IEEE/ACM Transactions on Networking

An Efficient and Fair MAC Protocol Robust to Reactive and Worst-Case Interference

Journal:	*IEEE/ACM Transactions on Networking*
Manuscript ID:	TNET-00529-2011.R2
Manuscript Type:	Original Article
Date Submitted by the Author:	n/a
Complete List of Authors:	Richa, Andrea; Arizona State University, SCIDSE Scheideler, Christian; Uni Paderborn, Computer Science Schmid, Stefan; TU Berlin & T-Labs, EECS Zhang, Jin; Arizona State University, SCIDSE
Keywords:	Medium Access, Robustness, Throughput, Algorithms, Anti-Jamming, Adversary
Note: The following files were submitted by the author for peer review, but cannot be converted to PDF. You must view these files (e.g. movies) online.	
reviewer.tgz	

An Efficient and Fair MAC Protocol Robust to Reactive Interference

Andrea Richa[1], Christian Scheideler[2], Stefan Schmid[3], Jin Zhang[1]

[1] Computer Science and Engineering, SCIDSE, Arizona State University
Tempe, Arizona, USA; {aricha,jzhang82}@asu.edu
[2] Department of Computer Science, University of Paderborn, D-33102 Paderborn, Germany; scheideler@upb.de
[3] Telekom Innovation Laboratories (T-Labs) & TU Berlin, D-10587 Berlin, Germany; stefan@net.t-labs.tu-berlin.de

***Abstract*—Interference constitutes a major challenge to availability for communication networks operating over a shared medium. This article proposes the medium access (MAC) protocol ANTIJAM which achieves a high and fair throughput even in harsh environments. Our protocol mitigates *internal* interference, requiring no knowledge about the number of participants in the network. It is also robust to intentional and unintentional *external* interference, e.g., due to coexisting networks or jammers. We model external interference using a powerful *reactive adversary* who can jam a $(1-\varepsilon)$-portion of the time steps, where $0 < \varepsilon \leq 1$ 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. ANTIJAM makes efficient use of the non-jammed time periods and achieves, if ε is constant, a $\Theta(1)$-competitive throughput. In addition, ANTIJAM features a low convergence time and has excellent fairness properties, such that channel access probabilities do not differ among nodes by more than a small constant factor.**

I. INTRODUCTION

Disruptions of communications over a shared medium due to interference—intentional, or not—are a central challenge to wireless computing. It is well-known that simple jamming attacks—requiring no 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 industry and academia. Accordingly, much progress has been made over the last years on how to deal with different jammer secnarios.

This article presents a very robust medium access (MAC) protocol, ANTIJAM that makes effective use of the few and arbitrarily distributed time periods whenever the wireless medium is available. We model the external interference (due to co-located networks, jamming, etc.) as an *adversary*, simply called the *jammer*.[1]

In contrast to related protocols which are robust to simple oblivious adversaries, this article makes an important step forward and studies MAC protocols against "smart" adversaries, i.e., robust to more complex forms of interference. In particular, our adversary may behave in an adaptive and reactive manner: *adaptive* in the sense that the 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 also part of the 802.11 standard) to learn whether the channel is currently idle or not, and jam the medium depending on these measurements.

Note, the study of reactive jamming is relevant beyond purely adversarial contexts. Interactions between spatially co-located networks may appear as a reactive "jamming" when "interference" happens during idle time periods: the nodes in co-located networks are likely to transmit if there are no ongoing transmissions in the other networks. Furthermore, transmissions from co-located networks during idle time periods constitute a non-trivial challenge to many MAC protocols whose operation and states depend on the idle time periods.

A. Related Work

Many *MAC layer strategies* have been devised to resolve unintentional and malicious interference in the literature, including coding strategies (e.g., [5]), channel surfing and spatial retreat (e.g., [1], [21]), or mechanisms to hide messages from a jammer, evade its search, and reduce the impact of corrupted messages (e.g., [20]). These methods however, do not help against an adaptive jammer with *full* information about the history of the protocol, which we consider in our work. An interesting related approach is used in *IdleSense* [12], a variation of the 802.11 Distributed Coordination Function: in IdleSense, all nodes use similar values of the contention window to achieve good short-term access fairness. The synchronization is achieved via monitoring the number of idle slots between transmission attempts to dynamically control the contention. Although simulations indicate a high throughput, no formal bounds are derived.

In the theory community, work on MAC protocols has mostly focused on efficiency. Many of these protocols are *random backoff protocols* (e.g., [4], [6], [7], [11], [16]) that do not take jamming activity into account and are not robust against it (see [2] for more details). Some theoretical work on *jamming* is known, however (e.g., [8] for a short overview) and there are two basic approaches in the literature. The first assumes randomly corrupted messages (e.g. [15]), 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

[1] Note, however, that our notion of adversary is intended (and limited to) describe arbitrary and worst-case patterns of external interference. In particular, our adversary does not try to capture any kind of malicious or Byzantine behavior, and cannot perform, e.g., insider attacks with additional information. The study of such malicious adversaries is complementary work, beyond the scope of this article.

budget (e.g. [10], [13]), or bounds the number of channels the adversary can jam (e.g. [9], [14]). The protocols in, e.g., [13] can tackle adversarial jamming at both the MAC and network layers, where the adversary may not only jam the channel but also introduce malicious (fake) messages (possibly with address spoofing). However, these solutions 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 Theorem 1 of [10] upper bounds the adversary's capability of disrupting communications with a budget of β messages, and then shows that the proposed protocol needs at least 2β rounds to terminate, which implies a jamming rate below 0.5. The handshaking mechanism in [13] requires an even lower jamming rate, because the faulty nodes can cause n_c collisions and n_s rounds of address spoofing, while the number of rounds needed to have a message successfully broadcasted is at least $2n_c + n_s$.)

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 article, 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-located 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] is 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.

The problem becomes significantly more challenging than the nonreactive version, due to the large number of possible strategies a jammer can pursue. First, the analysis is more involved as the nodes' aggregate sending probability varies in a larger range depending on the adversarial strategy. Technically, the reactive jamming renders it impossible to apply Chernoff bounds over the non-jammed time periods as their patterns are no longer random; rather, we have to argue over *all* time periods. Second, modifications to the protocol in [2] are needed. For instance, the ANTIJAM protocol seeks to synchronize the nodes' sending probabilities; this has the desirable side effect of achieving fairness: all nodes are basically granted the same channel access probabilities, which greatly improves 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* [17]) and to devise robust applications such as leader election protocols [18].

