Noname manuscript No.
(will be inserted by the editor)

Sade: Competitive MAC under Adversarial SINR

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

Abstract This paper considers the problem of how to efficiently share a wireless medium which is subject to harsh external interference or even jamming. So far, this problem is understood only in simplistic single-hop or unit disk graph models. We in this paper initiate the study of MAC protocols for the SINR interference model (a.k.a. physical model). This paper makes two contributions. First, we introduce a new adversarial SINR model which captures a wide range of interference phenomena. Concretely, we consider a powerful, adaptive adversary which can *jam* nodes at arbitrary times and which is only limited by some *energy budget*. Our second contribution is a distributed MAC protocol called *Sade* which provably achieves a constant competitive throughput in this environment: We show that, with high probability, the protocol ensures that a constant fraction of the non-blocked time periods is used for successful transmissions.

1 Introduction

The problem of coordinating the access to a shared medium is a central challenge in wireless networks. To efficiently share the wireless medium, a proper medium access control (MAC) protocol is needed. Ideally, such a protocol should not only be able to use the wireless medium as effectively as possible, but it should also be robust against a wide range of interference problems including jamming attacks. Currently, the most widely used model to capture interference problems is the SINR (signal-to-interference-and-noise ratio) model [24]. In this model, a message sent by node u is correctly received by node v if and only if

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

where $P_x(y)$ is the received power at node x of the signal transmitted by node y, $\mathcal{N}$ is the background noise, and S is the set of nodes $w \neq u$ that are transmitting at the same time as u. The threshold $\beta > 1$ depends on the desired rate, the modulation scheme, etc. When using the standard model for signal propagation, then this expression results in $P(u)/d(u,v)^\alpha/(\mathcal{N} + \sum_{w \in S} P(w)/d(w,v)^\alpha) \geq \beta$ where $P(x)$ is the strength of the signal transmitted by x, $d(x,y)$ is the Euclidean distance between x and y, and α is the path-loss exponent. In this paper, we will assume that all nodes transmit with some fixed signal strength P and that $\alpha > 2$, which is usually the case in an outdoors environment [38].

In most papers on MAC protocols, the background noise $\mathcal{N}$ is either ignored (i.e., $\mathcal{N} = 0$) or assumed to behave like a Gaussian variable. This, however, is an oversimplification of the real world. There are many sources of interference producing a non-Gaussian noise such as electrical devices, temporary obstacles, co-existing networks [42], or jamming attacks. Also, these sources can severely degrade the availability of the wireless medium which can put a significant stress on MAC protocols that have only been designed to handle interference from the nodes themselves. In order to capture a very broad range of noise phenomena, one of the main contributions of this work is the modeling of the background noise $\mathcal{N}$ (due to jamming or to environmental noise) with the aid of an adversary $\mathcal{ADV}(v)$ that has a fixed energy budget within a certain time frame for each

[1] Department of Computer Science, University of Paderborn, Germany; {adriano,scheideler}@upb.de
[2] Computer Science and Engineering, SCIDSE, Arizona State University, USA; {aricha,jzhang82}@asu.edu
[3] Aalborg University, Denmark; schmiste@cs.aau.dk

node v. More precisely, in our case, a message transmitted by a node u will be successfully received by node v if and only if

$$\frac{P/d(u,v)^\alpha}{\mathcal{ADV}(v) + \sum_{w \in S} P/d(w,v)^\alpha} \geq \beta \quad (1)$$

where $\mathcal{ADV}(v)$ is the current noise level created by the adversary at node v. Our goal will be to design a MAC protocol that allows the nodes to successfully transmit messages under this model as long as this is in principle possible. Prior to our work, no MAC protocol has been shown to have this property.

1.1 Model

We assume that we have a static set V of n wireless nodes that have arbitrary fixed positions in the 2-dimensional Euclidean plane so that no two nodes have the same position. The nodes communicate over a wireless medium with a single channel. We also assume that the nodes are backlogged in the sense that they always have something to broadcast. Each node sends at a fixed transmission power of P, and a message sent by u is correctly received by v if and only if

$$P/d(u,v)^\alpha / (\mathcal{ADV}(v) + \sum_{w \in S} P/d(w,v)^\alpha) \geq \beta$$

For our formal description and analysis, we assume a synchronized setting where time proceeds in synchronized *rounds*. In each round, a node u may either transmit a message or sense the channel, but it cannot do both. A node which is sensing the channel may either (i) sense an *idle* channel, (ii) sense a *busy* channel, or (iii) *receive* a packet. In order to distinguish between an idle and a busy channel, the nodes have built-in a fixed noise threshold ϑ: if the measured signal power exceeds ϑ, the channel is considered busy, otherwise idle. The threshold ϑ relates to the distance withing which nodes can communicate with each other, and we will derive bounds on the adversary in terms of ϑ (see below). A competitive throughput can only be achieved if ϑ is sufficiently large (see below). The basic strategy used in this paper is to approximate the "acceptable amount of noise" and set ϑ to the denominator of the SINR formula: the transmission range of a node v can accordingly be defined as the disk with center v and radius r such that $P/r^\alpha \geq \beta\vartheta$. Whether a message is indeed successfully received is determined by the SINR rule described above.

Physical carrier sensing is part of the 802.11 standard, and is provided by a Clear Channel Assessment (CCA) circuit. This circuit monitors the environment to determine when it is clear to transmit. The CCA functionality can be programmed to be a function of the Receive Signal Strength Indication (RSSI) and other parameters. The ability to manipulate the CCA rule allows the MAC layer to optimize the physical carrier sensing to its needs. Adaptive settings of the physical carrier sensing threshold have been used, for instance, in [43] to increase spatial reuse.

In addition to the nodes there is an *adversary* who controls the background noise. In order to cover a broad spectrum of noise phenomena, we allow this adversary to be adaptive, i.e., for each round t the adversary is allowed to know the state of all the nodes in the system at the beginning of t (i.e., before the nodes perform any actions at time t) and can set the noise level $\mathcal{ADV}(v)$ based on that for each node v. To leave some chance for the nodes to communicate, we restrict the adversary to be (B,T)-*bounded*: for each node v and time interval I of length T, a (B,T)-*bounded adversary* has an overall noise budget of $B \cdot T$ that it can use to increase the noise level at node v and that it can distribute among the rounds of I as it likes. This adversarial noise model is very general, since in addition to being adaptive, the adversary is allowed to make independent decisions on which nodes to jam at any point in time (provided that the adversary does not exceed its noise budget over a window of size T). In this way, many noise phenomena can be covered.

Our goal is to design a *symmetric local-control* MAC protocol (i.e., there is no central authority controlling the nodes, and all the nodes are executing the same protocol) that has a constant competitive throughput against any (B,T)-bounded adversary as long as certain conditions (on B etc.) are met. In order to define what we mean by "competitive", we need some notation. The *transmission range* of a node v is defined as the disk with center v and radius r with $P/r^\alpha \geq \beta\vartheta$, i.e., the range at which v's messages can still be received in the absence of adversarial noise. Given a constant $\epsilon > 0$, a round is called *potentially busy* at some node v if $\mathcal{ADV}(v) \geq (1-\epsilon)\vartheta$ (i.e., only a little bit of additional interference by the other nodes is needed so that v sees a busy channel). For a not potentially busy round at v, it is still possible that a message sent by a node u within v's transmission range is successfully received by v. Therefore, as long as the adversary is forced to offer not potentially busy rounds due to its limited budget and every node has a least one other node in its transmission range, it is in principle possible for the nodes to successfully transmit messages. To investigate that formally, we use the following notation. For any time frame F (a set of consecutive rounds) and node v, let $f_v(F)$ be the number of rounds in F that

are not potentially busy at v and let $s_v(F)$ be the number of rounds in which v successfully receives a message. We call a protocol *c-competitive* for some time frame F if $\sum_{v\in V} s_v(F) \geq c \sum_{v\in V} f_v(F)$. An adversary is *uniform* if at any round, $\mathcal{ADV}(v) = \mathcal{ADV}(w)$ for all nodes $v, w \in V$, which implies that $f_v(F) = f_w(F)$ for all nodes. Note that the scope of this paper is not restricted to the case of a uniform jammer (cf Theorem 1).

