3.5) are markedly different from the results derived in Chapter 3.

5.3.2 Iterative analysis

In this section, we derive iterative equations for relating the utilizations of neighboring pairs of nodes. The following lemma provides expressions for p_i and $\overline{r}_i$. The proof follows similar lines as the derivations of Equations (3.8) and (3.6) in Section 3.4.2.

Lemma 12. *For $i \geq 2$,*

1. $p_i = 1 - \mathrm{e}^{-u_{i-1}}(1 - u_{i-1})$. *(5.2)*
2. $\overline{r}_i = \sum_{r=1}^{R} p_i^{r-1}$. *(5.3)*

Using the above lemma, one can obtain an expression for the average utilization of a node with a stable queue.

Lemma 13. *Let $i \geq 2$ and suppose that the queue of node A_i is stable. Then its utilization is*

$$\overline{r}_i \lambda T = \rho \sum_{r=1}^{R} \left(1 - \mathrm{e}^{-u_{i-1}}(1 - u_{i-1})\right)^{r-1}. \tag{5.4}$$

We next provide an expression for the congested utilization of a node with an unstable queue.

Lemma 14. *Let $i \geq 2$ and suppose that the queue of node A_i is unstable. Then its congested utilization is*

$$\mu_i T = \frac{\sum_{r=1}^{R} p_i^{r-1} T}{\sum_{r=1}^{R} p_i^{r-1}(d_r^{(s)}(1 - p_i) + d_r^{(f)} p_i + T)},$$

where $d_r^{(s)}$, $d_r^{(f)}$ and p_i are given by Equations (2.3), (2.4) and (5.2) respectively.

Proof. Define the backoff cycle of a packet as the time it takes for that packet to be successfully transmitted during a back-off procedure or dropped after R failed retransmissions. We note that the lengths of backoff cycles of different packets are independent, due to Assumption 1 and the fact that the contention window is reset at the beginning of each cycle. Hence, the backoff process of consecutive packets forms a regenerative process (Ross, 1996), which implies that the average utilization of node A_i is the ratio of the average time during which node A_i transmits during a backoff cycle to the average length of a backoff cycle.

Now, the fact that node A_i retransmits a packet for the rth time implies that all the previous $r-1$ retransmissions failed due to packet collisions caused by a hidden node. Hence, the probability that node A_i transmits a packet at least r times is p_i^{r-1} and the average time that node A_i spends transmitting during a backoff cycle is

$$\sum_{r=1}^{R} p_i^{r-1} T. \tag{5.5}$$

The average time that node A_i spends on the rth retransmission is $d_r^{(s)}(1-p_i) + d_r^{(f)} p_i + T$. Hence, the average length of a backoff cycle is

$$\sum_{r=1}^{R} p_i^{r-1} (d_r^{(s)}(1-p_i) + d_r^{(f)} p_i + T). \tag{5.6}$$

Taking the ratio of Eq. (5.5) to Eq. (5.6) gives the result stated by the lemma. □

To simply notation in the rest of the analysis, we define the following functions based on Lemmas 12, 13 and 14:

$$P(u_{i-1}) \triangleq p_i = 1 - \mathrm{e}^{-u_{i-1}}(1-u_{i-1}); \tag{5.7}$$

$$U(u_{i-1}) \triangleq \overline{r}_i \lambda T = \rho \sum_{r=1}^{R} \left(1 - \mathrm{e}^{-u_{i-1}}(1-u_{i-1})\right)^{r-1}; \tag{5.8}$$

$$\begin{aligned} S(u_{i-1}) &\triangleq \mu_i T \\ &= \frac{\sum_{r=1}^{R} (p_i)^{r-1} T}{\sum_{r=1}^{R} ((p_i)^{r-1} (d_r^{(s)}(1-p_i) + d_r^{(f)} p_i + T))}. \end{aligned} \tag{5.9}$$

Substituting (5.8) and (5.9) into (5.1), we obtain the following relationship between the utilizations of nodes A_i and A_{i-1}:

$$u_i = \min\left\{U(u_{i-1}), S(u_{i-1})\right\}. \tag{5.10}$$

5.3.3 Limiting behavior and fixed points

We next characterize the limiting behavior of the sequence of utilizations, using the concept of fixed points. We then formalize the notion of a cascading DoS attack, and obtain a sufficient condition for preventing it.

Consider the possible limits of the utilization sequence $\{u_i\}_{i=1}^{\infty}$. These limits represent *fixed points* of the iteration (5.10).

Definition 4 (Fixed point)**.** *We say that $\omega \in [0,1]$ is a fixed point of (5.10) if*

$$\omega = \min\{U(\omega), S(\omega)\}. \tag{5.11}$$

We next define the two possible types of fixed points.

Definition 5 (Congested and uncongested fixed points)**.** *Let*

$$\check{\omega} = U(\check{\omega}). \tag{5.12}$$

If $\check{\omega}$ also satisfies (5.11), we say that $\check{\omega}$ is an uncongested fixed point. Likewise, let

$$\hat{\omega} = S(\hat{\omega}). \tag{5.13}$$

If $\hat{\omega}$ also satisfies (5.11), then we say that $\hat{\omega}$ is a congested fixed point.

Based on the property of a fixed point (i.e., congested or uncongested), we define next whether a network is congested or not.

Definition 6 (Network congestion)**.** *A network is said to be uncongested if the limit of the utilization sequence $\{u_i\}_{i=1}^{\infty}$ is an uncongested fixed point $\check{\omega}$. Otherwise, if the limit of the utilization sequence $\{u_i\}_{i=1}^{\infty}$ is a congested fixed point $\hat{\omega}$, then the network is said to be congested.*

Using the above notions, we now formally define a cascading DoS attack.

Definition 7 (Cascading DoS attack)**.** *A cascading DoS attack occurs when changing u_1 causes the network to change its state from uncongested to congested.*

We conclude that an attack is feasible only if the utilization sequence has both uncongested and congested fixed points. If for each possible value of the offered traffic load ρ, (5.11) has only one type of fixed points, then a cascading DoS attack can never be launched on the network (assuming that all the other network parameters remain fixed).

In the following, we show that the value of $\hat{\omega}$ plays a key role in determining the feasibility of launching a cascading DoS attack. Specifically, we show that if $\hat{\omega} \leq (3-\sqrt{5})/2$, then (5.11) has only one type of fixed points for each traffic load ρ and a cascading DoS attack is unfeasible. In Section 5.3.6, we further show that if $\hat{\omega} = (3-\sqrt{5})/2$, then the network achieves the highest possible congestion throughput.

5.3.4 Existence of fixed points

We now investigate the existence of the two types of fixed points (uncongested and congested) in Equation (5.11). We first show that if a congested fixed point exists, then it is unique.

Before prove that, we provide the following lemmas.

