A Jamming-Resistant MAC Protocol for Single-Hop Wireless Networks

Baruch Awerbuch*
Dept. of Computer Science
Johns Hopkins University
Baltimore, MD 21218, USA
baruch@cs.jhu.edu

Andrea Richa
Dept. of Computer Science
Arizona State University
Tempe, AZ, USA
aricha@asu.edu

Christian Scheideler†
Dept. of Computer Science
Technical University of Munich
85748 Garching, Germany
scheideler@in.tum.de

ABSTRACT

In this paper we consider the problem of designing a medium access control (MAC) protocol for single-hop wireless networks that is provably robust against adaptive adversarial jamming. The wireless network consists of a set of honest and reliable nodes that are within the transmission range of each other. In addition to these nodes there is an adversary. The adversary may know the protocol and its entire history and use this knowledge to jam the wireless channel at will at any time. It is allowed to jam a $(1-\epsilon)$-fraction of the time steps, for an arbitrary constant $\epsilon > 0$, but it has to make a jamming decision before it knows the actions of the nodes at the current step. The nodes cannot distinguish between the adversarial jamming or a collision of two or more messages that are sent at the same time. We demonstrate, for the first time, that there is a local-control MAC protocol requiring only very limited knowledge about the adversary and the network that achieves a constant throughput for the non-jammed time steps under any adversarial strategy above. We also show that our protocol is very energy efficient and that it can be extended to obtain a robust and efficient protocol for leader election and the fair use of the wireless channel.

Categories and Subject Descriptors

C.2.5 [**Computer-Communication Networks**]: Local and Wide-Area Networks—*Access schemes*; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Sequencing and scheduling*

General Terms

Algorithms, Reliability, Theory

Keywords

wireless ad-hoc networks, MAC protocols, jamming

*Supported by NSF CCF 0515080, ANIR-0240551, CCR-0311795, and CNS-0617883

†Supported by DFG grant SCHE 1592/1-1.

PODC'08, August 18–21, 2008, Toronto, Ontario, Canada.

1. INTRODUCTION

Jamming can disrupt wireless transmission and can occur either unintentionally in the form of interference, noise or collision at the receiver side or in the context of an attack. A jamming attack is easy to perform since (i) no special hardware is needed for it to be launched, (ii) it can be implemented by simply listening to the open medium and broadcasting in the same frequency band as the network, and (iii) if launched wisely, it can lead to significant disruptions with small incurred cost for the attacker. Jamming attacks usually aim at the physical layer and are realized by means of a high transmission power signal that corrupts a communication link or an area, but they may also occur at the medium access control (MAC) layer; an adversary may either corrupt control packets or reserve the channel for the maximum allowable number of slots, so that other nodes experience low throughput by not being able to access the channel.

Traditional defenses against jamming focus on the design of physical layer technologies, such as spread spectrum (e.g., [24, 19, 18]). Spread spectrum techniques are useful because if signals are widely spread, it becomes harder for the jammer to detect the start of a packet quickly enough in order to jam it. Unfortunately, protocols such as 802.11b use relatively narrow spreading [11]. The spreading factor for 1Mbps 802.11 is only a factor of 11. Other versions and rates in 802.11 spread signals by equal or smaller factors [4]. Hence, a jammer that can simultaneously block a relatively small number of frequencies would render spread spectrum techniques useless in these scenarios.

Besides defenses at the physical layer, it is also interesting to study defenses at the MAC layer since in contrast to the physical layer, the MAC layer is usually in software and can be changed, so that even wireless devices that do not have a built-in protection against jammers can be made robust against them. However, the 802.11 MAC protocol does not offer much protection here since recent results show that the 802.11 MAC protocol cannot efficiently handle even simple, oblivious jammers [2].

1.1 Our model

In this paper we consider the problem of designing a MAC protocol for single-hop wireless networks that is provably robust against adaptive adversarial jamming at the physical layer. The wireless network consists of a set of n honest and reliable nodes that are within the transmission range of each other. All of the nodes are continuously contending for sending a packet on the wireless channel. We assume that time proceeds in synchronous time steps and in each time step any node may decide to transmit a packet. 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 who is 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 $(1-\epsilon)$-fraction of the time steps, for an arbitrary constant $\epsilon > 0$, and it has to make a jamming decision before it knows the actions of the nodes at the current step.

We allow the adversary to perform bursty jamming. More formally, an adversary is called (T, λ)-*bounded* for some $T \in \mathbb{N}$ and $0 < \lambda < 1$ if for any time window of size $w \geq T$ the adversary can jam at most λw of the time steps in that window. A MAC protocol is called *c-competitive* against some (T, λ)-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 that is constant competitive against any $(T, 1-\epsilon)$-bounded adversary, i.e., there is no central authority controlling the nodes, and the nodes have symmetric roles at any point in time. The nodes do not know ϵ, but we do allow them to have a very rough upper bound of their number n and T. More specifically, we will assume that the nodes have a common parameter $\gamma = O(1/(\log T + \log\log n))$. Such an estimate leaves room for a superpolynomial change in n and a polynomial change in T over time, so it does not make the problem trivial (as would be the case if the nodes knew constant factor approximations of n or T). Next, we formally state our contributions before we go on discussing related work.

1.2 Our contribution

Suppose that $n \geq 2$, i.e., we have at least two honest nodes in the system. Let $N = \max\{T, n\}$. In this paper, we present the first MAC protocol that is constant competitive w.h.p. under any $(T, 1-\epsilon)$-bounded adversary if the protocol is executed for $\Omega(\frac{1}{\epsilon}\log N \max\{T, \frac{1}{\epsilon}(\log^3 N)(\log T + \log\log n)^2\})$ many time steps. It does not need to know ϵ, so ϵ can be an arbitrarily small constant (as long as $\epsilon = \Omega(1/\log^3 N)$). The only information it needs to be constant competitive is that the nodes have a common parameter $\gamma = O(1/(\log T + \log\log n))$. In practice, $\log T$ and $\log\log n$ are reasonably small so that this is not a serious constraint. Also, as mentioned earlier, such an estimate leaves room for a superpolynomial change in n and a polynomial change in T over time. The MAC protocol is very simple and symmetric, and it can recover quickly from any state. We also show that the MAC protocol is very energy efficient. In fact, it converges to a bounded amount of energy consumption under continuous adversarial jamming. In addition to this, we will show how to extend the MAC protocol in order to obtain a robust and efficient protocol for leader election and the fair use of the wireless channel. More specifically, our leader election protocol needs $O(\frac{1}{\epsilon}\log N \max\{T, \frac{1}{\epsilon}(\log^3 N)(\log T + \log\log n)^2\})$ steps until a leader is selected and all nodes are aware of that, and our fair channel use protocol essentially needs $O(n/\epsilon)$ many steps until a fair channel use is guaranteed. All runtime bounds hold with high probability.

1.3 Related Work

Wireless network jamming has been extensively studied in the applied networking domain (e.g., [28, 27, 17, 16, 5, 1, 26, 4, 19, 18, 20, 25]). Mechanisms for launching jamming attacks (e.g., [28, 17, 16, 5]) as well as defense mechanisms against these attacks (e.g., [17, 28, 1, 26, 5, 4, 19, 18]) have been proposed and validated through simulations and experiments.

There are many different forms of jammers, and detecting sophisticated jammers is not easy. Xu et al. [17], for example, observe that simple methods based on signal strength and carrier sensing are unable to conclusively detect the presence of a jammer. Also the packet delivery ratio cannot be used to clearly distinguish between link problems due to mobility, congestion or jamming. Hence, enhanced detection schemes are necessary. To address this need, the authors propose two enhanced detection protocols that employ consistency checking. While being more effective than the prior detection schemes, these protocols still leave room for ambiguities.

Traditional defenses against jamming primarily focus on the design of physical layer technologies, such as spread spectrum [24, 19, 18]. As argued in the introduction, while widely spread frequencies could potentially help in guarding against physical layer jamming, spread spectrum techniques cannot be used effectively in the relatively narrow frequency bands used by the 802.11 standard.