Since the MAC protocol presented in this paper will be randomized, our performance results typically hold *with high probability* (short: *w.h.p.*), that is, with a (polynomial) probability of at least $1-1/n$. Some intermediate results will only hold *with moderate probability (w.m.p.)*, that is, with a (poly-logarithmic) probability of at least $1 - 1/\log n$.

1.2 Our Contribution

The contribution of this paper is twofold. First of all, we introduce a novel extension of the SINR model in order to investigate MAC protocols that are robust against a broad range of interference phenomena. Second, we present a MAC protocol called *Sade*[1] which can achieve a c-competitive throughput where c only depends on ϵ and the path loss exponent α but not on the size of the network or other network parameters. (In practice, α is typically in the range $2 < \alpha < 5$, and thus c is a constant for fixed ϵ [38].) Let n be the number of nodes and let $N = \max\{n, T\}$. Concretely, we show:

Theorem 1 *When running* Sade *for at least* $\Omega((T \log N)/\epsilon + (\log N)^4/(\gamma\epsilon)^2)$ *rounds, for some* $\gamma \in O(1/(\log T + \log\log n))$, Sade *has a* $2^{-\mathrm{O}((1/\epsilon)^{2/(\alpha-2)})}$*-competitive throughput for any* $((1-\epsilon)\vartheta, T)$*-bounded adversary as long as (a) the adversary is uniform and the transmission range of every node contains at least one node, or (b) there are at least* $2/\epsilon$ *nodes within the transmission range of every node.*

On the other hand, we also show the following.

Theorem 2 *The nodes can be positioned so that the transmission range of every node is non-empty and yet no MAC protocol can achieve any throughput against a* (B, T)*-bounded adversary with* $B > \vartheta$*, even if it is uniform.*

The two theorems demonstrate that our *Sade* protocol is basically as robust as a MAC protocol can get within our model. However, it should be possible to improve the competitiveness. We conjecture that a polynomial dependency on $(1/\epsilon)$ is possible, but showing that formally seems to be hard. In fact, a different protocol than *Sade* would be needed for that.

While our main contribution lies on the conceptual and theoretical side, to complement our formal analysis and worst-case bounds, we also report on the results of a simple simulation study. This study confirms many of our theoretical results, but also shows that the actual performance is often better than in the worst-case. For instance, it depends to a lesser extent on ϵ.

1.3 Paper Organization

The remainder of this paper is organized as follows. We present our algorithm in Section 2, and subsequently analyze its performance in Section 3. Simulation results are presented in Section 4. After reviewing related work in Section 5, we conclude our paper with a discussion in Section 6.

2 Algorithm

The intuition behind *Sade* is simple: Each node v maintains a parameter p_v which specifies v's probability of accessing the channel at a given moment of time. That is, in each round, each node u decides to broadcast a message with probability p_v. (This is similar to classical 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 their p_v values over time in a multiplicative-increase multiplicative-decrease manner, i.e., the value is lowered in times when the channel is utilized (more specifically, we decrease p_v whenever a successful transmission occurs) or increased during times when the channel is idling. However, p_v will never exceed $\hat{p}$, for some constant $0 < \hat{p} < 1$ to be specified later.

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

With these intuitions in mind, we can describe *Sade* in full detail.

[1] Sade stands for SINR Jade, the SINR variant of the jamming defense protocol in [39].

Initially, every node v sets $T_v := 1$, $c_v := 1$, and $p_v := \hat{p}$. In order to distinguish between idle and busy rounds, each node uses a fixed noise threshold of ϑ.

The *Sade* protocol works in synchronized rounds. In every round, each node v decides with probability p_v to send a message. If it decides not to send a message, it checks the following two conditions:

- If v successfully receives a message, then $p_v := (1+\gamma)^{-1}p_v$.
- If v senses an idle channel (i.e., the total noise created by transmissions of other nodes and the adversary is less than ϑ), then $p_v := \min\{(1+\gamma)p_v, \hat{p}\}, T_v := \max\{1, T_v - 1\}$.

Afterwards, v sets $c_v := c_v + 1$. If $c_v > T_v$ then it does the following: v sets $c_v := 1$, and if there was no idle step among the past T_v rounds, then $p_v := (1+\gamma)^{-1}p_v$ and $T_v := T_v + 2$.

In order for *Sade* to be constant competitive in terms of throughput, the parameter γ needs to be a sufficiently small value that depends very loosely on n and T. Concretely, $\gamma \in O(1/(\log T + \log\log n))$.[2]

Our protocol *Sade* is an adaption of the Unit Disk Graphs specific MAC protocol described in [39], generalized to physical interference models. The main difference in the new protocol is that in order to use the concepts of idle and busy rounds, the nodes employ a fixed noise threshold ϑ to distinguish between idle (noise $< \vartheta$) and busy rounds (noise $\geq \vartheta$): in some scenarios the threshold may not be representative, in the sense that, since the success of a transmission depends on the noise at the receiving node and on β, it can happen that a node senses an idle or busy channel while *simultaneously* successfully receiving a message. In order to deal with this problem, *Sade* first checks whether a message is successfully received, and *only otherwise* takes into account whether a channel is idle or busy. Another change to the protocol in [39] is that we adapt T_v based on idle rounds which allows us to avoid the upper bound on T_v in the protocol in [39].

3 Analysis

While the MAC protocol *Sade* is very simple, its stochastic analysis requires an understanding of the intricate interplay of the nodes following their randomized protocol in a dependent manner. In particular, the nodes' interactions depend on their distances (the geometric setting). In order to study the throughput achieved by *Sade*, we will consider some fixed node $v \in V$ and will divide the area around v into three circular and concentric *zones*.

Let $D_R(v)$ denote the *disk* of radius R around a given node $v \in V$. In the following, we will sometimes think of $D_R(v)$ as the corresponding geometric area on the plane, but we will also denote by $D_R(v)$ the *set of nodes* located in this area. The exact meaning will be clear from the context. Moreover, whenever we omit R we will assume R_1 as radius, where R_1 is defined as in Definition 1.

Definition 1 (Zones) Given any node $v \in V$, our analysis considers three zones around v, henceforth referred to as Zone 1, Zone 2, and Zone 3: Zone 1 is the disk of radius R_1 around v, Zone 2 is the disk of radius R_2 around v *minus* Zone 1, and Zone 3 is the remaining part of the plane. Concretely:

1. Zone 1 covers the transmission range of v, i.e., its radius R_1 is chosen so that $P/R_1^\alpha \geq \beta\vartheta$, which implies that $R_1 = \sqrt[\alpha]{P/(\beta\vartheta)}$. Zone $Z_1(v)$ has the property that if there is at least one sender $u \in Z_1(v)$, then v will either successfully receive the message from u or sense a busy channel, and v will certainly receive the message from u if the overall interference caused by other nodes and the adversary is at most ϑ.
2. Zone 2 covers a range that we call the *(critical) interference range* of v. Its radius R_2 is chosen in a way so that if none of the nodes in Zone 1 and Zone 2 transmit a message, then the interference at any node $w \in Z_1(v)$ caused by transmitting nodes in Zone 3 is likely to be less than $\epsilon\vartheta$. Hence, if the current round is potentially non-busy at some $w \in Z_1(v)$ (i.e., $\mathcal{ADV}(w) \leq (1-\epsilon)\vartheta$), then the overall inference at w is less than ϑ, which means that w will see an idle round. It will turn out that R_2 can be chosen as $\mathrm{O}\big((1/\epsilon)^{1/(\alpha-1)}R_1\big)$.
3. Everything outside of Zone 2 is called *Zone 3*.

