

Competitive Throughput in Multi-Hop Wireless Networks Despite Adaptive Jamming

Andrea Richa, Christian Scheideler, Stefan Schmid, Jin Zhang

Received: date / Accepted: June 6, 2012

Abstract This article 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 nodes can perform collision detection (unless they are transmitting themselves), but that they cannot distinguish between jammed transmissions and collisions of regular

messages. Nevertheless, we show that JADE achieves an asymptotically optimal throughput by efficiently exploiting the unpredictable time periods in which the medium is available.

Keywords medium access, jamming, distributed algorithms, randomization

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 slots so that other nodes experience low throughput by not being able to access the channel. In this article 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*

A preliminary version of this article appeared at the 24th International Symposium on Distributed Computing (DISC), 2010. Research is partly supported by NSF grants CCF-1116368 and CCF-0830704, as well as by DFG grants DFG SFB 901 and DFG SCHE 1592/2-1.

Andrea Richa
Computer Science and Engineering, SCIDSE
Arizona State University
Tempe, AZ 85287, USA
E-mail: aricha@asu.edu

Christian Scheideler
Universität Paderborn
33102 Paderborn, Germany
E-mail: scheideler@upb.de

Stefan Schmid
Telekom Innovation Laboratories (T-Labs)
Technische Universität Berlin
10587 Berlin, Germany
E-mail: stefan@net.t-labs.tu-berlin.de

Jin Zhang
Computer Science and Engineering, SCIDSE
Arizona State University
Tempe, AZ 85287, USA
E-mail: jzhang82@asu.edu

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 article, 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. Moreover, we assume that a (receiving) node can detect collisions. Concretely, 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 (controlling 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 ($\epsilon > 0$ can be an arbitrarily small constant independent of n), among which at least an arbitrary constant fraction are *open*: 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). More formally, an adversary is $(T, 1 - \epsilon)$ -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 , and at least an arbitrary constant fraction of the non-jammed rounds at v are *open* in every time interval of size w . We, later in this paper, will also consider a stronger adversary that does not have the limitation of providing open rounds. While if not stated otherwise and by default, we will always refer to the adversary defined here, we will sometime explicitly use the adjective *weak* to distinguish this adversary from the stronger variant.

Given a node v and a time interval I of size $|I|$ (in terms of rounds), 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 (i.e., a *c-competitive* throughput where c does not depend on n) against any $(T, 1 - \epsilon)$ -bounded adversary in any multi-hop network that can be modeled as a UDG. Not only the nodes are distributed in space in our model, but also the adversary. Concretely, we introduce the concept of a *k-uniform adversary*, an adversary that can jam different nodes at different times. 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. Note that the adversary must hence not necessarily be geometrically constrained.

As already mentioned, this article also discusses a stronger adversary: we say that a *strong* adversary is $(T, 1 - \epsilon)$ -bounded, 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 , where $T \in \mathbb{N}$ and $0 < \epsilon < 1$. Note that this adversary is stronger as we only guarantee that an ϵ -fraction of the rounds at v are non-jammed, but not that during these rounds there exists at least one neighbor free to receive a message from v . While the nodes do not know ϵ , we do allow them to have a very rough upper bound of the values n and T .

Finally, let us emphasize that our notion of throughput is constrained to Layer 2 (the MAC layer), and measures the number of successful transmissions over “links”, i.e., pairs of nodes. That is, assuming a backlogged situation where packets are constantly submitted to the medium access layer from higher layers, we can schedule transmissions over Layer 2 links efficiently. In contrast to other throughput models in literature (e.g., [40]), we explicitly consider the receiver-side which we believe is much more meaningful: in a broadcast medium and in a distributed setting, the throughput computed by focusing on the sender only can be misleading as simply sending a packet out does not imply that it is also received (and by how many nodes). However, also note that a (MAC layer) link-based throughput does not imply any minimal end-to-end throughput between remote nodes on higher layers, e.g., on the transport layer (especially when using TCP with its flow and congestion control mechanisms), or throughput of flows. Moreover, note that we do not model any retransmissions that would happen on higher layers. Indeed, our MAC protocol has the nice property that it does not rely on any acknowledgements on the MAC layer to guarantee the throughput, and assumes that retransmission mechanisms are in place on higher layers.

In this article, 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, 26–28, 30, 31, 39, 41–43], both from the attacker’s perspective [6, 26, 27, 43] as well as from the defender’s perspective [1, 5, 6, 27, 28, 30, 41, 43]—also in multi-hop settings (e.g. [21, 32, 45–47]).

Traditionally, jamming defense mechanisms operate on the physical layer [28, 30, 38]. 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, 44], or mechanisms to hide messages from a jammer, evade its search, and reduce the impact of corrupted messages [41]. 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. [10–15, 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 [10–15, 29], a node can only access one chan-

nel 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. (Their approach has later been extended to reactive jamming environments [35] and applications such as leader election [36].)

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 article, we present a robust MAC protocol called JADE. JADE is a fairly simple protocol: it is based on a small set of rules and assumptions (e.g., collision detection at receivers), and has a minimal storage overhead. We can prove the following main theorem:

Theorem 1 *When running JADE for $\Omega([T + (\log^3 n)/(\gamma^2 \epsilon)] \cdot (\log n)/\epsilon)$ rounds it holds w.h.p. that JADE achieves a constant competitive throughput (i.e., independent of n) for any $(T, 1 - \epsilon)$ -bounded (weak) adversary, where n is the total number of nodes and $\gamma \in O(1/(\log T + \log \log n))$ is a parameter.*

Since $\log T$ and $\log \log n$ are small the assumption on γ is not too restrictive: 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. Also note that the (unrealistic and non-scalable) assumption that the nodes know constant factor approximations of n or T directly would render the prob-

lem trivial. (Whether a competitive MAC protocol exists without any assumptions on the magnitude of these parameters is an open question. We conjecture no such algorithm exists.)