B. Model

We study a wireless network that consists of n cooperative and reliable simple wireless devices (e.g., sensor nodes) that are within the transmission range of each other and which 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).

n	number of nodes
T	time window of adversary
N	$N = \max\{T, n\}$
ε	adversary leaves εT time steps non-jammed
γ	common parameter to adapt nodes' access probabilities
p_v	node v's access probability
c_v	counter variable used to keep track of time steps
T_v	node v's estimation of T
$\hat{p}$	maximum individual node access probability
p	aggregate probabilities of the network
$p_t(v)$	node v's probability at time step t
p_t	aggregate probabilities at time step t
I'	subframe used to analyze the protocol
f	size of I'
I	a time frame consisting of a polylogarithmic number of I'
F	size of I
k	number of useful time steps in I'
k_0	number of idle time steps in I'
k_1	number of time steps in I' with a successful transmission
k_1'	successful transmission with different sender
k_2	number of times aggregate probability decreased
k_3	number of times pass started at initial step
g	number of non-jammed time steps

TABLE I
IMPORTANT VARIABLES

In addition to these nodes there is arbitrary external interference which we model as an adversary. (Note that our notion of adversary is mainly meant as a model to describe external interference only. The adversary does not, e.g., read and modify packet contents.)

We allow the adversary to know the protocol and its entire history (in terms of idle, busy, and successful transmission events) 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.

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 non-idle (the channel is non-idle either because of a successful transmission, or because the channel is busy) and can instantly make a jamming decision based on that information. Those jammers arise in scenarios where, for example, encryption is 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-located networks, as many MAC protocols based on carrier sensing activate nodes during idle time periods.

We assume that the adversary is only allowed to jam a $(1-\varepsilon)$-fraction of the time steps, for an arbitrary constant $0 < \varepsilon \leq 1$. In addition, we allow the adversary to perform *bursty* jamming. Formally, an adversary is called $(T, 1-\varepsilon)$*-bounded* for some $T \in \mathbb{N}$ and $0 < \varepsilon \leq 1$ 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.[2]

This article studies *competitive* MAC protocols.

Definition I.1 (c-Competitive)**.** *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.*

In other words, in a c-competitive protocol, there is a successful transmission in the network every c-th non-jammed round on average.

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 $\Theta(1)$-competitive against any $(T, 1-\varepsilon)$-bounded reactive 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 $\gamma = O(1/(\log T + \log\log n))$. As $\log T$ and $\log\log n$ are small for all reasonable values of T and n, this is scalable and not a critical constraint, as it leaves room for a super-polynomial change in n and a polynomial change in T over time.[3] Thus, all we need for our formal performance result to hold is a very rough upper bound on γ. As we will see in our theorems there is a tradeoff between too low γ values (which causes the protocol to react too slowly to changes) and too high γ values (with which the aggregate probability may overshoot). In practice we expect that choosing a constant, sufficiently small γ yields a good performance for any practical network. Indeed, in our simulations $\gamma = 0.1$ results in a good throughput for a wide range of networks.

C. Our Contributions

This article presents a very simple and robust medium access protocol called ANTIJAM, together with a rigorous performance analysis. We are not aware of any other protocol with provable competitive throughput guarantees in a similarly harsh environment.

Concretely ANTIJAM is provably robust against adaptive and reactive external interference, or equivalently, an adversary (an outsider) that can jam the medium a constant fraction of the time. Despite this harsh environment, we can show that the ANTIJAM MAC protocol achieves a high throughput by exploiting any non-jammed time intervals effectively. The main theoretical contribution is a formal and rigorous derivation of the good throughput and fairness guarantees of our protocol. We show that ANTIJAM is competitive in the sense that a constant fraction of the non-jammed 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.

Theorem I.2. *Let $N = \max\{T, n\}$. The* ANTIJAM *protocol is constant-competitive, namely $e^{-\Theta(1/\varepsilon^2)}$-competitive w.h.p.*[4] *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 a 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.*

Our theoretical results are complemented by extensive simulations.

II. THE ANTIJAM MAC PROTOCOL

The basic idea of the ANTIJAM MAC protocol is simple. In an ANTIJAM network, each node v maintains a medium access value p_v which describes the probability that v transmits a

[2] In an adversarial context, this model can be motivated e.g., in sensor networks. Sensor network consist of simple wireless nodes usually running on a single frequency and which cannot benefit from more advanced anti-jamming techniques such as 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.)

[3] On the other hand, note that the assumption that the nodes know constant factor approximations of n or T directly renders the problem simple: if the set of n nodes is static, nodes can simply access the medium with probability $1/n$. This yields a high and fair throughput: If T is known, a time period of length T without idle and successful periods implies that the aggregate probability is too high. This information can be exploited by the algorithm. However, such assumptions are unrealistic and do not scale.

[4] *With high probability*, or short *w.h.p.*, denotes a probability of a least $1 - 1/n^c$ for some constant c.

message in a communication round. The nodes adapt and synchronize their p_v values over time (which as a side-effect also improves fairness) in a multiplicative-increase multiplicative-decrease manner in order to ensure maximum throughput. The p_v values tend to be lowered in times of high interference, and increased during times where the channel is idling. (This is similar to classic random backoff mechanisms where the next transmission time t is chosen uniformly at random from an interval of size $1/p_v$.) More precisely, the sending probabilities are changed by a factor of $(1+\gamma)$. However, we impose an upper bound of $\hat{p}$ on p_v, for some constant $0 < \hat{p} < 1$.[5] As we will see, unlike in most classic backoff protocols, our adaption rules for p_v ensure that the adversary cannot influence the p_v values much by jamming.

In addition, each node maintains two variables, a threshold variable T_v and a counter variable c_v. The threshold variable T_v is used to estimate the adversary's time window T: a good estimation of T can help the nodes recover from a situation where they experience high interference in the network. In times of high interference, T_v will be increased and the sending probability p_v will be decreased.

Initially, every node v may set $T_v := 1$, $c_v := 1$ and $p_v := \hat{p}$. However, as we will see, ANTIJAM converges quickly and hence works for arbitrary initial variable values. Afterwards, the protocol works in synchronized time steps. We assume synchronized time steps for the analysis, but a non-synchronized execution of the protocol would also work as long as all nodes operate at roughly the same speed.

ANTIJAM is based on the following ideas and concepts. Suppose that each node v decides to send a message at the current time step with probability p_v with $p_v \le \hat{p}$. 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 originally appeared in [2]. It states that if $q_0 = \Theta(q_1)$, then the *aggregate sending probability* p, i.e., the sum of all the nodes' individual sending probabilities p_v (which can be larger than one), is constant. This in turn implies that at any non-jammed 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 game for the probabilities p_v and by synchronizing all the nodes, both in terms of sending probabilities and their own estimates on the time window threshold estimate T_v's, at every successful transmission.

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

With these definitions and insights, we can now formally describe the ANTIJAM protocol, see Algorithm 1.

A summary of all our variables (including the ones from the analysis) is provided in Table I.