The key to proving a constant competitive throughput is the analysis of the cumulative probability (i.e., the sum of the individual sending probabilities p_v) of nodes in disks $D_1(v)$ and $D_2(v)$: We will show that the expected cumulative probabilities of $D_1(v)$ and $D_2(v)$, henceforth referred to by p_1 and p_2, are likely to be at most a constant. Moreover, our analysis shows that while the cumulative probability p_3 of the potentially infinitely large Zone 3 may certainly be unbounded (i.e., grow as a function of n), the aggregated power received

[2] While 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 without the need to update γ, estimating n in adversarial settings within any reasonable guarantees (e.g., constant or polynomial) is challenging. Whether we can eliminate any dependence on n and T for γ remains an open question.

at any node $w \in D_1(v)$ from all nodes in Zone 3 is also constant in expectation.

3.1 Zone 1

To show an upper bound on $p_1 = \sum_{u \in Z_1(v)} p_u$, i.e., the cumulative probability of the nodes in Zone 1 of v, we can follow a strategy similar to the one introduced for the Unit Disk Graph protocol [39].

In the following, we assume that the budget B of the adversary is limited by $(1-\epsilon')\vartheta$ for some constant $\epsilon' = 2\epsilon$. In this case, B is at most $(1-\epsilon)^2\vartheta$. We first look at a slightly weaker form of adversary. We say that a round t is *open* for a node v if v and at least one other node w within its transmission range are potentially non-busy, i.e., $\mathcal{ADV}(v) \leq (1-\epsilon)\vartheta$ and $\mathcal{ADV}(w) \leq (1-\epsilon)\vartheta$ (which also implies that v has at least one node within its transmission range). An adversary is *weak* (B,T)*-bounded* if it is (B,T)-bounded and in addition to this, at least a constant fraction of the potentially non-busy rounds at each node is open in every time interval of size T. The weak adversary is more constrained and hence indeed weaker.

Let us focus on a *time frame* I of size F consisting of $\delta \log N/\epsilon$ *subframes* I' of size $f = \delta[T + (\log^3 N)/(\gamma^2\epsilon)]$ each, i.e., $F = (\delta \log N/\epsilon)f$, where f is a multiple of T, δ is a sufficiently large constant, and $N = \max\{T, n\}$. Consider some fixed node v. We partition $Z_1(v)$ into six *sectors* of equal angles from v, $S_1, ..., S_6$. Note that for any sector S_i it holds that if a node $u \in S_i$ transmits a message, then its signal strength at any other node $u' \in S_i$ is at least $\beta\vartheta$. Fix a sector S and consider some fixed time frame F. Let us refer to the sum of the sending probabilities of the neighboring nodes of a given node $v \in S$ by $\bar{p}_v := \sum_{w \in S \setminus \{v\}} p_w$.

Our proof strategy to upper bound the sending probabilities in Zone 1 is as follows: First, in Lemma 1 we study any individual node's sending probability, and show that it decreases over time if it is high initially. Lemma 2 then studies cumulative probabilities, per *disk sector*, and shows that the cumulative probability is low at least one time during the interval. Lemmas 3 and 4 derive bounds on the cumulative probability throughout the interval. Finally, Lemma 5 shows that the number of "bad timeframes", in which cumulative probabilities are high, is small. Thus, since we have only 6 sectors per disk, it follows that during a significant amount of time, the cumulative probability is upper bounded in Zone 1.

Let us start. The following lemma shows that p_v will decrease dramatically if $\bar{p}_v$ is high throughout a certain time interval. The lemma is equivalent to Lemma 7 in [39], and we restate the proof for completeness.

Lemma 1 *Consider any node w in S. If $\bar{p}_w > 5 - \hat{p}$ during* all *rounds of a subframe I' of I and at the beginning of I', $T_w \leq \sqrt{F}$, then p_w will be at most $1/n^2$ at the end of I', w.h.p.*

Proof We call a round *useful* for node u if, from u's perspective, there is an idle channel or a successful transmission in that round (when ignoring the action of u); otherwise the round is called *non-useful*. Thus, the noise level during an idle useful round is a small constant (smaller than ϑ), however it can be very high in the presence of a successful transmission if two nodes are very close. Also note that in a non-useful round, according to our protocol, p_v will either decrease (if the threshold T_v is exceeded) or remain the same. On the other hand, in a *useful* round, p_v will increase (if v senses an idle channel), decrease (if v senses a successful transmission) or remain the same (if v sends a message). Hence, p_v can only increase during useful rounds of I. Let $\mathcal{U}$ be the set of useful rounds in I for our node v. We can distinguish between two cases, depending on the cardinality $|\mathcal{U}|$. In the following, let $p_v(0)$ denote the probability of v at the beginning of I (which is at most $\hat{p}$). Suppose that $f \geq 2[(3c \ln n)/\gamma]^2$ for a sufficiently large constant c.

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

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

Case 2: Next, suppose that $|\mathcal{U}| \geq (c \ln n)/\gamma$. We will show that many of these useful rounds will imply successful transmissions and consequently p_v decreases. Since due to our assumption, $p_v \leq \hat{p} \leq 1/24$ throughout I, it follows from the Chernoff bounds that w.h.p. v will sense the channel for at least a fraction of $2/3$ of the useful rounds w.h.p. Let this set of useful rounds be called $\mathcal{U}'$. Consider any round $t \in \mathcal{U}'$. Let q_0 be the probability that there is an idle channel at round t and let q_1 be the probability that there is a successful transmission at t. It holds that $q_0 + q_1 = 1$. From [39, Lemma 1] we also know that $q_1 \geq q_0 \cdot \bar{p}_v$. Since $\bar{p}_v > 5 - \hat{p}$ for all rounds in I, it follows that $q_1 \geq 4/5$ for every round in $\mathcal{U}'$. Thus, it follows from the Chernoff bounds that

for at least 2/3 of the rounds in $\mathcal{U}'$, v will sense a successful transmission w.h.p. Hence, at the end of I it holds w.h.p. that

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

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

Given this property of the individual probabilities, we can derive an upper bound for the cumulative probability of a Sector S. In order to compute $p_S = \sum_{v\in S} p_v$, we introduce three thresholds, a low one, $\rho_{green} = 5$, one in the middle, $\rho_{yellow} = 5e$, and a high one, $\rho_{red} = 5e^2$. The following three lemmas provide some important insights about these probabilities. The first lemma is shown in [39], and we restate it for completeness.

Lemma 2 *Consider any subframe I' in I. If at the beginning of I', $T_w \leq \sqrt{F}$ for all $w \in S$, then 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 1 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 3 *For any subframe I' in I 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. The probability bounds hold irrespective of the events outside of S.*

The proof is lengthy and technical. Thus, for readability reasons, it is deferred to the Appendix.

Lemma 4 *For any subframe I' in I 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., and the probability bound holds irrespective of the events outside of S.*

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 3 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 3 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 and Lemma 4 we can prove the following lemma (see also [39]).