Lemma 15. *If $b \geq a$, then the function $f(x) = \frac{a+xb}{1+x}$ is monotonically increasing in x.*

Proof. Let $b > a$. The derivative of function $f(x)$ is

$$f'(x) = \frac{b(1+x)-(a+xb)}{(1+x)^2} = \frac{b-a}{(1+x)^2} \geq 0.$$

□

Lemma 16. *Consider an arbitrary sequence $\{a_r\}_{r=1}^{R}$ such that $a_{r+1} \geq a_r$, if $p' > p$, then $\frac{\sum_{r=1}^{R}(p')^{r-1}a_r}{\sum_{r=1}^{R}(p')^{r-1}} \geq \frac{\sum_{r=1}^{R}(p)^{r-1}a_r}{\sum_{r=1}^{R}(p)^{r-1}}$.*

Proof. We use recursive method to prove this lemma. When $R = 2$, we have $\frac{a_1+pa_2}{1+p}$. Its derivative is

$$(\frac{a_1+pa_2}{1+p})' = \frac{a_2(1+p)-(a_1+pa_2)}{(1+p)^2} = \frac{a_2-a_1}{(1+p)^2} \geq 0.$$

Thus, when $R = 2$, the lemma is correct. We assume the lemma is correct at R. That is,

$$\frac{\sum_{r=1}^{R}(p')^{r-1}a_r}{\sum_{r=1}^{R}(p')^{r-1}} \geq \frac{\sum_{r=1}^{R}(p)^{r-1}a_r}{\sum_{r=1}^{R}(p)^{r-1}}.$$

We next to prove whether the lemma is also correct at $R+1$. If so, we can prove the lemma. At $R+1$, we have

$$\begin{aligned}
\frac{\sum_{r=1}^{R+1}(p')^{r-1}a_r}{\sum_{r=1}^{R+1}(p')^{r-1}} &= \frac{\sum_{r=1}^{R}(p')^{r-1}a_r + (p')^R a_{R+1}}{\sum_{r=1}^{R}(p')^{r-1} + (p')^R} \\
&= \frac{\frac{\sum_{r=1}^{R}(p')^{r-1}a_r}{\sum_{r=1}^{R}(p')^{r-1}} + \frac{(p')^R a_{R+1}}{\sum_{r=1}^{R}(p')^{r-1}}}{1 + \frac{(p')^R}{\sum_{r=1}^{R}(p')^{r-1}}} \\
&\geq \frac{\frac{\sum_{r=1}^{R}(p)^{r-1}a_r}{\sum_{r=1}^{R}(p)^{r-1}} + \frac{(p')^R a_{R+1}}{\sum_{r=1}^{R}(p')^{r-1}}}{1 + \frac{(p')^R}{\sum_{r=1}^{R}(p')^{r-1}}}
\end{aligned}$$

According to Lemma 15, since $\frac{(p')^R}{\sum_{r=1}^{R}(p')^{r-1}} > \frac{(p)^R}{\sum_{r=1}^{R}(p)^{r-1}}$ and $a_{R+1} \geq \frac{\sum_{r=1}^{R}(p)^{r-1}a_r}{\sum_{r=1}^{R}(p)^{r-1}}$, we have

$$\begin{aligned}
\frac{\sum_{r=1}^{R+1}(p')^{r-1}a_r}{\sum_{r=1}^{R+1}(p')^{r-1}} &\geq \frac{\frac{\sum_{r=1}^{R}(p)^{r-1}a_r}{\sum_{r=1}^{R}(p)^{r-1}} + \frac{(p')^R a_{R+1}}{\sum_{r=1}^{R}(p')^{r-1}}}{1 + \frac{(p')^R}{\sum_{r=1}^{R}(p')^{r-1}}} \\
&\geq \frac{\frac{\sum_{r=1}^{R}(p)^{r-1}a_r}{\sum_{r=1}^{R}(p)^{r-1}} + \frac{(p)^R a_{R+1}}{\sum_{r=1}^{R}(p)^{r-1}}}{1 + \frac{(p)^R}{\sum_{r=1}^{R}(p)^{r-1}}} \\
&= \frac{\sum_{r=1}^{R+1}(p)^{r-1}a_r}{\sum_{r=1}^{R+1}(p)^{r-1}}
\end{aligned}$$

□

Based on the above two lemmas, we present the following theorem to show the uniqueness of the value of $\hat{\omega}$.

Lemma 17. *Eq. (5.13) has a unique solution* $\hat{\omega}$.

Proof. We show that the function $F(\omega) \triangleq S(\omega) - \omega$ is continuous and strictly decreasing in the interval $[0, 1]$ with $F(0) > 0$ and $F(1) < 0$. Therefore, according to the intermediate value theorem (Szecsei, 2007), there exists a unique solution $F(\hat{\omega}) = 0$ (i.e., $S(\hat{\omega}) = \hat{\omega}$).

According to (5.7), $P(0) = 0$. Therefore,

$$\begin{aligned} F(0) &= S(0) - 0 \\ &= \frac{\sum_{r=1}^{R}(P(0)^{r-1}T)}{\sum_{r=1}^{R}((P(0))^{r-1}(d_r^{(s)}(1-P(0)) + d_r^{(f)}P(0) + T))} \\ &= \frac{T}{T + d_1^{(s)}} > 0. \end{aligned}$$

Since $S(\omega)$ is always strictly smaller than 1 (due to the MAC timing constants that only appear in the denominator), we have

$$F(1) = S(1) - 1 < 0.$$

It remains to prove that the derivative of $F(\omega)$ is always negative in the interval $[0, 1]$. That is,

$$\frac{d(S(\omega) - \omega)}{d\omega} = \frac{dS(\omega)}{dP(\omega)} \cdot \frac{dP(\omega)}{d\omega} - 1 < 0.$$

The derivative of $P(\omega)$ is

$$\frac{dP(\omega)}{d\omega} = e^{-\omega}(1-\omega) + e^{-\omega} = e^{-\omega}(2-\omega) > 0.$$

We next prove that $\frac{dS(\omega)}{dP(\omega)}$ is negative for all $\omega \in [0, 1]$, which proves the result. That is, $S(\omega)$ decreases as $P(\omega)$ increases.

Since $S_\omega \geq 0$, the function $S(\omega)$ is decreasing if and only if $S(\omega)^{-1}$ is increasing. We thus investigate $S(\omega)^{-1}$ which is

$$\begin{aligned} S(\omega)^{-1} &= \frac{\sum_{r=0}^{R-1}(P(\omega))^r(d_{s,r}(1-P(\omega)) + d_{f,r}P(\omega) + T)}{\sum_{r=0}^{R-1}(P(\omega))^r T} \\ &= \frac{\sum_{r=0}^{R-1}(P(\omega))^r d_{s,r}}{\sum_{r=0}^{R-1}(P(\omega))^r T} + \frac{\sum_{r=0}^{R-1}(P(\omega))^{r+1}(d_{f,r} - d_{s,r})}{\sum_{r=0}^{R-1}(P(\omega))^r T} + 1. \end{aligned} \tag{5.14}$$

The above equation shows that $S(\omega)^{-1}$ can be divided into three terms, $\frac{\sum_{r=0}^{R-1}(P(\omega))^r d_{s,r}}{\sum_{r=0}^{R-1}(P(\omega))^r T}$, $\frac{\sum_{r=0}^{R-1}(P(\omega))^{r+1}(d_{f,r}-d_{s,r})}{\sum_{r=0}^{R-1}(P(\omega))^r T}$ and 1. If all those three terms are non-increasing functions of $P(\omega)$ and at least one of them is increasing, then $S(\omega)^{-1}$ is increasing with $P(\omega)$. Since the term 1 is a constant, we start with the term $\frac{\sum_{r=0}^{R-1}(P(\omega))^{r+1}(d_{f,r}-d_{s,r})}{\sum_{r=0}^{R-1}(P(\omega))^r T}$. According to

(2.3) and (2.4), the value of $d_{f,r} - d_{s,r}$ is a positive constant. Then we simplify

$$\frac{\sum_{r=0}^{R-1}(P(\omega))^{r+1}(d_{f,r} - d_{s,r})}{\sum_{r=0}^{R-1}(P(\omega))^{r}T} = \frac{P(\omega)(d_{f,r} - d_{s,r})}{T} \tag{5.15}$$

which is increasing with $P(\omega)$. We finally investigate the term

$$\begin{aligned} & \frac{\sum_{r=0}^{R-1}(P(\omega))^{r}d_{s,r}}{\sum_{r=0}^{R-1}(P(\omega))^{r}T} \\ = \quad & \frac{\sum_{r=0}^{R-1}(P(\omega))^{r}(T_{\text{DIFS}} + \overline{T}_{\text{backoff},r} + T_{\text{SIFS}} + T_{\text{ACK}})}{\sum_{r=0}^{R-1}(P(\omega))^{r}T} \end{aligned}$$

Since $\overline{T}_{\text{backoff},r}$ is an increasing sequence in r and $T_{\text{DIFS}} + T_{\text{SIFS}} + T_{\text{ACK}}$ is constant, according to Lemma 16, the value of $\frac{\sum_{r=0}^{R-1}(P(\omega))^{r}d_{s,r}}{\sum_{r=0}^{R-1}(P(\omega))^{r}T}$ increases with $P(\omega)$. We derive that $S(\omega)^{-1}$ increases as $P(\omega)$ decreases, which proves the result. □

We next determine when a congested fixed point exists at $\hat{\omega}$, for a given traffic load ρ. Based on (5.11), such a fixed point must satisfy

$$\hat{\omega} \leq U(\hat{\omega}). \tag{5.16}$$

Let

$$G(\omega) \triangleq \frac{\rho\omega}{U(\omega)} = \frac{\omega}{\sum_{r=1}^{R}\left(1 - \mathrm{e}^{-\omega}(1-\omega)\right)^{r-1}}. \tag{5.17}$$

The following lemma follows directly from (5.16) and (5.17).

Lemma 18. *A congested fixed point exists at $\hat{\omega}$ if and only if $\rho \geq G(\hat{\omega})$.*

Proof. We construct a chain of iff implications, starting with the statement that the fixed point $\hat{\omega}$ exists. Based on Equation (5.11), this property holds iff

$$\hat{\omega} \leq U(\hat{\omega}).$$

From (5.17), this property holds iff

$$\rho \geq \frac{\rho\hat{\omega}}{U(\hat{\omega})} = G(\hat{\omega}).$$

□

The following lemma establishes when an uncongested fixed point exists.

Lemma 19. *An uncongested fixed point exists if and only if* $\rho \leq \max_{\omega\in[0,\hat{\omega}]} G(\omega)$.

Proof. We prove that the existence of an uncongested fixed point implies $\rho \leq \max_{\omega\in[0,\hat{\omega}]} G(\omega)$ and vice-versa. On one hand, suppose an uncongested fixed point $\breve{\omega}$ exists. According to Definition 5, the uncongested fixed point $\breve{\omega}$ satisfies

$$\breve{\omega} = min\{U(\breve{\omega}), S(\breve{\omega})\} = U(\breve{\omega}), \tag{5.18}$$

and the congested fixed point $\hat{\omega}$ satisfies

$$\hat{\omega} = min\{U(\hat{\omega}), S(\hat{\omega})\} = S(\hat{\omega}). \tag{5.19}$$

Since Lemma 17 shows that the congested fixed point $\hat{\omega}$ is unique, we replace $S(\breve{\omega})$ in (5.18) by $\hat{\omega}$. We thus have

$$\breve{\omega} = min\{U(\omega), \hat{\omega}\} = U(\omega) \leq \hat{\omega}. \tag{5.20}$$

That is, $0 \leq \breve{\omega} \leq \hat{\omega}$. Then the traffic load ρ must satisfy

$$\rho = \frac{\rho\breve{\omega}}{U(\breve{\omega})} = G(\breve{\omega}) \leq \max_{\omega\in[0,\hat{\omega}]} G(\omega). \tag{5.21}$$

On the other hand, let $\rho \leq \max_{\omega\in[0,\hat{\omega}]} G(\omega)$. Since $G(0) = 0$, for any $\omega \in [0, \hat{\omega}]$, it satisfies

$$0 \leq G(\omega) \leq \max_{\omega\in[0,\hat{\omega}]} G(\omega).$$

Besides, $G(\omega)$ is continuous at the interval $\omega \in [0, \hat{\omega}]$. According to the intermediate value theorem (Szecsei, 2007), if $\rho \in [0, \max_{\omega\in[0,\hat{\omega}]} G(\omega)]$, then there exists at least one $\omega \in [0, \hat{\omega}]$ such that $\rho = G(\omega)$. According to (5.17), that ω also satisfies $\omega = U(\omega)$, which represents an uncongested fixed point.

□

5.3.5 Avoidance of cascading DoS attacks

We next establish a sufficient condition to avoid a cascading DoS attack on a network. According to Definition 7, a cascading DoS attack is unfeasible if Equation (5.11) has only one type of fixed points (i.e., either uncongested or congested) for each ρ. Hence, we provide the following lemma.

Lemma 20. *If $G(\hat{\omega}) > G(\omega)$ for all $\omega \in [0, \hat{\omega})$, then Equation (5.11) has only one type of fixed points for each traffic load $\rho > 0$.*

Proof. The result follows directly from Lemma 18 and 19. When $\rho > G(\hat{\omega})$, only a congested fixed point exists, while when $\rho < G(\hat{\omega})$, only one (or more) uncongested fixed points exist. Note that in the special case $\rho = G(\hat{\omega})$, there exists a unique fixed point $\hat{\omega}$ that is both congested and uncongested since $U(\hat{\omega}) = S(\hat{\omega})$. This boundary case is similar to when the server load equals 1 in a queueing system. Nevertheless, since the fixed point is unique, an attacker cannot impact the limiting fixed point in that case either. □

Let

$$\alpha \triangleq \frac{3 - \sqrt{5}}{2} \approx 0.38. \tag{5.22}$$

We now state our first main result.

Theorem 21 (Prevention of cascading attacks)**.** *A cascading DoS attack is unfeasible if $\hat{\omega} \leq \alpha$, where $\hat{\omega}$ is the unique solution of (5.12) and α is given by (5.22).*

Proof. Using algebra, the function $G(\omega)$ can be shown to be strictly increasing in the interval $[0, \alpha]$. The result then follows by Lemma 20. □

The above theorem implies that an attacker cannot launch a cascading DoS attack, if $\hat{\omega}$ is kept sufficiently low.

5.3.6 Optimizing the congestion throughput

In this section, we optimize the packet duration to achieve the highest throughput performance when the network is congested. We remind that the throughput of node

A_i is defined as the average number of bits per second that it successfully transmits to node B_i (this quantity is also sometimes referred to as goodput in the literature). The *congestion throughput* is the throughput of a node when packets are always waiting in its queue (i.e., when the queue is unstable). The congestion throughput can be found by taking the product of the congested utilization with the probability that a packet does not get lost. As i get large (i.e., looking at a node far down in the chain), the congested utilization of node A_i converges to $S(\hat{\omega}) = \hat{\omega}$ and the packet loss probability converges to $P(\hat{\omega})$, where the functions $P(\cdot)$ and $S(\cdot)$ are defined in Eqs. (5.7) and (5.9), respectively. The congestion throughput is therefore given by

$$\begin{aligned} X(\hat{\omega}) &\triangleq (1 - P(\hat{\omega})) \cdot \hat{\omega} \\ &= e^{-\hat{\omega}}(1 - \hat{\omega}) \cdot \hat{\omega}. \end{aligned} \tag{5.23}$$

Eq. (5.23) implies that the congestion throughput $X(\hat{\omega})$ does not always increase with $\hat{\omega}$. The following theorem determines the value of $\hat{\omega}$ that optimizes $X(\hat{\omega})$.

Theorem 22 (Optimal congestion throughput)**.** *The maximum congestion throughput is achieved at $\hat{\omega} = \alpha$, where α is given by (5.22).*

Proof. Let $\hat{\omega} \in [0, 1]$. According to (5.23), the derivative of $X(\hat{\omega})$ is

$$X'(\hat{\omega}) = e^{-\hat{\omega}}(1 - 3\hat{\omega} + \hat{\omega}^2). \tag{5.24}$$

There exists a unique solution of the equation $X'(\hat{\omega}) = 0$ at $\hat{\omega} = \alpha$. Since the second order derivative of $X(\hat{\omega})$ is negative at $\hat{\omega} = \alpha$, that is,

$$X''(\alpha) = e^{-\alpha}(-4 + 5\alpha - \alpha^2) < 0,$$

we conclude that $X(\alpha)$ is the maximum of $X(\hat{\omega})$ in the interval $\hat{\omega} \in [0, 1]$. □

Combined with Theorem 21, we obtain the remarkable result that $\hat{\omega} = \alpha$ both prevents cascading DoS attacks and maximizes the congestion throughput.

By setting $\hat{\omega} = \alpha$, we can calculate the optimal packet duration T^* that maximizes the congestion throughput. Specifically substituting $\hat{\omega} = \alpha$ into (5.13) and using (5.9), we get

$$T^* = \frac{\alpha \sum_{r=1}^{R} (P(\alpha))^{r-1} (d_r^{(s)}(1 - P(\alpha)) + d_r^{(f)} P(\alpha))}{(1-\alpha) \sum_{r=1}^{R} (P(\alpha))^{r-1}}. \tag{5.25}$$

Note that for any bit rate, the optimal packet length can be found by multiplying the optimal packet duration with the bit rate.

According to (5.25), the optimal packet duration is affected by the MAC overhead parameters. In particular, the optimal packet duration in IEEE 802.11b networks is longer than in 802.11g/n networks. Using the parameters shown in Table 2.1 the optimal packet duration in IEEE 802.11b is $T^* = 1.10$ ms, while in IEEE 802.11g/n with long slot time $T^* = 0.65$ ms and with short slot time $T^* = 0.27$ ms.

5.4 Simulation Results

We next present simulation results using ns-3 (ns-3, 2018). We first demonstrate the importance of properly modeling MAC timing parameters in the context of cascading DoS attacks (the impact of the packet length was shown in Section 5.2.2). We then validate the accuracy of our analytical model in predicting the congested utilization of a network. Next, we verify Theorems 21 and 22, and compare the performance of our method (based on optimizing the packet duration) to an RTS/CTS-based method. Finally, we evaluate the performance of the mitigation in an office building scenario with cross traffic. All the simulations shown in this section assume that the retry limit R is set to 7 and nodes communicate using UDP. Each simulation is run for 200 seconds and the plotted results are averages computed over three independent runs.

5.4.1 Impact of MAC timing parameters

We compare the behavior of IEEE 802.11g/n networks using respectively a long slot time (i.e., $T_{\text{slot}} = 20$ μs) and a short slot time (i.e., $T_{\text{slot}} = 9$ μs). All the other system parameters are identical. The network contains 20 pairs of nodes (see Fig. 5·1). Each node A_i transmits 1500 bytes packets at 6 Mb/s bit rate to node B_i ($i = 1, 2, \ldots, 20$). The offered load of nodes A_i ($i \geq 2$) is set to $\rho = 0.14$.

The simulation results are shown in Fig. 5·4. When the network uses a short slot time, the utilization of node A_{20} jumps when the offered load of the attacker ρ_1 exceeds 0.5. Hence, a cascading DoS attack occurs in that case. However, when the network uses a long slot time, the utilization of node A_{20} is not affected. This result confirms that the MAC configuration has an important impact on the possible occurrence of a cascading DoS attack. Because a network using a short slot time has a higher congested utilization than a network using a long slot time it is more vulnerable to a cascading DoS attack, assuming that all the other parameters are fixed.

5.4.2 Model accuracy

We next check if the value of the congested fixed point $\hat{\omega}$, as given by Eq. (5.13), predicts well the limit of the sequence of node utilizations when the network is congested. An accurate estimation of $\hat{\omega}$ is crucial for Theorems 21 and 22.

We run ns-3 simulations with 50 pairs of nodes. To ensure that the network is congested, the offered load ρ is set to 0.98. Fig. 5·5 depicts the utilization of node A_{50} for different bit rates and packet lengths. Fig. 5·5(a) shows results for an IEEE 802.11b configuration while Fig. 5·5(b) shows results for an IEEE 802.11g/n with short slot time. Both figures show excellent match between the analytical and simulation results. In both cases, the congested utilization decreases with the bit rate

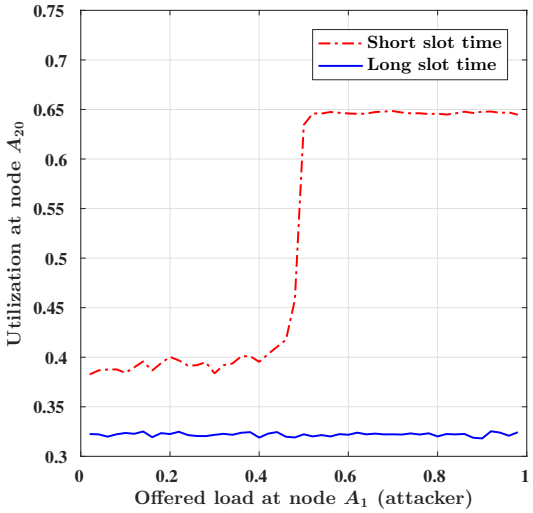

Figure 5·4: IEEE 802.11g/n networks under different MAC configurations. With a short slot time $T_{\text{slot}} = 9$ μs, a cascading DoS attack occurs. However, the attack does not occur if the network uses a long slot time $T_{\text{slot}} = 20$ μs.

but increases with the packet length. This is expected since the overhead of MAC timing parameters remains constant. Likewise, for a given bit rate and packet length, the congested utilization of IEEE 802.11g/n is higher than that of IEEE 802.11b, due to the lower MAC overhead of IEEE 802.11g/n. While such a property is generally viewed as desirable, it makes a network more vulnerable to a cascading DoS attack as explained previously.

5.4.3 Empirical validation of Theorems 21 and 22

We finally empirically validate our main results, namely that if $\hat{\omega} = \alpha$ then a cascading DoS attack is unfeasible for all traffic loads and the congestion throughput is maximized. To achieve the desired congested utilization α, we compute the theoretically optimal packet length by taking the product of the optimal packet duration given by Eq. (5.25) with the bit rate.

All our simulations, run for different bit rates and MAC configuration (e.g., IEEE 802.11b and IEEE 802.11g/n), show that no cascading attack occurs when the packet

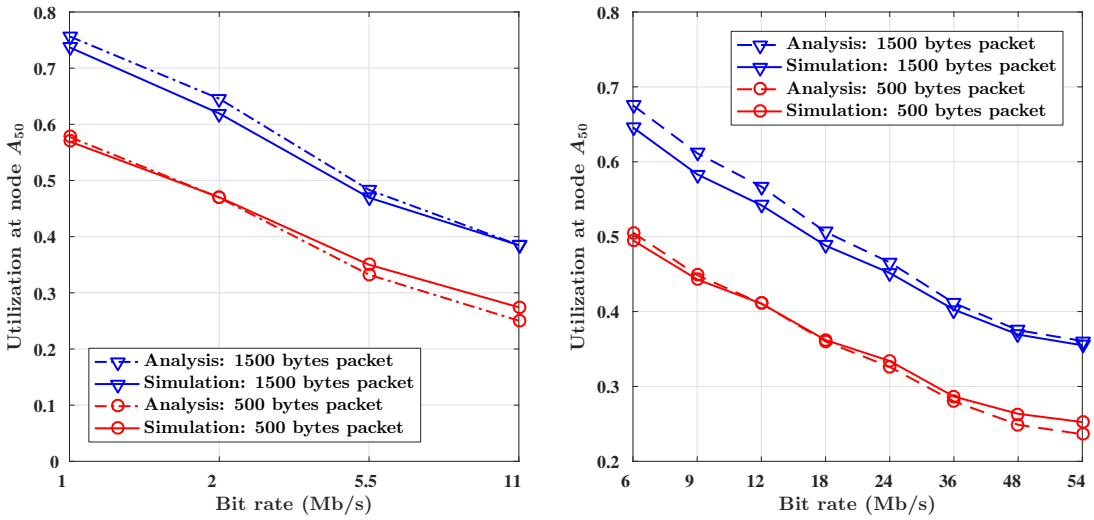

(a) 802.11b. **(b)** IEEE 802.11g/n.

