© 2013 Sang-Yoon Chang

SECURE PROTOCOLS FOR WIRELESS AVAILABILITY

BY

SANG-YOON CHANG

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2013

Urbana, Illinois

Doctoral Committee:

Associate Professor Yih-Chun Hu, Chair
Professor Nitin H. Vaidya
Associate Professor Nikita Borisov
Professor Carl Gunter

ABSTRACT

Since wireless networks share a communication medium, multiple transmissions on the same channel cause interference to each other and degrade the channel quality, much as multiple people talking at the same time make for inefficient meetings. To avoid transmission collision, the network divides the medium into multiple orthogonal channels (by interleaving the channel access in frequency or time) and often uses medium access control (MAC) to coordinate channel use. Alternatively (e.g., when the wireless users use the same physical channel), the network users can emulate such orthogonal channel access in processing by spreading and coding the signal. Building on such orthogonal access technology, this dissertation studies protocols that support the coexistence of wireless users and ensure wireless availability.

In contrast to other studies focusing on improving the *overall* efficiency of the network, I aim to achieve reliability at *all* times. Thus, to study the worst-case misbehavior, I pose the problem within a security framework and introduce an adversary who compromised the network and has insider access. In this dissertation, I propose three schemes for wireless availability: SimpleMAC, Ignore-False-Reservation MAC (IFR-MAC), and Redundancy Offset Narrow Spectrum (RONS). SimpleMAC and IFR-MAC build on MAC protocols that utilize explicit channel coordination in control communication. SimpleMAC counters MAC-aware adversary that uses the information being exchanged at the MAC layer to perform a more power efficient jamming attack. IFR-MAC nullifies the proactive attack of denial-of-service injection of false reservation control messages. Both SimpleMAC and IFR-MAC quickly outperform the Nash equilibrium of disabling MAC and converge to the capacity-optimal performance in worst-case failures. When the MAC fails to coordinate channel use for orthogonal access or in a single-channel setting (both cases of which, the attacker knows the exact frequency and time location of the victim's channel access), RONS introduces a physical-layer,

processing-based technique for interference mitigation. RONS is a narrow spectrum technology that bypasses the spreading cost and effectively counters the attacker's information-theoretically optimal strategy of correlated jamming.

To my family, for their unconditional love and support

ACKNOWLEDGMENTS

In Urbana-Champaign, I was fortunate to meet and interact with people who are not only highly accomplished in their respective fields but are also willing to share their insights and experience. Surrounded by such people, I grew substantially as an engineer and a researcher during graduate school, and I am grateful to the department and the university for providing me with the environment.

My advisor, Yih-Chun Hu, has provided me with invaluable guidance throughout my graduate school experience. Meeting with him provided great venues to practice critical thinking and other useful research skills (from bouncing off ideas to developing concrete projects). Yih-Chun further taught me of the joy of living as a researcher beyond the academic meetings and serves as a role model for my career.

I am grateful for my doctoral committee members Nitin Vaidya, Nikita Borisov, and Carl Gunter for their feedback on my dissertation work and the guidance that they provided for my post-doctoral life.

I would also like to thank my labmates and other colleagues with whom I shared my graduate school life. Jihyuk Choi and Dongho Kim, Jerry Chiang, and Jason Haas, my labmates in the Coordinated Science Laboratory at Illinois during my junior years, helped ease my transition to graduate school. My internship with Jerry Chiang at the Advanced Digital Sciences Center in Singapore facilitated me to contemplate and plan beyond graduation.

Lastly, as is apparent from the similarity between my career interest and his, my dad, Soo-Young Chang has been a big influence (for my career and everything beyond). He, my mom, and my family motivate me to become better.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Wireless systems offer an advantage in mobility and convenience (which are becoming increasingly important in many modern-day applications) but have a disadvantage that they inherently share a communication medium which is less reliable than the devoted cable medium of wired communication systems. To achieve the better of both ends, researchers in communications and computer networks have studied and designed schemes that will ensure reliable communication while embracing the coexistence of communication users within a shared medium.

Reliable transmission for wireless users depends on the availability of lower layers of the Open Systems Interconnection (OSI) model. Since more-capable wireless systems tend to use the same protocols as wired internet nodes at the network layer and above, and since less-capable wireless systems have substantial integration between application-level requirements and protocols used above the network layer, the most significant gap in reliable protocols for wireless systems lies at the physical and link layer. Therefore, I address reliability and robustness at the two lower layers.

To model the worst-case impact of misbehavior and ensure reliability at *all* times, I assume a malicious adversary. In contrast to previous work, I consider attackers that are insiders, intelligent, and adversarial. In particular, such adversaries are capable of reacting to legitimate user strategy ("intelligent"), they have the keys of one or more legitimate nodes ("insider"), and their goal is to minimize the throughput of legitimate users ("adversarial"). The adversarial model represents a worst-case scenario for wireless availability; a misbehaving user with equal capability that chooses any other strategy cannot result in worse legitimate user performance.

I discuss medium access control protocols that utilize explicit channel coordination in control communication in Chapter 2. Chapter 3 introduces SimpleMAC, a protocol that counters MAC-aware attacks where attackers

utilize MAC-layer information to have greater destructive impact on the network. SimpleMAC is comprised of Simple Signaling Scheme (SSS) and Simple Transmitter Strategy (STS). SimpleMAC counters two smart, power-efficient jamming attacks: SSS mitigates MAC-aware jamming attack on control communication (where the vulnerability comes from using a common, or known, control channel), and STS prevents MAC-facilitated jamming attack on data communication (where adversaries use the information being exchanged in control communication to focus their jamming on data channels that are being used). Then, in Chapter 4, the Ignore-False-Reservation MAC (IFR-MAC) that nullifies the proactive attack of Denial-of-Service injection of false reservation control messages is discussed. SimpleMAC and IFR-MAC, together, provide a DoS-secure MAC protocol.

In contrast to SimpleMAC and IFR-MAC, I also consider the case where MAC fails and the wireless system can not rely on orthogonal channel access for availability. Chapter 5 provides a primer for wireless communication and the physical-layer techniques used in practical settings, and Chapter 6 introduces RONS that is implemented at the physical layer and effectively suppresses the attacker-optimal interference of correlated jamming.

CHAPTER 2

MEDIUM ACCESS CONTROL BACKGROUND

2.1 Control-Based Medium Access Control

As wireless features are introduced into more and more electronic devices, it is becoming increasingly important to use scarce radio spectrum as efficiently as possible. An important part of efficient usage is effective coordination of user transmissions. Traditional protocols aim to avoid overlapping transmissions; the typical channel access schemes separate users' usage in some combination of time, frequency, and code. As networks increasingly carry data traffic, which is characterized by bursty arrivals, fixed channelization is being replaced by dynamic Medium Access Control (MAC) protocols that change user allocations from frame to frame. In a distributed MAC, each node announces its usage intentions, both to link a transmitter-receiver pair for communication and to help other transmitters minimize interference (since other nodes avoid making conflicting transmissions minimizing interference both to the node that has announced its intentions and to a node that cooperates by avoiding transmissions during the reserved slot). In this dissertation, the explicit messages containing channel use information are referred to as *reservations* and the task of exchanging reservations as *channel coordination*. In channel coordination, a network user reserves a channel by sending one or more control packets (that contain its channel usage intentions) on a *control channel*, and then uses the reserved *data channel* to send its data traffic.

Modern day-to-day communication widely makes use of such reservation-based MAC protocols to provide better performance. For example, IEEE 802.11 WLAN (also known as WiFi) virtual carrier sensing four-way handshaking protocol has transmitter-receiver pair exchange Request to Send (RTS) and Clear to Send (CTS) packets to reserve a channel; the broad-

casted packets notify the pair's pending use of the channel to the nearby users that are within the transmission range. On the other hand, Bluetooth and WiMax (based on IEEE 802.16 standard) have a centralized authority scheduling the channel use to the network users.

2.2 MACs Are Built for Collaborative Nodes

Channel coordination is only useful when other nodes respect reservations; such environments are called *collaborative environments*. In such environments, channel coordination protocols can provide substantial performance gains. In particular, if I characterize the channel capacity using *Shannon capacity* (as given by the Shannon-Hartley theorem), when two nodes with equal power levels share a band, coordinating nodes get capacity

$$\frac{W}{2}\log_2\left(1+\frac{S}{N}\right)$$

whereas when they do not cooperate, they get capacity

$$W\log_2\left(1+\frac{S}{S+N}\right)$$

since each node's transmission is interference to the every other node. Whenever the band's signal-to-noise ratio exceeds about 2 dB, coordination provides substantial gains.

However, when users are selfish, the Nash equilibrium is to disable MAC protocol and spread the transmission across the entire band [1,2]; such strategy results in the tragedy of the commons where the self-centered behavior over-exploits the shared network medium and the users end up performing worse than had they cooperated and complied to the protocol, since there is no reduction in collisions at the Nash equilibrium.

Even though a MAC is designed for collaborative environments I study the network behavior when a portion of the network deviates from the protocol (while the rest of the network is cooperative). A user might deviate for selfish reasons or a user's hardware may fail, leading to unpredictable results. This dissertation analyzed the protocol-deviance in the worst case, by considering the impact of an adversary whose sole goal is to minimize network

performance. Current protocols, when faced with an intelligent, insider jammer, will at best reach the Nash equilibrium, in which channel coordination is completely disabled and each message is spread across the entire band. That is, a misbehaving user can force the network's optimal behavior to give up the advantage of collaboration and disable MAC. Section 2.3 discusses threats that can reduce the optimal network behavior to turn off MAC.

2.3 Threat Overview on MAC

A wide variety of MAC-layer protocols have been proposed for various environments and applications. This section presents an outline of security vulnerabilities of existing wireless MAC protocols, where the attacker can jam control messages and can use control messages to jam more effectively.

2.3.1 Threat Overview and Related Work

To reduce the inefficiencies inherent in simultaneous channel usage, most wireless MAC-layer protocols (with few exceptions, such as ALOHA [3]) attempt to reserve a channel by exchanging channel coordination information. Traditionally, a common control channel is used to exchange channel coordination information among users. There are two important jamming attacks against a control channel: first, the attacker can jam the channel itself, and second, the attacker can use *jamming-relevant information* transmitted on the control channel (such as when and where data transmissions will take place) to facilitate effective jamming. In addition to jamming, an attacker can conduct a proactive attack of injecting incorrect control messages to deny availability to the network.

For denial-of-service (DoS) attack on wireless availability, previous literature describe attackers who can send excessive reservation messages to prevent legitimate nodes from using the channel [4–7]. Another form of adversarial behavior is channel jamming. Awerbuch et al. [8] propose a fair single-channel MAC protocol against a power-limited jammer that does not jam all of the time. Other papers propose mechanisms to avoid jamming [9–11] but these approaches are not secure against insider attacks; that is, when jammers are compromised network participants and thus have access to some of

the keys of the network nodes, jamming avoidance cannot be assured by this prior work.

2.3.2 Vulnerabilities in Current Protocols

The use of a common control channel, which is typical in currently available protocols, is vulnerable to jamming attacks. For example, in the IEEE 802.11 WiFi standard [12], nodes use *virtual carrier sense* in which they reserve the channel by exchanging Request to Send (RTS) and Clear to Send (CTS) messages; these messages can be jammed to reduce network performance. Though virtual carrier sense provides an effective way for a legitimate potential transmitter to avoid collisions with another transmitter, the very mechanism that allows them to mitigate interference also allows an attacker to jam every transmission. In particular, whenever an attacker senses another user's transmission, either through carrier sense or virtual carrier sense, the attacker can jam the corresponding data packet.

In addition to carrier sense and virtual carrier sense, WiFi also uses a Collision Avoidance mechanism in which a node transmitting a frame chooses a backoff interval. The node counts down the backoff interval whenever the channel is idle; this mechanism reduces the probability that two nodes will transmit simultaneously. Several researchers have investigated the attack wherein the attacker chooses incorrect backoff intervals [13–15].

In the Out-of-Band signaling scheme [16], each receiver sends a very narrowband busy tone whenever it receives data to indicate the channel is in use. A powerful adversary may be able to jam the busy tone, and even when the jammer is unable to remove the busy tone, a jammer that hears a busy tone knows that a receiver is active within its wireless transmission range. A jammer that jams the data channel whenever it hears a busy tone can effectively deny service to receivers within its interference range. An attacker can also falsely reserve the channel by continuously sending a busy tone.

IEEE 802.16 [17], commonly called WiMAX, uses a centralized scheduling algorithm in which the base station assigns time slots to each user. Since the base station broadcasts control messages, a jammer that knows the location of the control channel can either jam the control channel to disrupt the exchange of control messages or use the received channel scheduling information to jam

data transmissions at the assigned time slots (and frequency channels). In WiMAX, the control channel location is a published part of the standard, but even if it were not, an attacker that has compromised a legitimate node must know the location of the control channel. Furthermore, a node can request and be scheduled for time and frequency slots that it does not need, thus wasting time and bandwidth.

Another centralized protocol is Bluetooth [18], in which a master device sends control messages to each slave device in the network (called a *piconet*). An attacker who knows the frequency hopping pattern of a scheduled transmission can easily jam that transmission. In Bluetooth, these frequency hopping patterns are a public part of the standard, but even if it were not, an attacker that has compromised a legitimate node must know the location of the control channel. Furthermore, an attacker can become the master and have significant control over other legitimate users.

MAC protocols that do not perform channel coordination suffer from higher probability of collisions between simultaneous transmitters, resulting in more interference. Thus, a protocol that lacks channel coordination functionality yields lower SINR and therefore lower capacity.

In conclusion, currently implemented protocols either mitigate interference from legitimate nodes by regulating their channel usage, in which case jammers can effectively jam during legitimate node usage, or provide no channel coordination and suffer from increased interference.

CHAPTER 3

SimpleMAC

3.1 Chapter Overview

As discussed in Chapter 2, for efficient use of the shared network medium, wireless systems often use a Medium Access Control (MAC) protocol to perform channel coordination by having each node announce its usage intentions and other nodes avoid making conflicting transmissions. In a collaborative environment, such MAC protocols yield performance gain. However, when an attacking node can receive channel coordination information, such as when the attacker is a compromised network node, coordination information can also be used to jam more effectively, since jammers know exactly on which channel to focus their jamming power to disrupt data communication. Also, since the location of the control channel is pre-assigned and known to network users, attackers can jam the control channel, thereby eliminating any benefit legitimate users might gain from channel coordination.

MAC-layer Protocols, when faced with just one intelligent, insider jammer, will at best reach the Nash equilibrium, in which channel coordination is completely disabled and each message is spread across the entire band [1, 2]; at the Nash equilibrium, there is no reduction in collisions, which reflects a non-cooperative environment. In this chapter, I construct a theoretical framework to analyze the dynamics between the adversaries and legitimate users, then propose SimpleMAC, a MAC-layer protocol that performs channel coordination while mitigating the effects of jamming.

Section 3.3.1 describes the MAC-layer framework. A MAC protocol provides reduced probability of collision by exchanging a channel usage plan with other network users; this channel usage plan is *jamming-relevant information* because the plan allows legitimate users to avoid the transmitter but also allows jammers to intentionally collide with the transmitter. I divide

the scheme into two components. The *transmitter strategy* selects the set of network nodes with which to share the relevant control message. This set, which may vary for each packet, is called the *recipient list* and it is denoted with S. The *signaling scheme* delivers a control message to each node in the recipient list, and ensures that no other nodes are able to receive the control message. When the adversary is malicious, the ideal recipient list includes all legitimate users and no attackers.

SimpleMAC consists of the *Simple Signaling Scheme* (SSS) and the *Simple Transmitter Strategy* (STS). I develop a jamming-resistant signaling scheme to deliver jamming-relevant control message to exactly the set of nodes in the recipient list S, and a transmitter strategy that decides on the recipient list based on prior receiver feedback. In the transmitter strategy, the transmitter-receiver pair measures the performance of S after sending each packet and uses this information to adapt S for future packet transmissions. The long-term goal for the transmitter is to search for a set S that provides the optimal performance. As a general rule, the choice of $S = \emptyset$ (equivalent to disabling channel coordination) represents baseline performance; after a sufficient number of independent trials, any set that performs significantly worse must contain a jammer. Therefore, the transmitter can determine whether channel coordination has been compromised by comparing the performance of the recipient list with performance when $S = \emptyset$. However, since attackers are intelligent, and thus capable of dynamically changing their jamming strategy, a recipient list S with better performance than when $S = \emptyset$ does not necessarily mean that S excludes all attackers.

This thesis applies to both single-channel TDMA systems (with an *energy-limited* attacker, since a power-limited attacker would jam at all times and gain no advantage from channel coordination information) and multi-channel systems (with a *power-limited* attacker). For clarity of presentation, I present SimpleMAC as applied in multi-channel systems. Because I model legitimate users as cooperative and attackers as malicious, SimpleMAC must simultaneously allow legitimate users to avoid the transmissions and yet prevent attackers from coinciding with them. For this reason, each transmission in the protocol transfer is sent using Frequency Hopping Spread Spectrum (FHSS), in which each transmission is sent while the transmitter hops from one frequency band to another according to a pseudorandom hopping pattern. Channel coordination information thus consists of the time at which a

sender plans to send and the frequency hopping pattern the sender plans to use.

SimpleMAC quickly outperforms the case where channel coordination is disabled, eventually converges to the recipient list offering optimal performance, and forces the optimal jammer strategy to be jamming at full power all the time (even though jamming alerts the user and prompts it to stop sharing information with the compromised recipient lists).

The rest of the chapter is organized as follows. After presenting the model in Section 3.2 and setting up the theoretical framework in Section 3.3, I analyze the general jammer behavior in Section 3.4. I then introduce SimpleMAC in Section 3.5, and the jammer reaction to the SimpleMAC scheme in Section 3.6.1. Next, I mathematically analyze the performance of SimpleMAC in Section 3.6.2 and evaluate it using MATLAB simulations and WARP implementation in Section 3.7. Lastly, I present conclusions and open problems in Section 3.9.

3.2 System Model and Assumptions

I consider an environment with $T+1$ non-idle transmitters (each transmitter has T potential interference sources), each identified by an index $i \in \mathcal{T} = \{1, \ldots, T+1\}$, a subset $\mathcal{N} = \{i_1, \ldots, i_N\}$ of which are N jammers. All non-jammers are protocol-compliant. I assume a shared secret key between each pair of nodes, and that all nodes operate in shared spectrum divided into C channels, each with bandwidth W Hz. No online authority governs users. I consider a repeated game with infinite horizon; either the transmission never ends or the users do not know when the transmissions will end. I index the rounds of the game $r \in \{1, 2, 3, ...\}$.

I assume that each user has technical means to transmit on the spectrum, that the attacker ignores any legal prohibitions against interference, and that there is no way to a priori determine which node is trustworthy. Also, because of the possibility of jamming, I make the standard assumption [19, 20] that each transmitter sends data using fast frequency hopping on randomly generated hopping patterns chosen independently for each packet. Traditionally, fast frequency hopping is characterized by a hopping time of more than one hop per symbol; here, I only require that the hopping time be faster than

the jammer's reaction time.

At the physical layer, I assume there exists a known spreading gain at which any pair of neighbors can communicate with a suitably low bit error rate. Alternatively, I define a *neighbor* as a node that can be reached using a specific spreading gain. In SimpleMAC, described in Section 3.5, I use Direct Sequence Spread Spectrum (DSSS) for control communications and Frequency Hopping Spread Spectrum (FHSS) for data communications, and I communicate the frequency hopping pattern in the control message.

The communication between any transmitter-receiver pair is single-hop; that is, the transmitter does not rely on a third node to relay the message to the final destination. For simplicity, each user is a neighbor of every other user. Thus, when two nodes transmit on the same frequency at the same time, those transmissions interfere with each other, resulting in reduced capacity. SimpleMAC can be extended to hidden-terminal environments by having both sender and receiver repeat each channel coordination transmission, although the details of this approach are beyond the scope of this dissertation.

SimpleMAC is designed for unicast data transmissions, where performance affects only the receiver. Therefore, when a sender transmits to multiple receivers, the feedback of a malicious receiver does not affect the performance of other receivers. However, a malicious receiver may be able to induce a sender to choose $S = \emptyset$ for all transmissions to that receiver. The impact of this selection depends on the transmission priority scheme, so this attack is discussed further in Section 3.8.

All users, including attackers, share the same power constraint P_c. The case where each attacker is more powerful than a normal user can be modeled by increasing the fraction of nodes which are attackers.

3.2.1 Performance Metric

When user i transmits to user j, it does so on a frequency channel that varies with time according to a frequency hopping pattern known to user i and user j. At any point in time, the user transmits on frequency channel $c \in \{1, \ldots, C\}$. Assuming a flat fading channel with additive white Gaussian

noise and Gaussian signals, the channel capacity of the link $i \to j$ is:

$$\mathcal{R} = \int_{f_c - W/2}^{f_c + W/2} \log_2 \left[1 + \mathrm{SINR}_{i,j}(f)\right] df \tag{3.1}$$

where f_c is user i's carrier frequency, and SINR is the effective signal-to-interference-and-noise ratio at the receiver

$$\mathrm{SINR}_{i,j} = \frac{\gamma_{i,j}\widetilde{P}_i(f)}{\widetilde{N_0} + \sum_{\ell \neq i, \ell \in \mathcal{N}^c}[\gamma_{\ell,j}\widetilde{P}_\ell(f)] + \sum_{k \in \mathcal{N}}[\gamma_{k,j}\widetilde{J}_k(f)]} \tag{3.2}$$

In Equation 3.2, $\gamma_{a,b}$ is the channel gain between transmitter a and receiver b, $\widetilde{N_0}$ is the power spectral density of the noise, $\mathcal{N}^c$ are the indices of legitimate users, $\mathcal{N}$ are the indices of jammers, $\widetilde{P_\alpha}$ is transmitter α's power spectral density for some α, and $\widetilde{J}_k(f)$ is the jammer k's power spectral density.

Shannon channel capacity ($\mathcal{R}$) is an upper bound on communication rate performance. Channel capacity is a mathematically simple formulation and is tight in many practical environments; existing codes very nearly achieve channel capacity [21]. In order to separate MAC-layer issues from physical-layer decisions such as modulation and coding, I use both SINR and channel capacity as representative performance metrics in the mathematical analysis. I observe that Equation 3.1 exhibits two properties that I use in my analysis: it is decreasing and convex with respect to jamming power and monotonically increasing with respect to the user's signal power. Though I use SINR and capacity as representative measures of performance, my approach generalizes to any utility function that is convex in interference power and monotonically increasing in SINR. (In Section 3.7, the implementation testbed simulation results show the effective SINR at the receiver, because achieving channel capacity involves sophisticated coding and modulation, and because the instantaneous capacity is strictly monotonic in the instantaneous SINR.)

Channel capacity $\mathcal{R}$ (Equation 3.1) serves as the utility function for the legitimate transmitter i. The transmitter's aim is to maximize its capacity $\mathcal{R}$. As $\mathcal{R}$ is a monotonically increasing function of P_i, the transmitter will emit full power. To aggregate capacity (which is an instantaneous metric) over time, I compute its time-average. At time t, given $\{\mathcal{R}_{t'} \mid t' < t\}$, the

utility function is:

$$U = \frac{1}{N_t} \sum_{t'=t-N_t+1}^{t} \mathrm{E}[\mathcal{R}_{t'}] \tag{3.3}$$

where $\mathcal{R}_\gamma$ is the capacity measured at time γ for some γ. In an infinite-horizon game, Equation 3.3 is replaced by its limit as $N_t \to \infty$, where N_t represents the time duration of transmission.