Lemma 5 *For any sector S, the expected number of bad subframes I' in I is at most $1/polylog(N)$, and at most $\epsilon\beta'/6$ of the subframes I' in I are bad w.h.p., where the constant $\beta' > 0$ can be made arbitrarily small depending on the constant δ in f. The bounds hold irrespective of the events outside of S.*

Proof From Lemma 2 it follows that for every subframe I in F there is a time point $t \in I$ at which $p_D \leq \rho_{green}$ 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 $p_D \leq \rho_{green}$. Then it follows from Lemma 4 that for all rounds in I, $p_D \leq \rho_{red}$ w.m.p. (where the probability only depends on I and its preceding subframe), i.e., I is good. Hence, it follows from the Chernoff bounds that at most $\epsilon\beta/7$ of the odd-numbered as well as the even-numbered subframes after the first subframe in F are bad w.h.p. (if the constant α is sufficiently large). This implies that overall at most $\epsilon\beta/6$ of the subframes in F are bad w.h.p. □ □

Since we have exactly 6 sectors, it follows from Lemma 5 that apart from an $\epsilon\beta'$-fraction of the subframes, all subframes I' in I satisfy $\sum_{v\in Z_1(u)} p_v \leq 6\rho_{red}$ throughout I' w.h.p.

3.2 Zone 3

Next, we will show that although the cumulative probability of the nodes in Zone 3 may be high (for some distributions of nodes in the space it can actually be as high as $\Omega(n)$), their influence (or noise) at node v is limited if the radius of Zone 2 is sufficiently large. Thus, probabilities recover quickly in Zone 1 and there are many opportunities for successful receptions.

In order to bound the interference from Zone 3, we divide Zone 3 into two sub-zones: Z_3^-, which contains all nodes from Zone 3 up to a radius of $\mathrm{O}\big(\log^2 n\big)$, and Z_3^+, which contains all remaining nodes in Zone 3. For Zone Z_3^- we can prove the following lemma.

Lemma 6 *At most an $\epsilon\beta$-fraction of the subframes I' in I are bad for some R_1-disk in Zone Z_3^- w.h.p., where the constant $\beta' > 0$ can be made arbitrarily small depending on the constant δ in f.*

Proof The claim follows from the fact that the radius of Zone Z_3^- is $\mathrm{O}\big(\log^2 n\big)$ and hence $d = \mathrm{O}\big(\log^4 n\big)$

disks of radius R_1 are sufficient to cover the entire area of Z_3^-. According to Lemma 5, over all of these disks, the expected number of bad subframes is at most $1/polylog(N)$. Using similar techniques as for the proof of Lemma 5 in [39], it can also be shown that for each disk D, the probability for D to have k bad subframes is at most $1/polylog(N)^k$ irrespective of the events outside of D. Hence, one can use Chernoff bounds for sums of identically distributed geometric random variables to conclude that apart from an $\epsilon\beta'/d$-fraction of the subframes, all subframes I' in I satisfy $\sum_{v\in D} p_v \leq 6\rho_{red}$ throughout I' w.h.p. This directly implies the lemma. □

Suppose that $R_2 = c \cdot R_1$. Lemma 6 implies that in a good subframe the expected noise level at any node $w \in Z_1(v)$ created by transmissions in Zone Z_3^- is upper bounded by

$$6\rho_{red} \cdot \sum_{d=(c-1)}^{\mathrm{O}(\log^2 n)} \frac{2\pi(d+1)}{\sqrt{2}(dR_1)^\alpha} \leq \frac{12\pi\rho_{red}}{\alpha-1} \cdot \frac{1}{(c-2)^{\alpha-2}R_1^\alpha}$$

which is at most $\epsilon\vartheta/4$ if $c = O((1/\epsilon)^{1/(\alpha-2)})$ is sufficiently large. In order to bound the noise level at any node $w \in Z_1(v)$ from Zone Z_3^+, we prove the following claim.

Lemma 7 *Consider some fixed R_1-disk D. If at the beginning of time frame I, $T_w \leq \sqrt{F}$ for all $w \in D$, then for all rounds except for the first subframe in I, $p_D \in \mathrm{O}(\log n)$, w.h.p.*

Proof Lemma 2 implies that there must be a round t in the first subframe of I with $p_D \leq 6\rho_{green}$ w.h.p. Since for $p_D \in \Omega(\log n)$, at least a logarithmic number of nodes in D transmit and therefore every node sees a busy channel, w.h.p., and p_D can only increase if a node sees an idle channel, p_D is bounded by $\mathrm{O}(\log n)$ for the rest of I w.h.p. □

The claim immediately implies the following result.

Lemma 8 *If at the beginning of time frame I, $T_w \leq \sqrt{F}$ for all w, then for all rounds except for the first subframe in I, the interference at any node $w \in Z_1(v)$ due to transmissions in Z_3^+ is at most $\epsilon\vartheta/4$ w.h.p.*

Hence, we get:

Lemma 9 *If at the beginning of time frame I, $T_w \leq \sqrt{F}$ for all w, then at most an $\epsilon\beta$-fraction of the subframes in I contain rounds in which the expected interference at any node $w \in Z_1(v)$ due to transmissions in Zone 3 is at least $\epsilon\vartheta/2$.*

3.3 Zone 2

For Zone Z_2 we can prove the following lemma in the same way as Lemma 6.

Lemma 10 *At most an $\epsilon\beta$-fraction of the subframes I' in I are bad for some R_1-disk in Zone 2, w.h.p., where the constant $\beta > 0$ can be made arbitrarily small depending on the constant δ in f.*

3.4 Throughput

Given the upper bounds on the cumulative probabilities and interference, we are now ready to study the throughput of *Sade*. For this we first need to show an upper bound on T_v in order to avoid long periods of high p_v values. Let J be a time interval that has a quarter of the length of a time frame, i.e., $|J| = F/4$. We start with the following lemma whose proof is identical to Lemma III.6 in [40].

Lemma 11 *If in subframe I' the number of idle rounds at v is at most k, then node v increases T_v by 2 at most $k/2 + \sqrt{f}$ many times in I'.*

Next, we show the following lemma.

Lemma 12 *If at the beginning of J, $T_v \leq \sqrt{F}/2$ for all nodes v, then every node v has at least $2^{-\mathrm{O}((1/\epsilon)^{2/(\alpha-2)})}|J|$ rounds in J in which it senses an idle channel, w.h.p.*

Proof Fix some node v. Let us call a subframe I' in J *good* if in Zone 1 and in any R_1-disk in Zone 2 of v, the cumulative probability is upper bounded by a constant, and the expected interference due to transmissions at v induced from Zone 3 is at most $\epsilon\vartheta/2$ throughout I'. From Lemmas 5, 10, and 9 it follows that there is an $(1-\epsilon)$-fraction of good subframes in J. Since $R_2 = \mathrm{O}\big((1/\epsilon)^{1/(\alpha-2)}R_1\big)$, for any round t in a good subframe I' the total cumulative probability in Zones 1 and 2 of v is upper bounded by $\mathrm{O}\big((1/\epsilon)^{2/(\alpha-2)}\big)$. Hence, the probability that none of the nodes in Zones 1 and 2 of v transmits is given by

$$\sum_{w\in Z_1\cup Z_2}(1-p_w) \geq e^{-2\sum_{w\in Z_1\cup Z_2} p_w} = 2^{-\mathrm{O}\left((1/\epsilon)^{2/(\alpha-2)}\right)}$$

Due to the Markov inequality, the probability that the interference due to transmissions in Zone 3 is at least $\epsilon\vartheta$ is at most 1/2. These probability bounds hold independently of the other rounds in I'. Moreover, the total interference energy of the adversary in I' is bounded by $|I'|(1-\epsilon)^2\vartheta$, which implies that at most a $(1-\epsilon)$-fraction of the rounds in I' are potentially

busy, i.e., $\mathcal{ADV}(v) \geq (1-\epsilon)\vartheta$. Hence, for at least a $2^{-\mathrm{O}\left((1/\epsilon)^{2/(\alpha-2)}\right)}$-fraction of the rounds in I', the probability for v to sense an idle channel is a constant, which implies the lemma. □

This allows us to prove the following lemma.