Figure 5·5: Congested utilization: comparison of analytical and simulation results.

length is set optimally. For instance, for a bit rate of 6 Mb/s, the optimal packet length is 200 bytes. In that case, Fig. 5·3, which was introduced in Section 5.2.2, shows that the network experiences a cascading attack if the packet length is 1500 bytes and $\rho = 0.14$. On the other hand, an attacker cannot cause a cascading attack if the packet length is 200 bytes.

Next, we run simulations to evaluate the congestion throughput of the network using the optimal packet length. We set up a congested network consisting of 20 pairs of nodes with $\rho = 0.98$. We consider a 802.11g/n network using a long slot time. We compare the congestion throughput obtained using the theoretically optimal packet length, based on Eq. (5.25), with the maximum congestion throughput obtained empirically for 22 different packet lengths, that is, $100, 200, \ldots, 2200$ bytes. We also compare the results when enabling RTS/CTS with packets of length 500 bytes and 1500 bytes.

Figure 5·6 shows the congestion throughput of node A_{20} at different bit rates.

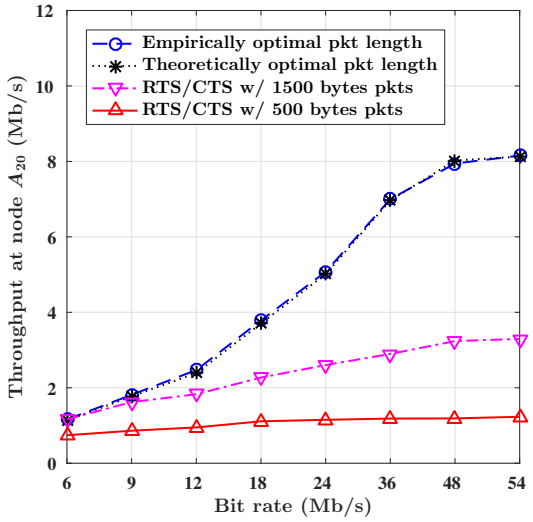

Figure 5·6: Comparison of congestion throughput in IEEE 802.11g/n, based on the theoretically optimal packet length, empirically optimal packet length, and RTS/CTS.

We observe that the congestion throughput obtained using the theoretically optimal packet length is close to the maximum congestion throughput obtained empirically over the 22 different packet lengths. Moreover, the congestion throughput is always higher than that obtained when using RTS/CTS and the difference becomes more significant as the bit rate increases. When the bit rate is 54 Mb/s, the congestion throughput obtained when using the optimal packet length is 2.5 times higher than that obtained when using 1500 bytes packets in conjunction with RTS/CTS.

5.4.4 Topology with cross traffic

We present the simulation results to evaluate the performance of the mitigation in an office building scenario with the cross traffic. We consider the communication within an office building using the same ns-3 building model as in Section 5.2.2. As shown in Figure 5·7, we create one floor of office building with 3 rows of rooms and each row has 11 rooms. We use location (x, y) to present each room where $x \in \{0, 1, 2\}$ and $y \in \{0, 1, \ldots, 10\}$.

We consider the cascading attacks on three transmission pairs (A_i, B_i), $i \in \{1, 2, 3\}$, which are positioned in every other room of the middle row. Node A_1 is the attacker and the other two pairs are victims. Additionally, there are four other transmission pairs (C_j, D_j), $j \in \{1, 2, 3, 4\}$. The transmitters C_j are positioned at rooms $(0, 2)$, $(0, 6)$, $(2, 4)$, $(2, 8)$ and their receivers D_j are positioned at rooms $(0, 4)$, $(0, 8)$, $(2, 2)$, $(2, 6)$, respectively. We vary the offered load at node A_1 and set the offered load at nodes A_2 and A_3 to 0.12 and nodes C_j to 0.06 for all $j \in \{1, 2, 3, 4\}$. All the nodes transmit UDP packets at 6 Mb/s bit rate. The running time of each simulation is 200 s.

We first set the packet length to 1500 bytes and illustrate the feasibility of the cascading attack in Figure 5·8. We observe that as node A_1 starts to transmit after 50 s, the utilization of node A_3 increases from 0.2 to 0.45. Its throughput drops from 0.73 Mb/s to less than 0.6 Mb/s. The utilization and throughput of node A_3 recovers once node A_1 stops transmitting after 150 s. The cascading attack is feasible in this simulation.

We then demonstrate the mitigation of cascading DoS attacks by changing the packet length to 200 bytes. When node A_1 is transmitting, we observe that the utilization of node A_1 does not increase as shown in Figure 5·8(a). Meanwhile, the throughput of node A_3 does not change shown in Figure 5·8(b). This result shows that cascading attacks are still feasible in the topology with cross traffic and our mitigation still works.

5.5 Mitigation in Experimental Testbed

In this section, we provide experiment results to show that shortening packet length can prevent the occurrence of the cascading attacks in Wi-Fi networks. The idea is to repeat two experiments on the experimental testbed with the same parameter

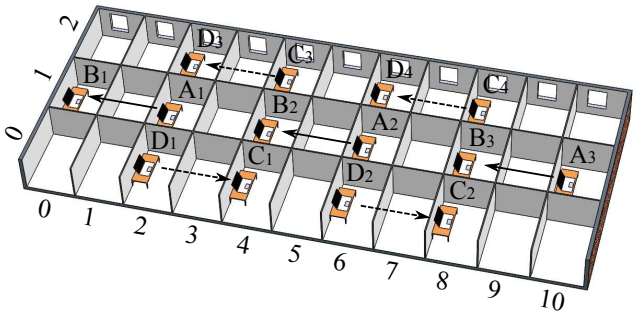

Figure 5·7: General network topology in an office building.

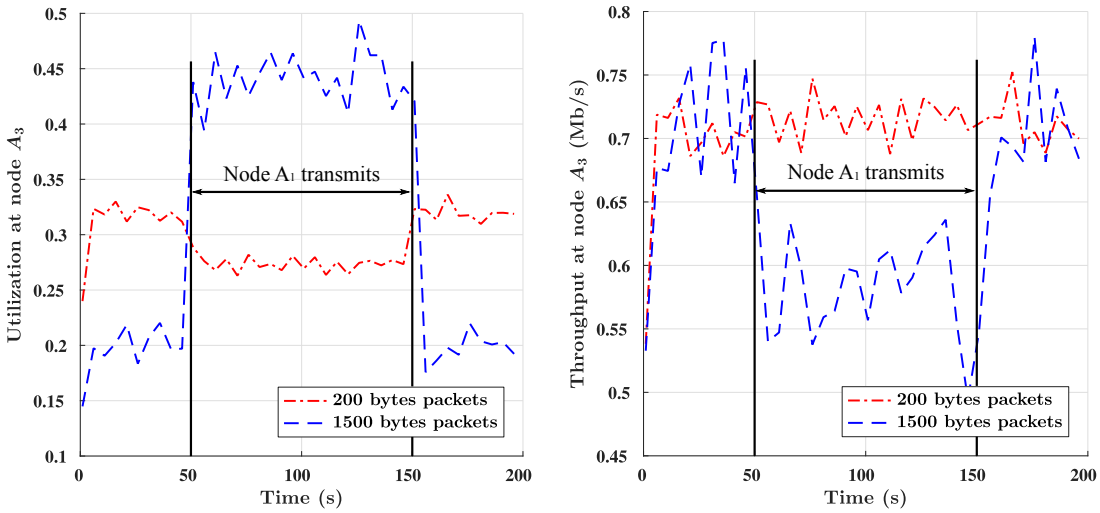

(a) Utilization. **(b)** Throughput.

Figure 5·8: Cascading attack and its mitigation in a network topology with cross traffic.

settings except the packet length. When the testbed uses the long packet length for transmissions, the cascading attacks are feasible. However, the attacks disappear when the testbed uses the short packet length instead.