3.2.2 Attacker Model

I consider a jammer that intends to minimize the utility function, Equation 3.3, subject to its power constraint:

$$\text{minimize } U \text{ subject to } \int_f \widetilde{J}_k(f)\, df \leq P_c,\ \forall k \in \mathcal{N} \tag{3.4}$$

I assume that jammers collude. Thus, if one jammer knows a user's frequency hopping pattern, then all jammers can make use of that information. Also, attackers know the protocol and can adaptively change their strategies according to the legitimate users' strategies.

I also consider reactive jammers that jam according to their observations on the target signal, in case they do not receive the user's channel coordination information. To counteract reactive jammers, the user can shorten the frequency hopping time so that the jammers do not have enough time to observe the spectrum and jam the used channel. I do not consider the very strong and sophisticated attack of correlated jamming, where an adversary mimics the target signal with a phase offset of π at equal amplitude, canceling the target signal (Chapter 6 considers such threat). Under this attack, assuming the attacker has at least as much power as the legitimate node at a target receiver, no physical layer can provide any throughput [22].

Attackers can choose between narrowband jamming (concentrating its power on one or a subset of frequency channels at a time) or wideband jamming (emitting power across the spectrum at a time) and can freely switch between these strategies. Because the jammer has so much flexibility, I do not consider legitimate user attempts to infer information about a jammer; however, this approach still converges to the optimal performance.

I also consider the possibility of non-Gaussian jamming, since Equation 3.1 holds only when all received signals are Gaussian signals. Given a Gaussian fading channel with Gaussian signals, where the overall signal power is much greater than the combined power of the jammer network, the optimal jammer strategy is to jam with Gaussian noise [22, 23]. Also, the transmitter can make any received jamming signal appear Gaussian by using a sufficiently long Direct Sequence Spread Spectrum (DSSS) code shared only with the receiver. When jammers do not know the code used by the transmitter, the received jamming signal looks like a Gaussian signal by the central limit theorem.

3.3 Theoretical Framework

3.3.1 Overview

I design SimpleMAC from the ground up without making any MAC-layer assumptions. The MAC-layer framework contains two parts: a transmitter strategy and a signaling scheme. For each packet, a transmitter strategy determines the set of users S that will receive the channel coordination information for that packet; this set is called the *recipient list*. The optimal transmitter strategy will prevent the attacker from gaining any advantage from its knowledge of insider network keys while minimizing interference from legitimate users. A signaling scheme delivers the control message to the recipient list and provides availability (that is, messages are not easily jammed) and confidentiality (that is, nodes not on the recipient list will not receive the control message).

Figure 3.1 depicts the MAC-layer framework. When sending a packet, the transmitter (1) chooses a subset S of network users, (2) transmits its frequency hopping information to S, (3) transmits the data packet using the previously reserved hopping pattern, and (4) determines the effectiveness of S based on the feedback that it receives from the receiver.

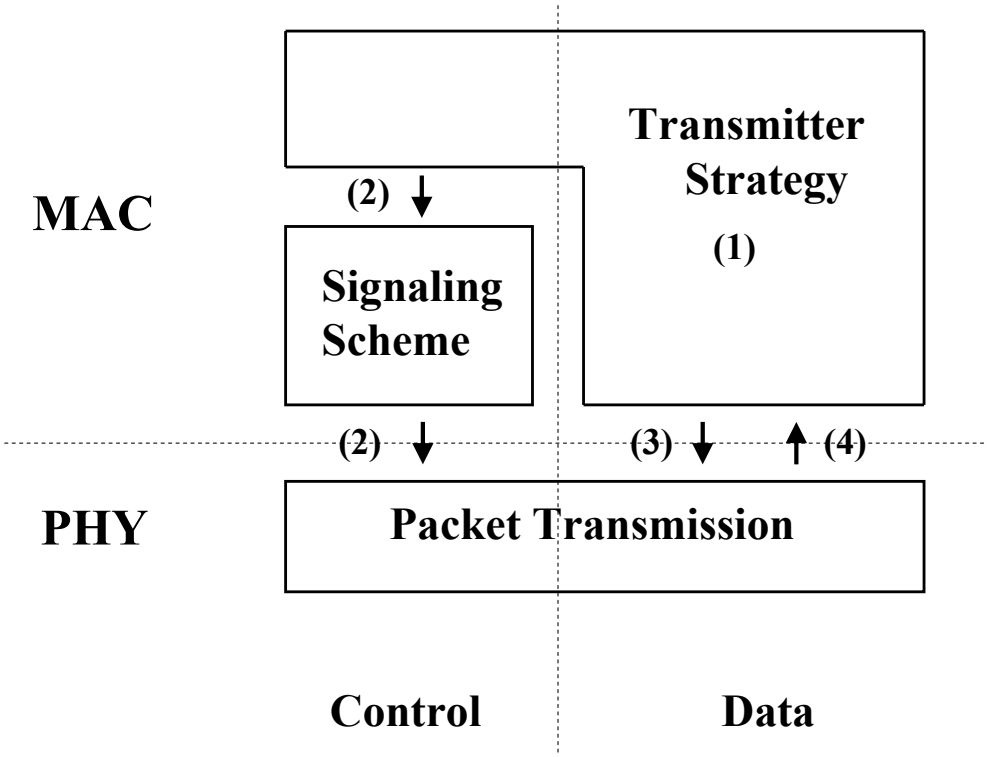

Figure 3.1: MAC protocol framework

3.3.2 Collision between Benign Users

Even when the spectrum is used very sparsely, a randomly selected frequency hopping pattern is likely to collide with the hopping patterns of other nodes. Channel coordination schemes are designed to reduce this inter-transmitter interference. When two nodes wish to use the same channel during the same time slot, they each determine which of the transmitters has priority, for example, based on the time at which each node claimed the channel. The node that does not have priority will not transmit at all, so the corresponding receiver will decode random data in this position. However, since the positions for these lost bits are known a priori, a sender mitigates the loss by using channel coding and forward error or erasure correction.

Two nodes may collide when neither node informed the other about its frequency hopping patterns, and, depending on the priority scheme, when only one of the two nodes disclosed its transmission intentions to the other. In the analysis, I assume that transmitter i has the highest priority for transmission, and is targeted by all the jammers. In other words, all nodes in the transmitter's recipient list will avoid interfering with the transmitter (I revisit this assumption in Section 3.8 and consider the case when all nodes have equal priority). Thus, increasing the number of benign transmitters in S reduces the number of potential interferers, increasing capacity.

3.3.3 Capacity Expression for the Framework

In this section, I mathematically derive the capacity of the system when all channels have equal gains and all users emit power uniformly across their chosen channel. Since legitimate users that receive the transmitter's channel coordination information will not interfere, the transmitter's capacity depends on its selection of recipient list S. Equation 3.1 can be simplified to:

$$\mathcal{R}(S) = W \cdot \log_2 \left[1 + \frac{P}{N_0 + \sum_{\ell \neq i, \ell \in (\mathcal{N}^c \cap S^c)} P_\ell + \sum_{k \in \mathcal{N}} J_k(S) \cdot P_c}\right]$$

where N_0 is the noise power in the channel, P is the transmitter's signal power, P_ℓ is the amount of user ℓ's power that interferes with the transmitter's signal, J_k is the jammer k's power normalized with respect to the power

constraint P_c, and

$$J_k(S) \leq \begin{cases} \frac{1}{C}, & \text{if } (S \cap \mathcal{N}) = \emptyset \\ 1, & \text{otherwise} \end{cases}$$

In the $(S \cap \mathcal{N}) = \emptyset$ case, the jammer does not receive the user's channel coordination information, and therefore can at best conduct wideband jamming across C channels, as described in Section 3.4.

If I further assume that legitimate users not in S emit at full power to maximize their own performance, then $E[P_\ell] = \frac{P_c}{C}, \forall \ell$, since there is a $\frac{1}{C}$ chance that any legitimate user not in S will interfere with the transmitter. Then, using Jensen's inequality, the expected capacity is bounded from below by:

$$\mathrm{E}[\mathcal{R}(S)] \geq W \cdot \log_2 \left[1 + \frac{P}{N_0 + |\mathcal{N}^c \cap S^c| \frac{P_c}{C} + \sum_{k \in \mathcal{N}} J_k(S) \cdot P_c}\right] \tag{3.5}$$

I use this expression in the analysis.

3.3.4 Transmitter Strategy

To make future recipient list decisions, each network user records historical performance data for each packet, including the recipient list used for that packet and the resulting performance. One natural choice for the recipient list is the set that has yielded the best average performance in the past. I call this the *Best so far* set and denote it with S_B: $S_B(t) = \underset{\sigma \in \{S(t'), \forall t' < t\}}{\operatorname{argmax}} \bar{\mathcal{R}}(\sigma)$, where $\bar{\mathcal{R}}$ is the time-average performance. When jammers jam at full power, the optimal set is the set that includes all legitimate nodes and excludes all jammer nodes; I denote this optimal set S^*. However, an attacker might choose not to jam when certain nodes are in the recipient list, so S^* may not have the maximum performance for a particular jammer strategy; however, the performance of S^* is optimal in the worst case. The scheme will converge to at least the performance of S^*, but if the attacker concedes better performance, the scheme can take advantage of the better-performing set.

In order to improve the Best so far recipient list, a sender must explore possible sets from time to time. To reach optimal performance, a transmitter strategy must eventually explore the optimal set. When I do not know the

jammer's strategy or the distribution for the number of jammers, the optimal set may be any set, because the attacker may choose to cease all jamming activities when a particular recipient list is chosen. Thus, to provide optimality against arbitrary attackers, a sender must be willing to explore all possible sets. In the framework, as well as SimpleMAC, convergence to the optimal set takes exponential time in the average case; however, I will show in Section 3.7 that SimpleMAC improves over the state-of-the-art within a single round in many cases, and fast convergence is not a goal of the SimpleMAC design.

3.4 Jammer Strategy Analysis

In this section, I assume that the attacker is purely adversarial, as described in Section 3.2.2. Attackers are capable of using a potentially nondeterministic, time-varying strategy to meet their goal of minimizing capacity. Since an attacker's strategy depends on whether or not it receives the channel coordination information, I study both cases.

3.4.1 Recipient List with No Jammer

If the recipient list S contains no jammers, then jammers do not learn any jamming-relevant information, and thus do not know which channel will be used for the user's transmission. This limits jammers to a much weaker attack, since they cannot use their compromised keys and gain no advantage from collusion. The only decision to be made in this case is whether to choose narrowband jamming or wideband jamming. I assume that a legitimate user i will uniformly choose any of the C channels, and I observe that $\mathcal{R}$ is a decreasing and convex function of $\widetilde{J}_k(f)$. By Jensen's inequality, the expected capacity $\mathrm{E}[\mathcal{R}]$, under the constraint of Equation 3.4, is minimized by choosing $\widetilde{J}_k(f) = \frac{P_c}{C \cdot W}$ for each jammer k. Thus, to minimize capacity, jammers will conduct wideband jamming when they do not know the frequency hopping pattern, but conduct narrowband jamming when they do have the information. In the analysis, I assume the jammer uses this strategy when it does not know the frequency hopping pattern.

3.4.2 Compromised Recipient List

I now analyze the jammer strategy when a jammer does receive the user's channel coordination information. In this scenario, jammers know where to concentrate their power to minimize transmitter capacity. However, in an infinite-horizon repeated game, jammers must also consider how their current action will affect future capacity. Equation 3.1 shows that jamming with higher power causes more interference and lowers capacity. However, since a user will avoid any set S that appears to contain jammers, jammers may not wish to strongly jam the transmission, hoping to reduce the user's suspicions that S contains a jammer. If the user converges on a new S_B that contains no jammers, then jammers can no longer influence capacity except by wideband jamming. A jammer may then want the Best so far set to include a jammer by abstaining from excessive jamming.

In the long run, the jammer knows that the transmitter will explore S^*, and will choose the best set. If the jammer allows another set S' to have better performance than S^*, then the transmitter will pick S', otherwise it will pick S^*. If the jammer's goal is to minimize capacity, they should not concede any additional long-run performance to the sender. Thus the sender will choose $S_B = S^*$, and the optimal jammer strategy will converge to full-power jamming.

Claim 1. *Given the general transmitter strategy in Section 3.3.4, jammer strategy converges to full power over time:*

$$\forall k \in \mathcal{N}, J_k(t) \to 1 \text{ as } t \to \infty$$

Proof. Proof is by contradiction. Let $J = \frac{1}{N}\sum_k J_k$ and t^* be the time when the legitimate user explores S^*. Suppose there exists an optimal jammer strategy $J(t)$ that does not converge to full-power over time: $\forall \epsilon > 0, \exists t > \epsilon, J(t) < 1$, yet yields minimum capacity. Since the legitimate user occasionally explores new recipient lists (as described in Section 3.3.4), it eventually explores S^* in finite time ($t^* < \infty$). Once the transmitter explores $S = S^*$, it will choose its Best so far recipient list S_B, so that the capacity performance is no worse than when $S = S^*$. Now let $\epsilon = t^*$ and compare J with a jamming strategy J' that jams with full power after t^*. $\forall t \geq \epsilon,\ J'(t) = 1$. In every time interval, J' results in performance at least as bad as J, be-

cause the recipient list in J is at least as good as S^*, and because the power used by J' is at least as high as that used by J. Because J' causes greater interference (power) at least once, J' results in lower performance than J. Therefore J is not an optimal strategy, establishing by contradiction that an optimal jammer strategy J must converge to full power over time. □

3.5 SimpleMAC

SimpleMAC protocol has two components: the Simple Transmitter Strategy (STS) and the Simple Signaling Scheme (SSS). Despite the simplicity of the schemes, from which SimpleMAC derives its name, SimpleMAC effectively combats intelligent attackers: it quickly outperforms the case where MAC protocol is disabled (which is the standard approach for securing MAC protocols) and has an easily analyzed optimal jammer strategy.

When selecting a recipient list, I determine the effectiveness of recipient list S by comparing the capacity when S is chosen as the recipient list to the capacity when no one knows the recipient list. In the latter case (i.e., when $S = \emptyset$), there is neither gain in capacity from legitimate nodes avoiding the transmitter nor loss in capacity from the jammers using the jamming-relevant information. Whenever the capacity is less than or equal to (with some error margin) the capacity when $S = \emptyset$, the transmitter chooses a new set S before the next transmission, because the current set S provides no advantage over $S = \emptyset$.

SimpleMAC does not try to infer which nodes are jammers and which ones are not; rather, it directly uses channel feedback to determine which recipient lists result in good performance. For example, when node A shares its information with a jammer (but does not cause interference itself), any recipient list with node A in it will have decreased performance, so the STS will avoid such list. Similarly, if node A jams only when node B is also in the recipient list, the STS will avoid lists that contain both A and B. Because SimpleMAC makes the recipient list decisions based on actual performance and not behavior, SimpleMAC is immune to collusion.

3.5.1 Simple Transmitter Strategy

In the STS, for each transmission t, a legitimate user has three options when choosing a recipient list S:

1. *Best so far* (B): the set with best average performance among explored sets, as described in Section 3.3.4.

2. *Randomly explore* (R): chosen uniformly at random among all possible sets.

3. *Empty set* (E): $S(t) = \emptyset$.

The transmitter always chooses one of these three strategies. The *Best so far* action, $S(t) = S_B$ corresponds to choosing the set that yielded the highest average capacity among all the recipient lists that have been tried through time $t-1$, which guarantees performance at least as good as $S = \emptyset$, since $S = \emptyset$ has been tried earlier. If jammers jam with sufficient power ($\sum_{k \in \mathcal{N}} J > \frac{T}{C}$), then the set S that yields the highest capacity is S^*, the set that contains all the legitimate users and excludes all the attackers. In this case, when the user explores sets occasionally (so that the user eventually visits all possible sets with probability one), the Best so far set S_B converges to S^*, since the probability that S^* has been previously chosen approaches one. The user chooses the *Randomly explore* action to search for a set that yields higher capacity than the previous Best so far set. Once such a set is found, the set S_R becomes the new Best so far until the node discovers another set that yields even higher capacity. The more often the user chooses to explore a random set, the more quickly S_B converges to S^*. The *Empty set* action establishes baseline performance during each time interval, so that slow time-variance in channel conditions do not bias set selection.

The STS operates in rounds. For each transmission t within round r, the user makes an independent random choice among the three options. The probabilities may vary with r, so that in expectation, round r contains $B(r)$ transmissions with the Best so far recipient list, $R(r)$ transmissions on a randomly chosen recipient list, and $E(r)$ transmissions using an empty recipient list. Round r lasts for $B(r) + R(r) + E(r)$ transmissions, and I do not rely on the secrecy of $B(r)$, $E(r)$, $R(r)$.

In order to converge to the optimal performance, I explore a user strategy where the user uses the Best so far set more and more often, while occasionally using Randomly explore and Empty set. One such user strategy

is:

$$B(r) = r^{\delta}, R(r) = 1, E(r) = 1, \forall r \tag{3.6}$$

In order to converge to the optimal performance for S^*, δ needs to be positive. A higher δ corresponds to more aggressive search for a better Best so far set and thus quicker convergence to S^*.

3.5.2 Modified and Hybrid Simple Transmitter Strategy

STS uses a uniform distribution for choosing a recipient list to explore new sets. Although the protocol quickly finds a recipient list that only contain legitimate users and thus outperforms $S = \emptyset$ (it takes 2^N rounds in expectation), it takes a long time to search for the optimal set S^* for large T (it takes 2^T rounds in expectation).

To improve the convergence rate to the optimal performance of using $S = S^*$, I develop the Modified Simple Transmitter Strategy (MSTS). MSTS is a modified version of STS where I use a deterministic exploration of recipient lists, as opposed to a random strategy for exploring. In particular, I use a brute-force approach searching large sets before smaller sets. Given the number of users T, it first tries $S = \mathcal{T}$ (i.e., broadcast to all entities in the network), then explores all possible sets that have $T - 1$ users, and then move on to sets that have $T - 2$ users, and so on. Compared to STS, MSTS is guaranteed to find S^* in $\sum_{i=0}^{N} \binom{T}{i} = O(T^N)$ rounds and quickly converges to the optimal performance for large T. However, MSTS does not find a jammer-free recipient list for the first $\sum_{i=0}^{N-1} \binom{T}{i}$ rounds, during which period, it has no gain from channel coordination. Therefore, STS and MSTS have a tradeoff between the rate of convergence to $S = S^*$ and convergence to improvement over $S = \emptyset$. Finally, I define Hybrid Simple Transmitter Strategy (HSTS), which interleaves the exploration approaches of the STS and the MSTS. In particular, in each exploration stage, I alternate between the random exploration strategy of STS and the deterministic strategy of MSTS.

3.5.3 Simple Signaling Scheme

In order to send a control message to exactly those nodes in a recipient list, SimpleMAC needs a signaling scheme that provides confidentiality against malicious entities and reliability in the presence of jamming. I make no attempt to design an efficient signaling scheme; because the overhead of a control message is amortized over the data frame, and because the wireless user can choose arbitrarily long data frames, the system can reach near-optimal overall protocol performance even with an extremely inefficient signaling scheme. Thus, SSS simply unicasts the control messages to each recipient in the recipient list. SSS provides confidentiality by encrypting messages with a symmetric key, and availability by using direct sequence spread spectrum (DSSS) using a chip sequence known only to the sender and receiver. In a 50-node network with 20-byte reservation messages (consisting of source address, destination address, and a seed for the hopping pattern), if each reservation covers 100 kB of data (for example, 66 packets each 1500-bytes long), SSS incurs an overhead of not more than 1%, and average overhead of 0.5%. Though the data rate may be higher than the control rate due to the use of DSSS for the control message, SSS can continue to keep the overhead low by covering more data with each control message, or by replacing repeated-unicast with a jamming-resilient broadcast protocol, of which several have been proposed [24–26]. SSS simply requires that each node has a pairwise shared key with every other node. Such keys can be established through Diffie-Hellman exchanges over a jamming-resilient broadcast protocol.

3.6 SimpleMAC Theoretical Analysis

3.6.1 Jammer Reaction to SimpleMAC

In Section 3.4, I studied the attacker strategy under the general framework and showed that the optimal attacker strategy converges to full power, even though an attacker may wish to avoid detection so that the legitimate user will use a compromised recipient list. In this section, I claim that against the STS, optimal jammers *jam at full power all the time*. The claim holds

because, unlike the user's selection of recipient lists, the user's choice of action (Best so far, Randomly explore, or Empty set) does not adapt to jammer strategy. Intuitively, the sender forms a partial order on recipient lists based on their past performance. An attacker that does not jam at full power can jam at a higher power level and yet maintain the same partial order of recipient lists (or a functional equivalent), which shows that any strategy that does not jam at full power cannot be optimal.

Claim 2. *Against the STS, the best jammer strategy is to emit at full power all the time, i.e.,*

$$\forall t, \forall k \in \mathcal{N}, J_k(S) = \begin{cases} \frac{1}{C}, & \textit{if } (S \cap \mathcal{N}) = \emptyset \\ 1, & \textit{else} \end{cases}$$

Proof. Proof is by contradiction. Suppose that an optimal jammer strategy $J = (J_1, J_2, ..., J_N)$ does not jam at full power at some time; I let Γ be the set of times at which J does not use full power. I now show that there exists a different jammer strategy J' that yields less capacity than J while preserving the legitimate user strategy. To find such J', I assume perfect knowledge about the recipient list S. (This does *not* mean that a jammer needs perfect information; rather, it shows that even a jammer with perfect information will still choose the simple strategy of full-power jamming, and therefore *any* attacker should do the same.) J' will only diverge from J when J does not emit at full power,

At time $t' \in \Gamma$, let the two best previously measured recipient lists be A and C, where A is the best and C is the second-best. Then either $S = A$ yields higher capacity than $S = C$ or both sets $S = A$ and $S = C$ yield the same performance. I study the two cases separately:

(i) If $\mathcal{R}(A) > \mathcal{R}(C)$, then pick $J'(t')$ such that $\mathcal{R}_{J'}(A) = \frac{\mathcal{R}_J(A)+\mathcal{R}_J(C)}{2}$. This choice preserves the performance order of recipient lists and thus does not change the user's choice of recipient list.

(ii) If $\mathcal{R}(A) = \mathcal{R}(C)$, then pick $J'(t')$ such that its corresponding performance is ϵ smaller than that of $J(t')$ for small ϵ. This breaks the tie between A and C since $A \neq C$. Though this changes the legitimate user's choice of recipient list (because the legitimate user will choose A over C for the *Best so far* set), the legitimate user strategy when jammer picks J' is functionally equivalent to the legitimate user strategy when jammer picks J (because

when the jammer picks J, it does not matter whether the user chooses A or C). Therefore, J' yields smaller capacity than J.

Since, in both cases J' yields lower capacity than J while preserving the order of recipient lists (and thus preserving the legitimate user strategy or its equivalent), J is not optimal and there is a contradiction. □

3.6.2 Performance Analysis

Under the STS, a legitimate user chooses the recipient list from among three options: Best so far (B), Randomly explore (R), and Empty set (E). I use Equation 3.1 to determine the expected capacity. Since attackers in S jam at full power, as shown in Section 3.6.1, $S = S^*$ (including no jammers but all other legitimate users) yields the optimal performance. Thus, the expected capacity varies with time (in units of rounds r) until it reaches the steady state where $S_B = S^*$. The steady-state expected capacity is shown in Equation 3.9, where α denotes the number of legitimate users outside S (who could potentially cause interference to the transmitter) and $\binom{a}{b} = \frac{a!}{b!(b-a)!}$.