Lemma 13 *If at the beginning of J, $T_v \leq \sqrt{F}/2$ for all v, then also $T_v \leq \sqrt{F}/2$ for all v at the end of J, w.h.p.*

Proof From the previous lemma we know that every node v senses an idle channel for $\Omega(|J|)$ rounds in J for any constants $\epsilon > 0$ and $\alpha > 2$. T_v is maximized at the end of J if all of these idle rounds happen at the beginning of J, which would eventually reduce T_v down to 1. Afterwards, T_v can rise to a value of at most t for the maximum t with $\sum_{i=1}^{t} 2i \leq |J|$ (because v increases T_v by 2 each time it sees no idle channel in the previous T_v steps), which is at most $\sqrt{|J|}$. Since $\sqrt{|J|} = \sqrt{|F|}/2$, the lemma follows. □

Since T_v can be increased at most $(F/4)\sqrt{F}/2$ many times in J, we get:

Lemma 14 *If at the beginning of a time frame I, $T_v \leq \sqrt{F}/2$ for all v, then throughout I, $T_v \leq \sqrt{F}$ for all v, and at the end of I, $T_v \leq \sqrt{F}/2$ for all v, w.h.p.*

Hence, the upper bounds on T_v that we assumed earlier are valid w.h.p. We are now ready to prove the following result which essentially implies Theorem 1 *a)*, as we will see below:

Theorem 3 *When running* Sade *for at least $\Omega((T\log N)/\epsilon' + (\log N)^4/(\gamma\epsilon')^2)$ rounds,* Sade *has a $2^{-\mathrm{O}\left((1/\epsilon')^{2/(\alpha-2)}\right)}$-competitive throughput for any weak $((1-\epsilon')\vartheta, T)$-bounded adversary.*

Recall that we allow here any $((1-\epsilon)\vartheta, T)$-bounded adversary.

Proof (of Theorem 3) Recall that a round is *open* for a node v if v and at least one other node in $Z_1(v)$ are not potentially busy. Let J be the set of all open rounds in I. Furthermore, let k_0 be the number of times v senses an idle channel in J and let k_1 be the number of times v receives a message in I. From Lemma 12 and the assumptions in Theorem 3 we know that $k_0 = 2^{-\mathrm{O}\left((1/\epsilon)^{2/(\alpha-2)}\right)}|I|$.

Case 1: $k_1 \geq k_0/6$. In this case, our protocol is $2^{-\mathrm{O}\left((1/\epsilon)^{2/(\alpha-2)}\right)}$-competitive for v and we are done.

Case 2: $k_1 < k_0/6$. In this case, we know from Lemma 11 that p_v is decreased at most $k_0/2+\sqrt{F}$ times in I due to $c_u > T_u$. In addition to this, p_v is decreased at most k_1 times in I due to a received message. On the other hand, p_v is increased at least k_0 times in J (if possible) due to an idle channel w.h.p. Also, we know from our protocol that at the beginning of I, $p_v = \hat{p}$. Hence, there must be at least $(1-1/2-1/6)k_0 - \sqrt{|F|} \geq k_0/4$ rounds in J w.h.p. at which $p_v = \hat{p}$. Now, recall the definition of a good subframe in the proof of Lemma 12. From Lemmas 5, 10, and 9 it follows that at most a $\epsilon\beta$-fraction of the subframes in I is bad. In the worst case, all of the rounds in these subframes are open rounds, which sums up to at most $k_0/8$ if β is sufficiently small. Hence, there are at least $k_0/8$ rounds in J that are in good subframes, w.h.p., and at which $p_v = \hat{p}$, which implies that the other not potentially busy node in $Z_1(v)$ has a constant probability of receiving a message from v. Using Chernoff bounds, at least $k_0/16$ rounds with successfully received transmissions can be identified for v, 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. It follows that our protocol is $2^{-\mathrm{O}\left((1/\epsilon)^{2/(\alpha-2)}\right)}$-competitive in F. □

We now complete the proof of Theorem 1.

Proof (of Theorem 1)

Case 1: the adversary is uniform and $\forall v : Z_1(v) \neq \emptyset$.

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 weak $((1-\epsilon)\vartheta, T)$-bounded adversary are satisfied. Therefore Theorem 3 applies, which completes the proof of Theorem 1 a).

Case 2: $\boldsymbol{|Z_1(v)| \geq 2/\epsilon}$ *for all* $\boldsymbol{v \in V}$.

Consider some fixed time interval I with $|I|$ being a multiple of T. For every node $v \in Z_1(u)$ let f_v be the number of non-jammed rounds at v in I and o_v be the number of open rounds at v in I. Let J be the set of rounds in I with at most one non-jammed node. Suppose that $|J| > (1-\epsilon/2)|I|$. Then every node in $Z_1(u)$ must have more than $(\epsilon/2)|I|$ of its non-jammed rounds in J. Note that these non-jammed rounds must be "serialized" (i.e., disjoint) in J to satisfy our requirement on J: if J_v is the set of non-jammed slots for $v \in Z_1(u)$ in J, there is no intersection between J_x and J_y for two nodes in $Z_1(u)$. Therefore, it holds that $|J| > \sum_{v \in Z_1(u)} (\epsilon/2)|I| \geq (2/\epsilon) \cdot (\epsilon/2)|I| = |I|$. Since this is impossible, it must hold that $|J| \leq (1-\epsilon/2)|I|$.

Thus, $\sum_{v\in Z_1(u)} o_v \geq (\sum_{v\in Z_1(u)} f_v) - |J| \geq (1/2)\sum_{v\in Z_1(u)} f_v$ because $\sum_{v\in Z_1(u)} f_v \geq (2/\epsilon)\cdot\epsilon|I| = 2|I|$. Let $D'(u)$ be the set of nodes $v \in Z_1(u)$ with $o_v \geq f_v/4$. That is, for each of these nodes, a constant fraction of the non-jammed rounds is open. Then $\sum_{v\in Z_1(u)\setminus D'(u)} o_v < (1/4)\sum_{v\in Z_1(u)} f_v$, so $\sum_{v\in D'(u)} o_v \geq (1/2)\sum_{v\in Z_1(u)} o_v \geq (1/4)\sum_{v\in Z_1(u)} f_v$.

Consider now a set $U \subseteq V$ of nodes so that $\bigcup_{u\in U} Z_1(u) = V$ and for every $v \in V$ there are at most 6 nodes $u \in U$ with $v \in Z_1(u)$. Note U is easy to construct in a greedy fashion for arbitrary UDGs, and therefore for $Z_1(u)$ in the SINR model, and also known as a *dominating set of constant density*. Let $V' = \bigcup_{u\in U} D'(u)$. Since $\sum_{v\in D'(u)} o_v \geq (1/4)\sum_{v\in Z_1(u)} f_v$ for every node $u \in U$, it follows that $\sum_{v\in V'} o_v \geq (1/6)\sum_{u\in U}\sum_{v\in D'(u)} o_v \geq (1/24)\sum_{u\in U}\sum_{v\in Z_1(u)} f_v \geq (1/24)\sum_{v\in V} f_v$. The fact that a constant fraction of rounds are open implies that, using Theorem 3, SADE is constant competitive w.r.t. the nodes in V'. This completes the proof of Theorem 1 b). □

3.5 Optimality

Obviously, if a jammer has a sufficiently high energy budget, it can essentially block all nodes all the time. In the following we call a network *dense* if $\forall v \in V : Z_1(v) \geq 1$.

Theorem 4 *The nodes can be positioned so that the transmission range of every node is non-empty and yet no MAC protocol can achieve any throughput against a (B,T)-bounded adversary with $B > \vartheta$, even if it is uniform.*

Proof Let us suppose that the jammer has an energy budget $B > \vartheta$. If every node v only has nodes right at the border of its disk $Z_1(v)$ and the adversary continuously sets $\mathcal{ADV}(v) = B$, then v will not be able to receive any messages according to the SINR model. Thus, the overall throughput in the system is 0. □

4 Simulations

To complement our formal analysis and to investigate the average-case behavior of our protocol, we conduct a simulation study. In the following, we consider two scenarios which differ in the way nodes are distributed in the 2-dimensional Euclidean space. In the first scenario, called UNI, the nodes are distributed *uniformly at random* in the 2-dimensional plane of size 25×25 units. In the second scenario, called HET, we first subdivide the 2-dimensional plane of size 25×25 units into 25 *sub-squares* of size 5×5 units. For each sub-square, we then choose the number of nodes λ uniformly at random from the interval $[20, 1000]$ and distribute said nodes (uniformly at random) in the corresponding sub-square. Consequently, each sub-square may provide a different density, where the attribute density represents the average number of nodes per unit in the plane of the corresponding scenario. In order to avoid boundary effects, for both UNI and HET, we assume that the Euclidean plane "wraps around", i.e., is a torus where distances are computed modulo the boundaries.

While our formal throughput results in Section 3 hold for *any* adversary which respects the jamming budget constraints, computing the best adversarial strategy (i.e., the strategy which minimizes the throughput of *Sade*) is difficult. Hence, in our simulations, we consider the following two types of adversaries:

1. *Regular (or random) jammer (*REG*):* given an energy budget B per node, a time interval T, and a specific $1 > \epsilon > 0$, the adversary randomly jams each node every ϵth round (on average) using exactly $\frac{B}{\epsilon}$ energy per node. Additionally we make sure that the overall budget B is perfectly used up by the end of T.
2. *Bursty (or deterministic) jammer (*BUR*):* For each time period T, the adversary jams *all* initial rounds at the node, until the budget B is used up. The remaining rounds in T are unjammed. In other words, the first ϵT many rounds are jammed by the adversary using exactly $\frac{B}{\epsilon}$ energy per node.

If not stated otherwise, we consider the jammer REG with parameters $\alpha = 3, \epsilon = \frac{1}{3}$, $\beta = 2$, $\Pi = 8$, $T = 60$, $B = (1-\epsilon)\cdot\beta$ and run the experiment for 3000 rounds. We will typically plot the percentage of successful message receptions, averaged over all nodes, with respect to the *unjammed rounds*. If not specified otherwise, we repeat each experiment ten times with different random seeds, both for the distribution of nodes in the plane as well as the decisions made by our MAC protocol. By default, our results show the *average* over these runs; the variance of the runs is low.

We employ discrete-event simulations and keep the simulator very simple: the simulator corresponds to our theoretical model investigated in the previous section. In particular, we stick to synchronous rounds and do not model the TCP/IP networking stack. For reproducibility reasons, we will make our simulation code publicly available together with this paper.

Impact of Scale and α. We first study the throughput as a function of the network size. Therefore

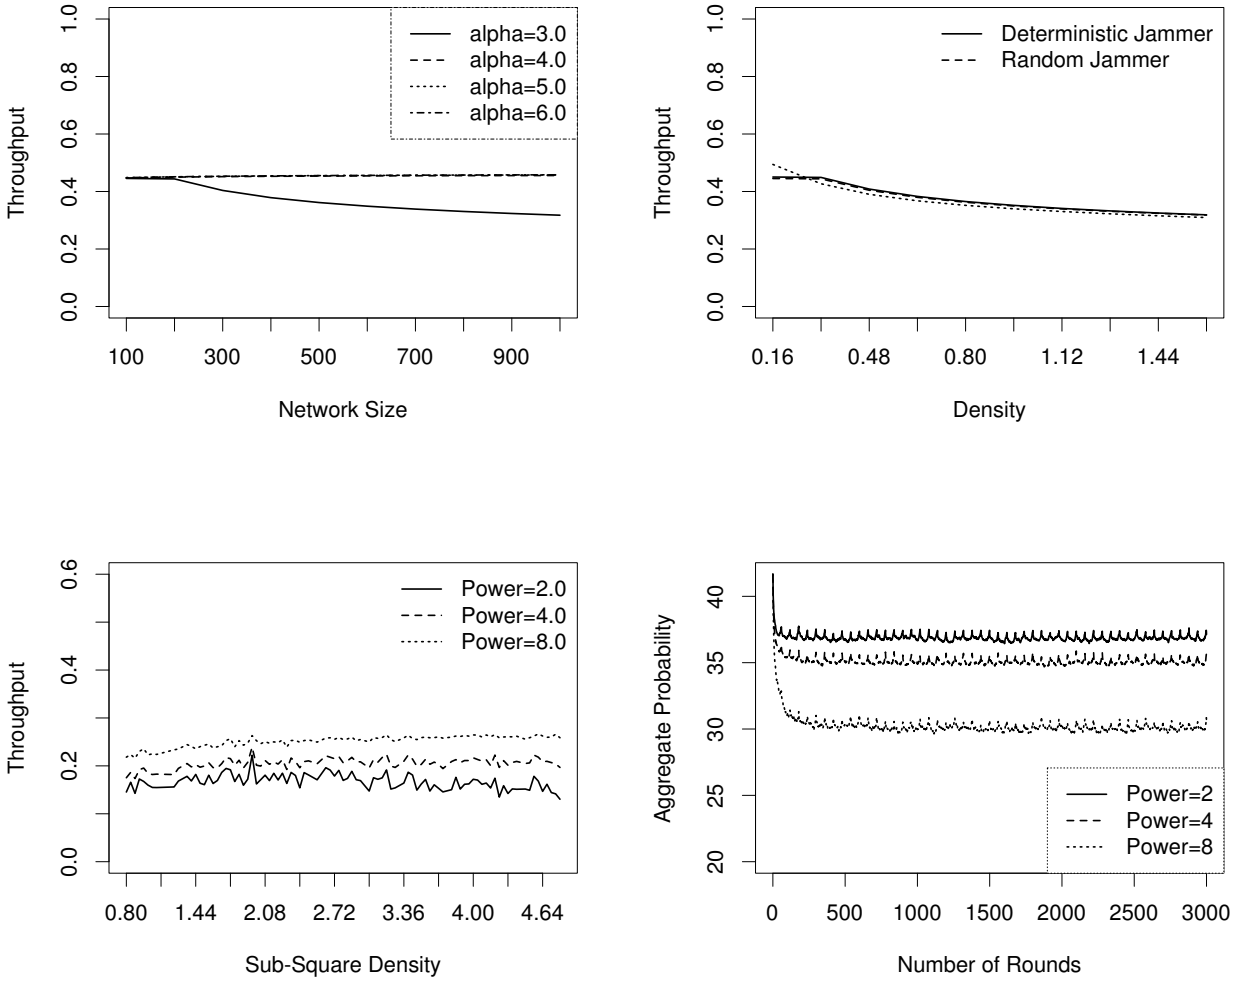

Fig. 1 Throughput as a function of network size (*top left*), density (*top right*), sub-grid density (*bottom left*), and power (*bottom right*).

we distribute n nodes uniformly in the $\sqrt{n} \times \sqrt{n}$ plane. Figure 1 (*top left*) shows our results under the Reg (or *random*) jammer and different α values. First, we can see that the competitive throughput is around 40%, which is higher than what we expect from our worst-case formal analysis. Interestingly, for $\alpha = 3$, we observe a small throughput decrease for larger networks; but for $\alpha > 3$, the throughput is almost independent of the network scale. (In the literature, α is typically set to 3 or 4.)

This is reminiscent of Theorem 1: a higher α renders the transmissions and power propagation more local. This locality can be exploited by *Sade*.

Impact of Density. Next, we investigate how the performance of *Sade* depends on the node density. We focus on $\alpha = 3$ and study both the Reg jammer as well as the Bur (deterministic) jammer. Figure 1 (*top right*) shows that results for the Uni scenario (n nodes distributed uniformly in the 25×25 plane, i.e., density $n/625$). The throughput is similar under both jammers, and slightly declines for denser networks. This effect is very similar to the effect of having larger (but equidistant) networks.

However, *Sade* suffers from more heterogenous densities. The results for the scenario Het are shown in Figure 1 (*bottom left*). While the throughput is generally lower, the specific sub-square density plays a minor role.

Convergence Time. *Sade* adapts quite fast to the given setting, as the nodes increase and decrease their

sending probabilities in a multiplicative manner. Being able to adapt quickly is an important feature, in particular in dynamic or mobile environments where nodes can join and leave over time, or where nodes are initialized with too high or low sending probabilities. Our distributed MAC protocol will adjust automatically.

Figure 1 (*bottom right*) shows representative executions over time and plots the cumulative probability. Initially, nodes have a maximum sending probability of $\hat{p} = 1/24$. This will initially lead to many collisions; however, very quickly, the senders back off and the overall sending probabilities (the *aggregated probability*) reduce almost exponentially, and we start observing successful message transmissions. (Observe that the aggregated "probability" can be higher than one, as it is simply the sum of the probabilities of the individual nodes.)

The sum of all sending probabilities also converges quickly for any other Π. However, for smaller powers, the overall probability is higher. This is in accordance with our goal: since for very large sending powers, also more remote nodes in the network will influence each other and interfere, it is important that there be only a small number of concurrent senders in the network at any time—the aggregated sending probability must be small. On the other hand, small powers allow for more local transmissions, and to achieve a high overall throughput, many senders should be active at the same time—the overall sending probability should be high.

5 Related Work

Traditional jamming defense mechanisms typically operate on the physical layer [32,34,44], and mechanisms have been designed to both *avoid* jamming as well as *detect* jamming. Especially spread spectrum technology is very effective to avoid jamming, as with widely spread signals, it becomes harder to detect the start of a packet quickly enough in order to jam it. Unfortunately, protocols such as IEEE 802.11b use relatively narrow spreading [26], and some other IEEE 802.11 variants spread signals by even smaller factors [8]. Therefore, a jammer that simultaneously blocks a small number of frequencies renders spread spectrum techniques useless in this case. As jamming strategies can come in many different flavors, detecting jamming activities by simple methods based on signal strength, carrier sensing, or packet delivery ratios has turned out to be quite difficult [31].

Recent work has investigated *MAC layer strategies* against jamming in more detail, for example coding strategies [9], channel surfing and spatial retreat [2,47], or mechanisms to hide messages from a jammer, evade its search, and reduce the impact of corrupted messages [46]. Unfortunately, these methods do not help against an adaptive jammer with *full* information about the history of the protocol, like the one considered in our work.

In the theory community, work on MAC protocols has mostly focused on efficiency. Many of these protocols are random backoff or tournament-based protocols [5,10,23,25,30,37] that do not take jamming activity into account and, in fact, are not robust against it (see [3] for more details). The same also holds for many MAC protocols that have been designed in the context of broadcasting [11] and clustering [29]. Also some work on jamming is known (e.g., [13] for a short overview). There are two basic approaches in the literature. The first assumes randomly corrupted messages (e.g. [36]), which is much easier to handle than adaptive adversarial jamming [4]. 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. [20, 28]) or bounds the number of channels the adversary can jam (e.g. [14–19,33]).

The protocols in [20,28] can tackle adversarial jamming at both the MAC and network layers, where the adversary may not only be jamming the channel but also introducing malicious (fake) messages (possibly with address spoofing). However, they depend on the fact that the adversarial jamming budget is finite, so it is not clear whether the protocols would work under heavy continuous jamming. (The result in [20] seems to imply that a jamming rate of 1/2 is the limit whereas the handshaking mechanisms in [28] seem to require 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 [14–19,33], 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 [15,18] and at least quadratic in the number of jammed channels for randomized protocols [16, 33] if the adversary can jam almost all channels at a time). Recent work [14] 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. [19] study robust information exchange in single-hop networks.

Our work is motivated by the work in [4] and [3]. In [4] 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. [3] initiated the

study of throughput-competitive MAC protocols under continuously running, adaptive jammers, but they only consider single-hop wireless networks. Their approach has later been extended to reactive jamming environments [40], co-existing networks [42] and applications such as leader election [41].

Several research groups have recently investigated similar models in different contexts [1,6,7,12,22,21,27, 45], e.g., in Byzantine and Sybil environments [1,7,22], in multi-channel environments [45], and in learning environments [12]: Dams et al. [12] introduced distributed algorithms based on no-regret learning which approximate the number of successful transmissions well. In contrast to our work, the authors do not need to make any assumptions on the number of nodes n. However, the paper does not provide any bounds regarding convergence time. Bender et al. [12] recently initiated a systematic study of "scalable" backoff protocols (assuming dynamic packet arrivals), identifying three natural properties: constant throughput, few failed access attempts, robustness (continue to work efficiently even if some of the access attempts fail for spurious reasons). The authors present a the RE-BACKOFF protocol guaranteeing expected constant throughput with dynamic process arrivals and requiring only an expected polylogarithmic number of access attempts per process.

The result closest to ours is the robust MAC protocol for Unit Disk Graphs presented in [39]. In contrast to [39], we initiate the study of the more relevant and realistic *physical interference model* [24] and show that a competitive throughput can still be achieved. Indeed, to the best of our knowledge, our paper is the first to consider jamming resistant protocols in the SINR model. As unlike in Unit Disk Graphs, in the SINR setting far-away communication can potentially interfere and there is no absolute notion of an idle medium, a new protocol is needed whose geometric properties must be understood. For the SINR setting, we also introduce a new adversarial model (namely the *energy budget adversary*).

Bibliographic Note. A first version of this paper was presented at the IEEE INFOCOM 2014 conference [35].

6 Conclusion

This paper has shown that robust MAC protocols achieving a constant competitive throughput exist even in the physical model. This concludes a series of research works in this area. Nevertheless, several interesting questions remain open. For example, while our theorems prove that *Sade* is as robust as a MAC protocol can get within our model and for constant ϵ, we conjecture that a throughput which is polynomial in $(1/\epsilon)$ is possible. However, we believe that such a claim is very difficult to prove. We also plan to explore the performance of *Sade* under specific node mobility patterns.

Remark: In order for the community to be able to reproduce our results, we will make the simulation code publicly available together with this paper.

Acknowledgments. The authors would like to thank Michael Meier from the University of Paderborn for his help with the evaluation of the protocol. This research is partly supported by the Danish Villum project *ReNet*.

References

1. D. Alistarh, S. Gilbert, R. Guerraoui, Z. Milosevic, and C. Newport. Securing your every bit: Reliable broadcast in byzantine wireless networks. In *Proc. Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 50–59, 2010.
2. G. Alnifie and R. Simon. A multi-channel defense against jamming attacks in wireless sensor networks. In *Proc. of Q2SWinet*, pages 95–104, 2007.
3. B. Awerbuch, A. Richa, and C. Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *Proc. ACM PODC*, 2008.
4. E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa. On the performance of IEEE 802.11 under jamming. In *Proc. IEEE INFOCOM*, pages 1265–1273, 2008.
5. M. A. Bender, M. Farach-Colton, S. He, B. C. Kuszmaul, and C. E. Leiserson. Adversarial contention resolution for simple channels. In *Proc. ACM SPAA*, 2005.
6. M. A. Bender, J. T. Fineman, S. Gilbert, and M. Young. How to scale exponential backoff: Constant throughput, polylog access attempts, and robustness. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 636–654, 2016.
7. M. Bertier, A.-M. Kermarrec, and G. Tan. Message-efficient byzantine fault-tolerant broadcast in a multi-hop wireless sensor network. In *Proc. IEEE 30th International Conference on Distributed Computing Systems (ICDCS)*, pages 408–417, 2010.
8. T. Brown, J. James, and A. Sethi. Jamming and sensing of encrypted wireless ad hoc networks. In *Proc. ACM International Symposium on Mobile Ad hoc Networking and Computing (MOBIHOC)*, pages 120–130, 2006.
9. J. Chiang and Y.-C. Hu. Cross-layer jamming detection and mitigation in wireless broadcast networks. In *Proc. MOBICOM*, pages 346–349, 2007.
10. B. S. Chlebus, D. R. Kowalski, and M. A. Rokicki. Adversarial queuing on the multiple-access channel. In *Proc. ACM PODC*, 2006.
11. A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60(2):115 – 143, 2006.
12. J. Dams, M. Hoefer, and T. Kesselheim. Jamming-resistant learning in wireless networks. In *Proc. International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 447–458, 2014.
13. 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.
14. S. Dolev, S. Gilbert, R. Guerraoui, F. Kuhn, and C. C. Newport. The wireless synchronization problem. In *Proc. 28th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 190–199, 2009.
15. 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.
16. 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.

17. S. Dolev, S. Gilbert, R. Guerraoui, and C. C. Newport. Gossiping in a multi-channel radio network. In *Proc. 21st International Symposium on Distributed Computing (DISC)*, pages 208–222, 2007.
18. 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.
19. S. Gilbert, R. Guerraoui, D. R. Kowalski, and C. C. Newport. Interference-resilient information exchange. In *Proc. 28th IEEE International Conference on Computer Communications (INFOCOM)*, pages 2249–2257, 2009.
20. S. Gilbert, R. Guerraoui, and C. Newport. Of malicious motes and suspicious sensors: On the efficiency of malicious interference in wireless networks. In *Proc. OPODIS*, 2006.
21. S. Gilbert, V. King, S. Pettie, E. Porat, J. Saia, and M. Young. (near) optimal resource-competitive broadcast with jamming. In *Proc. ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 257–266, 2014.
22. S. L. Gilbert and C. Zheng. Sybilcast: Broadcast on the open airwaves. In *Proc. ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 130–139, 2013.
23. L. A. Goldberg, P. D. Mackenzie, M. Paterson, and A. Srinivasan. Contention resolution with constant expected delay. *J. ACM*, 47(6), 2000.
24. P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388 –404, 2000.
25. J. Hastad, T. Leighton, and B. Rogoff. Analysis of backoff protocols for mulitiple accesschannels. *SIAM Journal on Computing*, 25(4), 1996.
26. IEEE. Medium access control (MAC) and physical specifications. In *IEEE P802.11/D10*, 1999.
27. V. King, J. Saia, and M. Young. Conflict on a communication channel. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*, pages 277–286, 2011.
28. C. Koo, V. Bhandari, J. Katz, and N. Vaidya. Reliable broadcast in radio networks: The bounded collision case. In *Proc. ACM PODC*, 2006.
29. F. Kuhn, T. Moscibroda, and R. Wattenhofer. Radio network clustering from scratch. In *Proc. ESA*, 2004.
30. B.-J. Kwak, N.-O. Song, and L. E. Miller. Performance analysis of exponential backoff. *IEEE/ACM Transactions on Networking*, 13(2):343–355, 2005.
31. M. Li, I. Koutsopoulos, and R. Poovendran. Optimal jamming attacks and network defense policies in wireless sensor networks. In *Proc. IEEE INFOCOM*, pages 1307–1315, 2007.
32. X. Liu, G. Noubir, R. Sundaram, and S. Tan. Spread: Foiling smart jammers using multi-layer agility. In *Proc. IEEE INFOCOM*, pages 2536–2540, 2007.
33. D. Meier, Y. A. Pignolet, S. Schmid, and R. Wattenhofer. Speed dating despite jammers. In *Proc. DCOSS*, June 2009.
34. V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein. Using channel hopping to increase 802.11 resilience to jamming attacks. In *Proc. IEEE INFOCOM*, pages 2526–2530, 2007.
35. A. Ogierman, A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Competitive mac under adversarial sinr. In *Proc. IEEE INFOCOM*, April 2014.
36. A. Pelc and D. Peleg. Feasibility and complexity of broadcasting with random transmission failures. In *Proc. ACM PODC*, 2005.
37. P. Raghavan and E. Upfal. Stochastic contention resolution with short delays. *SIAM Journal on Computing*, 28(2):709–719, 1999.
38. T. Rappaport. *Wireless communications*. Prentice Hall PTR Upper Saddle River, 2002.
39. A. Richa, C. Scheideler, S. Schmid, and J. Zhang. A Jamming-Resistant MAC Protocol for Multi-Hop Wireless Networks. In *Proc. DISC*, 2010.
40. A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Competitive and fair medium access despite reactive jamming. In *Proc. 31st International Conference on Distributed Computing Systems (ICDCS)*, 2011.
41. A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Self-stabilizing leader election for single-hop wireless networks despite jamming. In *Proc. 12th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2011.
42. A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Competitive and fair throughput for co-existing networks under adversarial interference. In *Proc. 31st Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 2012.
43. C. Scheideler, A. Richa, and P. Santi. An $O(\log n)$ Dominating Set Protocol for Wireless Ad-Hoc Networks under the Physical Interference Model. In *Proc. ACM International Symposium on Mobile Ad hoc Networking and Computing (MOBIHOC)*, 2008.
44. M. K. Simon, J. K. Omura, R. A. Schultz, and B. K. Levin. *Spread Spectrum Communications Handbook*. McGraw-Hill, 2001.
45. H. Tan, C. Wacek, C. Newport, and M. Sherr. A disruption-resistant mac layer for multichannel wireless networks. In *Proc. International Conference on Principles of Distributed Systems (OPODIS)*, pages 202–216, 2014.
46. A. Wood, J. Stankovic, and G. Zhou. DEEJAM: Defeating energy-efficient jamming in IEEE 802.15.4-based wireless networks. In *Proc. SECON*, 2007.
47. 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.

A Proof of Lemma 3

We prove the lemma for the case that initially $p_S \leq \rho_{green}$; 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 a total interference of less than ϑ in round t when ignoring the nodes in S. Similarly, let $p_S^{(1)}$ denote the cumulative probability of the nodes $w \in S$ with a single transmitting node in $Z_1(w) \setminus S$ and additionally an interference of less than ϑ in round t, and let $p_S^{(2)}$ be the cumulative probability of the nodes $w \in S$ that do not satisfy the first two cases (which implies that they will not experience an idle channel, no matter what the nodes in S will do). 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, let $q_1(S)$ be the probability that exactly one node in S is transmitting, and let $q_2(S) = 1 - q_0(S) - q_1(S)$ be the probability that at least two nodes in S are transmitting.

First, let us simplify our setting slightly and ignore the case that $c_v > T_v$ for a node $v \in S$ at round t. By examining the 9 different cases, we obtain the following result:

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

To give an example (the other cases are similar), we consider $q_0(S)$ and $p_S^{(1)}$, i.e., all nodes in S are silent and for all nodes in $w \in S$ accounted for in $p_S^{(1)}$ there is exactly one transmitting node in $Z_1(w) \setminus S$ and the remaining interference is less than ