Securing Wireless Medium Access Control Against Insider Denial-of-Service Attackers

Sang-Yoon Chang
Advanced Digital Sciences Center
Singapore 138632
sychg@adsc.com.sg

Yih-Chun Hu Zhuotao Liu
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
{yihchun,zliu48}@illinois.edu

***Abstract*—In a wireless network, users share a limited resource in bandwidth. To improve spectral efficiency, the network dynamically allocates channel resources and, to avoid collisions, has its users cooperate with each other using a medium access control (MAC) protocol. In a MAC protocol, the users exchange control messages to establish more efficient data communication, but such MAC assumes user compliance and can be detrimental when a user misbehaves. An attacker who compromised the network can launch a two-pronged denial-of-service (DoS) attack that is more devastating than an outsider attack: first, it can send excessive reservation requests to waste bandwidth, and second, it can focus its power on jamming those channels that it has not reserved. Furthermore, the attacker can falsify information to skew the network control decisions to its favor.**

To defend against such insider threats, we propose a resource-based channel access scheme that holds the attacker accountable for its channel reservation. Building on the randomization technology of spread spectrum to thwart outsider jamming, our solution comprises of a *bandwidth allocation* component to nullify excessive reservations, *bandwidth coordination* to resolve over-reserved and under-reserved spectrum, and *power attribution* to determine each node's contribution to the received power. We analyze our scheme theoretically and validate it with WARP-based testbed implementation and MATLAB simulations. Our results demonstrate superior performance over the typical solutions that bypass MAC control when faced against insider adversary, and our scheme effectively nullifies the insider attacker threats while retaining the MAC benefits between the collaborative users.

I. Introduction

From smartphones and wearable devices to Internet of Things (IoT)-based appliances, the demand for wireless communication keeps increasing. However, wireless communication consumes bandwidth, and the users inherently share a medium; therefore, one's signal becomes another's interference when they collide in channel access. To cope with the increased demand in wireless, the recent developments in radio technology such as cognitive radio and software-defined radio facilitate flexible and dynamic access and enable better adaptation to the ongoing traffic for greater spectral efficiency. These sophisticated technologies, however, increase the complexity of radio operations and network management, further necessitating a complementary platform to coordinate the channel access when supporting multiple users.

To cope with the dynamism in channel access and avoid inter-user collision, *medium access control* (MAC) protocols have long been designed to share a medium among multiple transmitters. Since wireless networks lack collision detection (in contrast to wired networks), they use MAC protocols that coordinate channel use through explicit messages. This process involves *control communication*, in which users *reserve* channels and notify the network of their transmission intentions before the data transmissions. The MAC protocol ensures collision avoidance among network users by ensuring that each user reserves channels separated in frequency, time, or processing/coding (we focus on frequency channel access in this work although our approach generalizes to other channelization techniques).

MAC is designed for protocol-compliant users. However, we study the network behavior when some network users deviate from the protocol. There are three types of deviations that we may contemplate: accidental failures, selfish users, and malicious users. Previous work has shown that the Nash equilibrium when all users are selfish is to disable MAC exchanges and have each user access the entire bandwidth all the time [1]. The success of network protocols such as WiFi and TCP demonstrates that selfishness is not as prominent in real life as game theorists fear; instead, most users are protocol-compliant, and protocols based on user cooperations can yield overall network gain. Thus, to take advantage of the cooperative nature of most users, we focus on a group of compliant users sharing spectrum with malicious users (whose goal is to disrupt the network operations and have the option of behaving like a greedy user if other attacks fail).

In the presence of malicious users, much prior work in wireless MAC relies on the defense at the (virtual) network perimeter, e.g., filtering and blacklisting. Even when the network is compromised, prior work [2]–[4] focuses on detecting and isolating attackers based on their identities/credentials to make the attacker's insider capabilities obsolete, effectively reducing them into outsiders (whose identities grant no or limited capabilities and rights). Such perimeter-based approach works well when the network boundaries and the user behaviors (including the attackers's) are relatively static. However, relying entirely on the perimeter defense for securing the network is becoming challenging as the wireless space becomes more complex in applications (e.g., IoT) and in operations (e.g., cognitive radio) and as attackers become more sophisticated to bypass the perimeter defense.

Thus, instead of depriving the attacker of its insider credentials, we adopt a real-time strategy (updating the assigned

978-1-4673-7876-5/15/$31.00 ©2015 IEEE

network resources in every control communication) to build an additional layer of resiliency after the network perimeter. Therefore, we consider the more sophisticated threat model where attackers retain insider capabilities e.g., of reading and writing the network's control messages.

In wireless MAC, an insider attacker can perform the following threats: *jamming*[1] (injecting artificial interference to flood the medium with noise), *false reservation injection* (initiating excessive MAC reservations and reserving the channel resources without using them), and *false feedback distribution* (reporting false information to skew the decisions on MAC control). False reservation denies bandwidth to legitimate users and takes relatively little attacker resources (power to transmit control messages) and consumes network resources disproportionate to attacker effort; it is thus generally more efficient threat than jamming. For our work, we consider an attacker that launches all three threats. However, the attacker is *power-limited* and wants to maximize its impact given a power constraint; such assumption is standard in wireless security because any jamming attacker without power limits can jam across all RF spectrum (from DC to light) at unlimited power, in which case none of the users can achieve any throughput.

We consider a *multi-channel environment* with power-limited users. The frequency spectrum is divided into multiple channels and each user competes for bandwidth on one channel at a time; our framework considers channels that have flexible bandwidth and varying center frequency. We also investigate both environments where a trusted entity exists and acts as an authority (centralized) and where such entity is absent (distributed); our analyses focus more on the distributed protocol that presents greater challenges.