Algorithm 1: ANTIJAM: for each node v

```
$roundcounter = 0$
while $true$ do
    $v$ decides with probability $p_v$ to send a message
    if $v$ decides to send a message then
        $v$ sends a message along with a triple: $(p_v, c_v, T_v)$
    else
        $v$ senses the channel
        if $v$ senses an idle channel then
            $p_v := \min\{(1+\gamma)p_v, \hat{p}\}$
            $T_v := \max\{T_v - 1, 1\}$
        else
            if $v$ successfully receives a message along with the triple of $(p_{new}, c_{new}, T_{new})$ then
                $p_v := (1+\gamma)^{-1} p_{new}$
                $c_v := c_{new}$
                $T_v := T_{new}$
    $c_v := c_v + 1$
    if $c_v > T_v$ then
        if there was no idle step among the past $T_v$ time steps then
            $p_v := (1+\gamma)^{-1} p_v$
            $T_v := T_v + 2$
    $roundcounter := roundcounter + 1$
```

III. ANALYSIS

Our analysis of Theorem I.2 unfolds in a number of lemmas. We show that given a certain sufficiently large initial aggregate probability p_t in a subframe, the aggregate probability cannot be smaller at the end of the subframe (Lemma III.6). 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 (Lemma III.10). Finally, we show that for any jamming strategy, ANTIJAM has an aggregate probability of $p_t \le \delta/\varepsilon^2$ for most of the time (Lemma III.13).

The analysis makes repeated use of the following well-known relation and the Chernoff bounds derived from [19].

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

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

Lemma III.2 (Chernoff Bounds [19])**.** *Consider any set of binary random variables $X_1, \ldots, X_n$. Suppose that there are values $p_1, \ldots, p_n \in [0,1]$ with $\mathbb{E}[\prod_{i \in S} X_i] \le \prod_{i \in S} p_i$ for every set $S \subseteq \{1, \ldots, n\}$. Then it holds for $X = \sum_{i=1}^n X_i$ and $\mu = \sum_{i=1}^n p_i$ and any $\delta > 0$ that*

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

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

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

[5] While our formal result is valid for any choice of constant $\hat{p}$, $\hat{p}$ should not be chosen too low in small networks. See also our discussion in the simulation section.

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.

We start with some simple facts which also provide some intuition for ANTIJAM. Fact III.3 states that the protocol synchronizes the sending probabilities of the nodes (up to a factor of $(1+\gamma)$) as well as the values c_v and T_v.

Fact III.3. *Right after a successful transmission of the triple (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 following facts study how the aggregate sending probability varies over time depending on the different events.

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:*

***1.** 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 $(1+\gamma)$.

We can now prove the following crucial lemma lower bounding the aggregate sending probability.

Lemma III.6. *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.*

Proof: We start with the following claim about the maximum number of times the nodes decrease their probabilities in I' due to $c_v > T_v$.

Claim III.7. *If in subframe I' the number of idle time steps is at most k_0, then every node v increases T_v by 2 at most $k_0/2 + \sqrt{f}$ many times.*

Proof: Only idle time steps reduce T_v. If there is no idle time step during the last T_v many steps, T_v is increased by 2. Suppose that $k_0 = 0$. Then the number of times a node v increases T_v by 2 is upper bounded by the largest possible ℓ so that $\sum_{i=0}^{\ell} T_v^0 + 2i \leq f$, where T_v^0 is the initial size of T_v. For any $T_v^0 \geq 1$, $\ell \leq \sqrt{f}$, so the claim is true for $k_0 = 0$. At best, each additional idle time step allows us to reduce all thresholds for v by 1, so we are searching for the maximum ℓ so that $\sum_{i=0}^{\ell} \max\{T_v^0 + 2i - k_0, 1\} \leq f$. This ℓ is upper bounded by $k_0/2 + \sqrt{f}$, which proves our claim. ■

This allows us to prove that p exceeds a certain minimal threshold in a subframe.

Claim III.8. *Suppose that for the first time step t_0 in I', $p_{t_0} \in [1/(f^2(1+\gamma)^{\sqrt{2f}}), 1/f^2]$. Then there is a time step t in I' with $p_t \geq 1/f^2$, w.h.p.*

Proof: Suppose that there are g non-jammed time steps in I'. Let k_0 be the number of these steps with an idle channel and k_1 be the number of these steps with a successful message transmission. Furthermore, let k_2 be the maximum number of times a node v increases T_v by 2 in I'. If all time steps t in I' satisfy $p_t < 1/f^2$, then it must hold that $k_0 - \log_{1+\gamma}(1/p_{t_0}) \leq k_1 + k_2$. This is because no v has reached a point with $p_t(v) = \hat{p}$ in this case, so Fact III.5 implies that for each time step t with an idle channel, $p_{t+1} = (1+\gamma)p_t$. Thus, at most $\log_{1+\gamma}(1/p_{t_0})$ time steps with an idle channel would be needed to get p_t to $1/f^2$, and then there would have to be a balance between further increases (that are guaranteed to be caused by an idle channel) and decreases (that might be caused by a successful transmission or the case $c_v > T_v$) of p_t in order to avoid the case $p_t \geq 1/f^2$. The number of times we can allow an idle channel is maximized if all successful transmissions and cases where $c_v > T_v$ cause a reduction of p_t. So we need $k_0 - \log_{1+\gamma}(1/p_{t_0}) \leq k_1 + k_2$ to hold to avoid the case $p_t \geq 1/f^2$ somewhere in I'.

We know from Claim III.7 that $k_2 \leq k_0/2 + \sqrt{f}$. Hence,

$$\begin{aligned} k_0 &\leq 2\log_{1+\gamma} f + \sqrt{f} + k_1 + k_0/2 + \sqrt{f} \\ \Rightarrow \quad k_0 &\leq 4\log_{1+\gamma} f + 2k_1 + 4\sqrt{f} \end{aligned}$$

Suppose that $4\log_{1+\gamma} f + 4\sqrt{f} \leq \varepsilon f/4$, which is true if $f = \Omega(1/\varepsilon^2)$ is sufficiently large (which is true for $\varepsilon = \Omega(1/\log^3 N)$). Since $g \geq \varepsilon f$ due to our adversarial model, it follows that we must satisfy $k_0 \leq 2k_1 + g/4$.

Certainly, for any time step t with $p_t \leq 1/f^2$,

$$\mathbb{P}[\geq 1 \text{ message transmitted at } t] \leq 1/f^2.$$

Suppose for the moment that no time step is jammed in I'. Then $\mathbb{E}[k_1] \leq (1/f^2)f = 1/f$. In order to prove a bound on

k_1 that holds w.h.p., we can use the general Chernoff bounds stated above. For any step t, let the binary random variable X_t be 1 if and only if at least one message is transmitted at time t and $p_t \leq 1/f^2$. Then

$$\begin{aligned} \mathbb{P}[X_t = 1] &= \mathbb{P}[p_t \leq 1/f^2] \cdot \mathbb{P}[\geq 1 \text{ msg sent} \mid p_t \leq 1/f^2] \\ &\leq 1/f^2. \end{aligned}$$

and it particularly holds that for any set S of time steps prior to some time step t that, if there are multiple message transmissions and since $p_t \leq 1/f^2$,

$$\mathbb{P}[X_t = 1 \mid \prod_{s \in S} X_s = 1] \leq 1/f^2.$$

Then, we have

$$\begin{aligned} \mathbb{P}[\prod_{s \in S} X_s = 1] &= \mathbb{P}[X_1 = 1] \cdot \mathbb{P}[X_2 = 1 | X_1 = 1] \\ &\cdot \quad \mathbb{P}[X_3 = 1 | \prod_{s=1,2} X_s = 1] \\ &\cdot \ldots \cdot \\ &\cdot \quad \mathbb{P}[X_{|S|} = 1 | \prod_{s=1,2,\ldots,|S|-1} X_s = 1] \\ &\leq (1/f^2)^{|S|} \end{aligned}$$

and

$$\mathbb{E}[\prod_{s \in S} X_s = 1] = \mathbb{P}[\prod_{s \in S} X_s = 1] \leq (1/f^2)^{|S|}.$$

