A Jamming-Resistant MAC Protocol for Multi-Hop Wireless Networks*

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

[1] Computer Science and Engineering, SCIDSE, Arizona State University
Tempe, AZ 85287, USA
{aricha,jzhang82}@asu.edu

[2] Department of Computer Science, University of Paderborn, D-33102 Paderborn, Germany
scheideler@upb.de

[3] Deutsche Telekom Laboratories, TU Berlin, D-10587 Berlin, Germany
stefan@net.t-labs.tu-berlin.de

Abstract. This paper presents a simple local medium access control protocol, called Jade, for multi-hop wireless networks with a single channel that is provably robust against adaptive adversarial jamming. The wireless network is modeled as a unit disk graph on a set of nodes distributed arbitrarily in the plane. In addition to these nodes, there are adversarial jammers that know the protocol and its entire history and that are allowed to jam the wireless channel at any node for an arbitrary $(1-\epsilon)$-fraction of the time steps, where $0 < \epsilon < 1$ is an arbitrary constant. We assume that the nodes cannot distinguish between jammed transmissions and collisions of regular messages. Nevertheless, we show that Jade achieves an asymptotically optimal throughput if there is a sufficiently dense distribution of nodes.

1 Introduction

The problem of coordinating the access to a shared medium is a central challenge in wireless networks. In order to solve this problem, a proper medium access control (MAC) protocol is needed. Ideally, such a protocol should not only be able to use the wireless medium as effectively as possible, but it should also be robust against attacks. Unfortunately, most of the MAC protocols today can be easily attacked. A particularly critical class of attacks are *jamming attacks* (i.e., denial-of-service attacks on the broadcast medium). Jamming attacks are typically easy to implement as the attacker does not need any special hardware. Attacks of this kind 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 MAC layer, where an adversary may either corrupt control packets or reserve the channel for the maximum allowable number of

* A full version of this article including all proofs appears on arXiv.org (identifier 1007.1189).

N.A. Lynch and A.A. Shvartsman (Eds.): DISC 2010, LNCS 6343, pp. 179–193, 2010.

slots so that other nodes experience low throughput by not being able to access the channel. In this paper we focus on jamming attacks at the physical layer, that is, the interference caused by the jammer will not allow the nodes to receive messages. The fundamental question that we are investigating is: *Is there a MAC protocol such that for any physical-layer jamming strategy, the protocol will still be able to achieve an asymptotically optimal throughput for the non-jammed time steps?* Such a protocol would *force* the jammer to jam all the time in order to prevent any successful message transmissions. Finding such a MAC protocol is not a trivial problem. In fact, the widely used IEEE 802.11 MAC protocol already fails to deliver any messages for very simple oblivious jammers that jam only a small fraction of the time steps [3]. On the positive side, Awerbuch et al. [2] have demonstrated that there are MAC protocols which are provably robust against even massive adaptive jamming, but their results only hold for single-hop wireless networks with a single jammer, that is, all nodes experience the same jamming sequence.

In this paper, we significantly extend the results in [2]. We present a MAC protocol called Jade (a short form of "jamming defense") that can achieve a constant fraction of the best possible throughput for a large class of jamming strategies in a large class of multi-hop networks where transmissions and interference can be modeled using unit-disk graphs. These jamming strategies include jamming patterns that can be completely different from node to node. It turns out that while Jade differs only slightly from the MAC protocol of [2], the proof techniques needed for the multi-hop setting significantly differ from the techniques in [2].

1.1 Model

We consider the problem of designing a robust MAC protocol for multi-hop wireless networks with a single wireless channel. The wireless network is modeled as a *unit disk graph* (UDG) $G = (V, E)$ where V represents a set of $n = |V|$ honest and reliable nodes and two nodes $u, v \in V$ are within each other's transmission range, i.e., $\{u, v\} \in E$, if and only if their (normalized) distance is at most 1. We assume that time proceeds in synchronous time steps called *rounds*. In each round, a node may either transmit a message or sense the channel, but it cannot do both. A node which is sensing the channel may either (*i*) sense an *idle* channel (if no other node in its transmission range is transmitting at that round and its channel is not jammed), (*ii*) sense a *busy* channel (if two or more nodes in its transmission range transmit at that round or its channel is jammed), or (*iii*) *receive* a packet (if exactly one node in its transmission range transmits at that round and its channel is not jammed).

In addition to these nodes there is an adversary (who may control any number of jamming devices). 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 round (i.e, the adversary is *adaptive*). However, like in [2], the adversary has to make a jamming decision *before* it knows the actions of the nodes at the current round. The adversary can jam the nodes individually at will, as long

as for every node v, at most a $(1-\epsilon)$-fraction of its rounds is jammed, where $\epsilon > 0$ can be an arbitrarily small constant. That is, v has the chance to receive a message in at least an ϵ-fraction of the rounds. More formally, an adversary is called $(T, 1-\epsilon)$-*bounded* for some $T \in \mathbb{N}$ and $0 < \epsilon < 1$, if for any time window of size $w \geq T$ and at any node v, the adversary can jam at most $(1-\epsilon)w$ of the w rounds at v.

Given a node v and a time interval I, we define $f_v(I)$ as the number of time steps in I that are non-jammed at v and $s_v(I)$ as the number of time steps in I in which v successfully receives a message. A MAC protocol is called *c-competitive* against some $(T, 1-\epsilon)$-bounded adversary if, for any time interval I with $|I| \geq K$ for a sufficiently large K (that may depend on T and n), $\sum_{v \in V} s_v(I) \geq c \cdot \sum_{v \in V} f_v(I)$. In other words, a c-competitive MAC protocol can achieve at least a c-fraction of the best possible throughput.

Our goal is to design a *symmetric local-control* MAC protocol (i.e., there is no central authority controlling the nodes, and all the nodes are executing the same protocol) that has a constant-competitive throughput against any $(T, 1-\epsilon)$-bounded adversary in any multi-hop network that can be modeled as a UDG. In order to obtain a more refined picture of the competitiveness of our protocol, we will also investigate so-called k-uniform adversaries. An adversary is *k-uniform* if the node set V can be partitioned into k subsets so that the jamming sequence is the same within each subset. In other words, we require that at all times, the nodes in a subset are either all jammed or all non-jammed. Thus, a 1-uniform jammer jams either everybody or nobody in a round whereas an n-uniform jammer can jam the nodes individually at will.

In this paper, we will say that a claim holds *with high probability (w.h.p.)* iff it holds with probability at least $1 - 1/n^c$ for any constant $c \geq 1$; it holds *with moderate probability (w.m.p.)* iff it holds with probability at least $1 - 1/(\log n)^c$ for any constant $c \geq 1$.

1.2 Related Work

Due to the topic's importance, wireless network jamming has been extensively studied in the applied research fields [1,5,6,22,26,27,28,30,31,37,38,39,40], both from the attacker's perspective [6,26,27,40] as well as from the defender's perspective [1,5,6,27,28,30,38,40]—also in multi-hop settings (e.g. [21,32,42,43,44]).

Traditionally, jamming defense mechanisms operate on the physical layer [28,30,36]. Mechanisms have been designed to *avoid* jamming as well as *detect* jamming. Spread spectrum technology has been shown to be very effective to avoid jamming as with widely spread signals, it becomes harder to detect the start of a packet quickly enough in order to jam it. Unfortunately, protocols such as IEEE 802.11b use relatively narrow spreading [20], and some other IEEE 802.11 variants spread signals by even smaller factors [5]. Therefore, a jammer that simultaneously blocks a small number of frequencies renders spread spectrum techniques useless in this case. As jamming strategies can come in many different flavors, detecting jamming activities by simple methods based on

signal strength, carrier sensing, or packet delivery ratios has turned out to be quite difficult [27].

Recent work has also studied *MAC layer strategies* against jamming, including coding strategies [6], channel surfing and spatial retreat [1,41], or mechanisms to hide messages from a jammer, evade its search, and reduce the impact of corrupted messages [38]. Unfortunately, these methods do not help against an adaptive jammer with *full* information about the history of the protocol, like the one considered in our work.

In the theory community, work on MAC protocols has mostly focused on efficiency. Many of these protocols are random backoff or tournament-based protocols [4,7,17,18,25,34] that do not take jamming activity into account and, in fact, are not robust against it (see [2] for more details). The same also holds for many MAC protocols that have been designed in the context of broadcasting [8] and clustering [24]. Also some work on jamming is known (e.g., [9] for a short overview). There are two basic approaches in the literature. The first assumes randomly corrupted messages (e.g. [33]), which is much easier to handle than adaptive adversarial jamming [3]. The second line of work either bounds the number of messages that the adversary can transmit or disrupt with a limited energy budget (e.g. [16,23]) or bounds the number of channels the adversary can jam (e.g. [10,11,12,13,14,15,29]).

The protocols in [16,23] can tackle adversarial jamming at both the MAC and network layers, where the adversary may not only be jamming the channel but also introducing malicious (fake) messages (possibly with address spoofing). However, they depend on the fact that the adversarial jamming budget is finite, so it is not clear whether the protocols would work under heavy continuous jamming. (The result in [16] seems to imply that a jamming rate of 1/2 is the limit whereas the handshaking mechanisms in [23] seem to require an even lower jamming rate.)

In the multi-channel version of the problem introduced in the theory community by Dolev [13] and also studied in [10,11,12,13,14,15,29], a node can only access one channel at a time, which results in protocols with a fairly large runtime (which can be exponential for deterministic protocols [11,14] and at least quadratic in the number of jammed channels for randomized protocols [12,29] if the adversary can jam almost all channels at a time). Recent work [10] also focuses on the wireless synchronization problem which requires devices to be activated at different times on a congested single-hop radio network to synchronize their round numbering while an adversary can disrupt a certain number of frequencies per round. Gilbert et al. [15] study robust information exchange in single-hop networks.

Our work is motivated by the work 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, but they only consider single-hop wireless networks. We go one step further by considering *multi-hop*

networks where different nodes can have different channel states at a time, e.g., a transmission may be received only by a fraction of the nodes. It turns out that while the MAC protocol of [2] can be adopted to the multi-hop setting with a small modification, the proof techniques cannot. We are not aware of any other theoretical work on MAC protocols for multi-hop networks with provable performance against adaptive jamming.

1.3 Our Contributions

In this paper, we present a robust MAC protocol called JADE. JADE is a fairly simple protocol: it is based on a very small set of assumptions and rules and has a minimal storage overhead. In fact, in JADE every node just stores a constant number of parameters, among them a fixed parameter γ that should be chosen so that the following main theorem holds:

Theorem 1. *When running* JADE *for at least* $\Omega((T \log n)/\epsilon + (\log n)^4/(\gamma\epsilon)^2)$ *time steps,* JADE *has a constant competitive throughput for any* $(T, 1-\epsilon)$*-bounded adversary and any UDG w.h.p. as long as* $\gamma = O(1/(\log T + \log\log n))$ *and (a) the adversary is 1-uniform and the UDG is connected, or (b) there are at least* $2/\epsilon$ *nodes within the transmission range of every node.*

Note that in practice, $\log T$ and $\log\log n$ are rather small so that our condition on γ is not too restrictive. Also, a conservative estimate on $\log T$ and $\log\log n$ would leave room for a superpolynomial change in n and a polynomial change in T over time.

On the other hand, we can also show the following result demonstrating that Theorem 1 essentially captures all the scenarios (within our notation) under which JADE can have a constant competitive throughput.

Theorem 2. *If (a) the UDG is not connected, or (b) the adversary is allowed to be 2-uniform and there are nodes with* $o(1/\epsilon)$ *nodes within their transmission range, then there are cases in which* JADE *is not constant competitive for any constant* c *independent of* ϵ*.*

Certainly, no MAC protocol can guarantee a constant competitive throughput if the UDG is not connected. However, it is still open whether there are simple MAC protocols that are constant competitive under non-uniform jamming strategies even if there are $o(1/\epsilon)$ nodes within the transmission range of a node.

2 Description of JADE

This section first gives a short motivation for our algorithmic approach and then presents the JADE protocol in detail.

2.1 Intuition

The intuition behind our MAC protocol is simple: in each round, each node u tries to send a message with probability p_u with $p_u \leq \hat{p}$ for some small constant $0 < \hat{p} < 1$. Consider the unit disk $D(u)$ around node u consisting of

u's neighboring nodes as well as u.[1] Moreover, let $N(u) = D(u) \setminus \{u\}$ and $p = \sum_{v \in N(u)} p_v$. Suppose that u is sensing the channel. Let q_0 be the probability that the channel is idle at u and let q_1 be the probability that exactly one node in $N(u)$ is sending a message. It holds that $q_0 = \prod_{v \in N(u)}(1 - p_v)$ and $q_1 = \sum_{v \in N(u)} p_v \prod_{w \in N(u) \setminus \{v\}}(1 - p_w)$. Hence,

$$q_1 \leq \sum_{v \in N(u)} p_v \frac{1}{1-\hat{p}} \prod_{w \in N(u)} (1-p_w) = \frac{q_0 \cdot p}{1-\hat{p}}, \quad q_1 \geq \sum_{v \in N(u)} p_v \prod_{w \in N(u)} (1-p_w) = q_0 \cdot p.$$

Thus we have the following lemma, which has also been derived in [2] for the single-hop case.

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

By Lemma 1, if a node v observes that the number of rounds in which the channel is idle is essentially equal to the number of rounds in which exactly one message is sent, then $p = \sum_{v \in N(v)} p_v$ is likely to be around 1 (if $\hat{p}$ is a sufficiently small constant), which would be ideal. Otherwise, the nodes know that they need to adapt their probabilities. Thus, 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 because at least two messages interfere with each other (see also [19] for a similar conclusion). Unfortunately, p can be very high for some reason, which requires a more sophisticated strategy for adjusting the access probabilities.

2.2 Protocol Description

In Jade, each node v maintains a probability value p_v, a threshold T_v and a counter c_v. The parameters $\hat{p}, \gamma > 0$ in the protocol are fixed and the same for every node. $\hat{p}$ may be set to any constant value so that $0 < \hat{p} \leq 1/24$, and γ should be small enough so that the condition in Theorem 1 is met.

Initially, every node v sets $T_v := 1$, $c_v := 1$ and $p_v := \hat{p}$. Afterwards, the Jade protocol works in synchronized rounds. In every round, each node v decides with probability p_v to send a message. If it decides not to send a message, it checks the following two conditions:

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

[1] In this paper, disks (and later sectors) will refer both to 2-dimensional areas in the plane as well as to the set of nodes in the respective areas. The exact meaning will become clear in the specific context.

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 round among the past T_v rounds in which v sensed a successful message transmission *or an idle channel*, then $p_v := (1+\gamma)^{-1} p_v$ and $T_v := \min\{T_v + 1, 2^{1/(4\gamma)}\}$.

As we will see in the upcoming section, the concept of using a multiplicative-increase-multiplicative-decrease mechanism for p_v and an additive-increase-additive-decrease mechanism for T_v, as well as the slight modifications of the protocol in [2], marked in italic above, are crucial for JADE to work.

3 Analysis of JADE

In contrast the description of JADE, its stochastic analysis is rather involved as it requires to shed light onto the complex interplay of the nodes all following their randomized protocol in a highly dependent manner. We first prove Theorem 1 (Sections 3.1 and 3.2) and then prove Theorem 2 (Section 3.3). In order to show the theorems, we will frequently use the following variant of the Chernoff bounds [2,35].

Lemma 2. *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

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

Throughout the section we assume that $\gamma = O(1/(\log T + \log\log n))$ is sufficiently small.

3.1 Proof of Theorem 1

We first look at a slightly weaker form of adversary. We say a round t is *open* for a node v if v and at least one other node in its neighborhood are non-jammed (which implies that v's neighborhood is non-empty). An adversary is *weakly* $(T, 1-\epsilon)$*-bounded* for some $T \in \mathbb{N}$ and $0 < \epsilon < 1$ if the adversary is $(T, 1-\epsilon)$-bounded and in addition to this, at least a constant fraction of the non-jammed rounds at each node are open in every time interval of size $w \geq T$.

Theorem 3. *When running* JADE *for $\Omega([T + (\log^3 n)/(\gamma^2 \epsilon)] \cdot (\log n)/\epsilon)$ rounds it holds w.h.p. that* JADE *is constant competitive for any weakly $(T, 1-\epsilon)$-bounded adversary.*

Proof. First, we focus on a *time frame* F consisting of $\alpha \log n/\epsilon$ *subframes* of size $f = \alpha[T + (\log^3 n)/(\gamma^2\epsilon)]$ each, where f is a multiple of T and α is a sufficiently large constant. The proof needs the following three lemmas. The first one is identical to Claim 2.5 in [2]. It is true because only successful message transmissions reduce T_u.

Lemma 3. *If in a time interval I the number of rounds in which a node u successfully receives a message is at most r, then u increases T_u in at most $r + \sqrt{2|I|}$ rounds in I.*

The second lemma holds for arbitrary (not just weakly) $(T, 1-\epsilon)$-bounded adversaries and will be shown in Section 3.2.

Lemma 4. *For every node u, $\sum_{v \in D(u)} p_v = O(1)$ for at least a $(1-\epsilon\beta)$-fraction of the rounds in time frame F, w.h.p., where the constant $\beta > 0$ can be made arbitrarily small.*

The third lemma just follows from some simple geometric argument.

Lemma 5. *A disk of radius 2 can be cut into at most 20 regions so that the distance between any two points in a region is at most 1.*

Consider some fixed node u. Let $J \subseteq F$ be the set of all non-jammed *open* rounds at u in time frame F (which are a constant fraction of the non-jammed rounds at u because we have a weakly $(T, 1-\epsilon)$-bounded adversary). Let p be a constant satisfying Lemma 4 (i.e., $\sum_{w \in D(v)} p_w \leq p$). Define $DD(u)$ to be the disk of radius 2 around u (i.e., it has twice the radius of $D(u)$). Cut $DD(u)$ into 20 regions $R_1, \ldots, R_{20}$ satisfying Lemma 5, and let v_i be any node in region R_i (if such a node exists), where $v_i = u$ if $u \in R_i$. According to Lemma 4 it holds for each i that at least a $(1-\epsilon\beta'/20)$-fraction of the rounds in F satisfy $\sum_{w \in D(v_i)} p_w \leq p$ for any constant $\beta' > 0$, w.h.p. Thus, at least a $(1-\epsilon\beta'')$-fraction of the rounds in F satisfy $\sum_{w \in D(v_i)} p_w \leq p$ for *every* i for any constant $\beta'' > 0$, w.h.p. As $D(v) \subseteq DD(u)$ for all $v \in D(u)$ and u has at least $\epsilon|F|$ non-jammed rounds in F, we get the following lemma, which also holds for arbitrary $(T, 1-\epsilon)$-bounded adversaries:

Lemma 6. *At least a $(1-\beta)$-fraction of the rounds in J satisfy $\sum_{v \in D(u)} p_v \leq p$ and $\sum_{w \in D(v)} p_w = O(p)$ for all nodes $v \in D(u)$ for any constant $\beta > 0$, w.h.p.*

Let us call these rounds *good*. Since the probability that u senses the channel is at least $1-\hat{p}$ and the probability that the channel at u is idle for $\sum_{w \in D(u)} p_w \leq p$ is equal to $\prod_{v \in N(u)}(1-p_v) \geq \prod_{v \in N(u)} e^{-2p_v} \geq e^{-2p}$, u senses an idle channel for at least $(1-\hat{p})(1-\beta)|J|e^{-2p} \geq 2\beta|J|$ many rounds in J on expectation if β is sufficiently small. This also holds w.h.p. when using the Chernoff bounds under the condition that at least $(1-\beta)|J|$ rounds in F are good (which also holds w.h.p.). Let k be the number of times u receives a message in F. We distinguish between two cases.

Case 1: $k \geq \beta|J|/6$. Then JADE is constant competitive for u and we are done.

Case 2: $k < \beta|J|/6$. Then we know from Lemma 3 that p_u is decreased at most $\beta|J|/6 + \sqrt{2|F|}$ times in F due to $c_u > T_u$. In addition to this, p_u is decreased at most $\beta|J|/6$ times in F due to a received message. On the other hand, p_u is increased at least $2\beta|J|$ times in J (if possible) due to an idle channel w.h.p. Also, we know from the JADE protocol that at the beginning of F, $p_u = \hat{p}$. Hence, there must be at least $\beta(2 - 1/6 - 1/6)|J| - \sqrt{2|F|} \geq (3/2)\beta|J|$ rounds in J w.h.p. at which $p_u = \hat{p}$. As there are at least $(1-\beta)|J|$ good rounds in J (w.h.p.), there are at least $\beta|J|/2$ good rounds in J w.h.p. in which $p_u = \hat{p}$. For these good rounds, u has a constant probability to transmit a message and every node $v \in D(u)$ has a constant probability of receiving it, so u successfully transmits $\Theta(|J|)$ messages to at least one of its non-jammed neighbors in F (on expectation and also w.h.p.).

If we charge 1/2 of each successfully transmitted message to the sender and 1/2 to the receiver, then a constant competitive throughput can be identified for every node in both cases above, so JADE is constant competitive in F.

It remains to show that Theorem 3 achieves constant competitiveness for any time interval exceeding $|F|$. First, note that all the proofs are valid as long as $\gamma \leq 1/[c(\log T + \log\log n)]$ for a constant $c \geq 2$, so we can increase T and thereby also $|F|$ as long as this inequality holds. So w.l.o.g. we may assume that $\gamma = 1/[2(\log T + \log\log n)]$. In this case, $2^{1/(4\gamma)} \leq \sqrt{|F|}$, so our rule of increasing T_v in JADE implies that $T_v \leq \sqrt{|F|}$ at any time, which is crucial for Lemma 4 to hold for a time frame starting at any time. This allows us to extend the competitive throughput result to any sequence of time frames, which completes the proof of Theorem 3. □

Now, let us consider the two cases of Theorem 1. Recall that we allow here any $(T, 1-\epsilon)$-bounded adversary and not just a weakly bounded.

Case 1: the adversary is 1-uniform and the UDG is connected. In this case, every node has a non-empty neighborhood and therefore *all* non-jammed rounds of the nodes are open. Hence, the conditions on a weakly $(T, 1-\epsilon)$-bounded adversary are satisfied. So Theorem 3 applies, which completes the proof of Theorem 1 a).

Case 2: $|D(v)| \geq 2/\epsilon$ for all $v \in V$. Consider some fixed time interval I with $|I|$ being a multiple of T. For every node $v \in D(u)$ let f_v be the number of non-jammed rounds at v in I and o_v be the number of open rounds at v in I. Let J be the set of rounds in I with at most one non-jammed node. Suppose that $|J| > (1-\epsilon/2)|I|$. Then every node in $D(u)$ must have more than $(\epsilon/2)|I|$ of its non-jammed rounds in J. As these non-jammed rounds must be serialized in J to satisfy our requirement on J, it holds that $|J| > \sum_{v\in D(u)}(\epsilon/2)|I| \geq (2/\epsilon)\cdot(\epsilon/2)|I| = |I|$. Since this is impossible, it must hold that $|J| \leq (1-\epsilon/2)|I|$.

Thus, $\sum_{v\in D(u)} o_v \geq (\sum_{v\in D(u)} f_v) - |J| \geq (1/2)\sum_{v\in D(u)} f_v$ because $\sum_{v\in D(u)} f_v \geq (2/\epsilon)\cdot\epsilon|I| = 2|I|$. Let $D'(u)$ be the set of nodes $v \in D(u)$

with $o_v \geq f_v/4$. That is, for each of these nodes, a constant fraction of the non-jammed time steps is open. Then $\sum_{v \in D(u) \setminus D'(u)} o_v < (1/4) \sum_{v \in D(u)} f_v$, so $\sum_{v \in D'(u)} o_v \geq (1/2) \sum_{v \in D(u)} o_v \geq (1/4) \sum_{v \in D(u)} f_v$.

Consider now a set $U \subseteq V$ of nodes so that $\bigcup_{u \in U} D(u) = V$ and for every $v \in V$ there are at most 6 nodes $u \in U$ with $v \in D(u)$ (U is easy to construct in a greedy fashion for arbitrary UDGs and also known as a *dominating set of constant density*). Let $V' = \bigcup_{u \in U} D'(u)$. Since $\sum_{v \in D'(u)} o_v \geq (1/4) \sum_{v \in D(u)} f_v$ for every node $u \in U$, it follows that $\sum_{v \in V'} o_v \geq (1/6) \sum_{u \in U} \sum_{v \in D'(u)} o_v \geq (1/24) \sum_{u \in U} \sum_{v \in D(u)} f_v \geq (1/24) \sum_{v \in V} f_v$. Using that together with Theorem 3, which implies that JADE is constant competitive w.r.t. the nodes in V', completes the proof of Theorem 1 b).

3.2 Proof of Lemma 4

In order to finish the proof of Theorem 1, it remains to prove Lemma 4. Consider any fixed node u. We partition u's unit disk $D(u)$ into six *sectors* of equal angles from u, $S_1, ..., S_6$. Note that all nodes within a sector S_i have distances of at most 1 from each other, so they can directly communicate with each other (in $D(u)$, distances can be up to 2). We will first explore properties of an arbitrary node in one sector, then consider the implications for a whole sector, and finally bound the cumulative sending probability in the entire unit disk.

Recall the definition of a time frame, a subframe and f in the proof of Theorem 3. Fix a sector S in $D(u)$ and consider some fixed time frame F. Let us refer to the sum of the probabilities of the neighboring nodes of a given node $v \in S$ by $\bar{p}_v := \sum_{w \in S \setminus \{v\}} p_w$. The following lemma shows that p_v will decrease dramatically if $\bar{p}_v$ is high throughout a certain time interval. It needs the fact that $max_v T_v \leq \sqrt{|F|}$ (not shown here).

Lemma 7. *Consider a node v in a unit disk $D(u)$. If $\bar{p}_v > 5 - \hat{p}$ during* all *rounds of a subframe I of F, then p_v will be at most $1/n^2$ at the end of I, w.h.p.*

We omit the proof here. Given this property of the individual probabilities, we can derive a bound for the cumulative probability of an entire sector S. In order to compute $p_S = \sum_{v \in S} p_v$, we introduce three thresholds, a low one, $\rho_{green} = 5$, one in the middle, $\rho_{yellow} = 5e$, and a high one, $\rho_{red} = 5e^2$. The following three lemmas provide some important insights about these probabilities. The proof of the second one is omitted here.

Lemma 8. *For any subframe I in F and any initial value of p_S in I there is at least one round in I with $p_S \leq \rho_{green}$ w.h.p.*

Proof. We prove the lemma by contradiction. Suppose that throughout the entire interval I, $p_S > \rho_{green}$. Then it holds for every node $v \in S$ that $\bar{p}_v > \rho_{green} - \hat{p}$ throughout I. In this case, however, we know from Lemma 7, that p_v will decrease to at most $1/n^2$ at the end of I w.h.p. Hence, all nodes $v \in S$ would decrease p_v to at most $1/n^2$ at the end of I w.h.p., which results in $p_S \leq 1/n$. This contradicts our assumption, so w.h.p. there must be a round t in I at which $p_S \leq \rho_{green}$. □

Lemma 9. *For any time interval I in F of size f and any sector S it holds that if $p_S \leq \rho_{green}$ at the beginning of I, then $p_S \leq \rho_{yellow}$ throughout I, w.m.p. Similarly, if $p_S \leq \rho_{yellow}$ at the beginning of I, then $p_S \leq \rho_{red}$ throughout I, w.m.p.*

Lemma 10. *For any subframe I in F it holds that if there has been at least one round during the past subframe where $p_S \leq \rho_{green}$, then throughout I, $p_S \leq \rho_{red}$ w.m.p.*

Proof. Suppose that there has been at least one round during the past subframe where $p_S \leq \rho_{green}$. Then we know from Lemma 9 that w.m.p. $p_S \leq \rho_{yellow}$ at the beginning of I. But if $p_S \leq \rho_{yellow}$ at the beginning of I, we also know from Lemma 9 that w.m.p. $p_S \leq \rho_{red}$ throughout I, which proves the lemma. □

Now, define a subframe I to be *good* if $p_S \leq \rho_{red}$ throughout I, and otherwise I is called *bad*. With the help of Lemma 8 and Lemma 10 we can prove the following lemma.

Lemma 11. *For any sector S, at most $\epsilon\beta/6$ of the subframes I in F are bad w.h.p., where the constant $\beta > 0$ can be made arbitrarily small depending on the constant α in f.*

From Lemma 11 it follows that apart from an $\epsilon\beta$-fraction of the subframes, all subframes I in F satisfy $\sum_{v \in D(u)} p_v \in O(1)$ throughout I, which completes the proof of Lemma 4.

3.3 Limitations of the Jade Protocol

One may ask whether a stronger throughput result than Theorem 1 can be shown. Ideally, we would like to use the following model. A MAC protocol is called *strongly c-competitive* against some $(T, 1-\epsilon)$-bounded adversary if, for any sufficiently large time interval and any node v, the number of rounds in which v successfully receives a message is at least a c-fraction of the total number of non-jammed rounds at v. In other words, a strongly c-competitive MAC protocol can achieve at least a c-fraction of the best possible throughput for every individual node. Unfortunately, such a protocol seems to be difficult to design. In fact, Jade is not strongly c-competitive for any constant $c > 0$, even if the node density is sufficiently high.

Theorem 4. *In general,* Jade *is not strongly c-competitive for a constant $c > 0$ if the adversary is allowed to be 2-uniform and $\epsilon \leq 1/3$.*

Proof. Suppose that (at some corner of the UDG) we have a set U of at least $1/\hat{p}$ nodes located closely to each other that are all within the transmission range of a node v. Initially, we assume that $\sum_{u \in U} p_u \geq 1$, $p_v = \hat{p}$ and $T_x = 1$ for all nodes $x \in U \cup \{v\}$. The time is partitioned into time intervals of size T. In each such time interval, called *T-interval*, the $(T, 1-\epsilon)$-bounded adversary jams all but the first ϵT rounds at U and all but the last ϵT rounds at v. It follows

directly from Section 2.3 of [2] that if $T = \Omega((\log^3 n)/(\gamma^2\epsilon))$, then for every node $u \in U$, $T_u \leq \alpha\sqrt{T \log n/\epsilon}$ w.h.p. for some sufficiently large constant α. Thus, $T_u \leq \gamma T/(\beta \log n)$ w.h.p. for any constant $\beta > 0$ if T is sufficiently large. Hence, between the last non-jammed round at U and the first non-jammed round at v in a T-interval, the values T_u are increased (and the values p_u are decreased) at least $\beta(\log n)/(6\gamma)$ times. Thus, at the first non-jammed round at v, it holds for every $u \in U$ that

$$p_u \leq \hat{p} \cdot (1+\gamma)^{-\beta(\log n)/(6\gamma)} \leq \hat{p} \cdot e^{-(\beta/6)\log n} \leq 1/n^{\beta/6}$$

and, therefore, $\sum_{u \in U} p_u = O(1/n^2)$ if $\beta \geq 18$. This cumulative probability will stay that low during all of v's non-jammed rounds as during these rounds the nodes in U are jammed. Hence, the probability that v receives any message during its non-jammed rounds of a T-interval is $O(1/n^2)$, so Jade is not c-competitive for v for any constant $c > 0$. □

Also, in our original model, Jade is not constant competitive if the node density is too low.

Theorem 5. *In general,* Jade *is not c-competitive for a constant c independent of* ϵ *if there are nodes u with* $|D(u)| = o(1/\epsilon)$ *and the adversary is allowed to be 2-uniform.*

Proof. Suppose that we have a set U of $k = o(1/\epsilon)$ nodes located closely to each other that are all within the transmission range of a node v. Let $T = \Omega((\log^3 n)/(\gamma^2\epsilon))$. In each T-interval, the adversary never jams v but jams all but the first ϵT rounds at U. Then Section 2.3 of [2] implies that for every node $u \in U$, $T_u \leq \gamma T/(\beta \log n)$ w.h.p. for any constant $\beta > 0$ if T is sufficiently large. The nodes in U continuously increase their T_u-values and thereby reduce their p_u values during their jammed time steps. Hence, the nodes in $U \cup \{v\}$ will receive at most $\epsilon T \cdot |U| + (\epsilon T + O(T/\log n)) = \epsilon T \cdot o(1/\epsilon) + (\epsilon + o(1))T = (\epsilon + o(1))T$ messages in each T-interval on expectation whereas the sum of non-jammed rounds over all nodes is more than T. □

This implies Theorem 2. Hence, Theorem 1 is essentially the best one can show for Jade (within our notation).

3.4 Simulations

We briefly report on some simulation results that complement the theoretical insights. We assume that initially, $p_v = \hat{p} = 1/24$ for all nodes v. The nodes are distributed over a square plane of 4×4 units, and are connected in a unit disk graph manner (multi-hop). In each round, a node is jammed independently with probability $(1 - \epsilon)$. We run the simulation for a sufficiently large number of time steps indicated by the Theorem 1, where $\epsilon = 0.3$, $T = 200$, and $\gamma = 0.1$. Figure 1 (*left*) shows the throughput competitiveness of Jade as a function of the network size for a scenario with a uniform node distribution and a scenario

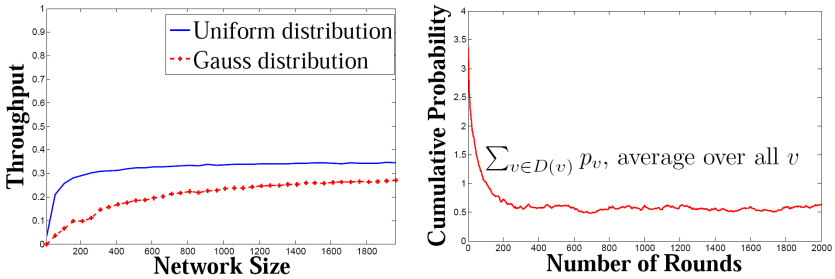

Fig. 1. *Left:* Throughput as a function of network size. *Right:* Convergence behavior for multi-hop networks (uniform distribution). For the plot, we used $n = 500$.

with a normal/Gaussian distribution. In both cases, the throughput is larger when the density is higher (20% to 40%), which corresponds to our formal insight that a constant competitive throughput is possible only if the node density exceeds a certain threshold. Moreover, we found that a constant throughput and a constant cumulative sending probability (per unit disk) is reached fast. See Figure 1 (*right*).

4 Conclusion

This paper has presented the first jamming-resistant MAC protocol with provably good performance in multi-hop networks exposed to an adaptive adversary. While we have focused on unit disk graphs, we believe that our stochastic analysis is also useful for more realistic wireless network models. Moreover, although our analysis is involved, our protocol is rather simple.

References

1. Alnifie, G., Simon, R.: A multi-channel defense against jamming attacks in wireless sensor networks. In: Proc. of Q2SWinet '07, pp. 95–104 (2007)
2. Awerbuch, B., Richa, A., Scheideler, C.: A jamming-resistant MAC protocol for single-hop wireless networks. In: Proc. of PODC '08 (2008)
3. Bayraktaroglu, E., King, C., Liu, X., Noubir, G., Rajaraman, R., Thapa, B.: On the performance of IEEE 802.11 under jamming. In: Proc. of IEEE Infocom '08, pp. 1265–1273 (2008)
4. Bender, M.A., Farach-Colton, M., He, S., Kuszmaul, B.C., Leiserson, C.E.: Adversarial contention resolution for simple channels. In: Proc. of SPAA '05 (2005)
5. Brown, T., James, J., Sethi, A.: Jamming and sensing of encrypted wireless ad hoc networks. In: Proc. of MobiHoc '06, pp. 120–130 (2006)
6. Chiang, J., Hu, Y.-C.: Cross-layer jamming detection and mitigation in wireless broadcast networks. In: Proc. of MobiCom '07, pp. 346–349 (2007)
7. Chlebus, B.S., Kowalski, D.R., Rokicki, M.A.: Adversarial queuing on the multiple-access channel. In: Proc. of PODC '06 (2006)

8. Czumaj, A., Rytter, W.: Broadcasting algorithms in radio networks with unknown topology. Journal of Algorithms 60(2), 115 (2006)
9. Dolev, S., Gilbert, S., Guerraoui, R., Kowalski, D., Newport, C., Kuhn, F., Lynch, N.: Reliable distributed computing on unreliable radio channels. In: Proc. 2009 MobiHoc S3 Workshop (2009)
10. Dolev, S., Gilbert, S., Guerraoui, R., Kuhn, F., Newport, C.C.: The wireless synchronization problem. In: Proc. 28th Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 190–199 (2009)
11. Dolev, S., Gilbert, S., Guerraoui, R., Newport, C.: Gossiping in a multi-channel radio network: An oblivious approach to coping with malicious interference. In: Pelc, A. (ed.) DISC 2007. LNCS, vol. 4731, pp. 208–222. Springer, Heidelberg (2007)
12. Dolev, S., Gilbert, S., Guerraoui, R., Newport, C.: Secure communication over radio channels. In: Proc. 27th ACM Symposium on Principles of Distributed Computing (PODC), pp. 105–114 (2008)
13. Dolev, S., Gilbert, S., Guerraoui, R., Newport, C.C.: Gossiping in a multi-channel radio network. In: Pelc, A. (ed.) DISC 2007. LNCS, vol. 4731, pp. 208–222. Springer, Heidelberg (2007)
14. Gilbert, S., Guerraoui, R., Kowalski, D., Newport, C.: Interference-resilient information exchange. In: Proc. of the 28th Conference on Computer Communications, IEEE Infocom 2009 (2009)
15. Gilbert, S., Guerraoui, R., Kowalski, D.R., Newport, C.C.: Interference-resilient information exchange. In: Proc. 28th IEEE International Conference on Computer Communications (INFOCOM), pp. 2249–2257 (2009)
16. Gilbert, S., Guerraoui, R., Newport, C.: Of malicious motes and suspicious sensors: On the efficiency of malicious interference in wireless networks. In: Shvartsman, M.M.A.A. (ed.) OPODIS 2006. LNCS, vol. 4305, pp. 215–229. Springer, Heidelberg (2006)
17. Goldberg, L.A., Mackenzie, P.D., Paterson, M., Srinivasan, A.: Contention resolution with constant expected delay. J. ACM 47(6) (2000)
18. Hastad, J., Leighton, T., Rogoff, B.: Analysis of backoff protocols for mulitiple accesschannels. SIAM Journal on Computing 25(4) (1996)
19. Heusse, M., Rousseau, F., Guillier, R., Duda, A.: Idle sense: An optimal access method for high throughput and fairness in rate diverse wireless lans. In: Proc. SIGCOMM (2005)
20. IEEE: Medium access control (MAC) and physical specifications. In: IEEE P802.11/D10 (1999)
21. Jain, K., Padhye, J., Padmanabhan, V.N., Qiu, L.: Impact of interference on multi-hop wireless network performance. In: Proc. 9th Annual International Conference on Mobile Computing and Networking (MobiCom), pp. 66–80 (2003)
22. Jiang, S., Xue, Y.: Providing survivability against jamming attack via joint dynamic routing and channel assigment. In: Proc. 7th Workshop on Design of Reliable Communication Networks, DRCN (2009)
23. Koo, C.Y., Bhandari, V., Katz, J., Vaidya, N.H.: Reliable broadcast in radio networks: The bounded collision case. In: Proc. of PODC '06 (2006)
24. Kuhn, F., Moscibroda, T., Wattenhofer, R.: Radio network clustering from scratch. In: Albers, S., Radzik, T. (eds.) ESA 2004. LNCS, vol. 3221, Springer, Heidelberg (2004)
25. Kwak, B.-J., Song, N.-O., Miller, L.E.: Performance analysis of exponential backoff. IEEE/ACM Transactions on Networking 13(2), 343–355 (2005)

26. Law, Y., van Hoesel, L., Doumen, J., Hartel, P., Havinga, P.: Energy-efficient link-layer jamming attacks against wireless sensor network mac protocols. In: Proc. of SASN '05, pp. 76–88 (2005)
27. Li, M., Koutsopoulos, I., Poovendran, R.: Optimal jamming attacks and network defense policies in wireless sensor networks. In: Proc. of Infocom '07, pp. 1307–1315 (2007)
28. Liu, X., Noubir, G., Sundaram, R., Tan, S.: Spread: Foiling smart jammers using multi-layer agility. In: Proc. of Infocom '07, pp. 2536–2540 (2007)
29. Meier, D., Pignolet, Y.A., Schmid, S., Wattenhofer, R.: Speed dating despite jammers. In: Krishnamachari, B., Suri, S., Heinzelman, W., Mitra, U. (eds.) DCOSS 2009. LNCS, vol. 5516, Springer, Heidelberg (2009)
30. Navda, V., Bohra, A., Ganguly, S., Rubenstein, D.: Using channel hopping to increase 802.11 resilience to jamming attacks. In: Proc. of Infocom '07, pp. 2526–2530 (2007)
31. Negi, R., Perrig, A.: Jamming analysis of MAC protocols. Technical report, Carnegie Mellon University (2003)
32. Noubir, G.: On connectivity in ad hoc networks under jamming using directional antennas and mobility. In: Langendoerfer, P., Liu, M., Matta, I., Tsaoussidis, V. (eds.) WWIC 2004. LNCS, vol. 2957, pp. 186–200. Springer, Heidelberg (2004)
33. Pelc, A., Peleg, D.: Feasibility and complexity of broadcasting with random transmission failures. In: Proc. of PODC '05 (2005)
34. Raghavan, P., Upfal, E.: Stochastic contention resolution with short delays. SIAM Journal on Computing 28(2), 709–719 (1999)
35. Schmidt, J., Siegel, A., Srinivasan, A.: Chernoff-Hoeffding bounds for applications with limited independence. SIAM Journal on Discrete Mathematics 8(2), 223–250 (1995)
36. Simon, M.K., Omura, J.K., Schultz, R.A., Levin, B.K.: Spread Spectrum Communications Handbook. McGraw-Hill, New York (2001)
37. Thuente, D., Acharya, M.: Intelligent jamming in wireless networks with applications to 802.11b and other networks. In: Proc. of MILCOM '06 (2006)
38. Wood, A., Stankovic, J., Zhou, G.: DEEJAM: Defeating energy-efficient jamming in IEEE 802.15.4-based wireless networks. In: Proc. of SECON '07 (2007)
39. Xu, W., Ma, K., Trappe, W., Zhang, Y.: Jamming sensor networks: attack and defense strategies. IEEE Network 20(3), 41–47 (2006)
40. Xu, W., Trappe, W., Zhang, Y., Wood, T.: The feasibility of launching and detecting jamming attacks in wireless networks. In: Proc. of MobiHoc '05, pp. 46–57 (2005)
41. Xu, W., Wood, T., Zhang, Y.: Channel surfing and spatial retreats: defenses against wireless denial of service. In: Proc. of Workshop on Wireless Security (2004)
42. Ye, S., Wang, Y., Tseng, Y.: A jamming-based MAC protocol for wireless multihop ad hoc networks. In: Proc. IEEE 58th Vehicular Technology Conference (2003)
43. Ye, S.-R., Wang, Y.-C., Tseng, Y.-C.: A jamming-based MAC protocol to improve the performance of wireless multihop ad-hoc networks. Wirel. Commun. Mob. Comput. 4(1), 75–84 (2004)
44. Zander, J.: Jamming in slotted ALOHA multihop packed radio networks. IEEE Transactions on Networking 39(10), 1525–1531 (1991)