We carefully model our work to produce fundamental results. For instance, our work offers a modular design that supports generality in physical-layer design (e.g., in modulation and coding), and the implementation-specific work of optimizing energy and in-device computation are not the focuses of our work. Our goal is to increase the achievable communication rate in wireless capacity, which applies to *all* system implementations. Therefore, we use the channel capacity as our performance metric, which allows us to simultaneously assess the impact of bandwidth and channel condition while abstracting away physical-layer system decisions (which may introduce additional vulnerabilities apart from those inherent in our MAC framework).

To provide a secure MAC, our scheme offers a cross-layer design between the physical and link layers and is comprised of four main components: *bandwidth allocation* that allocates bandwidth to the users, *randomization* that varies the channel access, *bandwidth coordination* that exchanges the access information across users and facilitates MAC adaptation for higher spectral efficiency, and *power attribution* that estimates each user's power. While we rely on FHSS for randomization against outsider threats (the more traditional threat model that is weaker than ours), we make major contributions in the rest of the components. The *bandwidth allocation* scheme allocates spectrum proportional to its demonstrated power, rather than the number of network identities demonstrated; it eliminates the false reservation threat, since a falsely reserving node will demonstrate minimal power. To subsequently mitigate the deficiencies of the randomization, we build *bandwidth coordination* that resolves collisions (on the overreserved bandwidth) and implements waterfilling (on the unreserved bandwidth) to ensure full bandwidth utilization. Afterward, the *power attribution* determines the amount of power that was sent by each node for data communication, even in the waterfilled region, while defending against false feedback distribution threat. Integrating the four components, our scheme provides a countermeasure solution against intelligent and insider adversaries (we reduce the optimal attacker strategy to that of an outsider) while retaining the benefit of legitimate user collaborations and achieving significant performance gain over the strategy of disabling MAC (which is the typical MAC-layer solution when the network is compromised).

[1]In Section II-C which describes our threat model, we discuss the insider advantage for jamming over the outsiders. Our analyses assume that the insiders will use any advantage they have to maximize their impact.

II. System Model

There are T non-idle transmitters, which form 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 and collaborative. The network is a single-hop network, in which users communicate directly without any need for relaying and each transmission is heard by all users. Thus, when two or more users operate on the same channel, they collide. The users do not favor any particular subset of spectrum, and every part experiences equal path loss in expectation. Furthermore, users operate in a repeated game with infinite-horizon; that is, the transmitters do not run out of queued packets. Also, all users are time-synchronized at the packet level, and they operate in the same phase in the protocol (e.g., control communication phase) at any give time.

We build our scheme on a pre-established key infrastructure, and each pair of nodes share a secret key. To prevent forgery of reservation messages, we timestamp and authenticate control packets either by using digital signatures or by authenticating them to an online trusted authority (the reservation messages need only be authenticated to that online authority). This authentication eliminates forged MAC control messages, thus ensuring that a user can be held responsible for the channels it has reserved. We further assume that each node knows which users are valid (for example, based on a certificate signed by an offline trusted authority), which prevents the Sybil attack in which one entity fakes multiple identities.

A. MAC Framework

In order to handle bursty traffic patterns characteristic of data transmissions, medium access control (MAC) protocols are dynamic, rapidly adapting resource allocations based on user demand. One common MAC approach is to have each

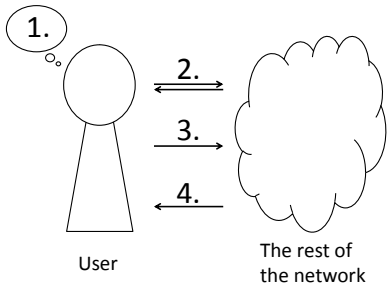

Fig. 1: Our MAC framework: 1. MAC control decision; 2. Control communication; 3. Data communication; 4. Feedback from receiver and network

node explicitly announce its channel usage intentions before transmitting data. Figure 1 illustrates a general control-communication-based MAC framework. To send a packet, the transmitter (1) makes a MAC-layer decision based on its observations and the history from previous transmission rounds, (2) reserves channels for data transmission and shares its channel usage intention with other users in a control packet, (3) transmits the data packet using the reserved channels, and (4) receives feedback from the receiver and the network. Such design is supported in wireless MACs, such as WiFi.

B. Performance Metric

Our analysis holds when using *any* metric as long as it exhibits the following three properties: it is decreasing and convex with jamming power, monotonically increasing with transmitter's signal power, and linear in available bandwidth. As a representation, we use the Shannon channel capacity limit[2] to construct our performance metric. Channel capacity is strongly correlated with both bandwidth and signal-to-interference-and-noise ratio (SINR), while abstracting away physical-layer decisions such as modulation and coding.

Whenever user i transmits to its destination user j, it does so on a frequency channel, whose location is known to user i and user j and whose bandwidth is W_i. Under a flat fading Gaussian channel with Gaussian signals and interference being treated as noise, the channel capacity of the link $i \to j$ is:

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

In Equation 1, N_0 is the noise power spectral density, $\mathcal{M}$ are the indices of jammers, $\mathcal{M}^c$ are the indices of legitimate users, $P_{x,y}$ is the *effective* or *received* signal power of the link $x \to y$, and the fraction inside of the parenthesis is the SINR.