Thus, the Chernoff bounds and our choice of f imply that either $\sum_{t \in I'} X_t < \varepsilon f/4$ and $p_t \leq 1/f^2$ throughout I' w.h.p., or there must be a time step t in I' with $p_t > 1/f^2$ which would finish the proof. Therefore, unless $p_t > 1/f^2$ at some point in I', $k_1 < \varepsilon f/4$ and $k_0 > (1 - \varepsilon/4)f$ w.h.p. As the reactive adversary can now reduce k_0 by at most $f - g$ when leaving g non-jammed steps, it follows that for any adversary, $k_0 > (1 - \varepsilon/4)f - (f - g) = g - (\varepsilon/4)f$. That, however, would violate our condition above that $k_0 \leq 2k_1 + g/4$ as that can only hold given the bounds on g and k_1 if $k_0 \leq g - (\varepsilon/4)f$.

Note that the choice of g is not oblivious as the adversary may *adaptively* decide to set g based on the history of events. Thus, we cannot assume that g is a fixed value, and the worst adaptive adversarial path is hard to assess. Therefore, we apply a union bound argument and sum up over all adversarial choices for g, showing that our claim holds for all g simultaneously. In order to show that none of them succeeds, observe that there are only f many possible values for g, and for each the claimed property holds w.h.p. (for all possible distributions of the g events); therefore, the claim holds simultaneously for the polynomially many options of g as well. ■

Similarly, we can also prove that once the aggregate probability exceeds a certain threshold, it cannot become too small again.

Claim III.9. *Suppose that for the first time step t_0 in I', $p_{t_0} \geq 1/f^2$. Then there is no time step t in I' with $p_t < \frac{1}{f^2(1+\gamma)^{\sqrt{2f}}}$, w.h.p.*

Proof: Consider some fixed time step t in I' and let $I'' = (t_0, t]$. Suppose that there are g non-jammed time steps in I''. If $g \leq \beta \log N$ for a (sufficiently large) constant β, then it follows for the probability p_t at the end of I'' due to Claim III.7 that

$$p_t \geq \frac{1}{f^2} \cdot (1+\gamma)^{-(2\beta \log N + \sqrt{f})} \geq \frac{1}{f^2 (1+\gamma)^{\sqrt{2f}}}$$

given that $\varepsilon = \Omega(1/\log^3 N)$, because in order to compute a pessimistic lower bound on p_t, assume that all g non-jammed steps are successful so at most $\beta \log N$ decreases of p_t can happen, or similarly, assume that all g non-jammed steps are idle, so at most $\beta \log N/2 + \sqrt{f}$ decreases of p_t can happen due to exceeding T_v; the total number of decreases is smaller than $\beta \log N + \beta \log N/2 + \sqrt{f} < 2\beta \log N + \sqrt{f}$.

So suppose that $g > \beta \log N$. Let k_0 be the number of these steps with an idle channel and k_1 be the number of these steps with a successful message transmission. Furthermore, let k_2 be the maximum number of times a node v increases T_v in I''. If $p_t < \frac{1}{f^2(1+\gamma)^{\sqrt{2f}}}$ then it must hold (deterministically) that $k_0 \leq k_1 + k_2$ because of our assumption that $p_{t_0} \geq 1/f^2$ (more idle rounds would yield higher p_t values).

Since $k_2 \leq k_0/2 + \sqrt{f}$, this implies that $k_0 \leq 2k_1 + 2\sqrt{f} \leq 2k_1 + g/4$. Thus, we are back to the case in the proof of Claim III.8, which shows that $k_0 \leq 2k_1 + g/4$ does not hold w.h.p., given that $g > \beta \log N$ and we never have the case in I'' that $p_t > 1/f^2$.

If there is a step t' in I'' with $p_{t'} > 1/f^2$, we prune I'' to the interval $(t', t]$ and repeat the case distinction above. As there are at most f time steps in I'', the claim follows. ■

Combining Claims III.8 and III.9 completes the proof of Lemma III.6. ■

Lemma III.10 establishes an important relationship between aggregate sending probability and throughput.

Lemma III.10. *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 non-jammed time steps in I', then* ANTIJAM *is at least $\frac{\delta}{8\varepsilon^2} e^{-\delta/(1-\hat{p})\varepsilon^2}$-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.11. *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 Corollary III.4, if $p_t \leq \delta/\varepsilon^2$ and $n \geq (1+\gamma)\delta/(\varepsilon^2\hat{p})$, then $p_v(t) \leq \hat{p}/(1+\gamma)$. 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 an aggregate probability of at least δ/ε^2. On the other hand, each of the k_1' successful transmissions reduces the aggregate probability by a factor of $(1+\gamma)$. Therefore, once the aggregate 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 aggregate 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.7 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,

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

Thus, we have a lower bound for k_1 that depends on k, and it remains to find a lower bound for k.

Claim III.12. *Let g be the number of non-jammed time steps t in I' with $p_t \leq \delta/\varepsilon^2$. If $g \geq \varepsilon f/2$ then*

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

w.h.p.

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

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

where p is the aggregate probability at the step. Since we only need to lower bound the number of useful time steps k, and $p_t \leq \delta/\varepsilon^2$, it follows that the probability of a non-idle time step to be useful (note that we are considering non-jammed time steps here) is at least

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

Thus,

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

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

Since our lower bound for the probability of a non-idle step to be useful holds independently for all non-jammed non-idle 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.12 it follows that

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

w.h.p., which completes the proof of Lemma III.10: if we divide the lower bound on k_1 by the number of non-jammed time steps εf (as $g \geq \varepsilon f/2$, $k_1 \geq k/3 - 2\sqrt{f}$ and as $-2\sqrt{f}$ is negligible). ■

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

Lemma III.13. *Suppose that at the beginning of I', $T_v \leq \sqrt{F}/2$ for every node v. Then at most half of the non-jammed time steps t can have the property that $p_t > \delta/\varepsilon^2$ w.h.p.*

Proof: Recall from Fact III.5 that as long as the access probabilities of the nodes do not hit $\hat{p}$, the aggregate probability only changes by a $(1+\gamma)$-factor in both directions. Suppose that δ is selected so that δ/ε^2 represents one of these values. Let H be the set of time steps $t \in I'$ with the property that either $p_t = \delta/\varepsilon^2$ and the channel is idle or $p_t \geq (1+\gamma)\delta/\varepsilon^2$. Now, we define a step t to be *useful* if $t \in H$ and there is either an idle channel or a successful transmission at t. Let k be the number of useful time steps in H. Furthermore, let k_0 be the number of useful time steps with an idle channel, k_1 be the number of useful time steps with a successful transmission and k_2 be the maximum number of times a node v reduces p_v in H because of $c_v > T_v$. It holds that $k = k_0 + k_1$.