For the transient expected capacity, $\mathrm{E}[\mathcal{R}|Random]$ and $\mathrm{E}[\mathcal{R}|Empty]$ are constant in time, whereas the expected capacity for Best so far varies with time. The user chooses $S_B = \emptyset$ at round r if all the previously explored sets contain jammers; otherwise, he chooses the set that contains no jammer and the most nodes (minimizing α). The term $\mathrm{E}[\mathcal{R}(r)|\mathrm{Best}]$ for the r^{th} round is expressed in Equation 3.10 where β corresponds to the number of times that the user found a jammer-free set, and $\mathcal{B}_i$ are independent binomial random variables with probability 0.5 ($p = 0.5$) and $T - N$ trials ($n = T - N$), since $T - N$ is the number of protocol-compliant users.

3.7 SimpleMAC Evaluation

In earlier sections, I have analyzed SimpleMAC theoretically, in a manner that is general and not limited to any particular system design. In this section, I evaluate SimpleMAC in practice both using MATLAB simulations and a testbed implementation on the WARP software radio platform [27]. As described in Section 3.2.1, I use SINR as the metric in this section both

$$
\begin{aligned}
\mathrm{E}[\mathcal{R}]_{\mathrm{SS}} &= \lim_{r\to\infty} \mathrm{E}[\mathcal{R}(r)] \\
&= \Pr[\text{Best}] \cdot \mathrm{E}[\mathcal{R}|\text{Best}] \\
&\quad +\Pr[\text{Random}] \cdot \mathrm{E}[\mathcal{R}|\text{Random}] \\
&\quad +\Pr[\text{Empty}] \cdot \mathrm{E}[\mathcal{R}|\text{Empty}] \\
&= \frac{B}{B+R+E} \cdot W \log_2\left(1+\frac{P}{N_0+\frac{N}{C}}\right) \\
&+ \frac{R}{B+R+E}\frac{1}{2^N}\sum_{\alpha=0}^{T-N}\left[\frac{\binom{T-N}{\alpha}}{\sum_{\gamma=0}^{T-N}\binom{T-N}{\gamma}}\right. \\
&\quad \left.\cdot W \log_2\left(1+\frac{P}{N_0+\frac{N}{C}+\frac{\alpha}{C}}\right)\right] \\
&+ \frac{R}{B+R+E}\frac{2^N-1}{2^N}\sum_{\alpha=0}^{T-N}\left[\frac{\binom{T-N}{\alpha}}{\sum_{\gamma=0}^{T-N}\binom{T-N}{\gamma}}\right. \\
&\quad \left.\cdot W \log_2\left(1+\frac{P}{N_0+N+\frac{\alpha}{C}}\right)\right] \\
&+ \frac{E}{B+R+E} \cdot W \log 2\left(1+\frac{P}{N_0+\frac{T-1}{C}}\right) \\
&= \frac{B}{B+R+E} \cdot W \log_2\left(1+\frac{P}{N_0+\frac{N}{C}}\right) \\
&+ \frac{R}{B+R+E}\sum_{\alpha=0}^{T-N}\left\{\frac{\binom{T-N}{\alpha}}{2^{T-N}}\right. \\
&\quad \cdot\left[\frac{1}{2^N}\cdot W \log_2\left(1+\frac{P}{N_0+\frac{N}{C}+\frac{\alpha}{C}}\right)\right. \\
&\quad \left.\left.+\frac{2^N-1}{2^N}\cdot W \log_2\left(1+\frac{P}{N_0+N+\frac{\alpha}{C}}\right)\right]\right\} \\
&+ \frac{E}{B+R+E} \cdot W \log_2\left(1+\frac{P}{N_0+\frac{T}{C}}\right)
\end{aligned}
\tag{3.9}
$$

$$
\begin{aligned}
\mathrm{E}[\mathcal{R}(r)|\mathrm{Best}] &= \Pr[\forall r' < r, S_R(r') \cap \mathcal{N} \neq \emptyset] \cdot \mathrm{E}[\mathcal{R}]_{S=\emptyset} + \Pr[\exists r' < r, S_R(r') \cap \mathcal{N} = \emptyset] \\
&\quad \cdot \sum_{\alpha=0}^{T-N} \Pr[|S| = T - N - \alpha] \cdot \mathrm{E}[\mathcal{R}]_{|S|=T-N-\alpha} \\
&= (1-2^{-N})^r \cdot W \cdot \log_2\left(1 + \frac{P}{N_0 + \frac{T}{C}P}\right) \;+\; \sum_{\beta=1}^{r} \binom{r}{\beta} (1-2^{-N})^{r-\beta} (2^{-N})^{\beta} \\
&\quad \cdot \left[\sum_{\alpha=0}^{T-N} \Pr[\min_{i=1\ldots\beta} \mathcal{B}_i = \alpha] \cdot W \cdot \log_2\left(1 + \frac{P}{N_0 + \frac{N+\alpha}{C}P}\right)\right] \\
&= (1-2^{-N})^r \cdot W \cdot \log_2\left(1 + \frac{P}{N_0 + \frac{T}{C}P}\right) \;+\; \sum_{\beta=1}^{r} \left\{ \binom{r}{\beta} (1-2^{-N})^{(r-\beta)} 2^{-T\beta} \right. \\
&\quad \left. \cdot \sum_{\alpha=0}^{T-N} \left[\left\{ \left(\sum_{\sigma=\alpha}^{T-N} \binom{T-N}{\sigma} \right)^{\beta} - \left(\sum_{\sigma=\alpha+1}^{T-N} \binom{T-N}{\sigma} \right)^{\beta} \right\} \cdot W \cdot \log_2\left(1 + \frac{P}{N_0 + \frac{N+\alpha}{C}P}\right) \right] \right\}
\end{aligned}
\tag{3.10}
$$

because capacity is strictly monotone in SINR and because I can evaluate SINR improvements without needing to make perfect modulation and coding choices that are necessary to achieve channel capacity.

3.7.1 Methodology and Metric

I built SimpleMAC implementation on the WARP (Wireless Open-Access Research) software-defined radio platform. I used four WARP boards: one acting as the source, one acting as the receiver, and the other two acting as co-existing transmitters. By using the MIMO capabilities of the boards, I built an environment consisting of one source, one receiver, and four other transmitters ($T = 4$), one of which is a jammer ($N = 1$). I divided the spectrum into five channels of equal bandwidth[1] ($C = 5$). Also, I manually calibrated the antenna locations so that the receiver observes approximately the same power from each transmitter.

For the purposes of the evaluation, I filled the queues at each node so that each transmitter transmits packets all the time. This is not a requirement of SimpleMAC; because the recipient list performance estimates will not be updated during periods without traffic, traffic is always present from the

[1]The evaluations focus on scenarios with relatively few channels; this is not a limitation of SimpleMAC, but is a performance optimization. SimpleMAC can improve performance regardless of the number of channels, but the optimal number of channels tends to be small relative to the number of transmitters, because from a capacity perspective, it is much better to have a legitimate-to-legitimate node collision than to let spectrum go unused.

perspective of the protocol. In fact, SimpleMAC works even in dynamic environments where the jammer and competing transmitters are sometimes present and sometimes absent; in Section 3.7.6, I show that SimpleMAC works even better in mobile environments.

At the physical layer, the system modulates data using differential quadrature phase-shift keying (DQPSK), and synchronize using a preamble which is a Barker sequence modulated using binary phase-shift keying (BPSK). The entire 12 MHz-wide spectrum is divided (centered at 2.452 GHz) into 300 OFDM subcarriers, so each channel contains 60 subcarriers. I send the control communication across the entire band (300 subcarriers). The frequency hopping scheme is to split each data message into frames of 60 symbols, which I simultaneously send on each of 60 subcarriers in the chosen channel. The user hops from channel to channel between frames.

The source transmitter sends random symbols to the receiver, and I observe the decoded symbols at the receiver. I compare these symbols to determine the error rate, and use that error rate to estimate the signal-to-interference-and-noise ratio (SINR) at the receiver. The expected bit error rate ($\overline{\text{BER}}$) and the expected SINR at the receiver ($\overline{\text{SINR}}$) have the following relationship for DQPSK modulation [28, 29]: $\overline{\text{BER}} = \frac{1}{2}\left(1 - \frac{\sqrt{2 \cdot \overline{\text{SINR}}}}{\sqrt{1+4 \cdot \overline{\text{SINR}}+2 \cdot \overline{\text{SINR}}^2}}\right)$.

I also validated the results using a MATLAB-based simulation. The simulator works on a per-packet basis: for each time slot, each transmitter chooses a recipient list according to the STS, and the channel selection according to a uniform random distribution. The channel model is an independent, identically distributed Rayleigh fading channel with AWGN noise. I then compute the number of interfering users (legitimate and jammer) and calculate the resultant SINR, which I then use as feedback for the next round.

To analyze the performance of STS in the implementation and simulation environments, I use the $S = \emptyset$ performance (corresponding to the no channel coordination) as the reference and study the performance gain over $S = \emptyset$. This gain represents the improvement over a protocol that does not reserve a channel prior to data communication. I study the *SINR gain* which is the SINR observed by the STS divided by the SINR when $S = \emptyset$. As shown in Equation 3.1, the instantaneous channel capacity (which I use in the theoretical analysis) is strictly monotone in instantaneous SINR, and assuming the optimal fixed strategy for jammers, this relationship extends

to the time-average SINR and the time-average capacity in Equation 3.3.

3.7.2 Without Attack

I first consider the performance of the protocol when all transmitters are protocol-compliant. In this scenario, the protocol minimizes unintentional interference, and the more nodes the recipient list includes, the better the performance. Figure 3.2(a) shows the estimated SINR and reflects a 14.1 dB SINR increase between $|S| = 0$ (no coordination) and $|S| = 3$ (full coordination). The performance under full coordination gives us an estimate of the SNR without interference (when one channel is used): $\overline{\text{SNR}} \approx 20\text{dB} = 100$.

3.7.3 Under Attack

I now consider protocol performance under jamming. If the recipient list is compromised and contains the jammer, the jammer can effectively jam the transmitter by following its hopping pattern. Otherwise, the jammer can either choose to jam across all five channels at reduced power per channel, or on a random channel at full power. Figure 3.2(a) displays the expected SINR at the receiver in each of the three cases. Including more legitimate users in the recipient list yields better SINR, as described in Section 3.7.2. I also observe a drastic drop in performance when S is compromised; whenever a set contains a jammer, its SINR is below 0 dB, since the jamming power is equal to the signal power, and other legitimate nodes may accidentally interfere. (When there is perfect coordination among legitimate nodes, the only additional noise is the receiver's thermal noise, so the SINR is very close to 0 dB in this case.) These cases with compromised recipient lists thus all perform worse than when disabling channel coordination ($S = \emptyset$), which provides 3.72 dB SINR. Furthermore, wideband jamming is more effective and yields lower SINR for the target transmitter than narrowband jamming, verifying the theoretical analysis in Section 3.4.1.

Despite the risk of possibly choosing the jammer, channel coordination is still potentially advantageous. Choosing a random recipient list $S = S_R$ has expected performance better than $S = \emptyset$ in expectation (computation shows an SINR gain of about 1 dB assuming wideband jamming for uncompromised

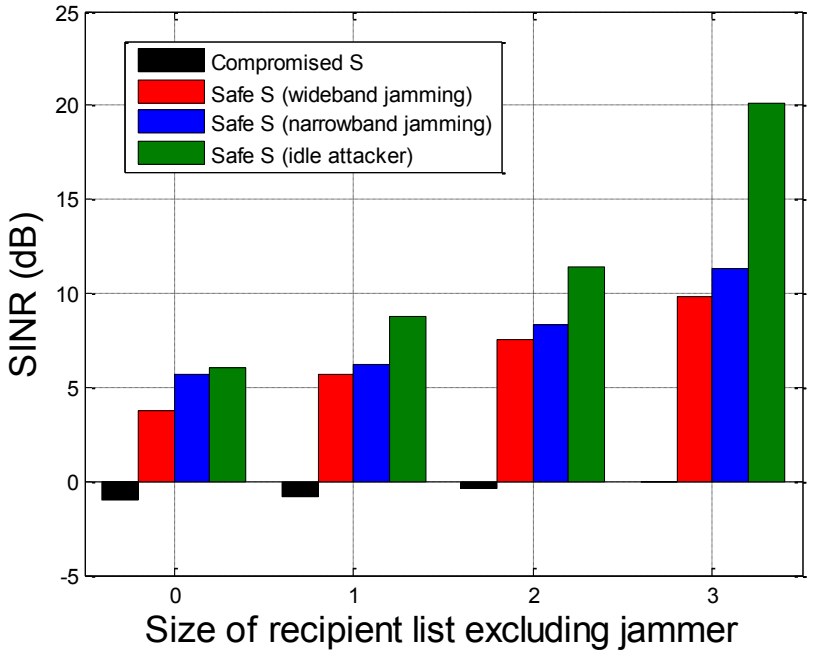

(a) Performance for S

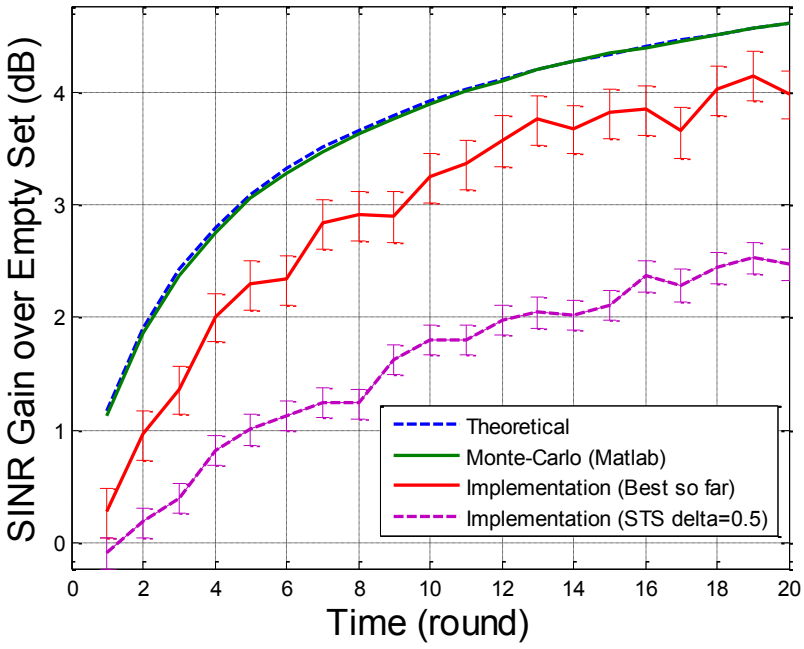

(b) STS performance with time

Figure 3.2: SimpleMAC STS performance

S). Furthermore, once the STS converges to the best possible set, I can reach an SINR of about 9.69 dB in spite of the wideband jamming, which reflects an SINR gain of 5.97 dB over the baseline performance of $S = \emptyset$.

3.7.4 Data Communication Using the STS

Now that I have established the performance of known-good and known-bad sets, I study the performance of the STS and explore its convergence behavior. For each round, the STS performs three actions (B, R, E) as described in Section 3.5. In the evaluation, the jammer uses the optimal strategy (full-power jamming using all available information).

Because the metric is SINR gain, and the baseline performance is the empty set, the Empty set has performance of 0 dB. The Randomly explore action chooses a recipient list at random with uniform probability. Therefore, the performance of Randomly explore is independent and has constant expectation across time. Assuming that the user randomly explores at least once per round, the Best so far performance is increasing in time, and converges to the optimal steady-state performance where $S = S^*$, as more sets are explored and the user has more sets from which to choose S_B.

In Figure 3.2(b), I plot the performance of the Best so far strategy under three evaluation environments: the theoretical analysis corresponding to Equation 3.10, the simulation, and the testbed implementation. For the implementation results, I also plot 95% confidence intervals, which are not shown for the simulation results because the simulation results included enough runs that the confidence intervals would not be visible. The results show that the performance predicted by the theoretical analysis coincides with the simulations. The implementation performance is worse than the theoretical and simulation results because the simulation assumes a perfect measurement of SNR, whereas the implementation infers it from the bit error rate; early in the run, when the number of observed bits is small, the BER measurement can deviate from the expected BER, and the STS may as a result make suboptimal choices. However, in later rounds, this performance difference decreases as the implementation gains better information. As a result, the maximum performance difference of 23% (of implementation performance) occurs at round one and decreases to 16% at round 20

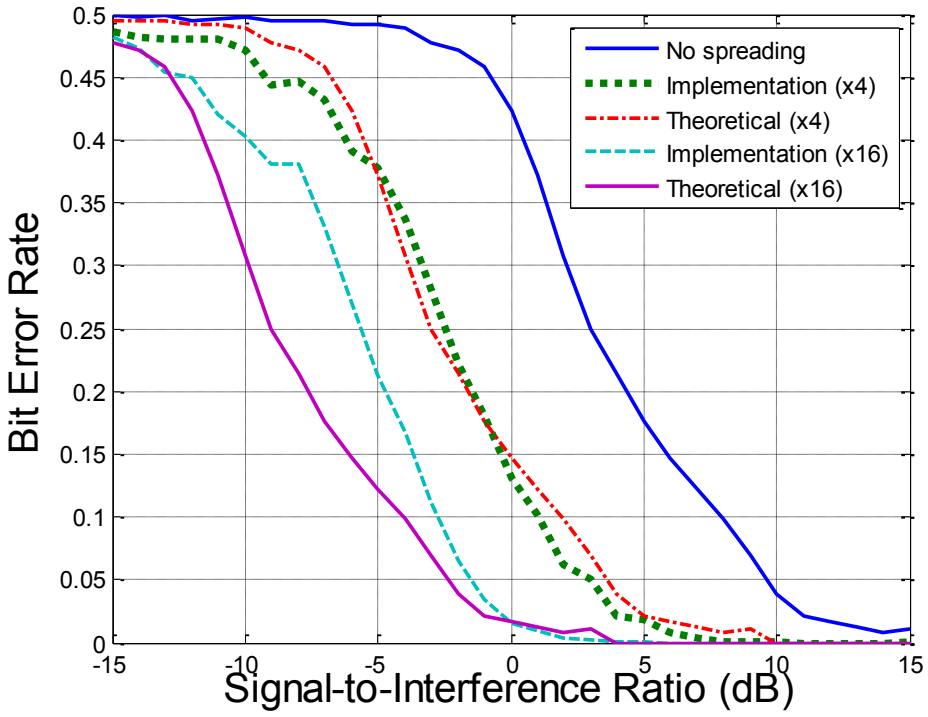

Figure 3.3: SimpleMAC SSS performance

in Figure 3.2(b). I also show the overall (as opposed to Best so far only) performance of STS under the implementation for $\delta = 0.5$. The performance of STS overall (including Empty and Random transmission sets) is monotonically increasing in time and begins outperforming the no-channel-coordination option after round one. I will later show that the STS converges to the Best so far performance in Section 3.7.6.

3.7.5 Control Communication Using the SSS

SimpleMAC relies on a robust signaling scheme that can reach each recipient on the recipient list. I implemented Simple Signaling Scheme in WARP using wideband communication and Direct Sequence Spread Spectrum (DSSS), as described in Section 3.5.3; the decoder is based on an analog correlator. In order to study the effect of spreading gain, I sent messages using three different code lengths: 1 (no spreading and thus, no redundancy), 4, and 16. I sent each chip on a separate, adjacent OFDM subcarrier. To get various signal-to-interference ratios, I fixed the signal power and varied the

interference power to obtain signal-to-interference ratios ranging from -15 dB to 15 dB. Figure 3.3 shows the relationship between the BER performance and the signal-to-interference ratio.

To show the effectiveness of spreading to avoid interference in SSS, I generated theoretical curves by shifting the "no spreading" result by the expected spreading gain. As expected, the implementation result aligned well with the theoretical. When the code length is small, the performance can be slightly better than the theoretical because the x-axis considers only interference but not noise; the processing gain filters both interference and noise, which means that the implementation slightly outperforms the shifted curve. For example, when the code length is 4, the implementation has better performance than the theoretical except between -5 and -1 dB. However, with increasing code length, I use an increasing number of adjacent subcarriers, increasing inter-carrier interference and degrading performance.

3.7.6 Simulations with Mobility

I have previously shown that without mobility, STS is effective. Using the simulator, I now show that introducing mobility makes STS *more effective*; that is, transmitter performance increases when it is surrounded with mobile users. In the mobile environment, at the beginning of each round, each mobile user is placed at a random position, and I compute the channel gains based on the positions and a path loss model (with path loss exponent 3) and Rayleigh fading. I fix the receiver's location and choose the position of each mobile user with a distribution such that the expected received power corresponds to unit channel gain; that is, the expected received power is the same in the mobile and static cases. I use the same parameter values as I did previously: $C = 5$, $T = 4$ (other than the source transmitter), $N = 1$, SNR (across the entire spectrum) = 13 dB. To allow for comparison between scenarios, I normalize all performance to the SINR achieved using a $S = \emptyset$ in the *static case*. I also ran for an increased number of rounds (10^4). I plot performance on a semi-log scale to better show the dynamics of convergence.

In each case, I consider *convergence to improvement*, which shows how quickly the performance exceeds the no-channel-coordination case, i.e., $S = \emptyset$, and *approximate convergence*, which shows when the system reaches within

10% of steady-state performance.

For comparison, the static case, where no user moves, is shown in Figure 3.4(a) (and reflects the data shown in Figure 3.2(b)). The scheme converges to improvement instantly at round one, since the expected SINR for Random set exceeds one (as I discussed in Section 3.7.3), and achieves approximate convergence in about 32 rounds and 316 rounds, respectively, for $\delta = 1$ and $\delta = 0.5$, for an eventual performance gain of 5.31 dB or 3.4.

I expect that mobility will improve performance *for any set* S because of Jensen's inequality. Using Equation 3.2, I observe that SINR is convex in channel gain γ of other users, so static channel gains are worse than random channel gains when the expected channel gain is equal. Adding mobility thus improves both SINR for any set S, *including* $S = \emptyset$. As more nodes become mobile (0 in Figure 3.4(a), 1 in Figure 3.4(c), 4 in Figure 3.4(b), and 5 in Figure 3.4(d)), the Empty set $\overline{\text{SINR}}$ performance monotonically increases. The values for $\overline{\text{SINR}}_{\text{Empty}}$ are shown in Figure 3.4.

To show the source of the additional SINR gain, I performed two additional set of experiments, one in which only legitimate users were mobile (Figure 3.4(b)) and one in which only jammers were mobile (Figure 3.4(c)). Adding mobility for legitimate users, as shown in Figure 3.4(b), has minimal impact on the steady-state performance $\overline{\text{SINR}}_{\text{Best}}$, because they avoid collision and create no interference in steady-state where $S_B = S^*$. However, when legitimate users move, convergence is much faster. In particular, *approximate convergence* occurs around 16 rounds and 188 rounds, compared to the static case of 32 rounds and 316 rounds, for $\delta = 1$ and $\delta = 0.5$, respectively. This increased convergence speed is because the jammers' noise contribution is more consistent and therefore has a larger impact on performance, making jammers more easily identifiable.

On the other hand, jammer mobility, as shown in Figure 3.4(c), results in a substantial improvement of $\overline{\text{SINR}}_{\text{Best}}$. When $S = \emptyset$, the majority of noise comes from legitimate nodes, so jammer has limited impact. However, after many rounds, S_B approaches S^*, so jammers' interference comprises most of the noise, so $\overline{\text{SINR}}_{\text{Best}}$ improves more than $\overline{\text{SINR}}_{\text{Empty}}$.

