Robust and Efficient Medium Access Despite Jamming

by

Jin Zhang

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved May 2012 by the
Graduate Supervisory Committee:

Andréa W. Richa, Chair
Christian Scheideler
Arunabha Sen
Guoliang Xue

ARIZONA STATE UNIVERSITY

August 2012

ABSTRACT

Interference constitutes a major challenge for communication networks operating over a shared medium where availability is imperative. This dissertation studies the problem of designing and analyzing efficient medium access protocols which are robust against strong adversarial jamming. More specifically, four medium access (MAC) protocols (i.e., JADE, ANTIJAM, COMAC, and SINRMAC) which aim to achieve high throughput despite jamming activities under a variety of network and adversary models are presented. We also propose a self-stabilizing leader election protocol, SELECT, that can effectively elect a leader in the network with the existence of a strong adversary.

Our protocols can not only deal with *internal* interference without the exact knowledge on the number of participants in the network, but they are also robust to unintentional or intentional *external* interference, e.g., due to co-existing networks or jammers. We model the external interference by a powerful *adaptive and/or reactive adversary* which can jam a $(1-\varepsilon)$-portion of the time steps, where $0 < \varepsilon \leq 1$ is an arbitrary constant. We allow the adversary to be adaptive and to have complete knowledge of the entire protocol history. Moreover, in case the adversary is also reactive, it uses carrier sensing to make informed decisions to disrupt communications.

Among the proposed protocols, JADE, ANTIJAM and COMAC are able to achieve $\Theta(1)$-competitive throughput with the presence of the strong adversary; while SINRMAC is the first attempt to apply SINR model (i.e., Signal to Interference plus Noise Ratio), in robust medium access protocols design; the derived principles are also useful to build applications on top of the MAC layer, and we present SELECT, which is an exemplary study for leader election, which is one of the most fundamental tasks in distributed computing.

Dedicated to my family

ACKNOWLEDGEMENTS

First of all, I would like to express my sincerest gratitude to my advisor, Dr. Andréa W. Richa, for her professional guidance, patience, support, and encouragement over the course of my PhD. I would like to thank all my committee members, Dr. Christian Scheideler, Dr. Arunabha Sen, and Dr. Guoliang Xue for their invaluable advice. I learned a lot from my committees about research either by writing papers with them or taking their classes. I really appreciate the opportunity to be able to spend two fruitful summers in University of Paderborn and Telekom Innovation Laboratories in Germany, working with Dr. Christian Scheideler, Dr. Stefan Schmid, and Dr. Andréa W. Richa on my dissertation, which I enjoyed very much. I would also like to acknowledge my coauthors (Andréa W. Richa, Christian Scheideler, Stefan Schmid, Guoliang Xue, Dejun Yang, Xi Fang, and Satyajayant Misra), and I am grateful for their collaboration and contribution.

My gratitude also goes to my friends. It is their help and support that makes my life at ASU more joyful. I would like to thank Dejun Yang, who is always willing to help and give his best suggestions. Many thanks to Dr. Stefan Schmid, for all the discussions and support regarding my research and dissertation. I also would like to thank the staff members in Department of Computer Science and Engineering (CSE) for providing a supportive learning and research environment. I really appreciate their kind help whenever I ask for them.

Finally, I would like to thank my dearest parents, who give me endless love and constant encouragement during my study at ASU and throughout my whole life.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figure Page

Figure Page

Figure Page

Chapter 1

INTRODUCTION

1.1 Overview

Designing efficient medium access protocols in a wireless network environment that are robust against different kinds of interference is one of the most relevant but also most complex problems in distributed computing. First, a wireless network requires distributed access coordination mechanisms which minimize the *internal* interference due to simultaneous transmissions from wireless devices in the same network. In addition, the availability of the wireless medium can vary significantly over time due to the *external* interference, e.g., due to disturbances from other sources such as microwaves, due to transmissions of coexisting (potentially mobile) networks, or due to intentional or even adversarial interruptions. For example, it is well-known that already simple jamming attacks—without using any special hardware—constitute a threat for the widely used IEEE 802.11 MAC protocol [7]. Due to the problem's relevance, there has been a significant effort to cope with such disruption problems both from the industry and the academia side and much progress has been made over the past few years.

This dissertation aims to design and analyze robust medium access protocols, so that even with the existence of a strong adversary the protocols can still manage to achieve provably high throughput. Note that we consider adversarial physical layer jamming only. Although we do not study malicious fake message jamming and other form of jamming activities which are above the physical layer, our physical adversarial jamming model works in conjunction with other adversary models at higher layers. The protocols studied here operate on the Medium Access Control (MAC) sub-layer of the data link layer (defined in the seven-layer OSI model), so we call them MAC protocols. The MAC protocols address the problem of how to decide who gets to use the shared medium when there is a contention for it. Many MAC protocols have been proposed, however, the presence of a strong adversary in the network makes the existing protocols vulnerable and inefficient to the jamming attack. Classic defense mechanisms operate at the physical layer [38, 42] and there exist approaches both to *avoid* as well as to *detect* jamming. Spread spectrum and frequency hopping technologies have been shown to be very effective to avoid jamming with widely spread signals. These physical layer solutions are orthogonal to our work, and can improve the robustness of the protocols further. However, the ISM frequency band used by IEEE 802.11 variants is too narrow to effectively apply spread spectrum techniques [10]. Also, as jamming strategies can come in many different flavors, detecting jamming activities by simple methods based on signal strength, carrier sensing, or packet delivery ratios has turned out to be quite difficult [37]. A more comprehensive overview of the related work is provided in Section 1.3.

We consider two types of adversaries in this dissertation: (i) **adaptive but non-reactive**, where the adversary has the complete knowledge of the protocol history and can use this knowledge to make jamming decisions. However, the adversary has to take actions *before* honest nodes decide whether to transmit a message or not; Clearly, the adaptive adversary is much stronger compared to oblivious or random adversaries; (ii) **adaptive and reactive**, where the adversary is adaptive, and moreover, it is reactive in the sense that the adversary can use carrier sensing to sense the channel and make jamming decisions *after* honest nodes made their transmission decisions, which makes the adversary even more powerful and effective. Let us consider the following scenario as an example: suppose that at the current time step no node in the network decides to transmit, then the adversary can quickly sense the channel and decide not to jam the channel, so that the energy can be saved to jam the channel when certain transmission activity is taking place.

Our work is motivated by the results in [7] and [6]. In [7] 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. [6] initiated the study of throughput-competitive MAC protocols in single-hop wireless networks under continuously running, adaptive jammers, and presented a protocol that achieves a high performance under adaptive but non-reactive jamming. In a single-hop network, all the nodes are within transmission and interference range of each other, i.e., the communication network is a complete graph. In contrast, in a multi-hop network, not all nodes are within each other's transmission and interference range. In order to get a message broadcasted in a multi-hop network, more than one hop of transmission is needed. We extend the results in [6] in the following ways:

1. JADE: The JADE protocol is designed for multi-hop wireless networks. Crucial modifications are made based on the protocol in [6], so that JADE achieves constant throughput in multi-hop wireless networks that can be modeled as unit disk graph (see Section 1.2.1), and is robust against an adaptive but non-reactive adversary. We discuss JADE in more detail in Chapter 2 (this work also appeared in [49]).
2. ANTIJAM: Although an adaptive but non-reactive adversary is much stronger than an oblivious or random adversary, not being able to make a jamming decision based on honest nodes' decisions in the current time step makes the adversary model less practical. Hence, we consider adaptive and reactive adversary, and propose ANTIJAM, where constant throughput can be achieved

with the presence of a strong adaptive and reactive adversary in the single-hop wireless networks We discuss ANTIJAM in detail in Chapter 3 (this work also appeared in [50]).

3. SELECT: Leader election is a classical problem in the field of distributed computing. Once a leader is elected, many coordination tasks are greatly simplified. We consider the problem of electing a leader in a harsh wireless network in order to coordinate access to a shared communication medium. We propose SELECT, a self-stabilizing leader election protocol that can always elect one and only one leader from a single-hop wireless network, no matter what the initial state is, and despite the existence of a strong adaptive and reactive adversary. We discuss SELECT in detail in Chapter 4 (this work also appeared in [51]).

4. COMAC: The problem of accessing the shared medium by different *co-existing* networks fairly and efficiently, especially in environments with uncontrollable external interference, such as jamming, is important and challenging. Nowadays, more and more devices belonging to co-existing networks share a chunk of the limited wireless spectrum resource simultaneously. We propose COMAC, which is able to achieve constant throughput and fairness, since it evenly distributes the number of successful transmissions for each individual network, up to a small multiplicative factor. The protocol is also robust against an adaptive but non-reactive adversary. We discuss COMAC in detail in Chapter 5 (this work also appeared in [53])).

5. SINRMAC: Designing a jamming-resistant MAC protocol under the widely used and more realistic *Signal-to-Interference-plus-Noise-Ratio* (*SINR*) would be the next big step forward. Hence, we explore the possibility to come up with such a MAC protocol, called SINRMAC, which can achieve high

throughput despite jamming. An initial study on this problem is presented in Chapter 6 (this work also appeared in [52]).

1.2 General Model

We specify the general network (interference), communication and adversary models here. Note that models for specific protocols may vary. Please refer to the corresponding chapters for the detailed models used by each protocol. We summarize the models used by different protocols in Table 1.1.

Protocol	Network Model	Communication Model	Adversary Model
JADE	UDG (Multi-hop)	Single Channel, Half-duplex	Adaptive but Non-Reactive
ANTIJAM	Single-hop	Single Channel, Half-duplex	Adaptive and Reactive
SELECT	Single-hop	Single Channel, Half-duplex	Adaptive and Reactive
COMAC	Single-hop	Single Channel, Half-duplex	Adaptive but Non-Reactive
SINRMAC	SINR (Multi-hop)	Single Channel, Half-duplex	Adaptive but Non-Reactive

Table 1.1: Different models for different protocols.

1.2.1 Network (Interference) Model:

1. *Single-hop*: The network consists of n honest and reliable nodes that are within the transmission and interference range of each other, which is equivalent to assuming that the network topology is a complete graph.

2. *UDG (Multi-hop)*: As an initial study of multi-hop wireless networks, we use the Unit Disk Graph (UDG) to model the network topology. More specifically, let the network be represented by a graph $G = (V, E)$ 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 and interference range, i.e., $\{u, v\} \in E$, if and only if their (normalized) distance is at most 1. Note that the transmission and interference range of the nodes are the same under the UDG model.

3. *SINR (Multi-hop)*: We assume n wireless nodes are distributed arbitrarily in the 2-dimensional Euclidean plane. The SINR model defines a parameter called minimum *signal-to-interference-plus-noise ratio* (SINR) at which a data frame can still be received with a reasonably low frame error rate. A message sent from u to v is received correctly if and only if

$$\frac{P_v(u)}{\mathscr{N} + \sum_{w \in S} P_v(w)} \geq \beta_1$$

where $P_v(u)$ is the received power at node v of the signal transmitted by node u, $\mathscr{N}$ captures the background noise (e.g., thermal), S is the subset of nodes in $V \setminus \{u, v\}$ that are concurrently transmitting, and β_1 is the *SINR threshold* that depends on the desired rate, the modulation scheme, etc.

1.2.2 Communication Model:

We assume a back-logged scenario where the nodes continuously contend for sending a packet on the wireless channel. A node may either transmit a message or sense the channel at a time step, but it cannot do both, and there is no immediate feedback mechanism telling a node whether its transmission was successful. When considering single-hop or UDG (multi-hop) network model, a node sensing the channel may come across one of the following three scenarios: (i) sense an *idle* channel (in case no other node in the transmitting range of the node transmits at that time); (ii) sense a *busy* channel (in case two or more nodes within the transmission range of the node transmit at the time step, or the adversary or the adversary disrupts the signal at the node); or (iii) *receive* a packet (in case exactly one node within the transmitting range of the node transmits at the time step). While considering SINR as the network model, a node does not have clear distinctions regarding *idle* and *busy* channels. A noise level threshold is introduced to resolve this issue. More details can be found in Chapter 6. The wireless channel considered in this dissertation has single frequency and is half-duplex.

1.2.3 Adversary Model:

We consider two types of adversaries:

1. *adaptive but non-reactive adversary*: **adaptive** in the sense that the adversary knows the protocol history, and can make jamming decisions based on it. The adversary has to make a jamming decision *before* honest nodes decide whether to transmit or not.

2. *adaptive and reactive adversary*: in addition to being **adaptive**, the adversary is **reactive** in the sense that it can perform physical carrier sensing to learn whether the channel is currently idle or not, and jam the medium depending on these measurements. Note that the adversary can sense the channel condition in the current time step and make a jamming decision instantly (i.e., the jamming decision is made *after* honest nodes decide whether to transmit or not).

Note that although being reactive gives the adversary more power by revealing some information about the nodes' random decisions at the current time step, an adaptive adversary is already much stronger than its oblivious and random adversaries counterparts.

1.3 Related Work

Due to the topic's importance, wireless network jamming has been extensively studied in the applied research fields [1, 10, 12, 31, 35, 37, 38, 42, 43, 58, 61, 62, 63], both from the attacker's perspective [12, 35, 37, 63] as well as from the defender's perspective [1, 10, 12, 37, 38, 42, 61, 63]—also in multi-hop settings (e.g. [29, 45, 65, 66, 67]).

Recent work has also studied *MAC layer strategies* against jamming, including coding strategies (e.g., [12]), channel surfing and spatial retreat (e.g., [1, 64]), or mechanisms to hide messages from a jammer, evade its search, and reduce the impact of corrupted messages (e.g., [61]). However, these methods do not help against an adaptive jammer with *full* information about the history of the protocol, like the one considered in our work.

In the theory community, work on MAC protocols has mostly focused on efficiency. Many of these protocols are *random backoff protocols* (e.g., [8, 13, 14, 25, 48]) that do not take jamming activity into account and, in fact, are not robust against it (see [6] for more details). But also some theoretical work on *jamming* is known (e.g., [16] for a short overview). There are two basic approaches in the literature. The first assumes randomly corrupted messages (e.g. [47]), which is much easier to handle than adaptive adversarial jamming [7]. 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. [23, 32]), or bounds the number of channels the adversary can jam (e.g. [22, 39]). The protocols in, e.g., [32] can tackle adversarial jamming at both the MAC and network layers, where the adversary may not only jam the channel but also introduce malicious (fake) messages (possibly with address spoofing). However, these solutions depend on the fact that the adversarial jamming budget is finite, so it is not clear whether the protocols would work under heavy continuous jamming. (The result in Theorem 1 of [23] upper bounds the adversary's capability of disrupting communications with a budget of β messages, and then shows that the proposed protocol needs at least 2β rounds to terminate, which implies a jamming rate below 0.5. The handshaking mechanism in [32] requires an even lower jamming rate.

In the multi-channel version of the problem introduced in the theory community by Dolev [17] and also studied in [19, 17, 18, 22, 39], 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 [17, 22] and at least quadratic in the number of jammed channels for randomized protocols [18, 39] if the adversary can jam almost all channels at a time). Recent work [19] 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. [22] study robust information exchange in single-hop networks.

There is also a chapter on the leader election application considered in this dissertation. Leader election is an evergreen in distributed algorithms research and there exist many theoretical and practical results [5, 20, 33, 36, 41, 44, 57, 60]. The following two book chapters provide a good introduction: Chapter 3 in [4] and Chapter 8 in [27]. A leader election algorithm should be as flexible as possible in the sense that a correct solution is computed *independently* of the initial network state. For instance, the algorithm should be able to react to a leader departure, or be able to cope with situations where for some reasons, multiple nodes consider themselves leaders. *Self-stabilization* [15] is an attractive concept to describe such self-repairing properties of an algorithm, and it has been intensively studied already, not only in terms of eventual stabilization but also in terms of guaranteed convergence times (see e.g., the works on time-adaptive self-stabilization such as [34]). Several self-stabilizing leader election protocols have been devised, e.g., [2, 11, 28] (see also the fault-contained solutions such as [21]). However, none of these approaches allows us to elect a leader in a wireless network that is exposed to harsh interference or even adaptive jamming. But such interruptions of communication are often unavoidable in wireless systems, and we believe that electing a leader can be particularly useful in such harsh environments.

When it comes to design robust and efficient MAC protocols for coexisting networks, the performance achieved by the MAC protocols described in [6, 49, 50], which are jamming-resistant in single network settings, drops sharply if multiple networks are collocated: This is due to the fact that in these protocols, each individual co-existing network will strive to achieve a constant competitive throughput in the non-jammed time periods, which requires a constant cumulative access probability *per co-existing network*. It is easy to see that this necessarily leads to a throughput which is exponentially small in the number of co-existing networks. More importantly, the algorithmic approach used [6, 49, 50] is doomed to fail in the context of co-existing networks, as nodes in different networks do not have a consistent view of the successful transmissions: in a remote network, a successful transmission cannot be distinguished from a collision or jammed time step.

It turns out that in a co-existing scenario, the nodes must strike a good balance between a less aggressive (more cooperative) medium access strategy while remaining robust against external interference. We will show that this can be achieved by monitoring the availability of the wireless medium over time and adjusting the sending probabilities or backoffs according to the fraction of observed *idle time periods*. (A similar approach is used in the *IdleSense* [26] Distributed Coordination Function to synchronize the nodes' contention windows.) Implicitly synchronizing access via idle time periods is also the key to enable fairness between co-existing networks. The performance analysis of such an algorithm however is involved, as the distributed and randomized decisions exhibit many non-trivial dependencies. Nevertheless, we are able to rigorously prove good competitive throughput and fairness properties, which is also confirmed by our simulation study.

Interestingly, although co-existing networks are ubiquitous and many different aspects are discussed intensively (e.g., the packet inter-arrival time and fairness in co-existing 802.11a/g and 802.11n networks [3], interference cancelation phenomena [54], transmission capacities in multi-antenna ad-hoc networks [30], or even explicit inter-network communication for frequency cooperation [68]) in different contexts (e.g., in the current debate on white space liberalization [46] where primary TV and microphone users announcing their reservations in a central database are given strict priority), we are not aware of any work on the design of MAC protocols for independent co-existing networks with rigorous formal competitive throughput and fairness guarantees.

1.4 Preliminaries

In this section, we present some important definitions and basic results, which will be used in the following chapters.

1.4.1 Intuition

We first explain the intuition used by our protocols in this dissertation.

Although serving for different purposes under different models, the intuition behind Jade, AniJam, Select, CoMac, and SinrMac is similar to the one presented in [6]. For Jade, AntiJam, CoMac and SinrMac, the goal is to achieve provably high throughput against adversarial jamming by adapting nodes' access probabilities based on the events of idle channel and successful transmission. For Select, the main goal is to have a leader election protocol that is self-stabilizing despite adversarial jamming. To accomplish this, nodes also need to adjust their probabilities based on idle channel or successful transmissions so that FOLLOWER and LEADER messages can go through despite jamming. Hence, how to adjust nodes' access probabilities appropriately is crucial to the design of robust and efficient medium access despite jamming. Next, we explain the intuition in more detail. We assume the network model is single-hop and single network for now, in case of multi-hop networks as well as coexisting networks, the same intuition still applies, as explained in Chapters 2 and 5.

Let $G = (V,E)$ be a single-hop network where $n = |V|$. Each node v maintains a medium access probability p_v which determines the probability that v transmits a message in a communication round. Let the cumulative probability $p = \sum_v p_v$, q_0 be the probability that the channel is idle, and q_1 be the probability that exactly one node is sending a message. We have the following lemma which was first proved in [6]:

Lemma 1.1 $q_0 \cdot p \leq q_1 \leq \frac{q_0}{1-\hat{p}} \cdot p$.

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

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

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

which implies the lemma.

According to Lemma 1.1, if $q_0 = \Theta(q_1)$, then the cumulative sending probability p is constant, which in turn implies that at any non-jammed time step we have constant probability of having a successful transmission. Hence our protocol aims at adjusting the sending probabilities p_v of the nodes such that $q_0 = \Theta(q_1)$, in spite of adversarial jamming activities. This could be achieved by adjusting nodes' access probabilities based on the events of idle channel and successful transmissions, more details of which are provided in the following chapters.

1.4.2 Mathematical Tools

In order to perform theoretical analysis on our protocols, we will frequently use the following two lemmas: Lemma 1.2 is the variant of the Chernoff bounds [6, 56]; Lemma 1.3 follows immediately from the Taylor series of the exponential function.

Lemma 1.2 *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*

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

Lemma 1.3 *For all* $0 < x < 1$ *it holds that*

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

Chapter 2

THE JADE PROTOCOL

In this chapter, we consider the problem of designing a robust MAC protocol for multi-hop wireless networks, with the existence of a strong adaptive but non-reactive adversary. We prove that the proposed protocol, JADE, can achieve constant competitive throughput, and the limitations of JADE is also discussed.

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 [6], 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 $(1 - \varepsilon)$-fraction of its rounds is jammed ($\varepsilon > 0$ 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 $(T, 1-\varepsilon)$-bounded for some $T \in \mathbb{N}$ and $0 < \varepsilon < 1$, if for any time window of size $w \geq T$ and at any node v, the adversary can jam at most $(1-\varepsilon)w$ 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. In this chapter, if not stated otherwise and by default, we always refer to the adversary defined here. We also consider a stronger adversary that does not have the limitation of providing open rounds. Note that we sometime explicitly use the adjective *weak* to distinguish the adversary defined previously from the stronger variant.

Next, we need to rigorously define *c-competitiveness* in a multi-hop wireless network setting. 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. Then, we have the following definition:

Definition 2.1 *A MAC protocol is called* c-competitive *against some* $(T, 1-\varepsilon)$-*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 $(T, 1-\varepsilon)$-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, we also consider a stronger adversary: we say that a *strong* adversary is $(T, 1-\varepsilon)$-bounded, if for any time window of size $w \geq T$ and at any node v, the adversary can jam at most $(1-\varepsilon)w$ of the w rounds at v, where $T \in \mathbb{N}$ and $0 < \varepsilon < 1$. 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., [59]), 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 dissertation, we 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 $c \geq 1$.

2.1 Contribution

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 2.2 *When running* JADE *for* $\Omega([T+(\log^3 n)/(\gamma^2\varepsilon)]\cdot(\log n)/\varepsilon)$ *rounds it holds w.h.p. that* JADE *achieves a constant competitive throughput (i.e., independent of n) for any* $(T, 1-\varepsilon)$*-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 problem 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 2.3 is satisfied.

Theorem 2.3 *When running* JADE *for* $\Omega((T \log n)/\varepsilon + (\log n)^4/(\gamma\varepsilon)^2)$ *rounds,* JADE *has a constant competitive throughput against any strong adversary that is* $(T, 1-\varepsilon)$*-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* $2/\varepsilon$ *nodes within the transmission range of every node.*

In Section 2.3.4, we show that Theorem 2.3 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. Let $D(u)$ denote the set of nodes around node u, consisting of u's neighboring nodes as well as u.

Theorem 2.4 *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* $\varepsilon \leq 1/3$. *Moreover,* JADE *is also not c-competitive for a constant c if there are nodes u with* $|D(u)| = o(1/\varepsilon)$ *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.

2.2 Description of JADE

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

2.2.1 Intuition

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 $1/p_v$.) 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 $0 < \hat{p} < 1$.

The intuition behind JADE follows the guideline illustrated in Section 1.4.1, although Lemma 1.1 needs to be verified in multi-hop scenario. We show this as follows.

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

[1] In this dissertation, 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.

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 prove Lemma 1.1 for the multi-hop case.

By Lemma 1.1, 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 [26] 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.

2.2.2 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.

Algorithm 1 JADE: for each node v

```
$roundcounter = 0$
$p_v := \hat{p}$
$c_v := 1$
$T_v := 1$ {JADE works in synchronized rounds}
while True do
    $v$ decides with probability $p_v$ to send a message
    if $v$ decides to send a message then
        $v$ sends a message
    else
        $v$ senses the channel
        if $v$ senses an idle channel then
            $p_v := \min\{(1+\gamma)p_v, \hat{p}\}$
        else if $v$ successfully receives a message then
            $p_v := (1+\gamma)^{-1}p_v$
            $T_v := \max\{T_v - 1, 1\}$
        end if
    end if
    $c_v := c_v + 1$
    if $c_v > T_v$ then
        $c_v := 1$
        if there was no successful transmission or an idle channel among the past
        $T_v$ time steps then
            $p_v := (1+\gamma)^{-1}p_v$
            $T_v := \min\{T_v + 1, 2^{1/(4\gamma)}\}$
        end if
    end if
    $roundcounter := roundcounter + 1$
end while
```

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 [6], 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 which is surrounded by many nodes with low p_v values which are in turn surrounded by nodes with high p_w values (and hence the middle nodes's p_v values stay low), will never see any successful transmissions (apart from u's own transmissions), and hence T_u may increase arbitrarily. Such high T_u values however are harmful to the fast recovery properties of the protocol.