We set up an experimental testbed as shown in Figure 5·9. We establish an IEEE 802.11n ad hoc network consisting of three transmission pairs (A_i, B_i), where $i \in \{1, 2, 3\}$. Each node A_i or B_i consists of a PC and a TP-LINK TL-WN772N Wi-Fi USB adapter. We use RF cables to link the nodes and use splinters to split and combine the Wi-Fi signals. As shown in the figure, we add 60 dB attenuators on the

links (A_i, B_{i+1}) and 70 dB attenuators on the links (A_i, B_i). We set the transmission power of all the nodes to 0 dBm. Thus, node A_i is a hidden node with respect to node A_{i+1}. Meanwhile, the received signal (interference) strength at node B_i from the hidden node A_{i-1} is stronger than the signal strength from node A_i. This makes sure that the interference due to hidden node A_{i-1} is able to corrupt the packet transmissions between A_i and B_i.

The goal is to test whether the increase of the traffic at the link (A_1, B_1) will affect the throughput at the link (A_3, B_3). The links (A_1, B_1) and (A_3, B_3) do not interfere each other directly. If we observe the throughput at the link (A_3, B_3) reduces due to the increase of the traffic at the link (A_1, B_1), then the cascading attack is feasible. Note that though we wire the nodes together, the nodes in the testbed can still receive the packets from outside. Therefore, the experiment results are also affected by the cross traffic.

In the experiments, we set the packet lengths to 500 bytes and 1500 bytes. Note that these two packet lengths are not optimal according to our mitigation. This is because the Wi-Fi cards use adaptive bit rates instead of fixed bit rate. However, our mitigation optimizes the packet duration. The optimal packet lengths vary by the bit rates. Since the Wi-Fi cards are closed-source, we are not able to set the optimal packet length at each bit rate.

The experiment results are illustrated in Figure 5·10. We first let node A_i transmit 1500 bytes packets to node B_i, where $i \in \{1, 2, 3\}$ and illustrate the results in Figure 5·10(a). We observe that when node A_1 starts to transmit after 50 s, the throughput of nodes A_2 and A_3 drops from 400 Kb/s to 100 Kb/s. When node A_1 stops transmitting after 100 s, the throughput of nodes A_2 and A_3 recovers. We then set the packet length to 500 bytes and repeat the experiment. As shown in Figure 5·10(b), the transmission at node A_1 does not have an impact on the throughput

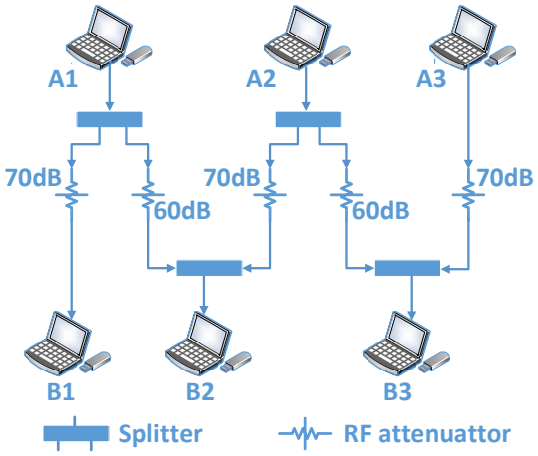

Figure 5·9: Experimental testbed.

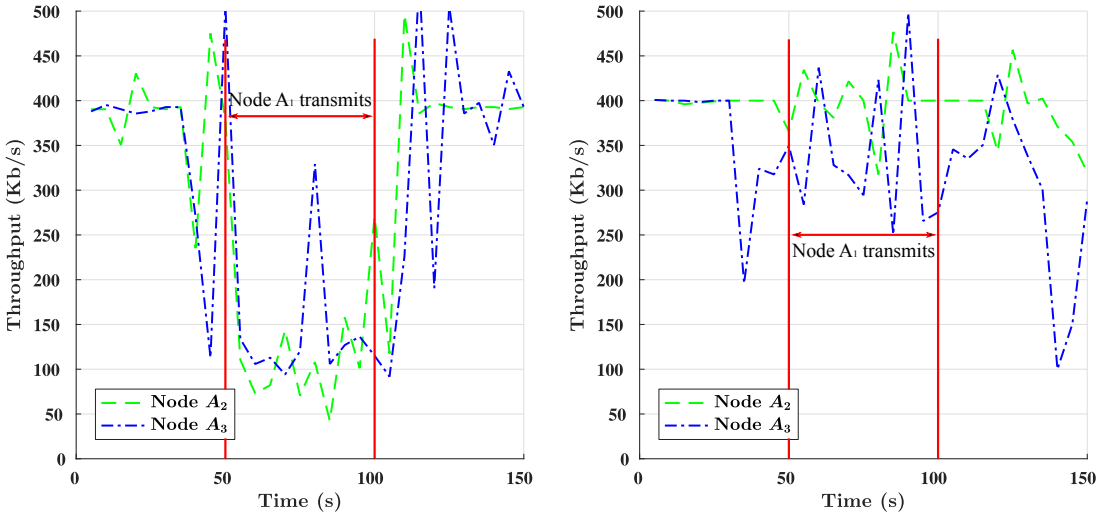

(a) Attack is feasible when nodes transmit 1500 bytes packets.

(b) Attack is unfeasible when nodes transmit 500 bytes packets.

Figure 5·10: Feasibility assessment of a cascading DoS attack in the experimental testbed.

of the other nodes. This result shows that the attack is feasible when the network transmits 1500 bytes packets but unfeasible when the packet length is shortened to 500 bytes. This result shows that the cascading attack can be avoided by shortening the packet length in real Wi-Fi networks.

5.6 Summary

In this chapter, we propose, analyze, and simulate a method to prevent cascading DoS attacks against Wi-Fi networks. When a cascading DoS attack is feasible, a small change in the offered load of the attacker can lead the network to suddenly transition from stability to instability. Our method derives the optimal packet length to prevent such change to ever occur for any traffic load. Moreover, for the same packet length, we show that the network achieves the maximum congestion throughput performance possible.

Specifically, we provide an analytical model to predict the feasibility of a cascading DoS attack. We develop an iterative analysis that characterizes the sequence of node utilizations, and use fixed point techniques to study its limiting behavior. We show that two types of fixed points may arise: uncongested fixed points and congested fixed points. We show that if the congested fixed point exists, it is unique. We further show that if the value of the congested fixed point $\hat{\omega}$ is lower or equal to $(3-\sqrt{5})/2 \approx 0.38$, then a cascading attack is unfeasible. In this case, the sequence of node utilizations can only converge to one type of fixed points, no matter what is the initial value of the sequence set by the attacker. The analysis captures the effect of MAC overhead parameters on the feasibility of launching a cascading DoS attack. For instance, with all other parameters kept fixed, we showed that an IEEE 802.11g/n network using a short slot time is more vulnerable to a cascading DoS attack than an IEEE 802.11g/n network using a long slot time.

Our mitigation method simultaneously optimizes the throughput performance of the network. Indeed, the analysis shows that when the congested utilization is $\hat{\omega} = (3 - \sqrt{5})/2$, the network achieves the highest congestion throughput. Our simulation results validate that the throughput performance of the network using the theoretically optimal packet length indeed approaches the highest possible through-

put and that it is higher (sometimes significantly) than the throughput obtained using RTS/CTS. We also provide the simulation results to investigate the impact of the cross traffic on the performance of our mitigation.

Chapter 6

Conclusion

In this dissertation, we investigate and demonstrate cascading DoS attacks in Wi-Fi networks. The attack starts with increasing the traffic load of one node in the network, the effect of which propagates through the network due to interference coupling, and results in global congestion of the network. Since this attack is remotely launched and protocol-compliant, it makes it harder to identify or locate the attacker.