More recent work has also focused on various MAC layer strategies in order to handle jamming, including coding strategies [5], channel surfing and spatial retreat [29, 1], or mechanisms to hide messages from a jammer, evade its search, and reduce the impact of corrupted messages [26]. Most of these strategies have only been evaluated experimentally and would not help against the jammers considered in this paper.

A recent study [2] shows both theoretically and experimentally that an adaptive jammer, such as the one proposed here, can dramatically reduce the throughput of the standard random backoff MAC protocol of the IEEE802.11 standard with only limited energy cost on the adversary side (please also refer to [2] for other references on jamming in 802.11).

Adversarial jamming has also been studied theoretically. There are two basic approaches in the literature. The first assumes that messages may be corrupted at random (e.g. [21]), and the second bounds the number of messages that the adversary can transmit or disrupt due to, for example, a limited energy budget (e.g. [12, 8]). In a single hop wireless network (like ours), messages will not be corrupted independently at random (every time the jammer transmits, all messages in that time step will be corrupted); moreover, an adaptive adversary seems more powerful than one that jams uniformly at random [2]. Hence, we focus on the second line of theoretical work since it is more relevant to the results in this paper.

The latest results in [8, 12] address 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). The results in [8] only consider the scenario that the nodes have one message to transmit (e.g., a broadcast operation). When translated to our continuous data stream scenario, the protocol presented in [8] would not be able to sustain a constant-competitive ratio if the adversary is allowed to jam more than half of the time steps (i.e., if $\epsilon < 1/2$), given the fact that their single message broadcast algorithm takes at least twice as many steps as the number of time steps utilized by the jammer. Moreover, [8] assumes that the nodes have knowledge of n and of the fact that the adversary has a bounded number of messages it can transmit (in contrast, we only need the nodes to have an estimate on $\log\log n$ and $\log T$).

In [12], the authors consider a wireless network in which node positions form a grid where multiple (at most t) adversarial nodes are allowed in the direct neighborhood of a node. If t is at most a suitably small constant, then they give a protocol for reliable broadcast of a single message given that there is a fixed bound on the number of time steps the adversary is disrupting communication (if t is large, no broadcast protocol is guaranteed to terminate). The authors only show that eventually the broadcast operation will be completed, but give no bounds on how long that will take. Moreover, their algorithms will clearly deplete the energy of the non-faulty nodes at a higher rate than that of the faulty nodes.

Most of the theoretical work on the design of efficient MAC protocols has focused on random backoff protocols (e.g., [3, 6, 10, 9, 15, 22]) that do not take jamming activity into account and therefore are not robust against it. MAC protocols have also been designed in the context of broadcasting (e.g., [7]) and clustering (e.g., [14]). Most of them use random backoff or tournaments in order to handle interference and thereby achieve a fast runtime.

In general terms, in a random backoff protocol, each node periodically attempts to transmit a message starting with a certain probability p. In case the message transmission is unsuccessful (due to interference), the node will retry sending the message in the next time steps with monotonically decreasing probabilities (for example, $p^2, p^4, p^8, \ldots$) until the message is successfully transmitted or the minimum allowable probability is reached. In a dense network (as in our single-hop scenario), an adversary with knowledge of the MAC protocol would simply wait until the nodes have reached transmission probabilities that are inversely proportional to the number of close-by nodes to start jamming the channel, forcing the nodes to lower their transmission probabilities by so much that a constant throughput is not achievable. In tournaments, local leader election is used to determine the node that is allowed to use the wireless medium for its message transmission. If the adversary jams the channel whenever a local leader is about to be selected, most protocols will fail and start all over, so that only rarely a message will get through. Also any work that relies on physical carrier sensing in order to adjust the transmission probabilities of the nodes (e.g., [13]) would fail in the presence of jamming as a blocked channel would be interpreted as a message collision. Hence, no solution is currently available that can provably handle the jammers considered here.

1.4 Structure of the paper

In Section 2 we will present and analyze our MAC protocol, and in Section 3 we will show how to extend it to robust leader election and the fair use of the wireless channel.

2. THE ROBUST MAC PROTOCOL

In this section we present and analyze our MAC protocol. We start with a description of our basic ideas behind the protocol then we formally describe the protocol and analyze its competitiveness. At the end of the section, we also study its energy efficiency.

2.1 Basic approach

Our MAC protocol is based on a simple idea. Suppose that each node v decides to send a message at the current time step with probability p_v with $p_v \leq \hat{p}$ for some small constant $0 < \hat{p} < 1$. 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. Then the following claim holds.

CLAIM 2.1. $q_0 \cdot p \leq q_1 \leq \frac{q_0}{1-\hat{p}} \cdot p$.

PROOF. It holds that $q_0 = \prod_v (1 - p_v)$ and $q_1 = \sum_v p_v \prod_{w \neq v} (1 - p_w)$. Hence,

$$q_1 \leq \sum_v p_v \frac{1}{1-\hat{p}} \prod_w (1 - p_w) = \frac{q_0 \cdot p}{1-\hat{p}} \quad \text{and}$$

$$q_1 \geq \sum_v p_v \prod_w (1 - p_w) = q_0 \cdot p$$

which implies the claim. □

Hence, if the nodes observe that the number of time steps in which the channel is idle is essentially equal to the number of time steps in which exactly one message is sent, then $p = \sum_v p_v$ is likely to be around 1. Otherwise, they know that they need to adapt their probabilities. Therefore, if we had sufficiently many cases in which an idle channel or exactly one message transmission is observed (which is the case if the adversary does not heavily jam the channel and p is not too large), then one can adapt the probabilities p_v just based on these two events and ignore all cases in which the wireless channel is blocked (either because the adversary is jamming it or at least two messages interfere with each other). Essentially, the following strategy could be used at every node for some small enough $\gamma > 0$:

In each time step, every node v is sending a message with probability p_v. If it decides not to send a message, it checks the following two cases:

- If the wireless channel is idle, then $p_v := (1 + \gamma)p_v$.
- If exactly one message is sent, then $p_v := (1 + \gamma)^{-1} p_v$.

The beauty of the algorithm is that it ignores blocked time steps, which makes it more robust against adversarial jamming. However, there is a catch to this strategy because it only works well as long as p does not get too high. If p is initially very high or by chance gets very high, it will be extremely unlikely for the nodes to observe one of the two cases above. Hence, further ideas are necessary.

Our idea is to use a threshold T_v for each node v that cuts its time into time intervals. If v does not observe a successful message transmission for T_v many steps, then p_v is decreased. In this way, eventually p will become small. However, since the algorithm is not aware of T, the time window of the adversary, p may be decreased too quickly or too slowly in this way. Hence, we need proper rules for adapting T_v over time. It turns out that the following rules work: whenever v senses a successful transmission, T_v is decreased by 1, and whenever v does not sense a successful transmission for T_v time steps, T_v is increased by 1 for the next time interval considered by v. One may ask why T_v should not be decreased as well if an idle channel is sensed, but interestingly this is not a good rule, as will come out in the analysis. Next, we give a formal description of our MAC protocol.

2.2 Description of the MAC protocol

In our MAC protocol, each node v maintains a probability value p_v, a threshold 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 $O(1/(\log T + \log\log n))$. Thus, we assume that the nodes have some polynomial estimate of T and even rougher estimate of n. 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 non-synchronized execution of the protocol would also work as long as all nodes operate at roughly the same speed.

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

1. If v senses an idle channel, then $p_v := \min\{(1+\gamma)p_v, \hat{p}\}$.
2. If v successfully receives a message, then $p_v := (1+\gamma)^{-1}p_v$ and $T_v := \max\{1, T_v - 1\}$.

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 step among the past T_v time steps in which v sensed a successful message transmission, then $p_v := (1+\gamma)^{-1}p_v$ and $T_v := T_v + 1$.

2.3 Robustness

Let $N = \max\{T, n\}$. In this section, we will prove the following theorem.

THEOREM 2.2. *For $n \geq 2$ the MAC protocol is constant competitive w.h.p. under any $(T, 1-\epsilon)$-bounded adversary if the protocol is executed for at least $\Theta(\frac{1}{\epsilon}\log N \max\{T, \frac{1}{\epsilon\gamma^2}\log^3 N\})$ many time steps.*

Notice that for $n = 1$ a node will never experience a time step with a successful transmission. Hence, it would just keep reducing its access probability in our protocol, thereby reaching a dormant state, which is the best it can do in this case as there is no one else to communicate with. Thus, it only makes sense to consider the case $n \geq 2$. More on energy efficiency will be discussed later.

The proof of the theorem will frequently use the following general form of the well-known Chernoff bounds, which may be of independent interest. They are derived from Chernoff bounds presented in [23].

LEMMA 2.3. *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 $\mathbb{E}[\prod_{i\in S} X_i] \leq \prod_{i\in S} p_i$ 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*

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

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 $S \subseteq \{1, \ldots, n\}$, then it holds for any $0 < \delta < 1$ that

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

Let V be the set of all nodes. For the proof of the theorem we will consider all possible decompositions of V into a single node v_0 and $U = V \setminus \{v_0\}$. Let $p_t(v)$ be node v's access probability p_v at the beginning of the t-th time step. Furthermore, let $p_t = \sum_{v\in U} p_t(v)$ (i.e., without node v_0) and $L = \Omega(\frac{1}{\epsilon}\log N \max\{T, \frac{1}{\epsilon\gamma^2}\log^3 N\})$ be the number of time steps for which we study the competitiveness of the protocol. If $L \geq N$, we will redefine N to $N = \max\{T, n, L\}$ in order to cover long runtimes. If we can prove a constant competitiveness for any such L, Theorem 2.2 follows.

We prove the theorem by induction over sufficiently large time frames. Let I be a time frame consisting of $\frac{\alpha}{\epsilon}\log N$ *subframes* I' of size $f = \max\{T, \frac{\alpha\beta^2}{\epsilon\gamma^2}\log^3 N\}$, where α and β are sufficiently large constants. Let $F = \frac{\alpha}{\epsilon}\log N \cdot f$ denote the size of I. We assume that at the beginning of I, $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ and $T_v \leq \sqrt{F}/2$ for every node v. Our goal is to show that in this case the MAC protocol is constant competitive for I w.r.t. every subset $U = V \setminus \{v_0\}$ and at the end of I, $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ and $T_v \leq \sqrt{F}/2$ for every node v with probability at least $1 - 1/N^c$ for any constant $c > 0$ (which we will also call *with high probability* or *w.h.p.* in the following). Since initially $T_v = 1$ and $p_v = \hat{p}$ for every v, this implies that the MAC protocol achieves a constant competitiveness in the first time frame, w.h.p., and due to the properties on T_v and p_v, this also holds for polynomially many time frames, w.h.p.

The proof for time frame I proceeds as follows. Consider some fixed subset $U = V \setminus \{v_0\}$. A time step t or subframe I' of I with starting time t is called *good* if $p_t \leq 9$. Otherwise, it is called *bad*. First, we show that for any subframe I' in which initially $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$, also afterwards $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$, w.h.p. (Lemma 2.4). Then we show that for any subframe I' with $T_v \leq (3/4)\sqrt{F}$ for every node $v \in U$ at the beginning of I', the subsequent subframe is good with probability at least $1 - 1/f^c$ for any constant $c > 0$ (which we will call *with moderate probability* or *w.m.p.*) (Lemma 2.7). Based on the insights gained in the proof, we show that in a good subframe I', all non-jammed time steps in I' are good w.m.p. (Corollary 2.11). After that, we prove that a constant fraction of the time steps in such a subframe also have probabilities lower bounded by a constant (Lemma 2.12), w.h.p., which implies that the MAC protocol is constant competitive for I' w.m.p. (Lemma 2.13). If at the beginning of frame I, $T_v \leq \sqrt{F}/2$ for every node $v \in U$, then during the first eighth of I, called J, $T_v \leq (3/4)\sqrt{F}$, no matter what happens to the nodes in J. This allows us to show that a constant fraction of the subframes of J are constant competitive w.h.p., which implies that the MAC protocol is constant competitive for J w.h.p. (Lemma 2.14). With that insight we can show that if at the beginning of J, $T_v \leq \sqrt{F}/2$ for every node $v \in U$, then this also holds at the end of J w.h.p. (Lemma 2.15). Hence, all eighths of I have a constant competitiveness, w.h.p., which implies that I has a constant competitiveness and at the end of I, $T_v \leq \sqrt{F}/2$ for every node v, w.h.p. Applying these results inductively over all time frames I yields Theorem 2.2.

At the end of this subsection, we also study the recovery properties of our MAC protocol (Theorem 2.16). It turns out that the MAC protocol can get quickly out of any set of (p_v, c_v, T_v)-values, which implies that it also works well if the nodes enter the network at arbitrary times and with arbitrary values instead of starting the protocol at the same time and with the same values, which is not realistic in practice.

LEMMA 2.4. *For any subframe I' in which initially $p_{t_0} \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$, the last time step t of I' satisfies $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$, w.h.p.*

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

CLAIM 2.5. *If in subframe I' the number of successful message transmissions is at most k, then every node v increases T_v at most $k + \sqrt{2f}$ many times.*

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

This claim allows us to show the following claim.

CLAIM 2.6. *Suppose that for the first time step t_0 in I', $p_{t_0} \in [1/(f^2(1+\gamma)^{2\sqrt{f}}), 1/f^2]$. 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 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, which implies that for each time step t' with an idle channel, $p_{t'+1} = (1+\gamma)p_{t'}$. Furthermore, at most $\log_{1+\gamma}(1/p_{t_0})$ increases of p_t due to 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 and decreases of p_t in order to avoid the case $p_t \geq 1/f^2$. We know from Claim 2.5 that $k_2 \leq k_1 + \sqrt{2f}$. Hence,

$$k_0 \leq 2\log_{1+\gamma} f + 2\sqrt{f} + 2k_1 + \sqrt{2f}$$

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

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

$$\begin{aligned}\mathbb{P}[\geq 1 \text{ message transmitted at } t] &\leq \sum_v p_v(t) = p_t + \hat{p} \\ &\leq 1/f^2 + \hat{p}\end{aligned}$$

where $\hat{p}$ is due to node v_0 not considered in p_t. Hence, $\mathbb{E}[k_0] \geq (1-1/f^2-\hat{p})g$ and $\mathbb{E}[k_1] \leq (1/f^2+\hat{p})g$. In order to prove bounds on k_0 and k_1 that hold 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 the channel is idle at step t or $p_t \geq 1/f^2$. Then

$$\begin{aligned}\mathbb{P}[X_t = 1] &= \mathbb{P}[\text{channel idle and } p_t \leq 1/f^2] + \mathbb{P}[p_t > 1/f^2] \\ &= \mathbb{P}[p_t \leq 1/f^2] \cdot \mathbb{P}[\text{channel idle} \mid p_t \leq 1/f^2] + \\ &\quad \mathbb{P}[p_t > 1/f^2] \\ &\geq \mathbb{P}[p_t \leq 1/f^2](1 - 1/f^2 - \hat{p}) + \mathbb{P}[p_t > 1/f^2] \\ &\geq 1 - 1/f^2 - \hat{p}\end{aligned}$$

and since this probability bound holds irrespective of prior steps and is *independent* of the adversarial jamming decision at time t, it follows for any set S of time steps prior to some time step t that

$$\mathbb{P}[X_t = 1 \mid \prod_{s \in S} X_s = 1] \geq 1 - 1/f^2 - \hat{p}$$

Thus, for any set of time steps S it holds that $\mathbb{E}[\prod_{s\in S} X_s] \geq (1 - 1/f^2 - \hat{p})^{|S|}$. Together with the fact that $g \geq \epsilon f \geq \alpha \log N$, the Chernoff bounds imply that, w.h.p., either $k_0 > 3g/4$ (given that $\hat{p} \leq 1/24$) or we have a time step t with $p_t \geq 1/f^2$.

On the other hand, let the binary random variable Y_t be 1 if and only if exactly one message is sent at time t and $p_t \leq 1/f^2$. Then

$$\begin{aligned}\mathbb{P}[Y_t = 1] &= \mathbb{P}[p_t \leq 1/f^2] \cdot \mathbb{P}[\text{one msg sent} \mid p_t \leq 1/f^2] \\ &\leq 1/f^2 + \hat{p}\end{aligned}$$

and it holds for any set S of time steps prior to some time step t that

$$\mathbb{P}[Y_t = 1 \mid \prod_{s \in S} Y_s = 1] \leq 1/f^2 + \hat{p}$$

Thus, the Chernoff bounds imply that $k_1 < g/8$, w.h.p. (given that $\hat{p} \leq 1/24$). That, however, would violate the condition that $k_0 \leq 2k_1 + g/2$.

Note that the choice of g is not oblivious as the adversary may *adaptively* decide to set g based on the history of events. Hence, we need to sum up the probabilities over all adversarial strategies of selecting g in order to show that none of them succeeds, but since there are only f many, and for each the claimed property holds w.h.p., the claim follows. □

So suppose that there is a time step t in I' with $p_t \geq 1/f^2$. If t belongs to one of the last $\beta \log N$ non-jammed steps in I', then it follows for the probability $p_{t'}$ at the end of I' that

$$p_{t'} \geq \frac{1}{f^2} \cdot (1+\gamma)^{-2\beta \log N + \sqrt{2f}} \geq \frac{1}{f^2(1+\gamma)^{2\sqrt{f}}}$$

given that $\epsilon = \Omega(1/\log^3 N)$ as at most $\beta \log N$ decreases of p_t can happen due to a successful transmission and at most $\beta \log N + \sqrt{2f}$ decreases of p_t can happen due to exceeding T_v.

Suppose, on the other hand, that there is no time step t among the last $\beta \log N$ non-jammed steps in I' with $p_t \geq 1/f^2$. In this case, we assume that a specific step t in I' outside of these last steps is the last time step with $p_t \geq 1/f^2$. When defining k_0, k_1 and k_2 as above but from that point on it follows that $p_{t'}$ at the end of I' is still bounded from below by $1/(f^2(1+\gamma)^{2\sqrt{f}})$ as long as $k_0 \geq k_1$. Our analysis above implies that this is true w.h.p. (see Claim 2.8 for similar arguments in the other direction), which finishes the proof of Lemma 2.4. □

LEMMA 2.7. *For any subframe I' with $T_v \leq (3/4)\sqrt{F}$ for all nodes v at the beginning of I', the last time step t of I' satisfies $p_t \leq 9$ w.m.p.*

PROOF. We first show that there is a time step t in I' with $p_t \leq 6$, w.h.p. Let the time steps in which the adversary does not jam the channel and at most one message is sent by the nodes be called *useful*. Suppose that there are g useful 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. In order to establish a relationship between k_0 and k_1 we need the following claims.

CLAIM 2.8. *If all time steps $t \in I'$ satisfy $p_t > 6$, then it holds for any $g \geq \delta \log N$ for a sufficiently large constant δ that $k_1 \geq k_0$ w.h.p.*

PROOF. Let $q_0(t)$ be the probability of an idle channel and $q_1(t)$ be the probability of a successful message transmission at a useful step t. If $p_t > 6$, then it follows from Claim 2.1 that

$$\begin{aligned}\mathbb{P}[\text{channel idle}] &= \frac{q_0(t)}{q_0(t) + q_1(t)} \leq \frac{q_0(t)}{q_0(t) + p_t \cdot q_0(t)} \\ &\leq \frac{1}{1+6} = \frac{1}{7}\end{aligned}$$

irrespective of what happened at previous time steps. Hence, $\mathbb{E}[k_0] \leq g/7$ under the assumption that all useful time steps t satisfy $p_t > 6$. Thus, our Chernoff bounds yield $k_0 \leq g/2$ w.h.p. (given that δ is a sufficiently large constant), which implies that $k_1 \geq k_0$. □

Now we are ready for the following claim.

CLAIM 2.9. *If all time steps in I' satisfy $p_t > 6$, then it must hold w.h.p. that*

$$k_1 - 2\log_{1+\gamma} N \leq (5/4)k_0$$

PROOF. If exactly one message is sent at a step t, then $p_{t+1} \geq (1+\gamma)^{-1}p_t$ and

$$p_{t+1} \leq (1+\gamma)^{-1}(p_t - \hat{p}) + \hat{p} \leq (1+\gamma)^{-1}p_t + \gamma(1+\gamma)^{-1}\hat{p}$$

because only the sending node does not decrease its probability, and for this node the maximum probability is $\hat{p}$. For $p_t > 6$ it follows that $p_{t+1} \in [(1+\gamma)^{-1}p_t, (1+\gamma)^{-4/5}p_t]$. From Claim 2.8 we now that after the first $\delta \log N$ useful steps, there must have been more steps with a successful transmission than with an idle channel for any one of the remaining useful steps, w.h.p, which implies that for each of them, $p_v < \hat{p}$ for all nodes v. Thus, whenever there is an idle channel for these steps, $p_{t+1} = (1+\gamma)p_t$. Hence, if we start with $p_t = 6$ after the first $\delta \log N$ useful steps, then in order to avoid a step t' with $p_{t'} \leq 6$ in I' we must have that $k_1 \leq (5/4)k_0$. Since p_t might be as high as $\hat{p}n$ initially, we can allow at most $(5/4)\log_{1+\gamma} N$ further events of a successful message transmission without having a step t' with $p_{t'} \leq 6$.

Since $\log_{1+\gamma} N = \omega(\log N)$, it holds that

$$\delta \log N + (5/4)\log_{1+\gamma} N \leq 2\log_{1+\gamma} N$$

for a sufficiently large N, which implies the claim. □

Also, $k_0 + k_1 = g$. Suppose that $g \geq \delta \log_{1+\gamma} N$ for a sufficiently large constant δ. It holds that

$$(g - k_0) - 2g/\delta \leq (5/4)k_0 \quad \Leftrightarrow \quad k_0 \geq (4/9)(1 - 2/\delta)g$$

We know from the proof of Claim 2.8 that for any useful step t with $p_t > 6$, $\mathbb{P}[\text{channel idle}] \leq \frac{1}{7}$. Hence, $\mathbb{E}[k_0] \leq g/7$. Since random decisions are made independently in each step, our Chernoff bounds imply that $k_0 < (4/9)(1 - 2/\delta)g$ w.h.p. if δ is sufficiently large.

Thus, if I' contains at least $\delta \log_{1+\gamma} N$ useful steps, we are done. Otherwise, notice that for every node v it follows from the MAC protocol and the choice of f and F that if initially $T_v \leq (3/4)\sqrt{F}$, then T_v can be at most $\sqrt{F}$ during I'. Let us cut I' into m intervals of size $2\sqrt{F}$ each. It is easy to check that if β in the definition of f is sufficiently large compared to δ, then $m \geq 3\delta \log_{1+\gamma} N$. If there are less than $\delta \log_{1+\gamma} N$ useful steps, then at least $2\delta \log_{1+\gamma} N$ of these intervals do not contain any useful step, which implies that p_v is reduced by at least $(1+\gamma)^{-1}$ by each v in each of these intervals.

Hence, altogether, every p_v gets reduced by a factor of at least $(1+\gamma)^{-2\delta \log_{1+\gamma} N}$ during I'. The useful time steps can only raise that by $(1+\gamma)^{\delta \log_{1+\gamma} N}$, so altogether we must have $p_t \leq 6$ at some time point during I', w.h.p.

In the following, let t_0 denote any time in I' with $p_{t_0} \leq 6$. We finally prove the following claim.

CLAIM 2.10. *For any useful time step t after a step t_0 in I' with $p_{t_0} \leq \phi$ for some $\phi \geq 6$ and any constant $\delta > 0$ it holds that*

$$\mathbb{P}[p_t \geq (1+\delta)\phi] \leq 8 \cdot (1+\delta)^{-1/(6\gamma)}$$

PROOF. Suppose that t_0 be the last useful time step before step t in I' with $p_{t_0} \leq \phi$. Let g be the number of useful time steps from t_0 to t. Then $g \geq \ln(1+\delta)/\ln(1+\gamma)$ because otherwise it is not possible that $p_t \geq (1+\delta)\phi$. Recall that for any useful step r with $p_r \geq 6$, $\mathbb{P}[p_{r+1} = (1+\gamma)p_r] \leq 1/7$. If exactly one message is sent at a useful step, then $p_{r+1} \in [(1+\gamma)^{-1}p_r, (1+\gamma)^{-4/5}p_r]$. Let k_0 be the number of useful steps with an idle channel and k_1 be the number of useful steps with a successful message transmission. It must hold that $k_0 \geq (4/5)k_1 + \ln(1+\delta)/\ln(1+\gamma)$ so that $p_t \geq (1+\delta)\phi$. Also, $k_0 + k_1 = g$. Hence, $k_0 \geq (4/9)g + (5/9)\ln(1+\delta)/\ln(1+\gamma) \geq \max\{(4/9)g, \ln(1+\delta)/\ln(1+\gamma)\}$. It holds that $\mathbb{E}[k_0] \leq g/7$, so the Chernoff bounds imply that

$$\begin{aligned}\mathbb{P}[k_0 \geq (4/9)g] &\leq \mathbb{P}[k_0 \geq (1+2)g/7] \\ &\leq e^{-[2^2/(2(1+2/3))](g/7)} = e^{-g/6}\end{aligned}$$

Hence,

$$\begin{aligned}\mathbb{P}[p_t \geq (1+\delta)\phi] &\leq \sum_{g \geq \frac{\ln(1+\delta)}{\ln(1+\gamma)}} \mathbb{P}[k_0 \geq (4/9)g] \leq \sum_{g \geq \frac{\ln(1+\delta)}{\ln(1+\gamma)}} e^{-g/6} \\ &\leq 8(1+\delta)^{-\frac{1}{6\ln(1+\gamma)}} \leq 8(1+\delta)^{-1/(6\gamma)}\end{aligned}$$

□

Since we assume that $\gamma = O(1/\log f)$, it follows that w.m.p., $p_t \leq (1+\delta)6$ for any particular time step t after t_0, resulting in the lemma with $\delta = 1/2$. □

Claim 2.10 with $\phi = 9$ and $\delta = 1/3$ implies the following result.

COROLLARY 2.11. *For any good subframe I', all non-jammed time steps t of I' satisfy $p_t \leq 12$ w.m.p.*

We also need to show that for a constant fraction of the non-jammed time steps in a good subframe, p_t is also lower bounded by a constant. Recall that $\hat{p} \leq 1/24$.

LEMMA 2.12. *For any subframe I' in which initially $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$, at least $1/8$ of the non-jammed steps t satisfy $p_t \geq \hat{p}$, w.h.p.*

PROOF. Let G be the set of all non-jammed time steps in I' and S be the set of all steps t in G with $p_t < \hat{p}$. Let $g = |G|$ and $s = |S|$. If $s \leq 7g/8$, we are done. Hence, consider the case that $s \geq 7g/8$.

Suppose that p_t must be increased k_0 many times to get from its initial value up to a value of $\hat{p}$ and that p_t is decreased k_1 many times in S due a successful message transmission. Furthermore, let k_2 be the maximum number of times a node v decreases p_v due to $c_v > T_v$ in the MAC protocol. For S to be feasible (i.e., probabilities can be assigned to each $t \in S$ so that $p_t < \hat{p}$) it must hold for the number ℓ of times in S in which the channel is idle that

$$\ell \leq k_0 + k_1 + k_2$$

For the special case that $k_0 = k_2 = 0$ this follows from the fact that whenever there is a successful message transmission, p_t is reduced to $p_{t+1} \geq (1+\gamma)^{-1}p_t$. On the other hand, whenever there is an idle channel, it holds that $p_{t+1} = (1+\gamma)p_t$ because of $p_t < \hat{p}$. Thus, if $\ell > k_1$, then one of the steps in S would have to have a probability of at least $\hat{p}$, violating the definition of S. k_0 comes into the formula due to the startup cost of getting to a value of $\hat{p}$, and k_2 comes into the formula since the reductions of the $p_t(v)$ values due to $c_v > T_v$ in the MAC protocol allow up to k_2 additional increases of p_t for S to stay feasible.

First, we bound ℓ. If $p_t < \hat{p}$, then $\mathbb{P}[\text{idle channel at step } t] \geq 1 - \hat{p} - \hat{p}$ (where the second $\hat{p}$ is due to node v_0), irrespective of prior time steps, Hence, $\mathbb{E}[\ell] \geq (1 - 2\hat{p})s$. For $\hat{p} \leq 1/24$ our Chernoff bounds imply because of $s \geq 7g/8 \geq (7/8)\epsilon f$ that $\ell \geq s/2$ w.h.p. If at the beginning of I', $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ then $k_0 \leq 2\log_{1+\gamma} f + 2\sqrt{f}$. Moreover, $k_2 \leq g/8 + k_1 + \sqrt{2f}$ because of Claim 2.5. Hence, $k_0 + k_1 + k_2 \leq 2\log_{1+\gamma} f + 2\sqrt{f} + 2k_1 +$

$g/8 + \sqrt{2f}$, which must be at least $s/2$ so that $\ell \leq k_0 + k_1 + k_2$ (given that $\ell \geq s/2$). Suppose that $2\log_{1+\gamma} f + 4\sqrt{f} \leq \epsilon f/16$ (which is true if $f = \Omega(1/\epsilon^2)$ is large enough). Then for this to be true it must hold that

$$2k_1 + g/8 + g/16 \geq (7g/8)/2 \quad \Leftrightarrow \quad k_1 \geq g/8$$

If $k_1 \geq g/8$ then also $k_1 \geq s/8$, so our goal will be to show that $k_1 < s/8$ w.h.p.

If $p_t < \hat{p}$, then $\mathbb{P}$[successful message transmission at step t] $\leq 2\hat{p}$, irrespective of prior time steps. Hence, $\mathbb{E}[k_1] \leq 2\hat{p}s$. Furthermore, for $\hat{p} \leq 1/24$ our Chernoff bounds imply because of $s \geq 7g/8 \geq (7/8)\epsilon f$ that $k_1 < s/8$ w.h.p. Since there are at most f^2 ways (for the adversary) of choosing g and s, this holds for any combination of g and s, which yields the lemma. □

Combining the results above, we get:

LEMMA 2.13. *For any good subframe I' the MAC protocol is constant competitive in I' w.m.p.*

PROOF. From Corollary 2.11 and Lemma 2.12 we know that in a good subframe at least $1/8$ of the non-jammed time steps t have a constant probability value p_t w.m.p. For these steps there is a constant probability that a message is successfully sent. Using the Chernoff bounds results in the lemma. □

Consider now the first eighth of frame I, called J.

LEMMA 2.14. *If at the beginning of J, $p_v \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ and $T_v \leq \sqrt{F}/2$ for all nodes v, then we also have $p_v \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ at the end of J for every v and the MAC protocol is constant competitive for J, w.h.p.*

PROOF. The bound for p_v at the end of J directly follows from Lemma 2.4.

Suppose, as a worst case, that initially $T_v = \sqrt{F}/2$ for some v. Clearly, T_v assumes the maximum possible value at the end of J if T_v is never decreased in J. Since T_v can be increased at most $(F/8)/(\sqrt{F}/2) = \sqrt{F}/4$ many times in J, T_v can reach a maximum value of at most $(3/4)\sqrt{F}$ inside of J, so we can apply Lemma 2.7.

Recall that J consists of $k = \frac{\alpha}{8\epsilon}\log N$ many subframes, numbered $I_1, \ldots, I_k$. For each I_i, let the binary random variable X_i be 1 if and only if I_i is good. From Lemma 2.7 it follows that for any $i \geq 1$ and any set $S \subseteq \{1, \ldots, i-1\}$,

$$\mathbb{P}[X_i = 1 \mid \prod_{j \in S} X_j = 1] \geq 1 - 1/f^c$$

for some constant c that can be made arbitrarily large. Hence, for any set $S \subseteq \{1, \ldots, k\}$, $\mathbb{E}[\prod_{i \in S} X_i] \geq (1 - 1/f^c)^{|S|}$. Our Chernoff bounds therefore imply that at most $(\alpha/24\epsilon)\log N$ of the subframes in J are bad, w.h.p, if α is sufficiently large. According to Lemma 2.13, each of the good subframes is constant competitive w.m.p., where the probability bounds are only based on events in the subframes themselves and therefore hold irrespective of the other subframes (given that each of them is good). So the Chernoff bounds imply that at most $(\alpha/24\epsilon)\log N$ of them do not result in a constant competitiveness of the MAC protocol, w.h.p. The remaining $(\alpha/24\epsilon)\log N$ subframes in J achieve constant competitiveness, which implies that the MAC protocol is constant competitive on J, w.h.p. □

We finally need the following lemma that bounds T_v. The proof of this lemma requires considering all possible decompositions of V into a node v_0 and $U = V \setminus \{v_0\}$ so that every node experiences many successful transmissions.

LEMMA 2.15. *If at the beginning of J, $T_v \leq \sqrt{F}/2$ for all v, then it holds that also $T_v \leq \sqrt{F}/2$ at the end of J, w.h.p.*

PROOF. We know from Lemma 2.14 that for any node v our protocol is constant competitive for $V \setminus \{v\}$ w.h.p. Hence, every node v notices $\Omega(\epsilon|J|)$ successful message transmissions in J w.h.p. T_v is maximized at the end of J if all of these successful transmissions happen at the beginning of J, which would get T_v down to 1. Afterwards, T_v can raise to a value of at most t for the maximum t with $\sum_{i=1}^{t} i \leq |J|$. Since such a t can be at most $\sqrt{2|J|}$, it follows that T_v can be at most $\sqrt{2F/8} = \sqrt{F}/2$ at the end of J, w.h.p. □

Inductively using Lemmas 2.13 and 2.15 on the eighths of frame I implies that our MAC protocol is constant competitive on I and at the end of I, $p_v \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ and $T_v \leq \sqrt{F}/2$ for all v w.h.p. Hence, our MAC protocol is constant competitive for L many time steps, w.h.p., for any $L = \Omega(\frac{1}{\epsilon}\log N \max\{T, \frac{1}{\epsilon\gamma^2}\log^3 N\})$, which implies Theorem 2.2.

Finally, we show that our protocol can quickly recover from any setting of the (T_v, c_v, p_v)-values.

THEOREM 2.16. *For any p_{t_0} and $\hat{T} = \max_v T_v$ it takes at most $O(\frac{1}{\epsilon}\log_{1+\gamma}(1/p_{t_0}) + \hat{T}^2)$ many time steps, w.h.p., until the MAC protocol satisfies again $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ and $\max_v T_v \leq \sqrt{F}/2$ for the original definitions of F and f above.*

PROOF. Suppose that $p_{t_0} < 1/(f^2(1+\gamma)^{2\sqrt{f}})$ for some time point t_0. Then it follows from the constraints of the adversary and the Chernoff bounds that it takes at most $\frac{\delta}{\epsilon}\log_{1+\gamma}(1/p_{t_0})$ steps for some sufficiently large constant δ to get the system from p_{t_0} up to $p_{t_0}^{1/2}$, w.h.p. (in fact, with a probability of at least $1 - p_{t_0}^c$ for any constant c, irrespective of $\hat{T}$). Another $\frac{\delta}{2\epsilon}\log_{1+\gamma}(1/p_{t_0})$ steps will then get the system from $p_{t_0}^{1/2}$ to $p_{t_0}^{1/4}$, w.h.p. (in fact, with probability at least $1 - (p_{t_0}^{1/2})^c$ for any constant c). Continuing these arguments in order to get from $p_{t_0}^{1/2^i}$ to $p_{t_0}^{1/2^{i+1}}$ it follows that altogether at most $\frac{2\delta}{\epsilon}\log_{1+\gamma}(1/p_{t_0})$ steps are needed to get the system from p_{t_0} to a probability $p_t \geq \frac{1}{f^2(1+\gamma)^{2\sqrt{f}}}$, w.h.p. (or more precisely, with probability at least $1 - 1/N^c$).

It remains to bound the time to get T_v down to $\sqrt{F}/2$ for every v. It holds that $\hat{T} \leq \sqrt{F}/2$ if and only if $F \geq 4\hat{T}^2$. Hence, consider a time frame I of size $F' = \max\{F, 4\hat{T}^2\}$ for the old definition of F above, where I starts at the point at which the probabilities p_v have recovered to $p_t \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$. Then all the proofs above go through and imply that I is constant competitive. Moreover, when cutting I into pieces of size $|I|/32$ instead of $|I|/8$, the proof of Lemma 2.15 implies that at the end of the first $1/32$-piece J of I, $T_v \leq \sqrt{F'}/4$, w.h.p. Hence, the time frames of the nodes shrank by a factor of at least 2 in J. Inductively using this bound, it follows that also at the end of I, $T_v \leq \sqrt{F'}/4$ for all v, w.h.p. This allows us to reduce F' by a factor of 2 for the next frame I. Also for this F', we get $T_v \leq \sqrt{F'}/4$ for all v, w.h.p., so we can keep shrinking I by a factor of 2 until $|I| = F$ for the original F considered in our proofs above. Altogether, the recovery to $\hat{T} \leq \sqrt{F}/2$ for all v takes at most $O(\hat{T}^2)$ time.

Combining the two upper bounds for the recovery time yields the theorem. □

2.4 Energy efficiency

Next, we show that our MAC protocol is very energy-efficient under adversarial attacks. The first lemma follows directly from our insights gained in the previous subsection.

LEMMA 2.17. *For any time frame I of size F as defined above, the total energy spent by all the nodes together on sending out messages is bounded by $O(F)$ w.h.p.*

If the adversary performs permanent jamming, the energy spent on message transmissions even converges, i.e., our MAC protocol reaches a dormant stage.

LEMMA 2.18. *Consider any time step t_0 with $\sum_v p_v \leq p$ and $\max_v T_v \leq \hat{T}$ for some values $p > 0$ and $\hat{T} \geq 1/\gamma$. Then for any continuous jamming attack starting at t_0 the total energy consumption of the nodes during the entire attack is at most $O(p \cdot \hat{T}/\gamma + \log N)$ w.h.p.*

PROOF. First, we determine the expected energy consumption of a single node v. Let $p_v(t)$ be the probability that v transmits a message in round $t_0 + t$. Due to our MAC protocol, $p_v(t)$ decreases by $(1+\gamma)^{-1}$ at latest for $t = \hat{T}$, then another time after $\hat{T}+1$ further steps, another time after $\hat{T}+2$ further steps, and so on. Hence, the total expected energy consumption of v for any continuous jamming attack is at most

$$\begin{aligned}
&\sum_{T_v \geq \hat{T}} T_v \cdot p_v(t_0)(1+\gamma)^{T_v - \hat{T}} \\
&= \quad p_v(t_0) \sum_{i \geq 0} (\hat{T}+i)(1+\gamma)^{-i} \\
&\leq \quad \frac{1+\gamma}{\gamma} \cdot \hat{T} \cdot p_v(t_0) + \left(\frac{1+\gamma}{\gamma}\right)^2 \cdot p_v(t_0) \\
&= \quad O(p_v(t_0)\hat{T}/\gamma)
\end{aligned}$$

Summing up over all nodes, we obtain a total expected energy consumption of $O(p \cdot \hat{T}/\gamma)$. Since all transmission decisions are done independently at random, the Chernoff bounds imply a total energy consumption of at most $O(p \cdot \hat{T}/\gamma + \log N)$ w.h.p. □

In our MAC protocol, beyond f steps after any initial choice of the access probabilities, $p = O(\log N)$, w.h.p. This is due to the proof of Lemma 2.7 and the fact that for $p \geq c \log N$, the probability that an idle channel is experienced is at most $1/N^c$, so further increasing p has a polynomially small probability. Furthermore, $\hat{T} = O(\log^2 N/\gamma)$ w.h.p. for any constant ϵ given that all nodes v start with $T_v = 1$. Hence, the total energy consumption of our MAC protocol under a permanent attack that starts after f steps would be bounded by $O(\log^3 N/\gamma^2)$ w.h.p.

3. APPLICATIONS OF THE MAC PROTOCOL

In this section we will demonstrate how our robust MAC protocol can be extended to perform robust leader election or to select fair access probabilities for the nodes.

3.1 Leader election

Consider the following adaptation of the MAC protocol. In addition to c_v, T_v and p_v, every node v maintains a counter s_v for successful transmissions. v also stores one of the states {unknown, leader, follower}. Initially, every node v sets $T_v := 1$, $c_v := 1$ and $p_v := \hat{p}$. Also, v sets s_v to 0 and its state to "unknown". Afterwards, v does the following in each step.

v decides with probability p_v to send a message. If it does so, its message is piggy-backed with s_v. If it decides not to send a message, it checks the following two conditions:

1. If v senses an idle channel, then $p_v := \min\{(1+\gamma)p_v, \hat{p}\}$.
2. If v successfully receives a message with some counter s_w, then $p_v := (1+\gamma)^{-1}p_v$ and $T_v := \max\{1, T_v - 1\}$. If v is still in the state "unknown", then v checks the following two cases: If $s_v \geq s_w$ then v becomes a "follower", otherwise v becomes a "leader". In any case, v sets $s_v := \max\{s_v, s_w\} + 1$.

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 step among the past T_v time steps in which v sensed a successful message transmission, then $p_v := (1+\gamma)^{-1}p_v$ and $T_v := T_v + 1$.

This protocol has the following performance.

THEOREM 3.1. *Within $O(\frac{1}{\epsilon}\log N \max\{T, \frac{1}{\epsilon\gamma^2}\log^3 N\})$ many steps, the leader election protocol reaches a state in which there is exactly one leader and the other nodes are followers, w.h.p.*

PROOF. At the beginning, all counters s_v are set to 0. Once the first node, say v, is able to successfully transmit a message, then all nodes $w \neq v$ will become a follower and set s_w to 1. v may then go on being successful for k more steps until the first node $w \neq v$ successfully transmits a message. When w transmits its message, it also sends $s_w = k+1$ which is greater than s_v since s_v is still set to 0. Hence, v will become a leader. According to the analysis of our original MAC protocol, which is embedded in our leader election protocol, it takes at most $O(\frac{1}{\epsilon}\log N \max\{T, \frac{1}{\epsilon\gamma^2}\log^3 N\})$ many steps until at least two nodes successfully transmit a message (as constant competitiveness is ensured for any set $U = V \setminus \{v_0\}$), w.h.p., which yields the theorem. □

Once a node becomes a leader, it may then select a fixed access probability of $\hat{p}$ (which, as we know from our analysis, does not cause problems for the competitiveness of the follower nodes) so that an effective coordination of the follower nodes is possible.

3.2 Establishing fairness

In our original MAC protocol, some probabilities may eventually dominate the others. This is due to the fact that whenever there is a successful message transmission, all nodes sensing the successful transmission are lowering their access probabilities while the access probability of the sending node stays the same. Since nodes with a larger access probability are more likely to transmit a message, there is a tendency towards preserving access probabilities of those nodes that already have large access probabilities so that the gap between large and small probabilities will increase over time. This would result in an unfair use of the channel among the nodes. In order to ensure fairness, we slightly modify our MAC protocol. In the new protocol, each node v maintains a counter s_v for successful transmissions and a counter m_v of the different nodes it has seen so far. It also maintains a state in {covered, uncovered} and memorizes in $olds$ the last counter it has seen so far. Initially, every node v sets $T_v := 1$, $c_v := 1$ and $p_v := \hat{p}$. Also, s_v and m_v are set to 0, $olds$ is set to -1, and the state is set to "uncovered". Afterwards, every node v does the following in each step.

v decides with probability p_v to send a message. If it does so, its message is piggy-backed with s_v and its state. If it decides not to send a message, it checks the following two conditions:

1. If v senses an idle channel, and v is still uncovered then $p_v := \max\{(1+\gamma)p_v, \hat{p}\}$.
2. If v successfully receives a message with some counter s_w, then v considers the following cases.

- If w is uncovered and $s_w \neq olds$ then $m_v := m_v + 1$. If v is covered then it sets $p_v := \hat{p}/m_v$.
- If v is uncovered and $s_w > s_v$ then v changes its state to "covered", sets $m_v := m_v + 1$ and $p_v := \hat{p}/m_v$.
- If v is uncovered and $s_w \leq s_v$ then v sets $p_v := (1+\gamma)^{-1} p_v$ and $T_v := \max\{1, T_v - 1\}$.

$olds := s_w$ and $s_v := \max\{s_v, s_w\} + 1$.

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 step among the past T_v time steps in which v sensed a successful message transmission, then $p_v := (1+\gamma)^{-1} p_v$ and $T_v := T_v + 1$.

We will prove the following result for this protocol:

THEOREM 3.2. *If $T \leq n^\delta$ for some constant $\delta < 1$ and $\hat{p} \leq 1/48$, then it takes at most $O(n/\epsilon)$ time steps until all nodes have an access probability of $\Theta(1/n)$, w.h.p.*

We first state some properties of s_v and m_v.

LEMMA 3.3. *At any time, s_v is equal to the number of successful transmissions performed so far, except for the most recent transmissions of v without a transmission of a node $w \neq v$ afterwards.*

PROOF. We prove the lemma by induction over the number of successful transmissions. Initially, the lemma is certainly true. So consider the situation that it is still true after the first k successful transmissions. Let v be the origin of the last message transmission. Then $s_w = k$ for all $w \neq v$ and $s_v = k - r_v$ where r_v is the number of most recent transmissions of v without a transmission of a node $w \neq v$ afterwards.

If the next node successfully transmitting a message is v, then all other nodes w receive a message with $s_v \leq s_w$ and therefore increase s_w by 1, which satisfies the lemma. If, on the other hand, some node $u \neq v$ transmits a message, then v receives a message with $s_u > s_v$, so it updates s_v to $s_u + 1 = k + 1$. All nodes $w \notin \{u, v\}$ satisfy $s_u \leq s_w$, so they increase s_w by 1. In both cases, the lemma holds again, which completes the proof. □

LEMMA 3.4. *A node is in the state "covered" if and only if it has already successfully sent a message and received a message from a node afterwards.*

PROOF. According to the protocol, a node v only becomes covered if $s_w > s_v$, so the lemma follows from Lemma 3.3. □

LEMMA 3.5. *m_v counts the number of different nodes that have successfully sent a message, except v itself if v has successfully sent messages without receiving a message from another node so far.*

PROOF. We prove the lemma above by induction over the number of successful transmissions. Initially, the lemma is certainly true. So suppose that it is true after the first k successful transmissions. Let v be the origin of the last message transmission. We distinguish between several cases for the $k+1$th message transmission.

Suppose that the next node successfully transmitting a message is v. Then $s_v = olds$ and $s_v \leq s_w$ for every other node w according to Lemma 3.3. Hence, no changes will happen to the m_w's. So suppose that the next node transmitting is $u \neq v$. Then $s_u > olds = s_v$ according to Lemma 3.3. Thus, if v was still uncovered, then v changes to "covered" and increases m_v by 1, which satisfies the lemma. Otherwise, v does nothing, which also satisfies our lemma as well. For all other nodes w, we consider the following cases. If u is uncovered, then each of these nodes increases m_w by 1 (because of $s_u > olds$), and otherwise, they leave m_w as before, which satisfies our claim. Putting all pieces together, the lemma follows. □

Lemmas 3.4 and 3.5 and the way the covered nodes set their access probabilities immediately yield the following result.

COROLLARY 3.6. *At any time, the set of covered nodes together have an access probability in $[(1 - 1/(m+1))\hat{p}, \hat{p}]$, where m is the number of nodes with successful transmissions so far, and this probability is shared evenly among them.*

Hence, once all nodes are covered, fairness is established among all nodes. The following lemma bounds the time necessary to cover all nodes.

LEMMA 3.7. *If $T \leq n^\delta$ for some constant $\delta < 1$ and $\hat{p} \leq 1/48$, then it takes at most $O(n/\epsilon)$ time steps until until all nodes are covered, w.h.p.*

PROOF. First, we establish the following claim.

CLAIM 3.8. *All nodes that have not been able to successfully send a message so far have the same access probabilities.*

PROOF. Notice that all nodes that have not been able to successfully send a message so far have the property that whenever there was an idle channel or a successful message transmission, all of them noticed that. Since all of them start with $c_v := 1$ and $T_v := 1$, this implies that their time frames are in synchrony and any changes in the access probabilities due to a channel condition or the case $c_v > T_v$ are done in synchrony as well. As all nodes initially start with $p_v = \hat{p}$, the claim follows. □

Notice that even if the nodes do not initialize p_v, c_v and T_v with the same values, the analysis of our original MAC protocol implies that as long as all nodes v initially satisfy $T_v \leq \sqrt{F}$ (for the parameter F in the previous section), it takes at most F steps until a point is reached at which all $T_v = 1$ for all v, so the non-successful nodes will operate in synchrony from that point on (though with different probability offsets). For simplicity, however, we will consider the case of Claim 3.8.

Now, it follows from the analysis of the original MAC protocol that the time needed for the first node to be covered is polylogarithmic in N w.h.p. Once the first node has been covered, the remaining nodes quickly become covered as well, as shown next.

CLAIM 3.9. *Consider any consecutive sequence of $\log n$ nodes that become covered during the algorithm after at least one node has been covered. The number of successful transmissions they need for that is $O(\log n)$ w.h.p.*

PROOF. Let C be the set of covered nodes and $m = |C|$. Moreover, let $p_0 = \prod_v (1 - p_v)$ be the probability that the channel is idle at a given time step. Since the covered nodes together have an access probability of at least $(1 - 1/(m+1))\hat{p}$ at any time (Corollary 3.6) and the least recently successful but not yet covered node, v, has an access probability of at most $\hat{p}$, it holds that

$$\mathbb{P}[\text{node in } C \text{ successful}] = \sum_{u \in C} p_u \prod_{w \neq v} (1 - p_w) \geq \frac{p_0 \cdot m}{m+1} \cdot \hat{p}$$

and

$$\mathbb{P}[\text{node } v \text{ successful}] = p_v \prod_{w \neq v} (1 - p_w) \leq \frac{p_0}{1 - \hat{p}} \cdot \hat{p}$$

Thus,

$$\mathbb{P}[\text{node in } C \text{ successful}] \geq \frac{(1-\hat{p})m}{m+1} \cdot \mathbb{P}[\text{node } v \text{ successful}]$$

which implies that the probability that k consecutive successful transmissions are due to v is at most $(1/(1+c))^k$ with $c = (1-\hat{p})m/(m+1)$. This is polynomially small if $k = \Omega(\log n)$. Furthermore, when considering a consecutive sequence of $O(\log n)$ nodes that become covered, it follows from the independence of the transmission attempts of the nodes that altogether the number of successful transmissions they need for that is $O(\log n)$ w.h.p. $\square$

It remains to bound the time until the uncovered nodes (at the time of the transmission) have had $\Omega(n)$ successful transmissions. Let $v_1, v_2, \ldots, v_n$ be the order in which the nodes become covered, i.e., v_i is the ith node with a successful transmission. Let $U_i = \{v_1, \ldots, v_i\}$ for all $i \geq 1$. Once v_i has had its first successful transmission, we consider the partition $(U_i, V \setminus U_i)$. For U_i we know that $\sum_{u \in U_i} p_u \leq 2\hat{p} \leq 1/24$ and at the time v_i had its first success, p_{v_i} is a $1/|V \setminus U_{i-1}| = 1/(n-i+1)$-fraction of $\sum_{v \in V \setminus U_i} p_v$. Hence, when switching from $(U_i, V \setminus U_i)$ to $(U_{i+1}, V \setminus U_{i+1})$, only a small fraction of the probability gets lost in the uncovered nodes, and the probability in U_i stays bounded by $1/24$. In fact, as long as there are still at least f uncovered nodes left, then the total reduction in the access probability over a subframe is at most $\prod_{g=f}^{2f}(1-1/g) \geq 1/e$. This is low enough so that the analysis of the original MAC protocol still applies, i.e. the protocol is constant competitive w.r.t. the still uncovered nodes within time frames of size F, w.h.p. Once there are less than f uncovered nodes, the analysis implies that at least one uncovered node gets covered within a time frame of size F, w.h.p. Combining that with Claim 3.9, it takes at most $O(n/\epsilon + f \cdot F)$ steps, w.h.p., for all nodes to become covered. When assuming that $T \leq n^\delta$ for some constant $\delta < 1$, the lemma follows. $\square$

4. CONCLUSIONS

In this paper we presented the first MAC protocol that is provably robust against adversarial jammers. In fact, our protocol can even handle adaptive jammers. Many open questions remain. Can the MAC protocol be extended to multi-hop networks? How can we adapt to join and leave behavior or mobility of the nodes, and which rate is sustainable without losing a constant competitiveness? Can the MAC protocol be modified so that no knowledge about T and n is required any more? We have tried several variants of our protocol that all had counterexamples. A constant γ appears to work fine under stochastic jammers, but it does not seem to work under adaptive jammers. What other applications than leader election and a fair use of the wireless channel can be considered?

5. 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] 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*, page 1265, 2008.

[3] Michael A. Bender, Martin Farach-Colton, Simai He, Bradley C. Kuszmaul, and Charles E. Leiserson. Adversarial contention resolution for simple channels. In *Proc. of SPAA '05*, pages 325–332, 2005.

[4] 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.

[5] J.T. 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] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz 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. 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.

[9] Leslie Ann Goldberg, Philip D. Mackenzie, Mike Paterson, and Aravind Srinivasan. Contention resolution with constant expected delay. *Journal of the ACM*, 47(6):1048–1096, 2000.

[10] Johan Hastad, Tom Leighton, and Brian Rogoff. Analysis of backoff protocols for mulitiple accesschannels. *SIAM Journal on Computing*, 25(4):740–774, 1996.

[11] IEEE. Medium access control (MAC) and physical specifications. In *IEEE P802.11/D10*, 1999.

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

[13] K. Kothapalli, C. Scheideler, M. Onus, and A. Richa. Constant density spanners for wireless ad hoc networks. In *Proc. of SPAA '05*, pages 116–125, 2005.

[14] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Radio Network Clustering from Scratch. In *Proc. of ESA '04*, 2004.

[15] Byung-Jae Kwak, Nah-Oak Song, and Leonard E. Miller. Performance analysis of exponential backoff. *IEEE/ACM Transactions on Networking*, 13(2):343–355, 2005.

[16] Y.W. 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] Xin Liu, Guevara Noubir, Ravi Sundaram, and San Tan. Spread: Foiling smart jammers using multi-layer agility. In *Proc. of Infocom '07*, pages 2536–2540, 2007.

[19] Vishnu Navda, Aniruddha Bohra, Samrat Ganguly, and Dan Rubenstein. Using channel hopping to increase 802.11 resilience to jamming attacks. In *Proc. of Infocom '07*, pages 2526–2530, 2007.

[20] R. Negi and A. Perrig. Jamming analysis of MAC protocols. Technical report, Carnegie Mellon University, 2003.

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

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

[23] J. Schmidt, A. Siegel, and A. Srinivasan. Chernoff-Hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995.

[24] M. K. Simon, J. K. Omura, R. A. Schultz, and B. K. Levin. *Spread Spectrum Communications Handbook*. McGraw-Hill, 2001.

[25] David Thuente and Mithun Acharya. Intelligent jamming in wireless networks with applications to 802.11b and other networks. In *Proc. of MILCOM '06*, 2006.

[26] A.D. Wood, J.A. Stankovic, and G. Zhou. DEEJAM: Defeating energy-efficient jamming in IEEE 802.15.4-based wireless networks. In *Proc. of SECON '07*, 2007.

[27] W. Xu, K. Ma, W. Trappe, and Y. Zhang. Jamming sensor networks: attack and defense strategies. *IEEE Network*, 20(3):41–47, 2006.

[28] 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.

[29] 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*, page 80Ű89, 2004.