When all nodes other than the source transmitter are mobile (Figure 3.4(d)), I get benefits from both user mobility and jammer mobility. leading to faster convergence and better SINR performance for all STS action choices of B,R,E as compared to the stationary scenario. The steady-state $\overline{\text{SINR}}_{\text{Best}}$ values

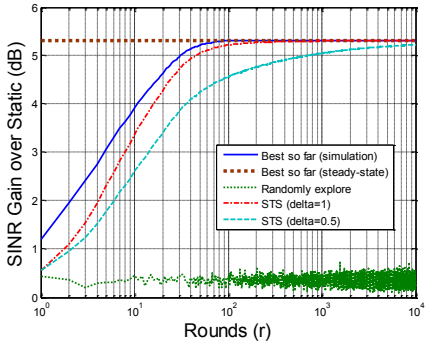

(a) All users are static. $\overline{\text{SINR}}_{\text{Empty}} = 1.26$

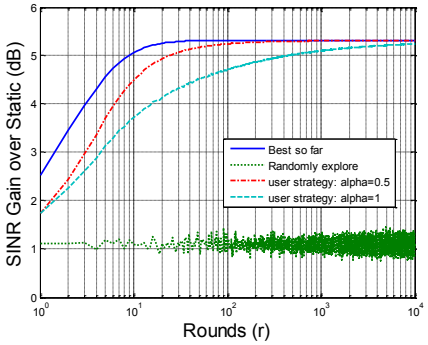

(b) Only legitimate users are mobile. $\overline{\text{SINR}}_{\text{Empty}} = 2.20$

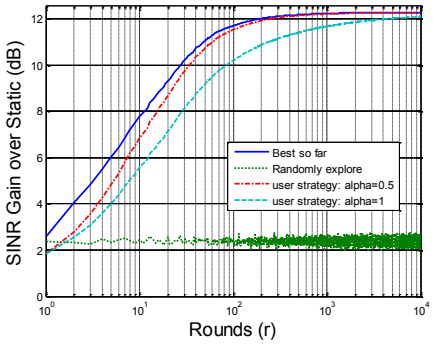

(c) Only jammers are mobile. $\overline{\text{SINR}}_{\text{Empty}} = 1.45$

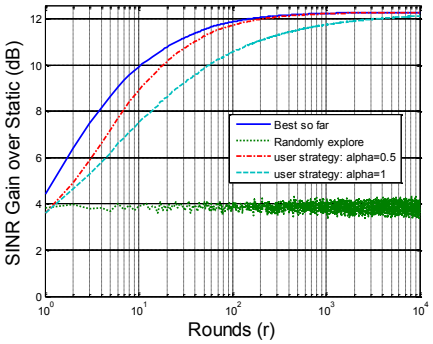

(d) All users are mobile. $\overline{\text{SINR}}_{\text{Empty}} = 3.33$

Figure 3.4: SimpleMAC simulations with mobility (gain is relative to static-case $S = \emptyset$ performance)

show the mobile case having 7 dB improvement, or about five times as much improvement as, compared the static case, which yields 95% capacity improvement over the static case. Thus SimpleMAC performs even better in mobile environments than in stationary ones.

3.7.7 MSTS and HSTS

As the network size grows, STS performance converges more slowly. In specific, convergence to improvement over $S = \emptyset$ and convergence to the optimal performance of $S = S^*$ occur in 2^N and 2^T rounds, respectively, in expectation. In this section, I study the convergence dynamics and demonstrate the effectiveness of Modified Simple Transmitter Strategy (MSTS) and Hy-

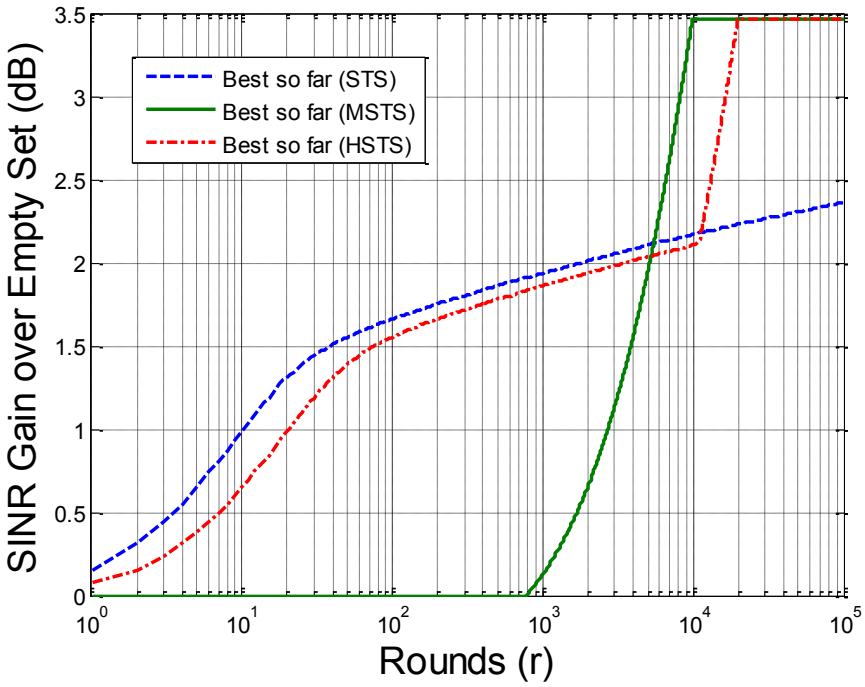

Figure 3.5: SimpleMAC STS and MSTS comparison

brid Simple Transmitter Strategy (HSTS), discussed in Section 3.5.2. To demonstrate that the schemes work for larger networks, the simulator uses the parameter set of $C = 120$, $T = 40$, $N = 3$ and is deployed in a static environment without mobility. Figure 3.5 compares the convergence of STS, MSTS, and HSTS and plots the expected SINR gain of the best so far recipient list. MSTS achieves steady-state performance much quicker than STS (S^* is guaranteed to be found in $\sum_{i=0}^{3} \binom{40}{i} = 10,701$ rounds) but achieves slower rate for convergence to improvement (the transmitter does not find a jammer-free recipient list before round $821 = \sum_{i=0}^{2} \binom{40}{i}$). HSTS interleaves the exploration approaches of STS and MSTS in each round and simultaneously provides fast convergence to $S = S^*$ (within $2 \cdot 10,701 = 21,402$ rounds) and fast convergence to improvement (one round in expectation).

3.8 Alternative Transmission Priorities

In Section 3.3.2, I considered a single user that has the highest priority for transmission on the channel, so all other legitimate users will avoid that node's transmissions, and I calculate the performance improvement of that node. In this section, I consider nodes that have equal priority. At the steady-state where $S_B = S^*$ for all legitimate users, and where all users have equal priority, each user will defer a $1 - \sigma$ fraction of its slots, where

$$\sigma = \sum_{k=0}^{T-N} \Pr_k \cdot \frac{1}{k+1}$$

where $\Pr_k$ is the probability that k other users are transmitting on the same channel and has a binomial distribution with parameters of $T - N$ trials and $\frac{1}{C}$ probability. To evaluate SimpleMAC's performance in this egalitarian regime, I multiply the previous performance measurements by a fraction of σ. Since the simulations with mobility showed a capacity gain of 112% for a single priority node (Section 4.6), the capacity gain is 56% when all nodes have equal priority. Similarly, without mobility, the capacity gain over baseline strategy is 51% with equal priority among all nodes.

Transmission priorities also affect how a node performs with a malicious receiver. If a malicious receiver forces a legitimate transmitter A to send with $S = \emptyset$, then the probability that any other node collides with transmitter A increases. In the worst case scenario, if transmitter A has absolute priority and is always transmitting, then A operates as a narrowband jammer. However, I can also consider the class of transmission priority schemes in which reserved channels always take priority over unreserved channels; under such schemes, in the steady-state, transmitter A will always defer to transmissions of other nodes, actually *increasing* the performance of other nodes. A more complete analysis of transmission priority schemes and malicious receivers is beyond the scope of this dissertation.

3.9 Chapter Summary

SimpleMAC is a MAC protocol that provides effective channel coordination to minimize interference among coexisting transmitters while simultaneously

resisting jammers that use channel coordination information to jam more effectively. SimpleMAC avoids control channel jamming and limits jamming-relevant information to a recipient list, adjusting the recipient list to optimize performance. SimpleMAC converges to the optimal performance and forces an optimal jammer to always jam at full power. I used a game-theoretical approach to counter intelligent attackers, and analyzed the effectiveness of the SimpleMAC scheme through theory, simulation, and implementation, and observed over 570% increases in SINR and over 50% increases in channel capacity gains in a realistic mobile environment.

CHAPTER 4

IGNORE-FALSE-RESERVATION MAC

4.1 Chapter Overview

In wireless networks, many MAC protocols allow a node to make a channel reservation to reduce the probability that its transmission will collide with another user's transmission. An adversary can launch a two-pronged Denial-of-Service attack against such MAC protocols: first, it can send excessive reservation requests (to consume network resources with minimal attacker resources), and second, it can jam those channels that it has not reserved. Since normal users do not attempt to transmit on channels reserved by others, the attacker can attack with greater power efficiency, reducing the number of channels it needs to jam by using reservations to cover a subset of the channels.

The *false reservation attack* reserves channel bandwidth without the intention of using it and is a power-efficient way to deny service to legitimate nodes. To best understand the effect of the false reservation attack, I consider a *multi-channel environment* with *power-limited attackers* (however, the work can also be applied to single-channel time division multiple access (TDMA) systems with energy-limited attackers). In a multi-channel environment, the frequency spectrum is divided into multiple channels and each user competes for bandwidth on one channel at a time; the framework considers channels that have flexible bandwidth and varying center frequency. It is a standard assumption that the attacker is power-limited because an attacker without power limits can jam from DC to light at unlimited power, in which case the normal user can receive no throughput. A false reservation is a strong attack in the multi-channel power-limited environment because the reservation message takes much less power than actually using the channel, and the multi-channel environment is appropriate because a power-limited

attacker could fully occupy a single channel even without the need to resort to the false reservation attack.

To counter the false reservation attack, I develop a spectrum channel allocation mechanism that assigns bandwidth to each network user to maximize overall performance in the presence of malicious entities. Because an attacker that uses its reserved bandwidth is indistinguishable from a normal user, I focus on attackers that request bandwidth without using it for transmission. This attack takes relatively little attacker effort (to transmit control messages) and takes resources out of the network disproportionately to attacker effort. I consider jamming legitimate transmissions as the attacker's best alternative attack; however, effective jamming demands greater attacker power and is thus less efficient than false reservation. Successful false reservation also allows attackers to focus their power on jamming the remaining, unreserved channels, causing greater degradation in channel quality for legitimate users than wideband jamming without false reservation. *Ignore-False-Reservation MAC* (IFR-MAC) is a channel allocation scheme that allocates bandwidth according to demonstrated resources, rather than according to identity; it allocates bandwidth to each user based on the power received from that user, and thus forcing attackers to consume power on the reserved channel to effectively make future reservation requests. In other words, I aim to drive the attacker's behavior to converge to that of a normal user (who uses the data channel that it has reserved) or a pure wideband jammer (which cannot be stopped) - but not both. In contrast, an identity-based channel reservation scheme assigns all network users equal bandwidth and is vulnerable to false reservation attack, which I consider to be the *baseline strategy* for IFR-MAC performance comparison.

In order to use the observed power information to effectively coordinate channel access, I design a protocol in which all users agree on and distribute the bandwidth assignment. IFR-MAC provides a centralized channel coordination scheme where there is an online trusted authority, such as the access point in WiFi, who coordinates the users' spectrum use by broadcasting each node's bandwidth assignment. I then devise *Distributed Ignore-False-Reservation MAC* (DIFR-MAC), in which the receivers reach consensus on bandwidth allocation in a distributed manner. Since it is hard to individually detect attackers and exclude attackers from contributing to the coordination process, DIFR-MAC introduces another vulnerability in which attackers can

share false information for coordination. With IFR-MAC and DIFR-MAC, when attackers comprise a minority of the network, the optimal DoS attackers' strategy usually becomes physical-layer jamming, as opposed to the more efficient false reservation attack. IFR-MAC and DIFR-MAC counter the false reservation attack while simultaneously improving spectral efficiency, maximizing the aggregate rate of the network.

The rest of the chapter is organized as follows. I establish the setup of the investigation in Section 4.2. In Section 4.2.2, I present the threat model and assumptions about attackers. I introduce the secure channel reservation scheme, IFR-MAC in Section 4.3 and its distributed counterpart in Section 4.4 and theoretically analyze the performance and the attacker's optimal strategy in Section 4.5. Afterward, I evaluate IFR-MAC with implementation testbed experiments in Section 4.6 and with simulations in Section 4.7 and conclude the chapter in Section 4.8.

4.2 System Model and Assumptions

I consider a scenario in which there are T non-idle transmitters, which compose the set $\mathcal{T}$ (each user is indexed with i where $i \in \mathcal{T} = \{1, 2, \ldots, T\}$), that share a frequency band with a total bandwidth W. In $\mathcal{T}$, there are M malicious attackers, each identified by an index $k \in \mathcal{M} = \{1, 2, \ldots, M\}$, and the rest of them are protocol-compliant. All users operate on open spectrum and communicate directly, i.e., no communication relies on a third node to relay the message to the final destination node. For the model, I also assume that all users are within transmission range of each other, so that any transmission is heard by every user. Thus, when two or more users operate on the same channel, they collide. However, users who operate on different frequency channels, i.e., non-overlapping portion of spectrum, do not interfere with each other.

All data communication is unicast, whereas control packets are broadcast (the literature contains broadcast proposals that ensure availability [24–26]). Because the overhead of a control message is amortized over the data frame, and because I can choose arbitrarily large data frames, I assume that the resource consumption of control communication is much smaller than that of data communication (this makes the false reservation attack more efficient,

as detailed in Section 4.2.2). Thus, I focus on the performance of data communication.

All users have equal priority for transmission and their contributions to the overall network performance are weighted equally. Thus, attackers do not target a specific group of users, and, if their attacking strategy is limited to jamming, then their optimal strategy becomes wideband jamming across the frequency band utilized by legitimate users, as detailed in Section 4.2.2.

In order to prevent forgery of reservation messages, I authenticate control packets containing channel reservation information by timestamping them and including an authenticator based on a secret shared between an offline central authority and the requester. This eliminates spoofing attacks or forgery of packets and ensures that a user is held responsible for the channels it has reserved. I further assume that the central authority knows which users are valid, which prevents the Sybil attack in which one entity fakes multiple identities. (Sybil avoidance is not necessary for IFR-MAC that implements resource-based channel allocation as opposed to entity-based allocation, but it is assumed in DIFR-MAC when users aggregate their observations to agree on a consensus bandwidth allocation.)

In contrast to the baseline strategy of entity-based bandwidth allocation, IFR-MAC diverges from the conventional slotted channelization approach (where the spectrum consists of channels with fixed bandwidth and static location). In particular, by allocating channels with varying bandwidth and center frequency, IFR-MAC can more effectively match the abilities of the users when assigning bandwidth and increase the system's spectral efficiency. Researchers, in a non-security framework, have already adopted the channelization approach with flexible boundaries for frequency channels [30–33]. IFR-MAC does not require that users have the hardware capability of non-contiguous frequency access [31–33], but I assume that attackers do have such capability to describe the optimal attacker strategy and consider the worst-case scenario as described in Section 4.2.2.

4.2.1 Performance Metric: Channel Capacity

The analysis for IFR-MAC holds when using *any* metric as long as it exhibits the two properties that I use in the analysis: it is decreasing and convex

Table 4.1: Variables and their meanings (listed in order of appearance)

Variable Notations (lowercase subscripts indicate a network entity index)	
$\mathcal{T}$	Set of non-idle transmitters
T	Size of $\mathcal{T}$
$\mathcal{M}$	Set of malicious transmitters
M	Size of $\mathcal{M}$
W_i	User i's bandwidth
$i \to j$	Link between transmitter i and receiver j
$\mathcal{R}_{i,j}$	Rate performance of link $i \to j$
N_0	Noise power spectral density
$P_{i,j}$	Effective power of link $i \to j$
P_x	User x's effective power constraint
I_l	User l's normalized interference signal power
J_k	User k's normalized jamming power
$\widetilde{P_x}$	User x's transmitted power spectral density
W_i^*	Bandwidth assigned to user i using IFR-MAC
P_i^*	Signal power on channel reserved by user i
$P_{\mathcal{T}}^*$	Total signal power across frequency band
W	Bandwidth of the entire frequency band
$\overline{P}$	Uniform power constraint, i.e., $\overline{P} \triangleq P_x,\ \forall x \in \mathcal{T}$
α	Attacker's power allocated for false reservation
$\mathcal{R}^C/\mathcal{R}^C$	User performance for centralized-IFR-MAC/DIFR-MAC
U^C/U^D	Network performance for centralized-IFR-MAC/DIFR-MAC
β	Bandwidth advantage of attacker due to collusion
$\mathcal{U}$	Aggregate network performance over time, i.e., $\mathcal{U} = \sum_t U$
α_t	Attacker's power for false reservation at time t
$\hat{\alpha}$	Optimal choice for α

with jamming power, and monotonically increasing with transmitter's signal power. In this chapter, I use the Shannon channel capacity to construct the performance metric.[1] As a reference, I list the notation in Table 4.1.

Whenever user i transmits to its destination user j, it does so on a fre-

[1]The channel capacity given by the Shannon-Hartley theorem provides an asymptotic upper bound for the communication rate of an independent additive white Gaussian (AWGN) channel. As is commonly used for evaluating performance and is generally considered tight (although researchers in information theory continue to pursue even tighter bounds in more complex and realistic channel models), Shannon channel capacity is not only simple and sufficiently expressive to demonstrate the performance of IFR-MAC and DIFR-MAC but it also allows us to abstract away physical-layer decisions such as modulation, coding, and channel estimation. Nevertheless, the results hold when using *any* metric that is monotone and convex, such as a more sophisticated capacity expression.

quency channel, the location of which is pre-shared between user i and user j. At any point in time, the user transmits on frequency channel with bandwidth W_i. Under a flat fading Gaussian channel with Gaussian signals and interference being treated as noise, the Shannon capacity of the link $i \to j$ is:

$$\mathcal{R}_i = W_i \log_2 \left[1 + \frac{P_{i,j}}{N_0 W_i + \sum\limits_{\ell \neq i, \ell \in \mathcal{M}^c} P_{\ell,j} + \sum\limits_{k \in \mathcal{M}} P_{k,j}} \right] \tag{4.1}$$

In Equation 4.1, N_0 is the power spectral density of noise, $\mathcal{M}$ are the indices of jammers, $\mathcal{M}^c$ are the indices of legitimate users, and $P_{x,y}$ is the *effective* or *received* signal power of the link $x \to y$ (the numerator is the transmitter's signal power while the denominator terms correspond to unintentional interference and jamming).

Since the rate performance is determined at the intended receiver, I reconstruct the performance metric using received power (as opposed to the transmitted power). First, I place finite power constraints to all users including attackers (the constraint is denoted with P_x where the subscript x denotes the user index):

$$E[P_{x,y}] \leq P_x < \infty, \ \ \forall x \in \mathcal{T}$$

Since $P_{x,y}$ is random due to the randomness in channel gain for $x \to y$, P_x is the upper constraint of the expected $P_{x,y}$. As $\mathcal{R}_i$ is a monotonically increasing function of P_i, the transmitter i (whose performance I am measuring) will emit at full power and maximize the signal power P_i on the channel, i.e., $E[P_{i,j}] = P_i$. For the analysis, I assume that the channel gain across all channels have the same expected value (users with better channel gains can be modeled by having a larger *effective* power constraint). From Equation 4.1, this yields the following expression for the rate bound:

$$E[\mathcal{R}_i] \leq W_i \log_2 \left[1 + \frac{P_i}{N_0 W_i + \sum\limits_{\ell \neq i, \ell \in \mathcal{M}^c} I_\ell P_\ell + \sum\limits_{k \in \mathcal{M}} J_k P_k} \right] \tag{4.2}$$

where P_i is the transmitter i's signal power, I_ℓ is the amount of user ℓ's power that interferes with the transmitter's signal normalized with respect to the power constraint P_ℓ, J_k is the attacker k's jamming power normalized to the

power constraint, P_k (that is, I_ℓ and J_k are control variables indicating the amount of power emitted by other users with respect to their constraints). I use Equation 4.2 for the performance of link $i \to j$.

The system goal is to maximize the performance of the overall network. I introduce a network utility function, U that is the aggregate rate of the users:

$$U = \sum_{i \in \mathcal{T}} E[\mathcal{R}_i] = \sum_{i \in \mathcal{M}^c} E[\mathcal{R}_i] \tag{4.3}$$

The second equality comes from the fact that the attackers make no contribution to the network; they, in fact, aim to degrade the network performance as described in Section 4.2.2.

4.2.2 Attack Model

Malicious attackers aim to degrade the network performance. Thus, I consider an attacker that intends to minimize the utility function subject to its power constraint:

$$\text{minimize } U \text{ subject to } J_k \leq 1,\ \forall k \in \mathcal{M} \tag{4.4}$$

and that there is collusion among jammers. Thus, attackers learn and share all information (including the secret key for control packets) through a secure, covert communication path. Also, attackers do not care more about a particular user than any other; that is, they do not target a specific subset of users.

I am mainly concerned with two types of DoS attacks: false reservation (wasting network resources) and jamming (injecting noise to decrease reliability of communication). False reservation is the more efficient attack of the two, since it allows an attacker to make a big impact while using less power. Each attacker can send a short reservation request message and reserve a channel for an extended period of time (which is supposedly to be used for data transmission), preventing legitimate users from using the resource. This requires only a small amount of power to deliver control packets (since there is no measure to check or regulate whether the channel is used during the reserved period) and attackers can use the majority of their power to jam

and disrupt the communication of legitimate users. Therefore, attackers will first falsely reserve the bandwidth as much as possible, and then jam the rest of the transmissions on the frequency band that is being used by legitimate users (attackers are capable of accessing non-contiguous frequency band); focusing their jamming power on smaller bandwidth will yield greater interference for users using the frequency band. In this case, attackers are successful in both wasting resources by falsely reserving portions of spectrum and degrading network performance by jamming the rest of the spectrum.

For the frequency band occupied by legitimate users, attackers need to decide whether to choose narrowband jamming or wideband jamming. In Equation 4.1, I observe that the capacity for link $i \rightarrow j$, R_i is a decreasing and convex function of the jamming power of attacker k, $P_{k,j}$. Hence, by Jensen's inequality, the minimum value of its statistical expectation $E[R_i]$ under the constraint in Equation 4.4, is attained by choosing $E[\widetilde{P_k}(f)]$ for every frequency f, where $\widetilde{P_k}(f)$ is the user k's *transmitted* power spectral density at frequency f, e.g., $P_{k,j} = \alpha \int_{f-W_k/2}^{f+W_k/2} \widetilde{P_k}(f)df$ for some scalar channel gain α in a flat fading channel. Thus, to minimize the communication rate, jammers will conduct wideband jamming across the channels that are being used by all legitimate users as opposed to targeting and jamming a specific set of users. Since attackers perform wideband jamming across the entire frequency band except for the portion that they have already falsely reserved (SWIFT [31] offers such wideband transmission avoiding some in-band narrowband channels), legitimate users do not benefit from spreading spectrum, e.g., via frequency hopping.