Regarding the strong adversary, we can show constant throughput only if one of the conditions in Theorem 2 is satisfied.

Theorem 2 *When running JADE for $\Omega((T \log n)/\epsilon + (\log n)^4/(\gamma \epsilon)^2)$ rounds, JADE has a constant competitive throughput against any strong adversary that is $(T, 1 - \epsilon)$ -bounded and in any UDG w.h.p., as long as (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.*

In Section 3.4, we show that Theorem 2 captures all the scenarios for which JADE can have a constant competitive throughput under a strong adversary.

Concretely, we will show the following limitations under a strong adversary.

Theorem 3 *In general, JADE is not strongly c -competitive for a constant $c > 0$ (independent of n) if the strong adversary is allowed to be 2-uniform and $\epsilon \leq 1/3$. Moreover, JADE is also not c -competitive for a constant c if there are nodes u with $|D(u)| = o(1/\epsilon)$ and the strong adversary is allowed to be 2-uniform.*

Here, *strongly c -competitive* refers to a stronger throughput model where we require that 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 .

1.4 Article Organization

The remainder of this article is organized as follows. Section 2 presents our MAC protocol, and the formal analysis is given in Section 3. Section 4 reports on the simulation results. The article is concluded in Section 5.

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: Each node u maintains a parameter p_u which describes u 's

probability of accessing the channel at a given moment of time. That is, in each round, each node u decides to broadcast a packet with probability p_u . (This is similar to classic random backoff mechanisms where the next transmission time t is chosen uniformly at random from an interval of size $1/p_u$.) The nodes adapt and synchronize their p_u values over time in a multiplicative increase multiplicative decrease manner, i.e., the value is lowered in times of high interference or increased during times where the channel is idling. However, p_u will never exceed \hat{p} , for some 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 $p = p(u) = \sum_{v \in N(u)} p_v$; henceforth, when u is clear from the context, we will often simply write p instead of $p(u)$. 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,

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

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

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 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 (e.g., due to high initial sending probabilities), which requires a more sophisticated strategy for adjusting the access probabilities.

¹ In this article, 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.

2.2 Protocol Description

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

Initially, every node v sets $c_v := 1$ and $p_v := \hat{p}$. Note however that while we provide some initial values for the variables in our description, our protocol is self-stabilizing and works for *any* initial variable values, as we will show in our proofs.

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 $T_v := \max\{T_v - 1, 1\}$.

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. If in the *Afterwards* part of the algorithm we did not include the “idle” condition, in a distributed setting, it could happen that a center node u with high p_u value will never see any successful transmissions: This can happen if the cumulative probability of u 's neighbors is low while the cumulative probability of u 's neighbors' neighbors is high, which makes u 's neighbors always stay at low probabilities. In this situation, T_u may increase arbitrarily. Such high T_u values however are harmful to the fast recovery properties of the protocol.

Our simulation study showed that this protocol change is crucial also in the average case, and not just in artificial scenarios.

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 dependent manner. We first prove Theorem 1 in Sections 3.1 and 3.2, and then derive Theorem 2 in Section 3.3. The limitations of JADE under the strong adversary are discussed in Section 3.4. In order to show the theorems, we will frequently use the following variant of the Chernoff bounds [2,37].

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

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

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

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

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

3.1 Proof of Theorem 1

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 following lemma even holds for a *strong* adversary 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 following lemma follows from simple geometric arguments.

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). 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, \dots, 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 1 also holds for larger time intervals than $|F|$. 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. This allows us to extend the competitive throughput result to any sequence of time frames. Let $J \subset l \cdot F$ be the set of all non-jammed open rounds at u overall time frames, where l is the number of frames considered here. Hence, Case 1 holds directly; as for Case 2, we have $\beta(2 - 1/6 - 1/6)|J| - \sqrt{2l|F|} \geq (3/2)\beta|J|$ rounds in J w.h.p. at which $p_u = \hat{p}$. Hence, the rest of the proof follows directly, which completes the proof of Theorem 1.

3.2 Proof of Lemma 4

This section is dedicated to the proof of Lemma 4 which is rather involved. Consider any fixed node u . We partition u 's unit disk $D(u)$ into six *sectors* of equal angles from u , S_1, \dots, 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 one another (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 1. 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.

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.*

Proof We say that a round is *useful* for node v if from v 's perspective there is an idle channel or a successful

transmission at that round (when ignoring the action of v); otherwise the round is called *non-useful*. Note that in a non-useful round, according to our protocol, p_v will either decrease (if the threshold T_v is exceeded) or remain the same. On the other hand, in a *useful* round, p_v will increase (if v senses an idle channel), decrease (if v senses a successful transmission) or remain the same (if v sends a message). Hence, p_v can only increase during useful rounds of I . Let \mathcal{U} be the set of useful rounds in I for our node v . We distinguish between two cases, depending on the cardinality $|\mathcal{U}|$. In the following, let $p_v(0)$ denote the probability of v at the beginning of I (which is at most \hat{p}). Suppose that $f \geq 2[(3c \ln n)/\gamma]^2$ for a sufficiently large constant c . (This lower bound coincides with our definition of f in the proof of Theorem 1.)