2.3 Analysis

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 2.2 in Sections 2.3.1 and 2.3.2, and then derive Theorem 2.3 in Section 2.3.3. The limitations of JADE under the strong adversary are discussed in Section 2.3.4.

The analysis makes repeated use of Lemma 1.3 and the Chernoff bounds in Lemma 1.2.

2.3.1 Proof of Theorem 2.2

First, we focus on a *time frame* F consisting of $(\alpha \log n)/\varepsilon$ *subframes* of size $f = \alpha[T + (\log^3 n)/(\gamma^2 \varepsilon)]$ 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 [6]. It is true because only successful message transmissions reduce T_u.

Lemma 2.5 *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 2.3.2.

Lemma 2.6 *For every node u, $\sum_{v\in D(u)} p_v = O(1)$ for at least a $(1-\varepsilon\beta)$-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 2.7 *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 2.6 (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 2.7, 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 2.6 it holds for each i that at least a $(1-\varepsilon\beta'/20)$-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 $(1-\varepsilon\beta'')$-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 $\varepsilon|F|$ non-jammed rounds in F, we get the following lemma, which also holds for arbitrary $(T, 1-\varepsilon)$-bounded adversaries.

Lemma 2.8 *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 2.5 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 2.2 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 2.2.

2.3.2 Proof of Lemma 2.6

This section is dedicated to the proof of Lemma 2.6 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, $S_1,...,S_6$. 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 2.2. 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 2.9 *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 $\mathscr{U}$ be the set of useful rounds in I for our node v. We distinguish between two cases, depending on the cardinality $|\mathscr{U}|$. 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 2.2.)

Case 1: Suppose that $|\mathscr{U}| < (c\ln n)/\gamma$, 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 $|\mathscr{U}| < (c\ln n)/\gamma$, 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,

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

Case 2: Next, suppose that $|\mathscr{U}| \geq (c\ln n)/\gamma$. 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 $\mathscr{U}'$. Consider any round $t \in \mathscr{U}'$. 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 1.1 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 $\mathscr{U}'$. Thus, it follows from the Chernoff bounds that for at least $2/3$ of the rounds in $\mathscr{U}'$, v will sense a successful transmission w.h.p. Hence, at the end of I it holds w.h.p. that

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

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

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 2.10 *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 2.9, 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 2.11 *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 1.1 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 2.12 *There is a $\phi \in [0,1]$ (depending on $p_S^{(0)}$, $p_S^{(1)}$ and $p_S^{(2)}$) so that*

$$\mathbb{E}[p'_S] \leq (1+\gamma)^{-\phi} p_S \quad \textit{and} \quad p'_S \leq (1+\gamma)^{2\phi} p_S. \tag{2.1}$$

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

$$((1/a)x_0 + (1/b)x_1 + (1/c)x_2)^2 (a^2 x_0 + b^2 x_1 + c^2 x_2) \leq 1 \tag{2.2}$$

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

Claim 2.13 *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 2.14 *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 (2.2) follows, which completes the proof.

∎

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 (2.1) 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 2.15 (Azuma Inequality) *Let $X_0, X_1, \ldots$ 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,

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

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. ∎

Lemma 2.16 *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 2.11 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 2.11 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 2.10 and Lemma 2.16 we can prove the following lemma.

Lemma 2.17 *For any sector S, at most $\varepsilon\beta/6$ 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 2.10 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 2.16 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 $\varepsilon\beta/7$ 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 $\varepsilon\beta/6$ of the subframes in F are bad w.h.p. ∎

From Lemma 2.17 it follows that apart from an $\varepsilon\beta$-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 2.6.

2.3.3 Proof of Theorem 2.3

Now, let us consider the two cases of Theorem 2.3 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 $(T, 1-\varepsilon)$-bounded adversary are satisfied. So Theorem 2.2 applies, which completes the proof of Theorem 2.3 a).

Case 2: $|D(v)| \geq 2/\varepsilon$ 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 $|J| > (1-\varepsilon/2)|I|$. Then every node in $D(u)$ must have more than $(\varepsilon/2)|I|$ 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 $|J| > \sum_{v \in D(u)} (\varepsilon/2)|I| \geq (2/\varepsilon) \cdot (\varepsilon/2)|I| = |I|$. Since this is impossible, it must hold that $|J| \leq (1-\varepsilon/2)|I|$.

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 $\sum_{v \in D(u)} f_v \geq (2/\varepsilon) \cdot \varepsilon|I| = 2|I|$. 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 2.2, which implies that JADE is constant competitive w.r.t. the nodes in V', completes the proof of Theorem 2.3 b).

2.3.4 Limitations under the Strong Adversary

One may ask whether a stronger throughput result than Theorem 2.3 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 $(T, 1-\varepsilon)$-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 2.4.

Lemma 2.18 *In general,* JADE *is not strongly c-competitive for a constant* $c > 0$ *if the* strong *adversary is allowed to be 2-uniform and* $\varepsilon \leq 1/3$.

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 $(T, 1-\varepsilon)$-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 [6] that if $T = \Omega((\log^3 n)/(\gamma^2 \varepsilon))$, then for every node $u \in U$, $T_u \leq \alpha \sqrt{T \log n / \varepsilon}$ 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.

Lemma 2.19 *In general,* JADE *is not c-competitive for a constant c independent of ε if there are nodes u with* $|D(u)| = o(1/\varepsilon)$ *and the* strong *adversary is allowed to be 2-uniform.*

Proof. Suppose that we have a set U of $k = o(1/\varepsilon)$ nodes located closely to each other that are all within the transmission range of a node v. Let $T = \Omega((\log^3 n)/(\gamma^2 \varepsilon))$. In each T-interval, the adversary never jams v but jams all but the first εT rounds at U. Then Section 2.3 of [6] 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 $\varepsilon T \cdot |U| + (\varepsilon T + O(T/\log n)) = \varepsilon T \cdot o(1/\varepsilon) + (\varepsilon + o(1))T = (\varepsilon + o(1))T$ messages in each T-interval on expectation whereas the sum of non-jammed rounds over all nodes is more than T. ∎

Hence, Theorem 2.3 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 $o(1/\varepsilon)$ nodes within the transmission range of a node.

2.4 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 $(1 - \varepsilon)$. 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 2.2, i.e., for $([T+(\log^3 n)/(\gamma^2\varepsilon)]\cdot(\log n)/\varepsilon$ rounds, where $\varepsilon=0.1$, $T=200$, and $\gamma=1/(\log T+\log\log N)$. Simulations with different combinations of $\varepsilon\in\{0.5,0.3,0.1\}$ and $T=\{50,100,150,200\}$ showed that $\varepsilon=0.1$ 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 2.1 (*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 in the 4×4 plane (density of 6.25), there are at least $6.25\pi\approx20\geq2/\varepsilon=20$ 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 2.1 (*bottom left*). The figure also illustrates the high correlation between success

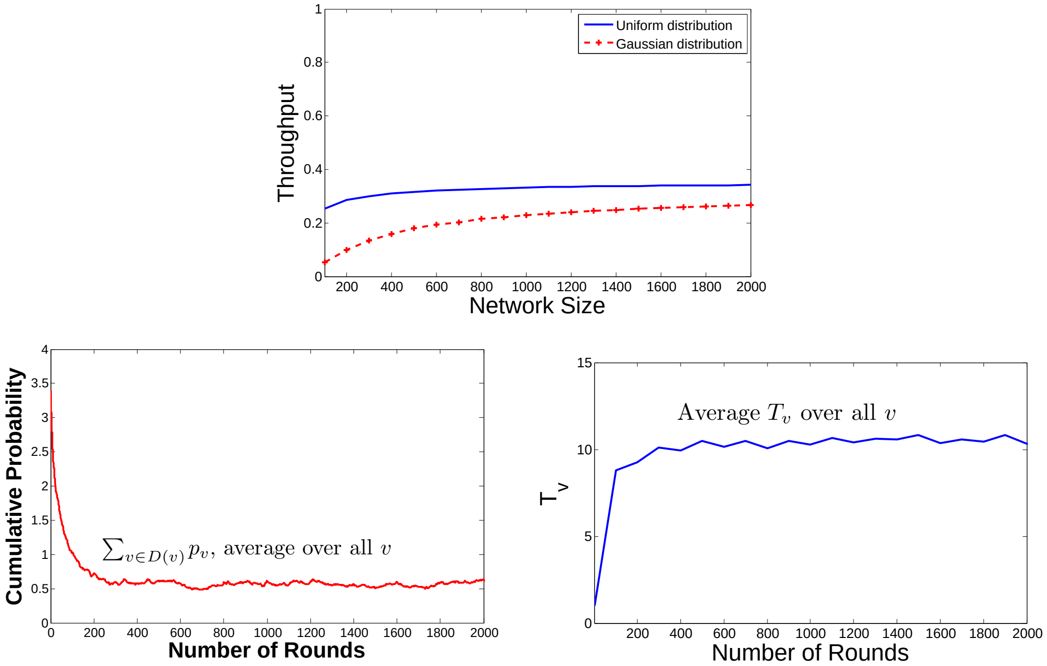

Figure 2.1: *Top:* Throughput as a function of network size, where $n \in [100, 2000]$, $\varepsilon = 0.1$, $T = 200$, and $\gamma = 1/(\log T + \log\log n)$. The result is averaged over 10 runs. *Bottom left:* Convergence behavior for multi-hop networks (uniform distribution). As a demonstration, we used $n = 500$, $\varepsilon = 0.1$, $T = 200$, 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 $n = 500$, $\varepsilon = 0.1$, $T = 200$, and $\gamma = 1/(\log T + \log\log N)$.

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 2.1 (*bottom right*).

2.5 Conclusion

To the best of our knowledge, JADE is the first jamming-resistant MAC protocol with provably good performance in multi-hop networks exposed to an adaptive but non-reactive 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. Also, there are several questions 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.

Chapter 3

THE ANTIJAM PROTOCOL

In this chapter, we study the problem of designing a robust MAC protocol that can achieve provably high competitive throughput despite a strong adaptive and reactive adversary.

The wireless network considered consists of n honest and reliable simple wireless devices (e.g., sensor nodes) that are within the transmission range of each other and which communicate over a single frequency (or a limited, narrow frequency band). We assume a back-logged scenario where the nodes continuously contend for sending a packet on the wireless channel. A node may either transmit a message or sense the channel at a time step, but it cannot do both, and there is no immediate feedback mechanism telling a node whether its transmission was successful. A node sensing the channel may either (i) sense an *idle channel* (in case no other node is transmitting at that time), (ii) sense a *busy channel* (in case two or more nodes transmit at the time step), or (iii) *receive* a packet (in case exactly one node transmits at the time step).

In addition to these nodes there is arbitrary external interference which we model as an adversary. Note that our notion of adversary is a model to describe external interference only; it does not, e.g., read and modify packet contents. We allow the adversary to know the protocol and its entire history (in terms of idle, busy, and successful transmission events) and to use this knowledge in order to jam the wireless channel at will at any time (i.e, the adversary is *adaptive*). Whenever it jams the channel, all nodes will notice a busy channel. However, the nodes cannot distinguish between the adversarial jamming or a collision of two or more messages that are sent at the same time.

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

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

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

Indeed, due to the large number of possible strategies a jammer can pursue, the problem becomes significantly more challenging than the non-reactive version. First, the analysis is more involved as the nodes' cumulative sending probability varies in a larger range depending on the adversarial strategy. Technically, the reactive jamming renders it impossible to apply Chernoff bounds over the non-jammed time periods as their patterns are no longer random; rather, we have to argue over *all* time periods. Second, modifications to the protocol in [6] are needed. For instance, the ANTIJAM protocol seeks to synchronize the nodes' sending probabilities; this has the desirable side effect of achieving fairness: all nodes are basically granted the same channel access probabilities, which greatly improves the unfair protocol of [6]. While our formal analysis confirms our expectations that the overall throughput under reactive jammers is lower than the throughput obtainable against non-reactive jammers, we are still able to prove a constant-competitive performance (for constant ε), which is also confirmed by our simulation study.

We study *competitive* MAC protocols.

Definition 3.1 (c-Competitive) *A MAC protocol is called c-competitive against some $(T, 1-\varepsilon)$-bounded adversary (with high probability or on expectation) if, for any sufficiently large number of time steps, the nodes manage to perform successful message transmissions in at least a c-fraction of the time steps not jammed by the adversary (with high probability or on expectation).*

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

Our goal is to design a *symmetric local-control* MAC protocol (i.e., there is no central authority controlling the nodes, and the nodes have symmetric roles at any point in time) that is fair and $O(1)$-competitive against any $(T, 1-\varepsilon)$-bounded adversary. The nodes do not know ε, but we do allow them to have a very rough upper bound of the number n and T. More specifically, we will assume that the nodes have a common parameter $\gamma = O(1/(\log T + \log\log n))$. As $\log T$ and $\log\log n$ are small for all reasonable values of T and n, this is scalable and not a critical constraint, as it leaves room for a super-polynomial change in n and a polynomial change in T over time.[1] Thus, all we need for our formal performance result to hold is a is very a rough upper bound on γ, and as we will see in our theorems there is a tradeoff between too low γ values (which causes the protocol to react too slowly to changes) and a too high γ values (with which the cumulative probability may overshoot). In practice we expect that choosing a constant, sufficiently small γ yields a good performance for any practical network; indeed, in our simulations $\gamma = 0.1$ results in a good throughput for a wide range of networks.

3.1 Contribution

This chapter presents a very simple medium access protocol called ANTIJAM. ANTIJAM is provably robust to a strong adaptive and reactive adversary that can block the medium a constant fraction of the time and thus models a large range of (intentional and unintentional) interference scenarios. Nevertheless, we can show that the ANTIJAM MAC protocol achieves a high throughput performance by exploiting any non-blocked time intervals effectively. The main theoretical contribution is a formal and rigorous derivation of the good throughput and fairness guarantees of our protocol. We show that ANTIJAM is competitive in the sense that a constant fraction of the non-jammed execution time is used for successful

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

transmissions, i.e., ANTIJAM is able to benefit from the rare and hard-to-predict time intervals where the shared medium is available. Our theoretical results are complemented by extensive simulations.

Theorem 3.2 *Let* $N = \max\{T, n\}$. *The* ANTIJAM *protocol is constant-competitive, namely* $e^{-\Theta(1/\varepsilon^2)}$*-competitive w.h.p., under any* $(T, 1-\varepsilon)$*-bounded* reactive *adversary if the protocol is executed for at least* $\Theta(\frac{1}{\varepsilon}\log N \max\{T, (e^{\delta/\varepsilon^2}/\varepsilon\gamma^2)\log^3 N\})$ *many time steps, where* $\varepsilon \in (0,1]$ *is a constant,* $\gamma = O(1/(\log T + \log\log n))$*, and where* δ *is a sufficiently large constant. Moreover,* ANTIJAM *achieves a high fairness: the channel access probabilities among nodes do not differ by more than a factor of* $(1+\gamma)$ *after the first message was sent successfully.*

3.2 Description of ANTIJAM

3.2.1 Intuition:

In the ANTIJAM protocol, each node v maintains a medium access probability p_v which determines the probability that v transmits a message in a communication round. The nodes adapt and synchronize their p_v values over time (which as a side-effect also improves fairness) in a multiplicative-increase multiplicative-decrease manner in order to ensure a throughput that is as good as possible. The p_v values tend to be lowered in times of high interference, and increased during times where the channel is idling. (This is similar to classic random backoff mechanisms where the next transmission time t is chosen uniformly at random from an interval of size $1/p_v$.) More precisely, the sending probabilities are changed by a factor of $(1+\gamma)$. However, we impose an upper bound of $\hat{p}$ on p_v, for some constant $0 < \hat{p} < 1/24$. As we will see, unlike in most classic backoff protocols, our adaption rules for p_v ensure that the adversary cannot influence the p_v values much by jamming.

In addition, each node maintains two variables, a threshold variable T_v and a counter variable c_v. 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 $T_v := 1$, $c_v := 1$ and $p_v := \hat{p}$; however, as we will see, ANTIJAM works for arbitrary variable values. Afterwards, the protocol works in synchronized time steps. We assume synchronized time steps for the analysis, but a nonsynchronized execution of the protocol would also work as long as all nodes operate at roughly the same speed.

ANTIJAM is based on the intuition presented in Section 1.4.1. However, since ANTIJAM aims to be jamming-resistant against an adaptive and reactive adversary, in order to still achieve constant cumulative probability, not only does the protocol need to use a multiplicative increase/decrease game for the probabilities p_v, but also it synchronizes all the nodes, both in terms of sending probabilities and their own estimates on the time window threshold estimate T_v's, at every successful transmission.

3.2.2 Protocol Description:

With these definitions and insights, we can now formally present the ANTIJAM protocol, see Algorithm 2.

A summary of all our variables (including the ones from the analysis) is provided in Table 3.1.

The most significant change in ANTIJAM compared to the protocol in [6] is that the nodes synchronize everything: (i) their p_v, c_v, and T_v values whenever a message is successfully received, and (ii) T_v is decreased only when the channel is idle, since idle channel is experienced by all the nodes. The reason that we seek synchronization is that the adversary we consider here is much stronger, i.e., adap-

Algorithm 2 ANTIJAM: for each node v

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

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

Table 3.1: Important Variables

tive and reactive, and hence could dramatically affect the cumulative probability of the network. By synchronizing the nodes, we could greatly simplify the proofs, and manage to show constant competitive throughput can still be achieved. As a by-product, AntiJam achieves fairness, as any two nodes' access probabilities do not differ by more than $(1+\gamma)$ factor.

3.3 Analysis

Our analysis of Theorem 3.2 unfolds in a number of lemmas. We first show that given a certain initial cumulative sending probability p, p stays high in the future, i.e., it cannot drop below this initial probability over time (Lemma 3.6). Lemma 3.10 then shows that a sufficiently large initial cumulative sending probability p implies a good throughput in time intervals where p often remains below a certain threshold. Finally, Lemma 3.13 proves that with high probability, p indeed does not increase beyond a certain threshold.

The analysis makes repeated use of Lemma 1.3 and the Chernoff bounds in Lemma 1.2.

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

We start with some simple facts which also provide some intuition for ANTIJAM. Fact 3.3 states that the protocol synchronizes the sending probabilities of the nodes (up to a factor of $(1+\gamma)$) as well as the values c_v and T_v.

Fact 3.3 *Right after a successful transmission of the triple* (p', c', T'), $(p_v, c_v, T_v) = ((1+\gamma)^{-1}p', c', T')$ *for all receiving nodes* v *and* $(p_u, c_u, T_u) = (p', c', T')$ *for the sending node* u*. In particular, for any time step* t *after a successful transmission by node* u, $(c_v, T_v) = (c_w, T_w)$ *for all nodes* $v, w \in V$.

Fact 3.3 also implies the following corollary.

Corollary 3.4 *After a successful transmission, the access probabilities* p_v *of the nodes* $v \in V$ *will never differ by more than a factor* $(1+\gamma)$ *in the future.*

The following facts study how the cumulative sending probability varies over time depending on the different events.

Fact 3.5 *For any time step* t *after a successful transmission or a well-initialized state of the protocol (in which* $(p_v, c_v, T_v) = (\hat{p}, 1, 1)$ *for all nodes* v*) it holds:*

1. *If the channel is* idle *at time* t *then* (i) *if* $p_v = \hat{p}$ *for all* v*, then* $p_{t+1} = p_t$*;* (ii) *if* $p_u = \hat{p}$ *and* $p_v = (1+\gamma)^{-1}\hat{p}$ *for all nodes* $v \neq u$*, then* $p_{t+1} = (1+\gamma - O(1/n))p_t$

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

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

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

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

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

Lemma 3.6 *For any subframe* I' *in which initially* $p_{t_0} \geq 1/(f^2(1+\gamma)^{\sqrt{2f}})$*, the last time step t of* I' *again satisfies* $p_t \geq 1/(f^2(1+\gamma)^{\sqrt{2f}})$*, w.h.p.*

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

Claim 3.7 *If in subframe I' the number of idle time steps is at most k, then every node v increases T_v by 2 at most $k/2+\sqrt{f}$ many times.*

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

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

Claim 3.8 *Suppose that for the first time step t_0 in I', $p_{t_0} \in [1/(f^2(1+\gamma)^{\sqrt{2f}}), 1/f^2]$. Then there is a time step t in I' with $p_t \ge 1/f^2$, w.h.p.*

Proof. Suppose that there are g non-jammed time steps in I'. Let k_0 be the number of these steps with an idle channel and k_1 be the number of these steps with a successful message transmission. Furthermore, let k_2 be the maximum number of times a node v increases T_v by 2 in I'. If all time steps t in I' satisfy $p_t < 1/f^2$, then it must hold that $k_0 - \log_{1+\gamma}(1/p_{t_0}) \le k_1 + k_2$. This is because no v has reached a point with $p_t(v) = \hat{p}$ in this case, so Fact 3.5 implies that for each time step t with an idle channel, $p_{t+1} = (1+\gamma)p_t$. Thus, at most $\log_{1+\gamma}(1/p_{t_0})$ time steps with an idle channel would be needed to get p_t to $1/f^2$, and then there would have to be a balance between further increases (that are guaranteed to be caused by an idle channel) and decreases (that might be caused by a successful transmission or the case $c_v > T_v$) of p_t in order to avoid the case $p_t \geq 1/f^2$. The number of times we can allow an idle channel is maximized if all successful transmissions and cases where $c_v > T_v$ cause a reduction of p_t. So we need $k_0 - \log_{1+\gamma}(1/p_{t_0}) \leq k_1 + k_2$ to hold to avoid the case $p_t \geq 1/f^2$ somewhere in I'.

We know from Claim 3.7 that $k_2 \leq k_0/2 + \sqrt{f}$. Hence,

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

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

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

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

Suppose for the moment that no time step is jammed in I'. Then $\mathbb{E}[k_1] \leq (1/f^2)f = 1/f$. In order to prove a bound on k_1 that holds w.h.p., we can use the general Chernoff bounds stated above. For any step t, let the binary random variable X_t be 1 if and only if at least one message is transmitted at time t and $p_t \leq 1/f^2$. Then

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

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

that, if there are multiple message transmissions and since $p_t \leq 1/f^2$,

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

Then, we have

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

and

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

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

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

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

Claim 3.9 *Suppose that for the first time step t_0 in I', $p_{t_0} \geq 1/f^2$. Then there is no time step t in I' with $p_t < \frac{1}{f^2(1+\gamma)^{\sqrt{2f}}}$, w.h.p.*

Proof. Consider some fixed time step t in I' and let $I'' = (t_0, t]$. Suppose that there are g non-jammed time steps in I''. If $g \leq \beta \log N$ for a (sufficiently large) constant β, then it follows for the probability p_t at the end of I'' due to Claim 3.7 that

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

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

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

because of our assumption that $p_{t_0} \geq 1/f^2$ (more idle rounds would yield higher p_t values).

Since $k_2 \leq k_0/2 + \sqrt{f}$, this implies that $k_0 \leq 2k_1 + 2\sqrt{f} \leq 2k_1 + g/4$. Thus, we are back to the case in the proof of Claim 3.8, which shows that $k_0 \leq 2k_1 + g/4$ does not hold w.h.p., given that $g > \beta \log N$ and we never have the case in I'' that $p_t > 1/f^2$.

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

Combining Claims 3.8 and 3.9 completes the proof of Lemma 3.6.

Lemma 3.10 establishes an important relationship between cumulative sending probability and throughput.

Lemma 3.10 *Consider any subframe I', and let $\delta > 1$ be a sufficiently large constant. Suppose that at the beginning of I', $p_{t_0} \geq 1/(f^2(1+\gamma)^{\sqrt{2f}})$ and $T_v \leq \sqrt{F}/2$ for every node v. If $p_t \leq \delta/\varepsilon^2$ for at least half of the non-jammed time steps in I', then* ANTIJAM *is at least $\frac{\delta}{8\varepsilon^2} e^{-\delta/(1-\hat{p})\varepsilon^2}$-competitive in I'.*

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

Claim 3.11 *If $n \geq (1+\gamma)\delta/(\varepsilon^2 \hat{p})$, then*

$$k_0 - \log_{1+\gamma}(\delta/(\varepsilon^2 \cdot p_{t_0})) \leq k_1' + k_2$$

where k_1' is the number of useful time steps with a successful transmission in which the sender is different from the previously successful sender.

