

Competitive and Fair Medium Access despite Reactive Jamming

Andrea Richa	Christian Scheideler	Stefan Schmid	Jin Zhang
Arizona State University	University of Paderborn	T-Labs/TU Berlin	Arizona State University
Tempe, Arizona, USA	D-33102 Paderborn	D-10587 Berlin	Tempe, Arizona, USA
aricha@asu.edu	scheideler@upb.de	stefan@net.t-labs.tu-berlin.de	jzhang82@asu.edu

Abstract—Intentional interference constitutes a major threat for communication networks operating over a shared medium where availability is imperative. Jamming attacks are often simple and cheap to implement. Today’s jammers can perform physical carrier sensing in order to disrupt communication more efficiently, especially in a network of simple wireless devices such as sensor nodes, which usually operate over a single frequency (or a limited frequency band) and which cannot benefit from the use of spread spectrum or other more advanced technologies. This paper proposes the medium access (MAC) protocol ANTIJAM which is provably robust against a powerful reactive adversary who can jam a $(1 - \varepsilon)$ -portion of the time steps, where ε is an arbitrary constant. The adversary uses carrier sensing to make informed decisions on when it is most harmful to disrupt communications. Moreover, we allow the adversary to be adaptive and to have complete knowledge of the entire protocol history. Our MAC protocol is able to make efficient use of the nonjammed time periods and achieves a $\Theta(1)$ -competitive throughput in this harsh scenario, if ε is constant. In addition, ANTIJAM features a low convergence time and has excellent fairness properties in the sense that channel access probabilities among nodes do not differ by more than a small constant factor.

I. INTRODUCTION

Disruptions of the communication over a shared medium—either because of interference of concurrent transmissions or intentionally—are a central challenge in wireless computing. It is well-known that already simple jamming attacks—without any special hardware—constitute a threat for the widely used IEEE 802.11 MAC protocol. Due to the problem’s relevance, there has been a significant effort to cope with such disruption problems both from the industry and the academia side, and much progress has been made over the last years on how to deal with different jammer types.

While simple oblivious jammers are well-understood today and many countermeasures exist, this paper goes an important step further and studies MAC protocols against smart jammers. In particular,

we argue that adversaries may behave in an adaptive and reactive manner: *adaptive* in the sense that their decisions on whether to jam at a certain moment in time can depend on the protocol history; and *reactive* in the sense that the adversary can perform physical carrier sensing (which is part, e.g., of the 802.11 standard) to learn whether the channel is currently idle or not, and jam the medium depending on these measurements.

This paper presents the medium access (MAC) protocol ANTIJAM that makes effective use of the few and arbitrarily distributed nonjammed time periods, to achieve fairness and a provable throughput despite the presence of such a strong reactive jammer. As we will see, the throughput is competitive, i.e., a constant fraction of the nonjammed time period is used for successful transmissions. Besides this interesting theoretical result, our protocol is simple to implement and performs well also in the average case. Also, worth to note is that our approach is at the MAC level and may be used in conjunction with some of the anti-jamming techniques developed at the physical layer (e.g., frequency hopping, spread spectrum).

A. Related Work

Researchers have studied the problem of unintentional and malicious interference in wireless networks for several years now, e.g., [6], [14], [16], [17], [18], [20], [21], [27], [29]. Classic defense mechanisms operate on the physical layer [18], [20] and there exist approaches both to *avoid* as well as to *detect* jamming. Spread spectrum and frequency hopping technologies have been shown to be very effective to avoid jamming with widely spread signals. However, the ISM frequency band used by IEEE 802.11 variants is too narrow to effectively apply spread spectrum techniques [5]. Also, as jamming strategies can come in many different flavors, detecting jamming activities by simple methods based on signal strength, carrier sensing, or packet delivery

ratios has turned out to be quite difficult [17].

Recent work has also studied *MAC layer strategies* against jamming, including coding strategies (e.g., [6]), channel surfing and spatial retreat (e.g., [1], [30]), or mechanisms to hide messages from a jammer, evade its search, and reduce the impact of corrupted messages (e.g., [28]). However, these methods do not help against an adaptive jammer with *full* information about the history of the protocol, like the one considered in our work.

In the theory community, work on MAC protocols has mostly focused on efficiency. Many of these protocols are *random backoff protocols* (e.g., [4], [7], [8], [12], [23]) that do not take jamming activity into account and, in fact, are not robust against it (see [2] for more details). But also some theoretical work on *jamming* is known (e.g., [9] for a short overview). There are two basic approaches in the literature. The first assumes randomly corrupted messages (e.g. [22]), which is much easier to handle than adaptive adversarial jamming [3]. The second line of work either bounds the number of messages that the adversary can transmit or disrupt with a limited energy budget (e.g. [11], [15]), or bounds the number of channels the adversary can jam (e.g. [10], [19]). The protocols in, e.g., [15] can tackle adversarial jamming at both the MAC and network layers, where the adversary may not only be jamming the channel but also introducing malicious (fake) messages (possibly with address spoofing). However, they depend on the fact that the adversarial jamming budget is finite, so it is not clear whether the protocols would work under heavy continuous jamming. (The result in [11] seems to imply that a jamming rate of 0.5 is the limit whereas the handshaking mechanisms in [15] seem to require an even lower jamming rate.)

Our work is motivated by the results in [3] and [2]. In [3] it is shown that an adaptive jammer can dramatically reduce the throughput of the standard MAC protocol used in IEEE 802.11 with only limited energy cost on the adversary side. Awerbuch et al. [2] initiated the study of throughput-competitive MAC protocols under continuously running, adaptive jammers, and presented a protocol that achieves a high performance under adaptive jamming.