We first show that a cascading DoS attack can be launched by exploiting the vulnerability of the CSMA mechanism in the MAC layer of IEEE 802.11 protocols. Since the RTS/CTS mechanism is generally turned off by default, the CSMA mechanism is vulnerable to the hidden node problem. The hidden nodes create an interference coupling effect between neighboring cells which facilitates the propagation of the attack. We show that the parameter retry limit plays an important role in the feasibility of the attack. The attack is feasible only when the retry limit is greater than 7.

We then unveil another vulnerability, namely the receiver capture effect at the physical layer of Wi-Fi. Under the receiver capture effect, a receiver aligns its state machine with information provided by the PHY header of the first transmission, before the second packet arrives. By exploiting the receiver capture effect, weak interference due to a hidden node can corrupt packet transmissions with strong power. This creates an interference coupling effect between neighboring cells. A cascading DoS attack becomes then feasible, forcing nodes in the network to operate at low bit rate.

We last propose a method to prevent cascading DoS attacks against Wi-Fi networks. Our method analytically derives the optimal packet length to prevent such change to ever occur for any traffic load. Moreover, for the same packet length, we show that the network achieves the maximum congestion throughput performance possible.

6.1 Future work

In the following, we discuss future directions of the research extending from this dissertation.

6.1.1 Coupling vulnerability

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 (Tews et al., 2007) and MD5 (Black et al., 2006)). In our case, our simulations for ring topologies indicate that the presence of a cycle in the topology could reinforce cascading DoS attacks, a result that warrants further investigations.

6.1.2 Mitigation of attacks due to strong interferers

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 Chapter 2. 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 include using collision-aware rate

adaptation algorithms, dynamic channel selection, and full-duplex radios. We leave the investigation and comparison of these mitigation techniques as possible areas for future work.

6.1.3 Mitigation of attacks due to weak interferers

Besides the existing RTS/CTS solution and shortening packet transmissions (Xin et al., 2018), we envision that several other approaches, likely more efficient, are possible. First, additional conditions could be added into the transitions from the CS/CCA state to the Rx state and back. For instance, when the receiver detects a valid PLCP preamble and header but determines that the probability of a packet loss is high (e.g., based on the received signal strength), it could decide to stay in the CS/CCA state (or move back from the Rx state to the CS/CCA state). Second, information about the destination could be added into the PLCP header. Instead of checking the receiver address at the MAC layer after the entire packet was received, this job could be performed at the physical layer, perhaps at least partially, to prevent the receiver capture effect.

References

Afaqui, M. S., Garcia-Villegas, E., Lopez-Aguilera, E., Smith, G., and Camps, D. (2015). Evaluation of dynamic sensitivity control algorithm for IEEE 802.11 ax. In *2015 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1060–1065. IEEE.

Aziz, A., Starobinski, D., and Thiran, P. (2011). Understanding and tackling the root causes of instability in wireless mesh networks. *IEEE/ACM Transactions on Networking*, 19(4):1178–1193.

Aziz, A., Starobinski, D., Thiran, P., and El Fawal, A. (2009). EZ-Flow: Removing turbulence in IEEE 802.11 wireless mesh networks without message passing. In *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, pages 73–84. ACM.

Bellardo, J. and Savage, S. (2003). 802.11 denial-of-service attacks: Real vulnerabilities and practical solutions. In *USENIX security*, pages 15–28.

Berg, J. (2016). Minstrel. `https://wireless.wiki.kernel.org/en/developers/documentation/mac80211/ratecontrol/minstrel`.

Bertsekas, D. and Gallager, R. (1992). Data networks. 1992. *PrenticeHall, Englewood Cliffs, NJ*.

Bianchi, G. (2000). Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communication*, 18(3):535–547.

Biaz, S. and Wu, S. (2008). Rate adaptation algorithms for ieee 802.11 networks: A survey and comparison. In *IEEE Symposium on Computers and Communications. ISCC 2008.*, pages 130–136. IEEE.

Bicakci, K. and Tavli, B. (2009). Denial-of-service attacks and countermeasures in IEEE 802.11 wireless networks. *Computer Standards & Interfaces*, 31(5):931–941.

Black, J., Cochran, M., and Highland, T. (2006). A study of the MD5 attacks: Insights and improvements. In *Fast Software Encryption*, pages 262–277. Springer.

Broustis, I., Eriksson, J., Krishnamurthy, S., and Faloutsos, M. (2007). Implications of power control in wireless networks: A quantitative study. *Passive and Active Network Measurement*, pages 83–93.

Calì, F., Conti, M., and Gregori, E. (2000). Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit. *IEEE/ACM Transactions on Networking (ToN)*, 8(6):785–799.

Chen, C., Luo, H., Seo, E., Vaidya, N. H., and Wang, X. (2007). Rate-adaptive framing for interfered wireless networks. In *IEEE INFOCOM 2007. 26th IEEE International Conference on Computer Communications*.

Cisco Systems, Inc. (2016). Cisco cleanair technology. `http://www.cisco.com/c/en/us/solutions/enterprise-networks/cleanair-technology/index.html`.

Cisco Systems, Inc. (2017). Cisco visual networking index: Global mobile data traffic forecast update, 2016 2021 white paper. `https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html`.

D-link (2016). D-link dir-855 user manual.

Dai, L. and Sun, X. (2013). A unified analysis of IEEE 802.11 DCF networks: Stability, throughput, and delay. *IEEE Transactions on Mobile Computing*, 12(8):1558–1572.

Daneshgaran, F., Laddomada, M., Mesiti, F., and Mondin, M. (2008). Unsaturated throughput analysis of IEEE 802.11 in presence of non ideal transmission channel and capture effects. *IEEE Transactions on Wireless Communications*, 7(4):1276–1286.

Duda, A. et al. (2008). Understanding the performance of 802.11 networks. In *2008 IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications*, volume 8. doi: 10.1109/PIMRC.2008.4699942.

Durvy, M., Dousse, O., and Thiran, P. (2007). Modeling the 802.11 protocol under different capture and sensing capabilities. In *IEEE INFOCOM 2007. 26th IEEE International Conference on Computer Communications*, pages 2356–2360. IEEE.

Foh, C. H. and Tantra, J. W. (2005). Comments on IEEE 802.11 saturation throughput analysis with freezing of backoff counters. *IEEE Communications Letters*, 9(2):130–132.

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

Gast, M. (2005). *802.11 wireless networks: the definitive guide*. O'Reilly Media, Inc.

Giustiniano, D., Malone, D., Leith, D. J., and Papagiannaki, K. (2007). Estimating link quality in 802.11 WLANs. Technical Report. Hamilton Institute, Maynooth, University. `http://eprints.maynoothuniversity.ie/2215/`.

Gummadi, R., Wetherall, D., Greenstein, B., and Seshan, S. (2007). Understanding and mitigating the impact of RF interference on 802.11 networks. *ACM SIGCOMM Computer Communication Review*, 37(4):385–396.

Hadzi-Velkov, Z. and Spasenovski, B. (2003). Capture effect with diversity in IEEE 802.11 b DCF. In *Proceedings. Eighth IEEE International Symposium on Computers and Communication, 2003 (ISCC 2003)*, pages 699–704. IEEE.

Haenggi, M., Andrews, J. G., Baccelli, F., Dousse, O., and Franceschetti, M. (2009). Stochastic geometry and random graphs for the analysis and design of wireless networks. *IEEE Journal on Selected Areas in Communications*, 27(7):1029–1046.

IEEE 802.11 Working Group and others (2004). Amendment 6: Medium access control (MAC) security enhancements. *IEEE Standard*, 802.

IEEE Standards Association and others (2012). 802.11-2012-IEEE standard for information technology–telecommunications and information exchange between systems local and metropolitan area networks–specific requirements part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications. *Retrieved from* ***http://standards.ieee.org/about/get/802/802.11.html***.

Intel (2017). Different Wi-Fi protocols and data rates. `https://www.intel.com/content/www/us/en/support/network-and-i-o/wireless-networking/000005725.html`.

Jiang, L. B. and Liew, S. C. (2007). Hidden-node removal and its application in cellular WiFi networks. *IEEE Transactions on Vehicular Technology*, 56(5):2641–2654.