Proof. According to Corollary 3.4, if $p_t \leq \delta/\varepsilon^2$ and $n \geq (1+\gamma)\delta/(\varepsilon^2\hat{p})$, then $p_v(t) \leq \hat{p}/(1+\gamma)$. This implies that whenever there is a useful time step $t \in I$ with an idle channel, then $p_{t+1} = (1+\gamma)p_t$. Thus, it takes at most $\log_{1+\gamma}(\delta/(\varepsilon^2 \cdot p_{t_0}))$ many useful time steps with an idle channel to get from p_{t_0} to a cumulative probability of at least δ/ε^2. On the other hand, each of the k_1' successful transmissions reduces the cumulative probability by a factor of $(1+\gamma)$. Therefore, once the cumulative probability is at δ/ε^2, we must have $k_0 \leq k_1' + k_2$ since otherwise there must be at least one useful time step where the cumulative probability is more than δ/ε^2, which contradicts the definition of a useful time step.

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

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

From Lemma 3.7 we also know that $k_2 \leq k_0/2 + \sqrt{f}$. Hence,

$$\begin{aligned} k_0 &\leq 2k_1' + 2 \cdot \log_{1+\gamma}(\delta f^2/\varepsilon^2) + 2 \cdot (\sqrt{f} + \sqrt{2f}) \\ &\leq 2k_1' + 6\sqrt{f} \end{aligned}$$

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

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

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

Claim 3.12 *Let g be the number of non-jammed time steps t in I' with $p_t \leq \delta/\varepsilon^2$. If $g \geq \varepsilon f/2$ then*

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

w.h.p.

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

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

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

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

Thus,

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

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

Since our lower bound for the probability of a non-idle step to be useful holds independently for all non-jammed non-idle steps t with $p_t \leq \delta/\varepsilon^2$ and $E[k] \geq$

$\alpha \log N$ for our choice of g, it follows from the Chernoff bounds that $k \geq \mathbb{E}[k]/2$ w.h.p.

From Claim 3.12 it follows that

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

w.h.p., which completes the proof of Lemma 3.10: if we divide the lower bound on k_1 by the number of non-jammed time steps εf (as $g \geq \varepsilon f/2$, $k_1 \geq k/3 - 2\sqrt{f}$ and as $-2\sqrt{f}$ is negligible).

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

Lemma 3.13 *Suppose that at the beginning of I', $T_v \leq \sqrt{F}/2$ for every node v. Then at most half of the non-jammed time steps t can have the property that $p_t > \delta/\varepsilon^2$ w.h.p.*

Proof. Recall from Fact 3.5 that as long as the access probabilities of the nodes do not hit $\hat{p}$, the cumulative probability only changes by a $(1+\gamma)$-factor in both directions. Suppose that δ is selected so that δ/ε^2 represents one of these values. Let H be the set of time steps $t \in I'$ with the property that either $p_t = \delta/\varepsilon^2$ and the channel is idle or $p_t \geq (1+\gamma)\delta/\varepsilon^2$. Now, we define a step t to be *useful* if $t \in H$ and there is either an idle channel or a successful transmission at t. Let k be the number of useful time steps in H. Furthermore, let k_0 be the number of useful time steps with an idle channel, k_1 be the number of useful time steps with a successful transmission and k_2 be the maximum number of times a node v reduces p_v in H because of $c_v > T_v$. It holds that $k = k_0 + k_1$.

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

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

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

Claim 3.14 *For any collection P of passes, w.h.p., $k_0 \geq k_1 - \Delta - \Theta(1)$ where k_0 and k_1 are the number of idle time steps and the number of successful transmissions in P.*

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

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

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

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

Then, let $E_i = 1$ denote the event that the sender of the i-th successful transmission is the same as the sender of the previous successful transmission. We show that the probability that $\sum_i E_i \geq c$ (c is a constant) given k_1 is extremely small. According to Corollary 3.4, the nodes' access probabilities do not differ by more than a $(1+\gamma)$-factor after the first successful transmission. Hence, each node has almost the same probability of transmitting a message at any given time step, which implies that $\mathbb{P}[E_i = 1] \leq (1+\gamma)/n$.

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

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

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

Claim 3.15

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

where k is the number of useful steps in H.

Proof. If at the beginning of I', $T_v \leq \sqrt{F}/2$ for every node v, then according to Claim 3.7, $T_v \leq \sqrt{F}$ for every node v at any time during I'. Hence, after at most

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

For the calculations below recall the definition of f with the constants α and β that are assumed to be sufficiently large. If $k \leq \alpha \log N$, then it follows from Claim 3.15 that, for large enough δ,

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

where $N = \max\{n, T\}$. Thus, the number of non-jammed time steps in H is also at most $\varepsilon f/\beta$, and since β can be arbitrarily large, Lemma 3.13 follows, as the steps in H fulfill this property ($\beta \geq 2$ yields half of the steps).

It remains to consider the case that $k > \alpha \log N$. Let us assume that H contains at least $\varepsilon f/2$ non-jammed time steps, otherwise the claim certainly holds. Our goal is to contradict that statement in order to show that the lemma is true. For this we will show that Claim 3.14 is violated w.h.p.

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

For a cumulative probability $p \geq \Phi$, $\mathbb{P}[\text{idle} \mid p] \leq e^{-\Phi} = (\frac{\log N}{f})^{1-\hat{p}}$ and $\mathbb{P}[\text{success} \mid p] \leq \frac{\Phi}{1-\hat{p}} \cdot e^{-\Phi} \leq \ln(f/\log N) \cdot (\frac{\log N}{f})^{1-\hat{p}}$. Hence, by multiplying these probabilities by the $|H| \leq f$ steps, we get that $k \leq f^{\hat{p}} \cdot \ln f \cdot \log^{1-\hat{p}} N$ on expectation, and from the Chernoff bounds it follows that $k \leq 2f^{\hat{p}} \cdot \ln f \cdot \log^{1-\hat{p}} N$ w.h.p., so Claim 3.15 implies that the number of time steps in I' with cumulative probability

$p \geq \Phi$ is at most

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

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

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

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

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

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

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

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

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

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

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

It implies that w.h.p.,

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

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

This is a contradiction to Claim 3.14, and hence completes the proof of Lemma 3.13.

In order to proceed, we need the following claim.

Claim 3.16 *For any collection P of passes it holds that*

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

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

Proof. Because of Fact 3.5, the probability that a successful transmission is done by a node different from the node of the last successful transmission is equal to

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

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

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

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

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

Lemma 3.17 *For any time frame I in which initially $T_v \leq \sqrt{F}/2$ for all v, also $T_v \leq \sqrt{F}/2$ for all v at the end of I w.h.p.*

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

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

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

$$k_0 \geq k_1' + k_2 - \log_{1+\gamma}((\delta/\varepsilon^2) \cdot f^2(1+\gamma)^{\sqrt{f}}).$$

Because of Claims 3.7 and 3.16 it follows that

$$k_0 \geq k_1/3$$

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

With Lemma 3.17, we show that Lemma 3.13 is true for a polynomial number of subframes. Then, Lemma 3.13 and Lemma 3.17 together imply that Lemma 3.10 holds for a polynomial number of subframes. Hence, our main Theorem 3.2 follows. Along the same line as in [6], we can show that ANTIJAM is self-stabilizing, so the throughput result can be extended to an arbitrary sequence of time frames.

3.4 Simulation

We have implemented a simulator to study additional properties of our protocol and to complement our formal insights. Our focus here is on the qualitative nature of the performance of ANTIJAM, and we did not optimize the parameters to obtain the best constants. We consider three different jamming strategies for a reactive jammer that is $(T, 1-\varepsilon)$-bounded, for different ε values and where $T = 100$: (1) one that jams non-idle steps with probability $(1-\varepsilon)$; (2) one that jams non-idle steps deterministically (as long the jamming budget is not used up); (3) one that jams idle steps deterministically (as long as the jamming budget is not used up). Intuitively, it seems that jamming non-idle steps is more harmful than jamming idle steps. However, note that jamming idle steps may be an effective strategy to steer the protocol into bad states; moreover, it may capture scenarios where nodes in co-existent networks start sending in quiet times.

We define throughput as the number of successful transmissions over the number of non-jammed time steps. Moreover, for networks larger than 100, we choose $\hat(p) = 1/24$, whereas for smaller networks we choose $\hat(p) = 1/2$. As a general guideline, it is always better to choose larger $\hat(p)$ values, as this avoids capping the throughput in small networks artificially. A smaller $\hat(p)$ can make sense for bootstrapping large networks, but due to the fast convergence times of the protocol (see Section 3.4.2), this is unproblematic.

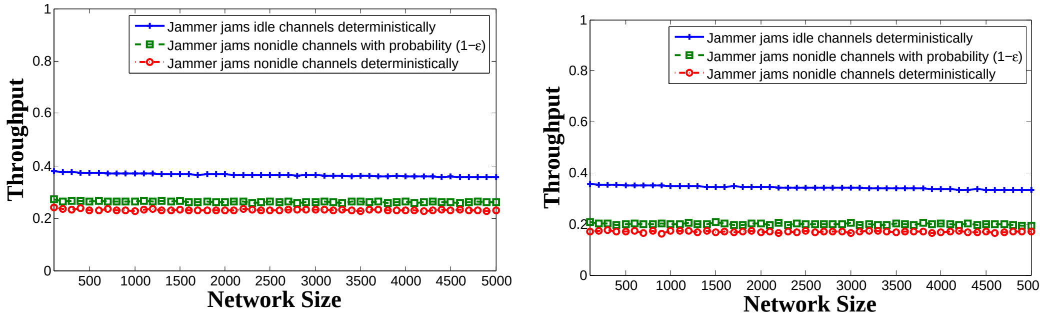

Figure 3.1: Throughput under three different jamming strategies as a function of the network size (large) and ε, where $\hat{p} = 1/24$ (averaged over 10 runs) (*left:* $\varepsilon = 0.5$, *right:* $\varepsilon = 0.3$)

.

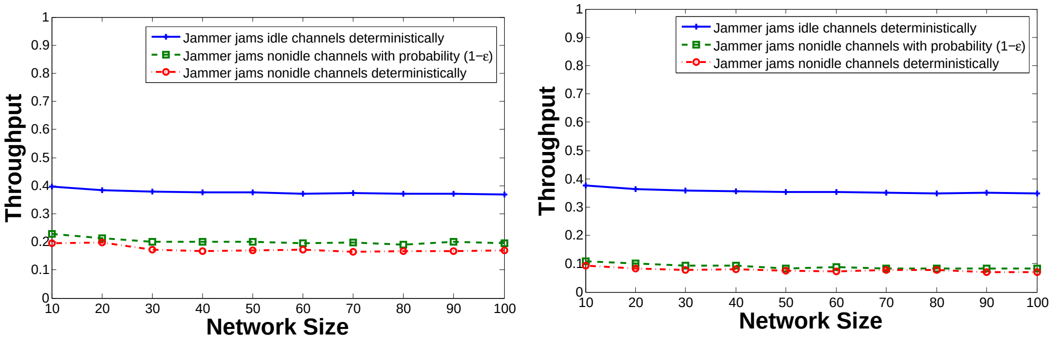

Figure 3.2: Throughput under three different jamming strategies as a function of the network size (small) and of ε, where $\hat{p} = 1/2$ (averaged over 10 runs) (*left:* $\varepsilon = 0.5$, *right:* $\varepsilon = 0.3$)

3.4.1 Throughput

In a first set of experiments we study the throughput as a function of the network size and ε. We evaluate the throughput performance for each type of adversary introduced above, see Figure 3.1. For all three strategies, the throughput is basically constant, independently of the network size; this is in accordance with our theoretical insight of Theorem 3.2. We can see that given our conditions on ε and T, the strategy that jams non-idle channels deterministically results in the lowest throughput. Hence, in the remaining experiments described in this section, we will focus on this particular strategy. As expected, jamming idle channels does not affect the protocol behavior much. In our simulations, ANTIJAM makes effective use of the non-jammed time periods, yielding $20\% - 40\%$ successful transmissions even without optimizing the protocol parameters. Having shown the protocol scales well for large network size, we also study the throughput results when the network size is *small*, see Figure 3.2. We observe that the results for small and large scale networks are comparable, but the throughput in the small scale networks can be slightly lower under an adversary that jams non-idle channels deterministically or with probability $(1-\varepsilon)$.

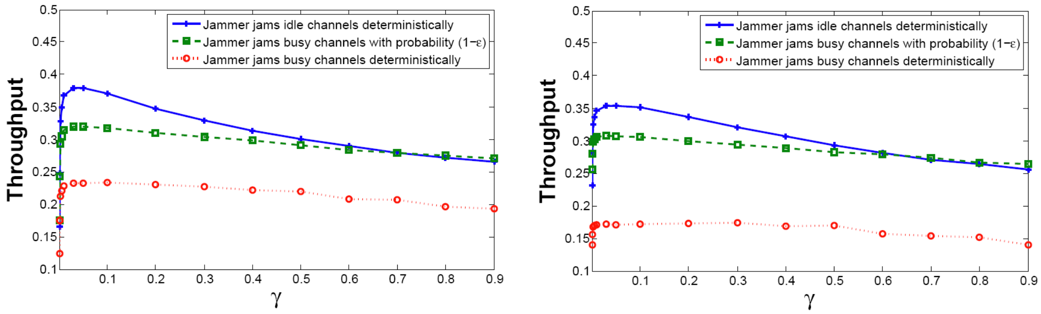

Figure 3.3: Throughput as a function of γ under three different jamming strategies, when $n = 1000$, and results are averaged over 10 runs

(*left:* $\varepsilon = 0.5$, *right:* $\varepsilon = 0.3$).

In additional experiments we also studied the throughput as a function of γ, see Figure 3.3. As expected, the throughput declines slightly for large γ, but this effect is small. (Note that for very small γ, the convergence time becomes large and the experiments need run for a long time in order not to underestimate the real throughput.)

3.4.2 Convergence Time

Besides a high throughput, fast convergence is the most important performance criterion of a MAC protocol. The traces in Figure 3.4 (*top left*) show the evolution of the cumulative probability over time. It can be seen that the protocol converges quickly to constant access probabilities. (Note the logarithmic scale.) If the initial probability for each node is high, the protocol needs more time to bring down the low-constant cumulative probability. Moreover, the ratio of the time period the cumulative probability is in the range of $[\frac{1}{2\varepsilon}, \frac{2}{\varepsilon}]$ to the time period the protocol being executed is 92.98% when $\hat{p} = 1/24$, and 89.52% when $\hat{p} = 1/2$. This implies that for a sufficiently large time period, the cumulative probability is well bounded most of the time, which corresponds to our theoretical insights (cf Lemma 3.6 and 3.13). Figure 3.4 (*top right*) studies the convergence time for different network sizes. We ran the protocol 50 times, and assume that the execution has converged when the cumulative probability p satisfies $p \in [1,5]$, for at least 5 consecutive rounds. The simulation result also qualitatively confirms our theoretical analysis in Theorem 3.2, as the number of rounds needed to converge the execution is bounded by $\Theta(\frac{1}{\varepsilon} \log N \max\{T, \frac{1}{\varepsilon\gamma^2} \log^3 N\})$. (Of course, the concrete convergence time can depend on the scenario, and may be faster than expected in the general case.)

Figure 3.4 (*bottom left*) indicates that independently of the initial values $\hat{p}$ and T_v, the throughput rises quickly (up above 20%) and stays there afterwards.

3.4.3 Fairness

As ANTIJAM synchronizes c_v, T_v, and p_v values upon message reception, the nodes are expected to transmit roughly the same amount of messages; in other words, our protocol is fair. Figure 3.4 (*bottom right*) presents a histogram showing how the successful transmissions are distributed among the nodes. More specifically, we

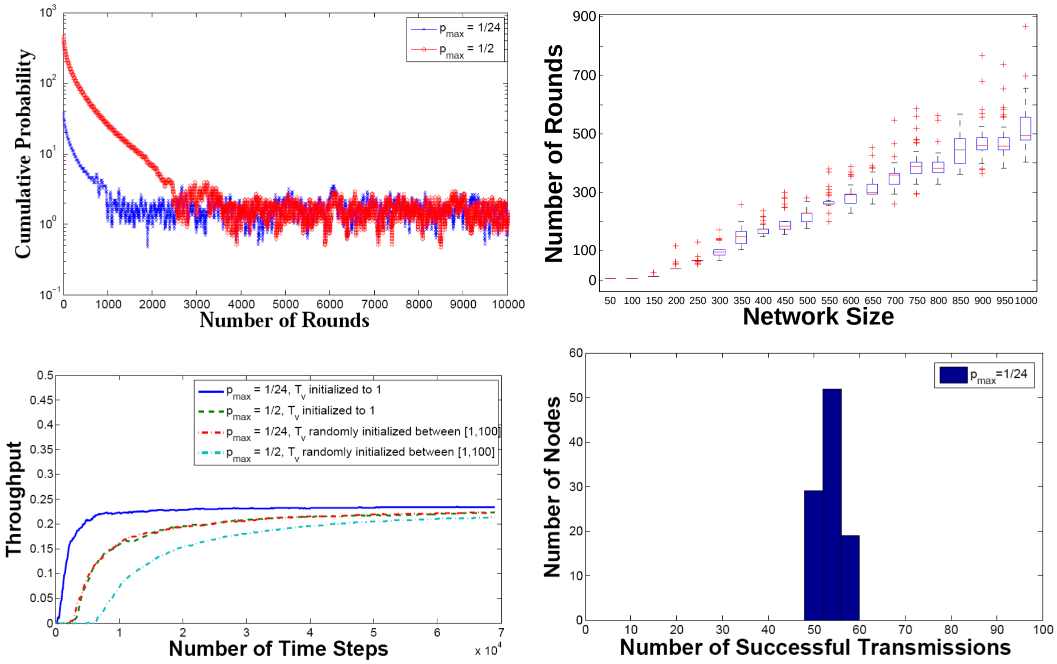

Figure 3.4: *Top left:* Evolution of cumulative probability over time (network size is 1000 nodes, and $\varepsilon = 0.5$). Note that the plot has logarithmic scale. *Top right:* Boxplot of ANTIJAM runtime as a function of network size for $\hat{p} = 1/24$, and $\varepsilon = 0.5$. *Bottom left:* Convergence in a network of 1000 nodes where $\varepsilon = 0.5$. *Bottom right:* Fairness in a network of 1000 nodes, where $\varepsilon = 0.5$, and $\hat{p} = 1/24$ (averaged over 10 runs).

partition the number of successful transmissions into intervals of size 4. Then, all the transmissions are grouped according to those intervals in the histogram.

3.4.4 Comparison

Finally, to put ANTIJAM into perspective, as a comparison, we implemented the MAC protocol proposed in [6], as well as a simplified version of the widely used 802.11 MAC protocol (with a focus on 802.11a).

The configurations for the simulation are the following: (1) the jammer is reactive and $(T, 1-\varepsilon)$-bounded; (2) the unit slot time for 802.11 is set to $50\mu s$; for simplicity, we define one time step for ANTIJAM to be $50\mu s$ also; (3) we run

ANTIJAM, the MAC protocol in [6], and 802.11 for 4 min, which is equal to $4.8 \cdot 10^6$ time steps in our simulation; (4) the backoff timer of the 802.11 MAC protocol implemented here uses units of $50\mu s$; (5) we omit SIFS, DIFS, and RTS/CTS/ACK.

A comparison is summarized in Figure 3.5. The throughput achieved by ANTIJAM and the MAC protocol in [6] are significantly higher than the one by the 802.11 MAC protocol, specially for lower values of ε, when the 802.11 MAC protocol basically fails to deliver any successful message. Note that the throughput results between ANTIJAM and the MAC protocol in [6] are similar in the simulations, but ANTIJAM is slightly better for the most ε.

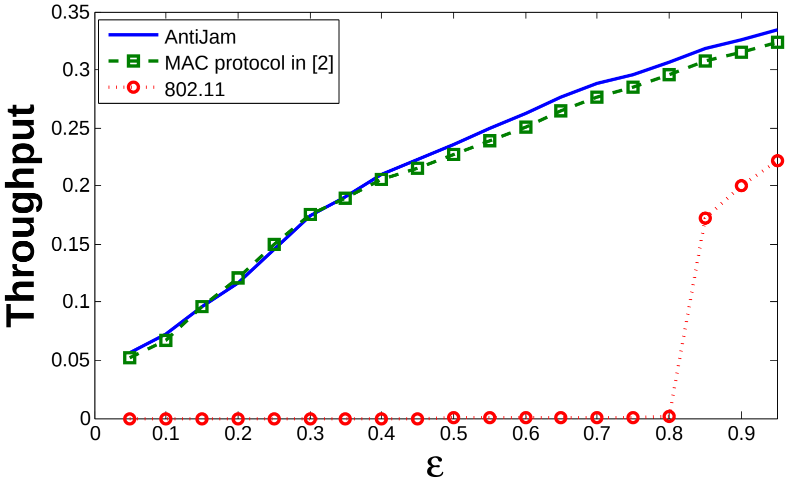

Figure 3.5: Throughput as a function of $\varepsilon \in [0.05, 0.95]$, compared to the MAC protocol in [6] and

802.11, averaged over 10 runs, where $\hat{p} = 1/24$.

3.5 Conclusion

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

Chapter 4

THE SELECT PROTOCOL

In this chapter, we consider the problem of designing a self-stabilizing distributed protocol to elect a leader among a set V of n simple wireless nodes (e.g., nodes of a sensor network) that are within each other's transmission range and communicate over a single channel. For our formal analysis, we assume that the time proceeds in synchronous *rounds* (or *steps*).[1] The general communication model specified earlier (in 1.2.2) applies to SELECT also.

In addition to these nodes there is an adversary. We allow the adversary to know the protocol and its entire history and to use this knowledge in order to jam the wireless channel at will at any round. Such an adversary is called *adaptive*. If in addition to that the adversary also knows (through physical carrier sensing) the *current* channel state, we call it *reactive*. That is, a reactive adversary can distinguish between the channel being currently idle (no node transmits) or busy (either because of a successful transmission, a collision of transmissions, or too much background noise) and can instantly make a jamming decision based on that information. Whenever the adversary jams the channel, all nodes will notice a busy channel. The nodes cannot distinguish between the adversarial jamming and a collision of two or more messages that are sent at the same time.

In order to study the degree of jamming activity needed by the adversary to prevent successful message transmissions, we use the notion of a $(T, 1-\varepsilon)$-bounded adversary. An adversary is called $(T, 1-\varepsilon)$-*bounded* for some $T \in N$ and $0 < \varepsilon < 1$ if for any time window of size $w \geq T$ the adversary can jam at most $(1-\varepsilon)w$ of the time steps in that window. Moreover we assume that the n

[1] A round may represent the time needed to send a message, e.g., a multiple of the 50μs unit in 802.11, depending on the message size.

nodes use an encryption mechanism that prevents the adversary from inspecting their messages.

As mentioned earlier, our goal is to design a leader election protocol that is self-stabilizing despite adversarial jamming. Following the usual notation in the self-stabilization literature, the *system state* is determined by the state of all *variables* in the system. That is, the protocol and any constants used by the protocol are assumed to be immutable and not part of the system state. A system is called *self-stabilizing* if and only if (1) when starting from any state, it is guaranteed to eventually reach a legal state (convergence) and (2) given that the system is in a legal state, it is guaranteed to stay in a legal state (closure), provided that there are no faults or membership changes in the system. In our case, roughly speaking, the legal state is the state in which we have exactly one leader. We will define the set of legal states more formally when we introduce our protocol. While our protocol is randomized and the leader election has to be performed under adversarial jamming, our protocol is still guaranteed to eventually elect exactly one leader from any initial state.

4.1 Contribution

This chapter presents SELECT ("SElf-stabilizing Leader EleCTion"), a protocol that solves the leader election problem in harsh environments—namely in wireless networks under adversarial *reactive* jamming—and in a self-stabilizing manner, independently of the initial network state. We believe that self-stabilization is a crucial feature in real networks where membership is often dynamic. Although our algorithm is randomized, we will present a formal proof that its correctness holds deterministically. Moreover, while our analysis is rather involved, the SELECT protocol itself is simple and hence easy to implement.

Concretely, in this chapter we will derive the following theorem.

Theorem 4.1 *Given an arbitrary initial configuration and in the absence of state faults, our leader election protocol reaches a state where there is exactly one leader and $n-1$ followers, despite a reactive $(T, 1-\varepsilon)$-bounded jammer, for any T and any constant $\varepsilon > 0$.*

In SELECT, the nodes do not have to know anything about the system for the protocol to work. The only assumption that we need is that some fixed common parameter γ used by the nodes satisfies $\gamma = O(1/(\log T + \log\log n))$. As $\log T$ and $\log\log n$ are small for all reasonable values of T and n, this is scalable and not a critical constraint, as it leaves room for a super-polynomial change in n and a polynomial change in T over time.[2] Thus, in practice we expect that choosing γ to be a sufficiently small constant yields a good performance for any practical network, which is confirmed by our simulations.