In this paper, we extend the model and result from [2] in two crucial ways. (1) We allow the jammer to be *reactive*, i.e., to listen to the current channel state in order to make smarter jamming decisions. Note that a reactive model is not only meaningful in the context of jamming: for example,

in many MAC protocols based on carrier sensing, nodes become active during idle time periods and hence, a MAC protocol in the reactive model also performs well in scenarios with *co-existent networks*. (2) We design a fair protocol in the sense that channel access probabilities among nodes do not differ by more than a small constant factor. The protocol in [2] was inherently unfair, as confirmed by our theoretical and simulation results. We believe that the reactive jammer model is much more realistic and hence that our study is of practical importance. For example, by sensing the channel, the adversary may avoid wasting energy by not jamming idle rounds. Note however that depending on the protocol, it may still make sense for the adversary to jam idle rounds, e.g., to influence the protocol execution. Indeed, due to the large number of possible strategies a jammer can pursue, the problem becomes significantly more challenging than the nonreactive version: Not only is the analysis more involved, but also key modifications to the protocol in [2] were needed. While we still build upon the algorithmic ideas presented in [2], our ANTIJAM protocol seeks to synchronize the nodes' sending probabilities. This has the desirable effect of achieving fairness, where all nodes are basically granted the same channel access probabilities, greatly improving the unfair protocol of [2]. While our formal analysis confirms our expectations that the overall throughput under reactive jammers is lower than the throughput obtainable against nonreactive jammers, we are still able to prove a constant-competitive performance (for constant ϵ), which is also confirmed by our simulation study. Finally, our first insights indicate that ANTIJAM-like strategies can also be used in multi-hop settings (see also the recent extension of [2] to *unit disk graphs* [25]) and to devise robust applications such as leader election protocols [26].

B. Model

We study a wireless network that consists of n honest and reliable simple wireless devices (e.g., sensor nodes) that are within the transmission range of each other and that communicate over a single frequency (or a limited, narrow frequency band). We assume a back-logged scenario where the nodes continuously contend for sending a packet on the wireless channel. A node may either transmit a message or sense the channel at a time step, but it cannot do both, and there is no immediate feedback mechanism telling a node whether its transmission was successful. A node sensing the channel may

either (i) sense an *idle channel* (in case no other node is transmitting at that time), (ii) sense a *busy channel* (in case two or more nodes transmit at the time step), or (iii) *receive* a packet (in case exactly one node transmits at the time step).

In addition to these nodes there is an adversary. We allow the adversary to know the protocol and its entire history and to use this knowledge in order to jam the wireless channel at will at any time (i.e., the adversary is *adaptive*). Whenever it jams the channel, all nodes will notice a busy channel. However, the nodes cannot distinguish between the adversarial jamming or a collision of two or more messages that are sent at the same time. We assume that the adversary is only allowed to jam a $(1 - \varepsilon)$ -fraction of the time steps, for an arbitrary constant $0 < \varepsilon \leq 1$.

Moreover, we allow the jammer to be *reactive*: it is allowed to make a jamming decision based on the actions of the nodes at the *current* step. In other words, reactive jammers can determine (through physical carrier sensing) whether the channel is currently idle or busy (either because of a successful transmission, a collision of transmissions, or too much background noise) and can instantly make a jamming decision based on that information. Those jammers arise in scenarios where, for example, encryption is being used for communication and where the jammer cannot distinguish between an encrypted package and noise in the channel. Note that robustness in the reactive model is relevant beyond jamming, e.g., in situations with co-existent networks, as many MAC protocols based on carrier sensing activate nodes during idle time periods.

In addition, we allow the adversary to perform *bursty* jamming. More formally, an adversary is called $(T, 1 - \varepsilon)$ -bounded for some $T \in \mathbb{N}$ and $0 < \varepsilon < 1$ if for any time window of size $w \geq T$ the adversary can jam at most $(1 - \varepsilon)w$ of the time steps in that window.

The network scenario described above arises, for example, in sensor networks, which consist of simple wireless nodes usually running on a single frequency and which cannot benefit from more advanced anti-jamming techniques such as frequency hopping or spread spectrum. In such scenarios, a jammer will also most probably run on power-constrained devices (e.g., solar-powered batteries), and hence will not have enough power to continuously jam over time. (The time window threshold T can be chosen large enough to accommodate the respective jamming pattern.)

This paper studies *competitive* MAC protocols. A MAC protocol is called c -competitive against some $(T, 1 - \varepsilon)$ -bounded adversary (with high probability or on expectation) if, for any sufficiently large number of time steps, the nodes manage to perform successful message transmissions in at least a c -fraction of the time steps not jammed by the adversary (with high probability or on expectation).

Our goal is to design a *symmetric local-control* MAC protocol (i.e., there is no central authority controlling the nodes, and the nodes have symmetric roles at any point in time) that is fair and $O(1)$ -competitive against any $(T, 1 - \varepsilon)$ -bounded adversary. The nodes do not know ε , but we do allow them to have a very rough upper bound of the number n and T . More specifically, we will assume that the nodes have a common parameter $\gamma = O(1/(\log T + \log \log n))$. This is still scalable, since such an estimate leaves room for a super-polynomial change in n and a polynomial change in T over time, so it does not make the problem trivial (as it would be the case if the nodes knew constant factor approximations for n or T).

C. Our Contributions

This paper introduces and analyzes the medium access protocol ANTIJAM. ANTIJAM is robust to a strong adaptive and reactive jammer, who can block the medium a constant fraction of the time and thus models a large range of (intentional and unintentional) interference models. Nevertheless, we can show that the ANTIJAM MAC protocol achieves a high throughput performance by exploiting any nonblocked time intervals effectively. The main theoretical contribution is the derivation of the following theorem that shows that ANTIJAM is competitive in the sense that a constant fraction of the nonjammed execution time is used for successful transmissions, i.e., ANTIJAM is able to benefit from the rare and hard-to-predict time intervals where the shared medium is available. Moreover, ANTIJAM converges fast and allocates the shared medium *fairly* to the nodes.