Let us cut the time steps in H into *passes* where each pass (t, p, S) starting at time t consists of a sequence of all (not necessarily consecutive) non-idle time steps $t' > t$ with $p_{t'} = (1+\gamma)p$ following t until a time step t'' is reached in which $p_{t''} = p$, or the end of I' is reached if there is no such step, where t'' is either due to $c_v > T_v$ or a successful transmission. The time step t is such that either $p_t = p$ and there is an idle channel at t, or t is the beginning of I' if there is no such idle channel to mark the beginning of S in I'. (Note that for two different passes (t, p, S) and (t', p', S') and $p \neq p'$, $S \cap S' = \emptyset$.)

Although passes defined like this could be nested, we additionally require that for any pair of passes (t, p, S) and (t', p', S') with $p' = p$ and final time step t'' in S, $(t' \cup S') \cap [t, t''] = \emptyset$, but passes with $p \neq p'$ are allowed to violate this (by one being nested into the other). It is not difficult to see that for any distribution of aggregate probabilities over the time steps of I' one can organize the time steps in H into passes as demanded above. Based on that, the following claim can be easily shown, where $k_1' \leq k_1$ is the number of useful time steps with a successful transmission by a node different from the previously successful node.

Let P be any collection of passes in H, and Δ be the number of distinct possible values of the aggregate probability p in P. We have the following claim.

Claim III.14. *For any collection P of passes, w.h.p., $k_0 \geq k_1 - \Delta - \Theta(1)$ where k_0 and k_1 are the number of idle time steps and the number of successful transmissions in P.*

Proof: We first show that $k_0 \geq k_1' - \Delta$. Recall that k_1' is the number of successful transmissions in which the sender is different from the previously successful sender. Moreover, we define k_2 as the number of times that the aggregate probability decreased due to $c_v > T_v$; we define k_3 as the number of times a pass started at the initial step of I' (i.e., the pass started at a non-idle time step). Clearly, we have $k_2 \geq 0$, and $k_3 \leq \Delta$. Since P is any collection of passes in H, it implies that the aggregate probability $p \geq \delta/\varepsilon^2$ throughout P. Hence, we have the following inequality:

$$k_0 + k_3 \geq k_1' + k_2$$

Together with the fact that $k_2 \geq 0$, and $k_3 \leq \Delta$, we have

$$k_0 \geq k_1' - \Delta$$

Then, let $E_i = 1$ denote the event that the sender of the i-th successful transmission is the same as the sender of the previous successful transmission. We show that the probability that $\sum_i E_i \geq c$ (c is a constant) given k_1 is extremely small. According to Corollary III.4, the nodes' access probabilities do not differ by more than a $(1+\gamma)$-factor after the first successful transmission. Hence, each node has almost the same probability of transmitting a message at any given time step, which implies that $\mathbb{P}[E_i = 1] \leq (1+\gamma)/n$.

$$\mathbb{P}[\sum_i E_i \geq c \mid k_1] \leq \binom{k_1}{c} \cdot (\frac{1+\gamma}{n})^c \leq \binom{f}{c} \cdot (\frac{1+\gamma}{n})^c$$

Since f is polynomially smaller than n, $\mathbb{P}[\sum_i E_i \geq c \mid k_1]$ becomes very small even for small c, which implies that $E_i = 1$ happens at most a constant number of times during P w.h.p. Hence, the claim holds. ■

We have the following upper bound on the number of such steps in H.

Claim III.15.

$$|H| \leq (k + \log_{1+\gamma} \max\{p_0/(\delta/\varepsilon^2), 1\})\sqrt{F}$$

where k is the number of useful steps in H.

Proof: If at the beginning of I', $T_v \leq \sqrt{F}/2$ for every node v, then according to Claim III.7, $T_v \leq \sqrt{F}$ for every node v at any time during I'. Hence, after at most $2\sqrt{F}$ nonuseful steps we run into the situation that $c_v > T_v$ for every node v, which reduces the aggregate probability by a factor of $(1+\gamma)$. Given that we only have k useful steps and we may initially start with a probability $p_0 > \delta/\varepsilon^2$, there can be at most $(k + \log_{1+\gamma} \max\{p_0/(\delta/\varepsilon^2), 1\})\sqrt{F}$ time steps in H; k are the useful ones, and the nonuseful ones are the non-idle and non-successful steps in which the aggregate probability is reduced: every $\sqrt{F}$ nonuseful steps give one reduction of p). This proves the claim. ■

For the calculations below recall the definition of f with the constants α and β that are assumed to be sufficiently large. If $k \leq \alpha \log N$, then it follows from Claim III.15 that, for large enough δ,

$$|H| \leq (\alpha \log N + \log_{1+\gamma} N)\sqrt{F} \leq \varepsilon f/\beta$$

where $N = \max\{n, T\}$. Thus, the number of non-jammed time steps in H is also at most $\varepsilon f/\beta$, and since β can be arbitrarily large, Lemma III.13 follows, as the steps in H fulfill this property ($\beta \geq 2$ yields half of the steps).

It remains to consider the case that $k > \alpha \log N$. Let us assume that H contains at least $\varepsilon f/2$ non-jammed time steps, otherwise the claim certainly holds. Our goal is to contradict that statement in order to show that the lemma is true. For this we will show that Claim III.14 is violated w.h.p.

Let T_p be the number of all time steps covered by passes (t', p', S') with $p' = p$. Certainly, $\sum_{p \geq \delta/\varepsilon^2} T_p = |H|$. Let $\phi = \delta/\varepsilon^2$, and $\Phi = (1-\hat{p})\ln(f/\log N)$.

For an aggregate probability $p \geq \Phi$, $\mathbb{P}[\text{idle} \mid p] \leq e^{-\Phi} = (\frac{\log N}{f})^{1-\hat{p}}$ and $\mathbb{P}[\text{success} \mid p] \leq \frac{\Phi}{1-\hat{p}} \cdot e^{-\Phi} \leq \ln(f/\log N) \cdot (\frac{\log N}{f})^{1-\hat{p}}$. Hence, by multiplying these probabilities by the $|H| \leq f$ steps, we get that $k \leq f^{\hat{p}} \cdot \ln f \cdot \log^{1-\hat{p}} N$ on expectation, and from the Chernoff bounds it follows that $k \leq 2f^{\hat{p}} \cdot \ln f \cdot \log^{1-\hat{p}} N$ w.h.p., so Claim III.15 implies that the number of time steps in I' with aggregate probability $p \geq \Phi$ is at most

$$(2f^{\hat{p}} \cdot \ln f \cdot \log^{1-\hat{p}} N + \log_{1+\gamma} N)\sqrt{F} \leq \varepsilon f/\beta, \text{w.h.p.}$$

Since β can be arbitrarily large, we can only focus on the time steps when $\phi \leq p < \Phi$.