4.2 The SELECT Protocol

4.2.1 Intuition:

SELECT is based on the following idea. Each node v maintains a parameter p_v which describes v's probability of accessing the medium at a given moment of time. That is, in each round, each node v decides to transmit a message with probability p_v (e.g., in an attempt to become a leader). (This is similar to classic random backoff mechanisms where the next transmission time t is chosen uniformly at random from an interval of size $1/p_v$.) The nodes adapt and synchronize their p_v 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_v will never exceed $\hat{p}$, for some constant $0 < \hat{p} < 1$.

[2] On the other hand, note that the assumption that the nodes know constant factor approximations of n or T directly would render the problem trivial. Moreover, such an assumption is unrealistic and non-scalable.

Algorithm 1 Leader Election: Follower

```
$mc := c_v \mod b$
if $mc = 0$ then
    $ls_1 := ls'_0, ls_2 := ls'_1, ls_3 := ls'_2, ls_4 := ls'_3$
    $s_v := s'_v$
end if
if ($ls_3 =$ undefined) or ($mc \neq ls_1$ and $mc \neq ls_2$ and $mc \neq ls_3$ and $mc \neq ls_4$) then
    $v$ decides with $p_v$ to send a follower message
    if $v$ sends a follower message then
        the message contains:
        $cc_1 := ls'_0, cc_2 := ls'_1, cc_3 := ls'_2, cc_4 := ls'_3, c_{new} := c_v, T_{new} := T_v, p_{new} := p_v$
    end if
end if
if $v$ does not send a follower message then
    $v$ senses the channel
    if channel is idle then
        if $mc = ls_3$ then
            $s'_v := 1$
            $p_v := \hat{p}$
        else
            $p_v := \min\{(1+\gamma)p_v, \hat{p}\}$
        end if
    else if $v$ receives 'LEADER' then
        $s'_v := 0$
        $ls_3 := undefined$
        $ls'_2 := undefined$
    else if $v$ receives a tuple of $\{cc_1, cc_2, cc_3, cc_4, c_{new}, T_{new}, p_{new}\}$ then
        $T_v := T_{new}$
        $p_v := (1+\gamma)^{-1}p_{new}$
        $c_v := c_{new}$
        $ls'_0 := random(0, b-1)$
        $ls'_1 := cc_1, ls'_2 := cc_2, ls'_3 := cc_3, ls'_4 := cc_4$
    end if
end if
$c_v := c_v + 1$
if $c_v \geq b \cdot T_v$ then
    $c_v := 0$
    if (not CONDITION) then
        $p_v := (1+\gamma)^{-1}p_v, T_v := T_v + 1$
        $ls'_0 := undefined, ls'_1 := undefined, ls'_2 := undefined, ls'_3 := undefined, ls'_4 := undefined$
    else
        $T_v := \max\{T_v - 1, 4\}$
    end if
end if
```

Algorithm 2 Leader Election: Leader

```
$mc := c_v \mod b$
if $mc = 0$ then
    $ls_1 := ls'_1, ls_2 := ls'_2, ls_3 := ls'_3, ls_4 := ls'_4$
end if
if $mc = ls_1$ or $mc = ls_2$ or $mc = ls_3$ or $mc = ls_4$ then
    $v$ sends the leader message 'LEADER'
else
    $v$ decides with $p_v$ to send 'LEADER'
    if $v$ does not send 'LEADER' then
        $v$ senses the channel
        if channel is idle then
            $p_v := \min\{(1+\gamma)^2 p_v, \hat{p}\}$
        else if $v$ receives a message then
            $p_v := (1+\gamma)^{-1}p_v$
            if message is 'LEADER' then
                $s_v := 0, s'_v := 0$
                $ls_3 := undefined, ls'_2 := undefined$
            else if message is a follower message, i.e., a tuple of $\{cc_1, cc_2, cc_3, cc_4, c_{new}, T_{new}, p_{new}\}$ then
                $c_v := c_{new}, T_v := T_{new}$
                $ls'_1 := cc_1, ls'_2 := cc_2, ls'_3 := cc_3, ls'_4 := cc_4$
            end if
        end if
    end if
end if
$c_v := c_v + 1$
if $c_v \geq b \cdot T_v$ then
    $c_v := 0$
    if (not CONDITION) then
        $p_v := (1+\gamma)^{-1}p_v, T_v := T_v + 1$
        $ls'_0 := undefined, ls'_1 := undefined, ls'_2 := undefined, ls'_3 := undefined, ls'_4 := undefined$
    else
        $T_v := \max\{T_v - 1, 4\}$
    end if
end if
```

Figure 4.1: Algorithm for followers (*left*) and leaders (*right*).

In addition, each node maintains two variables, a threshold variable T_v and a counter variable c_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.

SELECT distinguishes between two node roles: *follower* and *leader*. We use s_v to indicate the role of the node: $s_v = 1$ means that node v is a leader, whereas $s_v = 0$ means v is a follower. The basic idea of our protocol is to divide time into intervals of a small number of rounds specified by the constant parameter $b > 5$ (we use the variable mc as a modulo counter); in the following, we will refer to a sequence of rounds between two consecutive $mc = 0$ events as a *b-interval*. (Of course, it can happen that all b slots of an interval are jammed.)

Our protocol is based on the concept of so-called *leader slots*, special rounds—in each b-interval through which SELECT cycles—in which leaders are obliged to send an alive message (a so-called *leader message*) and in which followers keep silent. The idea is that the followers learn that the leader has left in case of an idling medium during a leader slot (of course, the leader slots may be jammed!) and a new election is triggered automatically.

SELECT uses four *leader slots*:[3] ls_1, ls_2, ls_3 and ls_4. Of course, in the beginning, all nodes may have different ls values and may disagree on which slots during the b-interval are leader slots. However, over time, the nodes synchronize their states and a consistent view emerges. For the synchronization, five temporary variables ls'_0, ls'_1, ls'_2, ls'_3, and ls'_4 are used, which store future ls values.

Depending on whether the node is of type follower or leader, the leader slots are updated differently: At the beginning of a new b-interval, a leader copies its ls'_i values to the ls_i values. A follower on the other hand copies the ls' values "diagonally" in the sense that ls'_i is copied to ls'_{i+1} for $i \in \{0,1,2,3\}$. As we will see, this mechanism ensures that an elected leader covers the leader slot ls_3 *of each follower*. (SELECT guarantees that the reactive adversary has no knowledge about the ls_3 slots at all until it is already too late to prevent a successful election.) An-

[3] It is an open question whether a protocol with less leader slots can be devised.

other special slot besides ls_3 is ls'_0 which is a random seed to mix the execution for increased robustness.

4.2.2 Description of SELECT

In Figure 4.1 we give the detailed formal description of the follower and the leader protocol, respectively. Recall that our algorithms can tolerate any initial values of mc, p_v, T_v, c_v, s_v, s'_v, ls_1, ls_2, ls_3, ls_4, ls'_0, ls'_1, ls'_2, ls'_3, ls'_4. For instance, in the beginning, all nodes v may be leaders and for all v, $s_v = 1$. However, the fixed parameters used by the algorithms, namely $\hat{p}, \gamma$, or b, are assumed to be immutable.

Both the follower and the leader algorithm consist of three main parts. The b-interval wise update (Lines $2-4$) makes sure that ls values are refreshed frequently. Lines $6-33$ (in case of a follower) and Lines $5-24$ (in case of a leader) are used for medium access in order to synchronize the nodes' states (by a message that includes c_v, T_v, and p_v values) and give nodes the chance to become or remain leader (by a 'LEADER' message). The last sections of the algorithms are used to react to high interference (by reducing p_v) and to reset leader slots. The reason for checking whether ls_3 is undefined in Line 6 of the follower protocol is to keep the leader slots hidden from the reactive adversary until it is already too late to prevent a successful leader election.[4]

Both the follower and the leader protocol depend on the following crucial CONDITION.

Definition 4.2 (CONDITION) *We define* CONDITION *(Line* 37 *for followers, and Line* 28 *for leaders) as the event that at least one 'LEADER' message was received during the past $b \cdot T_v$ steps.*

[4] This check would not be necessary against a non-reactive adversary.

The idea is that if CONDITION is fulfilled, we know that the protocol is already in a good state. Moreover, we will see that the adversary cannot prevent CONDITION to become true for a long time as the T_v values would continue to increase.

Finally, also note that leaders increase p_v faster (i.e., by larger multiplicative factors) during idle rounds than followers. With this mechanism, SELECT improves the likelihood that a 'LEADER' message gets through and hence that a unique leader is elected.

4.3 Analysis

This section shows that the randomized SELECT protocol is guaranteed to eventually reach a situation where there is exactly one leader and $n-1$ followers. We make use of the following definitions. First, we define the system state.

Definition 4.3 (State and System State) *The* state of node v *is determined by the state of the variables* p_v, T_v, c_v, s_v, s'_v, mc, ls'_0, ls_1, ls'_1, ls_2, ls'_2, ls_3, ls'_3, ls_4 *and* ls'_4. *The* state of the system *is the set of the states of all nodes.*

We use the following LS_L set to describe the union of all possible leader slot values present in the system.

Definition 4.4 (The LS_L State Set) *For any given system state, let* $LS_L = \{ls_1(v), ls_2(v), ls_3(v), ls_4(v) \; |v \text{ is leader}\} \setminus \{undefined\}$.

The system can be in several special states which are formalized next: follower states, pre-leader states, and leader states. Let $[b] = \{0, \ldots, b-1\}$.

Definition 4.5 (Follower State) *A state S is called a* follower state*, denoted by $S \in FOLLOWER$, if all the following conditions hold.* ***(i)*** *All nodes are followers* ($\forall v \in V : \; s_v = 0$)*;* ***(ii)*** *for every node v:* $ls_1(v), ls_2(v), ls_3(v)$, $ls_4(v) \in [b] \cup \{undefined\}$,

*$ls'_1(v)$, $ls'_2(v)$, $ls'_3(v)$, $ls'_4(v) \in [b] \cup \{undefined\}$, $ls'_0(v) \in [b]$; **(iii)** the follower nodes can be partitioned into two sets $\{v\}$ and $V \setminus \{v\}$, according to their ls' values (v is the node that successfully sent the last follower message); for each $w \in V \setminus \{v\}$: $ls'_1(w) = ls'_0(v)$, $ls'_2(w) = ls'_1(v)$, $ls'_3(w) = ls'_2(v)$, $ls'_4(w) = ls'_3(v)$, and $ls_2(w) = ls_1(v)$, $ls_3(w) = ls_2(v)$, and $ls_4(w) = ls_3(v)$; **(iv)** for any pair of follower nodes $v, w \in V$ with $ls'_2(v) \in [b]$ and $ls_3(v) \in [b]$, $c_v = c_w$ and $T_v = T_w$.*

We use the concept of so-called *pre-leader states*, i.e., states that result from follower states before some nodes become leaders.

Definition 4.6 (Pre-leader State) *A state S is called a* pre- leader state*, denoted by $S \in PRE - LEADER$, if it is a follower state, and at least one follower node v has $s'_v = 1$.*

While in the beginning, the leader sets may be large as each node regards different slots during the *b*-interval as the "leader slots", over time the values synchronize and the *LS* sets become smaller. This facilitates a fast leader (re-) election.

Definition 4.7 (Leader State) *A state S is called a* leader state*, denoted by $S \in LEADER$, if all the following conditions are satisfied:*

***(i)** There is at least one leader, i.e., $|\{v | v \in V : s_v = 1\}| \geq 1$; **(ii)** for every node v, $ls_1(v), ls_2(v), ls_3(v), ls_4(v) \in [b] \cup \{undefined\}$, $ls'_1(v), ls'_2(v), ls'_3(v), ls'_4(v) \in [b] \cup \{undefined\}$, $ls'_0(v) \in [b]$; **(iii)** let v be any follower and let w be any follower or leader, then $ls_3(v) \in \{ls_1(w), ls_2(w), ls_3(w), ls_4(w)\} \cup \{undefined\}$, $ls'_2(v) \in \{ls'_0(w), ls'_1(w), ls'_2(w), ls'_3(w)\} \cup \{undefined\}$; **(iv)** $|LS_L| \leq 5$; **(v)** for every follower w with $ls_3(w) \in [b]$ or $ls'_2(w) \in [b]$, $c_w = c_v$ and $T_w = T_v$ for any leader v.*

So in a leader state, it holds that any follower's ls_3 and ls'_2 slots are covered by either another follower's ls and ls' slots, or a leader's ls and ls slots (cf Condition (iii)).

Finally, it is useful to define safe and legal states.

Definition 4.8 (Safe and Legal State) *A system state S is called* safe *(denoted by $S \in SAFE$) if $S \in FOLLOWER$ or $S \in LEADER$, and* legal *(denoted by $S \in LEGAL$) if S is safe and there is exactly one node v with $s_v = 1$.*

Thus, according to our definitions, any legal state is also a safe state. In the following, let S be the set of all possible system states, $SAFE \subset S$ be the set of all *safe* system states and $LEGAL \subset SAFE$ be the set of all *legal* system states.

The proof of Theorem 4.1 unfolds in a number of lemmas. An interesting property of our randomized algorithm is that it is *guaranteed* to be correct, in the sense that deterministically exactly one leader is elected; only the runtime is probabilistic (i.e., depends on the random choices made by SELECT).

First, we study leader messages.

Lemma 4.9 *For any network state it holds that if a leader successfully transmits a 'LEADER' message, the system will immediately enter a legal state.*

Proof. When a node (either follower or leader) receives a 'LEADER' message, it sets ls_3 and ls'_2 to $undefined$ (Lines $22-25$ in Figure 4.1 *left*; after Lines $15-17$ of Figure 4.1 *right*), and considers itself a follower. Thus, in the new state, there is exactly one leader (the sender of the 'LEADER' message) and $n-1$ followers. The state is also a safe state, namely a leader state: Conditions (i) and (ii) are fulfilled trivially. Condition (iv) also holds as there is only one leader that has four slots. Condition (iii) is fulfilled because nodes receiving a 'LEADER' message reset their slots ls_3 and ls'_2; since ls_3 and ls'_2 are undefined for a follower, also Condition (v) holds.

We next consider what happens if nodes hear a message sent by a follower.

Lemma 4.10 *For any network state it holds that when a follower successfully transmits a message, the system is guaranteed to enter a safe state at the beginning of the next b-interval.*

Proof. First note that if a leader message gets through before the next b-interval, the claim holds trivially due to Lemma 4.9.

Otherwise we distinguish two cases: (A) For every node v, $s'_v = 0$ (not pre-leader) and $s_v = 0$ (not leader) by the end of current b-interval. (B) There is at least one node v with either $s'_v = 1$ (pre-leader) or $s_v = 1$ (leader) by the end of current b-interval.

In Case (A), after the follower message has been successfully sent, there are still n followers and no leaders or pre-leaders. We will show that the system enters the follower state at the beginning of the next b-interval. Let us refer to the follower node that sent the message by v and to any remaining node by w. When w receives the message from v (Lines 26 – 32 in Figure 4.1 *left*), it sets $ls'_1(w) := ls'_0(v)$, $ls'_2(w) := ls'_1(v)$, $ls'_3(w) := ls'_2(v)$, and $ls'_4(w) := ls'_3(v)$. The c values become the same ($c_w = c_v$), and $T_w := T_v$. The new state therefore fulfills the follower state conditions: Clearly, Conditions $(i), (ii)$, and (iv) are fulfilled immediately, and Condition (iii) holds as well, as for all followers w that did not send a message and follower v which sent a message, at the beginning of the next b-interval: $ls_3(w) = ls'_2(w) = ls'_1(v) = ls_2(v)$, $ls_3(v) = ls'_2(v) = ls'_3(w) = ls_4(w)$, and $ls_1(v) = ls_2(w) = ls'_0(v) = ls'_1(w)$.

For Case (B), observe that during the remainder of the b-interval the number of pre-leader nodes with $s'_v = 1$ cannot decrease, and hence there will be at least one leader at the beginning of the next b-interval. We now show that the new state will indeed be a leader state as nodes "synchronize" with the follower node that sent the message. Without loss of generality, assume that node u is the last follower that successfully sent a follower message in the current b-interval. Let us refer to the other follower nodes by v_1 and to the leader nodes or the pre-leader nodes (i.e., the followers v with $s'_v = 1$) by v_2. Again, Conditions (i) and (ii) are fulfilled trivially. As for Condition (iii), we need to consider two sub-cases:

(Case 1) No node experienced an idle channel in its ls_3 slot after the message has been successfully sent. If this is the case and follower u is not a pre-leader, it holds that for follower v_1: $ls'_2(v_1) = ls'_2(v_2) = ls'_1(u)$ in the current b-interval, and $ls_3(v_1) = ls_2(v_2) = ls_2(u)$ at the beginning of the next b-interval; on the other hand, if follower u is a pre-leader, then in the current b-interval it holds that for follower v_1: $ls'_2(v_1) = ls'_2(v_2) = ls'_1(u)$, and $ls_3(v_1) = ls_2(v_2) = ls_1(u)$ at the beginning of the next b-interval. Hence, Condition (iii) holds. Regarding the cardinality of the leader set LS_L, observe that at the beginning of the next b-interval, if u is not a pre-leader, all leaders will have $ls_1 = ls'_0(u), ls_2 = ls'_1(u), ls_3 = ls'_2(u), ls_4 = ls'_3(u)$, and hence $LS_L = \{ls'_0(u), ls'_1(u), ls'_2(u), ls'_3(u)\}$, therefore $|LS_L| \leq 5$; otherwise, if u is a pre-leader, then $LS_L = \{ls'_0(u), ls'_1(u), ls'_2(u), ls'_3(u), ls'_4(u)\}$, therefore $|LS_L| \leq 5$.

(Case 2) One or more nodes experienced an idle channel in their ls_3 slots after the message has been successfully sent. In the following, we prove this case correct assuming that u is a follower and not a pre-leader. If u is a pre-leader, the proof is analogous.

1. If v_1 experienced the idle channel at its ls_3 time slot, and became a pre-leader: Note that a node v_1 may experience an idle channel after receiving the message from u and hence become a pre-leader, however Condition (iii) is still satisfied, as it holds that for follower u: $ls'_2(u) = ls'_3(v_2) = ls'_3(v_1)$ in the current b-interval and $ls_3(u) = ls_3(v_2) = ls_3(v_1)$ at the beginning of the next b-interval. As for the cardinality of the leader set LS_L, observe that at the beginning of the next b-interval, all leaders will have $ls_1 = ls'_0(u), ls_2 = ls'_1(u), ls_3 = ls'_2(u), ls_4 = ls'_3(u)$, and hence $LS_L = \{ls'_0(u), ls'_1(u), ls'_2(u), ls'_3(u)\}$, therefore $|LS_L| \leq 5$.

2. If u experienced the idle channel at its ls_3 time slot, and became a pre-leader: If node u experienced an idle channel after successfully sending the message, u became a pre-leader, and we have for a follower v_1, $ls'_2(v_1) = ls'_2(v_2) = ls'_1(u)$ in the current b-interval and $ls_3(v_1) = ls_2(v_2) = ls_1(u)$ at the beginning of the next b-interval. Hence, Condition (iii) is satisfied. As for $|LS_L|$, observe that at the beginning of the next b-interval, for a leader v_2, $ls_1 = ls'_0(u), ls_2 = ls'_1(u), ls_3 = ls'_2(u), ls_4 = ls'_3(u)$, while for the remaining leader u, it holds that $ls_1 = ls'_1(u), ls_2 = ls'_2(u), ls_3 = ls'_3(u), ls_4 = ls'_4(u)$. Hence, also in this case, we have that $|LS_L| \leq 5$.

Finally, Condition (v) is true for both of the sub-cases, because the c_v and T_v values are "synchronized" when the follower message is received (Lines 27 and 29 in Figure 4.1 *left*; Line 19 in Figure 4.1 *right*).

An important property of SELECT is that once it is in a safe state, it will remain so in future (given that there are no external changes). Similar properties can be derived for other states, as we will see.

Lemma 4.11 *Once the system is in a safe state, it will remain in a safe state in the future.*

Proof. We study what can happen in one round, and show that in each case, the safety properties are maintained. In a round, (A) either a 'LEADER' message is successfully sent, (B) a follower message is successfully sent, (C) there are collisions or the channel is jammed, or (D) there is an idle channel.

In Case (A), the claim directly follows from Lemma 4.9 and from the fact that safe states are a super set of the legal states ($SAFE \supset LEGAL$). In Case (B), the claim follows from Lemma 4.10 and by the fact that the system is in the safe state already.

In Case (C), if the channel is blocked, follower nodes (even those which sent a message in this round) do not change their state except for the synchronized rounds in Lines $35-43$, and similarly for the leaders in Lines $26-34$. Our protocols guarantee that the leaders have the same c_v and T_v values as the followers when ls_3 and ls'_2 are valid, and since the leaders experience the same number of successful transmissions and idle time steps as the followers do (single-hop network), the claim follows.

If there is an idle channel (Case (D)), all nodes v for which $ls_3(v) = mc$ will set $s'_v = 1$ in the current b-interval, while other values remain the same. It is clear that from this point on until the end of the current b-interval, the claim holds. Moreover, as we show next, the claim is still true at the beginning of next b-interval. If $ls_3(v)$ is undefined, then the claim holds trivially, as no states will change in this case. If $ls_3(v) = mc$ for any node v and the nodes experience an idle channel, there is no leader since, if there was a leader, according to Condition (iii) of the leader state definition (Definition 4.7), a follower's ls_3 slot would always be covered by a leader slot of a leader, which yields the contradiction. Hence, the current safe state must be a pre-leader state. Let v denote the followers that have $s'_v = 0$ (i.e., they are not pre-leaders); let u denote the followers with $s'_u = 1$ (pre-leaders). In the current b-interval, we have $ls'_2(v) \in \{ls'_0(u), ls'_1(u), ls'_2(u), ls'_3(u)\} \cup \{undefined\}$, which is true according to Condition (iii) of the follower state definition (Definition 4.5). Then, at the beginning of next b-interval, u will become a leader, and hence we have $ls_3(v) = ls'_2(v)$, $ls_1(u) = ls'_1(u)$, $ls_2(u) = ls'_2(u)$, and $ls_3(u) = ls'_3(u)$. This implies that $ls_3(v) \in \{ls'_0(u), ls_1(u), ls_2(u), ls_3(u)\}$, which satisfies Condition (iii) of the leader state Definition 4.7. Conditions (i) and (ii) are clearly satisfied. Condition (iv) holds simply because we have shown (in Lemma 4.10, Case (B)), when there is an idle time step, $|LS_L| \leq 5$. Condition (v) is true because we always synchronize the c_v and T_v values.

Lemma 4.12 *Once a system is in a leader state, it will remain in a leader state in the future.*

Proof. Lemma 4.11 tells us that the system will never leave a safe state. Therefore, it remains to prove that there will always be at least one node v with $s_v = 1$. This clearly holds as the only way a leader can become a follower again is by receiving a 'LEADER' message (see Lines $15 - 17$), which of course implies that another leader is still active and remains to be a leader. Also, since we are in a leader state, Condition (v) holds and it further implies that leaders will never invalidate their ls slots before the followers. This guarantees that the protocol will never get out of a leader state.

Lemma 4.13 *Once a system is in a legal state, it will remain in a legal state in the future.*

Proof. By Lemma 4.11, we know that our system will never leave a safe state again, and hence, we only need to prove that there will always be exactly one node v with $s_v = 1$. This is true because in the safe state, a follower node w can never become a leader, as its $ls_3(w)$ slot is covered by the leader v: $ls_3(w) \in \{ls_1(v), ls_2(v), ls_3(v), ls_4(v)\}$ and $ls'_2(w) \in \{ls'_0(v), ls'_1(v), ls'_2(v), ls'_3(v)\}$ (Condition (iii) of leader state). Since a follower will never send a 'LEADER' message, v will remain a leader forever, which proves the claim.

Regarding convergence, note that the system quickly enters a safe state, deterministically.

Lemma 4.14 *For any initial system state with $\hat{T} = \max_v T_v$, it takes at most $b \cdot \hat{T}$ rounds until the system is in a safe state.*

Proof. We distinguish three cases: if a leader message gets through sometimes in these rounds, then the claim holds by Lemma 4.9; if a follower message gets through, then the claim holds by Lemma 4.10. If within $\max_v T_v b$ rounds neither a follower message nor a leader message gets through, all nodes will have to reset their ls slots (since CONDITION in Line 37 (Figure 4.1 *left*) resp. Line 28 (Figure 4.1 *right*) is not met). This however constitutes the safe state (all conditions fulfilled trivially), which is maintained according to Lemma 4.11.

Armed with these results, we can prove convergence.

Lemma 4.15 *For any safe state,* SELECT *will eventually reach a legal state.*

Proof. We divide the proof in two phases: the phase where the protocol transitions to the leader state from the follower state, and the phase where it transitions to the legal state from the leader state.

1. Follower state to leader state

   If CONDITION is fulfilled, we know that a 'LEADER' message got through and the system is in a legal state (and hence also in a leader state). As long as CONDITION is not fulfilled, T_v is increasing for each node v. So eventually, $\hat{T} = \max_v T_v \geq 2T/b$. We can also provide a lower bound on the cumulative probability p. W.l.o.g. suppose that $T \geq (3/\varepsilon)\log_{1+\gamma} n$ (a smaller T will only make the jammer less flexible and weaker). Suppose that p is at most $\varepsilon/4$ throughout some T-interval I. Then it follows from the standard Chernoff bounds that there are at most $\varepsilon T/3$ busy steps in I with high probability.[5] If this is true, then no matter how the adversary jams during I, at least $(1-\varepsilon/3)T - (1-\varepsilon)T = 2\varepsilon T/3$ non-jammed steps will be idle, which implies that the cumulative probability at the end of I will be by a factor of at least $(1+\gamma)^{\varepsilon T/3} \geq n^3$ higher than at the beginning of I. Using this insight, it follows that eventually a T-interval is reached with $p > \varepsilon/4$. Once such a T-interval has been reached, it is easy to show that p will not get below $1/n^2$ any more w.h.p. so that for every T-interval afterwards there is a time point t with $p > \varepsilon/4$ w.h.p. So infinitely often the following event can take place with some lower-bounded, positive probability:

   Consider two consecutive T-intervals I_1 and I_2 starting at a time when $c_v = 0$ for every node v. Suppose that I_1 just consists of busy steps and I_2 just consists of idle time steps. Then the adversary has to leave εT busy time steps in I_1 non-jammed and εT idle time steps in I_2 non-jammed. For I_1, there is a positive probability in this case that exactly 3 messages from different nodes are successfully sent in 3 different b-intervals. In this case, all but one follower respect the leader slots (as their ls_3-value is defined) while the follower

[5] "With high probability", or short "w.h.p.", means a probability of at least $1 - 1/n^c$ for any constant $c > 0$.

that sent the last successful message may still send out messages at *all* time steps (as its ls_3-value is still undefined, see Line 6 of the follower protocol). Thus, it is indeed possible that all time steps in I_1 are busy. Up to that point, the adversary has not learned anything about the leader slots. In I_2, there is also a positive probability that none of the followers transmits a message throughout I_2 so that all time steps are idle. As the adversary does not know which of them is a leader slot and has to leave εT non-jammed, there is a positive probability that ls_3 is non-jammed, and some of the followers become pre-leaders and then leaders.

Thus, the expected time to get from a follower to a leader state is finite.

2. Leader state to legal state

   If there is only one leader in the leader state, the system is already in a legal state by definition. If there is more than one leader, then we distinguish between the following cases. If CONDITION is fulfilled, we know that a 'LEADER' message got through and the system is in a legal state. Otherwise, the leaders will invalidate all of their ls slots once their c_v values are reset to 0. At this point there is a positive probability that for the next T steps a 'LEADER' message is successfully sent. As the adversary has to leave εT time steps non-jammed, at least one 'LEADER' message will be successfully transmitted within these T steps so that the system reaches a legal state.

   Analogous to the followers in the previous case, one can lower bound the cumulative probability of the leaders (in fact, the leaders will eventually reach a time point with a cumulative probability of $\Omega(\varepsilon)$ as they increase their probabilities in case of an idle channel more aggressively than the followers) so that the chance above of successfully transmitting a 'LEADER' message repeats itself infinitely often with a lower-bounded positive probability. Thus, the expected time to get from a leader to a legal state is finite as well.

From these cases, the lemma follows.

4.4 Experiments

We conducted several simulations to study the behavior of SELECT under different types of jammers and interference.

4.4.1 Performance under Jamming

For our formal analysis, we introduced the notion of a $(T, 1-\varepsilon)$-bounded adversary for some $T \in \mathbb{N}$ and $0 < \varepsilon < 1$ which denotes that for any time window of size $w \geq T$ the adversary can jam at most $(1-\varepsilon)w$ of the time steps in that window. While our protocol is provably robust to *any* adversary meeting these constraints, for our simulations, we will need to focus on specific instantiations. For example, we will consider an adversary that reactively jams all non-idle time periods only (as long as the budget is not used up), in order not to waste energy jamming idling periods.

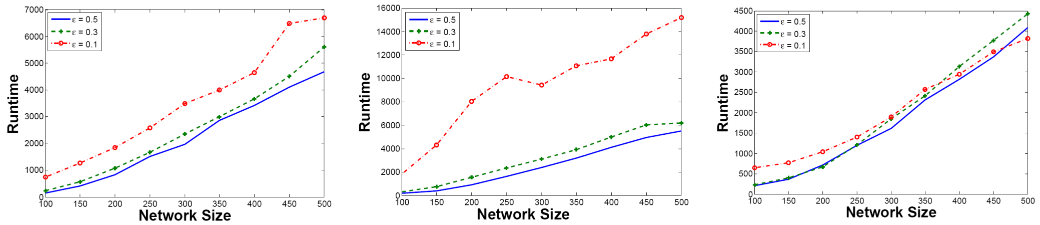

Figure 4.2: *Left:* Convergence time from safe state to legal state, where the adversary ADV_{rand} jams the channel. We ran our protocol until exactly one leader is elected. *Middle:* Convergence time from safe state to legal state, where a reactive adversary ADV_{busy} jams the channel when one or more nodes are transmitting. We ran SELECT until only one leader is elected. *Right:* Convergence time from safe state to legal state under the reactive ADV_{idle} adversary.

We consider jammers of different powers, one that can block the channel 90% of the entire time, one that blocks 70% of the time, and a "weak" one that blocks 50% of the time (i.e., $\varepsilon \in \{0.1, 0.3, 0.5\}$, resp.). We set $T = 100$ and consider a b-interval (see Figure 4.1) with parameter $b = 15$ (smaller b values are possible as well). Experiments are repeated 50 times for each individual setting, and average values are recorded correspondingly. We run each experiment until one and only one leader is elected.

We conducted experiments with different types of reactive jammers: jammers ADV_{rand} which interrupt transmissions *at random*, jammers ADV_{busy} which only jam busy periods where one or more nodes transmit, and jammers ADV_{idle} which jam the channel whenever it is idle. Concretely, for ADV_{busy} and ADV_{idle} we assume that the adversary will jam each busy resp. idle time period until the "jamming budget" is used up for this T-period. For ADV_{rand} we set the jamming probability per round equal to $(1 - \varepsilon)$. ADV_{idle} may appear less challenging to the deal with. However, note that an adversary may be able to lead a protocol to suboptimal states by jamming idle time periods. Moreover, this scenario also describes interference from co-existing networks where nodes are activated in quiet times. Hence, this adversary constitutes an interesting case that should not be neglected in the analysis.

Recall from Lemma 4.14 that from any initial state, the safe state is reached quickly, and hence, we are mainly interested in the convergence time from the safe state to the legal state. Figure 4.2 (*left*) plots the corresponding convergence times. At first sight the runtime may appear to be rather high. For example, under an adversary ADV_{rand} that jams 90% of the entire time, it takes a few thousand time steps. However, note that this result implies that during the merely a few hundred non-jammed time steps, the five hundred nodes are able to successfully coordinate the medium access among themselves—without being able to distinguish between time periods with collisions and time periods that are jammed!— and use the computed access probabilities to elect a leader. We believe that when taking this into account, and although we do not have any lower bounds, the convergence time is very good and probably cannot be improved much with alternative schemes.

Figure 4.2 (*middle* and right) presents the corresponding convergence times for the reactive jammers ADV_{busy} and ADV_{idle}. As expected, jamming the busy channel yields higher convergence times, also when comparing these results to our experiments with ADV_{rand}. In contrast, interestingly, for ADV_{idle}, the runtime is fairly independent of the adversarial power: a reactive jammer blocking idle channels gives similar results as ADV_{rand}. Clearly, among the scenarios we investigated, the most effective strategy for the adversary is to reactively jam the busy time periods as long as the total number of jammed time steps does not exceed $(1-\varepsilon)\cdot T$. Figure 4.3 complements Figure 4.2 by studying the execution times in smaller networks.

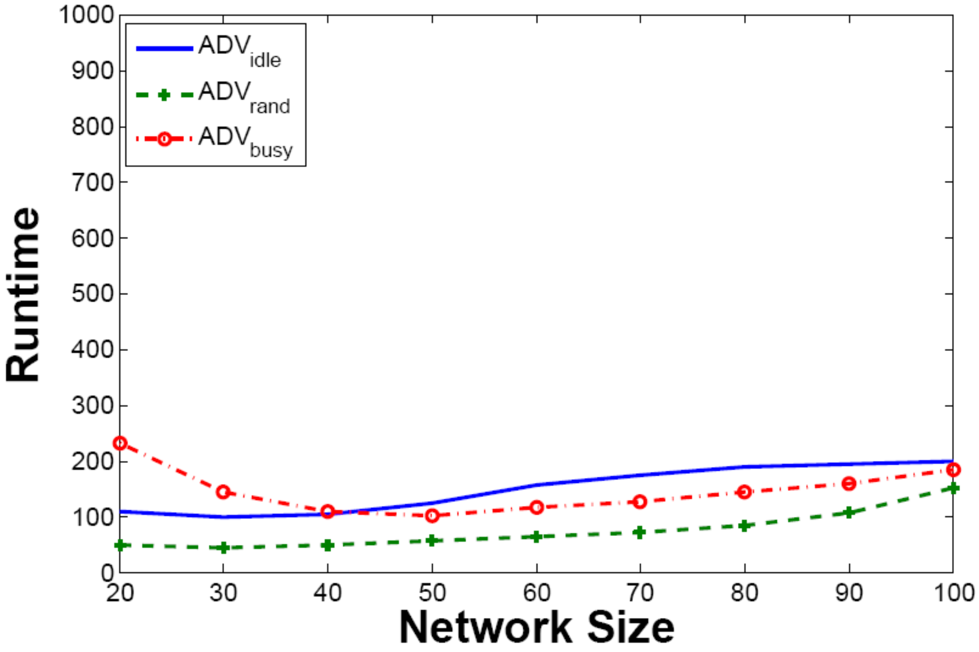

Figure 4.3: Like Figure 4.2, but for smaller networks and using $\varepsilon = 0.5$ and $\hat{p} = 1/24$.

Our protocol aims to quickly reach a cumulative sending probability around a small constant, such that on expectation, roughly one node will try to transmit a message in a non-jammed step. Thus, given the constant probability of having a successful transmission, a follower messages will get through soon, the nodes synchronize, and the *ls* slots are defined as well. Since the leaders' sending probabilities reach higher values more quickly than the sending probabilities of the followers (according to Line 12 of Figure 4.1 *right*), a leader message gets through soon, yielding a legal state. We consider two initial states, a "well-initialized one" where all nodes have the maximum access probability $\hat{p}$ (in simulation we set $\hat{p} = 1/24$), where there is no leader in the network, and where the *ls* and *ls*$'$ slots are all invalidated (according to Definition 4.5, this implies that we are in the follower state); and one with "arbitrary initialization" where the roles and variables are chosen at random (each node is either follower or leader, p_v is chosen uniformly at random between 0 and 1, and the ls values uniformly at random between 0 and $b-1$). Our experiments show that both scenarios yield similar results, which indicates that the convergence time of self-stabilization if fairly independent of the initial state.

Figure 4.4 (*left*) shows a typical trace of the cumulative probabilities over time when the protocol is well initialized, i.e., the protocol starts from the follower state. Initially, all nodes are followers, and we will denote the cumulative sending probability of the followers by p_F, and the cumulative sending probability of the leaders by p_L. At beginning, $p_F > 500 \cdot \hat{p} > 10$ while $p_L = 0$. As time goes on, p_F decreases quickly until it falls in an interval of small constant range (i.e., $p_F < 10$), and multiple successful transmissions happen which synchronize the nodes' ls and ls' values. Next, multiple leaders are elected because many followers sense an idle time step in their ls_3 slot. That is why p_L emerges at the same point in time as p_F decreases dramatically to a value between 0 and 1. Then, the nodes continue to adjust their transmission probabilities depending on the channel state, until the first leader message gets through and all the other leaders become followers; this yields the quick decrease of p_L and increase of p_F accordingly. One and only one leader is elected after this point. Subsequently, both p_F and p_L remain within a small constant range. Figure 4.4 (*right*) shows the cumulative probabilities when the protocol starts from an "arbitrary" state (p_v, T_v, leaders and follower roles, etc. chosen at random). In the beginning, there are both followers and leaders in the network. It can be seen that SELECT converges fast, similarly to the well-initialized case. After the legal state is reached, both p_F and p_L also remain in a small constant range.

4.4.2 Co-existing Networks

Our leader election protocol is robust to arbitrary (but bounded with respect to time) interruptions of the availability of the medium, and it is convenient to regard these

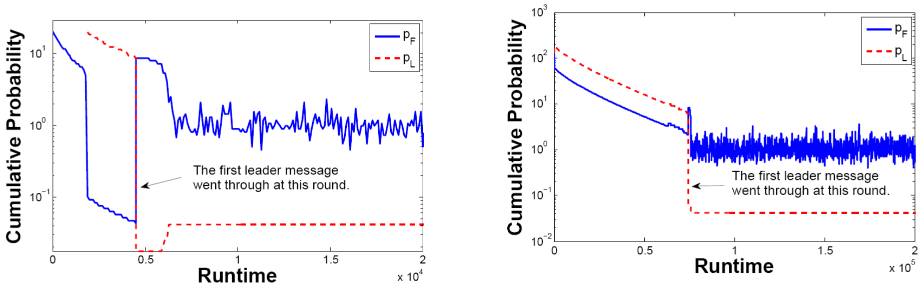

Figure 4.4: *Left:* Fast convergence of p_F and p_L ($\varepsilon = 0.5$, $T = 100$, network with 500 nodes) under ADV_{busy} when the protocol starts from a safe state. *Right:* Corresponding convergence of p_F and p_L when the protocol starts from an arbitrary state.

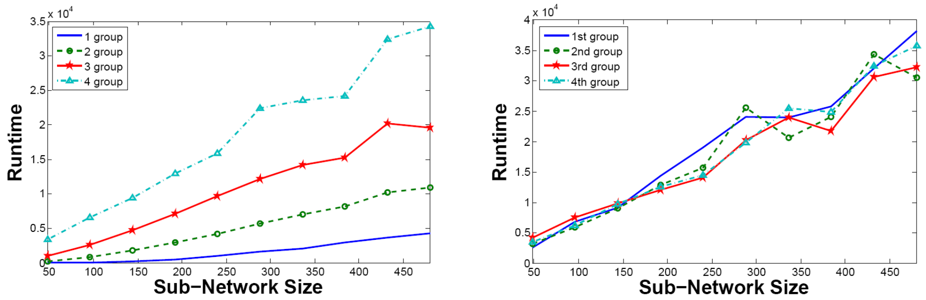

Figure 4.5: *Left:* Convergence time of co-existing networks (as a function of their individual sizes) performing the leader election algorithm. *Right:* Fair convergence time among co-existing networks.

interruptions as caused by a malicious adversary. However, there are many other forms of interference to which our protocol is resilient and under which the few available time slots can be exploited effectively. In the following, we briefly report on one more source of interference, namely co-existing protocol instances. Concretely, we remove the jammer from the network and we compare the performance of our leader election protocol when run alone to situations where additional networks (of the same size) are concurrently trying to elect a leader and interfere with the other protocol instances accordingly.

Figure 4.5 (*left*) plots the averaged runtime until successful leader election for one, two, three and four co-existing networks, as a function of the corresponding sub-network sizes (i.e., four co-existing networks imply a four times larger total number of nodes). Our results indicate that each additional interfering network increases the runtime by a factor corresponding to the additional nodes. The convergence time among the co-existing networks exhibits a high fairness, as can be seen in Figure 4.5 (*right*): in all networks, a leader is elected almost at the same time.

4.5 Conclusion

We introduce the first self-stabilizing leader election protocol, SELECT, for wireless networks operating in harsh environments, e.g., environments with hard-to-predict interference from co-existing networks or environments subject to (both adaptive and reactive) adversarial jamming. Although the nodes are not able to distinguish between collisions due to external interference or jamming and concurrent transmissions of other nodes in the network, they are able to coordinate access to the medium in the few and arbitrary time periods without external interference, and subsequently elect a leader in a robust manner. Although our protocol is randomized, it yields deterministic guarantees.

There are several important open directions for future research. For example, the formal study of convergence times under different adversaries is an open problem. Another open problem is the generalization of our algorithms to multi-hop networks where leaders need to be elected in different regions (e.g., in order to construct a sparse backbone).

Chapter 5

THE CoMac PROTOCOL

The decentralized allocation of a communication medium among a set of wireless nodes does not only constitute one of the most fundamental theoretical problems in distributed computing, but is also of direct practical relevance. Today, a chunk of the wireless spectrum is often simultaneously used by many devices belonging to different, so-called *co-existing networks*. It is expected that the popularity of wireless mobile devices will further increase the resource sharing by such networks in the future.

Interestingly, not much is known today on how a given spectrum can be shared efficiently and fairly among co-existing networks, especially in environments with uncontrollable external interference. Existing distributed MAC protocols (typically based on random backoff schemes) are either not resistent to the unpredictable unavailability of the medium at all, or are optimized towards a single network only, in the sense that the nodes of a network collaboratively seek to coordinate the access among themselves [50]. However, the state-of-the-art protocols fail if multiple networks are collocated (as illustrated, for example, in our simulation study in Section 5.4).

This chapter presents (and rigorously prove the performance of) a robust MAC protocol suited for co-existing networks exposed to a harsh environment with unpredictable or even adversarial interference.

We attend to a simplified scenario where a set of n wireless nodes V are located within transmission range of each other and need to communicate over a single shared channel. The wireless nodes belong to K co-existing networks N_i with node sets V_i, i.e., $V = V_1 \cup V_2 \cup \ldots \cup V_K$, for some constant K (which is of unknown to the nodes). For simplicity we will assume that these networks are node disjoint. However, by emulating multiple instances, a node may also participate in several networks simultaneously; the performance guarantees derived in this chapter would still hold.

We aim to design a distributed MAC protocol for these wireless nodes. Although the protocol is used by all nodes $v \in V$, it should not depend on any knowledge of how many nodes n there are in total, on the number of co-existing networks K, or on the size of the co-existing network v belongs to. Moreover, it should ensure that the K networks are independent in the sense that no communication is required between different networks.

Co-existing wireless networks appear in many scenarios where different wireless networks share the same wireless medium. For example, consider a major conference, e.g., organized by the United Nations, where participants from different countries use their hand-held devices to communicate with the other representatives of their country. We assume that the different networks only share the same medium access protocol, but are otherwise different and inter-network communication may not be desired or possible (except, e.g., for multi-national participants). Another scenario where ensuring fairness among co-existing networks is crucial are emergency response networks, where many emergency response services, such as fire squads, police, and paramedics, all arrive simultaneously at some accident or disaster scene and have to share the wireless medium in a fair and even manner in order to establish their own separate communication networks.[1]

We present a robust and fair medium access (MAC) protocol CoMAC that makes effective use of the few and arbitrarily distributed time periods where a wire-

[1] Whereas in some scenarios it may be desirable that messages are broadcast across all emergency unit networks, for better immediate response action to a disaster/accident, in the longer run, it is still important to be able to differentiate among the different ad-hoc networks established.

less medium is available. We model interference—due to simultaneous transmissions, co-existing networks, changes in the environment that affect the wireless medium, etc., and, when applicable, intentional jamming—generally as an *adversary*, which we may sometimes simply refer to as the *jammer* (even when a malicious jammer is not present in the environment and interference may be caused by other factors). Our adversary may behave in an *adaptive* manner: we assume that the adversary has full knowledge of the protocol and its history, and that it uses this knowledge to decide on whether to jam at a certain moment in time.

Let us use the simplifying notation $N(v)$ to denote the network node $v \in V$ belongs to. We assume that a node v can distinguish among the following events at some time t: (1) idle channel (no node in V transmits and there is no outside interference, including jamming activity, at time t); (2) successful transmission of a packet in network $N(v)$ (which occurs every time a single node in $N(v)$ transmits, and no other node in V nor the adversary transmits); and (3) medium busy (due to a transmission by a node in some co-existing network different from network $N(v)$, or to simultaneous transmissions by two or more nodes in $N(v)$, or to external interference or jamming).

How to design such a distributed medium access protocol which shares the bandwidth fairly among the K networks, without sacrificing performance? At first sight this may seem impossible: as the total number of co-existing networks and the number of devices is not known, a node cannot guess its fair share of the channel time. We show that this is indeed possible, even in the presence of a powerful adaptive adversarial jammer, referred to as a $(T, 1-\varepsilon)$-*bounded* (adaptive) adversary, which can jam the medium an arbitrary $(1-\varepsilon)$ fraction of the time for an arbitrarily small constant $\varepsilon > 0$ and which hence models a wide range of external interference scenarios or jammers. For the ease of presentation, we assume a synchronous environment where time proceeds in *rounds* (also called *steps*). Formally, the $(T, 1-\varepsilon)$-bounded adversary is defined as follows: for some $T \in \mathbb{N}$ and a constant $0 < \varepsilon < 1$, the adversary may jam at most $(1-\varepsilon)w$ of the time steps, for any time window of size $w \geq T$. In the following, we will use the notation $N = \max\{T, n\}$ to denote the maximum over the adversarial window size and n.

Assuming backlogged traffic at the wireless devices, we require that our MAC protocol fulfill the following properties: (1) *c-competitiveness*: Given a time interval I, we define $g(I)$ as the number of time steps in I that are non-jammed, and $s(I)$ as the total number of time steps in I in which a successful transmission happens in any network. A MAC protocol is called *c-competitive* against some $(T, 1-\varepsilon)$-bounded adversary if, for any sufficiently large time interval I, $s(I) \geq c \cdot g(I)$. (2) *Fairness*: The probabilities of having a successful transmission in any two networks N_i and N_j, where $i, j \in [1, K]$, do not differ by much; moreover, the nodes inside a network share the bandwidth fairly as well.

Note that the nodes have no knowledge of how many nodes are there in the same network as itself, nor do the nodes know how many other networks are co-existing and how many nodes are there in each of these co-existing networks, respectively. However, we assume that the nodes have a common parameter $\gamma \in O(1/(\log T + \log\log n))$. The assumption that nodes know γ is not critical for the scalability of our protocol, as it requires only a polynomial estimate of T and an even rougher estimate of n.

Although the presented COMAC protocol converges fast and is therefore expected to work well under continuously entering and leaving nodes, in this chapter we will just focus on a synchronous setting where nodes do not join or leave.

5.1 Contribution

To the best of our knowledge, we are the first to present a robust medium access protocol which provably performs well in an environment with co-existing networks. The COMAC protocol features a guaranteed competitive throughput in the presence of co-existing networks as well as a wide range of external interference patterns that can be subsumed and modeled as a $(T, 1-\varepsilon)$-bounded adaptive adversary blocking the medium a $(1-\varepsilon)$ fraction of all time. Moreover, it features fairness among co-existing networks and within an individual network. Finally, the protocol is attractive for its simple design. Our main theoretical result is summarized in the following theorem.

Theorem 5.1 *The* COMAC *medium access protocol guarantees that in a backlogged scenario, if executed for* $\Omega(\frac{1}{\varepsilon}\log N \max\{T, \frac{1}{\varepsilon\gamma^2}\log^3 N\})$ *many time steps,* COMAC *achieves a competitive throughput of* $\Omega(\varepsilon^2 \min\{\varepsilon, 1/poly(K)\})$ *w.h.p., despite the arbitrarily distributed non-jammed time periods left by the* $(T, 1-\varepsilon)$*-bounded adaptive adversary that arbitrarily jams the medium up to an* $(1-\varepsilon)$ *fraction of the time, and which has complete knowledge of the protocol history. Moreover, the cumulative probabilities among different networks, as well as the access probabilities of individual nodes within the same network, differ only by a small factor.*

To complement our theoretical asymptotic bounds, we also report on a comparative simulation study.

5.2 Description of COMAC

Before presenting the formal MAC algorithm, we explain its variables and provide some intuition.