Theorem I.1. *Let $N = \max\{T, n\}$. The ANTIJAM protocol is constant-competitive, namely $e^{-\Theta(1/\varepsilon^2)}$ -competitive w.h.p., under any $(T, 1 - \varepsilon)$ -bounded reactive adversary if the protocol is executed for at least $\Theta(\frac{1}{\varepsilon} \log N \max\{T, (e^{\delta/\varepsilon^2}/\varepsilon\gamma^2) \log^3 N\})$ many time steps, where $\varepsilon \in (0, 1)$ is a constant, $\gamma = O(1/(\log T + \log \log n))$, and where δ is a sufficiently large constant. Moreover, ANTIJAM*

achieves high fairness: the channel access probabilities among nodes do not differ by more than a factor of $(1 + \gamma)$ after the first message was sent successfully.

We believe that ANTIJAM is interesting also from a practical point of view, as—in contrast to the analysis—the basic protocol is very simple.

D. Paper Organization

The remainder of this paper is organized as follows. Section II introduces the ANTIJAM MAC protocol. Subsequently, we present a formal analysis of the throughput performance under reactive jamming (Section III). Section IV reports on the insights gained from our simulation experiments. The paper is concluded in Section V. Due to space constraints, not all proofs and simulation results are included in this article, and the interested reader is referred to the ArXiv Technical Report 1007.4389 [24].

II. THE ANTIJAM MAC PROTOCOL

The basic ideas of the ANTIJAM MAC protocol are inspired by slotted ALOHA schemes where nodes change their access probabilities over time, in particular the protocols described in [13] (which also uses access probabilities depending on the ratio between idling and successful time slots) and particularly [2]. However, the algorithm in [2] does not achieve a provable performance under reactive jammers, which is due to the asymmetric access probabilities. Therefore, in our protocol, we explicitly try to equalize access probabilities, which also improves fairness among the nodes.

Each node v maintains a time window threshold estimate T_v and a counter c_v . The parameter γ is the same for every node and is set to some sufficiently small value in $O(1/(\log T + \log \log n))$. Note that this assumption is not critical as it allows for a high scalability: the nodes only need some polynomial estimate of T and even rougher estimate of n .¹ Let \hat{p} be any constant so that $0 < \hat{p} \leq 1/24$. Initially, every node v sets $T_v := 1$, $c_v := 1$ and $p_v := \hat{p}$. Afterwards, the protocol works in synchronized time steps. We assume synchronized time steps for the analysis, but a nonsynchronized execution of the protocol would also work as long as all nodes operate at roughly the same speed.

The basic protocol idea is simple. Suppose that each node v decides to send a message at the current

¹The assumption of having a constant factor approximation of T and n would render the problem simple.

time step with probability p_v with $p_v \leq \hat{p}$. Let $p = \sum_v p_v$, q_0 be the probability that the channel is idle and q_1 be the probability that exactly one node is sending a message. The following claim appeared originally in [2]. It states that if $q_0 = \Theta(q_1)$, then the cumulative sending probability p is constant, which in turn implies that at any nonjammed time step we have constant probability of having a successful transmission. Hence our protocol aims at adjusting the sending probabilities p_v of the nodes such that $q_0 = \Theta(q_1)$, in spite of the reactive adversarial jamming activity. This will be achieved by using a multiplicative increase/decrease mechanism for the probabilities p_v and by synchronizing all the nodes, both in terms of sending probabilities and their own estimates T_v on the time window T , at every successful transmission.

Claim II.1. $q_0 \cdot p \leq q_1 \leq \frac{q_0}{1-\hat{p}} \cdot p$.

Now we present our ANTIJAM protocol:

In each step, each node v does the following. v decides with probability p_v to send a message along with a tuple: (p_v, c_v, T_v) . If it decides not to send a message, it checks the following two conditions:

- 1) If v senses an idle channel, then $p_v := \min\{(1 + \gamma)p_v, \hat{p}\}$ and $T_v := T_v - 1$.
- 2) If v successfully receives a message along with the tuple of $(p_{new}, c_{new}, T_{new})$, then $p_v := (1 + \gamma)^{-1}p_{new}$, $c_v := c_{new}$, and $T_v := T_{new}$.

Afterwards, v sets $c_v := c_v + 1$. If $c_v > T_v$ then it does the following: v sets $c_v := 1$, and if there was no idle step among the past T_v time steps, then $p_v := (1 + \gamma)^{-1}p_v$ and $T_v := T_v + 2$.

III. ANALYSIS

Let V be the set of all nodes. Let $p_t(v)$ be node v 's access probability p_v at the beginning of the t -th time step. Furthermore, let $p_t = \sum_{v \in V} p_t(v)$. Let I be a time frame consisting of $\frac{\alpha}{\varepsilon} \log N$ subframes I' of size $f = \max\{T, \frac{\alpha\beta^2}{\varepsilon\gamma^2} e^{\delta/\varepsilon^2} \log^3 N\}$, where α , β and δ are sufficiently large constants. Let $F = \frac{\alpha}{\varepsilon} \log N \cdot f$ denote the size of I .

First, we will derive some simple facts on the behavior of ANTIJAM. We then show that given a certain sufficiently large initial cumulative probability p_t in a subframe, the cumulative probability cannot be smaller at the end of the subframe. We proceed to show that ANTIJAM performs well in time periods in which p_t is upper bounded by δ/ε^2 for some constant δ . Finally, we show that for

any jamming strategy, ANTIJAM has a cumulative probability of $p_t \leq \delta/\varepsilon^2$ for most of the time, which yields our main theorem.

In our analysis, we will use the following relation.

Lemma III.1. *For all $0 < x < 1$ it holds that*

$$e^{-x/(1-x)} \leq 1 - x \leq e^{-x}$$

Moreover, we make extensive use of the following Chernoff bounds.