Case 1: Suppose that $|\mathcal{U}| < (c \ln n)/\gamma$, that is, many rounds are blocked and p_v can increase only rarely. As there are at least $(3c \ln n)/\gamma$ occasions in I in which $c_v > T_v$ and $|\mathcal{U}| < (c \ln n)/\gamma$, in at least $(2c \ln n)/\gamma$ of these occasions v only saw blocked channels for T_v consecutive rounds and therefore decides to increase T_v and decrease p_v . Hence, at the end of I ,

$$\begin{aligned} p_v &\leq (1 + \gamma)^{|\mathcal{U}| - 2c \ln n / \gamma} p_v(0) \\ &\leq (1 + \gamma)^{-c \ln n / \gamma} p_v(0) \\ &\leq e^{-c \ln n} = 1/n^c. \end{aligned}$$

Case 2: Next, suppose that $|\mathcal{U}| \geq (c \ln n)/\gamma$. We will show that many of these useful rounds will be successful such that p_v decreases. Since $p_v \leq \hat{p} \leq 1/24$ throughout I , it follows from the Chernoff bounds that w.h.p. v will sense the channel for at least a fraction of $2/3$ of the useful rounds w.h.p. Let this set of useful rounds be called \mathcal{U}' . Consider any round $t \in \mathcal{U}'$. Let q_0 be the probability that there is an idle channel at round t and q_1 be the probability that there is a successful transmission at t . It holds that $q_0 + q_1 = 1$. From Lemma 1 we also know that $q_1 \geq q_0 \cdot \bar{p}_v$. Since $\bar{p}_v > 5 - \hat{p}$ for all rounds in I , it follows that $q_1 \geq 4/5$ for every round in \mathcal{U}' . Thus, it follows from the Chernoff bounds that for at least $2/3$ of the rounds in \mathcal{U}' , v will sense a successful transmission w.h.p. Hence, at the end of I it holds w.h.p. that

$$\begin{aligned} p_v &\leq (1 + \gamma)^{-(1/3) \cdot |\mathcal{U}'|} p_v(0) \\ &\leq (1 + \gamma)^{-(1/3) \cdot (2c/3) \ln n / \gamma} p_v(0) \\ &\leq e^{-(2c/9) \ln n} = 1/n^{2c/9}. \end{aligned}$$

Combining the two cases with $c \geq 9$ results in the lemma. \square

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.

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 $p_S \leq \rho_{green}$. \square

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.*

Proof It suffices to prove the lemma for the case that initially $p_S \leq \rho_{green}$ as the other case is analogous. Consider some fixed round t in I . Let p_S be the cumulative probability at the beginning of t and p'_S be the cumulative probability at the end of t . Moreover, let $p_S^{(0)}$ denote the cumulative probability of the nodes $w \in S$ with no transmitting node in $D(w) \setminus S$ in round t . Similarly, let $p_S^{(1)}$ denote the cumulative probability of the nodes $w \in S$ with a single transmitting node in $D(w) \setminus S$, and let $p_S^{(2)}$ be the cumulative probability of the nodes $w \in S$ that experience a blocked round either because they are jammed or at least two nodes in $D(w) \setminus S$ are transmitting at t . Certainly, $p_S = p_S^{(0)} + p_S^{(1)} + p_S^{(2)}$. Our goal is to determine p'_S in this case. Let $q_0(S)$ be the probability that all nodes in S stay silent, $q_1(S)$ be the probability that exactly one node in S is transmitting, and $q_2(S) = 1 - q_0(S) - q_1(S)$ be the probability that at least two nodes in S are transmitting.

When ignoring the case that $c_v > T_v$ for a node $v \in S$ at round t , it holds:

$$\begin{aligned} \mathbb{E}[p'_S] &= q_0(S) \cdot [(1 + \gamma)p_S^{(0)} + (1 + \gamma)^{-1}p_S^{(1)} + p_S^{(2)}] \\ &\quad + q_1(S) \cdot [(1 + \gamma)^{-1}p_S^{(0)} + p_S^{(1)} + p_S^{(2)}] \\ &\quad + q_2(S) \cdot [p_S^{(0)} + p_S^{(1)} + p_S^{(2)}] \end{aligned}$$

This is certainly also an upper bound for $\mathbb{E}[p'_S]$ if $c_v > T_v$ for a node $v \in S$ because p_v will never be increased (but possibly decreased) in this case. Now, consider the event E_2 that at least two nodes in S transmit

a message. If E_2 holds, then $\mathbb{E}[p'_S] = p'_S = p_S$, so there is no change in the system. On the other hand, assume that E_2 does not hold. Let $q'_0(S) = q_0(S)/(1 - q_2(S))$ and $q'_1(S) = q_1(S)/(1 - q_2(S))$ be the probabilities $q_0(S)$ and $q_1(S)$ under the condition of $\neg E_2$. Then we distinguish between three cases.

Case 1: $p_S^{(0)} = p_S$. Then

$$\begin{aligned} \mathbb{E}[p'_S] &\leq q'_0(S) \cdot (1 + \gamma)p_S + q'_1(S) \cdot (1 + \gamma)^{-1}p_S \\ &= ((1 + \gamma)q'_0(S) + (1 + \gamma)^{-1}q'_1(S))p_S. \end{aligned}$$

From Lemma 1 we know that $q_0(S) \leq q_1(S)/p_S$, so $q'_0(S) \leq q'_1(S)/p_S$. If $p_S \geq \rho_{green}$, then $q'_0(S) \leq q'_1(S)/5$. Hence,

$$\mathbb{E}[p'_S] \leq ((1 + \gamma)/6 + (1 + \gamma)^{-1}5/6)p_S \leq (1 + \gamma)^{-1/2}p_S$$

since $\gamma = o(1)$. On the other hand, $p'_S \leq (1 + \gamma)p_S$ in any case.