5.2.1 Intuition

In the COMAC protocol, each node v maintains a medium access probability p_v which determines the probability that v transmits a message in a communication round. The nodes adapt and synchronize (inside a co-existing network) their p_v values over time (which as a side-effect also improves fairness) in a multiplicative-increase multiplicative-decrease manner in order to ensure a throughput that is as good as possible. More precisely, the sending probabilities are changed by a factor of $(1+\gamma)$. Moreover, we impose an upper bound of $\hat{p}$ on p_v, for some constant $0 < \hat{p} < 1$. As we will see, unlike in most classic backoff protocols, our adaption rules for p_v ensure that the adversary cannot influence p_v much by adaptive jamming.

In addition, each node maintains two variables, a threshold variable T_v and a counter variable c_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.

While these concepts have already been used in our other protocols in [6, 49, 50], they are not sufficient to ensure a jamming-resistant protocol that also works well in case of co-existing networks. The basic problem lies in the fact that all of these protocols aim at reaching a constant cumulative probability, irrespective of the adversarial jamming, so that a good throughput can be obtained in those steps that are not jammed. In co-existing networks, however, this is not a good idea: Suppose that we have K co-existing networks that each have a constant cumulative probability. Then the overall cumulative probability would be $\Theta(K)$ and therefore, the probability of having a successful transmission in any network would be as low as $\Theta(K)e^{-\Theta(K)}$, which is *exponentially* low in K.

Hence, a less aggressive approach than the one pursued in [6, 49, 50] is needed. Ideally, this approach should also make sure that the available bandwidth is shared in a fair way among the networks. Surprisingly, a relatively simple change in the protocol in [50] can achieve jamming-resistance, a good throughput in co-existing networks, and also fairness. The basic idea behind this change is to remember the latest idle time step, and whenever there is a new idle time step, then *with a probability q_v that is inversely proportional to the time difference to the previous idle time step*, p_v and T_v are adapted. (The protocol in [50] would *always* adapt p_v and T_v in case of an idle channel.) Since this probabilistic rule turned out to be very hard to analyze, we transformed it into a deterministic rule that shows the same performance in the experiments.

5.2.2 Protocol Description

Now we are ready to provide the detailed and formal description of the COMAC see Algorithm 3. Initially, each node v sets $p_v = \hat{p}$ ($\hat{p} \leq 1/24$), $c_v = T_v = 1$, and $q_v = 0$. In the following, $L_v \geq 1$ is the time that went by from v's viewpoint since the last idle time step. (If there has not yet been an idle time step, $L_v = \infty$.)

5.3 Analysis

For the analysis of our protocol we will use the following notation. We are given $K \geq 2$ co-existing networks denoted by $N_1, \ldots, N_K$. Each network N_i consists of a node set V_i where $n_i = |V_i| \geq 2$ (otherwise, the network would be irrelevant). The cumulative probability due to nodes in N_i is given by $P_i = \sum_{v \in V_i} p_v$, and the cumulative probability over all co-existing networks is given by $P = \sum_{i=1}^{K} P_i$. Whenever we consider some specific time step t, $P_i(t)$ is the value of P_i at time t and $P(t)$ is the value of P at time t.

Algorithm 3 CoMac: for each node v

```
$roundcounter = 0$
$p_v := \hat{p}$
$c_v := 1$
$T_v := 1$
$L_v := \infty$ {CoMac works in synchronized rounds}
while True do
    $v$ decides with probability $p_v$ to send a message
    if $v$ decides to send a message then
        $v$ sends a message along with a triple: $(p_v, c_v, T_v)$.
    else
        $v$ senses the channel
        if $v$ senses an idle channel then
            $q_v := q_v + 1/L_v$
            if $q_v \geq 1$ then
                $p_v := \min\{(1+\gamma)p_v, \hat{p}\}$
                $T_v := T_v - 1$
                $q_v := q_v - 1$
                update $L_v$
            end if
        else if $v$ successfully receives a message along with the triple of
            $(p_{new}, c_{new}, T_{new})$ then
            $p_v := (1+\gamma)^{-1} p_{new}$
            $c_v := c_{new}$
            $T_v := T_{new}$
        end if
    end if
    $c_v := c_v + 1$
    if $c_v > T_v$ then
        $c_v := 1$
        if there was no idle step among the past $T_v$ time steps then
            $p_v := (1+\gamma)^{-1} p_v$
            $T_v := T_v + 2$
        end if
    end if
    $roundcounter := roundcounter + 1$
end while
```

5.3.1 Basic Observations

Given that we have a single-hop network, any idle time period is observed by all nodes in all co-existing networks. Hence, the q_v and L_v values of all nodes are identical if all start at the same time (otherwise, two idle time steps suffice to synchronize the L_v values so that the increase of the q_v's is synchronized from that point on, which would also be sufficient for our analysis to go through). Henceforth, we will drop the subscript v from q_v and L_v. Since after the first successful transmission in N_i, the T_v and c_v values are synchronized among the nodes in N_i, we arrive at the following fact, which establishes fairness within a network.

Fact 5.2 *After the first successful transmission in network N_i, the access probabilities p_v of the nodes $v \in V_i$ differ by a factor of at most $(1+\gamma)$.*

Throughout our analysis, we will make use of Lemma 1.3 and Chernoff bounds from Lemma 1.2 repeatedly.

Based on Lemma 1.3, we prove the following lemma.

Lemma 5.3 *For any non-jammed time step,*

$$e^{-\frac{P}{1-\hat{p}}} \leq \mathbb{P}[\textit{channel is idle}] \leq e^{-P} \quad \textit{and}$$

$$P_i \cdot e^{-\frac{P}{1-\hat{p}}} \leq \mathbb{P}[\textit{successful msg transmission in } N_i] \leq \frac{P_i}{1-\hat{p}} \cdot e^{-P}$$

Proof. For $\mathbb{P}[\text{idle}]$ it holds that

$$\mathbb{P}[\text{idle}] \;=\; \prod_{v\in V}(1-p_v) \geq \prod_{v\in V} e^{-\frac{p_v}{1-p_v}} \geq \prod_{v\in V} e^{-\frac{p_v}{1-\hat{p}}} = e^{-\frac{P}{1-\hat{p}}}$$

and

$$\mathbb{P}[\text{idle}] \;=\; \prod_{v\in V}(1-p_v) \leq \prod_{v\in V} e^{-p_v} = e^{-P}$$

Next we show a lower bound on the probability of a successful transmission in some given co-existing network N_i:

$$\begin{aligned}
\mathbb{P}[\text{successful in } N_i] &= \sum_{v\in V_i} p_v \cdot \prod_{w\in V\setminus\{v\}} (1-p_w) \geq \sum_{v\in V_i} p_v \cdot \prod_{w\in V}(1-p_w) \\
&\geq \sum_{v\in V_i} p_v \cdot \prod_{w\in V} e^{-\frac{p_w}{1-p_w}} \geq \sum_{v\in V_i} p_v \cdot \prod_{w\in V} e^{-\frac{p_w}{1-\hat{p}}} \\
&= \sum_{v\in V_i} p_v \cdot e^{-\frac{P}{1-\hat{p}}} = p_i \cdot e^{-\frac{P}{1-\hat{p}}}
\end{aligned}$$

Finally, we derive an upper bound on $\mathbb{P}[\text{successful at network } i]$:

$$\begin{aligned}
\mathbb{P}[\text{successful in } N_i] &= \sum_{v\in V_i} p_v \cdot \prod_{w\in V\setminus\{v\}} (1-p_w) \leq \frac{1}{1-\hat{p}} \cdot \sum_{v\in V_i} p_v \cdot \prod_{w\in V}(1-p_w) \\
&\leq \frac{1}{1-\hat{p}} \cdot \sum_{v\in V_i} p_v \cdot \prod_{w\in V} e^{-p_w} = p_i \cdot \frac{e^{-P}}{1-\hat{p}}
\end{aligned}$$

5.3.2 Cumulative Probability

In the following, we will derive the first fundamental property of our protocol: we show that the overall cumulative probability $P = \sum_{i=1}^{K} P_i$ converges to some range of values so that the contention on the wireless medium is moderate. This is a necessary condition for a good performance. Our proof framework basically follows the framework of [6] but the proof arguments significantly differ in various places when it comes to analyzing the specifics of our new protocol. We refer to Section 2 of [6] for a comparison.

The proof works by induction over sufficiently large time frames. Let I be a time frame consisting of $\frac{\alpha}{\varepsilon}\log N$ *subframes* I' of size $f = \max\{T, \frac{\alpha\beta^2}{\varepsilon\gamma^2}\log^3 N\}$ rounds, where α and β are sufficiently large constants and $N = \max\{T, n\}$. Let $F = \frac{\alpha}{\varepsilon}\log N \cdot f$ denote the size of I.

First, we show that for any subframe I' in which initially the overall cumulative probability is at least $1/(f^2(1+\gamma)^{2\sqrt{f}})$, also afterwards this cumulative probability is at least $1/(f^2(1+\gamma)^{2\sqrt{f}})$, w.h.p.

Lemma 5.4 *For any subframe $I' = [t_0, t_1)$ in which $P(t_0) \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$, also $P(t_1) \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ w.h.p.*

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

Claim 5.5 *If in subframe I', T_v is decreased at most k times, then node v increases T_v by 2 at most $k/2+\sqrt{f}$ many times.*

Proof. Only an idle time step can potentially reduce T_v by 1. If there is no idle time step during the last T_v many steps, T_v is increased by 2. Suppose that $k = 0$. Then the number of times a node v increases T_v by 2 is upper bounded by the largest possible ℓ so that $\sum_{i=0}^{\ell} T_v^0 + 2i \leq f$, where T_v^0 is the initial value of T_v. For any $T_v^0 \geq 1$, $\ell \leq \sqrt{f}$, so the claim is true for $k = 0$. For each decrease of T_v, the current T_v as well as all subsequent values of T_v (until a T_v is reached with $T_v = 1$) get reduced by one. Hence, for an arbitrary value of $k \geq 0$ we are searching for the maximum ℓ so that $\sum_{i=0}^{\ell} \max\{T_v^0 + 2i - k, 1\} \leq f$. This ℓ is at most $k/2+\sqrt{f}$, which proves our claim.

This claim allows us to prove that the overall cumulative probability P will exceed a certain threshold in a subframe w.h.p.

Claim 5.6 *Suppose that in $I' = [t_0, t_1)$, $P(t_0) \in [1/(f^2(1+\gamma)^{\sqrt{2f}}), 1/f^2]$. Then there is a time step t in I' with $P(t) \geq 1/f^2$, w.h.p.*

Proof. Suppose that there are g non-jammed time steps in I'. Let k_0 be the number of these steps with an idle channel and k_1 be the number of these steps with a successful message transmission in any of the co-existing networks. Let the binary random variable X_i be 1 if and only if the nodes increase their access probabilities in the i-th idle time step in I', and let $X = \sum_{i=1}^{k_0} X_i$. Furthermore, let k_2 be the maximum number of times a node v increases T_v by 2 in I'.

Suppose for the moment that $P(t_0) = 1/f^2$. If all time steps t in I' satisfy $P(t) \leq 1/f^2$, then it must hold that the total decrease of $P(t)$ in I' (due to successful transmissions and cases in which access probabilities are decreased when $c_v > T_v$), which is at most $(1+\gamma)^{k_1+k_2}$, has to be at least as large as the total increase of $P(t)$ (due to idle time steps), which is equal to $(1+\gamma)^X$. Hence, we must have that $X \leq k_1 + k_2$. For an arbitrary initial probability $P(t_0) \leq 1/f^2$, we must therefore have

$$X - \log_{1+\gamma}((1/f^2)/P(t_0)) \leq k_1 + k_2 \tag{5.1}$$

to avoid a time step t in I' with $P(t) > 1/f^2$. Our goal is to show that this inequality is violated w.h.p., which implies that I' has a time step t with $P(t) > 1/f^2$ w.h.p.

Next, we focus on k_2. Consider some fixed $k_0 \geq 2$ (as we will see later, $k_0 \geq 2$ w.h.p.). Let L_i be the L-value of the nodes at the i-th idle time step (note that they are all the same) and let $q_i = 1/L_i$ denote the increase of the q-values of the nodes in the i-th idle time step. Also, let $\bar{q} = \frac{1}{k_0-1}\sum_{i=2}^{k_0} q_i$. Certainly, the number of times any node v decreases T_v in I' is bounded by the number of times q is at least 1, which is at most $\lceil \sum_{i=1}^{k_0} q_i \rceil \leq \lceil 1 + (k_0-1)\bar{q} \rceil$. Hence, it follows from Claim 5.5 that

$$k_2 \leq \lceil \bar{q}(k_0-1)+1 \rceil/2 + \sqrt{f} \tag{5.2}$$

On the other hand, the number of times any node v increases p_v in I' is at least $\lfloor \sum_{i=2}^{k_0} q_i \rfloor = \lfloor (k_0-1)\bar{q} \rfloor$ (because due to Fact 5.2 it follows from $P(t) \leq 1/f^2$ that $p_v(t) < \hat{p}$ for all v). Plugging this together with (5.2) into (5.1) and using the fact that $P(t_0) \geq 1/(f^2(1+\gamma)^{\sqrt{2f}})$, we obtain

$$\begin{aligned} \lfloor (k_0-1)\bar{q} \rfloor - \lceil (k_0-1)\bar{q}+1 \rceil/2 &\leq \sqrt{2f} + k_1 + \sqrt{f} \\ \Rightarrow \quad (k_0-1)\bar{q}/2 &\leq k_1 + 4\sqrt{f} \end{aligned} \tag{5.3}$$

given that f is large enough. It remains to lower bound $\bar{q}$ and k_0 and to upper bound k_1 in order to arrive at a contradiction.

We start with $\bar{q}$. Let $\bar{L} = \frac{1}{k_0-1}\sum_{i=2}^{k_0} L_i$. Since $\sum_{i=2}^{k_0} L_i < f$, it holds that $\bar{L} < \frac{f}{k_0-1}$. Moreover, we make use of the following well-known fact.

Fact 5.7 *For any sequence of positive numbers $x_1, \ldots, x_n$ it holds for its arithmetic mean $A = (1/n)\sum_{i=1}^{n} x_i$ and its harmonic mean $H = ((1/n)\sum_{i=1}^{n} 1/x_i)^{-1}$ that $A \geq H$.*

Hence, it follows that $\bar{L} \geq 1/(\frac{1}{k_0-1}\sum_{i=2}^{k_0} 1/L_i)$ and therefore, $\frac{1}{k_0-1}\sum_{i=2}^{k_0} 1/L_i \geq 1/\bar{L}$. This in turn implies that

$$\bar{q} \geq 1/\bar{L} \geq \frac{k_0 - 1}{f}$$

Next we provide an upper bound for k_1 that holds w.h.p. Certainly, for any time step t with $P(t) \leq 1/f^2$,

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

Hence, $\mathbb{E}[k_1] \leq g \cdot (1/f^2) \leq 1/f$. In order to prove an upper bound on k_1 that holds w.h.p., we can use the general Chernoff bounds stated in Lemma 1.2. For any step t let the binary random variable Y_t be 1 if and only if at least one message is transmitted successfully at time t and $P(t) \leq 1/f^2$. Then

$$\begin{aligned} \mathbb{P}[Y_t = 1] &= \mathbb{P}[P(t) \leq 1/f^2] \cdot \\ &\quad\ \mathbb{P}[\text{successful msg transmission} \mid P(t) \leq 1/f^2] \\ &\leq 1/f^2. \end{aligned}$$

Moreover, it certainly holds for any set S of time steps prior to some time step t that

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

Therefore, we have

$$\begin{aligned} &\mathbb{P}[\prod_{s \in S} Y_s = 1] \\ &= \mathbb{P}[Y_1 = 1] \cdot \mathbb{P}[Y_2 = 1 | Y_1 = 1] \cdot \mathbb{P}[Y_3 = 1 | \prod_{s=1,2} Y_s = 1] \cdot \ldots \\ &\quad \cdot \ \mathbb{P}[Y_{|S|} = 1 | \prod_{s=1,2,\ldots,|S|-1} Y_s = 1] \\ &\leq (1/f^2)^{|S|} \end{aligned}$$

and

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

Thus, the Chernoff bounds and our choice of f imply that w.h.p. either $\sum_{t \in I'} Y_t < \varepsilon^2 f/8$ and $P(t) \leq 1/f^2$ throughout I', or there must be a time step t in I' with $P(t) > 1/f^2$, which would finish the proof. Therefore, unless $P(t) > 1/f^2$ at some point in I', $k_1 < \varepsilon^2 f/8$ w.h.p.

Next we prove a lower bound on k_0 that holds w.h.p. For any time step t with $P(t) \leq 1/f^2$ it holds that

$$\mathbb{P}[\text{channel is idle}] \geq e^{-P(t)/(1-\hat{p})} \geq 1 - \frac{P(t)}{1-\hat{p}} \geq 1 - 1/f$$

Hence, $\mathbb{E}[k_0] \geq g \cdot (1 - 1/f) \geq \varepsilon f (1 - 1/f)$. Using similar arguments as for k_1, it follows that $k_0 > (7/8)\varepsilon f$ w.h.p. unless $P(t) > 1/f^2$ at some point in I'. When combining the bounds for $\bar{q}$ and k_0, we obtain

$$\begin{aligned} (k_0 - 1)\bar{q}/2 \ &\geq \ \frac{(k_0 - 1)^2}{2f} \ \geq \ (7/8)^2 \varepsilon^2 f/2 \\ &> \ \varepsilon^2 f/8 + 4\sqrt{f} \ > \ k_1 + 4\sqrt{f} \end{aligned}$$

w.h.p., if f is large enough, which violates Inequality (5.3) and therefore completes the proof of Claim 5.6.

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

Claim 5.8 *Suppose that for the first time step t_0 in I', $P(t_0) \geq 1/f^2$. Then there is no time step t in I' with $P(t) < \frac{1}{f^2(1+\gamma)^{\sqrt{2f}}}$, w.h.p.*

Proof. Consider some fixed subinterval $I'' = [t_1, t_2)$ in I' with the property that $P(t_1) \geq 1/f^2$ and $P(t) \leq 1/f^2$ for all other t in I'' (i.e., we will use conditional probabilities based on $P(t) \leq 1/f^2$ like in the bound for k_1 in the proof of Claim 5.6).

Suppose that there are g non-jammed time steps in I''. If $g \leq \beta \log N$ for a (sufficiently large) constant β, then it follows for the probability $P(t_2)$ at the end of I'' that

$$P(t_2) \geq \frac{1}{f^2} \cdot (1+\gamma)^{-((3/2)\beta \log N + \sqrt{f})} \geq \frac{1}{f^2(1+\gamma)^{\sqrt{2f}}}$$

given that f is large enough (i.e., $\varepsilon = \Omega(1/\log^3 N)$). This is because in the worst case for the decrease of $P(t)$ all non-jammed time steps are successful. In this case, $P(t)$ is decreased at most $\beta \log N$ times due to these steps. Moreover, from Claim 5.5 it follows that $P(t)$ can be decreased another at most $\beta \log N/2 + \sqrt{f}$ times due to $c_v > T_v$.

So suppose that $g > \beta \log N$. Let X be the number of time steps in I'' in which $P(t)$ increases and k_1 be the maximum number of time steps in I'' (over all networks) with a successful message transmission. Furthermore, let k_2 be the maximum number of times a node v increases T_v in I''. If $P(t_2) < \frac{1}{f^2(1+\gamma)^{\sqrt{2f}}}$ then it must hold that the total increase in $P(t)$ (which is equal to $(1+\gamma)^X$) is at most the total decrease in $P(t)$ (which is at most $(1+\gamma)^{k_1+k_2}$), or in other words,

$$X \leq k_1 + k_2.$$

From the previous claim we know that this is not true w.h.p. given that $P(t) \leq 1/f^2$ for all $t > t_1$ in I'' and the constant β is sufficiently large to achieve polynomially small probability bounds. Since there are at most f^2 possible values for t_1 and t_2, there is no time step t_2 in I' with $P(t_2) < \frac{1}{f^2(1+\gamma)^{\sqrt{2f}}}$ w.h.p., which completes the proof.

Combining Claims 5.6 and 5.8 completes the proof of Lemma 5.4. ∎

Next we show an upper bound for $P(t)$. In the following, $K' = O(K)$ is a sufficiently large constant $\geq K$.

Lemma 5.9 *For any subframe $I' = [t_0, t_1)$ with $T_v \leq (3/4)\sqrt{F}$ for all nodes v at the beginning of I', $P(t_1) \leq 12 \ln K'$ w.m.p.*

Proof. First, we will show that if $P(t) \geq 4 \ln K'$ throughout I', then for each N_i, there must be a step t' with $P_i(t') \leq (2 \ln K')/K'$ w.h.p., and once such a step is reached, we show that $P_i(t'') < (4 \ln K')/K'$ w.m.p. for all time steps t'' following t'. Hence, there must be a time step t'' in I' with $P_i(t'') < (4 \ln K')/K'$ for all i, w.m.p., contradicting the assumption that $P(t) \geq 4 \ln K'$ throughout I'. Once we have that, we will show that at the end of I', $P(t_1) \leq 12 \ln K'$ w.m.p.

Consider some fixed network i. Let k_0 be the number of idle steps in I' and k_1 be the number of successful time steps for network i. Moreover, let X be the total number of times $P_i(t)$ is increased by $(1+\gamma)$ due to an idle channel in I'. For N_i to avoid a time step t' in I' with $P_i(t') \leq (2 \ln K')/K'$, we must have that the total increase of $P_i(t)$ (which is equal to $(1+\gamma)^X$) is at least the total decrease of $P_i(t)$ once we have reached a point t with $P_i(t) = (2 \ln K')/K'$, which is the case after at most $\log_{1+\gamma}(n_i \cdot \hat{p})$ reductions of $P_i(t)$. Hence, we must have

$$X \quad \geq \quad k_1' - \log_{1+\gamma}(n_i \cdot \hat{p}) \tag{5.4}$$

where k_1' is the total decrease (in the exponent) of $P_i(t)$ due to successful transmissions to avoid a time step t' in I' with $P_i(t') \leq (2 \ln K')/K'$. Notice that k_1' is not equal to k_1 because if, for example, a node successfully transmits twice in a row, $P_i(t)$ does not get decreased the second time.

In order to contradict this bound, we first need to have a closer look at what happens when there is a successful transmission in N_i.

Claim 5.10 *If the node v successfully transmitting a message in N_i at time t is different from the node that previously successfully transmitted a message in N_i,* *then* $P_i(t+1) \in [\frac{1}{1+\gamma}P_i(t), \frac{1}{\sqrt{1+\gamma}}P_i(t)]$ *for any* $n_i \geq 2$.

Proof. The lower bound is obvious. Moreover, it follows from the protocol that

$$\begin{aligned} P_i(t+1) &= p_{v,t} + \sum_{w \in V_i \setminus \{v\}} \frac{1}{1+\gamma} \cdot p_{v,t} \\ &= \frac{1}{1+\gamma} \cdot P_i(t) + \frac{\gamma}{1+\gamma} \cdot p_{v,t} \\ &\leq \frac{1}{1+\gamma} \cdot P_i(t) + \frac{\gamma}{1+\gamma} \cdot \frac{P_i(t)}{n_i} \\ &= \frac{1}{1+\gamma}\left(1+\frac{\gamma}{n_i}\right) P_i(t) \\ &\leq \frac{1}{1+\gamma}(1+\gamma)^{1/n_i} P_i(t) \leq \frac{1}{\sqrt{1+\gamma}} P_i(t) \end{aligned}$$

given that $n_i \geq 2$.

If the same node v successfully transmits again at time t, then $P_i(t+1) = P_i(t)$, which only happens with probability at most $(1+\gamma)/n_i$ because in this case the transmitting node has an access probability that is by a $(1+\gamma)$ factor larger than the other access probabilities in N_i. Hence, on expectation, at least $1/3$ of the time steps with successful transmission, $P_i(t)$ is reduced by at least $(1+\gamma)^{1/2}$, which implies that $\mathbb{E}[k_1'] \geq k_1/6$.

Based on this insight, the next claim shows that under certain conditions, Inequality (5.4) is not true w.h.p. Let g_i be the number of *useful* time steps for N_i, which are time steps that are either idle or successful for N_i in I'.