Lemma III.2. *Consider any set of binary random variables X_1, \dots, X_n . Suppose that there are values $p_1, \dots, p_n \in [0, 1]$ with $\mathbb{E}[\prod_{i \in S} X_i] \leq \prod_{i \in S} p_i$ for every set $S \subseteq \{1, \dots, n\}$. Then it holds for $X = \sum_{i=1}^n X_i$ and $\mu = \sum_{i=1}^n p_i$ and any $\delta > 0$ that*

$$\mathbb{P}[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu \leq e^{-\frac{\delta^2 \mu}{2(1+\delta/3)}}.$$

If, on the other hand, it holds that $\mathbb{E}[\prod_{i \in S} X_i] \geq \prod_{i \in S} p_i$ for every set $S \subseteq \{1, \dots, n\}$, then it holds for any $0 < \delta < 1$ that

$$\mathbb{P}[X \leq (1 - \delta)\mu] \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu \leq e^{-\delta^2 \mu/2}.$$

We start with some simple facts. Fact III.3 shows that the protocol synchronizes the sending probabilities of the nodes (up to a factor of $(1 + \gamma)$), and that all values c_v and T_v are also synchronized.

Fact III.3. *Right after a successful transmission of the tuple (p', c', T') , $(p_v, c_v, T_v) = ((1 + \gamma)^{-1} p', c', T')$ for all receiving nodes v and $(p_u, c_u, T_u) = (p', c', T')$ for the sending node u . In particular, for any time step t after a successful transmission by node u , $(c_v, T_v) = (c_w, T_w)$ for all nodes $v, w \in V$.*

Fact III.3 also implies the following corollary.

Corollary III.4. *After a successful transmission, the access probabilities p_v of the nodes $v \in V$ will never differ by more than a factor $(1 + \gamma)$ in the future.*

The next fact follows from the protocol and Fact III.3, and they help one understand how the cumulative probabilities vary over time with successful transmissions, idle time steps, etc.

Fact III.5. *For any time step t after a successful transmission or a well-initialized state of the protocol (in which $(p_v, c_v, T_v) = (\hat{p}, 1, 1)$ for all nodes v) it holds:*

I. *If the channel is idle at time t then (i) if $p_v = \hat{p}$ for all v , then $p_{t+1} = p_t$; (ii) if $p_u = \hat{p}$ and*

$p_v = (1 + \gamma)^{-1} \hat{p}$ for all nodes $v \neq u$, then $p_{t+1} = (1 + \gamma - O(1/n)) p_t$ (because all nodes except for u increase their sending probability by a factor $(1 + \gamma)$ from $\hat{p}/(1 + \gamma)$); or (iii) if $p_v < \hat{p}$ for all nodes v , then $p_{t+1} = (1 + \gamma) p_t$.

2. *If there is a successful transmission at time t , and if $c_v \leq T_v$ or there was an idle time step in the previous T_v rounds, then (i) if the sender is the same as the last successful sender, then $p_{t+1} = p_t$ (because for the sender u , $p_u(t+1) = p_u(t)$, and the other nodes remain at $p_u(t+1)/(1 + \gamma) = p_u(t)/(1 + \gamma)$); if (ii) the sender w is different from the last successful sender u and $p_v = \hat{p}$ for all nodes v (including u and w), then $p_{t+1} = (1 + \gamma - O(1/n))^{-1} p_t$ (all nodes except w reduce their sending probability); or (iii) if the sender w is different from the last successful sender u and $p_v < \hat{p}$ for at least one node v (including u and w), then $p_{t+1} = (1 + \gamma)^{-1} p_t$ (because at time t , for all nodes $v \neq u$: $p_v(t) = p_u(t)/(1 + \gamma)$; subsequently, $p_w(t+1) = p_w(t)$ and for all nodes $v \neq w$: $p_v(t+1) = p_w(t+1)/(1 + \gamma)$).*

3. *If the channel is busy at time t , then $p_{t+1} = p_t$ when ignoring the case that $c_v > T_v$.*

Whenever $c_v > T_v$ and there has not been an idle time step during the past T_v steps, then p_{t+1} is, in addition to the actions specified in the two cases above, reduced by a factor of $(1 + \gamma)$.

The next two lemmas bound how many times T_v can be increased in a subframe and the minimum cumulative probability we should expect w.h.p., given some sufficiently large initial cumulative probability at the start of a subframe.

Lemma III.6. *If in subframe I' the number of idle time steps is at most k , then every node v increases T_v by 2 at most $k/2 + \sqrt{f}$ many times in I' .*

Lemma III.7. *For any subframe I' in which initially $p_{t_0} \geq 1/(f^2(1 + \gamma)^{\sqrt{2f}})$, the last time step t of I' again satisfies $p_t \geq 1/(f^2(1 + \gamma)^{\sqrt{2f}})$, w.h.p.*

The proofs of both lemmas follow from the arguments in the proofs of similar lemmas in [2].

Lemma III.8 shows that for times of low cumulative probabilities, ANTIJAM yields a good performance.

Lemma III.8. *Consider any subframe I' , and let $\delta > 1$ be a sufficiently large constant. Suppose that at the beginning of I' , $p_{t_0} \geq 1/(f^2(1 + \gamma)^{\sqrt{2f}})$ and $T_v \leq \sqrt{F}/2$ for every node v . If $p_t \leq \delta/\varepsilon^2$ for at least half of the nonjammed time steps in I' , then ANTIJAM is at least $\frac{\delta}{8(1-\hat{p})\varepsilon^2} e^{-\delta/(1-\hat{p})\varepsilon^2}$.*

competitive in I' .

Proof: A time step t in I is called *useful* if we either have an idle channel or a successful transmission at time t (i.e., the time step is not jammed and there are no collisions) and $p_t \leq \delta/\varepsilon^2$. Let k be the number of useful time steps in I' . Furthermore, let k_0 be the number of useful time steps in I' with an idle channel, k_1 be the number of useful time steps in I' with a successful transmission and k_2 be the maximum number of times a node v reduces p_v in I' because of $c_v > T_v$. Recall that $k = k_0 + k_1$. Moreover, the following claim holds:

Claim III.9. *If $n \geq (1 + \gamma)\delta/(\varepsilon^2\hat{p})$, then*

$$k_0 - \log_{1+\gamma}(\delta/(\varepsilon^2 \cdot p_{t_0})) \leq k'_1 + k_2$$

where k'_1 is the number of useful time steps with a successful transmission in which the sender is different from the previously successful sender.

Proof: According to Fact III.5, $p_v \in [(1 + \gamma)^{-1}p, p]$ for some access probability p for all time steps in I' . Hence, if $p_t \leq \delta/\varepsilon^2$ and $n \geq (1 + \gamma)\delta/(\varepsilon^2\hat{p})$, then $p_v(t) \leq \hat{p}/(1 + \gamma)$ for all v . This implies that whenever there is a useful time step $t \in I$ with an idle channel, then $p_{t+1} = (1 + \gamma)p_t$. Thus, it takes at most $\log_{1+\gamma}(\delta/(\varepsilon^2 \cdot p_{t_0}))$ many useful time steps with an idle channel to get from p_{t_0} to a cumulative probability of at least δ/ε^2 . On the other hand, each of the k'_1 successful transmissions reduces the cumulative probability by $(1 + \gamma)$. Therefore, once the cumulative probability is at δ/ε^2 , we must have $k_0 \leq k'_1 + k_2$ since otherwise there must be at least one useful time step where the cumulative probability is more than δ/ε^2 , which contradicts the definition of a useful time step. ■

Since $p_{t_0} \geq 1/(f^2(1 + \gamma)^{\sqrt{2f}})$ it holds that

$$\log_{1+\gamma}(\delta/(\varepsilon^2 \cdot p_{t_0})) \leq \log_{1+\gamma}(\delta f^2/\varepsilon^2) + \sqrt{2f}$$

From Lemma III.6 we also know that $k_2 \leq k_0/2 + \sqrt{f}$. Hence,

$$\begin{aligned} k_0 &\leq 2k'_1 + 2 \cdot \log_{1+\gamma}(\delta f^2/\varepsilon^2) + 2 \cdot (\sqrt{f} + \sqrt{2f}) \\ &\leq 2k'_1 + 6\sqrt{f} \end{aligned}$$

if f is sufficiently large. Also, $k_0 = k - k_1$ and $k'_1 \leq k_1$. Therefore, $k - k_1 \leq 2k_1 + 6\sqrt{f}$ or equivalently,

$$k_1 \geq k/3 - 2\sqrt{f}$$

It remains to find a lower bound for k .

Claim III.10. *Let g be the number of nonjammed time steps t in I' with $p_t \leq \delta/\varepsilon^2$. If $g \geq \varepsilon f/2$ then*

$$k \geq \frac{\delta}{2(1 - \hat{p})\varepsilon^2} e^{-\delta/(1 - \hat{p})\varepsilon^2} \cdot g$$

w.h.p.

Proof: Consider any $(T, 1 - \varepsilon)$ -bounded jammer for I' . Suppose that of the nonjammed time steps t with $p_t \leq \delta/\varepsilon^2$, s_0 have an idle channel and s_1 have a busy channel. It holds that $s_0 + s_1 = g \geq \varepsilon f/2$. For any one of the nonjammed time steps with an idle channel, the probability that it is useful is one, and for any one of the nonjammed time steps with a busy channel, the probability that it is useful (in this case, that it has a successful transmission) is at least

$$\begin{aligned} \sum_v p_v \prod_{w \neq v} (1 - p_w) &\geq \frac{1}{1 - \hat{p}} \sum_v p_v \prod_w (1 - p_w) \\ &\geq \frac{1}{1 - \hat{p}} \sum_v p_v \prod_w e^{-p_w/(1 - \hat{p})} \\ &= \frac{1}{1 - \hat{p}} \sum_v p_v e^{-p/(1 - \hat{p})} \\ &= \frac{p}{1 - \hat{p}} e^{-p/(1 - \hat{p})} \end{aligned}$$

where p is the cumulative probability at the step. Since $p_t \leq \delta/\varepsilon^2$, it follows that the probability of a busy time step to be useful is at least

$$\frac{\delta}{(1 - \hat{p})\varepsilon^2} e^{-\delta/(1 - \hat{p})\varepsilon^2}$$

Thus,

$$\begin{aligned} \mathbb{E}[k] &\geq s_0 + \frac{\delta}{(1 - \hat{p})\varepsilon^2} e^{-\delta/(1 - \hat{p})\varepsilon^2} s_1 \\ &\geq \frac{\delta}{(1 - \hat{p})\varepsilon^2} e^{-\delta/(1 - \hat{p})\varepsilon^2} \cdot g \end{aligned}$$

since k is minimized for $s_0 = 0$ and $s_1 = g$.

Since our lower bound for the probability of a busy step to be useful holds independently for all nonjammed busy steps t with $p_t \leq \delta/\varepsilon^2$ and $E[k] \geq \alpha \log N$ for our choice of g , it follows from the Chernoff bounds that $k \geq \mathbb{E}[k]/2$ w.h.p. ■

From Claim III.10 it follows that

$$k_1 \geq \left(\frac{\delta}{2(1 - \hat{p})\varepsilon^2} e^{-\delta/(1 - \hat{p})\varepsilon^2} \cdot g \right) / 3 - 2\sqrt{f}$$

w.h.p., which completes the proof of Lemma III.8. ■

It remains to consider the case that for less than half of the nonjammed time steps t in I' , $p_t \leq \delta/\varepsilon^2$. Fortunately, this does not happen w.h.p.

Lemma III.11. *Suppose that at the beginning of I' , $T_v \leq \sqrt{F}/2$ for every node v . Then at most half of the nonjammed time steps t can have the property that $p_t > \delta/\varepsilon^2$ w.h.p.*

The full proof is quite complex and appears in the ArXiv Technical Report 1007.4389 [24].