For each transmission period, all users need to agree on the same bandwidth allocation to be coordinated. In scenarios where a centralized entity does not exist, all users need an algorithm to reach consensus. This introduces another vulnerability which attackers can exploit by attempting to shift the coordination result to their advantages. This secondary attack on distributed channel coordination is discussed in Section 4.4.1.

4.3 Ignore-False-Reservation MAC

For secure channel reservation, Ignore-False-Reservation MAC (IFR-MAC) provides countermeasure for false reservation attack with two properties:

first, it provides the optimal performance in terms of spectral efficiency, and second, it generally causes the optimal power-finite attacker strategy to become a jamming-only strategy. The IFR-MAC protocol thus performs substantially better than the case in which no countermeasures are deployed, since in those environments, an optimal attacker can simultaneously perform false reservation and jamming, as described in Section 4.2.2.

The IFR-MAC bandwidth allocation scheme assigns bandwidth to a user according to the received power from the user. I design two mechanisms for determining and distributing this channel allocation: a centralized scheme run by a trusted entity (IFR-MAC) and a distributed scheme in which users cooperate to determine the bandwidth allocation (DIFR-MAC in Section 4.4).

4.3.1 IFR-MAC Bandwidth Allocation

IFR-MAC assigns bandwidth according to the observed received power; that is, each user is assigned bandwidth proportional to the amount of the user's effective power on the corresponding receiver. In other words, receiver only respects a user's channel request proportionally to the amount of power that user emits on the data channel. Therefore, an attacker needs to emit power on the reserved data channel in order to make a valid bandwidth reservation request for the following transmission. Thus, attackers' limited power capabilities force them to split their power between reserving channels and jamming legitimate transmissions.

When attackers use their entire power to transmit on their reserved data channels, I do not distinguish them from legitimate users. An attacker does not care about its own bandwidth under the attacker model but instead intends to minimize the bandwidth available to legitimate users. Thus, attackers' data transmissions do not contribute to the overall network performance. However, the problem of detecting that an attacker's data packets are content-free is difficult (potentially impossible under certain encryption schemes), so IFR-MAC does not attempt to identify attackers or their traffic.

Let W_i be the amount of transmission bandwidth assigned for user i. Suppose the IFR-MAC bandwidth allocation scheme using observations on the received power results in $\overrightarrow{W^*} = (W_1^*, ..., W_i^*, ..., W_T^*)$ where W_i^* corresponds

to the bandwidth allocated to user i, i.e., $W_i^* \propto P_i^*$, where P_i^* denotes user i's power on its reserved bandwidth W_i. ($P_i^* \neq P_i$, since P_i^* corresponds to the actual power used for data communication on the reserved bandwidth whereas P_i corresponds the power constraint; for example, an attacker can use some of its power budget on jamming other channels.) Also, let $P_{\mathcal{T}}^*$ be the total network power on the reserved channels, i.e., $P_{\mathcal{T}}^* = \sum_{i \in \mathcal{T}} P_i^*$, and W be the total bandwidth available for the network of users $\mathcal{T}$. Since I do not want to waste bandwidth, I allocate the entire frequency band, i.e., $\sum_{i \in \mathcal{T}} W_i^* = W$. Then, the scheme yields that $W_i^* = W \cdot \frac{P_i^*}{P_{\mathcal{T}}^*}, \quad \forall i \in \mathcal{T}$. Now, I claim that the IFR-MAC bandwidth allocation scheme also yields optimal performance (that is, it maximizes the network utility function, Equation 4.3).

Claim 3.

$$\underset{(W_1,\ldots,W_T)}{\operatorname{argmax}} U = \overrightarrow{W^*} = \left(W \cdot \frac{P_1^*}{P_{\mathcal{T}}^*}, \ \ldots, W \cdot \frac{P_T^*}{P_{\mathcal{T}}^*} \right) \tag{4.5}$$

Proof. To simplify the proof, I examine the case where the channel gain is equal across all channels, the noise power has equal power spectral density (N_0) over all frequencies, and that the attackers choose a fixed strategy (which I later show to be the optimal strategy in Section 4.5.2) between false reservation and jamming with their limited power budget. If attackers choose to jam, they conduct wideband jamming over narrowband jamming, as described to be the more efficient attack in Section 4.2.2.

Using Equation 4.1 and Equation 4.3, I observe that U (which is a summation of $E[\mathcal{R}_i]$) is a concave function with respect to $P_{i,j}$, and in a flat fading environment, with respect to $\widetilde{P}_i(f)$. Jensen's inequality yields $E[U(\widetilde{P}_i)] \leq U(E[\widetilde{P}_i])$. Therefore, in order to maximize the expected utility function, the network needs to have constant power spectral density across the entire frequency bandwidth W, i.e., $\widetilde{P}_i(f) = E[\widetilde{P}_i(f)] = \frac{P_{\mathcal{T}}^*}{W}, \forall f, \forall i$. This is equivalent to users' transmission bandwidth being proportional to their power capability, and yields $W_i = W \cdot \frac{P_i^*}{P_{\mathcal{T}}^*} = W_i^*$.

In one-user case, where user i is the sole member of the network, $W_i^* = W$ since $P_i^* = P_{\mathcal{T}}^*$, i.e., being the only member in the bandwidth allocation process, user i consumes all the bandwidth available. The case for a bigger network and more entities participating in the bandwidth allocation is equiv-

alent to having a single user that has the same power capability as the sum of that of the network (i.e., $P_i^* = \sum_{x \in \mathcal{T}} P_x^*$), as the destination receiver does not distinguish the source of the received power. Since I have the bandwidth allocation non-overlapping, the proof on the previous paragraph extends to the multi-user case.

The proof holds whether or not an attacker reserves a channel or jams the frequency band. This is because attackers who transmit on their reserved bandwidth for successful channel reservation, i.e., $\exists i \in \mathcal{M},\ P_i^* \neq 0$, are indistinguishable from legitimate users, while attackers who wideband jam the channels reserved by legitimate users increase the interference power spectral density by the same amount for all frequency that they jam.

The claim continues to hold in expectation regardless of channel gain or the distribution of the power spectral density of noise, e.g., frequency-selective fading. This is intuitive because in expectation, "waterfilling" the spectrum with power is known to maximize spectral efficiency. □

4.3.2 IFR-MAC Channel Coordination

The bandwidth allocation scheme in Section 4.3.1 is performed individually by each user. Before data transmission, all users need to agree on the same bandwidth allocation, which process I call *channel coordination*. Otherwise, their transmissions will overlap and cause collisions, resulting in poor performance. In a centralized scheme, a trusted entity can simply determine a bandwidth assignment (the trusted user performs channel coordination within itself) and broadcast that assignment to the users. Section 4.4 introduces a distributed counterpart of IFR-MAC channel coordination that involves inter-user communication.

4.4 Distributed IFR-MAC

Distributed Ignore-False-Reservation MAC (DIFR-MAC) builds on the bandwidth allocation scheme of IFR-MAC (described in Section 4.3.1) and extends the scheme for distributed channel coordination. In scenarios where a trusted entity does not exist or is offline, the users need to share the observations and reach consensus for channel coordination to avoid data collision. In DIFR-

MAC, users exchange their *power reports* that contain their received power observations. After gathering all users' power reports using secure broadcast [24–26], users choose the median of the power observations. Since all users choose the median from the same set of values, they reach consensus (the solution to the Byzantine general's problem using signed messages [34] also provides the same result). I also propose a commit-and-reveal protocol that conceals the power reports until all values are exchanged and gathered by every users, the motivation and procedure for which is detailed in Section 4.4.1.

4.4.1 DIFR-MAC against False Report Attack

Distributed channel coordination to reach consensus on bandwidth allocation is vulnerable to an attack where attackers attempt to distort the consensus to their advantage. I call this attack *false report attack* on distributed channel coordination where attackers tamper the coordination process by reporting false power reports. Such attack is inherent in a distributed algorithm where attackers participate and contribute to the outcome. DIFR-MAC channel coordination mitigates false report attack by using the median, an input-resilient metric for coordination [35] (a non-robust metric such as mean gives much control in channel coordination to attackers who report false power observations). Nevertheless, since attackers know each other, they can still distort the median by reporting favorable values (i.e., high received power observations) for fellow attackers and low power for others. The false report attack is difficult to detect because variable channel conditions by fading causes the reported observations of received power to vary between all users. The result of the attack is that the consensus median value will be shifted toward the value that the attackers report. Also, the use of median introduces an additional constraint that must be placed on the number of attackers: $M < \frac{T}{2}$ (this constraint is not necessary in a centralized scheme); otherwise, the attackers outnumber legitimate users and has direct control over the median and thus the channel coordination outcome.

Attackers that know the legitimate users' received power observations know exactly the median value and the maximum possible distortion. I hide the legitimate users' reports of received power values by committing the reported

values and then revealing them after all reports are gathered. Using a one-way hash function, a commit-and-reveal protocol conceals the power observations until all reports are broadcasted and gathered. In DIFR-MAC, the committed value is the hash of the power report and the corresponding hash function while the revealed value is the power report itself; only after all the committed values are broadcasted will users reveal the received power.

An attacker detection scheme using thresholds can prevent attackers from reporting extremely large or small values. However, even with a scheme that hides the legitimate users' reports, such a threshold-based detection scheme can be defeated by attackers who infer other legitimate users' observations based on the past reports. Attackers can then decide how much to distort the median by reporting moderately biased values while avoiding being detected. Not only is a threshold-based detection scheme ineffective, it can also result in performance-degrading false positives; when the attacker is intelligent, such a detection scheme will mostly detect legitimate, unintentional outliers, and may degrade the future performance by punishing benign users. Therefore, I do not consider such detection schemes.

4.5 Theoretical Analysis

IFR-MAC reduces the problem of false reservations to a two-party game between the *legitimate user network* (users who wish to maximize network utility) and the *attacker network* (entities who want to degrade network performance), because I assume cooperative behavior among benign users and collusion among attackers, and because the bandwidth allocation depends only on the received power.

IFR-MAC allocates bandwidth proportional to each node's power level. The IFR-MAC's behavior, and the attacker's optimal strategy, therefore depend on the relative power capabilities of the legitimate and attacker networks, rather than on the number of users. In the theoretical analysis, I consider nodes with equal power constraints; that is, all individual users, including attackers, have the same power constraint $\bar{P}$, i.e., $P_i = \bar{P}$, $\forall i \in \mathcal{T}$. Then, the power capability ratio of the legitimate user network to that of the attacker network is $\frac{T-M}{M}$, so the power capabilities of the two groups can be controlled by varying the number of users (T) and attackers (M). I denote

the expected performance of individual users as $\mathcal{R}^C = E[\mathcal{R}_i]$, $\forall i \in \mathcal{T}$, in the centralized IFR-MAC scheme, and with $\mathcal{R}^D$ in the DIFR-MAC scheme. For simplicity of analysis, I introduce a variable α to represent the fraction of attacker power expended on channels reserved by the attacker, so that $1-\alpha$ represents the fraction of attacker power used for jamming other channels.

4.5.1 IFR-MAC Performance

In the centralized scheme, the trusted entity makes its own observations and assigns bandwidth to users accordingly; thus control traffic flows only between the trusted entity and the nodes. In this section, I study the network performance of the centralized IFR-MAC scheme. Since the attacker network uses $(\alpha \cdot \bar{P} \cdot M)$ power for false reservation (and $[(1-\alpha) \cdot \bar{P} \cdot M]$ power for jamming), a legitimate user can reserve bandwidth $W_i = \frac{W}{T-M+M\alpha}$, where the denominator $T-(1-\alpha)M$ represents the number of valid reservations. This reduces Equation 4.2 into:

$$\begin{aligned}\mathcal{R}^C &= \frac{W}{T-M+M\alpha} \cdot \log_2 \left[1 + \frac{\bar{P}}{\frac{W \cdot N_0}{T-M+M\alpha} + \frac{M(1-\alpha)\bar{P}}{(T-M)}}\right] \\ &= \frac{W}{T-M+M\alpha} \log_2 \left[1 + \frac{\text{SNR}}{\frac{T}{T-M+M\alpha} + \frac{M(1-\alpha)}{(T-M)}\text{SNR}}\right]\end{aligned}$$

where the signal-to-noise ratio SNR is the ratio between the network power capability (including that of the insider attackers) and the natural noise on the entire frequency band ($\text{SNR} = \frac{T \cdot \bar{P}}{W \cdot N_0}$). Using Equation 4.3,

$$U^C = (T-M) \cdot \mathcal{R}^C$$

Combining both equations,

$$U^C = \frac{(T-M) \cdot W}{T-M+M\alpha} \log_2 \left[1 + \frac{\text{SNR}}{\frac{T}{T-M+M\alpha} + \frac{M(1-\alpha)}{(T-M)}\text{SNR}}\right]$$

4.5.2 Attacker Reaction to IFR-MAC

Section 4.3 describes IFR-MAC, which forces attackers to use a channel in order to reserve channels in the future. This section presents the following result about the optimal attacker strategy in IFR-MAC: attackers will use a static strategy (α is constant in time) to minimize the network utility. Let $\mathcal{U}$ be the aggregate network utility over time, i.e., $\mathcal{U} = \sum_t U_t$ where U_t is the network performance (expressed in Equation 4.3) at time t, and let α_t be the amount of power attackers use to reserve a channel rather than jamming at time t.

Claim 4. *Given $\hat{\alpha}$ that yields minimum U,*
$\alpha_t = \hat{\alpha}$, $\forall t$, yields the minimum performance, $\mathcal{U}$.

Proof. I use the network utility function expression for the centralized scheme U^C in this proof (the proof for distributed scheme follows the same procedure). From Equation 4.6, both $\frac{(T-M)\cdot W}{T-M+M\alpha}$ and $\log_2\left[1+\frac{\text{SNR}}{\frac{T}{T-M+M\alpha}+\frac{M(1-\alpha)}{(T-M)}\text{SNR}}\right]$ are convex, monotonic, and positive for all possible α. Therefore, the product, U_c is also convex with respect to α. By using Jensen's inequality, $\alpha_t = E[\hat{\alpha}] = \hat{\alpha}$, $\forall t$ yields the minimum network performance, $\mathcal{U}$. □

4.5.3 DIFR-MAC Performance

Without an online trusted authority, channel coordination becomes a distributed problem. In this section, I describe DIFR-MAC, which takes the median of each user's measurement to obtain a consensus bandwidth allocation, as described in Section 4.3.2. I also consider the false reporting attack described in Section 4.4.1, where attackers report receiving less power from legitimate users and more power received from colluding attackers. The user is assigned channel bandwidth proportional to the median of the power reports. I use β to denote the attacker's *bandwidth advantage* over a legitimate user; as discussed in Section 4.4.1, attackers can reserve more bandwidth than legitimate users with the same amount of power ($\beta \geq 1$) because attackers can collude while legitimate users report truthfully.

Because the power a node receives from another node fluctuates due to channel condition, legitimate nodes report different power levels for the same transmission. An attacker can shift the median by reporting an extreme

value. Without any attackers, the median returns the 50th percentile measurement for each transmitter. When assessing the data transmission power of colluding attacker, attackers report large power observations, shifting the observed median upward to the $100 \cdot \frac{0.5 \cdot T}{T-M} > 50$ percentile of legitimate observations. Also, for a legitimate user's power report, the attackers report low power to shift the median downward to the $100 \cdot \frac{0.5 \cdot T - M}{T-M} < 50$ percentile of legitimate observations. In contrast, legitimate users report their true observations that include random wireless channel fluctuation. Assuming a channel model with independent and identically distributed channels for all users, where each channel characterized by a cumulative distribution function CDF, the bandwidth advantage of an attacker is:

$$\beta = \frac{\text{CDF}^{-1}(\frac{0.5 \cdot T}{T-M})}{\text{CDF}^{-1}(\frac{0.5 \cdot T - M}{T-M})} \tag{4.6}$$

I consider three channel fading models, where channel characteristics vary in Rician fading with parameters ν and σ, where ν^2 is the power of the line-of-sight path[2] and $2\sigma^2$ is the power of the other, scattered paths. In specific, I study three cases for channel fading characteristics: *strong line-of-sight* ($\frac{\nu}{\sigma} = 10$), *weak line-of-sight* ($\frac{\nu}{\sigma} = 2$), *no line-of-sight* ($\frac{\nu}{\sigma} = 0$). The *no line-of-sight* case is equivalent to the Rayleigh fading model, a typical model used in highly dynamic environments, such as a mobile application in an urban setting.

I plot the attacker bandwidth advantage β under each of the three channel fading models in Figure 4.1. Increasing the number of attackers results in greater error in the computed median, resulting in less-correct bandwidth assignment and greater bandwidth advantage. Because channel fluctuation affects the randomness within the power reports, the attacker's bandwidth advantage β depends on $\frac{\nu}{\sigma}$; as the line-of-sight path becomes less dominant, the attackers gain additional bandwidth.

Under the false reporting attack, one legitimate node's bandwidth is one part in $T-M+\alpha\beta M$, since $T-M$ is the number of legitimate users and $\alpha\beta M$ is the effective requests made by attackers, resulting in per-user performance ($\mathcal{R}^D$) and network performance (U^D) of:

[2]As is common in wireless communications, the term *line-of-sight* path refers to the most dominant channel path, and not necessarily the straight-line path between the two nodes.

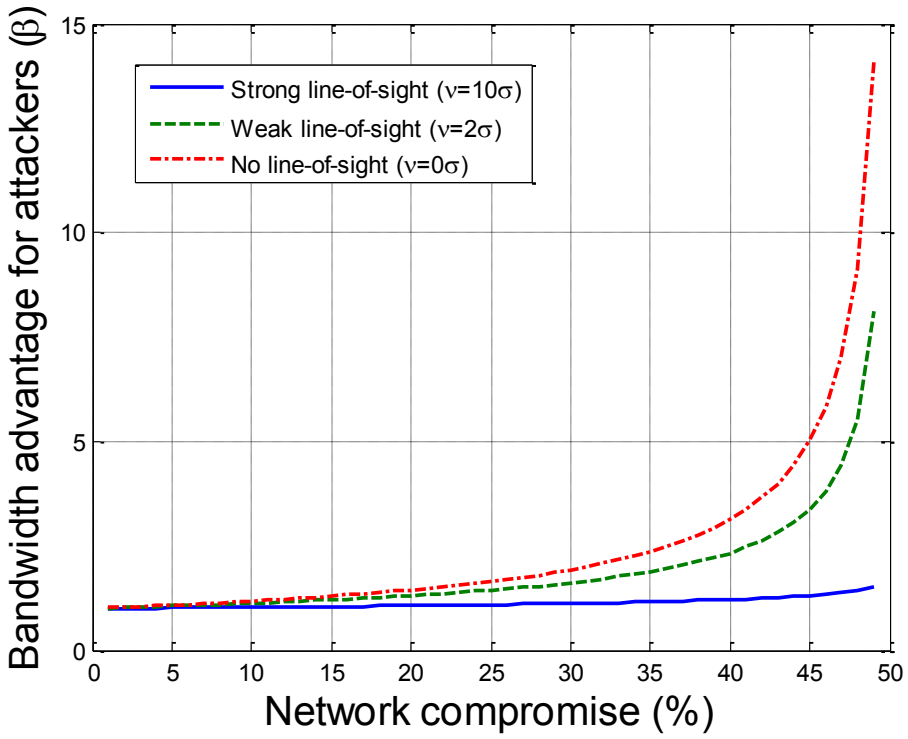

Figure 4.1: Ratio of attacker's and legitimate user's bandwidth under false report attack against DIFR-MAC

$$\mathcal{R}^D = \frac{W}{T - M + M\alpha\beta} \log_2 \left[1 + \frac{\text{SNR}}{\frac{T}{T-M+M\alpha\beta} + \frac{M(1-\alpha)}{(T-M)}\text{SNR}} \right]$$

$$U^D = \frac{(T - M) \cdot W}{T - M + M\alpha\beta} \log_2 \left[1 + \frac{\text{SNR}}{\frac{T}{T-M+M\alpha\beta} + \frac{M(1-\alpha)}{(T-M)}\text{SNR}} \right]$$

Compared to the centralized scheme, the distributed scheme, DIFR-MAC gives attackers an additional factor of β more bandwidth.

4.6 Testbed Evaluations

To complement the theoretical analysis in Section 4.5, I use testbed implementation and computer simulation to study the effectiveness of the IFR-MAC bandwidth allocation and the DIFR-MAC channel coordination scheme. The testbed implementation is built on the WARP software-defined radio platforms [27] (Section 4.6), and the simulation is built on MATLAB (Section 4.7).

4.6.1 Evaluation Methodology

I compare the performance of IFR-MAC to a baseline protocol with no countermeasure. In the baseline protocol, attackers can reserve channels without using power on the channels they are allocated; the frequency band is divided equally into multiple channels; and each user gets one channel. Optimal strategy against this baseline protocol, as detailed in Section 4.2.2, is for attackers to reserve channels, wasting $\frac{M}{T}$ of the entire network bandwidth without consuming much power (I assume that control packets are considerably smaller than data packets, so that the power required to transmit control packets is negligible). For the baseline strategy, $\alpha = 0$ represents the *optimal baseline strategy*, since α does not affect channel reservation (as discussed in Section 4.6.4).

In implementation, I use four WARP (Wireless Open Access Research) software-defined radio platforms [27]. Using the MIMO (multiple input mul-

tiple output) capability of the platform, I simulate a network comprised of four transmitters: two legitimate transmitters, one attacker, and one *identity-only attacker* with zero power. The legitimate transmitters and the attacker each have equal power budget (except when I vary the attacker power budget in Section 4.6.4). For the centralized IFR-MAC analysis, I fix one receiver to be the destination of the data packets and the trusted authority that performs IFR-MAC bandwidth allocation and channel coordination. Furthermore, I manually calibrate the antenna locations so that the receiver observes approximately the same power from each transmitter. For DIFR-MAC, I consider three equal-power transmitters and pair each transmitter with a receiver on a separate WARP board to simulate full duplex.[3] (Section 4.6.5 gives a detailed description of the DIFR-MAC implementation.)

To maximize network utility U, each node maintains full queues and continuously transmits. At the physical layer, nodes use differential quadrature phase-shift keying (DQPSK) modulation with a BPSK-modulated Barker sequence preamble. The total network bandwidth is 12 MHz wide and divided into 300 subchannels using orthogonal frequency division multiplexing (OFDM). Each transmitter sends random bits to its receiver, and its receiver demodulates the received signal and uses the bit error rate to estimate the effective signal-to-interference-and-noise ratio (SINR) at the receiver, using the equation from [28, 29]:

$$\overline{\mathrm{BER}} = \frac{1}{2}\left(1 - \frac{\sqrt{2}\cdot\overline{\mathrm{SINR}}}{\sqrt{1+4\cdot\overline{\mathrm{SINR}}+2\cdot\overline{\mathrm{SINR}}^2}}\right)$$

For performance evaluation, I compute the network's aggregate capacity from the effective SINR using Equation 4.1 and Equation 4.3 and then the network's spectral efficiency by dividing the aggregate capacity by the total network bandwidth (although the compromised node can waste some of the bandwidth).

[3]WARP version 2 does not support duplex because the antennas and the radio front-ends on a board partially share the processing chain preventing them from transmitting and receiving at the same time.

4.6.2 IFR-MAC Bandwidth Convergence

I study the bandwidth allocations under the two attacker strategies of $\alpha = 0$ (jamming) and $\alpha = 1$ (effective channel reservation). IFR-MAC allocates the entire network bandwidth, and each user's channel bandwidth is proportional to the observation of the power originated from the user. I model an optimal attacker, which means the attacker uses a fixed strategy (as described in Section 4.5.2). Figure 4.2(a) shows the expected normalized bandwidth allocation to the four transmitters over time. Beginning from the baseline strategy of equal-bandwidth allocation (i.e., each of the four entities occupy $\frac{1}{4}$ of the network bandwidth), IFR-MAC quickly converges to the steady-state bandwidth allocation, where the delay is caused by the noise in the spectrum reserved by attackers (noise power monotonically increases with bandwidth). Within two rounds of updates, IFR-MAC converges to the steady-state bandwidth allocation. I also plot *legitimate-user bandwidth*, the fraction of bandwidth utilized by the legitimate network and thus contribute to the network throughput, which converges to 1 and to $\frac{2}{3}$, respectively, for $\alpha = 0$ and for $\alpha = 1$, validating the steady-state theoretical results. Furthermore, the identity-only attacker quickly converges to zero bandwidth as it emits no power, and thus has no impact on network performance under IFR-MAC. The identity-only attacker, whose strategy is independent of α (because it has no power), gets bandwidth identical to that of the power-capable attacker with $\alpha = 0$. Figure 4.2(a) shows that even when attackers collude by having the jammer emit power on the identity-only attacker's reserved spectrum, IFR-MAC still converges quickly.

4.6.3 Attacker Power-Splitting Strategy

Against IFR-MAC, an attacker with a finite power budget can choose to jam ($\alpha = 0$), make effective bandwidth reservations ($\alpha = 1$), or split power between the two strategies ($0 < \alpha < 1$). I study the network performance with varying attacker strategy and present the results in Figure 4.2(b). These results show that for IFR-MAC (and DIFR-MAC) bandwidth allocation, jamming is more detrimental than consuming power on the reserved channel for effective reservation, so the optimal attacker strategy is jamming ($\alpha = 0$). However, in Section 4.7 where I study the optimal attacker strategy

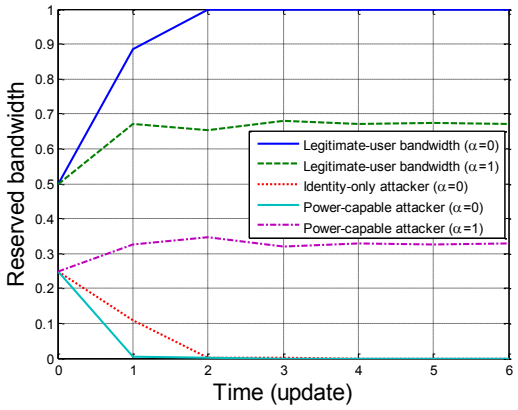

(a) Bandwidth utilization normalized by the total network bandwidth

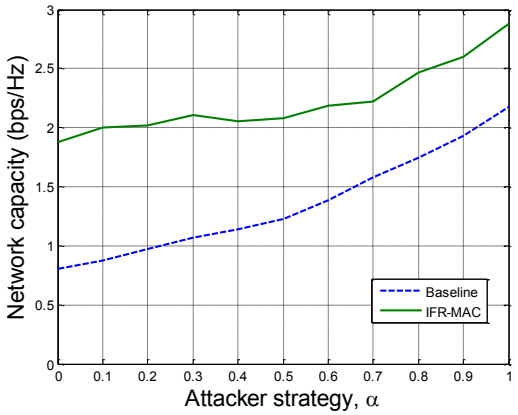

(b) IFR-MAC performance with α

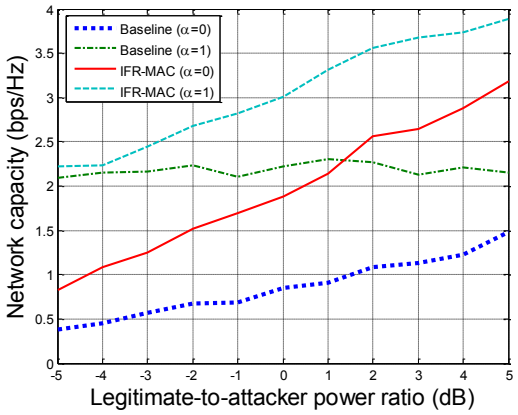

(c) Varying legitimate power relative to attacker power for $\alpha = 0$ and $\alpha = 1$

Figure 4.2: IFR-MAC implementation testbed results

with more complex network settings, I see that there are cases that the optimal attacker strategy diverges from only jamming. Here I sketch an intuition for IFR-MAC. For a given α, the ratio between legitimate user power and jammer power on each channel is constant regardless of legitimate user bandwidth. Thus, in high SNR regions, the performance gain of IFR-MAC over the baseline strategy is proportional to the legitimate network bandwidth utilization and decreases with increasing α.

4.6.4 Attacker Power Budget

In this section, I vary the attacker power budget while fixing the legitimate user's power (the identity-only attacker retains zero power budget). Since the optimal attacker uses all of its power budget to degrade network performance, varying the power budget directly affects the *legitimate-to-attacker power ratio* experienced by each receiver. I consider the ratio of power between a single legitimate user and the single attacker. Figure 4.2(c) plots the aggregate capacity performance against the legitimate-to-attacker power ratio. Larger legitimate-to-attacker power ratios result in better performance; when the attacker is jamming ($\alpha = 0$), this performance increase comes from reduced interference, and when the attacker is reserving ($\alpha = 1$), this performance increase comes from increased bandwidth. Fixing the MAC strategy (e.g., either baseline or IFR-MAC), $\alpha = 0$ provides worse performance than $\alpha = 1$, as predicted in Section 4.6.3. Results show that incremental power increases are best spent on jamming: nonzero values of α show less impact with increasing power. Since attacker can do much greater damage by jamming, $\alpha = 0$ is the *optimal baseline strategy* for the attackers.

4.6.5 DIFR-MAC Channel Coordination

In contrast to the centralized IFR-MAC, where a trusted entity (such as the receiver of uplink traffic) decides on the bandwidth allocation and announces it to the transmitters, in DIFR-MAC, the transmitters need to determine the bandwidth allocation among themselves. DIFR-MAC allocates bandwidth by computing the median of the observed power levels; each transmitter observes power levels by using full-duplex radio techniques [36]. The imple-

mentation pairs two radio chains from distinct WARP boards to simulate duplex. Whereas current work in full duplex wireless attempts to maximally cancel the sender signal to minimize self-interference, DIFR-MAC wants to achieve a fixed attenuation in order to estimate the relative power levels of the transmitters. The techniques for achieving fixed attenuation, or for digitally re-creating the eliminated interference, is beyond the scope of the dissertation. I bypass such calibration through the use of separate radios for transmission and reception, and by placing the receiving antennas so that they experience approximately the same amount of transmitter power from all transmitters. An alternative algorithm, which is also reflected by the testbed setup, is that each transmitter node has its distinct, designated receiver node, that helps in channel coordination by deciding on the power observation and having the transmitter relay the power report (in which case, the power report is compromised if the transmitter or the receiver is compromised).

In distributed channel coordination, attackers are vulnerable to the false reporting attack. Even though median-based DIFR-MAC channel coordination mitigates the false power report attack, as discussed in Section 4.4.1, it performs worse than the centralized IFR-MAC channel coordination which is unaffected by false reports. To show the persisting effect of the false reporting attack on DIFR-MAC, Figure 4.3(a) compares the bandwidth utilization of IFR-MAC and DIFR-MAC when attackers focus either on making successful false reservations ($\alpha = 1$) or jamming ($\alpha = 0$), and Figure 4.3(b) compares the capacity performance of IFR-MAC and DIFR-MAC. Comparing the two results, I observe that the performance difference between DIFR-MAC and IFR-MAC arises mostly from the impact of false reports on channel bandwidth coordination. The results show that DIFR-MAC performs better than the baseline despite the persistent effect of false power report attack. Finally, I observe that DIFR-MAC shifts the attacker's optimal power-splitting strategy from $\alpha = 0$, the optimal strategy against IFR-MAC. Section 4.7 continues the evaluation of DIFR-MAC in a mobile environment.

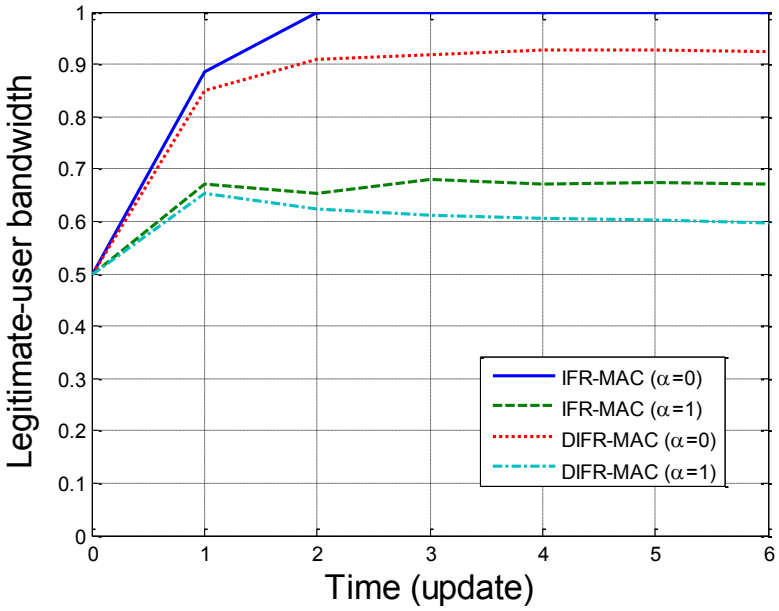

(a) Legitimate-user bandwidth utilization normalized by the total network bandwidth

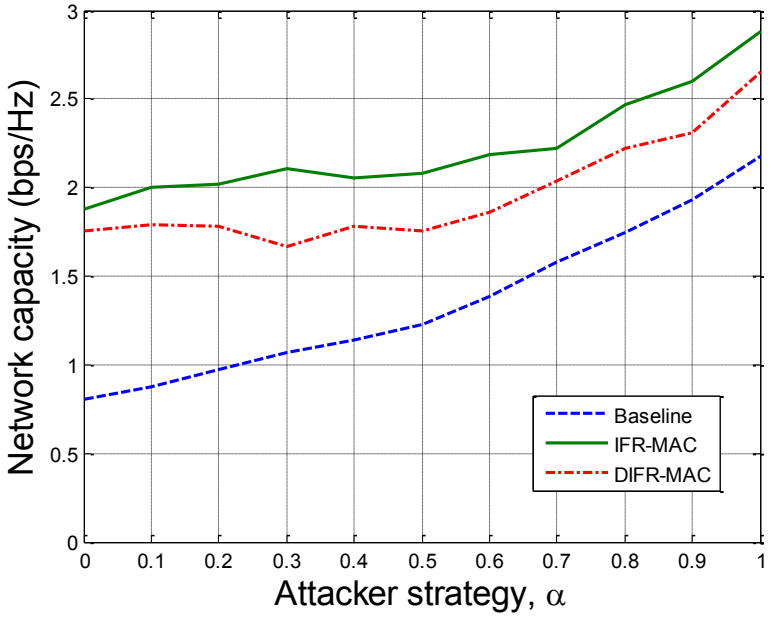

(b) Capacity performance with α

Figure 4.3: IFR-MAC and DIFR-MAC comparison

4.7 Simulation Evaluations

In Section 4.6, I showed the results from implementing IFR-MAC and DIFR-MAC on software-defined radios; I demonstrated that the schemes were practical and that they provide a performance advantage over the naïve baseline strategy of identity-based channel allocation; I also confirmed the theoretical determinations about optimal attacker strategy. To demonstrate scalability, I use MATLAB to simulate a more complicated network topology and channel fading environment. I simulate a network with equal-power transmitters and compute the bandwidth utilization and SINR, from which I derive the aggregate capacity performance (based on the theoretical analysis in Section 4.5). To capture the game between the legitimate network and the attacker network, I vary the network parameters and study their effect on IFR-MAC and DIFR-MAC. The parameters are: $T = 100$ transmitters, SNR (without interference) = 15 dB, W = 20 MHz, and the number of malicious entities M is varied and shown in the plots. In the DIFR-MAC study, network users are capable of full duplex transmissions, and I model the channel as an i.i.d. AWGN channel with Rayleigh fading, which emulates a highly dynamic environment, such as mobile applications in urban setting.

As discussed in Section 4.5, the relative power capabilities of legitimate user network and attacker network affects the performance of the IFR-MAC bandwidth allocation scheme. I observe that attacker network has a $\frac{M}{T}$ fraction of the entire power consumption of the network, while legitimate users have $\frac{T-M}{T}$.

Using the performance metric of network capacity (mathematically expressed in Equation 4.3), I analyze the *performance ratio* between the aggregate rate of IFR-MAC/DIFR-MAC and the aggregate rate of the baseline performance. In particular, I use the performance ratio over the optimal baseline strategy, which is that attackers jam ($\alpha = 0$). Since I use performance ratio, the number of legitimate users ($T - M$) becomes irrelevant; rather, the metric corresponds to the improvement observed by each user as compared to the baseline strategy with no countermeasure. I also study the optimal attacker reactions to IFR-MAC and DIFR-MAC where $\hat{\alpha}$ denotes the optimal power-splitting strategy.

4.7.1 IFR-MAC Performance

Unsurprisingly, the rate performance suffers as the attacker network increases its power capability of M, given T. On the other hand, Figure 4.4(a) plots performance ratio over the optimal baseline strategy while varying the attacker power-splitting strategy of α (the optimal baseline strategy has a constant performance across the x-axis since attackers only perform jamming and $\alpha = 0$), IFR-MAC not only outperforms the baseline performance in all possible scenarios (performance ratio is always greater than one) but also provides greater resistance to attacks with increasing attacker capability (performance ratio increases as M increases).

When the power capability of the attacker network is smaller than that of the legitimate user network, attackers' optimal strategy is to jam with all their power ($\alpha = 0$). Thus, the optimal attacker strategy $\hat{\alpha}$ is to jam with all of its power ($\hat{\alpha} = 0$), which agrees with out testbed studies in Section 4.6.3. However, as aggregate attacker power reaches and exceeds aggregate legitimate user power, $\hat{\alpha} > 0$ (using some power to reserve channels) becomes the optimal attacker strategy. For example, in Figure 4.4(a), when attackers have as much power capability as the legitimate users, i.e., $M = 50$ out of $T = 100$, attackers will transmit on data channels with $\hat{\alpha} = 0.2$ of their power capability, in order to obtain some valid channel reservations, and use the rest of the power $1 - \hat{\alpha} = 0.8$ to jam the rest of the channels, on which the legitimate users transmit. The optimal attacker strategy diverges from $\alpha = 0$, since jamming has a logarithmic impact on network performance while reservations have a linear impact on network performance. Therefore, as attackers' power capabilities grow, the marginal impact of reserving and consuming bandwidth exceeds that of jamming legitimate transmissions.

In order to better compare the optimal attacker strategy and pure jamming, Figure 4.4(c), varies the fraction of attacker nodes in the network and studies the *attack optimality of pure jamming* by plotting the performance ratio between the optimal attacker strategy (that minimizes the network performance), $\alpha = \hat{\alpha}$, and that when attackers jam at full power, $\alpha = 0$. Thus, the metric indicates how much pure jamming underachieves the attacker's goal of degrading the network performance compared to the optimal power-splitting strategy. When normal users outnumber malicious users, and thus legitimate user channels have sufficiently good quality, jamming ($\alpha = 0$) is

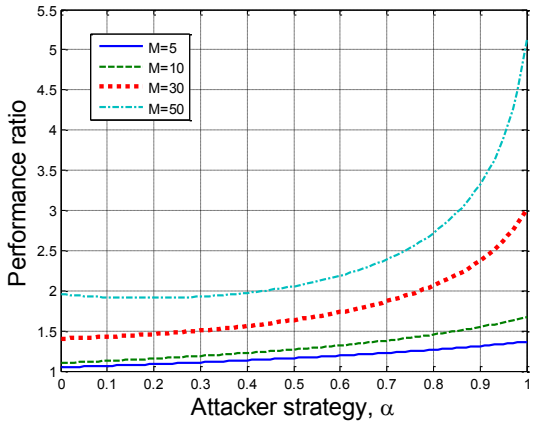

(a) IFR-MAC performance over baseline strategy

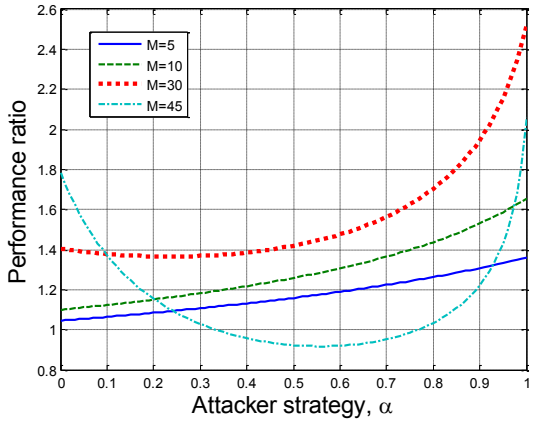

(b) DIFR-MAC performance over baseline strategy

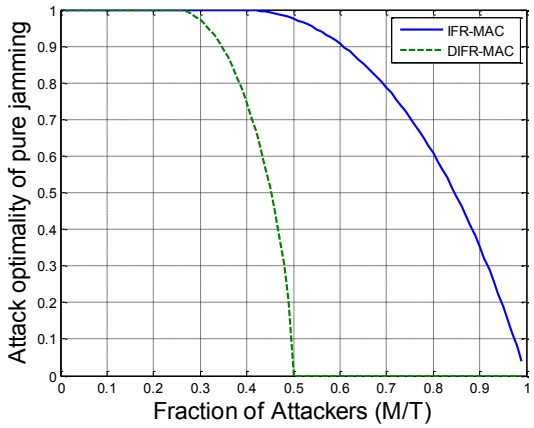

(c) Comparison between optimal strategy ($\alpha = \hat{\alpha}$) and only jamming ($\alpha = 0$)

Figure 4.4: IFR-MAC computer simulation evaluations

an optimal or near-optimal strategy. In particular, $\hat{\alpha} = 0$ until about 45% of the network nodes are compromised, and there is only 2.5% difference in performance between $\alpha = 0$ and $\alpha = \hat{\alpha} = 0.2$ when half the nodes are malicious ($M = T - M = 50$). Therefore, in most practical scenarios (where attackers do not substantially outnumber legitimate users) $\alpha = 0$ is the optimal jammer strategy or is negligibly suboptimal.

4.7.2 DIFR-MAC Performance

DIFR-MAC scheme has similar properties to those of its centralized counterpart. However, from Figure 4.4(b), I observe that the optimal attacker strategy diverges from $\alpha = 0$ more than in the centralized scheme due to the false reporting attack on distributed channel coordination, described in Section 4.4.1. For example, when $M = 30$, $\alpha = 0.24$ is the optimal strategy for distributed scheme (Figure 4.4(b)) whereas $\alpha = 0$ is the optimal attacker behavior in the centralized scheme (Figure 4.4(a)). From Figure 4.4(c) that displays the attack optimality of pure jamming, $\hat{\alpha} = 0$ if less than 27% of network is compromised by attackers. Also, as discussed in Section 4.4.1, reserving *any* bandwidth ($\alpha > 0$) yields complete control on bandwidth allocation to the attacker if $\frac{M}{T} \leq 0.5$.

When compared to IFR-MAC, DIFR-MAC has less performance gain over the baseline performance (as shown in Figure 4.4(b)) due to the false power reporting attack. Also unlike the the centralized scheme, DIFR-MAC's performance can be worse than the baseline strategy as the number of attackers (M) increases. For example, in Figure 4.4(b), $M = 45$ results in performance ratio of less than one for some α. This reduced performance arises from the bandwidth advantage that comes from the false reporting attack, which grows quickly as M increases. As shown in Figure 4.1, when $M = 45$, the attacker can reserve five times as much bandwidth than legitimate users if the same amount of power is used for reservation ($\beta = 5$). Therefore, as the number of attackers approaches the number of legitimate users, DIFR-MAC becomes less effective, and attackers have total control over the scheme if they have as many as or are greater in number than legitimate users, as described in Section 4.4.1. However, DIFR-MAC only exhibits poor performance when the number of malicious users approaches 50%; before DIFR-MAC breaks down

due to the false-reporting attack, DIFR-MAC effectively prevents the false reservation. For example, even when 30% of nodes are attackers, DIFR-MAC provides nearly a 40% performance improvement over the baseline strategy.

4.8 Chapter Summary

This chapter studies the *false reservation* attack, where attacker reserves bandwidth without the intention of using it and uses its remaining power to jam legitimate transmissions. IFR-MAC defends against the false reservation attack by allocating channel bandwidth based on received power observations, and not only effectively counters false reservations, but also yields optimal performance. I design both centralized IFR-MAC and its distributed counterpart DIFR-MAC to perform power-fair bandwidth allocation. The evaluations show that, in practical scenarios, both IFR-MAC and DIFR-MAC force the rational attacker into a jamming-only strategy, or result in performance similar to the jamming-only scenario. These protocols thus provide considerable improvement over allowing false channel reservations in the naïve identity-based channel allocation scheme.

CHAPTER 5

PHYSICAL LAYER BACKGROUND

This chapter provides a primer for wireless communication with focus on the physical layer, where logical symbols (typically bits in computer applications) get converted into physical signal that is suitable for propagation on the communication medium. It also presents the physical-layer framework and defines the terms that I use to describe RONS in Chapter 6.

5.1 Adding Redundancy (in Information-Theoretical Sense)

In coding, communication systems *add redundancy* by generating multiple bits that contain duplicate messages to mitigate the impact of failed delivery on communication reliability. After coding, at the physical layer (where discrete-time communication systems are limited in sampling rate), even though the information theoretically optimal strategy is to have all samples carry discrete, non-overlapping information content,[1] system designers further *add redundancy* in real-life communication practice by having redundant samples that carry overlapping information content as opposed to having all samples contain distinct information content. Transmitters add redundancy by spreading symbols over multiple samples via one-to-many mapping; that is, an input symbol entering the physical layer becomes mapped into multiple samples (and eventually to analog continuous signal) when leaving the transmitter chain. Although adding redundancy increases the processing load, it is commonly used to effectively deal with the real-life physical characteristics of the channel medium: to combat noise and incorporate error control,

[1] In a high SINR-regime, the capacity grows linearly with the transmission rate but grows logarithmically with SINR, and the maximum benefit of adding redundant symbols increases the SINR linearly by the number of samples that the information content spreads across, example of which technique is discussed in Section 6.5.1

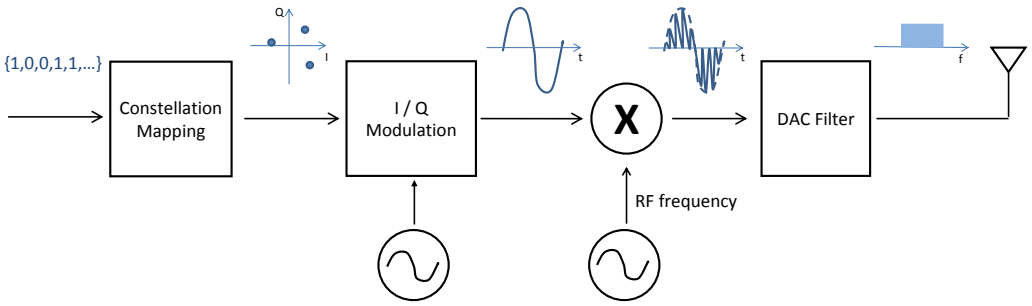

Figure 5.1: Typical transmitter processing chain at the physical layer

to nullify channel fading and synchronization imprecision/error, and to fit the transmissions to the channel constraints, e.g., frequency bandwidth. The channel constraints also come from sharing the medium with other communication users and can either be decided at the application layer by legislative enforcement or at the link layer from a medium access control (MAC) protocol. Thus, I extend the notion of *redundancy* beyond coding (with its typical error control and reliability purpose) and use the term in the information-theoretical sense; the extra samples that carry overlapping or duplicate information content are redundant in information theory but may be *necessary* to meet the real-life constraints.

I call the physical-layer processing blocks that perform one-to-many mapping, and therefore dilute the amount of information content per sample, *redundancy blocks* and the number of outputs of the redundancy block per input *redundancy rate*. Redundancy blocks perform two operations: up-conversion and profile mixing. *Up-conversion*, or *oversampling* maps an input into multiple outputs by repetition. The block then mixes the signal with a *profile*, which defines the redundancy block mapping. In Section 5.2, I present a typical radio chain design and present examples of redundancy blocks.

5.2 Basic Transmitter Design

Figure 5.1 shows a typical standard block design of a radio chain that uses phase modulation (PM) and/or amplitude modulation (AM) as the parameter control choice for the modulation scheme. In the chronological order, a basic transmitter processing chain consists of blocks that perform the follow-

ing functionalities: constellation mapping, I/Q modulation (or quadrature modulation), baseband-to-passband mixer, a digital-to-analog filter, and a radio-frequency (RF) frontend. The constellation mapping block performs bit-to-symbol mappings. Called *M-ary modulation*, it takes $\log_2(M)$ number of bits and convert them into M distinct symbols (greater alphabet size of M provides better bit-throughput rate but also increases the symbol susceptibility to error caused by channel/circuit noise). The next block is the quadrature modulation which takes two quadrature carriers for modulation; typically a sine wave generates the real (I) channel while the cosine wave generates the complex (Q) channel (sinusoidal carriers are used due to their ease of generation). Afterward, a mixer that modulates baseband envelope signal with another sinusoidal with RF frequency converts the signal from baseband to passband and adjusts the transmission's center frequency. Then, the signal goes through a digital-to-analog (DAC) filter, which modulates the signal with a pulse that is designed to fit the signal into the given channel bandwidth, before entering the analog domain of RF frontend (which outputs the signal as an electromagnetic propagation).

Even though it is theoretically impossible to have a time-constrained signal constrained in frequency and vice versa, the pulse shaping DAC filter is designed to minimize the transmission impact beyond the channel. Using a filter to control the bandwidth overspill is very typical among legitimate devices, which adhere to the FCC regulations on spectral mask that controls the bandwidth overspill beyond the channel, which specify the minimum power attenuation outside the accessed frequency band. Due to its common use, I focus on applying RONS on the redundancy block of DAC pulse-shaping filter for evaluations of RONS in Section 6.6.

In this standard design of the transmitter chain, there are three redundancy blocks: the quadrature (I/Q) modulation block, the DAC filter, and, optionally (if transmitter hardware sampling can support additional redundancy), the RF mixer. The redundancy profile for the quadrature (I/Q) modulation is the sinusoidal with the local frequency and that for the DAC filter is the pulse specified by the transmitter.

5.3 Receiver Design

I assume linear receiver and treat interference (both self-interference from fading and external interference from other transmitters) as random noise. At a high level, given a bit-to-samples mapping of the transmitter chain, the corresponding correct receiver that results in zero error with certain amount of channel uncertainty (degree of which depends on the redundancy added, for example, for error correction) performs an inverse mapping to the transmitter chain. With the receiver processing blocks operating in the reverse-chronological order as their counterparts on the transmitter chain, the notion of inverse mapping is straightforward for the processing blocks that perform an injective one-to-one mapping (in other words, they are not redundancy blocks). For the inversion of the other redundancy blocks, the receiver uses a soft-decision correlator and minimum mean squared error (MMSE) decision rule for samples-to-symbol mapping. When noise and interference's statistics are invariant of time, MMSE reduces into *matched filter* (the receiver performs the inverse mapping using the same profile that has been used by the transmitter). Matched filter is also SINR-optimal in Gaussian channels and is independent of both channel state (e.g., does not require channel estimation) and the interferers' strategies. Thus, to demodulate and decode the received signal, the receiver needs to know now not only the exact transmitter chain/strategy but also the profiles that the transmitter uses for the mapping.

CHAPTER 6

REDUNDANCY OFFSET NARROW SPECTRUM

6.1 Chapter Overview

Due to the inherent nature of sharing the medium, wireless communication is vulnerable to signal injection. Correlated jamming, introduced in the 1980s as the information-theoretically optimal interference signal, aims to cancel the target victim signal in contrast to the more traditional jamming approach of adding noise-like interference. The recent surge of antenna-cancellation based technology in a non-security context (including full duplex radio technology using MIMO antennas) has reignited interest in correlated jamming in wireless security. Successful attack of complete cancellation yields zero information about the source transmission signal to the victim receiver and the receiver strategy of recovering bits reduces into coin toss with equal weight. The information about the victim transmission that makes correlated jamming possible also yields easy access to the messages that has been relayed from above the physical layer and compromise the message integrity.

Correlated jamming utilizes antenna-based signal cancellation. In a non-security framework, the field of full duplex with multiple antennas uses such signal cancelation technique; they cancel the signal being transmitted at the receiver location, so that it does not interfere with the receiver reception [36]. In wireless security, others have used the technique in a white-hat approach where *friendly jamming* is used as a defense mechanism for confidentiality against eavesdroppers [37]; correlated jamming, on the other hand, assumes a malicious adversary who injects wireless interference to disrupt communication. As has been demonstrated in prior work in full duplex and friendly jamming, one of the key challenges for signal cancellation is synchronization between the jammer and the target transmitter. Thus, I study the impact of synchronization offsets and compare correlated jamming, *coded jamming*

(that does not need to follow the target transmission real-time), and Gaussian jamming. However, to devise a countermeasure against wireless interference, I assume the strongest threat of correlated jamming.

Typical spreading spectrum solutions against jamming assumes dividing the medium into multiple orthogonal channels and involves channel access randomization [19, 20], so that the choice of accessed channel (among many channel options) is random against attackers; a reactive attacker that observes the victim's channel access and adjust its strategy accordingly can be thwarted by switching channels and having the access duration on a channel be smaller than the attacker's reaction time. Spreading spectrum assumes channelization that provides orthogonal channel access by interleaving the channel use either by time, frequency, or code (processing) and can be effective in ensuring both confidentiality and availability by having the random spreading code/key (from which the channel access information is derived) known only between the source-destination pair involved in the communication. However, spreading spectrum bears a *spreading cost*. In other words, the wireless users consume more resource than no spreading by a factor that is proportional to the number of channel options that the users have for channel access, because the process of spreading symbols entitles either transmitting redundant information (in case of code-based spreading) or reserving more resource than the user uses at a time (in cases of frequency or time-based spreading), and thus has a negative impact on the throughput rate performance.

I introduce a novel physical-layer technology, Redundancy Offset Narrow Spectrum (RONS), that effectively counters both passive and active wireless attacks. RONS is *narrowband spectrum* since it does not require the spreading cost of consuming wireless resource proportionally to spreading gain; it uses the built-in physical-layer blocks of the communication chain but only adds phase offsets or cyclic delays (which values are only known among the legitimate key holders). Fully implemented at the physical layer, RONS also does not rely on randomization of the physical channel access. In other words, RONS counters threats even when the attacker knows the physical channel location of the signal transmission; in fact, I assume that the attacker does not waste its power accessing other channel to model the worst-case impact.

6.2 System Model and Assumptions

Due to the wide adoption in wireless communication community, I design RONS based on the basic communication design described in Section 5.2 and the receiver strategy in Section 5.3. The transmitter-receiver pair a priori agrees on a secure key [24, 26]. However, there is no collision-preventing channel coordination between the simultaneous transmitters at the medium access control (MAC) layer (MAC-layer approach to mitigate interference is an active field [8] and its physical layer counterparts for orthogonal medium access are described in Section 6.5.1).

Assuming an additive white Gaussian wireless channel model with numerous noise sources and limited fading with clear line-of-sight channel path (e.g., for evaluation in Section 6.6, RONS uses filtering that is robust to fading), the source transmitter coexists with $n-1$ other transmitters, consisting of a network of n users, sharing a bandwidth of W. In this framework, the user accesses the entire bandwidth by outputting samples at the rate of W before the RF frontend and transmits at all time with full queue. In contrast, currently typical channelization schemes discussed in Section 6.5.1 have an average application-layer goodput rate of $\frac{W}{n}$ at best, which requires correct and orthogonal channelization at MAC-layer and above. The single-channel setting (where the entire bandwidth is accessed) also models the worst-case collision-behavior among multiple coexisting transmitters.

I use the effective signal-to-interference-and-noise ratio (SINR) for the performance metric, since the greater the effective SINR at the receiver the better the reliability and rate performance. For instance, Shannon-Hartley theorem provides the theoretical upper bound on communication rate performance (R) in information theory: $R = W \log [1 + \text{SINR}]$. The effective SINR metric both enables us to abstract away from the particulars of the physical layer design such as the modulation and coding scheme and reduces the problem by a degree of freedom, since I no longer need to consider how many transmitters are coexisting but rather what their collective impact on the receiver is (e.g., the transmitted power on the channel); for example, one interfer with five times the power budget has the same impact on the receiver as five interfers that have identical channel with equal power budget.

6.3 Attack Model

I aim for security by design and consider a strong attacker model where the attacker knows not only the transmitter chain strategy but also the physical frequency and time location of the source transmitter's medium access (that is, all of the attacker's emitted power impacts the transmission as interference). In specific, I distinguish three-level of jamming attackers by their capabilities and information advantage that they have on the victim communication system (while assuming that the adversary has the capability to maximally use the information advantage for their jamming strategies). The strongest is the *correlated jammer* who not only knows the victim transmitter's physical-layer processing chain information but also the data input that has been relayed from the upper layers (or equivalently, correlated jammer is a processing-powerful reactive jammer with negligible reaction time). If an attacker does not know the data input but only the transmitter's physical-layer strategy, then it becomes a *coded jammer*. Lastly, an attacker who does not have any information injects interference signal that is independent to the victim transmission signal, in which case Gaussian signaling has the most detrimental effect on the victim transmission [22, 23], and thus the attacker is a *Gaussian jammer*. I establish the attacker model in this section, analyze and compare the jamming effect of the three jamming strategies in Section 6.4. Afterward, I introduce RONS that effectively counters the strongest attack of correlated jamming.

6.3.1 Active Attack on Availability

To model the worst-case scenario for interference, I model the interferer to be a malicious jammer, an interferer whose sole goal is to degrade the performance of the source transmitter. The optimal jammer strategy, or the interference signal that results in the minimum capacity rate performance, is linearly correlated with the target source signal if the jammer has a power budget comparable to that of the target source transmitter [22,23,38]. Based on these studies in information theory, I consider such *correlated jamming* attack. A correlated jammer cancels the source signal by injecting the same signal but only inverted (it *cancels* the signal by causing destructive interference, in contrast to the more conventional use of jamming to add noise).

If successful, the received signal becomes uncorrelated with the transmitted signal (and the received signal does not contain any information about the transmitted signal) and the capacity becomes zero. Correlated jamming has also recently been studied in a system-oriented work [39].

A power-constrained attacker needs to only match the transmission power in order to force zero information transfer at the receiver. In fact, inverted transmission that exceeds beyond cancellation leaks information to the victim receiver and may also be helpful in detecting the attacker, as I study in Section 6.4.1. Therefore, a power-efficient correlated jammer only matches the transmission power.

Correlated jamming is the most dangerous when frequency, phase, and amplitudes are matched with the victim's transmit signal. Natural frequency and phase drift and jitter can be matched by the use of aggressive locking mechanism such as by using phase-locked loop; this attacker model challenges the notion of (natural) indelible marks on transmissions using sinusoidal signals [25]. Such strong attacker needs to precisely know about the transmission and be able to quickly react. A weaker attack is *coded jamming* where an attacker only knows about the physical-layer transmission strategy (such as the redundancy profiles) and do not actively listen to the transmission. Nevertheless, a coded jammer is not independent to the transmission but is independent to the data input, whereas a correlated jammer is dependent on both the data input and the transmission chain (and thus correlated to the output transmission signal). A coded jammer, which has a much less stringent requirement than a correlated jammer, is much stronger than simple Gaussian jamming in real-world communication practices due to adding redundancy and the corresponding receiver strategy to use the redundancy to decode the symbol (as is explained in Chapter 5). Section 6.4 studies interference and compares between correlated, coded, and Gaussian jamming.

6.3.2 Passive Attack on Confidentiality

Since they know the victim transmitter's physical-layer strategy, both correlated jammer and coded jammer can correctly decode the message and breach privacy. The Gaussian jammer, on the other hand, does not know the pro-

cessing chain and cannot decode the message. RONS effectively preserves privacy against the advantageous attackers, as I study in Section 6.6.1.

6.4 Jamming Interference Analysis

For interference analysis, I perform Monte-Carlo simulations using MATLAB. While assuming the system model in Section 6.2 and the attack model in Section 6.3, I use binary phase shift keying (BPSK) modulation and a root raised cosine filter (RRCF) with a filter order of 256. The natural SNR (without interference) is 10 dB.

6.4.1 The Usefulness of Information Advantage for Attackers

In Figure 6.1(a), I compare the three jamming strategies of correlated, coded, and Gaussian against wireless availability and observe that knowing the victim transmitter strategy gives the jammer advantage and capability to inflict more damage on the network. While varying the attacker power budget with respect to that of the legitimate user and assuming perfect synchronization in phase and frequency, correlated jamming that cancels the signal power has the biggest impact on wireless availability by yielding the highest error rate for the legitimate system; coded jamming also has a more detrimental effect than Gaussian jamming. When the attacker power is matched to the legitimate transmitter's power, correlated jamming results in an error rate of 0.5 (i.e., no information due to complete cancellation of the transmitted signal) while coded jamming results in an error rate of 0.25 (since coded jamming sends random symbols imitating the physical-layer chain of the source transmitter and since the source transmitter uses BPSK with alphabet size of two, there is 0.5 chance of coded jammer sending conflicting symbols and, when that happens, there is a conditional probability of 0.5 for the event that the receiver tunes into and decodes the symbol that coded jammer sent; thus, $0.5 \cdot 0.5 = 0.25$). On the other hand, the legitimate user performs very well against transmission-independent Gaussian jamming since I incorporate redundancy at the physical-layer and the receiver uses soft-decoding to use the information of multiple samples to decode a symbol (as described in Section 5.3) and thus effectively mitigates transmission-independent noise;

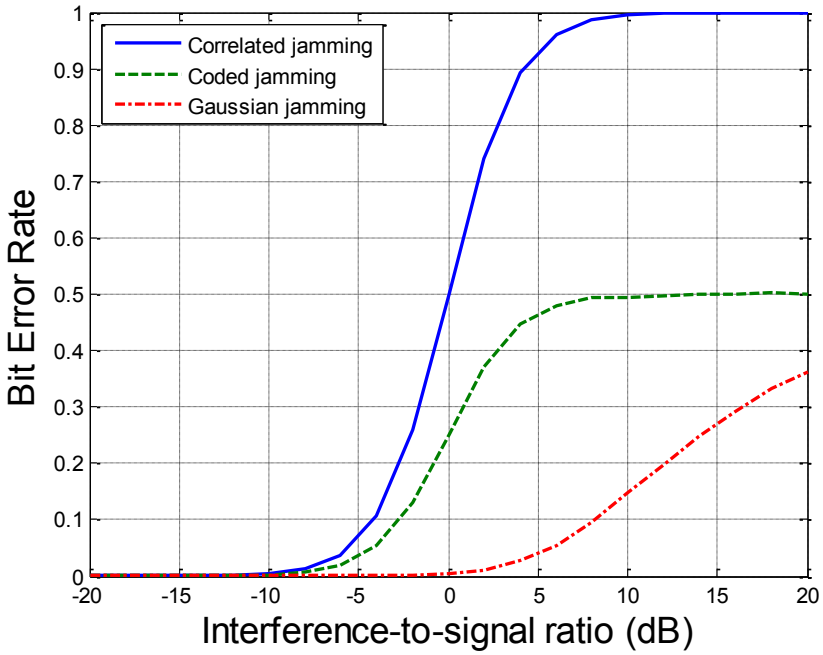

(a) Comparison between correlated, coded, and Gaussian jamming while varying attacker power budget

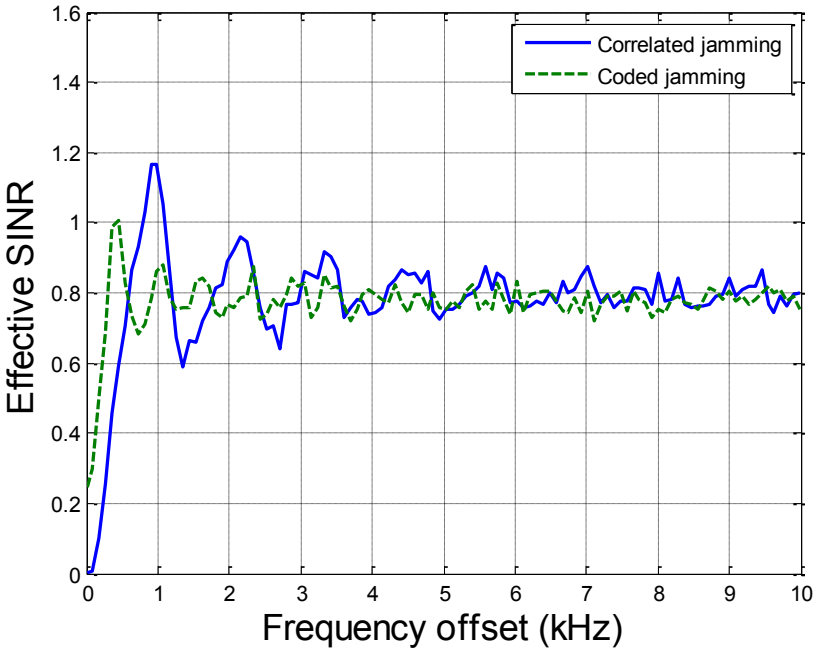

(b) Correlated jamming becomes like coded jamming when frequency is not synchronized

Figure 6.1: Interference analysis

coded and correlated jamming, on the other hand, knows the MMSE-based receiver strategy and customizes its signal injection accordingly. Bakr and Mudumbai [40] also suggests that Gaussian jamming is less effective than the transmission-customized jamming (their white-hat approach uses jamming for defense against an eavesdropping attacker).

As the interference power grows, the error rate for correlated jamming converges to one and that for coded jamming converges to 0.5 because jamming transmission dominates the channel and the receiver tunes to the jamming symbol (for correlated jamming, the inverted symbol is the other symbol of BPSK and for coded jamming, the random symbol dominates the transmitted symbol). While an error rate of 0.5 effectively reduces into a coin flip and corresponds to no-information, an error rate exceeding 0.5 actually yields information to the legitimate user since the correlated jammer sends an (inversely) correlated symbol to the transmission. If a correlated jammer is concerned about yielding any information (e.g., legitimate user uses high error rate for correlated jammer detection), then the correlated jammer can adjust its transmission power to match that of the legitimate user.

6.4.2 The Effect of Frequency Offset on Correlated Jamming

I study the case when the jamming signal and the victim transmission signal is not synchronized. While I vary the power amplitude in Section 6.4.1, I now vary the frequency offset between the two signals, as hardware oscillators naturally operate at different frequencies and have unique frequency drift and jitter. Figure 6.1(b) displays the result while the jamming power amplitude is matched to that of the legitimate transmitter (transmission-independent Gaussian jamming has a constant SINR of 0.909 and is not plotted). When perfectly frequency-synchronized, correlated jamming yields zero SINR since it completely cancels the legitimate source signal; correlated jamming yields some information about the source transmission and the effective SINR is 0.25. Within 0.5 kHz offset, the jamming effect gets substantially decreases and eventually settle at an effective SINR of 0.8, which performance is still better than Gaussian jamming in the attacker's malicious perspective. As the frequency offset grows, correlated jamming converges to coded jamming because the transmission signal does not effectively get cancelled. To accom-

modate the difference in operating frequencies between hardware oscillators, IEEE 802.11 allows a center frequency error of ±20 ppm [41]. When operating in GHz-band, such frequency offsets are enough for correlated and coded jamming to reach the steady-state effective SINR of 0.8; for instance, IEEE 802.11a channel 165 at carrier frequency 5.825 GHz tolerates frequency offset of 233 kHz around the center frequency. Although all oscillators operate at their own unique frequencies and it is natural to have frequency offsets, an attacker can use a frequency locking scheme with an aggressive use of phase-locked loops to synchronize frequency and perform the likes of correlated jamming for signal cancellation.

6.5 Redundancy Offset Narrow Spectrum (RONS)

6.5.1 Motivation and Comparison with Spreading Spectrum

To embrace the coexistence of simultaneously transmitting wireless systems, communication researchers perform channelization and divide the medium into multiple channels. The channels are designed to be orthogonal to each other, so that the transmissions using different channels do not result in collision and interfere with each other. Typically, the community ensures orthogonality of channels by interleaving their access by time, frequency, and code (processing).

Unfortunately, the current approaches that implement orthogonal channelization negatively affect the individual user's data rate performance. By sharing the medium with other users (with equal priority) via orthogonal channelization, the expected individual user's rate performance is inversely proportional to the number of coexisting users, in the best-case scenario, while the overall network performance (the sum aggregation of individual user performances) remains the same. This is because frequency or time-based channelization divides the respective resources by the number of users or more, while all the medium resources could have been used by the source transmitter if other transmitters were not present.

On the other hand, code-based channelization introduces additional redundancy and consumes more medium resources than it would have needed if there were no channel division. For instance, a typical realization of code

division multiple access (CDMA) involves a redundancy block called direct sequence spread spectrum (DSSS); DSSS temporally spreads the symbol by mapping a sample into multiple chips, and the number of chips per symbol is called *processing gain* or *spreading gain* (which is the redundancy rate of the DSSS redundancy block). If the system retains the sampling rate within the block, then it consumes the amount of time that is larger than if there were no DSSS block to transmit the same amount of information (the time increase is proportional to the processing gain). Alternatively, DSSS increases the sampling rate (so that the chip rate is greater than the input symbol rate by the processing gain), which consumes a proportionally larger bandwidth. DSSS increases the SNR by the processing gain and effectively mitigates physical interference by first, carefully choosing the set of DSSS profiles (chip sequences that map the symbol to chips) so that they are orthogonal to each other, and second, the receiver combining the information of multiple chips to decode the corresponding sample. By sharing the random spreading code only among the participating parties, DSSS can be helpful in both availability (interference, like noise, gets mitigated) confidentiality (correct code is necessary to decode the message). In contrast to the spreading spectrum technology which entitles the spreading cost in rate, RONS minimizes intra-channel interference by emulating orthogonal access without the drawback in rate performance.

6.5.2 Redundancy Offset Narrow Spectrum Scheme

The goal of RONS is to mitigate interference without the spreading cost in data rate. RONS is similar to DSSS in that it is processing based, but it uses the redundancy blocks that are already in place of the transmitter chain, as opposed to introducing a new set of redundancy as DSSS does. Given the profile of a pre-existing redundancy block, RONS creates multiple profiles by adding *cyclic phase offsets* (or *cyclic delay*), which I denote with ϕ.[1] With the offset values chosen so that the generated redundancy profiles have zero correlation with one another (so that they are statistically orthogonal to each other), the use of a profile for signal processing yields statistically

[1] If the redundancy profile of the block that RONS is deployed is odd and periodic, then $\phi = \pi$ is equivalent to the signal that a correlated jammer will transmit once the victim transmission is present.

independent channel path from using any other profiles. I call the set of profiles generated using such phase offsets *RONS channels*. In other words, the cross correlation between *any* two signals using distinct RONS channels is very small.

Deciding on RONS channels depends on the processing operation of the redundancy block. The phase offset selection for RONS channel generation is straightforward when the redundancy block only performs upconversion and element-by-element mapping, in which case, I can observe the correlation between the generated redundancy profiles after adding the cyclic phase offsets. One common use of RONS is the quadrature modulation, described in Section 5.2, where one channel (the I channel) uses sine profile and the other channel (the Q channel) uses sine with $\phi = \pi/2$. However, in contrast to quadrature modulation and DSSS, for redundancy blocks that do not merely perform element-by-element mapping, such as the DAC filter (which involves convolution), the RONS phase offset selection not only depends on the redundancy profiles but also with the input of the redundancy block.

After deciding on the set of pair-wise mutually uncorrelated RONS channels, the transmitter-receiver pair choose a random RONS channel (which can be derived from the secret key), so that a correlated jammer (and an eavesdropper) cannot target the correct RONS channel to compromise the signal. For reactive jammers who sense the channel, I incorporate randomization by channel hopping across RONS channels. Since the RONS phase offsets can be added per symbol basis, this requires that the attacker cannot respond within a symbol.

After sharing the initial key, RONS uses the data transmission content to update the key; specifically, it uses a hash function of a packet consisted of multiple symbols (and thus spread across multiple RONS channels). As the packet size (in symbols) increases, the probability of attacker guessing the correct hopping sequence and extracting the data decreases exponentially (in Section 6.6.1, where I study RONS effectiveness in confidentiality, I see that tuning in another RONS channel leads to incorrect decoding).

I also introduce a design parameter τ that controls the tradeoff between the statistical orthogonality between RONS channels and the number of RONS channels that the system can afford, since more RONS channels will make the channel-guessing attacker more difficult to make a correct guess. In other words, I allow that the inter-RONS-channel correlation to be as great as τ

in selecting RONS channels.

Furthermore, RONS uses natural binary code when it maps from bits to symbols rather than the more popular Gray code; in natural binary code, the number index of the bits proportionally increases with phase. By having the adjacent symbols only differ by a bit when mapping the bits into symbols, Gray code is also popularly used along with error correction since, if symbol error occurred, it is more likely that the random noise yields closeby symbols than symbols that are further away on the constellation diagram. However, correlated jamming injects inverted signal with a phase offset of π (on the opposite side of the constellation diagram) and not random noise, and thus, it is more likely that the decoded symbol is further away from the transmitted symbol, especially when correlated jammer transmits with higher power than the legitimate transmission. When the attacker power dominates, the coding scheme does not affect the symbol error rate (there will almost always be errors because the receiver tunes on the attacker signal) but it affects the bit error rate, and RONS uses natural binary coding to lower the bit error rate after the symbol-to-bit mapping on the receiver. Figure 6.2 displays the bit error rate varying the coding alphabet size while assuming a modulation scheme that, for every symbol representing a message, there is another symbol representing distinct message with a phase offset of π; the correlated jammer dominates and the receiver decodes the inverted symbol. With RONS using natural binary code, the bit error rate becomes inversely proportional to the log_2 of the alphabet size, or the number of bits that gets mapped to the symbol.

6.6 RONS Evaluation

For RONS evaluation, I perform Monte-Carlo simulation using MATLAB. For consistency, I use the same parameters as I did in Section 6.4. Namely, I use binary phase shift keying (BPSK) modulation and a root raised cosine filter (RRCF) with a filter order of 256; the natural SNR (without interference) is 10 dB.

I apply RONS on the RRCF pulse-shaping filter, as was mentioned in Section 5.2. RRCF filter is suitable for RONS because it is a digital filter operating in discrete-time domain (easier to implement, stable, and has a

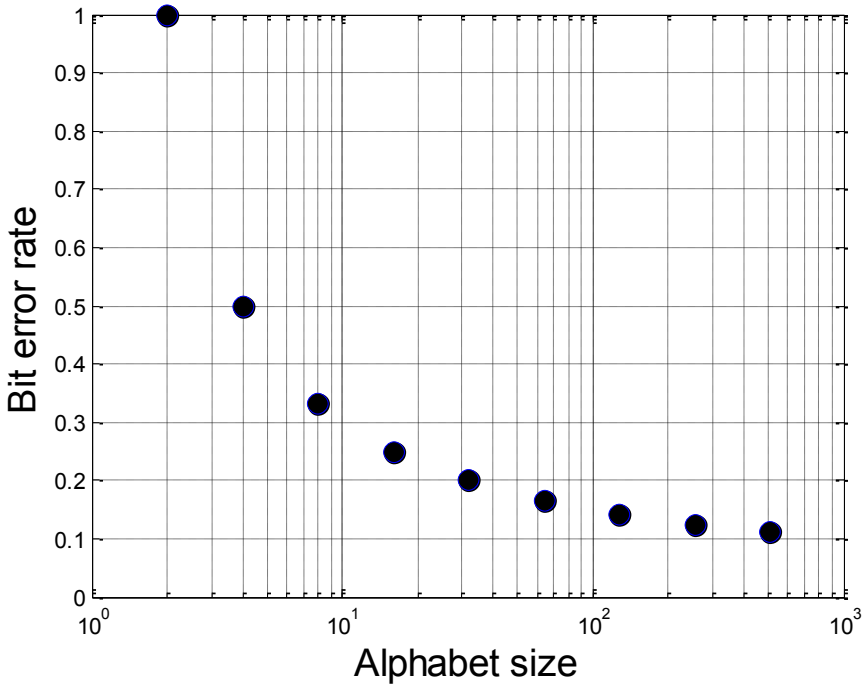

Figure 6.2: Bit error rate with the bit-to-symbol alphabet size when correlated jammer dominates

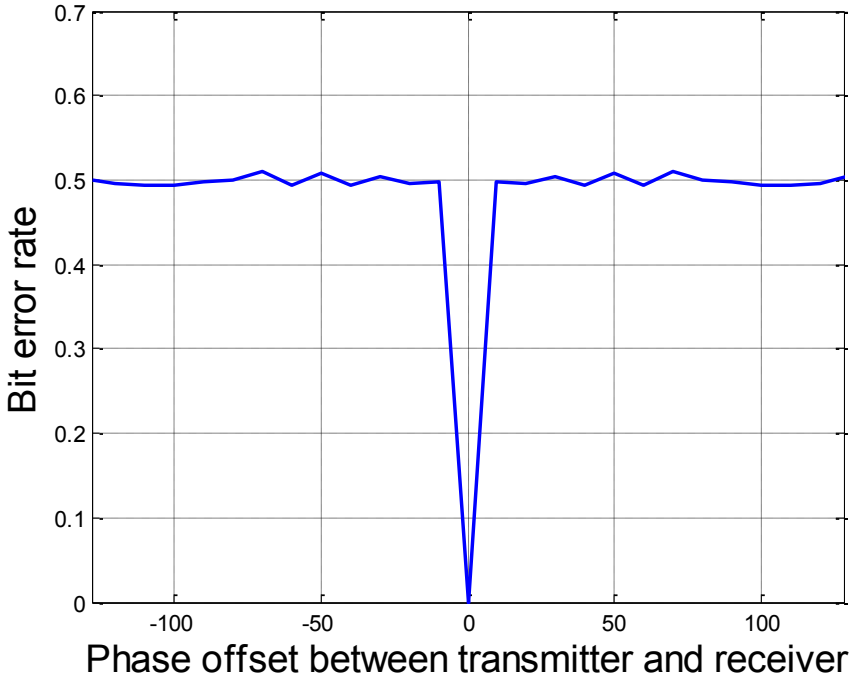

Figure 6.3: RONS channels and correct decoding

linear phase characteristic) and is robust to fading (and thus agrees with the channel model). I use correlation threshold τ of 0.01 and the number of RONS channels is 8. Random RONS channel hopping is simulated, and the correlated jammer guesses a RONS channel.

6.6.1 RONS for Confidentiality

By incorporating randomization in the phase offsets (the RONS channels), RONS can be helpful for protecting confidentiality. Figure 6.3 shows communication reliability while varying the phase offsets between a transmitter and a (potentially malicious) receiver. Without knowing the correct phase offset, the receiver does not have capability of decoding the transmission, as the error rate quickly approaches 0.5, which yields no information about the transmission. Therefore, if the source transmitter can keep the RONS channel random and secret, then the receivers failing to tune into the correct RONS channel fails to decode the transmitted data correctly.

6.6.2 RONS for Availability

I study RONS performance against correlated jamming. In Figure 6.4(a), I study the effective SINR while varying the signal-to-interference ratio (SIR) (where the interference model is the worst-case of corrrelated jamming). Without RONS, the effective SINR becomes zero when interference power is greater than the signal power meaning that the received signal has zero information about the transmitted signal, since it effectively got cancelled. Unsurprisingly, the performance monotonically increases as the SIR increases. Enabling RONS prevents the transmitted signal from getting completely cancelled and outperforms the case when RONS is disabled. Figure 6.4(b) plots the effective SINR gain of enabling RONS compared to that when RONS is disabled and the jamming attacker knows the transmission and is synchronized (for example, gain of 100% indicates that the performance of enabling RONS is twice as good as that of disabling RONS). With the gain being ∞ when SIR $<$ 0 dB (since the reference case of disabling RONS yields zero performance), the gain decreases as SIR increases. Even though RONS performance eventually converges to the correlated jamming performance both when SIR grows (as the remaining signal after cancellation is still big) and when SIR shrinks (as interference power simply overwhelms the signal power), RONS performance against correlated jamming is the most effective when they have comparable amount of power, i.e., SIR is close to 0 dB, and thus when correlated jamming has the most detrimental impact on the victim without a countermeasure (as shown in Section 6.4.1).

6.7 Chapter Summary

RONS provides a novel rate-efficient scheme to incorporate randomization for wireless security. Against a very strong attack of frequency-synchronized correlated jamming, it provides throughput even when the correlated jammer's power budget exceeds that of the source transmitter (with no countermeasure, the correlated jammer forces zero capacity). RONS is the most effective when correlated jamming is the most effective (without a countermeasure deployed) with its transmitter power comparable to that of the source transmitter; when the jammer power budget exceeds -2 dB of that of the source transmitter, the effective gain is greater than 70%.

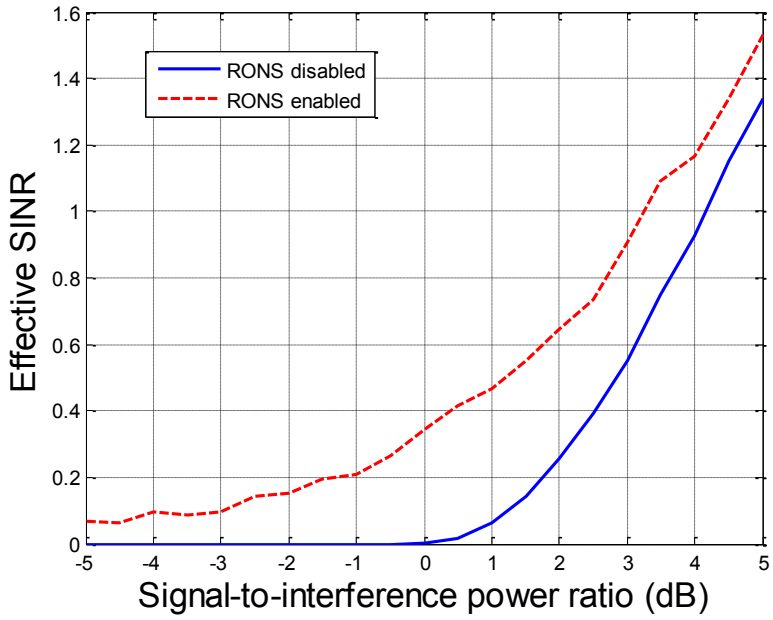

(a) The effect of correlated jamming and RONS' mitigation

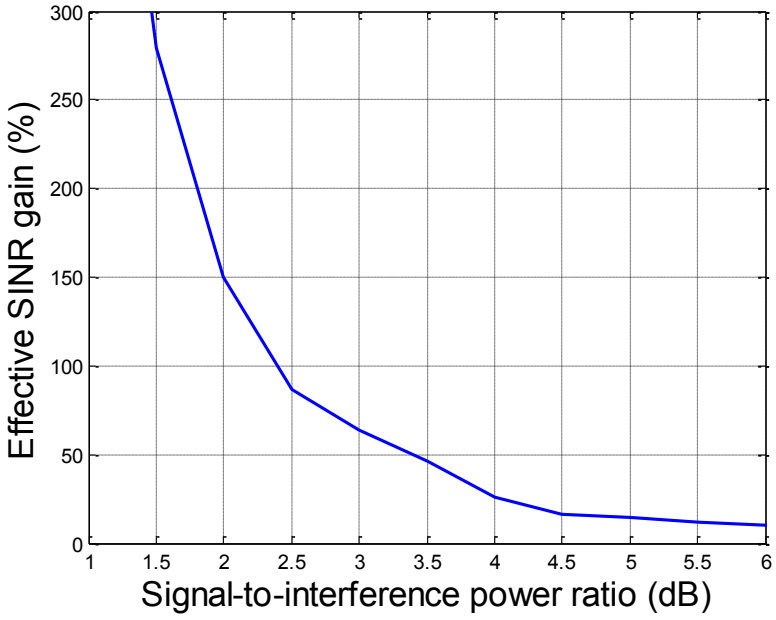

(b) RONS' performance gain over no countermeasure

Figure 6.4: RONS' performance for jamming mitigation

CHAPTER 7

DISSERTATION SUMMARY

MAC-layer protocols, used for efficiently supporting multiple coexisting transmissions given a shared channel resource, are vulnerable to node failure or misbehaving. In particular, I investigate the vulnerability of reservation-based MAC protocols where users first send control messages to reserve channels and then transmit data on those reserved channels (for example, in IEEE 802.11 four-way handshaking protocol for virtual carrier sensing, the pair of request-to-send (RTS) and clear-to-send (CTS) are control messages that notify the channel reservations to other simultaneously existing transmitters). When the network is compromised and a misbehaving user has as much capability as the protocol-complying users, I consider three threats on such MAC protocols to model failure: jamming on known control channel (disrupting the control communication and disabling channel reservation), using control message to facilitate power-efficient jamming on data channel (only targeting the channels that are occupied by legitimate transmissions, as opposed to wasting power on empty channels), and injecting bogus control messages (reserving channels without using it and thus wasting resources). The colluding attackers also compromise the network and have the same access as any other legitimate member.

To counter the first two attacks of control-channel jamming and control-aware data jamming, SimpleMAC relies on jamming-resistant unicast transmission and then uses a control communication mechanism that multicasts (as opposed to broadcasting) and shares the control messages containing the reservation information to subset of users. In other words, SimpleMAC diverges from the traditional binary approach of sharing the control message with either everybody or nobody and adjusts the strategy according to the user behaviors (and the performance outcomes); it minimizes interference from both protocol-abiding benign users and misbehaving attackers.

For the threat on MAC protocols that involves sending bogus control mes-

sages, IFR-MAC holds the users accountable by checking their use of reserved channels via feedback and adapting the channel allocation accordingly. In other words, assuming a flexible channelization scheme with varying carrier frequency and flexible bandwidth access, IFR-MAC assigns channel resource according to the users' truthfulness in reservations. I also study the distributed case where all users agree on a channelization scheme, in which case, the channel fluctuation (e.g., fading) negatively impacts the performance of the protocol.

When MAC fails to support orthogonal channel access between network users or in a single-channel scenario when the network medium is not divided into multiple orthogonal channels (and thus, the attacker knows the exact location of the victim user's accessed channel), RONS counters the information-theoretically optimal correlated jamming and further helps ensuring confidentiality. It operates on built-in redundancy blocks (and thus, does not incur additional cost in rate performance) and adds cyclic phase shifts to emulate randomization on channel access and spreading spectrum.

REFERENCES

[1] S. T. Chung, S. J. Kim, J. Lee, and J. Cioffi, "A game-theoretic approach to power allocation in frequency-selective gaussian interference channels," in *IEEE ISIT*, June 2003, pp. 316–316.

[2] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," *IEEE JSAC*, vol. 25, no. 3, p. 517, Apr. 2007.

[3] L. G. Roberts, "ALOHA packet system with and without slots and capture," *SIGCOMM CCR*, vol. 5, no. 2, pp. 28–42, Apr. 1975.

[4] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: Real vulnerabilities and practical solutions," in *USENIX Security Symposium*, Aug. 2003, pp. 15–28.

[5] Y. Hu, A. Perrig, and D. Johnson, "Ariadne: A secure on-demand routing protocol for ad hoc networks," *Wireless Networks*, vol. 11, no. 1, pp. 21–38, Jan. 2005.

[6] R. Negi and A. Rajeswaran, "DoS analysis of reservation based MAC protocols," in *IEEE ICC*, 2005.

[7] V. Gupta, S. Krishnamurthy, and M. Faloutsos, "Denial of service attacks at the MAC layer in wireless ad hoc networks," in *MILCOM*, vol. 2, Oct. 2002, pp. 1118–1123.

[8] B. Awerbuch, A. Richa, and C. Scheideler, "A jamming-resistant MAC protocol for single-hop wireless networks," in *PODC*, Aug. 2008, pp. 45–54.

[9] G. Alnifie and R. Simon, "A multi-channel defense against jamming attacks in wireless sensor networks," in *Q2SWinet*, Oct. 2007, pp. 95–104.

[10] W. Xu, W. Trappe, and Y. Zhang, "Channel surfing: Defending wireless sensor networks from interference," in *ACM IPSN*, Apr. 2007, pp. 499–508.

[11] K. Firouzbakht, G. Noubir, and M. Salehi, "On the capacity of rate-adaptive packetized wireless communication links under jamming," in *ACM WiSec*, Apr. 2012, pp. 3–14.

[12] "IEEE std 802.11-2007," *IEEE Std. Association*, 2007. [Online]. Available: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=4248378

[13] A. Cardenas, S. Radosavac, and J. Baras, "Performance comparison of detection schemes for MAC layer misbehavior," in *IEEE INFOCOM*, May 2007, pp. 1496–1504.

[14] P. Kyasanur and N. Vaidya, "Selfish MAC layer misbehavior in wireless networks," *IEEE Transactions on Mobile Computing*, vol. 4, no. 5, pp. 502–516, 2005.

[15] M. Raya, J.-P. Hubaux, and I. Aad, "DOMINO: A system to detect greedy behavior in IEEE 802.11 hotspots," in *MobiSys*, June 2004, pp. 84–97.

[16] L. Dryburgh and J. Hewitt, "Signalling system no. 7 (SS7/C7): Protocol, architecture, and services," *Cisco Press*, Aug. 2004.

[17] "IEEE std 802.16-2009," *IEEE Std. Association*, 2009. [Online]. Available: http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=5062428

[18] "IEEE std 802.15.1-2005," *IEEE Std. Association*, 2005. [Online]. Available: http://standards.ieee.org/findstds/standard/802.15.1-2005.html

[19] R. Pickholtz, D. Schilling, and L. Milstein, "Theory of spread-spectrum communications–A tutorial," *IEEE Transactions on Communications*, pp. 855–884, May 1982.

[20] M. Simon, J. Omura, R. Scholtz, and B. Levitt, *Spread Spectrum Communications Handbook*. McGraw-Hill: New York, Mar. 1994.

[21] H. Li and I. Marsland, "A comparison of rateless codes at short block lengths," in *IEEE ICC*, May 2008, pp. 4483–4488.

[22] T. Basar, "The Gaussian test channel with an intelligent jammer," *IEEE Trans. Info. Theory*, vol. 29, no. 1, pp. 152–157, Jan. 1983.

[23] A. Kashyap, T. Basar, and R. Srikant, "Correlated jamming on MIMO Gaussian fading channels," in *IEEE ICC*, vol. 1, June 2004, pp. 458–462.

[24] J. Chiang and Y. Hu, "Dynamic jamming mitigation for wireless broadcast networks," in *IEEE INFOCOM*, Apr. 2008, pp. 1211–1219.

[25] L. Baird, W. Bahn, M. Collins, M. Carlisle, and S. Butler, "Keyless jam resistance," in *Information Assurance and Security Workshop*, June 2007, pp. 143–150.

[26] M. Strasser, S. Capkun, C. Popper, and M. Cagalj, "Jamming-resistant key establishment using uncoordinated frequency hopping," in *IEEE SSP*, May 2008, pp. 64–78.

[27] P. Murphy, A. Sabharwal, and B. Aazhang, "Design of WARP: A flexible wireless open-access research platform," in *Proceedings of EUSIPCO*, Sep. 2006, pp. 53–54.

[28] C. Tellambura and V. Bhargava, "Unified error analysis of DQPSK in fading channels," *Electronics Letters*, vol. 30, no. 25, pp. 2110–2111, Dec. 1994.

[29] T. Tjhung, C. Loo, and N. Secord, "BER performance of DQPSK in slow Rician fading," *Electronics Letters*, vol. 28, no. 18, pp. 1763 –1765, Aug. 1992.

[30] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White space networking with wi-fi like connectivity," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 4, pp. 27–38, 2009.

[31] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat, "Learning to share: Narrowband-friendly wideband networks," in *ACM SIGCOMM 2008*, Seattle, WA, August 2008.

[32] L. Yang, W. Hou, L. Cao, B. Y. Zhao, and H. Zheng, "Supporting demanding wireless applications with frequency-agile radios," in *Proc. of NSDI*, 2010.

[33] K. Chintalapudi, B. Radunovic, V. Balan, M. Buettener, S. Yerramalli, V. Navda, and R. Ramjee, "WiFi-NC: WiFi over narrow channels," in *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI'12. Berkeley, CA, USA: USENIX, 2012.

[34] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.

[35] H. Chan, H.-C. Hsiao, A. Perrig, and D. Song, "Secure distributed data aggregation," *Found. Trends Databases*, vol. 3, no. 3, pp. 149–201, Mar. 2011.

[36] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving single channel, full duplex wireless communication," in *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '10. New York, NY, USA: ACM, 2010, pp. 1–12.

[37] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, "They can hear your heartbeats: Non-invasive security for implantable medical devices," in *Proceedings of the ACM SIGCOMM 2011 Conference*, ser. SIGCOMM '11. New York, NY, USA: ACM, 2011. [Online]. Available: http://doi.acm.org/10.1145/2018436.2018438 pp. 2–13.

[38] M. Medard, "Capacity of correlated jamming channels," in *Allerton Conference on Communications, Computing and Control*, 1998.

[39] C. Pöpper, N. O. Tippenhauer, B. Danev, and S. Capkun, "Investigation of signal and message manipulations on the wireless channel," in *Proceedings of the 16th European Conference on Research in Computer Security*, ser. ESORICS'11. Berlin, Heidelberg: Springer-Verlag, 2011. [Online]. Available: http://dl.acm.org/citation.cfm?id=2041225.2041229 pp. 40–59.

[40] O. Bakr and R. Mudumbai, "A new jamming technique for secrecy in multi-antenna wireless networks," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, 2010, pp. 2513–2517.

[41] "IEEE std 802.11a," *IEEE Std. Association*. [Online]. Available: http://standards.ieee.org/getieee802/download/802.11a-1999.pdf