Jamming-Resistant Rate Control in Wi-Fi Networks

Cankut Orakcal
Dept. of Electrical and Computer Eng.
Boston University
Boston, MA, 02215, USA
Email: orakcal@bu.edu

David Starobinski
Dept. of Electrical and Computer Eng.
Boston University
Boston, MA, 02215, USA
Email: staro@bu.edu

***Abstract*—Recent experimental studies reveal that several well-known and widely deployed rate adaptation algorithms (RAAs) in 802.11 WLANs are vulnerable to selective jamming attacks. However, previous work resorts to complex jamming strategies that are hard to implement and does not provide applicable solutions to this problem. In this work, we analyze the vulnerabilities of existing RAAs to simple jamming attacks and propose judicious use of randomization to address this problem. We introduce a theoretical framework based on a bursty periodic jamming model to analyze the vulnerabilities of popular RAAs, such as ARF and SampleRate. Our parameterized analysis shows that a jamming rate of 10% or below is sufficient to bring the throughput of these algorithms below the base rate of 1 Mb/s. Thereafter, we propose Randomized ARF (RARF), which has higher resistance to jamming attacks. We derive a closed-form lower bound on the minimum jamming rate required to keep the RARF throughput below the base rate. Finally, we conduct ns-3 simulations implementing various RAAs and jamming strategies for an IEEE 802.11g WLAN. Our simulations validate jamming strategies under different channel models and show that the minimum jamming rate required against RARF is about 33%.**

I. Introduction

Wireless local area networks (WLANs), based on the IEEE 802.11 (Wi-Fi) family of standards, form a large portion of the current Internet infrastructure. According to [9], over 100,000 public Wi-Fi hotspots are deployed in the U.S. In addition, several cellular network providers are deploying Wi-Fi hotspots to offload congestion in crowded areas [1].

IEEE 802.11 supports data transmission at multiple rates. Several rate adaptation algorithms (RAAs) have been proposed to this effect. The purpose of an RAA is to adapt the transmission rate and the modulation scheme in order to maximize performance (e.g., throughput) based on current wireless channel conditions. The networking community has put great effort on devising efficient RAAs [4], [10], [12], [13], and several of these algorithms have been commercialized.

Security represents a major challenge to wireless services due to the broadcast nature of the wireless channel. A major part of attacks against Wi-Fi consists of jamming. Commercial off-the-shelf jamming equipment is at reach of anyone's hands [8]. Recent studies, such as [7], [20], [21], demonstrate that various jamming attacks described in the literature can be implemented efficiently. Typically, a jammer aims to cause maximum damage using a minimum number of transmissions to avoid getting detected and to save energy.

The main motivation behind jamming is to cause Denial of Service (DoS) [18]. Jamming can be used for personal (deny communication between some parties), economic (competing companies) or governmental purposes (cyber warfare) [19]. Jammers can also perform Reduction of Quality (RoQ) attacks using intelligent jamming patterns [6], [18]. Typically, RoQ attacks exploit vulnerabilities above the physical layer. The result is degradation of throughput and prolonged delays, which are not tolerable in time-critical applications.

RAAs are not designed to operate against malicious entities. Many of them fail to distinguish between packet losses due to fluctuations in channel conditions and those due to interference. Thus, recent experimental studies [14], [17] reveal that several well-known and widely deployed rate adaptation algorithms used in 802.11 WLANs are vulnerable to jamming attacks (cf. Section II for a detailed discussion). However, these works neither offer simple methods to construct effective jammers nor present solutions with theoretical guarantees.

Our contributions in this context are as follows: First, we introduce a theoretical framework to analyze the vulnerabilities of several existing RAAs to jamming attacks. We propose a jamming model, called *bursty periodic jamming*, under which an adversary periodically jams the channel with a burst of packets. For this model, we constructively determine strategies and corresponding jamming rates to keep the throughput of RAAs below the base rate (i.e., the lowest bit-rate). For default parameters, we show that low jamming rates of about 9% for the early Automatic Rate Fallback (ARF) algorithm and 4% for the newer SampleRate algorithm, are sufficient to achieve this goal. Our analysis provides expressions for general parameter settings.

Next, we propose Randomized ARF (RARF) as a means to improve the resistance of RAAs to jamming attacks. RARF performs *comparably* to ARF under *no jamming*, but *much better under targeted jamming attacks* due to its randomized nature. We derive a closed-form lower bound on the minimum jamming rate required to keep RARF throughput below the base rate. For default parameters, this value is about 20% (at least twice the jamming rate required for other RAAs).