We bound the power of all users including attackers such that user x has a bound of P_x: $E[P_{x,y}] \leq P_x < \infty, \ \forall x \in \mathcal{T}$, where $P_{x,y}$ is random due to the channel $x \to y$. As $\mathcal{R}_i$ monotonically increases with P_i, each transmitter emits at full power to maximize the signal power P_i at the receiver. Users with better channel gains can be modeled with larger power constraints. Jensen's inequality yields the capacity of:

$$E[\mathcal{R}_i] \leq W_i \log_2 \left[1 + \frac{P_i}{N_0 W_i + \sum\limits_{\ell \neq i, \ell \in \mathcal{M}^c} I_\ell P_\ell + \sum\limits_{k \in \mathcal{M}} J_k P_k} \right] \quad (2)$$

where I_ℓ is the amount of benign user ℓ's power that interferes with the transmitter's signal normalized with respect to the power constraint P_ℓ, and 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 on the channel). We use Equation 2 as our performance metric for the link $i \to j$. For our goal of maximizing the performance of the overall network, we introduce a network utility function U, which is the aggregate rate of the legitimate users:

$$U = \sum_{i \in \mathcal{T}} E[\mathcal{R}_i] = \sum_{i \in \mathcal{M}^c} E[\mathcal{R}_i] \quad (3)$$

C. Threat Model

An insider attacker launches a denial-of-service (DoS) attack on the network to reduce the network performance. We consider the worst-case attacker who minimizes the utility:

$$\text{minimize } U \text{ subject to } J_k \leq 1, \ \forall k \in \mathcal{M} \quad (4)$$

It is in the attackers' best interest to fully utilize the power budget. We also consider a strong threat where attackers within the compromised network collude and share all information through a secure, covert channel with unlimited bandwidth.

We focus on threats that exploit vulnerabilities that are insufficiently addressed by prior work in wireless MAC security. In specific, we are concerned with the following attacks: *false reservation injection*, which wastes network bandwidth, and *jamming* on the remaining bandwidth. If successful, false reservation is the more power-efficient attack of the two, since it allows an attacker to reduce bandwidth available to legitimate nodes at nearly no power cost. Each attacker can send a short reservation request message and reserve a channel for an extended period of time, supposedly for data transmission, preventing legitimate users from using the bandwidth resource. This requires only small amount of power to deliver control packets, and attackers can use the majority of their power to jam and disrupt the communication of legitimate users. In this case, attackers are successful in both wasting resource by falsely reserving portions of spectrum and degrading the channel conditions of the rest of the spectrum by jamming. To realize this, attackers are capable of accessing non-contiguous frequency band, e.g., SWIFT [5].

Attackers can also target the feedback communication and perform *false feedback distribution* to manipulate MAC parameters and influence the user's decisions on MAC control. Attackers can do so in two ways: they can attack the aggregation of the power estimations (for bandwidth allocation in distributed settings) or affect the users' power sensing by over-claiming transmissions (especially for those band that are occupied by more than one user). We discuss each of these threats in greater details and present our corresponding countermeasures in Section III-A and Section III-C, respectively.

As a separate class of attackers, we also consider more traditional threats of *outsider jammers*, whose strategies the

[2]The channel capacity given by Shannon-Hartley Theorem provides an asymptotic upper bound for the communication rate of an independent AWGN channel. This bound is commonly used for evaluating performance and is generally considered tight (information theorists continue to pursue even tighter bounds in more complex and realistic channel models).

insider attackers can choose to emulate (in fact, our contribution lies in countering the insider threats and nullifying their insider advantages, reducing the insider's optimal strategy to that of an outsider). By definition, outsiders do not have the insider advantages; any outsider that colludes with an insider is an extension of the insider adversarial network. Since they do not know who reserved which channel, outsider attackers spread their jamming power across all channels, having less impact on the transmission than the insider attackers who focus their jamming on the channels that are being used by legitimate users. Although the insider jammers have the jamming-relevant information (about which channels are being used by the legitimate users) and outsiders do not, both minimize the network capacity by choosing wideband jamming (spreading power to affect all users) over narrowband jamming (to have higher per-user impact on subset of users) [6].

D. Bases for Our Work and Related Literature

Our work builds on prior work, and we review them in this section. Our approach diverges from the conventional slotted channelization (where the spectrum is divided into channels with fixed bandwidth and static location), the typical approach when studying security in wireless MAC. Instead, we allocate channels with varying bandwidth and center frequency to more effectively match the power of each user, increasing our system's spectral efficiency. Researchers have used flexible channelization in non-security contexts [5], [7], [8].

The resource consumption of control communication is much smaller than that of data communication (this actually helps the attacker and makes the false reservation threat more efficient, as described in Section II-C) because the overhead of a control message can be amortized over data frame, and we can choose arbitrarily large data frames. Thus, we focus on the performance of data communication when evaluating our proposed scheme. Furthermore, the low rate in control communication helps in building link reliability, as wireless researchers widely use the gain from time and processing redundancy to increase noise and interference resistance [2], [6], [9], [10]; thus, we rely on such techniques for the availability of control channel.

To build resistance against outsider jammers targeting specific users, we incorporate randomization and build on the traditional anti-jamming technique of spread spectrum, where the spreading code is known to only the sender and receiver [10]. Using frequency hopping spread spectrum (FHSS), each user transmits data using frequency hopping on randomly generated hopping patterns chosen independently for each packet. To avoid reactive jamming, where an attacker senses the channel usage and jams as soon as it detects the victim's transmission, we use fast hopping where the hopping duration is less than the attacker's reaction time. Even though incorporating FHSS randomization complicates our scheme, it is crucial to add robustness against outsider jammers; otherwise, their relatively inefficient strategy is already effective in disrupting transmission and our more advanced insider threat model is not necessary (although they can still be useful in saving the attackers' power cost).