Let $\bar{J}_p$ be the number of non-jammed time steps in T_p. We consider the case where $\bar{J}_p < \frac{2}{\mathbb{P}[\text{idle}|p]}$. Let $k_{1,p}$ be the number of successful time steps associated with p-passes (i.e., at aggregate probability $(1+\gamma)p$). Then, $\mathbb{E}[k_{1,p}] = \mathbb{P}[\text{success} \mid p] \cdot \bar{J}_p < 2$. If we sum up over all possible probabilities p with $\phi \leq p < \Phi$, the number of non-jammed time steps covered by all $\bar{J}_p$ such that $\bar{J}_p < \frac{2}{\mathbb{P}[\text{idle}|p]}$ is at most

$$\sum_{i=0}^{\log_{1+\gamma}\Phi} 2/e^{-(1+\gamma)^i} \leq 4 \cdot f/\log N = o(f)$$

many time steps, since the p values always differ by factors $(1+\gamma)$ (recall that $e^{-(1+\gamma)^i}$ is the corresponding probability of an idle step).

Hence, we can ignore all the passes where $\bar{J}_p < \frac{2}{\mathbb{P}[\text{idle}|p]}$. We denote the time steps that are ignored by H'. Since we assumed $|H| \geq \varepsilon f/2$, we have that $f \geq |H \setminus H'| \geq \frac{\varepsilon f}{2\eta} = \Theta(f)$, where η is a constant. Let N_p be the number of time steps in $H \setminus H'$ with aggregate probability p. Let X_t be a random variable, where $X_t = 1$ iff there is a successful transmission at time step t. This implies that $k_1 = \sum_{t \in H \setminus H'} X_t$, then:

$$\begin{aligned}\mathbb{E}[k_1] &= \sum_{p=\delta/\varepsilon^2}^{\Phi} N_p \cdot \mathbb{P}[\text{success} \mid p] \\ &\geq \sum_{p=\delta/\varepsilon^2}^{\Phi} N_p \cdot p \cdot e^{-\frac{p}{1-\hat{p}}} \geq \frac{\varepsilon f}{2\eta} \cdot \Phi \cdot e^{-\frac{\Phi}{1-\hat{p}}} \\ &= (1-\hat{p}) \cdot \frac{\varepsilon f}{2\eta} \cdot (\ln f - \ln \log N) \cdot \frac{\log N}{f} \\ &= \Omega(\log N)\end{aligned}$$

Applying Chernoff bounds, we have w.h.p., $k_1 \geq (1 - c_1)\mathbb{E}[k_1]$ where $0 < c_1 \leq 1$.

Similarly, let Y_t be a random variable, where $Y_t = 1$ iff the channel is idle at t. Then, $k_0 = \sum_{t \in H \setminus H'} Y_t$.

$$\begin{aligned}\mathbb{E}[k_0] &= \sum_{p=\delta/\varepsilon^2}^{\Phi} N_p \cdot \mathbb{P}[\text{idle} \mid p] \geq \sum_{p=\delta/\varepsilon^2}^{\Phi} N_p \cdot e^{-\frac{p}{1-\hat{p}}} \\ &\geq \frac{\varepsilon f}{2\eta} \cdot e^{-\frac{\Phi}{1-\hat{p}}} = \frac{\varepsilon}{2\eta} \cdot \log N = \Omega(\log N)\end{aligned}$$

Applying Chernoff bounds, we have w.h.p., $k_0 \leq c_2 \cdot \mathbb{E}[k_0]$ where $c_2 \geq 0$ is a large enough constant.

It implies that w.h.p.,

$$\begin{aligned}k_1 - k_0 &\geq (1-c_1)\mathbb{E}[k_1] - c_2 \cdot \mathbb{E}[k_0] \\ &\geq \sum_{p=\delta/\varepsilon^2}^{\Phi} N_p((1-c_1) \cdot p \cdot e^{-\frac{p}{1-\hat{p}}} - c_2 \cdot e^{-p}) \\ &\geq \frac{\varepsilon f}{2\eta} \cdot ((1-c_1) \cdot \Phi \cdot \frac{\log N}{f} - c_2 \cdot e^{-\phi}) \\ &\geq \frac{\varepsilon f}{2\eta} \cdot ((1-c_1) \cdot \Phi \cdot \frac{\log N}{f} - c_2 \cdot e^{-\delta/\varepsilon^2}) \\ &= \frac{\varepsilon}{2\eta} \cdot \log N((1-c_1) \cdot \Phi - \frac{c_3}{\log N}) \\ &> \log_{1+\gamma} \Phi \\ &> \Delta + \Omega(1)\end{aligned}$$

Note that $c_3 = c_2 \cdot e^{-\delta/\varepsilon^2}$ is a constant, since both δ and ε are constants. Moreover, the number of different p values in $[\phi, \Phi)$ associated with a pass is at most $\Delta = \log_{1+\gamma} \Phi - \log_{1+\gamma} \phi$. Hence, $\log_{1+\gamma} \Phi > \Delta + \Omega(1)$. This inequality holds w.h.p. when the constant c_1 is small enough, and N is sufficiently large.

This is a contradiction to Claim III.14, and hence completes the proof of Lemma III.13. ■

In order to proceed, we need the following claim.

Claim III.16. *For any collection P of passes it holds that*

$$\mathbb{E}[k'_1] \geq (1 - (1+\gamma)/n)k_1$$

where k_1 and k'_1 are defined w.r.t. P.

Proof: Because of Fact III.5, the probability that a successful transmission is done by a node different from the node of the last successful transmission is equal to

$$1 - \frac{(1+\gamma)p}{(n+\gamma)p} \geq 1 - \frac{1+\gamma}{n}.$$

To see this, observe that among the aggregate probability p, if the last sender u has a share $p_u(t) = x$, all other nodes v have a share $x/(1+\gamma)$, and

$$\frac{p_u(t)}{\sum_{v \in V} p_v(t)} = \frac{x}{(n-1) \cdot \frac{x}{1+\gamma} + x} = \frac{1+\gamma}{n+\gamma}.$$

Hence, $\mathbb{E}[k'_1] \geq (1 - (1+\gamma)/n)k_1$. ■

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.17. *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 non-jammed 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

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

Because of Claims III.7 and III.16 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.17, we show that Lemma III.13 is true for a polynomial number of subframes. Then, Lemma III.13 and Lemma III.17 together imply that Lemma III.10 holds for a polynomial number of subframes. Hence, our main Theorem I.2 follows. 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 and to complement our worst-case bounds. Our focus here is on the qualitative nature of the performance of ANTIJAM, and we did not optimize the parameters to obtain the best constants. We consider three different jamming strategies for a reactive jammer that is $(T, 1-\varepsilon)$-bounded, for different ε values and where $T = 100$: (1) one that jams non-idle steps with probability $(1-\varepsilon)$; (2) one that jams non-idle steps deterministically (as long the jamming budget is not used up); (3) one that jams idle steps deterministically (as long as the jamming budget is not used up). Intuitively, it seems that jamming non-idle steps is more harmful than jamming idle steps. However, note that jamming idle steps may be an effective strategy to steer the protocol into