Finally, we conduct ns-3 simulations implementing various RAAs and jamming strategies for an IEEE 802.11g WLAN. Our simulations validate the jamming strategies under different channel models. Furthermore, we see that the minimum jamming rate required against RARF is about 33% in practice.

This work was supported in part by NSF under grant CNS-1012910.

The rest of this paper is organized as follows. In Section II, we review related work. Next, we introduce our theoretical model in Section III. We analyze the impact of bursty periodic jamming on several RAAs in Section IV. Then, we propose and analyze a randomized jamming-resistant approach in Section V. We present the results of our ns-3 simulations in Section VI and conclude the paper in Section VII. Due to space limitations, some of the proofs of lemmas, propositions and theorems, and pseudo-codes of RAAs are deferred to [16].

II. Related Work

Several RAAs have been proposed for 802.11 WLANs. Notable examples include ARF [10], AARF [11], Onoe [13], SampleRate [4], and Minstrel [12], all of which have been used in commercial equipment. RAAs adaptively pick the best possible rate, based on changing channel conditions. However, many proposed RAAs fail to distinguish packet losses due to channel conditions from those due to interference.

Broustis et al. [5] investigate jamming attacks that exploit MAC protocols. The authors demonstrate that attacking a single node degrades the entire WLAN performance due to performance anomaly. They propose *FIJI*, a defense mechanism to identify the node under attack and prevent the other nodes to be affected. In our work, we analyze RoQ attacks that exploit RAAs directly, and propose a defense mechanism that increases the performance of the node under attack.

To our knowledge, Pelechrinis et al. [17] are the first to study the effect of jamming on RAA performance. They employ a random jammer that alternates between uniformly distributed jamming and idle periods. For several popular RAAs, system performance reduces drastically under select jamming attacks, whereas fixed rate transmission provides higher throughput. Thus, the authors propose an anti-jamming scheme *ARES* that uses rate adaptation when the jammer is idle, and uses fixed rate transmission otherwise. *ARES* adjusts the carrier sense threshold, so that packets can be received even when a jammer transmits. This scheme assumes that there exists a perfect method to detect a jammer, which is a non-trivial problem. Moreover, adjusting the carrier sense threshold usually works only if the transmission power of the jammer is lower than the transmission power of non-malicious nodes.

The recent work of Noubir et al. [14] investigates the vulnerability of several RAAs against smart (selective) jamming attacks, and shows the existence of effective attacks to degrade system performance. A jammer sniffs the header of each packet to retrieve the transmission rate. Based on this information, the jammer instantly decides on whether to jam the packet or not. In contrast to our paper, [14] does not analyze the performance of each RAA under jamming. There is no explicit construction for the jammer model and they do not provide a feasible and tested solution. Furthermore, the need of interpreting packet information for every transmission has high computational complexity. As shown in our work, such a complex jammer is unnecessary to significantly degrade the performance. Since many RAAs are deterministic, an adversary knows how an RAA behaves without retrieving the transmission rate. Lastly, randomization is mentioned in [14] as a possible solution, but no concrete algorithm is presented.

III. Models and Notation

A. Channel Model

We assume that n possible transmission rates exist, denoted as $R_1, R_2, \ldots, R_n$, where $R_1 < R_2 < \ldots < R_n$. For instance, IEEE 802.11g standard allows transmission at $n = 12$ different bit-rates. Let α_i denote the long-run proportion of packets transmitted at the bit-rate R_i, and f_i denote the long-run proportion of packet losses at rate R_i in the presence of a jammer. The *steady-state throughput* is then defined as:

$$Thr = \sum_{i=1}^{n} \alpha_i \left(1 - f_i\right) R_i. \tag{1}$$

To keep the analysis tractable, we ignore all control packets, back-off retransmissions, and inter-frame spacings.

B. Jamming Model

In this paper, we consider an easy to implement jamming model called *bursty periodic*. Implementation of a similar jamming model is demonstrated to be feasible by Bayraktaroglu et al. [3]. Under such a model, a jammer is capable of jamming a consecutive packets out of every T packets. We refer to T as the *jamming period* and to a as the *jamming burst size*. The value of a can be any positive integer, and T must always be greater than a. This model is a special case of the (T, λ) model introduced in the work of Awerbuch et al. [2].

The jammer has no knowledge of the real-time transmission rate or a history of rates used. The only information available is the RAA implemented on the target system. The jammer is reactive, i.e., it employs carrier sensing in order to jam the channel only if there is an ongoing packet transmission.

The *Rate of Jamming*, abbreviated RoJ, is the main metric of interest in this paper. It is defined as the ratio of number of jammed packets to the total number of transmitted packets. Given a and T, the jamming rate is $RoJ = a/T$. For each studied RAA in this paper, our goal is to find the minimum value of RoJ (or a bound on it) to keep the throughput of the RAA below the base rate R_1. A low RoJ implies that an RAA is highly vulnerable to jamming attacks, while a high RoJ implies that the RAA is resilient.

IV. Analysis

In this section, we analyze the effects of jamming on the throughput of ARF and SampleRate. Note that several other RAA algorithms, such as AARF and Onoe, are also vulnerable to periodic jamming. Their analysis can be found in our technical report [16]. We provide upper bounds on the jamming rates required to keep the throughput of each algorithm below the base rate. Our analysis applies to perfect or lossy channels. In the following, RoJ_{RAA} denotes the jamming rate required to keep the throughput of RAA scheme below R_1, and Thr_{RAA} denotes the corresponding throughput.

A. Automatic Rate Fallback (ARF)

ARF is the first documented RAA [10]. It keeps track of the number of consecutive packet transmissions and failures at the current rate. If s consecutive packets are correctly acknowledged, a probe packet is sent at the next higher rate (if available). If the probe packet succeeds, then the next higher rate is used for subsequent frame transmissions. Otherwise, ARF returns back to the previous rate. [10] refers to a probe packet failure as an *immediate fallback*. On the other hand, if f consecutive packet transmissions are not correctly acknowledged, the next lower bit-rate (if available) is used for subsequent frames. ARF is initiated from the lowest rate R_1. The default parameter values are $s = 10$ and $f = 2$. A possible strategy to keep the throughput below R_1 is to jam all probe packets sent at R_2. The jamming rate and the resulting throughput value are calculated in the proof of Proposition 1.

Proposition 1: The throughput of ARF can be kept below R_1 by using a bursty periodic jammer with jamming rate $RoJ_{ARF} = (s{+}1)^{-1}$. For default parameter values (i.e., $s = 10$ and $f = 2$), $RoJ_{ARF} \approx 9.1\%$.

Proof: We begin our proof by assuming perfect channel conditions. The jamming strategy under consideration allows s consecutive successful transmissions at R_1, but jams each probe packet sent at R_2. For this purpose, a jammer can set a jamming period of $T = s + 1$ packets and burst size $a = 1$. Since each probe packet sent at R_2 is jammed, ARF is never able to switch to R_2 for further transmissions, resulting in:

- $\alpha_1 = s(s+1)^{-1}$, $\alpha_2 = RoJ_{\text{ARF}} = (s+1)^{-1}$,
- $f_1 = 0$, $f_2 = 1$.

Using Eq. (1), we can calculate $Thr_{\text{ARF}} = s(s+1)^{-1}R_1$. This jamming strategy works also for *lossy channel* conditions. If any packet transmission at R_1 is lost within a jamming period, the system is not able to get s consecutive successful transmissions and does not even attempt to transmit at R_2. ■

B. SampleRate

SampleRate [4] estimates the expected per-packet transmission time at each bit-rate, and selects the bit-rate that is predicted to achieve the highest throughput. To get the estimates, it periodically sends packets at transmission rates other than the current one and records the transmission times. SampleRate switches to another bit-rate if the estimated average per-packet transmission time at that rate is smaller than that at the current bit-rate. Furthermore, bit-rates expected to perform worse, i.e. with minimum transmission times higher than the average transmission time of the current bit-rate, are not sampled. Results of transmissions that occurred over $updWin$ (default 10) seconds ago are discarded.

If no packets have been acknowledged at the current bit-rate, SampleRate picks the highest bit-rate that has not had four consecutive failures. Furthermore, a rate that had four consecutive failures is blacklisted, i.e. SampleRate does not pick it for $updWin$ seconds. Thus, preventing any transmission at rates higher than R_1 will cause all rates but R_1 to be blacklisted, keeping the system at R_1 for $updWin$ seconds.

A possible jamming strategy to keep the throughput of SampleRate below R_1 is to jam every packet transmitted at a bit-rate higher than R_1. The jamming rate and the resulting throughput value of this strategy are calculated in the proof of Proposition 2. We assume the packet length is set to $pktL$.

Proposition 2: The throughput of SampleRate can be kept below R_1 by a bursty periodic jammer with jamming rate:

$$RoJ_{SampleRate} = \frac{4(n-1)pktL}{4(n-1)pktL + updWin \times R_1} . \qquad (2)$$

For default parameter values (i.e., $n = 12$, $pktL = 10000$ bits, $updWin = 10$ sec and $R_1 = 1$ Mb/s), $RoJ_{SampleRate} \approx 4.2\%$.

Proof: We begin our proof by assuming perfect channel conditions. Since four packet failures are necessary to blacklist any rate, one needs to jam $a = 4(n-1)$ packets consecutively to blacklist all bit-rates higher than R_1. This causes the system to get stuck at R_1 for $updWin$ seconds. Since the jamming period consists of transmissions performed at rate R_1 for $updWin$ seconds and $4(n-1)$ packet transmissions at higher rates, $T = R_1\ updWin(pktL)^{-1} + 4(n-1)$, leading to the RoJ expression given by Eq. (2). The resulting values are:

- $\alpha_1 = 1 - RoJ_{\text{SampleRate}}, \sum_{i=2}^{n} \alpha_i = RoJ_{\text{SampleRate}},$
- $f_1 = 0$, $f_2 = f_3 = \ldots = f_n = 1$.

Using Eq. (1), we can calculate the throughput:

$$Thr_{\text{SampleRate}} = \left[\frac{R_1 \times updWin}{4(n-1)pktL + R_1 \times updWin}\right] R_1 \ .$$

This jamming strategy works also for *lossy channel* conditions. Once SampleRate is forced down to R_1, all bit-rates higher than R_1 are blacklisted. Packet losses at R_1 within $updWin$ seconds do not affect the behavior of SampleRate, since only R_1 is available. As soon as higher bit-rates are available, SampleRate switches to those rates regardless of any packet loss that might have occurred. ■

V. Jamming-Resistant Rate Adaptation

In this section, we introduce *Randomized ARF* (RARF) and analyze its performance under a bursty periodic jammer. We derive a closed-form lower bound on the minimum jamming rate required to keep RARF throughput below R_1. This bound is much higher than the jamming rate sufficient to keep the throughput of ARF and SampleRate below R_1.

A. Randomized ARF (RARF)

In Section IV-A, we have shown that ARF throughput can be kept below R_1 if $RoJ = (s+1)^{-1}$ due to its deterministic nature. Since the adversary knows when ARF jumps to R_2, employing a jamming strategy that keeps ARF at R_1 is simple. However, randomizing the location of these jumps prevents the adversary to decide on which packets to jam. Thus, instead of switching to the next higher rate after s successful transmissions, RARF switches with probability s^{-1} after each successful transmission. The failure mechanism is the same, i.e. RARF switches to the next lower rate after f consecutive failures. RARF does not make use of probe packets, however.

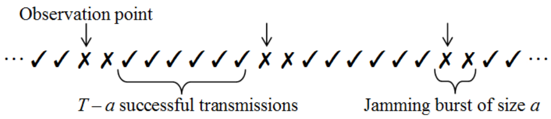

Fig. 1: Diagram of the observation process

B. Throughput of RARF

In this section we derive the expected throughput of RARF under jamming. In Section V-C, we use it to derive a lower bound on the minimum jamming rate to keep RARF throughput below R_1. We assume that RARF operates only over two rates (R_1 and R_2). We observe the system just before the jamming burst in every jamming period as in Fig. 1. The steady state probability of finding the system at R_i is denoted π_i.

We consider two cases for the adversary: $a < f$ and $a \geq f$. For $a < f$, the jammer cannot force the system down to R_1 once it switches to R_2. Thus, the system always transmits at R_2 in the steady state, leading to the throughput $(1-a/T)R_2$.

Next, we consider $a \geq f$. This time, RARF is guaranteed to be at R_1 after each jamming burst. If the system is observed at R_1, all $T-a$ packets before that must have been transmitted at R_1 as in Fig. 2(a). On the other hand, if the current rate is R_2, some of the last $T-a$ packets were transmitted at R_1 and the rest were transmitted at R_2 as in Fig. 2(b). X_1 denotes the number of packets transmitted at R_1 in a jamming period, given that the rate right before the jamming burst is R_2.

Lemma 1: $\mathbb{E}[X_1] = s - \dfrac{(T-a)\left(1-s^{-1}\right)^{T-a}}{1-\left(1-s^{-1}\right)^{T-a}}$.

Next, we provide expressions for the probability of finding the system at rate R_1 or rate R_2, at the time of an observation.

Lemma 2: Under a bursty periodic jammer with parameters (a, T) and $a \geq f$, the steady state probabilities for two-rate RARF system are $\pi_1 = \left(1-s^{-1}\right)^{T-a}$ and $\pi_2 = 1 - \pi_1$.

Expected throughput can be calculated as:

$$\mathbb{E}[Thr_{\text{RARF}}] = \pi_1\left(1-a/T\right)R_1 + \pi_2\left(\frac{\mathbb{E}[X_1]}{T}R_1 + \frac{T-a-\mathbb{E}[X_1]}{T}R_2\right). \quad (3)$$

Applying Lemmas 1 and 2, we obtain the expression for the system throughput as given by Eq. (4).

Theorem 1: Under a bursty periodic jammer with parameters (a, T), the throughput of two-rate RARF is

For $a < f$: $\quad \mathbb{E}[Thr_{RARF}] = (1-a/T)R_2$.

For $a \geq f$:

$$\mathbb{E}[Thr_{RARF}] = (1-a/T)R_2 - \left[1-\left(1-s^{-1}\right)^{T-a}\right]\frac{s\left(R_2-R_1\right)}{T}. \quad (4)$$

C. Jamming Strategies Against RARF

In this section, we derive a lower bound on the minimum jamming rate needed for a bursty periodic jammer to keep RARF throughput below R_1. First, we analyze $a < f$.

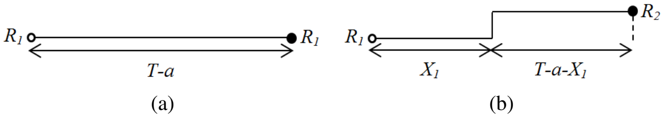

Fig. 2: Possible scenarios in a jamming period when $a \geq f$.

Proposition 3: For $a < f$, the throughput of RARF can be kept below R_1 by using a bursty periodic jammer with jamming rate $RoJ_{RARF} = 1 - R_1/R_2$. For default parameter values (i.e., $R_1 = 1$ Mb/s and $R_2 = 2$ Mb/s), $RoJ_{RARF} = 50\%$.

Unless R_2 is close to R_1, choosing $a < f$ requires a high jamming rate. Thus, the case $a \geq f$ usually results in a more efficient strategy. For $a \geq f$, we use a lower bound on RARF throughput that in turn will be used to derive a lower bound on the minimum jamming rate. This bound is based on:

Lemma 3: $\mathbb{E}[X_1] \leq (T-a+1)/2$.

Lemma 4: $\pi_1 \leq s[e\,(T-a)]^{-1}$, for $s \geq 2$ and $T-a \geq 1$, where e is the exponent number.

Applying Lemmas 3 and 4 to Eq. (3) yields the following lower bound on the throughput of RARF when $a \geq f$:

$$\mathbb{E}[Thr_{\text{RARF}}] \geq \frac{T-a}{2T}\Bigg[R_2 + R_1 - (R_2-R_1)\left(\frac{1}{T-a} + \frac{s}{e(T-a)} - \frac{s}{e(T-a)^2}\right)\Bigg]. \quad (5)$$

Using that lower bound, one can establish the following:

Lemma 5: For a bursty periodic jammer (a, T) with a fixed jamming rate and for $a \geq f$, a lower bound on the throughput of two-rate RARF (i.e. Eq. (5)) is minimized when $a = f$.

The following theorem provides a closed-form expression for a lower bound on the minimum jamming rate for RARF.

Theorem 2: To keep the throughput of RARF below R_1, a bursty periodic jammer with $a \geq f$ must satisfy the following:

$$RoJ_{RARF} \geq \frac{f}{\frac{b+\sqrt{b^2-4es}}{2e} + f}\,, \quad \text{where } b = e + s + \frac{2ef}{\frac{R_2}{R_1} - 1}\,.$$

For default parameter values (i.e. $s = 10$, $f = 2$, $R_1 = 1$ Mb/s and $R_2 = 2$ Mb/s), we get $b \approx 23.59$ and $RoJ_{RARF} \geq 19.5\%$.

Proof: Applying Lemma 5 on the lower bound given by Eq. (5) and setting this bound below R_1 yields an upper bound on T. Let $x = T - f$. Then,

$$\frac{x}{2(x+f)}\left[R_2 + R_1 - (R_2 - R_1)\left(\frac{1}{x} + \frac{s}{e\,x} - \frac{s}{e\,x^2}\right)\right] \leq R_1,$$

leading to the following quadratic expression:

$$ex^2 - \underbrace{\left(e + s + \frac{2ef}{\frac{R_2}{R_1} - 1}\right)}_{b} x + s \leq 0\,. \quad (6)$$

The only feasible root of Eq. (6) is $x = (b + \sqrt{b^2 - 4es})/2e$, resulting in the bound given in Theorem 2. ■

Through a coupling argument, one can show that the lower bound in Theorem 2 applies to any n-rate system as in [16].

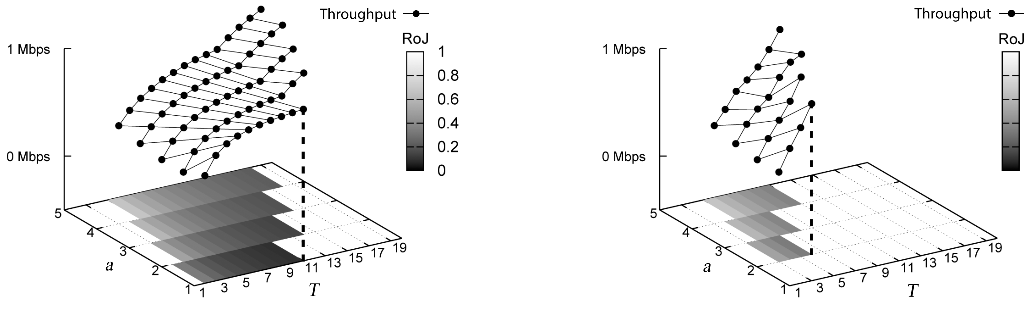

(a)

(b)

Fig. 3: Throughput of (a) ARF and (b) RARF under periodic jamming with parameters a and T. RoJ is depicted as a contour plot at the bottom of each graph. Low RoJ values are plotted in dark colors to indicate the severity of the vulnerability. Among the strategies to keep the throughput below 1 Mb/s, the one with the lowest RoJ is indicated.

VI. NS-3 Simulations

In this section, we present the results of ns-3 [15] simulations of IEEE 802.11g WLANs. Our goals are to monitor the rate used for each packet and to measure the throughput of a system that employs a specific RAA under a given bursty periodic jamming strategy. We use standard ns-3 libraries whenever possible. We build new ns-3 modules for SampleRate and RARF. In all our simulations, we set the length of each DATA packet to 1250 bytes. We use 802.11g in ad-hoc mode since we consider two stations and wish to avoid beacons. Throughput values are based on our definition in Section III-A.

A. ARF and RARF

In Fig. 3, throughput values of ARF and RARF under a bursty periodic jammer are plotted with a and T taken as parameters. $T \in \{1, 2, \ldots, 20\}$ and $a \in \{1, 2, \ldots, 5\}$. Simulations are run for 100 seconds under a perfect channel. RoJ is illustrated as a contour plot for valid jamming strategies at the bottom of each figure. Note that each data point indicates a strategy that manages to keep the throughput below R_1. The optimal strategy should have the lowest RoJ among those.

In Fig. 3(a) for ARF, RoJ is minimized when $a = 1$ and $T = 11$, as given by Proposition 1. On the other hand, Fig. 3(b) reveals the optimal jamming strategy against RARF has the parameters $a = 2$ and $T = 6$, resulting in a jamming rate of 33.3%. Thus, any jamming strategy with RoJ below 33.3% fails to keep RARF throughput below 1 Mb/s. Table I shows that RoJ_{RARF} is sizably higher than the lower bound derived in Theorem 2, since twelve rates are available in IEEE 802.11g. On the other hand, ns-3 simulations of RARF with two bit-rates yield $a = 2$ and $T = 9$ as the optimal values. This leads to an RoJ of 22.2%, which is close to the lower bound of 19.5% derived in Theorem 2.

B. SampleRate

We consider the jamming strategy given by Theorem 2 for SampleRate. Simulation is run for 100 seconds under a perfect channel, and the first 50 seconds are plotted in Fig. 4. Each data point indicates the rate used for a DATA packet transmission. We use an initial jamming phase forcing the rate down to 1 Mb/s. The bursty periodic jammer corrupts 44 consecutive packets after every 10 seconds spent at 1 Mb/s. Since the transmission of acknowledgments take a non-negligible amount of time, the system transmits 890 packets in 10 seconds at 1 Mb/s rather than 1000 packets. The jammer parameters are $a = 44$ and $T = 934$, leading to an RoJ value of 4.7%, close to the value of 4.2% predicted by Theorem 2.

RAA	Analytical Results	Simulation Results
ARF	$RoJ = 9.1\%$	$RoJ = 9.1\%$
SampleRate	$RoJ = 4.2\%$	$RoJ = 4.7\%$
RARF	$RoJ \geq 19.5\%$	$RoJ = 33.3\%$

TABLE I: Analytical and simulation results for minimum jamming rate to keep each RAA throughput below 1 Mb/s.

C. Lossy Channel

In this section, we perform simulations using the `ns3::LogDistancePropagationLossModel` of ns-3. The reception power is $L = L_0 + 10\, n\, \log_{10}(d/d_0)$, where n is the path loss distance exponent, d_0 is the reference distance (m), L_0 is the path loss at reference distance (dB), d is the distance (m), and L is the path loss (dB). The default channel parameter values are $n = 3$, $d_0 = 1\,m$, and $L_0 = 46.677$ dB. Under this channel model, we have evaluated the performances of ARF and RARF. The simulations are run for 100 seconds with no jammer for $d = 70\ m$ and default channel parameters. Fig. 5 plots the rates used for the first 200 DATA packet transmissions for each RAA. One can observe that ARF and RARF perform similarly, leading to close throughput values.

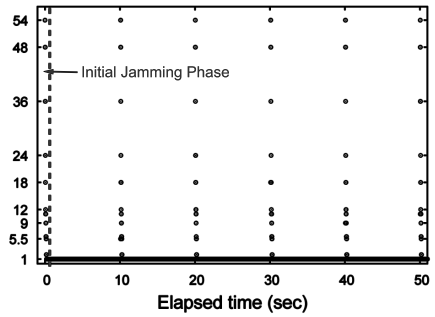

Fig. 4: Performance of SampleRate under bursty periodic jamming with $a = 44$ and $T = 934$. $Thr = 0.955$ Mb/s.

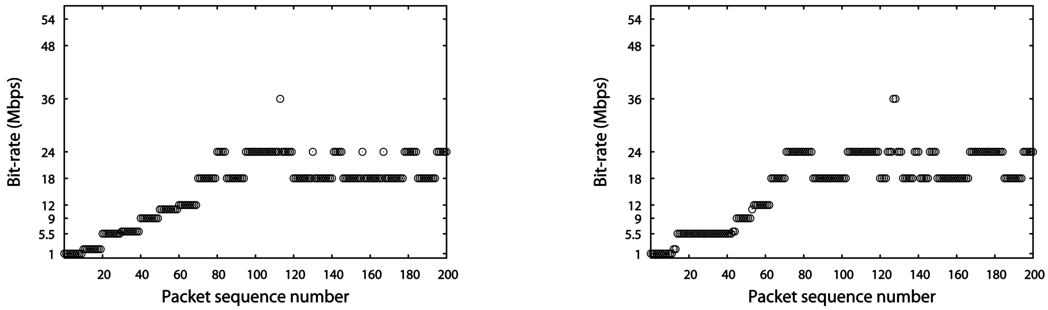

(a)

(b)

Fig. 5: Performance under lossy channel with no jamming: (a) ARF, $Thr = 21.77\,$Mb/s, (b) RARF, $Thr = 23.22\,$Mb/s. ARF and RARF perform similarly under lossy channel conditions in the absence of a jammer.

Next, we have implemented the corresponding effective jamming strategies for each RAA with $d \in \{10, 20, \ldots, 200\}$, $n \in \{1, 2, \ldots, 5\}$, and default values of d_0 and L_0. The results for $d = 100\,m$ and $n = 3$ are given in Table II.

RAA	Jammer Parameters	Thr
ARF	$a = 1$, $T = 11$, $RoJ = 9.1\%$	0.909 Mb/s
SampleRate	$a = 44$, $T = 934$, $RoJ = 4.7\%$	0.951Mb/s
RARF	$a = 2$, $T = 6$, $RoJ = 33.3\%$	0.895 Mb/s

TABLE II: Throughput values of various RAAs with corresponding effective jamming strategies under a lossy channel. The strategies considered for perfect channel keep the throughput below 1 Mb/s under lossy channel conditions as well.

VII. Conclusion

In this paper, we introduced a theoretical framework that employs a bursty periodic jamming model and a rate of jamming metric to analyze the vulnerabilities of widely used RAAs. We proved that the jamming rate required to keep the throughput below the base rate is low for ARF (9.1%), and even lower for SampleRate (4.2%). Thereafter, we proposed a randomized ARF (RARF) and analyzed its performance under the same jamming model. We proved that one needs a jamming rate of at least 19.5% to keep the RARF throughput below the base rate. This bound is tight for two-rate RARF. Using ns-3 simulations, we corroborated our analysis and observed that for IEEE 802.11g, RARF throughput falls below the base rate only for RoJ values above 33%. The strategies considered can be employed for perfect or lossy channels.

Our aim was to demonstrate that a randomized approach in probing or changing transmission rates can provide jamming resistance without drastically altering system performance in the absence of jamming. Although we considered bursty periodic jamming, randomization approach can be applied against any jammer that is unaware of the real-time transmission rate.

Our analytical findings and simulations show that state-of-the-art RAAs are vulnerable to simple jamming attacks. Randomization is an effective means to achieve more robust rate adaptation. A recently proposed RAA called Minstrel [12] includes some level of randomization. Future work could aim at analysis of Minstrel, adopting alternative jamming models and investigating RAA performance under general topologies.

References

[1] AT&T. Wi-Fi services for stadiums. http://www.business.att.com/enterprise/online_campaign/wi-fi-stadiums/.

[2] B. Awerbuch, A. Richa, and C. Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *PODC*, 2008.

[3] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa. On the performance of IEEE 802.11 under jamming. In *INFOCOM*, 2008.

[4] J. C. Bicket. Bit-rate selection in wireless networks. Master's thesis, Massachusetts Intitute of Technology, 2005.

[5] I. Broustis, K. Pelechrinis, D. Syrivelis, S. V. Krishnamurthy, and L. Tassiulas. FIJI: Fighting implicit jamming in 802.11 WLANs. In *SecureComm*, 2009.

[6] W. Chen, Y. Zhang, and Y. Wei. The feasibility of launching reduction of quality (RoQ) attacks in 802.11 wireless networks. In *ICPADS*, 2008.

[7] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan. Understanding and mitigating the impact of RF interference on 802.11 networks. In *SIGCOMM*, 2007.

[8] Jammer-Store. http://www.jammer-store.com/.

[9] JiWire. Global Wi-Fi finder. http://v4.jiwire.com/.

[10] A. Kamerman and L. Monteban. WaveLAN: A high-performance wireless LAN for the unlicensed band. *Bell Labs Technical Journal*, 2(3):118–133, 1997.

[11] M. Lacage, M. H. Manshaei, and T. Turletti. IEEE 802.11 rate adaptation: a practical approach. In *MSWiM*, 2004.

[12] MadWifi. Minstrel rate control.
http://madwifi-project.org/browser/madwifi/trunk/ath_rate/minstrel.

[13] MadWifi. Onoe rate control.
http://madwifi-project.org/browser/madwifi/trunk/ath_rate/onoe.

[14] G. Noubir, R. Rajaraman, B. Sheng, and B. Thapa. On the robustness of IEEE802.11 rate adaptation algorithms against smart jamming. In *WiSec*, 2011.

[15] ns 3 Network Simulator. http://www.nsnam.org/.

[16] C. Orakcal and D. Starobinski. Jamming-resistant rate control in IEEE 802.11 WLANs. CISE Technical Report 2011-IR-0021, Boston University, 2011. Also available as http://www.bu.edu/phpbin/cise/download.php?publication_id=1129.

[17] K. Pelechrinis, I. Broustis, S. V. Krishnamurthy, and C. Gkantsidis. ARES: An anti-jamming reinforcement system for 802.11 networks. In *CoNEXT*, 2009.

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

[19] A. Scott, T.J. Hardy, R.K. Martin, and R.W. Thomas. What are the roles of electronic and cyber warfare in cognitive radio security? In *MWSCAS*, 2011.

[20] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders. Short paper: reactive jamming in wireless networks: how realistic is the threat? In *WiSec*, 2011.

[21] W. Xu, W. Trappe, Y. Zhang, and T. Wood. The feasibility of launching and detecting jamming attacks in wireless networks. In *MOBIHOC*, 2005.