Claim 5.11 *If all time steps* $t \in I'$ *satisfy* $P(t) \geq 4 \ln K'$ *and* $g_i \geq \delta \log_{1+\gamma} N$ *for a sufficiently large constant* δ*, then* $X + \log_{1+\gamma} n_i < k_1'$ *w.h.p.*

Proof. It is easy to see that for any useful time step t,

$$\mathbb{P}[t \text{ successful for } N_i] \quad \geq \quad P_i(t) \cdot \mathbb{P}[t \text{ idle}] \tag{5.5}$$

and therefore $\mathbb{E}[k_1] \geq \frac{2\ln K'}{K'}\mathbb{E}[k_0]$ unless there is a time step t with $P_i(t) < (2\ln K')/K'$. For a given number of useful time steps g_i, since $k_0 + k_1 = g_i$ and therefore also $\mathbb{E}[k_0] + \mathbb{E}[k_1] = g_i$, $\mathbb{E}[k_1] \geq \frac{2\ln K'}{K'}(g_i - \mathbb{E}[k_1])$, which implies that $\mathbb{E}[k_1] \geq \frac{\ln K'}{K'} \cdot g_i$ if $K' = O(K)$ is a sufficiently large constant. Since $\mathbb{E}[k_1'] \geq k_1/6$, $g_i = \Omega(\log_{1+\gamma} N)$, and for each useful time step there is an independent probability whether this time step is idle or successful, it follows from the Chernoff bounds that $k_1' \geq \frac{\ln K'}{8K'} g_i$ w.h.p.

Next we bound X. Let the binary random variable X_j denote the increase of $P_i(t)$ by $(1+\gamma)^{X_j}$ in the j-th idle time step. Then $X = \sum_{j=1}^{k_0} X_j$. Moreover, let L_j be the number of time steps between the $(j-1)$-th and j-th idle time steps. It holds that

$$\mathbb{P}[t \text{ idle}] \leq e^{-P(t)} \leq 1/(K')^4$$

for every $t \in I'$ given that $P(t) \geq 4\ln K'$. Hence,

$$\begin{aligned}
\mathbb{E}[X_j] &= \sum_{\ell \geq 1} \mathbb{P}[L_j = \ell] \cdot 1/\ell \leq \sum_{\ell \geq 1} \frac{1}{(K')^4}\left(1 - \frac{1}{(K')^4}\right)^{\ell-1} \cdot \frac{1}{\ell} \\
&\leq \frac{1}{(K')^4 - 1} \sum_{\ell \geq 1} e^{-\ell/(K')^4}/\ell \leq \frac{1}{(K')^4 - 1} \cdot 2\ln(K')^4 \\
&= \frac{4\ln K'}{(K')^4 - 1}
\end{aligned}$$

and therefore, $\mathbb{E}[X] \leq \frac{4\ln K'}{(K')^4 - 1} \cdot k_0 \leq \frac{4\ln K'}{(K')^4 - 1} \cdot g_i$. Since the upper bound on $\mathbb{E}[X_j]$ holds independently for each j, it follows from the Chernoff bounds that $X \leq \frac{6\ln K'}{(K')^4} \cdot g_i$ w.h.p.

Since $g_i = \Omega(\log_{1+\gamma} N)$, $X + \log_{1+\gamma} n_i < k_1'$ w.h.p. if $K' = O(K)$ is sufficiently large, which completes the proof of the claim.

Otherwise, suppose that $g_i < \delta \log_{1+\gamma} N$. For every node v it follows from the COMAC protocol and the choice of f and F that if initially $T_v \leq (3/4)\sqrt{F}$, then

T_v can be at most $\sqrt{F}$ during I'. Let us cut I' into m intervals of size $2\sqrt{F}$ each. It is easy to check that if β in the definition of f is sufficiently large compared to δ, then $m \geq 3\delta \log_{1+\gamma} N$. Since there are less than $\delta \log_{1+\gamma} N$ useful steps in N_i in I', at least $2\delta \log_{1+\gamma} N$ of these intervals do not contain any useful step, which implies that p_v is reduced by $(1+\gamma)$ by each $v \in V_i$ in each of these intervals.

Hence, altogether, every p_v gets reduced by a factor of at least $(1+\gamma)^{-2\delta \log_{1+\gamma} N}$ during I' in N_i. The useful time steps can only raise that by at most $(1+\gamma)^{\delta \log_{1+\gamma} N}$, so altogether we must have $P_i(t') \leq (2\ln K')/K'$ at some time point t' in I', w.h.p.

Next we prove the following claim, which implies that for all $t'' > t'$ in I', $P_i(t'') < (4\ln K')/K'$ w.m.p.

Claim 5.12 *If all time steps $t \in I'$ satisfy $P(t) \geq 4\ln K'$ and initially $P_i(t) \leq (2\ln K')/K'$, then for all steps $t \in I'$, $P_i(t) \leq (4\ln K')/K'$ w.m.p.*

Proof. Consider some fixed subinterval $I'' = [t_1, t_2)$ in I' with the property that $P_i(t_1) \leq (2\ln K')/K'$ and $P_i(t) \geq (2\ln K')/K'$ for all other t in I''. Suppose that there are g_i useful time steps in I''. If $g_i \leq \ln_{1+\gamma} 2$, then it follows for the probability $P_i(t_2)$ at the end of I'' that $P_i(t_2) \leq \frac{2\ln K'}{K'} \cdot (1+\gamma)^{\ln_{1+\gamma} 2} \leq \frac{4\ln K'}{K'}$. Otherwise, suppose that $g_i > \ln_{1+\gamma} 2$, which is at least $1/(2\gamma) = \Omega(\ln f)$. Let X be the number of time steps in I'' in which $P_i(t)$ increases and k_1 be the number of time steps in I'' with a successful transmission in N_i. Furthermore, let k_2 be the maximum number of times a node $v \in V_i$ increases T_v in I''. If $P(t_2) > (4\ln K')/K'$ then it must hold that the total increase in $P_i(t)$ (which is equal to $(1+\gamma)^X$) is at least the total decrease in $P(t)$ (which is at most $(1+\gamma)^{k_1+k_2}$) plus $\ln_{1+\gamma} 2$, or formally,

$$X \geq k_1' + \ln_{1+\gamma} 2 \tag{5.6}$$

where k'_1 is the total decrease (in the exponent) of $P_i(t)$ due to successful transmissions. We know that $\mathbb{E}[k'_1] \geq k_1/6$. Also, from the proof of the previous claim it follows that $\mathbb{E}[k_1] \geq \frac{\ln K'}{K'} g_i$ if $K' = O(K)$ is a sufficiently large constant, unless there is a time step t in I' with $P_i(t) < (2\ln K')/K'$. Since $g_i = \Omega(\ln f)$, it follows from the Chernoff bounds that $k'_1 \geq \frac{\ln K'}{8K'} g_i$ w.m.p. On the other hand, it follows from the proof of the previous claim that $X \leq \frac{6\ln K'}{(K')^4} \cdot g_i$ w.m.p. Hence, inequality (5.6) is violated w.m.p., which implies that $P_i(t_2) \leq \frac{4\ln K'}{K'}$ w.m.p. Since there are at most f^2 different values of t_1 and t_2, there is no time step t_2 in I' with $P_i(t_2) > \frac{4\ln K'}{K'}$ w.m.p., which completes the proof.

Combining the insights above, it follows that there must be a time step t in I' with $P(t) < 4\ln K'$ w.m.p. To finish the proof, we need the following claim.

Claim 5.13 *If for the first time step t_0 in I', $P(t_0) \leq 4\ln K'$, then $P(t) \leq 12\ln K'$ for all time steps t in I' w.m.p.*

Proof. Consider some subinterval $I'' = [t_1, t_2)$ in I' with the property that $P(t_1) \leq 4\ln K'$ and $P(t) \geq 4\ln K'$ for all $t > t_1$ in I''. Suppose that there are g useful time steps in I'', where a time step is useful if there was either a successful transmission in some network or the channel is idle. If $g \leq \log_{1+\gamma} 2$, then certainly $P(t) \leq 12\ln K'$ for all t in I'. So suppose that $g > \log_{1+\gamma} 2$. Consider some fixed network N_i. Let X be the number of time steps in I'' in which $P_i(t)$ increases and k_1 be the number of time steps in I'' with a successful message transmission in N_i. Furthermore, let k_2 be the maximum number of times a node $v \in V_i$ increases T_v in I''. If $P(t_2) > 12\ln K'$ then there must be a network N_i with $P_i(t_2) > \max\{(8\ln K')/K', 2P_i(t_1)\}$. To see this, let I_1 be the set of all i with $P_i(t_1) < (4\ln K')/K'$ and I_2 be the set of all other i. As long as for all i, $P_i(t_2) \leq \max\{(8\ln K')/K', 2P_i(t_1)\}$, it must hold that $P(t_2) \leq$

$\sum_{i\in I_1}(8\ln K')/K' + \sum_{i\in I_2} 2P_i(t_1) \leq (8\ln K')/K' \cdot K + 2P(t_1) \leq 12\ln K'$ if $K' = O(K)$ is sufficiently large.

First, consider the case that for some i with $P_i(t_1) \geq (4lnK')/K'$, $P_i(t_2) > 2P_i(t_1)$. Then the total increase of $P_i(t)$ in I'' (which is equal to $(1+\gamma)^X$ is at least the total decrease in $P_i(t)$ plus $\log_{1+\gamma} 2$. Hence,

$$X \geq k_1' + \log_{1+\gamma} 2 \tag{5.7}$$

where k_1' is the total decrease (in the exponent) of $P(t)$ due to successful transmissions in N_i. From Inequality (5.5) we know that $\mathbb{E}[k_1] \geq \frac{4\ln K'}{K'} \cdot \mathbb{E}[k_0]$ and therefore $\mathbb{E}[k_1] \geq \frac{2\ln K'}{K'} \cdot g$ if $K' = O(K)$ is large enough. Since $\mathbb{E}[k_1'] \geq k_1/6$ and $g = \Omega(\ln f)$ it follows from the Chernoff bounds that $k_1' \geq \frac{\ln K'}{4K'} \cdot g$ w.m.p. On the other hand, we also know that $X \leq \frac{6\ln K'}{(K')^4} \cdot g$ w.m.p., which implies that Inequality (5.7) is violated w.m.p. Hence, $P_i(t_2) \leq 2P_i(t_1)$ w.m.p.

For the case that $P_i(t_1) < (4\ln K')/K'$ let t_1' be the first step in I'' with $P_i(t_1') \geq (4\ln K')/K'$. If t_1' does not exist, we are done, and otherwise we prove in the same way as above that w.m.p. $P_i(t_2) \leq (12\ln K')/K'$.

Since there are at most f^2 ways of choosing t_1 and t_2, there is no time step t in I' with $P(t) \leq 12\ln K'$ w.m.p., which completes the proof.

All claims combined imply Lemma 5.9. ∎

A proof similar to Lemma 5.9 also implies the following result.

Corollary 5.14 *For any subframe I' that satisfies $P(t) \leq 12 \ln K'$ at the beginning of I', all time steps t of I' satisfy $P(t) \leq 36 \ln K'$ w.m.p.*

We also need to show that for a constant fraction of the non-jammed time steps in a subframe where initially $P(t) \leq 12 \ln K'$, $P(t)$ is also lower bounded by a constant for a sufficiently large fraction of time steps t.

Lemma 5.15 *For any subframe I' in which initially $P(t_0) \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$, at least $\varepsilon/8$ of the non-jammed steps t satisfy $P(t) \geq \varepsilon \hat{p}/4$, w.h.p.*

Proof. Let G be the set of all non-jammed time steps in I' and S be the set of all steps t in G with $P(t) < \varepsilon \hat{p}/4$. Let $g = |G|$ and $s = |S|$. If $s \leq (1-\varepsilon/8)g$, we are done. Hence, consider the case that $s \geq (1-\varepsilon/8)g$.

Suppose that $P(t)$ must be increased ℓ many times to get from its initial value up to a value of $\varepsilon \hat{p}/4$. (If $P(t_0) \geq \varepsilon \hat{p}/4$ then $\ell = 0$.) Let k_0 be the number of time steps in S with an idle channel and k_1 be the number of time steps in S with a successful message transmission in any of the co-existing networks. Let the binary random variable X_i be 1 if and only if the nodes increase their access probabilities in the i-th idle time step in S, and let $X = \sum_{i=1}^{\ell} X_i$. Furthermore, let k_2 be the maximum number of times a node v decreases p_v due to $c_v > T_v$ in I'. For S to be feasible (i.e., probabilities can be assigned to each $t \in S$ so that $P(t) < \varepsilon \hat{p}/4$), we must have

$$X \quad \leq \quad \ell + k_1 + k_2 \tag{5.8}$$

For the special case that $\ell = k_2 = 0$ this follows from the fact that whenever there is a successful message transmission, $P(t)$ is reduced by $(1+\gamma)^{-1}$, at most. On the other hand, whenever the nodes decide to increase $P(t)$ for some $t \in S$, $P(t)$ can indeed increase because of $P(t) < \varepsilon\hat{p}/4$ and therefore $p_v < \hat{p}$ for all v. Thus, if $X > k_1$, then one of the steps in S would have to have a probability of at least $\varepsilon\hat{p}/4$, violating the definition of S. ℓ comes into the formula due to the startup cost of getting to a value of $\varepsilon\hat{p}/4$, and k_2 comes into the formula since the reductions of the $p_v(t)$ values due to $c_v > T_v$ allow up to k_2 additional decreases of $P(t)$ for S to stay feasible.

Certainly, $\ell \leq 2\log_{1+\gamma} f + 2\sqrt{f}$. Moreover, for k_1 it holds that $\mathbb{E}[k_1] \leq \varepsilon\hat{p}/4 \cdot s$ and therefore, $k_1 \leq \varepsilon\hat{p}/2 \cdot s$ w.h.p. For k_2 it holds that $k_2 \leq (X + \varepsilon g/8)/2 + \sqrt{f}$. Hence, Inequality (5.8) implies that

$$\begin{aligned} X &\leq 2\log_{1+\gamma} f + 2\sqrt{f} + \varepsilon\hat{p}s/2 + (X + \varepsilon g/8)/2 + \sqrt{f} \\ \Rightarrow \quad X &\leq (\hat{p} + 1/16)\varepsilon g + 8\sqrt{f} \end{aligned} \tag{5.9}$$

if f is sufficiently large. It remains to compute a lower bound for X.

Let X' be the total number of times $P(t)$ is increased over all time steps in G, k_0' be the number of idle time steps in G, and $\bar{q}$ be the average increase of the q_v-values in I'. From the proof of Claim 5.6 we know that $\bar{q} \geq (k_0' - 1)/f$ and that $X' \geq \lfloor (k_0' - 1)\bar{q} \rfloor$. Moreover, $X \geq X' - \varepsilon g/8$. Hence, $X \geq \lfloor (k_0 - 1)^2/f \rfloor - \varepsilon g/8$. We know that $\mathbb{E}[k_0] \geq (1 - \varepsilon\hat{p}/4)s$ and therefore, $k_0 \geq 3g/4$ w.h.p. Hence, $X \geq g^2/(4f) - \varepsilon g/8 \geq \varepsilon g/8$ w.h.p. Since this violates Inequality (5.9), the lemma follows.

In the following, let us call a subframe I' *good* if its initial step t_0 satisfies $P(t_0) \leq 12\ln K'$. Combining the results above, we get:

Lemma 5.16 *For any good subframe I', there are at least $\varepsilon^2 f/8$ non-jammed time steps t in I' with $P(t) \in [\varepsilon\hat{p}/4, 36\ln K']$ w.m.p.*

Consider now the first eighth of frame I, called J. The following lemma follows directly from Lemma 2.14 in [6].

Lemma 5.17 *If at the beginning of J, $p_v \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ and $T_v \leq \sqrt{F}/2$ for all nodes v, then we also have $p_v \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ at the end of J for every v and the number of non-jammed time steps t in I' with $P(t) \in [\varepsilon\hat{p}/4, 36\ln K']$ is at least $\varepsilon^2 f/16$ w.h.p.*

We finally need the following lemma, which follows from Lemma 2.15 in [6].

Lemma 5.18 *If at the beginning of J, $T_v \leq \sqrt{F}/2$ for all v, then it holds that also $T_v \leq \sqrt{F}/2$ at the end of J w.h.p.*

Inductively using Lemmas 5.17 and 5.18 on the eighths of frame I implies that COMAC satisfies the property of Lemmas 5.17 for the entire I and at the end of I, $p_v \geq 1/(f^2(1+\gamma)^{2\sqrt{f}})$ and $T_v \leq \sqrt{F}/2$ for all v w.h.p. Since our results hold with high probability, we can also extend them to any polynomial number of frames.

5.3.3 Throughput

Summarizing the results above, we obtain the following result for the throughput.

Theorem 5.19 *For any polynomial sequence of time steps of length at least F,* CoMAC *achieves a competitive throughput of $\Omega(\varepsilon^2 \min\{\varepsilon, 1/poly(K)\})$ for any constants ε and K.*

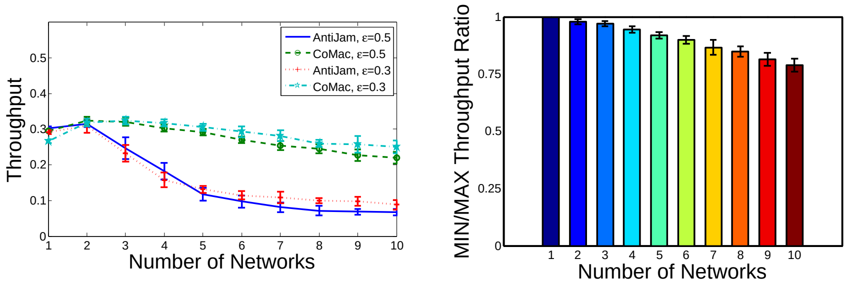

Figure 5.1: *Left:* Throughput of CoMAC and AntiJam [50] as a function of the number of co-existing networks and for two different adversaries ($\varepsilon = \{0.5, 0.3\}$). The total number of nodes for each $K = 1, \ldots, 10$ is 500, and each co-existing network has the same size (up to an additive node due to rounding). The protocol is executed for 7000 rounds, and the result is averaged over 10 runs. The adversary is modeled in a simplified manner and simply jams each round with independent probability $1 - \varepsilon$. *Right:* Fairness as the min/max competitive throughput ratio for $\varepsilon = 0.3$.

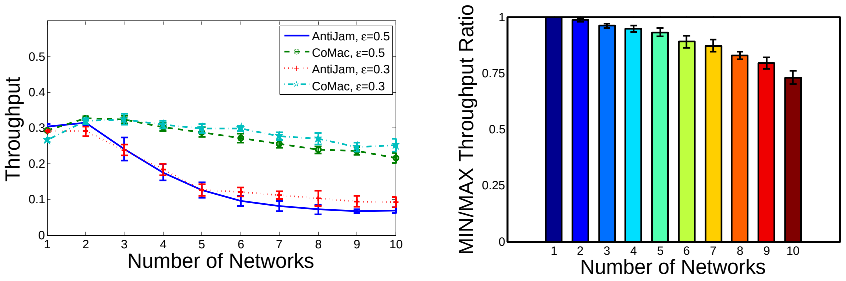

Figure 5.2: *Left:* Throughput and fairness of CoMAC and AntiJam [50] for a setting like in Figure 5.1 but where the size of the co-existing networks is heterogenous, i.e., the i-th largest network is roughly 1.5 times the size of $(i+1)$-largest network. *Right:* Fairness as the min/max competitive throughput ratio for $\varepsilon = 0.3$.

5.3.4 Fairness

Finally, we show that CoMAC also ensures a limited degree of fairness. Note that by Lemma 5.3, we can directly bound the probabilities of having a successful transmission within networks N_i and N_j by their respective cumulative probabilities, which we bound on the following theorem.

Theorem 5.20 *If all nodes v initially start with access probability $\hat{p}$, then it takes at most F time steps until a time step is reached in which the difference between minimum and maximum cumulative probability of a network is at most $O(K^2)$.*

Proof. Consider the potential function $\Phi = \sum_i |x_i - x_{\min}|$ where $x_i = \ln_{1+\gamma} P_i$ and $x_{\min} = \min_i x_i$. We focus on the events with a successful transmission, since only successful transmissions can change the difference among individual network probabilities. Assume a successful transmission occurred in network N_i, if $x_i > x_{\min}$, then the change in Φ, denoted by $\Delta\Phi$, satisfies $\Delta\Phi = -1$. If $x_i = x_{\min}$, then $\Delta\Phi \leq K$. Hence, $\mathbb{E}[\Delta\Phi] \leq -\mathbb{P}[x_i > x_{\min}$ & successful$] + K\mathbb{P}[x_i = x_{\min}$ & successful$]$. Suppose that $x_{\max} \geq x_{\min} + \log_{1+\gamma}(2K^2)$. Then, $\mathbb{P}[x_i > x_{\min}$ & successful$] \geq 2K \cdot \mathbb{P}[x_i = x_{\min}$ & successful$]$ as there can be up to $K-1$ many N_i with $x_i = x_{\min}$. Certainly, $\mathbb{P}[x_i > x_{\min}$ & successful$] + \mathbb{P}[x_i = x_{\min}$ & successful$] = 1$ given that there is a successful transmission. Hence in this case, $\mathbb{P}[x_i > x_{\min}$ & successful$] \geq \frac{2K}{2K+1}$, which implies that $\mathbb{E}[\Delta\Phi] \leq -\frac{2K}{2K+1} + \frac{K}{2K+1} = -\frac{K}{2K+1} \leq -1/3$, whenever there is a successful transmission.

Now, let us define the random variable X_t as follows for the t-th successful transmission:

- $X_t = 1$ if either $x_{\max} < x_{\min} + \log_{1+\gamma}(2K^2)$ (i.e., we reached our goal) or the successful transmission is from a network N_i with $x_i > x_{\min}$, and
- $X_t = -K$ otherwise.

Suppose that there are s successful transmissions across all networks. Let $X = \sum_{t=1}^{s} X_t$. Then it holds that $\mathbb{E}[X] \geq s/3$. In order to apply Chernoff bounds, let us define $Y_t = (X_t + K)/(K+1)$ and $Y = \sum_{t=1}^{s} Y_t$. Then Y_t is a binary random variable with $\mathbb{E}[Y_t] \geq (K+1/3)/(K+1)$ and therefore $\mathbb{E}[Y] \geq s(K+1/3)/(K+1)$. Since the upper bound on $\mathbb{E}[Y_t]$ holds irrespective of previous Y_j's, it follows from the Chernoff bounds that $\mathbb{P}[Y \leq (1-\delta)s(K+1/3)/(K+1)] \leq e^{-\delta^2 s/3}$, for any $0 < \delta < 1$. Since $Y = (X + s \cdot K)/(K+1)$, we get $\mathbb{P}[X \leq (1-\delta)s/3 - \delta sK] \leq e^{-\delta^2 s/3}$. If we choose $\delta = 1/(6(K+1/3))$ then $\mathbb{P}[Y \leq (1-\delta)s(K+1/3)/(K+1)] = \mathbb{P}[X \leq s/6]$ and hence, $\mathbb{P}[X \leq s/6] \leq e^{-\delta^2 s/3}$. Now, from Theorem **??** we know that $s = \Omega(\varepsilon^2 \min\{\varepsilon, 1/poly(K)\}F)$ w.h.p., so $s = \omega(K \log N)$. This implies that when running the protocol for F time steps, $X > K \log N$ w.h.p. Thus, if the initial value of the potential Φ_0 is at most $K \log N$, we must have reached a point where $x_{\max} < x_{\min} + \log_{1+\gamma}(2K^2)$ as otherwise we would end up with a negative potential. It remains to bound Φ_0.

Given that all nodes start with the same access probability $\hat{p}$, the maximum initial difference between P_i and P_j for any i and j is N and therefore, $x_{\max} < x_{\min} + \log_{1+\gamma} N$. Hence, $\Phi_0 \leq K \log_{1+\gamma} N$, which implies the theorem.

Fact 5.2 ensures that the access probabilities of the nodes within a network differs by at most a $(1+\gamma)$ factor, ensuring fairness within each network N_i.

5.4 Simulation

Although the focus of this chapter is on the formal, asymptotic and worst-case performance guarantees achieved by COMAC, we also briefly report on some of our quantitative insights from a simulation study. We are interested in: (i) how the competitive throughput of all the networks changes when the number of networks varies, where the competitive throughput of all the networks is defined as the fraction of non-jammed time steps that are used for successful transmissions among all K networks; (ii) the fairness of COMAC, i.e., whether the successful transmissions are evenly distributed among all the networks. Also, we compare COMAC to the state-of-the-art jamming resistent MAC protocol ANTIJAM in [50], and find that COMAC indeed better suits co-existing networks.

There is a total of 500 nodes among all the co-existing networks, and the number of networks K ranges from 1 to 10. All the results are averaged over 10 runs, and the confidence intervals are provided as well. More specifically, we conduct competitive throughput and fairness experiments in two different scenarios.