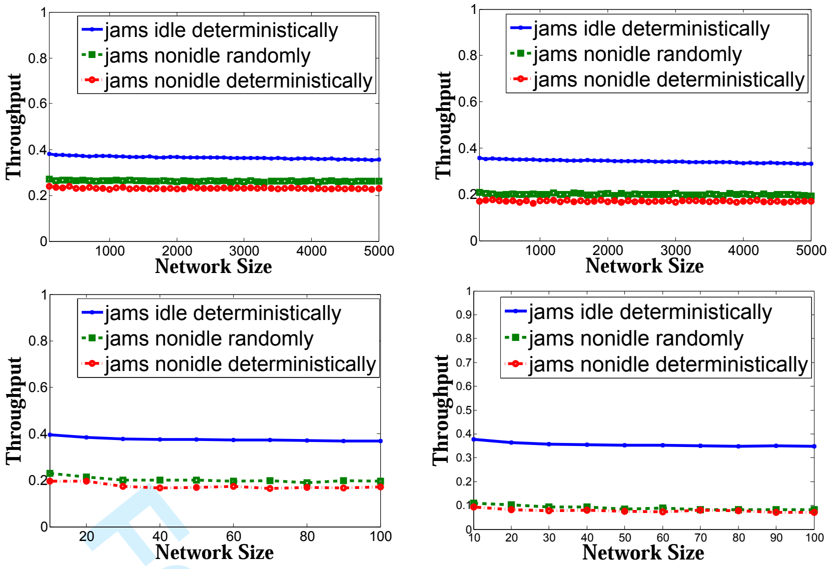

Fig. 1. *Top:* Throughput under three different jamming strategies as a function of the network size (large) and ε, where $\hat{p} = 1/24$ (averaged over 10 runs) (*top left:* $\varepsilon = 0.5$, *top right:* $\varepsilon = 0.3$). *Bottom:* Throughput under three different jamming strategies as a function of the network size (small) and of ε, where $\hat{p} = 1/2$ (averaged over 10 runs) (*bottom left:* $\varepsilon = 0.5$, *bottom right:* $\varepsilon = 0.3$)

bad states. Moreover, it may capture scenarios where nodes in co-located networks start sending in quiet times.

We define throughput as the number of successful transmissions over the number of non-jammed time steps. For networks larger than 100, we choose $\hat{p} = 1/24$, whereas for smaller networks we choose $\hat{p} = 1/2$. As a general guideline, it is always better to choose larger $\hat{p}$ values, as this avoids capping the throughput in small networks artificially. A smaller $\hat{p}$ can make sense for bootstrapping large networks, but due to the fast convergence times of the protocol (see Section IV-B), this is unproblematic.

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, see Figure 1 (*top*). For all three strategies, the throughput is basically constant, independently of the network size. This is in accordance with our theoretical insight of Theorem I.2. We can see that given our conditions on ε and T, the strategy that jams non-idle 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 non-jammed time periods, yielding $20\% - 40\%$ successful transmissions even without optimizing the protocol parameters. Having shown the protocol scales well for large network size, we also study the throughput results when the network size is *small*, see Figure 1 (*bottom*). We observe that the results for small and large scale networks are comparable, but the throughput in the small scale networks can be slightly lower under an adversary that jams non-idle channels deterministically or with probability $(1 - \varepsilon)$.

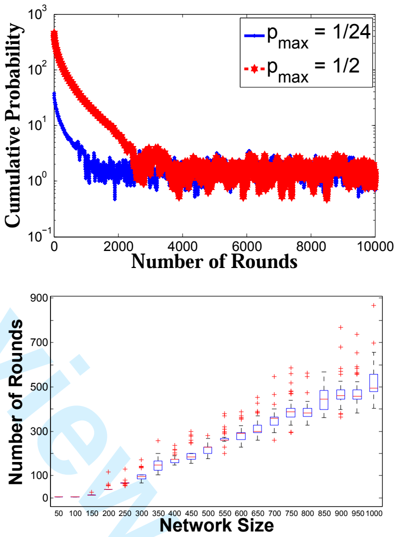

Fig. 3. *Top:* Evolution of aggregate probability over time (network size is 1000 nodes, and $\varepsilon = 0.5$). Note the logarithmic scale. *Bottom:* Boxplot of ANTIJAM runtime as a function of network size for $\hat{p} = 1/24$, and $\varepsilon = 0.5$.

In additional experiments we also studied the throughput as a function of γ, see Figure 2. 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 hence the simulations expensive if one wants to avoid wrong results of the real throughput.)

B. Convergence Time

Besides a high throughput, fast convergence is the most important performance criterion of a MAC protocol. The traces in Figure 3 (*top*) show the evolution of the aggregate probability over time. It can be seen that the protocol converges

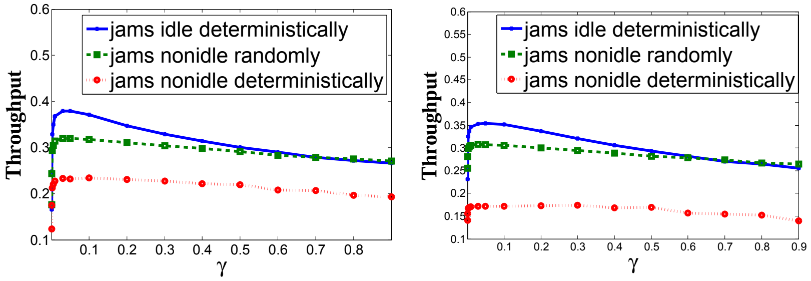

Fig. 2. Throughput as a function of γ under three different jamming strategies, when $n = 1000$, and results are averaged over 10 runs ($\textit{left:}$ $\varepsilon = 0.5$, $\textit{right:}$ $\varepsilon = 0.3$).

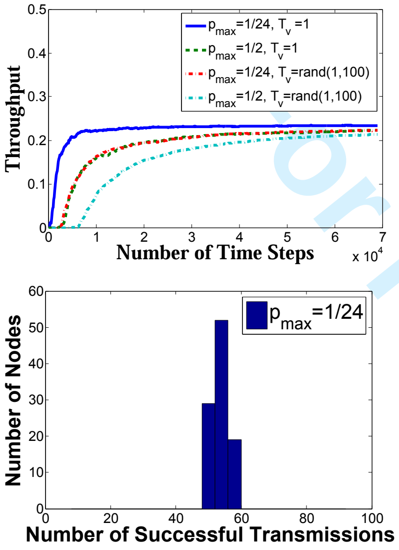

Fig. 4. *Top:* Convergence in a network of 1000 nodes where $\varepsilon = 0.5$. *Bottom:* Fairness in a network of 1000 nodes, where $\varepsilon = 0.5$, and $\hat{p} = 1/24$ (averaged over 10 runs).