Kamerman, A. and Monteban, L. (1997). Wavelan®-II: a high-performance wireless lan for the unlicensed band. *Bell Labs technical journal*, 2(3):118–133.

Kinney, R., Crucitti, P., Albert, R., and Latora, V. (2005). Modeling cascading failures in the north american power grid. *The European Physical Journal B-Condensed Matter and Complex Systems*, 46(1):101–107.

Kleinrock, L. (1975). *Queueing systems, Vol. 1.* Wiley, New York.

Kleinrock, L., Tobagi, F., et al. (1975). Packet switching in radio channels: Part I–carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Transactions on Communications*, 23(12):1400–1416.

Kong, Z. and Yeh, E. M. Wireless network resilience to degree-dependent and cascading node failures. In *7th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2009. WiOPT 2009.* IEEE.

Kumar, A., Altman, E., Miorandi, D., and Goyal, M. (2007). New insights from a fixed-point analysis of single cell IEEE 802.11 WLANs. *IEEE/ACM Transactions on Networking (TON)*, 15(3):588–601.

Lacage, M., Manshaei, M. H., and Turletti, T. (2004). IEEE 802.11 rate adaptation: a practical approach. In *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pages 126–134. ACM.

Li, X. and Zeng, Q. (2006). Capture effect in the IEEE 802.11 WLANs with rayleigh fading, shadowing, and path loss. In *IEEE International Conference on Wireless and Mobile Computing, Networking and Communications, 2006. (WiMob'2006)*, pages 110–115. IEEE.

Lin, G. and Noubir, G. (2005). On link layer denial of service in data wireless LANs. *Wireless Communications and Mobile Computing*, 5(3):273–284.

Linksys (2016). Wireless tab / advanced wireless settings. `http://ui.linksys.com/WAG300N/1.01.01/help/h_AdvWSettings.htm`.

Lynch, S. (2004). *Dynamical systems with applications using MATLAB*. Springer.

Magistretti, E., Chintalapudi, K. K., Radunovic, B., and Ramjee, R. (2011). WiFi-Nano: reclaiming WiFi efficiency through 800 ns slots. In *Proceedings of the 17th annual international conference on Mobile computing and networking*, pages 37–48. ACM.

Medhi, J. J. (2012). *Stochastic processes*. New Academic Science Limited, Tunbridge Wells, UK, 3rd edition.

Netgear (2016). Rangemax wifi range extender wpn824ext user manual. `http://documentation.netgear.com/WPN824EXT/enu/202-10310-02/WPN824EXT_UG-4-6.html`.

Noubir, G., Rajaraman, R., Sheng, B., and Thapa, B. (2011). On the robustness of IEEE 802.11 rate adaptation algorithms against smart jamming. In *Proceedings of the fourth ACM conference on Wireless network security*, pages 97–108. ACM.

ns-3 (2018). ns-3 documentation. `https://www.nsnam.org/documentation/`.

Nyandoro, A., Libman, L., and Hassan, M. (2007). Service differentiation using the capture effect in 802.11 wireless lans. *IEEE Transactions on Wireless Communications*, 6(8).

Orakcal, C. and Starobinski, D. (2014). Jamming-resistant rate adaptation in Wi-Fi networks. *Performance Evaluation*, 75:50–68.

Pei, G. and Henderson, T. (2009). Validation of ns-3 802.11 b PHY model. *Online: http://www. nsnam. org/˜ pei/80211b. pdf*.

Pei, G. and Henderson, T. R. (2010). Validation of OFDM error rate model in ns-3. *Boeing Research & Technology*, pages 1–15. `https://www.nsnam.org/~pei/80211ofdm.pdf`.

Pelechrinis, K., Iliofotou, M., and Krishnamurthy, S. V. (2011). Denial of service attacks in wireless networks: The case of jammers. *IEEE Communications Surveys & Tutorials*, 13(2):245–257.

Poisel, R. (2011). *Modern communications jamming principles and techniques*. Artech House Publishers.

Project, M. (2018). Bit-rate selection algorithms. `http://madwifi-project.org/wiki/UserDocs/RateControl`.

Ray, S., Carruthers, J. B., and Starobinski, D. (2003). RTS/CTS-induced congestion in ad hoc wireless LANs. In *Wireless Communications and Networking, WCNC 2003.*, volume 3, pages 1516–1521. IEEE.

Ray, S., Carruthers, J. B., and Starobinski, D. (2005a). Evaluation of the masked node problem in ad hoc wireless LANs. *IEEE Transactions on Mobile Computing*, 4(5):430–442.

Ray, S. and Starobinski, D. (2007). On false blocking in RTS/CTS-based multihop wireless networks. *IEEE Transactions on Vehicular Technology*, 56(2):849–862.

Ray, S., Starobinski, D., and Carruthers, J. B. (2005b). Performance of wireless networks with hidden nodes: a queuing-theoretic analysis. *Computer Communications*, 28(10):1179–1192.

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

Riley, G. F. and Henderson, T. R. (2010). The ns-3 network simulator. In *Modeling and tools for network simulation*, pages 15–34. Springer.

Rong, B. and Ephremides, A. (2009). Protocol-level cooperation in wireless networks: Stable throughput and delay analysis. In *7th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2009. WiOPT 2009.* IEEE.

Ross, S. M. (1996). *Stochastic processes*. Wiley, New York.

Saligrama, V. and Starobinski, D. (2006). On the macroscopic effects of local interactions in multi-hop wireless networks. In *4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, 2006*. IEEE.

Singh, A. and Starobinski, D. (2007). A semi markov-based analysis of rate adaptation algorithms in wireless LANs. In *4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON'07*, pages 371–380. IEEE.

Soltan, S., Mazauric, D., and Zussman, G. (2014). Cascading failures in power grids: analysis and algorithms. In *Proceedings of the 5th international conference on Future energy systems*, pages 195–206. ACM.

Stein, J. C. (1998). Indoor radio WLAN performance part II: Range performance in a dense office environment. *Intersil Corporation*, 2401 Palm Bay, Florida, 32905.

Sun, X. and Dai, L. (2015). Backoff design for IEEE 802.11 DCF networks: Fundamental tradeoff and design criterion. *IEEE/ACM Transactions on Networking (TON)*, 23(1):300–316.

Szecsei, D. (2007). *Calculus*. Homework helpers (Career Press Inc.). Career Press, Pompton Plains, N.J.

Tews, E., Weinmann, R.-P., and Pyshkin, A. (2007). Breaking 104 bit WEP in less than 60 seconds. In *Information Security Applications*, pages 188–202. Springer.

TP-Link (2016). Download center. `http://www.tp-link.us/support/download-center`.

Xia, D., Hart, J., and Fu, Q. Evaluation of the minstrel rate adaptation algorithm in IEEE 802.11g WLANs. In *2013 IEEE International Conference on Communications (ICC)*.

Xin, L. and Starobinski, D. (2018). Cascading attacks on Wi-Fi networks with weak interferers. In *Proceedings of the ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM'18)*.

Xin, L., Starobinski, D., and Noubir, G. (2016). Cascading denial of service attacks on Wi-Fi networks. In *2016 IEEE Conference on Communications and Network Security (CNS)*.

Xin, L., Starobinski, D., and Noubir, G. (2018). Mitigation of cascading denial of service attacks on Wi-Fi networks. In *2018 IEEE Conference on Communications and Network Security (CNS)*.

Xu, K., Gerla, M., and Bae, S. (2003). Effectiveness of RTS/CTS handshake in IEEE 802.11 based ad hoc networks. *Ad hoc networks*, 1(1):107–123.

Zou, Y., Zhu, J., Wang, X., and Hanzo, L. (2016). A survey on wireless security: Technical challenges, recent advances, and future trends. *Proceedings of the IEEE*, pages 1–39.

CURRICULUM VITAE