Scenario 1: The size of individual networks are the same, namely $|V_i| \in \{\lfloor 500/K \rfloor, \lceil 500/K \rceil\}$. In Figure 5.1 (*left*) we study the *competitive throughput*, i.e., the fraction of non-jammed time steps that are used for successful transmissions among all K networks. We observe that for a single network ($K = 1$) the competitive throughput of COMAC is relatively worse compared to ANTIJAM as p_v is raised more strictly when the channel is idle. However, COMAC is always better than ANTIJAM when there is more than one network ($K > 1$) as the additional interference introduced by co-existing networks is bounded. For example, when $K = 10$, the competitive throughput of COMAC is still above 20% even when adversary can jam 70% of all time steps, while the competitive throughput of ANTIJAM is below 10%. Note that there is a trend towards smaller competitiveness for larger K, as expected from our formal worst-case analysis. Figure 5.1 (*right*) studies the fairness of COMAC in terms of min/max competitive throughput ratio, where the minimum and maximum competitive throughput are selected from the K co-existing networks. The closer this ratio is to 1, the fairer the protocol. Obviously COMAC is fair in a sense that even when $K = 10$, the min/max competitive throughput ratio is above 0.78.

Scenario 2: The size of i-th largest network is roughly 1.5 times the size of

$(i+1)$-th largest network. Figure 5.2 shows that even when the size of individual networks vary a lot, COMAC still achieves a better competitive throughput (above 20% when $K = 10$) compared to ANTIJAM (below 10% when $K = 10$), and more importantly, COMAC is still fair in a sense that the min/max competitive throughput ratio when $K = 10$ is still above 0.73.

5.5 Conclusion

Motivated by our observation that MAC algorithms optimized for a single network often yield a poor performance in scenarios with multiple co-existing networks due to too high sending probabilities, we present the first protocol for provably robust, efficient and fair medium allocation among a set of co-existing networks (e.g., of a multi-nation conference or of an emergency network). Interestingly, with simple adaption, our protocol could even be used in scenarios where the throughput is required to be distributed according to some specific proportions among the co-existing networks (not necessarily fair). For instance, a spectrum owner may require the co-existing networks to use only a share of the medium that corresponds to the negotiated or auctioned share.

We believe that our work raises a series of interesting questions for future research. For example, we have assumed a rather naive interference model and it would be interesting to generalize our results for the SINR physical interference model.

Chapter 6

THE SINRMAC PROTOCOL

The protocols proposed in the previous four chapters (i.e., JADE, ANTIJAM, SELECT, and COMAC) work under the protocol interference model, i.e., the interference is modeled by either single-hop or UDG networks. In the protocol interference model, the impact of interference from neighboring nodes is binary and completely depends on whether or not the node falls within the interference range of non-intended transmitters. In particular, if the UDG is used, then the transmission range and the interference range are the same, and equal to 1 (note that the distance is normalized). Hence, under the UDG model, interference from the nodes outside the receiver node's interference range can be ignored completely, which greatly simplifies the theoretical analysis of the protocols.

However, it seems difficult to go beyond these simplistic interference models. The next big step forward would certainly be a result on the widely used and more realistic *Signal-to-Interference-plus-Noise-Ratio (SINR)* model. A crucial difference from the previous models such as the UDG model is the fact that in the SINR model, nodes cannot always objectively distinguish an idle medium from a busy one. This however was a central assumption of the MAC protocols presented so far as it was used to adjust the nodes' backoff periods: in times of an idling medium, the medium access probability was increased, and in times of a successful transmission, the medium access probability was decreased.

We report on our endeavor to generalize our previous results to the SINR model. Concretely, we describe a first algorithm where each node maintains a noise threshold to determine whether the channel is idle or busy, and then adjust its access probability and noise threshold accordingly in an adaptive fashion.

We assume that the wireless nodes V ($n = |V|$ many) are distributed arbitrarily in the 2-dimensional Euclidean plane, and that they communicate over a wireless network with a single channel. We also assume the nodes are backlogged in the sense that they always have something to broadcast. The SINR model defines a parameter called minimum *signal-to-interference-plus-noise ratio* (SINR) at which a data frame can still be received with a reasonably low frame error rate.[1] In other words, these SINR values specify the transmission range of the data transmission mechanism, i.e., the maximum range within which data frames can still be received correctly. The SINR model is first introduced in [24], which accounts for the SINR at the receiver end of a communication link to determine whether the transmission is successful. More specifically, a message sent from node u to node v is successfully received by node v if and only if

$$\frac{P_v(u)}{\mathcal{N} + \sum_{w \in S} P_v(w)} \geq \beta_1$$

where $P_v(u)$ is the received power at node v of the signal transmitted by node u, $\mathcal{N}$ captures the background noise (e.g., thermal), S is the subset of nodes in $V \setminus \{u, v\}$ that are concurrently transmitting, and β_1 is the *SINR threshold* that depends on the desired rate, the modulation scheme, etc.

In wireless communications, the value of the received signal power at a node r of a signal transmitted by node y, i.e., $P_x(y)$, is a decreasing function of the distance $d(x,y)$ between node x and node y. More specifically,

$$P_x(y) = \frac{P_y}{d(x,y)^{\alpha}}, \tag{6.1}$$

where P_y is the sending power of node y, and $\alpha \geq 2$ is the pass-loss exponent (in [40]), which is a constant between 2 and 6, and depends on external conditions of

[1] For example, according to [55], the minimum SINR for 802.11b are 10dB for 11Mbps down to 4dB for 1Mbps.

the medium as well as the exact sender-receiver distance. Although, in practise the received signal power may be different from 6.1, due to the reason that obstacles in the transmission medium may have an impact on the signal power at the receiver end. To better approximate the path loss, both [55] and [40] use the more generalized signal propagation model. However, as an initial study of designing jamming-resistant MAC protocols under SINR, we focus on the most basic signal propagation model which is indicated by 6.1.

For our formal description and analysis, we assume a synchronized setting where time proceeds in time steps called *rounds*. In each round, a node u may either transmit a message (at a certain power level) or sense the channel, but it cannot do both. A node which is sensing the channel may either (i) sense an *idle* channel, (ii) sense a *busy* channel, or (iii) *receive* a packet.

In the UDG model, the three cases can easily be distinguished in the following manner: idle means no other node in a node u's transmission range is transmitting at that round and the channel is not jammed, busy means two or more nodes in u's transmission range transmit at that round or the channel is jammed, and successful reception occurs if exactly one node in u's transmission range transmits at that round and the channel is not jammed. In the SINR model, things are more complicated. In order to distinguish between an idle and a busy channel, a node may use a certain threshold β_2: if the measured signal power exceeds β_2, a channel is considered busy, otherwise idle. Whether a message is successfully received is determined by the SINR rule described above. (There is at most one successful reception at any moment of time.)

We assume that in addition to the nodes there is an *adversary*: the idea is that our conservative definition of adversary subsumes many different forms of intentional and unintentional interference. Concretely, like in [6], we want to 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, unlike in previous works [6], the adversary is not bounded over time in the sense that it can only jam a subset of the time periods, but *with respect to energy*: for each time period of length T, the adversary has a certain energy budget to disrupt communications. Rather than assuming some jammer locations in the Euclidean plane from which it can transmit at different energy levels, we propose a model where the jammer has a certain budget B_v *for each wireless node* $v \in V$. Henceforth, we assume that this budget is the same for every node and we will simply refer to it by B. Such a jammer is called a (B,T)*-bounded adversary*: in every time interval of size $w \geq T$, the adversary can add $B \cdot w/T$ to the noise level $\mathcal{N}$ of each node.

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 "competitive" throughput against any (B,T)-bounded adversary in any multi-hop network that can be modeled by SINR. Intuitively, we want to call a MAC protocol competitive if the number of successful message receptions at the nodes is a "large" fraction of the messages that would have been received if the adversarial contributions to the noise $\mathcal{N}$ are subtracted in the SINR formula for the corresponding time steps.

6.1 Description of SINRMAC

Basically, the SINRMAC protocol we propose is a *random backoff protocol*, but with a twist: the nodes do not only backoff once their messages collide, but maintain a "backoff counter" which is adapted over time and reflects the current channel state (see also [6]). Rather than storing the backoff counter itself, each node v in SINRMAC stores a medium access probability p_v (between 0 and some upper bound $\hat{p} < 1$). The idea is that in times of an idling channel, p_v is increased (message transmissions become more likely), whereas in times of a busy medium, p_v is decreased. Unfortunately, unlike in the UDG model, such a distinction is not possible in the SINR model, because absolute silence on the channel no longer exists due to background noise and the jammer. Hence, it is hard to tell from a node's point of view that the noise it senses at a particular time step is due to background noise, message collisions, adversarial jamming, or any combination of these.

In SINRMAC, each node v maintains p_v (in some sense, the inverse of a random backoff timer), a noise threshold estimate τ_v to distinguish between idle and non-idle time periods, plus a time window threshold T_v, and a counter c_v. (The threshold T_v is necessary since an accurate estimation of T allows v to adjust its p_v correctly and in a timely manner.) Finally, the nodes share a common small factor γ with which the cumulative sending probabilities are adjusted, and a constant value c, which is used to additively adjust τ_v. In the following, let N_v be the *noise level* (background noise plus concurrent transmissions plus jamming) at node v.

In order to find a good equilibrium and achieve a high throughput, the p_v and τ_v values need to converge to meaningful values quickly. This constitutes a non-trivial challenge. If there are no successful message transmissions, a node v cannot decide whether τ_v is too high or too low. Fortunately, however, in practice one may determine some reasonable upper bound $\hat{\tau}$ for τ_v, as, e.g., (1) the RSSI register (i.e., *Received Signal Strength Indicator* which measures the power of a received radio signal) is of limited size and constitutes a natural upper bound, or as (2) according to [9], a constant density of transmitter nodes in the network implies that interference from far-away nodes can be bounded by a constant. Given such an upper bound, it seems feasible to come up with MAC protocols which find a good equilibrium (in terms of p_v and τ_v values in a certain region), even in the presence of adversaries.

Our solution, the SinrMacprotocol, is formally described in Algorithm 4. The algorithm is essentially interpreting any noise floor smaller than τ_v as an idle channel and increases the sending probabilities accordingly; if on the other hand the noise is relatively high, the sending probabilities are reduced, but only after T_v rounds where the channel was not idle.

In SinrMac, each node adapts τ_v additively and p_v multiplicatively, based on the channel states. Concretely, we decrease τ_v by $2c$ if there is not much noise ($N_v < \tau_v$), but only increase it by c otherwise: thus, in an equilibrium, we strive for a 2 : 1 ratio of busy to idle time periods.

Algorithm 4 SinrMac

```
Initially, every node $v$ sets $T_v := 1$, $c_v := 1$, $p_v := \hat{p}$, and $\tau_v := 0.1$.
Afterwards, the protocol proceeds in synchronized rounds:
$v$ decides with probability $p_v$ to send a message
if $v$ decides not to send a message then
    $v$ senses the channel
    if a message is successfully received then
        $p_v = p_v/(1+\gamma)$
    else if $N_v < \tau_v$ then
        $\tau_v := \max\{\tau_v - 2c, 0\}$
        $p_v := \min\{(1+\gamma)p_v, \hat{p}\}$
        $T_v := \max\{T_v - 1, 1\}$
    else if $N_v \geq \tau_v$ then
        $\tau_v := \min\{\tau_v + c, \hat{\tau}\}$
        if $c_v \geq T_v$ then
            $c_v := 1$
            if no idle channel in past $T_v$ rounds then
                $p_v := p_v/(1+\gamma)$
                $T_v := T_v + 2$
            end if
        end if
    end if
end if
```

6.2 First Results

Although intuitively, adapting τ_v seems to be crucial to accurately react to the channel states and converge to a good throughput, our first experiments indicate that static τ_v values (fixed at the maximal possible reception power) are better, if the fixed value for τ_v is chosen appropriately. In the following, we report on our preliminary simulation study to evaluate the performance of our protocol in terms of throughput and as a function of the network size. We define throughput as the number of messages successfully received in the whole network per round per node. In our network, nodes are distributed uniformly in a two-dimensional plane of size 7×7. The number of nodes $n \in [10, 200]$. We implement a (B, T)-bounded adversary which jams the channel using a random amount of energy from its remaining budget. The transmission power for all nodes is set to 4, the SINR ratio is $\beta_1 = 6$, and $T = 50$. We set $c = 0.1$, and consider $\hat{p} = 1/24$.

We evaluate four different schemes for adapting τ_v: the first one initializes $\tau_v = 1$ and adapts τ_v based on "idle" and "busy" channel states afterwards (see Algorithm 4); the other three schemes use a fixed τ_v (from $\{1, 4, 40\}$).

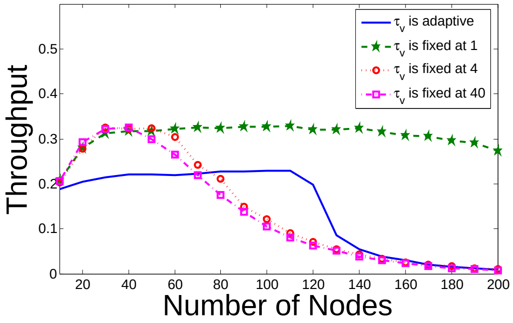

Figure 6.1: Normalized throughput as a function of the network size and under different τ_v adaption schemes. The result is averaged over 10 runs.

Figure 6.1 shows an exemplary dependency of the throughput on the different τ_v schemes when $\hat{p} := 1/24$. We see that fixing τ_v at 1 produce the best throughput result. More specifically, when the network density is low (i.e., $n \leq 70$), fixing τ_v at $\{1,4,40\}$ results better throughput than adapting τ_v. This is because when node density is low, the network needs to have increase cumulative probability in order to experience more successful transmissions, and in order to achieve this, there has to be sufficient number of "idle" channels sensed by the nodes. Fixing τ_v at a relatively high value (note that when τ_v is adaptive, it is initialized to 0.1), nodes would sense "idle" channels more frequently. However, as the network size grows, adaptively adjusting τ_v produces better result than fixing τ_v at $\{4,40\}$, simply because if the network density is high enough, the value of τ_v should be lower, so that by having more "busy" channels, the cumulative probability of the network would be deceased accordingly to maintain at an appropriate level. Although based on our intuition, adapting τ_v should give us better throughput when node density is high, fixing τ_v at 1 makes the throughput result remains around 30%, and is much better than our adaptive strategy. The throughput produced by our adaptive strategy for τ_v drops below 10% when $n \geq 130$. Here, being able to identify the busy channels and decrease access probabilities accordingly is crucial for the protocol to achieve a good throughput. As we can see from the simulation result, our adaptive approach for τ_v still needs improvement, because it cannot always reach to an appropriate value so that the throughput is maximized.

6.3 Conclusion

We propose a preliminary MAC protocol for the SINR model under jamming activities. We found that our adaptive idle/busy threshold adaption strategy scales better than a static strategy.

In our future work, we plan to rigorously evaluate different adapting schemes for τ_v, and study our algorithm under more sophisticated and worst-case adversaries, not only empirically but hopefully also by deriving performance proofs. Obviously, in this process, changes to the protocol presented here may be required.

Chapter 7

CONCLUSION AND FUTURE WORK

In this dissertation, we study the problem of designing and analyzing efficient MAC protocols that are robust against strong adversarial jamming. How to efficiently access the wireless medium, which is a limited resource, is one of the most important problems in wireless computing.

Four jamming-resistant MAC protocols and a leader election protocol, which work under different interference and network models, are presented. More specifically, JADE can achieve constant competitive throughput against the adaptive but non-reactive adversary in multi-hop wireless networks that can be modeled as UDG; ANTIJAM can achieve constant competitive throughput against the adaptive and reactive adversary in single-hop wireless networks; SELECT is a self-stabilizing leader election protocol that is also robust against adaptive and reactive jamming; COMAC can achieve constant competitive throughput as well as fairness for K co-existing networks in a single-hop wireless network environment; SINRMAC is our first attempt to explore the possibility of designing jamming-resistant MAC protocols under SINR model.

The next natural step would be to design a jamming-resistant MAC protocol under SINR model that can achieve provably high throughput. Also, in all the protocols presented in this dissertation, we assume the nodes have a common parameter $\gamma = O(\frac{1}{\log T + \log\log n})$. Although such estimate on $\log\log n$ and $\log T$ still allows for a superpolynomial increase in n and a polynomial increase in T without violating the assumptions on γ, it would of course be more desirable if γ could be set to a constant.

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] Gheorghe Antonoiu, Gheorghe Antonoiu, Pradip K Srimani, and Pradip K Srimani. A self-stabilizing leader election algorithm for tree graphs. *Journal of Parallel and Distributed Computing*, 34:227–232, 1996.

[3] Hirochika Asai, Kensuke Fukuda, and Hiroshi Esaki. Towards characterization of wireless traffic in coexisting 802.11a/g and 802.11n network. In *Proc. ACM CoNEXT Student Workshop*, 2010.

[4] Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics (Chapter 3).* John Wiley & Sons, 2004.

[5] B. Awerbuch. Optimal distributed algorithms for minimum weight spanning tree, counting, leader election, and related problems. In *Proc. STOC*, 1987.

[6] B. Awerbuch, A. Richa, and C. Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *Proc. of PODC '08*, 2008.

[7] 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.

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

[9] D. Blough, C. Canali, G. Resta, and P. Santi. On the impact of far-away interference on evaluations of wireless multihop networks. In *Proc. of the ACM International Symposium on Modeling, Analysis, and Simulation of Wireless and Mobile Systems (MSWIM)*, pages 90–95, 2009.

[10] 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.

[11] Shukai Cai, Taisuke Izumi, and Koichi Wada. Space complexity of self-stabilizing leader election in passively-mobile anonymous agents. In *Proc.*

16th International Colloquium on Structural Information and Communication Complexity (SIROCCO), 2009.

[12] J.T. Chiang and Y.-C. Hu. Cross-layer jamming detection and mitigation in wireless broadcast networks. In *Proc. of MobiCom '07*, pages 346–349, 2007.

[13] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. Adversarial queuing on the multiple-access channel. In *Proc. of PODC '06*, pages 92–101, 2006.

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

[15] Edsger W. Dijkstra. Self-stabilization in spite of distributed control. *Communications of the ACM*, 17(11):643–644, 1974.

[16] 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.

[17] 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.

[18] 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.

[19] Shlomi Dolev, Seth Gilbert, Rachid Guerraoui, Fabian Kuhn, and Calvin C. Newport. The wireless synchronization problem. In *Proc. 28th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 190–199, 2009.

[20] R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.*, 5(1):66–77, 1983.

[21] Sukumar Ghosh and Arobinda Gupta. An exercise in fault-containment: self-stabilizing leader election. *Inf. Process. Lett.*, 59(5):281–288, 1996.

[22] 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.

[23] 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.

[24] Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE TRANSACTIONS ON INFORMATION THEORY*, 46(2):388–404, 2000.

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

[26] Martin Heusse, Franck Rousseau, Romaric Guillier, and Andrzej Duda. Idle sense: An optimal access method for high throughput and fairness in rate diverse wireless lans. In *Proc. Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 121–132, 2005.

[27] J. Hromkovic, R. Klasing, A. Pelc, P. Ruzicka, and W. Unger. *Dissemination of Information in Communication Networks: Broadcasting, Gossiping, Leader Election, and Fault-Tolerance (Chapter 8).* Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.

[28] Gene Itkis, Chengdian Lin, and Janos Simon. Deterministic, constant space, self-stabilizing leader election on uniform rings. In *Proc. 9th International Workshop on Distributed Algorithms (WDAG)*, pages 288–302, 1995.

[29] Kamal Jain, Jitendra Padhye, Venkata N. Padmanabhan, and Lili 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.

[30] Jianbo Ji and Wen Chen. Transmission capacity of two co-existing wireless ad hoc networks with multiple antennas. In *Proc. IEEE International Conference on Communications (ICC)*, pages 1 –6, 2011.

[31] Shanshan Jiang and Yuan Xue. Providing survivability against jamming attack via joint dynamic routing and channel assigment. In *Proc. 7th Workshop on Design of Reliable Communication Networks (DRCN)*, 2009.

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

[33] E. Korach, S. Kutten, and S. Moran. A modular technique for the design of efficient distributed leader finding algorithms. *ACM Trans. Program. Lang. Syst.*, 12(1):84–101, 1990.

[34] Shay Kutten and Boaz Patt-Shamir. Time-adaptive self stabilization. In *Proc. PODC*, 1997.

[35] Y.W. Law, L. van Hoesel, J. Doumen, P. Hartel, and P. Havinga. Energy-efficient link-layer jamming attacks against wireless sensor network MAC protocols. In *Proc. of SASN '05*, pages 76–88, 2005.

[36] Seungjoon Lee, Dave Levin, Vijay Gopalakrishnan, and Bobby Bhattacharjee. Backbone construction in selfish wireless networks. In *Proc. 2007 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 121–132, 2007.

[37] 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.

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

[39] Dominic Meier, Yvonne Anne Pignolet, Stefan Schmid, and Roger Wattenhofer. Speed dating despite jammers. In *Proc. DCOSS '09*, June 2009.

[40] Thomas Moscibroda, Rogert Wattenhofer, and Aaron Zollinger. Topology control meets sinr: the scheduling complexity of arbitrary topologies. In *Proceedings of the 7th ACM international symposium on Mobile ad hoc networking and computing*, MobiHoc '06, pages 310–321, New York, NY, USA, 2006. ACM.

[41] Koji Nakano and Stephan Olariu. Randomized leader election protocols in radio networks with no collision detection. In *ISAAC '00: Proceedings of the 11th International Conference on Algorithms and Computation*, pages 362–373, London, UK, 2000. Springer-Verlag.

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

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

[44] Koji Nikano and Stephan Olariu. Uniform leader election protocols for radio networks. *IEEE Trans. Parallel Distrib. Syst.*, 13(5):516–526, 2002.

[45] Guevara 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.

[46] George Nychis, Ranveer Chandra, Thomas Moscibroda, Ivan Tashev, and Peter Steenkiste. Reclaiming the white spaces: spectrum efficient coexistence with primary users. In *Proc. 7th Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2011.

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

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

[49] A. Richa, C. Scheideler, S. Schmid, and J. Zhang. A jamming-resistant mac protocol for multi-hop wireless networks. In *Proc. 24th International Symposium on Distributed Computing (DISC)*, 2010.

[50] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive and fair medium access despite reactive jamming. In *Proc. 31st IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2011.

[51] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin 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.

[52] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Towards jamming-resistant and competitive medium access in the sinr model. In *Proceedings of the 3rd ACM workshop on Wireless of the students, by the students, for the students*, S3 '11, pages 33–36, New York, NY, USA, 2011. ACM.

[53] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive and fair throughput for co-existing networks under adversarial interference. In *Proc. of PODC '12*, 2012.

[54] Agus Santoso, Yang Tang, Branka Vucetic, Abbas Jamalipour, and Yonghui Li. Interference cancellation in coexisting wireless local area networks. In *Proc. 10th IEEE Singapore International Conference on Communication Systems*, pages 1 –7, 2006.

[55] Christian Scheideler, Andrea Richa, and Paolo Santi. An $O(\log n)$ Dominating Set Protocol for Wireless Ad-Hoc Networks under the Physical Interference Model. In *Proc. of MOBIHOC*, 2008.

[56] 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.

[57] Georgios Smaragdakis, Ibrahim Matta, and Azer Bestavros. SEP: A Stable Election Protocol for clustered heterogeneous wireless sensor networks. In *Proc. 2nd International Workshop on Sensor and Actor Network Protocols and Applications (SANPA)*, 2004.

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

[59] P. M. van de Ven, A. J. E. M. Janssen, J. S. H. van Leeuwaarden, and S. C. Borst. Achieving target throughputs in random-access networks. *Perform. Eval.*, 68:1103–1117, November 2011.

[60] Dan E Willard. Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM J. Comput.*, 15(2):468–477, 1986.

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

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

[63] 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.

[64] 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.

[65] Shiang-Rung Ye, You-Chiun Wang, and Yu-Chee Tseng. A jamming-based MAC protocol for wireless multihop ad hoc networks. In *Proc. IEEE 58th Vehicular Technology Conference*, 2003.

[66] Shiang-Rung Ye, You-Chiun Wang, and Yu-Chee 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.

[67] J. Zander. Jamming in slotted ALOHA multihp packed radio networks. *IEEE Transactions on Networking*, 39(10):1525–1531, 1991.

[68] Gang Zhou, John A. Stankovic, and Sang H. Son. Crowded spectrum in wireless sensor networks. In *Proc. 3rd Workshop on Embedded Networked Sensors (EmNets)*, 2006.