Notice that by the choice of f and F , T_v never exceeds $\sqrt{F}/2$ for any v when initially $T_v = 1$ for all v . Hence, the prerequisites of the lemmas are satisfied. We can also show the following lemma, which shows that T_v remains bounded over time.

Lemma III.12. *For any time frame I in which initially $T_v \leq \sqrt{F}/2$ for all v , also $T_v \leq \sqrt{F}/2$ for all v at the end of I w.h.p.*

Proof: We already know that in each subframe I' in I , at least $\varepsilon f/2$ of the nonjammed time steps t in I' satisfy $p_t \leq \delta/\varepsilon^2$ w.h.p. Hence, for all $(T, 1 - \varepsilon)$ -bounded jamming strategies, there are at least

$$(\delta/\varepsilon^2) \cdot e^{-\delta/\varepsilon^2} \cdot \varepsilon f/2$$

useful time steps in I' w.h.p. Due to the lower bound of $p_t \geq 1/(f^2(1+\gamma)^{\sqrt{f}})$ for all time steps in I w.h.p. we can also conclude that

$$k_0 \geq k'_1 + k_2 - \log_{1+\gamma}((\delta/\varepsilon^2) \cdot f^2(1+\gamma)^{\sqrt{f}}).$$

From the technical report [24], it follows that

$$k_0 \geq k_1/3$$

w.h.p. Since $k_0 + k_1 = k$ and $k \geq (\delta/\varepsilon^2) \cdot e^{-\delta/\varepsilon^2} \cdot \varepsilon f/2$ it follows that $k_0 = \Omega(f)$. Therefore, there must be at least one time point in I' with $T_v = 1$ for all $v \in V$. This in turn ensures that $T_v \leq \sqrt{F}/2$ for all v at the end of I w.h.p. ■

With Lemma III.12, we show that Lemma III.11 is true for a polynomial number of subframes. Then, Lemma III.11 and Lemma III.12 together imply that Lemma III.8 holds for a polynomial number of subframes. Together with the fairness result (Corollary III.4), our main Theorem I.1 follows. Finally, note that along the same line as in [2], we can show that ANTIJAM is self-stabilizing, so the throughput result can be extended to an arbitrary sequence of time frames.

IV. SIMULATION

We have implemented a simulator to study additional properties of our protocol. This section reports on some of our results. The focus here is on the qualitative nature of the performance of ANTIJAM,

and we did not optimize the parameters to obtain the best constants.

Note that the formal guarantees derived in the previous section hold for *any* type of reactive jammer that falls within our $(T, 1 - \varepsilon)$ -bounded adversary model. However, for the simulations, specific instantiations need to be considered. Concretely, we use three different jamming strategies, for different ε values and where $T = 100$: (1) one that jams busy channels with probability $(1 - \varepsilon)$; (2) one that jams busy channels deterministically (as long the jamming budget is not used up); (3) one that jams *idle* channels deterministically (as long as the jamming budget is not used up). (At first sight, the scenario where only idle time periods are jammed seems to be of less relevance. However, observe that this attack may be effective against certain protocols to steer the execution into inefficient states, e.g., against protocols that only increase sending probabilities in idle rounds.)

We define throughput as the number of successful transmissions over the number of nonjammed time steps.

A. Throughput

In a first set of experiments we study the throughput as a function of the network size and ε . We evaluate the throughput performance for each type of adversary introduced above. We find that for all three strategies, the throughput is basically constant (i.e., independent of the network size) and, depending on the scenario, between 20 and 40 %. (Please refer to the ArXiv Technical Report 1007.4389 [24] for the corresponding figures and a more detailed discussion.) This is in accordance with our theoretical insight of Theorem I.1. We observed that given our conditions on ε and T , the strategy that jams busy channels deterministically results in the lowest throughput. Hence, in the remaining experiments described in this section, we will focus on this particular strategy. As expected, jamming idle channels does not affect the protocol behavior much.

In our simulations, ANTIJAM makes effective use of the nonjammed time periods, yielding 20%-40% successful transmissions even without optimizing the protocol parameters. In additional experiments we also studied the throughput as a function of γ , see Figure 1. As expected, the throughput declines slightly for large γ , but this effect is small. (Note that for very small γ , the convergence time becomes large and the experiments need run for a long time in order not to underestimate the real throughput.)




Figure 1. Throughput as a function of γ under three different jamming strategies. *Left:* $\varepsilon = 0.2$, *Right:* $\varepsilon = 0.5$).

B. Convergence Time

Besides a high throughput, fast convergence is the most important performance criterion of a MAC protocol. The traces in Figure 2 show the evolution of the cumulative probability over time. It can be seen that the protocol converges quickly to constant access probabilities. (Note the logarithmic scale.) If the initial probability for each node is high, the protocol needs more time to bring down the low-constant cumulative probability. Moreover, the ratio of the time period the cumulative probability is in the range of $[\frac{1}{2\varepsilon}, \frac{2}{\varepsilon}]$ to the time period the protocol being executed is 92.98% when $\hat{p} = 1/24$, and 89.52% when $\hat{p} = 1/2$. This implies that for a sufficiently large time period, the cumulative probability is well bounded most of the time, which corresponds to our theoretical insights. Figure 3 studies the convergence time for different network sizes. We ran the protocol 50 times, and assume that the execution has converged when the cumulative probability $p \in [0.1, 10]$, for at least 5 consecutive rounds. The simulation result also confirms our theoretical analysis in Theorem I.1, as the number of rounds needed to converge the execution is bounded by $\Theta(\frac{1}{\varepsilon} \log N \max\{T, \frac{1}{\varepsilon\gamma^2} \log^3 N\})$.

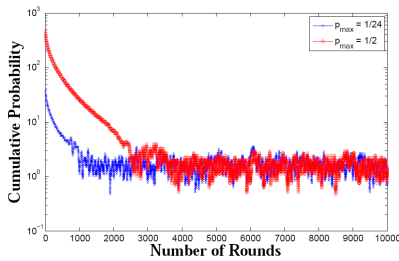


Figure 2. Evolution of cumulative probability over time (network size is 1000 nodes, and $\varepsilon = 0.5$). Note that the plot has logarithmic scale.

Figure 4 indicates that independently of the initial




Figure 3. ANTIJAM runtime as a function of network size for $\hat{p} = 1/24$, and $\varepsilon = 0.5$.

values \hat{p} and T_v , the throughput rises quickly (up above 20%) and stays there afterwards.

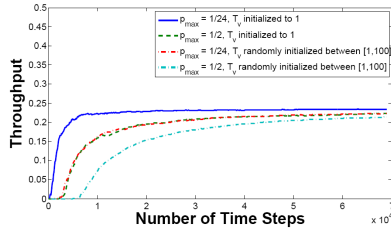


Figure 4. Convergence in a network of 1000 nodes where $\varepsilon = 0.5$.

C. Fairness

As ANTIJAM synchronizes c_v , T_v , and p_v values upon message reception, the nodes are expected to transmit roughly the same amount of messages; in other words, our protocol is fair. Figure 5 presents a histogram showing how the successful transmissions are distributed among the nodes. More specifically, we partition the number of successful transmissions into intervals of size 4. Then, all the transmissions are grouped according to those intervals in the histogram.

We also compare minimum throughputs achieved by individual nodes, when using the MAC protocol

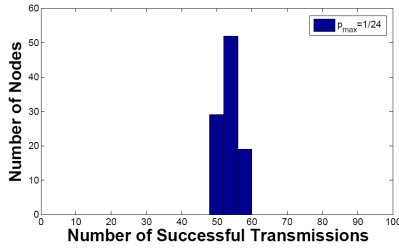


Figure 5. Fairness in a network of 1000 nodes, where $\varepsilon = 0.5$, and $\hat{p} = 1/24$ (averaged over 10 runs).

in [2] and ANTIJAM, so as to demonstrate the fairness property of the latter protocol. Figure 6 clearly shows that ANTIJAM achieves much higher fairness than the MAC protocol in [2]: since the minimum throughput produced by ANTIJAM is significantly higher than that of [2]. Moreover, according to Figure 7, the $\frac{MIN}{MAX}$ throughput ratio of ANTIJAM is also significantly higher than that of [2]. This indicates the difference between the minimum and maximum throughput achieved by ANTIJAM is much smaller, and hence ANTIJAM is fair while [2] is not.

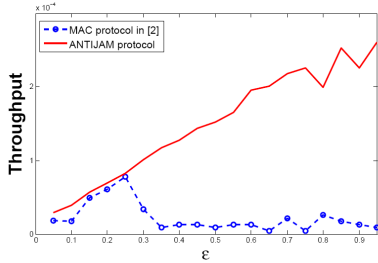


Figure 6. Minimum throughput as a function of $\varepsilon \in [0.05, 0.95]$, compared to the MAC protocol in [2], averaged over 10 runs, where $\hat{p} = 1/24$.

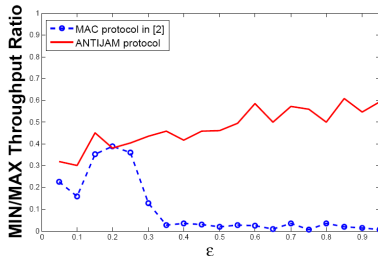


Figure 7. MIN/MAX throughput ratio as a function of $\varepsilon \in [0.05, 0.95]$, compared to the MAC protocol in [2], averaged over 10 runs, where $\hat{p} = 1/24$.

D. Comparison to 802.11

Finally, to put ANTIJAM into perspective, as a comparison, we implemented a simplified version of the widely used 802.11 MAC protocol (with a focus on 802.11a).

The configurations for the simulation are the following: (1) the jammer is reactive and $(T, 1 - \varepsilon)$ -bounded; (2) the unit slot time for 802.11 is set to $50\mu s$; for simplicity, we define one time step for ANTIJAM to be $50\mu s$ also; (3) we run ANTIJAM and 802.11 for 4 min, which is equal to $4.8 \cdot 10^6$ time steps in our simulation; (4) the backoff timer of the 802.11 MAC protocol implemented here uses units of $50\mu s$; (5) we omit SIFS, DIFS, and RTS/CTS/ACK since they are not relevant to this implementation.

A comparison is summarized in Figure 8. The throughput achieved by ANTIJAM is significantly higher than the one by the 802.11 MAC protocol, specially for lower values of ε , when the 802.11 MAC protocol basically fails to deliver any successful message.




Figure 8. Throughput as a function of $\varepsilon \in [0.05, 0.95]$, compared to 802.11, averaged over 10 runs, where $\hat{p} = 1/24$.

V. CONCLUSION

This paper presented a simple, fair and self-stabilizing distributed MAC protocol called ANTIJAM that is able to make efficient use of a shared communication medium whose availability is changing quickly and in a hard to predict manner over time. In particular, we proved that our protocol achieves a constant competitive throughput if ε is constant.

It is an open question whether there is a protocol where the throughput is independent of ε , as it exists for adaptive, nonreactive jammers [2]. At least for protocols that are based on cumulative probabilities, it seems that this is impossible, for the following reason: Note that information about the current state allows a $(1 - \varepsilon)$ -bounded reactive adversary to jam

all time steps with transmissions, given that there is at least an ε -fraction of idle time steps. This indicates that a reactive jammer is much stronger than a jammer which is only adaptive to the past: only when the probability of seeing an idle channel drops below ε , the nodes have a chance to successfully transmit messages. This however implies that the cumulative sending probability must be at least $\log 1/\varepsilon$, which means that if the access probabilities of all nodes are equal, the probability of successfully transmitting a message is at most $\varepsilon \cdot \log 1/\varepsilon$. Hence, at best one might be able to achieve a competitiveness that drops linearly with ε while we can only prove a competitiveness that drops exponentially with ε , so further research is needed to close that gap.

ACKNOWLEDGMENTS

This work was supported in part by NSF awards CCF-0830791 and CCF-0830704, and by the DFG project SCHE 1592/1-1.

REFERENCES

- [1] G. Alnifie and R. Simon. A multi-channel defense against jamming attacks in wireless sensor networks. In *Proc. of Q2SWinet '07*, pages 95–104, 2007.
- [2] B. Awerbuch, A. Richa, and C. Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *Proc. of PODC '08*, 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 *Proc. of IEEE Infocom '08*, pages 1265–1273, 2008.
- [4] M. A. Bender, M. Farach-Colton, S. He, B. C. Kuszmaul, and C. E. Leiserson. Adversarial contention resolution for simple channels. In *Proc. of SPAA '05*, pages 325–332, 2005.
- [5] T. Brown, J. James, and A. Sethi. Jamming and sensing of encrypted wireless ad hoc networks. In *Proc. of MobiHoc '06*, pages 120–130, 2006.
- [6] J. Chiang and Y.-C. Hu. Cross-layer jamming detection and mitigation in wireless broadcast networks. In *Proc. of MobiCom '07*, pages 346–349, 2007.
- [7] B. S. Chlebus, D. R. Kowalski, and M. A. Rokicki. Adversarial queuing on the multiple-access channel. In *Proc. of PODC '06*, pages 92–101, 2006.
- [8] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60(2):115 – 143, 2006.
- [9] S. Dolev, S. Gilbert, R. Guerraoui, D. Kowalski, C. Newport, F. Kuhn, and N. Lynch. Reliable distributed computing on unreliable radio channels. In *Proc. 2009 MobiHoc S3 Workshop*, 2009.
- [10] S. Gilbert, R. Guerraoui, D. R. Kowalski, and C. C. Newport. Interference-resilient information exchange. In *Proc. 28th IEEE International Conference on Computer Communications (INFOCOM)*, 2009.
- [11] S. Gilbert, R. Guerraoui, and C. Newport. Of malicious motes and suspicious sensors: On the efficiency of malicious interference in wireless networks. In *Proc. of OPODIS '06*, 2006.
- [12] J. Hastad, T. Leighton, and B. Rogoff. Analysis of backoff protocols for multiple access channels. *SIAM Journal on Computing*, 25(4):740–774, 1996.
- [13] M. Heusse, F. Rousseau, R. Guillier, and A. Duda. Idle sense: An optimal access method for high throughput and fairness in rate diverse wireless lans. In *Proc. Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 121–132, 2005.
- [14] S. Jiang and Y. Xue. Providing survivability against jamming attack via joint dynamic routing and channel assignment. In *Proc. 7th Workshop on Design of Reliable Communication Networks (DRCN)*, 2009.
- [15] C. Koo, V. Bhandari, J. Katz, and N. Vaidya. Reliable broadcast in radio networks: The bounded collision case. In *Proc. of PODC '06*, 2006.
- [16] Y. Law, L. van Hoesel, J. Doumen, P. Hartel, and P. Havinga. Energy-efficient link-layer jamming attacks against wireless sensor network MAC protocols. In *Proc. of SASN '05*, pages 76–88, 2005.
- [17] M. Li, I. Koutsopoulos, and R. Poovendran. Optimal jamming attacks and network defense policies in wireless sensor networks. In *Proc. of Infocom '07*, pages 1307–1315, 2007.
- [18] X. Liu, G. Noubir, R. Sundaram, and S. Tan. Spread: Foiling smart jammers using multi-layer agility. In *Proc. of Infocom '07*, pages 2536–2540, 2007.
- [19] D. Meier, Y. A. Pignolet, S. Schmid, and R. Wattenhofer. Speed dating despite jammers. In *Proc. DCOSS '09*, June 2009.
- [20] V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein. Using channel hopping to increase 802.11 resilience to jamming attacks. In *Proc. of Infocom '07*, 2007.
- [21] R. Negi and A. Perrig. Jamming analysis of MAC protocols. Technical report, Carnegie Mellon University, 2003.
- [22] A. Pelc and D. Peleg. Feasibility and complexity of broadcasting with random transmission failures. In *Proc. of PODC '05*, 2005.
- [23] P. Raghavan and E. Upfal. Stochastic contention resolution with short delays. *SIAM Journal on Computing*, 28(2):709–719, 1999.
- [24] A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Antijam: Efficient medium access despite adaptive and reactive jamming. In *ArXiv Report 1007.4389*, 2010.
- [25] A. Richa, C. Scheideler, S. Schmid, and J. Zhang. A jamming-resistant mac protocol for multi-hop wireless networks. In *Proc. 24th International Symposium on Distributed Computing (DISC)*, 2010.
- [26] A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Self-stabilizing leader election for single-hop wireless networks despite jamming. In *Proc. 12th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2011.
- [27] D. Thuente and M. Acharya. Intelligent jamming in wireless networks with applications to 802.11b and other networks. In *Proc. of MILCOM '06*, 2006.
- [28] A. Wood, J. Stankovic, and G. Zhou. DEEJAM: Defeating energy-efficient jamming in IEEE 802.15.4-based wireless networks. In *Proc. of SECON '07*, 2007.
- [29] W. Xu, W. Trappe, Y. Zhang, and T. Wood. The feasibility of launching and detecting jamming attacks in wireless networks. In *Proc. of MobiHoc '05*, pages 46–57, 2005.
- [30] W. Xu, T. Wood, and Y. Zhang. Channel surfing and spatial retreats: defenses against wireless denial of service. In *Proc. of Workshop on Wireless Security*, 2004.