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 aggregate probability. Moreover, the fraction of time in which the aggregate probability is in the range of $[\frac{1}{2\varepsilon}, \frac{2}{\varepsilon}]$ 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 aggregate probability is well bounded most of the time, which corresponds to our theoretical insights (cf Lemmas III.6 and III.13). Figure 3 (*bottom*) studies the convergence time for different network sizes. We performed 50 repetitions of each run, and assume that the execution has converged when the aggregate probability p satisfies $p \in [1, 5]$, for at least 5 consecutive rounds. The simulation result qualitatively confirms Theorem I.2, as the number of rounds needed to converge the execution is bounded by $\Theta(\frac{1}{\varepsilon} \log N \max\{T, \frac{1}{\varepsilon\gamma^2} \log^3 N\})$. (Of course, the concrete convergence time can depend on the scenario, and may be faster than expected in the general case.)

Figure 4 (*top*) indicates that independently of the initial values $\hat{p}$ and T_v, the throughput rises quickly (up above 20%) and stays there afterwards.

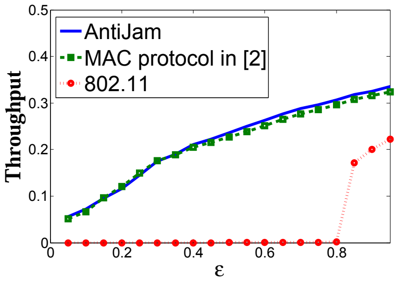

Fig. 5. Throughput as a function of $\varepsilon \in [0.05, 0.95]$, compared to the MAC protocol in [2] and 802.11, averaged over 10 runs, where $\hat{p} = 1/24$.

C. Fairness

As the nodes in the ANTIJAM network synchronize their c_v, T_v, and p_v values upon message reception, they 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.

D. Comparison

Finally, to put ANTIJAM into perspective, as a comparison, we implemented the MAC protocol proposed in [2], as well as 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, the MAC protocol in [2], and 802.11 for 4 min, which is equal to $4.8 \cdot 10M$ 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.

A comparison is summarized in Figure 4 (*bottom*). The throughput achieved by ANTIJAM and the MAC protocol in [2] are 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. Note that the throughput results between ANTIJAM and the MAC protocol in [2] are similar in the simulations, but ANTIJAM is slightly better for the most ε.

V. CONCLUSION

This article presents 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 quickly and unpredictably changes over time. In particular, we prove that our protocol achieves a constant competitive throughput if ε is constant.

We are not aware of any other protocol achieving a competitive throughput in similarly harsh environments. We regard our work as an important step forward towards the design and formal analysis of MAC protocols with provable performance guarantees in more general environments with arbitrary reactive (but *power-constrained*) interference.

ACKNOWLEDGMENTS

We thank Dan Levin from Telekom Innovation Laboratories for proof-reading this article. This work was supported in part by NSF awards CCF-0830791 and CCF-0830704, and by DFG projects SCHE 1592/2-1 and SFB 901.

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] 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.

[6] 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.

[7] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60(2):115 – 143, 2006.

[8] 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.

[9] 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.

[10] 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.

[11] J. Hastad, T. Leighton, and B. Rogoff. Analysis of backoff protocols for mulitiple accesschannels. *SIAM Journal on Computing*, 25(4):740–774, 1996.

[12] 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. *SIGCOMM Comput. Commun. Rev.*, 35(4):121–132, 2005.

[13] C. Koo, V. Bhandari, J. Katz, and N. Vaidya. Reliable broadcast in radio networks: The bounded collision case. In *Proc. of PODC '06*, 2006.

[14] D. Meier, Y. A. Pignolet, S. Schmid, and R. Wattenhofer. Speed dating despite jammers. In *Proc. DCOSS '09*, June 2009.

[15] A. Pelc and D. Peleg. Feasibility and complexity of broadcasting with random transmission failures. In *Proc. of PODC '05*, 2005.

[16] P. Raghavan and E. Upfal. Stochastic contention resolution with short delays. *SIAM Journal on Computing*, 28(2):709–719, 1999.

[17] 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.

[18] 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.

[19] J. P. Schmidt, A. Siegel, and A. Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. In *Proc. Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 331–340, 1993.

[20] 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.

[21] 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.

Prof. Dr. Andrea Richa is an Associate Professor of Computer Science at the School of Computing, Informatics, and Decision Systems Engineering at Arizona State University (ASU), Tempe, AZ. She joined ASU as an Assistant Professor in August 1998. Prof. Richa received her M.S. and Ph.D. degrees from the School of Computer Science at Carnegie Mellon University, in 1995 and 1998, respectively. She also earned an M.S. degree in Computer Systems from the Graduate School in Engineering (COPPE), and a B.S. degree in Computer Science, both at the Federal University of Rio de Janeiro, Brazil, in 1992 and 1990, respectively. Prof. Richa's main area of research is in network algorithms. Some of the topics Prof. Richa has worked on include packet scheduling, distributed load balancing, packet routing, wireless network modeling and topology control, wireless jamming, and distributed hash tables (DHT). Prof. Richa's DHT algorithm has been widely recognized as the first benchmark algorithm for the development of distributed databases in peer-to-peer networking, having being referenced by over 1000 academic journal or conference publications to date, and being implemented as part of two of the first projects in peer-to-peer networking. Prof. Richa was the recipient of an NSF CAREER Award in 1999. For a selected list of her publications, CV, and current research projects, please visit http:// www.public.asu.edu/˜aricha.

Prof. Dr. Christian Scheideler is a Full Professor in the Computer Science Department at the University of Paderborn. Before that, he has been a postdoc at the Weizmann Institute, Israel, for a year, an assistant professor at the Johns Hopkins University, USA, for five years, and an associate professor at the Technical University of Munich for three years. He is (co-)author of more than 100 publications in refereed conferences and journals and has served on more than 50 conference committees. His main focus is on network theory (in particular, peer-to-peer systems, mobile ad-hoc networks and sensors networks), the design and analysis of scalable distributed algorithms and data structures, and the design of algorithms and architectures for robust and secure distributed systems.

Dr. Stefan Schmid is a Senior Research Scientist at Deutsche Telekom Laboratories (T-Labs) and TU Berlin, Germany. He works in the Internet Network Architectures group headed by Prof. Dr. Anja Feldmann where he is developing the CloudNet prototype architecture that connects cloud resources with virtual networks. Before joining T-Labs, he was a postdoc at TU Munich and the University of Paderborn, Germany (with Prof. Dr. Christian Scheideler), and a Ph.D. student in Prof. Dr. Roger Wattenhofer's group at ETH Zurich, Switzerland. Stefan Schmid is interested in distributed systems, and especially in the design of robust and dynamic networks.

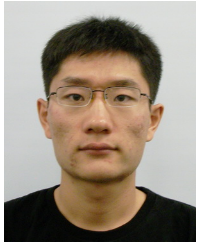

Jin Zhang is working toward his Ph.D. degree in computer science, under the supervision of Prof. Andrea Richa, in the Department of Computer Science and Engineering at Arizona State University, Tempe. Before that, he received his B.E. degree from Beijing University of Posts and Telecommunications, China, in 2008. His research interest is in the area of designing and analyzing efficient medium access protocols that are robust against adversarial jamming in wireless networks.