Case 2: $p_S^{(1)} = p_S$. Then

$$\begin{aligned} \mathbb{E}[p'_S] &\leq q'_0(S) \cdot (1 + \gamma)^{-1}p_S + q'_1(S)p_S \\ &= (q'_0(S)/(1 + \gamma) + (1 - q'_0(S)))p_S \\ &= (1 - q'_0(S)\gamma/(1 + \gamma))p_S. \end{aligned}$$

Now, it holds that $1 - x\gamma/(1 + \gamma) \leq (1 + \gamma)^{-x/2}$ for all $x \in [0, 1]$ because from the Taylor series of e^x and $\ln(1 + x)$ it follows that

$$(1 + \gamma)^{-x/2} \geq 1 - (x \ln(1 + \gamma))/2 \geq 1 - (x(1 - \gamma/2)\gamma)/2$$

and

$$1 - x\gamma/(1 + \gamma) \leq 1 - (x(1 - \gamma/2)\gamma)/2$$

for all $x, \gamma \in [0, 1]$ as is easy to check. Therefore, when defining $\varphi = q'_0(S)$, we get $\mathbb{E}[p'_S] \leq (1 + \gamma)^{-\varphi/2}p_S$. On the other hand, $p'_S \leq p_S \leq (1 + \gamma)^\varphi p_S$.

Case 3: $p_S^{(2)} = p_S$. Then for $\varphi = 0$, $\mathbb{E}[p'_S] \leq p_S = (1 + \gamma)^{-\varphi/2}p_S$ and $p'_S \leq p_S = (1 + \gamma)^\varphi p_S$.

Combining the three cases and taking into account that $p_S^{(0)} + p_S^{(1)} + p_S^{(2)} = p_S$, we obtain the following result.

Lemma 10 *There is a $\phi \in [0, 1]$ (depending on $p_S^{(0)}$, $p_S^{(1)}$ and $p_S^{(2)}$) so that*

$$\mathbb{E}[p'_S] \leq (1 + \gamma)^{-\phi}p_S \quad \text{and} \quad p'_S \leq (1 + \gamma)^{2\phi}p_S. \quad (1)$$

Proof Let $a = (1 + \gamma)^{1/2}$, $b = (1 + \gamma)^{\varphi/2}$ for the φ defined in Case 2, and $c = 1$. Furthermore, let $x_0 = p_S^{(0)}/p_S$, $x_1 = p_S^{(1)}/p_S$ and $x_2 = p_S^{(2)}/p_S$. Define $\phi = -\log_{1+\gamma}((1/a)x_0 + (1/b)x_1 + (1/c)x_2)$. Then we have

$$\begin{aligned} \mathbb{E}[p'_S] &\leq (1 + \gamma)^{-1/2}p_S^{(0)} + (1 + \gamma)^{-\varphi/2}p_S^{(1)} + p_S^{(2)} \\ &= (1 + \gamma)^{-\phi}p_S. \end{aligned}$$

We need to show that for this ϕ , also $p'_S \leq (1+\gamma)^{2\phi} p_S$. As $p'_S \leq (1+\gamma)p_S^{(0)} + (1+\gamma)^\varphi p_S^{(1)} + p_S^{(2)}$, this is true if

$$a^2 x_0 + b^2 x_1 + c^2 x_2 \leq \frac{1}{((1/a)x_0 + (1/b)x_1 + (1/c)x_2)^2}$$

or

$$((1/a)x_0 + (1/b)x_1 + (1/c)x_2)^2 (a^2 x_0 + b^2 x_1 + c^2 x_2) \leq 1 \quad (2)$$

To prove this, we need two claims whose proofs are tedious but follow from standard math.

Claim For any $a, b, c > 0$ and any $x_0, x_1, x_2 > 0$ with $x_0 + x_1 + x_2 = 1$,

$$(ax_0 + bx_1 + cx_2)^2 \leq (a^2 x_0 + b^2 x_1 + c^2 x_2)$$

Claim For any $a, b, c > 0$ and any $x_0, x_1, x_2 > 0$ with $x_0 + x_1 + x_2 = 1$,

$$((1/a)x_0 + (1/b)x_1 + (1/c)x_2)(ax_0 + bx_1 + cx_2) \leq 1$$

Combining the claims, Equation (2) follows, which completes the proof. \square

Hence, for any outcome of E_2 , $\mathbb{E}[p'_S] \leq (1+\gamma)^{-\varphi} p_S$ and $p'_S \leq (1+\gamma)^{2\varphi} p_S$ for some $\varphi \in [0, 1]$. If we define $q_S = \log_{1+\gamma} p_S$, then it holds that $\mathbb{E}[q'_S] \leq q_S - \varphi$. For any time t in I , let q_t be equal to q_S at time t and φ_t be defined as φ at time t . Our calculations above imply that as long as $p_S \in [\rho_{green}, \rho_{yellow}]$, $\mathbb{E}[q_{t+1}] \leq q_t - \varphi_t$ and $q_{t+1} \leq q_t + 2\varphi_t$.

Now, suppose that within subframe I we reach a point t when $p_S > \rho_{yellow}$. Since we start with $p_S \leq \rho_{green}$, there must be a time interval $I' \subseteq I$ so that right before I' , $p_S \leq \rho_{green}$, during I' we always have $\rho_{green} < p_S \leq \rho_{yellow}$, and at the end of I' , $p_S > \rho_{yellow}$. We want to bound the probability for this to happen.

Consider some fixed interval I' with the properties above, i.e., with $p_S \leq \rho_{green}$ right before I' and $p_S \geq \rho_{green}$ at the first round of I' , so initially, $p_S \in [\rho_{green}, (1+\gamma)\rho_{green}]$. We use martingale theory to bound the probability that in this case, the properties defined above for I' hold. Consider the rounds in I' to be numbered from 1 to $|I'|$, let q_t and φ_t be defined as above, and let $q'_t = q_t + \sum_{i=1}^{t-1} \varphi_i$. It holds that

$$\begin{aligned} \mathbb{E}[q'_{t+1}] &= \mathbb{E}[q_{t+1} + \sum_{i=1}^t \varphi_i] \\ &= \mathbb{E}[q_{t+1}] + \sum_{i=1}^t \varphi_i \leq q_t - \varphi_t + \sum_{i=1}^t \varphi_i \\ &= q_t + \sum_{i=1}^{t-1} \varphi_i \\ &= q'_t. \end{aligned}$$

Moreover, it follows from Inequality (1) that for any round t , $p'_S \leq (1+\gamma)^{2\varphi_t} p_S$. Therefore, $q_{t+1} \leq q_t + 2\varphi_t$, which implies that $q'_{t+1} \leq q'_t + \varphi_t$. Hence, we can define a martingale $(X_t)_{t \in I'}$ with $\mathbb{E}[X_{t+1}] = X_t$ and $X_{t+1} \leq X_t + \varphi_t$ that stochastically dominates q'_t . Recall that a random variable Y_t *stochastically dominates* a random variable Z_t if for any z , $\mathbb{P}[Y_t \geq z] \geq \mathbb{P}[Z_t \geq z]$. In that case, it is also straightforward to show that $\sum_i Y_i$ stochastically dominates $\sum_i Z_i$, which we will need in the following. Let $T = |I'|$. We will make use of Azuma's inequality to bound X_T .

Fact 4 (Azuma Inequality) *Let X_0, X_1, \dots be a martingale satisfying the property that $X_i \leq X_{i-1} + c_i$ for all $i \geq 1$. Then for any $\delta \geq 0$,*

$$\mathbb{P}[X_T > X_0 + \delta] \leq e^{-\delta^2 / (2 \sum_{i=1}^T c_i^2)}.$$

Thus, for $\delta = 1/\gamma + \sum_{i=1}^T \varphi_i$ it holds in our case that

$$\mathbb{P}[X_T > X_0 + \delta] \leq e^{-\delta^2 / (2 \sum_{i=1}^T \varphi_i^2)}.$$

This implies that

$$\mathbb{P}[q'_T > q'_0 + \delta] \leq e^{-\delta^2 / (2 \sum_{i=1}^T \varphi_i^2)},$$

for several reasons. First of all, stochastic dominance holds as long as $p_S \in [\rho_{green}, \rho_{yellow}]$, and whenever this is violated, we can stop the process as the requirements on I' would be violated, so we would not have to count that probability towards I' . Therefore,

$$\mathbb{P}[q_T > q_0 + 1/\gamma] \leq e^{-\delta^2 / (2 \sum_{i=1}^T \varphi_i^2)}.$$

Notice that $q_T > q_0 + 1/\gamma$ is required so that $p_S > \rho_{yellow}$ at the end of I' , so the probability bound above is exactly what we need. Let $\varphi = \sum_{i=1}^T \varphi_i$. Since $\varphi_i \leq 1$ for all i , $\varphi \geq \sum_{i=1}^T \varphi_i^2$. Hence,

$$\frac{\delta^2}{2 \sum_{i=1}^T \varphi_i^2} \geq \frac{(1/\gamma + \varphi)^2}{2\varphi} \geq \left(\frac{1}{2\varphi\gamma^2} + \frac{\varphi}{2} \right).$$

This is minimized for $1/(2\varphi\gamma^2) = \varphi/2$ or equivalently, $\varphi = 1/\gamma$. Thus,

$$\mathbb{P}[q_T > q_0 + 1/\gamma] \leq e^{-1/\gamma}$$

Since there are at most $\binom{f}{2}$ ways of selecting $I' \subseteq I$, the probability that there exists an interval I' with the properties above is at most

$$\binom{f}{2} e^{-1/\gamma} \leq f^2 e^{-1/\gamma} \leq \frac{1}{\log^c n}$$

for any constant c if $\gamma = O(1/(\log T + \log \log n))$ is small enough. \square

Lemma 11 *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. \square

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 11 we can prove the following lemma.

Lemma 12 *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 .*

Proof From Lemma 8 it follows that for every subframe I in F there is a time point $t \in I$ at which $p_S \leq \rho_{green}$ w.h.p. Consider now some fixed subframe I in F that is not the first one and suppose that the previous subframe in F had at least one round with $p_S \leq \rho_{green}$. Then it follows from Lemma 11 that for all rounds in I , $p_S \leq \rho_{red}$ w.m.p. (where the probability only depends on I and its preceding subframe), i.e., I is good. Hence, it follows from the Chernoff bounds that at most $\epsilon\beta/7$ of the odd-numbered as well as the even-numbered subframes after the first subframe in F are bad w.h.p. (if the constant α is sufficiently large). This implies that overall at most $\epsilon\beta/6$ of the subframes in F are bad w.h.p. \square

From Lemma 12 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 Proof of Theorem 2

Now, let us consider the two cases of Theorem 2 under the *strong* adversary.

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 $(T, 1 - \epsilon)$ -bounded adversary are satisfied. So Theorem 1 applies, which completes the proof of Theorem 2 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 1, which implies that JADE is constant competitive w.r.t. the nodes in V' , completes the proof of Theorem 2 b).

3.4 Limitations under the Strong Adversary

One may ask whether a stronger throughput result than Theorem 2 can be shown for the *strong* adversary. 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. We can prove the following lemmas which imply Theorem 3.

Lemma 13 *In general, JADE is not strongly c -competitive for a constant $c > 0$ if the strong 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 $c > 0$. \square

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

Lemma 14 *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 strong 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 . \square

Hence, Theorem 2 is the best one can show for JADE (within our notation). More generally, of course, no MAC protocol can guarantee a constant competitive throughput if the UDG is not connected. However, it is still an open question 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.

4 Simulations

In order to complement our theoretical insights, we report on our simulation results. First, we present our throughput results for a sufficiently large time interval, and then we discuss the convergence behavior. For our simulations, as in our formal analysis, we assume that initially all nodes $v \in V$ have a high sending probability of $p_v = \hat{p} = 1/24$. The nodes are distributed at random over a square plane of 4×4 units, and are connected in a unit disk graph manner (multi-hop). We simulate the jamming activity in the following way: for each round, a node is jammed independently with probability $(1 - \epsilon)$. Note that in the terminology we introduced, this adversary is strong (as rounds do not need to be open) and n -uniform (as nodes are jammed independently). The reason for studying this rather simplistic randomized “adversary” is twofold. First, although our formal results hold for arbitrary adversaries, it is not clear how to constructively compute such a worst adversarial strategy; second, a random adversary also complements our formal results better as it may capture the “average case” behavior.

We run the simulation for a sufficiently large number of time steps indicated by the Theorem 1, i.e., for $(\lceil T + (\log^3 n)/(\gamma^2 \epsilon) \rceil \cdot (\log n)/\epsilon)$ rounds, where $\epsilon = 0.1$, $T = 200$, and $\gamma = 1/(\log T + \log \log N)$. Simulations with different combinations of $\epsilon \in \{0.5, 0.3, 0.1\}$ and $T = \{50, 100, 150, 200\}$ showed that $\epsilon = 0.1$ and $T = 200$ yields the lowest throughput (and the strongest adversary), and hence, in the following, we will focus on this most challenging case. (The parameter γ is set to a value satisfy its definition, i.e., $\gamma = O(1/(\log T + \log \log N))$.)

Figure 1 (*top*) shows the throughput competitiveness of JADE for a scenario where different numbers of nodes (i.e., $n \in [100, 2000]$) are distributed uniformly at random over the plane and a scenario where the nodes are distributed according to a normal/Gaussian distribution $\mathcal{N}(0, 1)$. In both cases, the throughput is larger when the density is higher. This corresponds to our formal insight that a constant competitive throughput is possible only if the node density exceeds a certain threshold. For example, in a scenario with 100 nodes

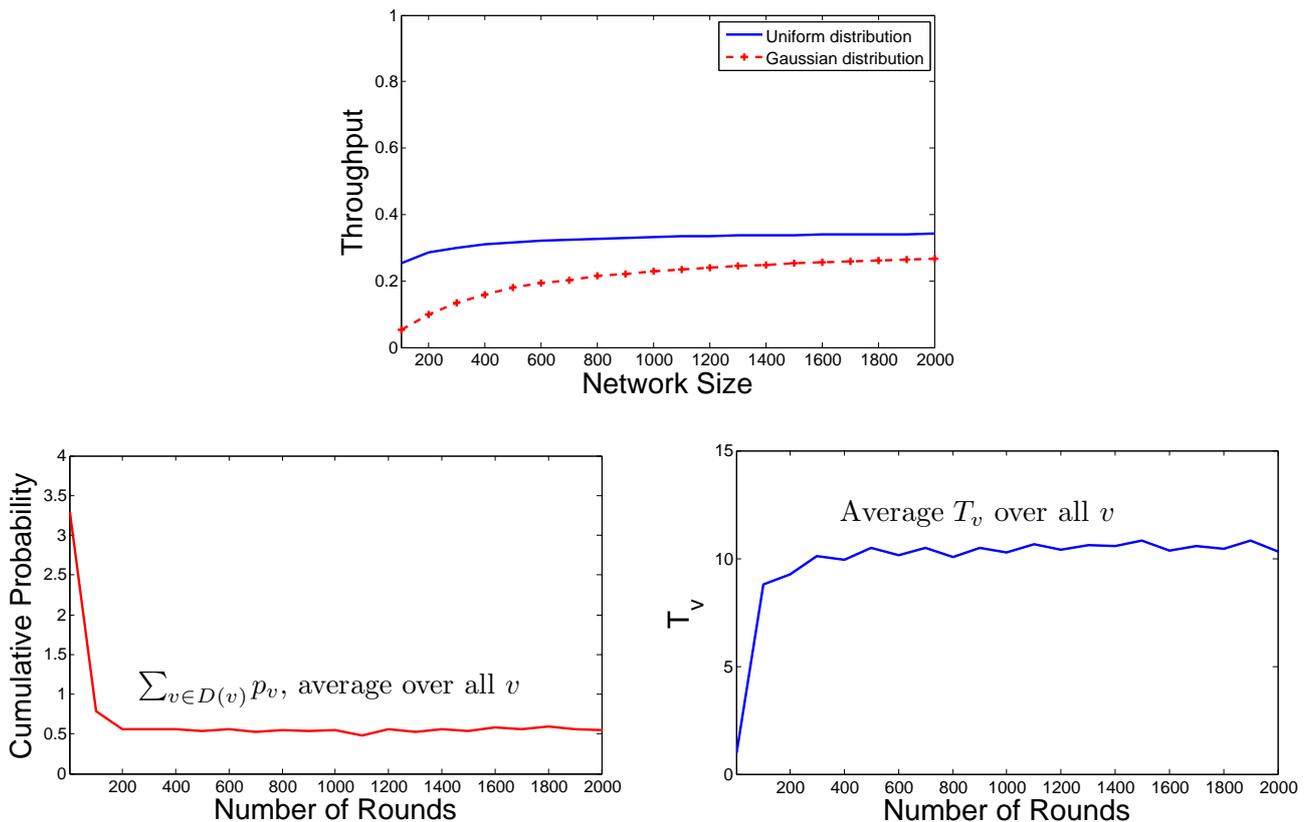


Fig. 1 *Top*: Throughput as a function of network size, where $n \in [100, 2000]$, $\epsilon = 0.1$, $T = 200$, and $\gamma = 1/(\log T + \log \log n)$. The result is averaged over 10 runs. *Bottom left*: Convergence behavior for multi-hop networks (uniform distribution). As a demonstration, we used $n = 500$, $\epsilon = 0.1$, $T = 200$, and $\gamma = 1/(\log T + \log \log N)$. Note that the start-up phase where the sending probabilities are high is short (no more than 50 rounds). *Bottom right*: Convergence of T_v for multi-hop networks (uniform distribution). For demonstration, we used $n = 500$, $\epsilon = 0.1$, $T = 200$, and $\gamma = 1/(\log T + \log \log N)$.

in the 4×4 plane (density of 6.25), there are at least $6.25\pi \approx 20 \geq 2/\epsilon = 20$ uniformly distributed nodes in one unit disk. As can be seen in the figure, when the number of nodes is larger than 600, the throughput falls between 20% and 40% for both uniform distribution and Gaussian distribution.

Convergence time is the second most important evaluation criterion. We found that already after a short time, a constant throughput is achieved; in particular, the total sending probability per unit disk approaches a constant value quickly. This is due to the nodes' ability to adapt their sending probabilities fast, see Figure 1 (*bottom left*). The figure also illustrates the high correlation between success ratio and aggregated sending probability.

Finally, we have also studied the average of the T_v values over time. The average quickly stabilizes to a value around 10, as shown in Figure 1 (*bottom right*).

5 Conclusion

This article 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.

There are several questions that remain open. For instance, we assumed a common parameter γ which is known by all nodes and which depends on n and T . Although the estimations on these parameters we need are very rough and scalable, it remains an open question whether this limitation can be relaxed, and e.g., a local value $\gamma_v = 1/\log T_v$ would also work.

References

1. G. Alnifie and R. Simon. A multi-channel defense against jamming attacks in wireless sensor networks. In *Proc. of*

- Q2SWinet '07*, pages 95–104, 2007.
2. B. Awerbuch, A. Richa, and C. Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *Proc. of PODC '08*, 2008.
 3. E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa. On the performance of IEEE 802.11 under jamming. In *Proc. of IEEE Infocom '08*, pages 1265–1273, 2008.
 4. M. A. Bender, M. Farach-Colton, S. He, B. C. Kuszmaul, and C. E. Leiserson. Adversarial contention resolution for simple channels. In *Proc. of SPAA '05*, 2005.
 5. T. Brown, J. James, and A. Sethi. Jamming and sensing of encrypted wireless ad hoc networks. In *Proc. of MobiHoc '06*, pages 120–130, 2006.
 6. J. Chiang and Y.-C. Hu. Cross-layer jamming detection and mitigation in wireless broadcast networks. In *Proc. of MobiCom '07*, pages 346–349, 2007.
 7. B. S. Chlebus, D. R. Kowalski, and M. A. Rokicki. Adversarial queuing on the multiple-access channel. In *Proc. of PODC '06*, 2006.
 8. A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60(2):115 – 143, 2006.
 9. S. Dolev, S. Gilbert, R. Guerraoui, D. Kowalski, C. Newport, F. Kuhn, and N. Lynch. Reliable distributed computing on unreliable radio channels. In *Proc. 2009 MobiHoc S3 Workshop*, 2009.
 10. S. Dolev, S. Gilbert, R. Guerraoui, F. Kuhn, and C. C. Newport. The wireless synchronization problem. In *Proc. 28th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 190–199, 2009.
 11. S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Gossiping in a multi-channel radio network: An oblivious approach to coping with malicious interference. In *Proc. of the Symposium on Distributed Computing (DISC)*, 2007.
 12. S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Secure communication over radio channels. In *Proc. 27th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 105–114, 2008.
 13. S. Dolev, S. Gilbert, R. Guerraoui, and C. C. Newport. Gossiping in a multi-channel radio network. In *Proc. 21st International Symposium on Distributed Computing (DISC)*, pages 208–222, 2007.
 14. S. Gilbert, R. Guerraoui, D. Kowalski, and C. Newport. Interference-resilient information exchange. In *Proc. of the 28th Conference on Computer Communications. IEEE Infocom 2009.*, 2009.
 15. S. Gilbert, R. Guerraoui, D. R. Kowalski, and C. C. Newport. Interference-resilient information exchange. In *Proc. 28th IEEE International Conference on Computer Communications (INFOCOM)*, pages 2249–2257, 2009.
 16. 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.
 17. L. A. Goldberg, P. D. Mackenzie, M. Paterson, and A. Srinivasan. Contention resolution with constant expected delay. *J. ACM*, 47(6), 2000.
 18. J. Hastad, T. Leighton, and B. Rogoff. Analysis of backoff protocols for multiple access channels. *SIAM Journal on Computing*, 25(4), 1996.
 19. M. Heusse, F. Rousseau, R. Guillier, and A. Duda. Idle sense: An optimal access method for high throughput and fairness in rate diverse wireless lans. In *Proc. SIGCOMM*, 2005.
 20. IEEE. Medium access control (MAC) and physical specifications. In *IEEE P802.11/D10*, 1999.
 21. K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proc. 9th Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 66–80, 2003.
 22. S. Jiang and Y. Xue. Providing survivability against jamming attack via joint dynamic routing and channel assignment. In *Proc. 7th Workshop on Design of Reliable Communication Networks (DRCN)*, 2009.
 23. C. Koo, V. Bhandari, J. Katz, and N. Vaidya. Reliable broadcast in radio networks: The bounded collision case. In *Proc. of PODC '06*, 2006.
 24. F. Kuhn, T. Moscibroda, and R. Wattenhofer. Radio network clustering from scratch. In *Proc. of ESA '04*, 2004.
 25. B.-J. Kwak, N.-O. Song, and L. E. Miller. Performance analysis of exponential backoff. *IEEE/ACM Transactions on Networking*, 13(2):343–355, 2005.
 26. Y. Law, L. van Hoesel, J. Doumen, P. Hartel, and P. Havinga. Energy-efficient link-layer jamming attacks against wireless sensor network mac protocols. In *Proc. of SASN '05*, pages 76–88, 2005.
 27. 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.
 28. X. Liu, G. Noubir, R. Sundaram, and S. Tan. Spread: Foiling smart jammers using multi-layer agility. In *Proc. of Infocom '07*, pages 2536–2540, 2007.
 29. D. Meier, Y. A. Pignolet, S. Schmid, and R. Wattenhofer. Speed dating despite jammers. In *Proc. DCOSS '09*, June 2009.
 30. V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein. Using channel hopping to increase 802.11 resilience to jamming attacks. In *Proc. of Infocom '07*, pages 2526–2530, 2007.
 31. R. Negi and A. Perrig. Jamming analysis of MAC protocols. Technical report, Carnegie Mellon University, 2003.
 32. G. Noubir. On connectivity in ad hoc networks under jamming using directional antennas and mobility. In *Proc. 2nd International Conference on Wired/Wireless Internet Communications (WWIC)*, pages 186–200, 2004.
 33. A. Peleg and D. Peleg. Feasibility and complexity of broadcasting with random transmission failures. In *Proc. of PODC '05*, 2005.
 34. P. Raghavan and E. Upfal. Stochastic contention resolution with short delays. *SIAM Journal on Computing*, 28(2):709–719, 1999.
 35. A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Competitive and fair medium access despite reactive jamming. In *Proc. 31st International Conference on Distributed Computing Systems (ICDCS)*, 2011.
 36. A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Self-stabilizing leader election for single-hop wireless networks despite jamming. In *Proc. 12th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2011.
 37. 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.
 38. M. K. Simon, J. K. Omura, R. A. Schultz, and B. K. Levin. *Spread Spectrum Communications Handbook*. McGraw-Hill, 2001.
 39. D. Thunte and M. Acharya. Intelligent jamming in wireless networks with applications to 802.11b and other networks. In *Proc. of MILCOM '06*, 2006.
 40. P. M. van de Ven, A. J. E. M. Janssen, J. S. H. van Leeuwen, and S. C. Borst. Achieving target throughputs in random-access networks. *Perform. Eval.*, 68:1103–1117, November 2011.

41. A. Wood, J. Stankovic, and G. Zhou. DEEJAM: Defeating energy-efficient jamming in IEEE 802.15.4-based wireless networks. In *Proc. of SECON '07*, 2007.
42. W. Xu, K. Ma, W. Trappe, and Y. Zhang. Jamming sensor networks: attack and defense strategies. *IEEE Network*, 20(3):41–47, 2006.
43. 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.
44. W. Xu, T. Wood, and Y. Zhang. Channel surfing and spatial retreats: defenses against wireless denial of service. In *Proc. of Workshop on Wireless Security*, 2004.
45. S. Ye, Y. Wang, and Y. Tseng. A jamming-based MAC protocol for wireless multihop ad hoc networks. In *Proc. IEEE 58th Vehicular Technology Conference*, 2003.
46. S.-R. Ye, Y.-C. Wang, and Y.-C. Tseng. A jamming-based MAC protocol to improve the performance of wireless multihop ad-hoc networks. *Wirel. Commun. Mob. Comput.*, 4(1):75–84, 2004.
47. J. Zander. Jamming in slotted ALOHA multihop packed radio networks. *IEEE Transactions on Networking*, 39(10):1525–1531, 1991.