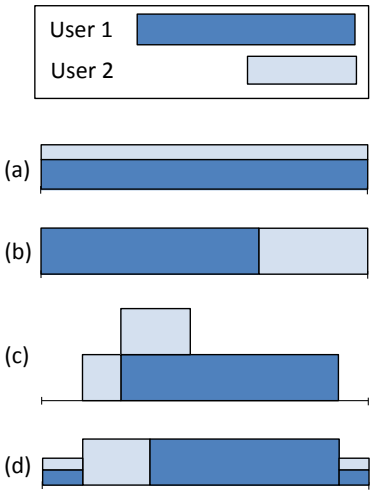

Fig. 2: We illustrate the spectrum of: (a) Nash equilibrium (wideband access); (b) Our scheme after bandwidth allocation; (c) After randomization; (d) After bandwidth coordination. Frequency in horizontal axis and power spectral density (PSD) in vertical axis.

Other researchers in wireless security have also considered and developed countermeasures for a denial-of-service (DoS) attacker capable of either jamming [2], [4], [6], sending bogus requests (e.g., to reserve channels) [11]–[13], or falsifying information at the communication feedback [14]. However, these prior work largely focus on their respective threats and remain vulnerable when facing a more comprehensive threat model that introduces an attacker capable of performing all of the aforementioned threats, which we study and counter.

III. Our Scheme

Our scheme is comprised of four components: bandwidth allocation, randomization, bandwidth coordination, and power attribution. Unlike the other components, power attribution occurs after data communication and affects the future control decisions. *Bandwidth allocation* determines the amount of bandwidth to allocate to each user based on prior power measurements. The goal of bandwidth allocation is to allocate spectrum bandwidth to each user proportional to its power capability, ensuring a constant power spectral density, known to be optimal in channel capacity [15]. Bandwidth allocation decisions are made once per round. Within a round, channel access based on allocations cannot be based on fixed center frequencies because a static allocation of channels is vulnerable to outsider narrowband jammers. We thus have a *randomization* component, which implements frequency hopping spread spectrum (FHSS) while maintaining the bandwidth allocation. (Unlike the other components, we do not further describe FHSS randomization in this section since it is well-studied in prior literature and we provide an overview in Section II-D.) Randomization in channel access results in collisions in some parts of frequency band and vacancy in others. The *bandwidth coordination* addresses these problems by sharing the bandwidth allocation and the randomization results, resolving known conflicting reservations arising from randomization-induced collisions, and allocating transmission to regions that would otherwise be underutilized. Data com-

Bandwidth allocation	
1. *Allocation*	Assign bandwidth according to the received power
2. *Data communication*	Transmit goodput data
3. *Feedback*	Report the receiver's observations in performance and power

TABLE I: Bandwidth allocation, data communication, and feedback

munication for goodput delivery follows. Finally, after a round of data transmission is complete, each node performs *power attribution* to determine the amount of power contributed for data communication by each node while leveraging commit-and-reveal to build resilience against the manipulation of power attribution. These transmission power estimates are then used for MAC control in the next round. Figure 2 illustrates the Nash equilibrium of wideband access and our scheme (the effect of each components); two users, one of which has twice as much bandwidth as the other, share the band.

A. Bandwidth Allocation

Our bandwidth allocation counters the false reservation attack with three properties: first, it provides power-fairness across users (from the receivers' perspective); second, it achieves the optimal performance in terms of spectral efficiency; and third, it prevents the attacker from simultaneously making effective reservations and using all of its power for jamming. In fact, as we will see in Section V-A, the optimal power-limited strategy is to forgo false reservation and exclusively focus on jamming. Our protocol thus performs substantially better than previous protocols that do not counter false reservations since, in those protocols, an optimal attacker can simultaneously perform false reservation and jamming, as described in Section II-C. Table I presents our scheme when only considering false reservation threats.

Our scheme assigns bandwidth to a user proportional to the received power from the user; only the current bandwidth allocation (which is necessary to perform power attribution on each users) and the update (proportional to the users' power attribution results) affect the upcoming allocation. To determine and disseminate the bandwidth allocation, the scheme has two stages: first, each node performs *individual bandwidth allocation* and then, for environments that lack an online central authority, the nodes perform *distributed bandwidth allocation* to decide and agree on the bandwidth allocation while minimizing the effect of colluding adversaries. In contrast, in the presence of a trusted authority, the authority broadcasts its own individual bandwidth allocation results to assign bandwidths, and we can bypass the distributed bandwidth allocation.

The *individual bandwidth allocation* assigns bandwidth proportional to the observed received power; all users have unique observations because they are in different spatial locations (which affects the wireless propagation attenuation) and wireless channels naturally fluctuate (e.g., fading).

The *distributed bandwidth allocation* aggregates the individual allocation decisions in a distributed manner, so that users agree on the bandwidth allocation. Each user disseminates its bandwidth allocations using the Byzantine General's algorithm with signed messages [16]. All users then compute the median bandwidth allocation for each node (median is known to be an attack-resistent aggregation mechanism [17]), and use these

Our protocol	
1. *Allocation*	Assign bandwidth according to the received power
2. *Commit*	Commit waveform and nonce for attribution
3. *Coordination*	Disclose initial channel reservations and randomization results Adjust bandwidth and waterfill
4. *Data communication*	Transmit goodput data
5. *Reveal*	Reveal waveform and nonce for attribution
6. *Power attribution*	Observe the spectrum and determine users' power levels
7. *Feedback*	Report the receiver's observations in performance and power

TABLE II: Overview of our protocol

values as the network-wide consensus. Because each node computes the median over the same set of data, each node arrives at the same bandwidth allocations. Thus, the bandwidth allocation scheme is resilient to Byzantine failure and requires only one round of message delivery.

However, because the attackers compromise a fraction of the network and have legitimate rights to vote, the distributed bandwidth allocation is vulnerable to an attack where attackers attempt to distort the consensus to their advantage. In this false feedback distribution attack, attackers report false bandwidth allocations to distort the outcome of the distributed allocation. Because the distributed bandwidth allocation uses the median, it is somewhat resilient to such attacks [17]. Nevertheless, since attackers know each other, they can still distort the median by reporting favorable values for fellow attackers and discredit legitimate nodes by claiming low power observations. As a result, the consensus median value will be shifted towards the value that the attackers report. Furthermore, if the number of attackers exceeds half the network population, the attackers gain total control over the distributed scheme.

We do not attempt to detect and isolate false reporting attackers because the variable channel conditions caused by wireless fading adds randomness to each user's received power observations and makes detection difficult. A threshold-based 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 detection. Because such detection approaches may be ineffective, yet add complexity and a new attack vector (e.g., false positive creation), we do not use them in our scheme.

B. Bandwidth Coordination

We focus on the distributed environment. (In a centralized environment, the trusted authority can facilitate channel orthogonality in randomization, e.g., by taking the individual bandwidth allocation from Section III-A and permuting them across the users.) Since the randomization process (FHSS) selects hopping patterns individually without regarding others, it causes *collisions* (channels in which multiple users make a reservation) and *under-utilization* (channels in which no users make a reservation). To avoid spectral inefficiency caused by collision and underutilization, legitimate nodes perform bandwidth coordination. In bandwidth coordination, the network users exchange the center frequencies of their reserved channels and adjust the channel access to minimize mutual interference (for impending collision, the users divide the collided bandwidth so that each node obtains an amount of bandwidth proportional to their respective bandwidth allocations);

furthermore, the users utilize the under-utilized bandwidth by waterfilling such bandwidth with their transmissions (the user divides its power between its solely-operating reserved channel and the under-utilized channel to maximize its performance based on its estimate of interference power on both the reserved channel and the under-utilized channel). The delivery of the bandwidth coordination results involves a single one-way broadcast communication in message. Table II presents an overview; the coordination (Section III-B) and the power attribution (Section III-C) are shaded to contrast with Table I; these are the additional complexities required to incorporate randomization to thwart outsider jamming and build resilience against false feedback distribution.

C. Power Attribution & Commit-and-Reveal

As the bandwidth allocation is based on the power observed from each transmitter, we provide a physical-layer supplement to our MAC protocol that attributes power to users given a received signal. Separating out the power from each user requires two parts: first, we need to know where each user is transmitting at any time, and second, for bandwidth regions where multiple users transmit at the same frequency and the same time, we need to be able to determine the power of each user. After the coordination in Section III-B, for regions where one user has sole access, we can trivially determine that user's power level by filtering and observing the amount of power in the user's reserved channel. However, for bands where multiple users transmit, such passband-based attribution schemes are ineffective. Thus, users *commit* to the waveform signature that they will transmit prior to data transmission.

In every round, our protocol involves five steps: two before the data is sent, and three afterwards. First, the user chooses a randomization pattern for hopping and commits to the data, a random nonce, and the waveform signature. Second, the user sends a control message to each node and includes the randomization pattern in that message for coordination. Then the user sends the data using its randomization pattern. Third, each user reveals their waveform signature and the nonce from initial commitment. Fourth, each user combines the random numbers (e.g., using XOR) to determine a network-wide random number, and uses that random number to determine a short portion of its data transmission to reveal (e.g., using a PRF keyed with the random number); this interval could be identical system-wide, or it could be determined on a per-transmitter basis. Finally, each user A reconstructs the perspective of each other user B during B's revealed time t_B, and determines the amount of power transmitted by B. To do so, A considers the reservations that B has sent and received, determines the output of B's coordination, determines the data that B will send, and obtains the waveform that B sent at time t_B. User A then takes the cross-correlation between the signal received by A at time t_B and the waveform that B sent at time t_B to determine B's contribution on the A's received signal. This correlator-based approach for attribution is widely used in communication such as in signal detection (e.g., preamble for synchronization), matched filter, and direct sequence spread spectrum (DSSS).

IV. Theoretical Analysis

In this section, we present our theoretical analyses results and establish the bases for the testbed evaluation in Section V.

A. Two-Party Game for Allocation

Our protocol reduces the problem of false reservations to a two-party game between the *legitimate user network* and the *attacker network*, because we assume cooperative behavior among benign users and collusion among attackers, and because the bandwidth allocation depends only on the received power. Specifically, the users' behavior and the attacker's optimal strategy depend on the relative power capabilities of the legitimate and attacker networks, rather than on the number of users. In our theoretical analysis, we 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 we control the power capabilities of the two groups by varying the number of users (T) and attackers (M). Because all users have equal power, they have the same expected performance $\mathcal{R}$, i.e., $\mathcal{R} = E[\mathcal{R}_i],\ \forall i \in \mathcal{T}$. We introduce 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. Since the attacker network uses $(\alpha \cdot \bar{P} \cdot M)$ power for false reservation, Equation 2 yields:

$$\mathcal{R} = \frac{W}{T - M + M\alpha} \log_2 \left[1 + \frac{\text{SNR}}{\frac{T}{T-M+M\alpha} + \frac{M(1-\alpha)}{(T-M)}\text{SNR}} \right] \quad (5)$$

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

B. Attacker's Lead on Distributed Protocol

Our distributed allocation scheme takes the median of each user's bandwidth allocations to reach a consensus bandwidth allocation. In this section, we study the impact of wireless channel fluctuations when our protocol is under attack by false-feedback-distribution attackers. We use β to denote the attacker's *bandwidth advantage* over a legitimate user. As discussed in Section III-A, attackers can reserve more bandwidth than legitimate users with the same power ($\beta \geq 1$) because attackers collude while legitimate users report truthfully. Due to channel diversity and fading, 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 a colluding attacker, attackers report large power observations, shifting the observed median upward to the $100 \cdot \frac{0.5 \cdot T}{T-M} > 50$th percentile of legitimate observations. Also, for a legitimate user, the attackers report low power to shift the median downward to the $100 \cdot \frac{0.5 \cdot T - M}{T-M} < 50$th percentile of observations. In contrast, legitimate users report their true observations that include random wireless channel fluctuation. Assuming an iid channel for all users with each channel characterized by a cumulative distribution function CDF, the attacker's advantage is: $\beta = \frac{\text{CDF}^{-1}(\frac{0.5 \cdot T}{T-M})}{\text{CDF}^{-1}(\frac{0.5 \cdot T - M}{T-M})}$.

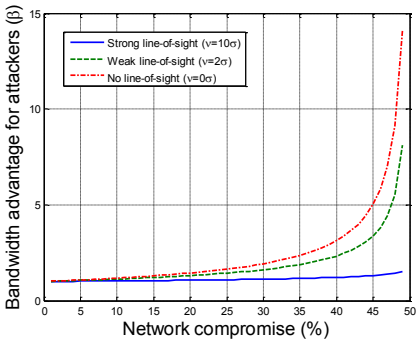

Fig. 3: Ratio of attacker's and legitimate user's bandwidth under false feedback distribution attack on bandwidth allocation.

In Figure 3, we plot the attacker bandwidth advantage β under varying channel fading where channel characteristics vary in Rician fading with ν and σ, where ν^2 is the power of the line-of-sight path[3] and $2\sigma^2$ is the power of the other scattered paths. In particular, we study three choices of 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, suitable for a highly dynamic environments, e.g., a mobile application in the cities. Increasing the number of attackers results in greater error in the computed median and greater attacker bandwidth advantage. Because channel fluctuation affects the randomness within the power reports, the attacker's bandwidth advantage β depends on $\frac{\nu}{\sigma}$. The advantage increases as the line-of-sight path becomes less dominant.

Under the false feedback distribution 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 of: $\mathcal{R} = \frac{W}{T-M+M\alpha\beta} \log_2 \left[1 + \frac{\text{SNR}}{\frac{T}{T-M+M\alpha\beta} + \frac{M(1-\alpha)}{(T-M)}\text{SNR}}\right]$. Thus, compared to the performance of a centralized scheme (Equation 5), our distributed scheme gives attackers an additional factor of β more bandwidth.

C. Static α Strategy for Attackers

We study the optimal attacker strategy and show that the optimal α is static in time. Let $\mathcal{U}$ be the aggregate utility over time, i.e., $\mathcal{U} = \sum_t U_t$, where U_t is the network performance (Equation 3) at time t; let $\hat{\alpha}$ be the static-game α strategy minimizing U; and let α_t be the amount of power used to reserve channels at time t.

Theorem 1: Given $\hat{\alpha}$, $\alpha_t = \hat{\alpha}$, $\forall t$, minimizes $\mathcal{U}$.

Proof: We provide a sketch of the proof and overlook the impact of fading (and β) here. From Equation 5, 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, $\alpha_t = E[\hat{\alpha}] = \hat{\alpha}$, $\forall t$ minimizes $\mathcal{U}$.

[3]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.

V. Testbed Evaluations

We take a modular approach in analyzing each of the components and then implement the protocol as a whole; the results for the combined system is summarized in Section VI because its behavior follows the results and insights presented in this section. Our implementation uses four WARP software-defined radio platforms [18], each of which has two antenna chains. Using the MIMO (multiple input multiple output) capability of the platform, we build a network with four transmitters and four receivers. Each transmitter has equal power budget unless otherwise noted (e.g., Section V-A introduces an identity-only attacker with zero power transmission). We manually calibrate the antenna locations so that each receiver observes approximately the same power from each transmitter. In the absence of interference, the channel experiences a SNR of approximately 16 dB for any transmitter-receiver pair. For power attribution, we assume that each transmitter can learn its power level relative to other transmitters, either through full-duplex radio techniques or by, for each transmitter node, designating a single receiver node trusted by that transmitter.

Each node continuously transmits to maximize network utility U. At the physical layer, we use DQPSK modulation with a BPSK-modulated Barker sequence preamble. We use 12 MHz of network bandwidth divided into 300 subcarriers using OFDM. For example, if the bandwidth is allocated equally among n registered users (the baseline allocation strategy, as defined in Section V-A), each user will use $\frac{300}{n}$ subcarriers. The experiment results are averaged over 1000 runs. Each run is for a single round, and each round lasts for 6 hops. Each transmitter sends random bits to its receiver, and its receiver demodulates the received signal and uses the BER to estimate the SINR at the receiver, using the equation from [19], [20]: $\overline{\text{BER}} = \frac{1}{2}\left(1 - \frac{\sqrt{2}\cdot\overline{\text{SINR}}}{\sqrt{1+4\cdot\overline{\text{SINR}}+2\cdot\overline{\text{SINR}}^2}}\right)$. We then determine the capacity based on the observed SINR using Equation 1.

A. Bandwidth Allocation

We compare our protocol to a baseline protocol. In the *baseline protocol*, the frequency band is divided equally into multiple channels and each user gets one channel, regardless of whether they use the allocated channel or not. Optimal attacker strategy against this baseline protocol is to use no power in the spectrum allocated to them ($\alpha = 0$), wasting $\frac{M}{T}$ of the entire network bandwidth for free, and focus all its power on jamming ($1-\alpha = 1$). In our evaluation, we consider two legitimate transmitters, one attacker, and one *identity-only attacker* (which has zero power budget).

We study individual bandwidth allocation under the two attacker strategies of $\alpha = 0$ (using all power for jamming) and $\alpha = 1$ (using all power for effective reservation). Figure 4(a) shows the expected normalized bandwidth allocation to the four transmitters. Beginning from the baseline strategy of equal-bandwidth allocation (i.e., each of the four entities occupy 0.25 of the network bandwidth), our scheme quickly converges to the steady-state bandwidth allocation in two rounds, where the delay is caused by noise in the spectrum reserved by attackers. We plot the *legitimate network bandwidth*,

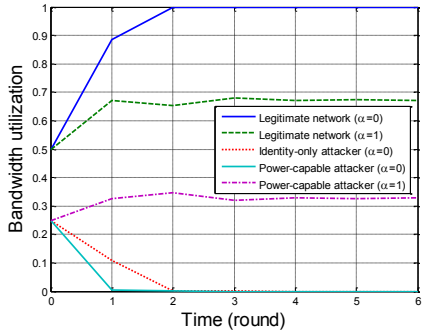

(a) Normalized BW utilization

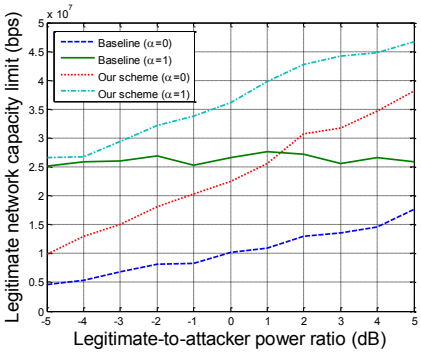

(b) Varying power budget

Fig. 4: Individual bandwidth allocation performance

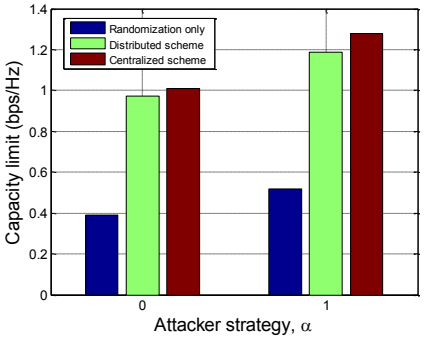

Fig. 5: The impact of bandwidth coordination

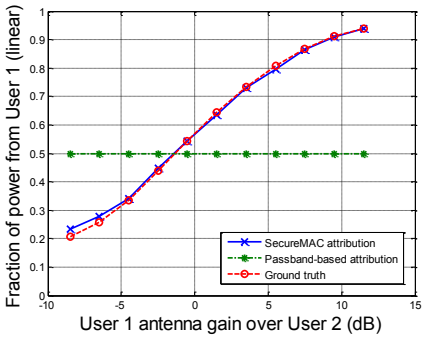

Fig. 6: Power attribution (overlapping channel case)

the fraction of bandwidth allocated to legitimate network users, which converges to 1 and to $\frac{2}{3}$, respectively, for $\alpha = 0$ and for $\alpha = 1$, validating our steady-state theoretical results. The $\alpha = 0$ attacker is quickly found to be emitting no power in reserved spectrum. Furthermore, the identity-only attacker also quickly converges to zero bandwidth as it emits no power and thus has no impact on network performance under our scheme.

We also vary the attacker's power budget relative to the legitimate user's power (the identity-only attacker retains zero power budget) in Figure 4(b). As expected, larger legitimate-to-attacker power ratios result in better performance, where the performance increase comes from reduced interference for $\alpha = 0$ (jamming) and from increased bandwidth for $\alpha = 1$ (reserving). For each MAC, $\alpha = 0$ represents a stronger attack than $\alpha = 1$. Power increments are best spent on jamming; spending power to make effective reservations shows less impact with increasing power. Thus, the attacker chooses to jam rather than to spend power to make effective channel reservations under our scheme; we also verify these results in detailed MATLAB simulations in a technical report [21] where $0 < \alpha < 1$ cases are also analyzed.

We also implement the distributed bandwidth allocation and summarize our results here. First, the distributed bandwidth allocation performs worse than the individual allocation due to the impact of the false feedback distribution attack; for instance, at steady-state, distributed allocation achieves 92.4% and 88.7% of the individual allocation performance when $\alpha = 0$ and $\alpha = 1$, respectively. Second, despite the persistent effect of the false feedback, the distributed allocation performs better than the baseline strategy of entity-fair allocation. Third, the distributed allocation shifts the attacker's optimal power-splitting strategy from $\alpha = 0$, but the impact difference is marginal from when $\alpha = 0$. We further study these phenomenons in MATLAB simulations in a technical report [21].

B. Bandwidth Coordination

To isolate the behavior of randomization and coordination from allocation, we set each user's bandwidth to 0.25 of the total network bandwidth and conduct our experiments with four equal-power users, one of which acts as an attacker; the attacker strategy of $\alpha = 1$ does not help with its objective in Equation 4 since we adopt the baseline allocation scheme and fix the bandwidth. In Figure 5, we compare the capacity limit performance of three schemes: the naïve frequency hopping ("Randomization only"), our proposed randomization and coordination ("Distributed scheme"), and the centralized scheme that offers fully orthogonal access ("Centralized scheme"). Our scheme is strong compared to, and consistently outperforms, the random frequency hopping; it outperforms randomization-only by 129% when attackers adopt $\alpha = 1$ and the performance advantage increases to over 148% when $\alpha = 0$. Furthermore, our protocol compares well with the perfectly-orthogonal centralized approach and performs within 5% of the optimal coordination.

C. Physical-Layer Power Attribution

In this section, for simplicity of the presentation and since we study a physical-layer phenomenon, we have two transmitters transmitting at the same time and sharing the medium. Our implementation samples one hop in each round. We compare

our power attribution scheme, the simple *passband-based* power attribution (which, to attribute power within a frequency band, filters each band and evenly divides the power between the users who reserved the band), and the *ground truth* (which assumes a priori knowledge about the exact transmission waveform that leaves the transmitter antenna and by using soft correlation with the signal at the receiver antenna). We study two scenarios: one in which the reserved channels do not overlap and another in which reserved channels completely overlap, only the latter of which is shown in Figure 6. Any scenario is a linear combination of these two scenarios. Compared to the ground truth, our attribution scheme and passband observation both provide good performance and follow the ground truth attribution's behavior when the channels have no overlap; both schemes perform well because the entire signal power in the passband originates from a single sender. However, when the two users choose completely overlapping channels (after randomization and coordination), the passband observation observes the same channel for each user and divides the power in half, resulting in equal power attributions for each user, regardless of the actual power. Our scheme uses the actual waveform transmitted and is much more accurate; as seen in Figure 6, across all relative powers studied (in the x-axis), the maximum error in our power attribution is 0.56 dB while the maximum error in passband attribution is 3.86 dB. Also, because passband attribution gives a constant 50% attribution to each user, the error in relative power increases as the difference in transmission power level increases.

VI. Conclusion

This paper studies the inherent vulnerabilities of MAC against attackers who have the credentials of legitimately registered users. Threats that have been largely left unresolved in such environments include false reservation injection, false feedback distribution, and jamming. Our scheme defends against such threats using a combination of four mechanisms: bandwidth allocation that allocates channels based on the usage in previously reserved spectrum, randomization to defend against reactive and outsider jamming, coordination to resolve collisions caused by randomization, and power attribution to make future MAC control decisions. Our evaluations show that, in practical scenarios, both centralized and distributed versions of our work are successful in nullifying the attackers' advantages of compromising the network while having the benign users retain the benefit of user collaboration in MAC. In particular, in our implementation environment, our work outperforms security-oblivious MAC with entity-fair channelization by 159%, FHSS (without coordination) and entity-fair channelization by 149%, and the Nash equilibrium of wideband access by 76%.

VII. Acknowledgments

This study is partially supported by NSF under Contract No. NSF CNS-0953600 and by the Human-Centered Cyber-physical Systems Programme at the Advanced Digital Sciences Center from Singapores Agency for Science, Technology and Research (A*STAR).

References

[1] R. Etkin, A. Parekh, and D. Tse, "Spectrum Sharing for Unlicensed Bands," vol. 25, no. 3, p. 517, 2007.

[2] J. Chiang and Y. Hu, "Dynamic Jamming Mitigation for Wireless Broadcast Networks," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, 2008, pp. 1211–1219.

[3] L. Li, S. Zhu, D. Torrieri, and S. Jajodia, "Self-healing wireless networks under insider jamming attacks," in *Communications and Network Security (CNS), 2014 IEEE Conference on*, Oct 2014, pp. 220–228.

[4] Z. Liu, H. Liu, W. Xu, and Y. Chen, "Exploiting jamming-caused neighbor changes for jammer localization," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 23, no. 3, pp. 547–555, 2012.

[5] 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.

[6] S.-Y. Chang, Y.-C. Hu, and N. Laurenti, "SimpleMAC: a jamming-resilient MAC-layer protocol for wireless channel coordination," in *Proceedings of the 18th annual international conference on Mobile computing and networking*, ser. Mobicom '12, 2012, pp. 77–88.

[7] 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.

[8] 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.

[9] S.-Y. Chang, Y.-C. Hu, J. Chiang, and S.-Y. Chang, "Redundancy offset narrow spectrum: countermeasure for signal-cancellation based jamming," in *Proceedings of the 11th ACM international symposium on Mobility management and wireless access*, ser. MobiWac '13. New York, NY, USA: ACM, 2013, pp. 51–58. [Online]. Available: http://doi.acm.org/10.1145/2508222.2508233

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

[11] R. Negi and A. Rajeswaran, "DoS attacks on a reservation based MAC protocol," in *IEEE ICC*, 2005.

[12] V. Gupta, S. Krishnamurthy, and M. Faloutsos, "Denial of Service Attacks at the MAC Layer in Wireless Ad Hoc Networks," in *MILCOM*, vol. 2, 2002, pp. 1118–1123.

[13] N. Golde, K. Redon, and J.-P. Seifert, "Let me answer that for you: Exploiting broadcast information in cellular networks," in *Proceedings of the 22Nd USENIX Conference on Security*, ser. SEC'13. Berkeley, CA, USA: USENIX Association, 2013, pp. 33–48. [Online]. Available: http://dl.acm.org/citation.cfm?id=2534766.2534770

[14] Y.-C. Tung, S. Han, D. Chen, and K. G. Shin, "Vulnerability and protection of channel state information in multiuser mimo networks," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '14. New York, NY, USA: ACM, 2014, pp. 775–786. [Online]. Available: http://doi.acm.org/10.1145/2660267.2660272

[15] D. N. C. Tse, "Optimal power allocation over parallel gaussian broadcast channels," 1997, p. 27.

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

[17] B. Przydatek, D. Song, and A. Perrig, "SIA: secure information aggregation in sensor networks," in *Proceedings of the 1st international conference on Embedded networked sensor systems*, ser. SenSys '03. New York, NY, USA: ACM, 2003, pp. 255–265. [Online]. Available: http://doi.acm.org/10.1145/958491.958521

[18] 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.

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

[20] 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.

[21] S.-Y. Chang and Y.-C. Hu, "Secure Channel Reservation for Wireless Networks," in *Technical Report*. [Online]. Available: https://www.ideals.illinois.edu/bitstream/handle/2142/17098/2010-2509_Chang.pdf?sequence=2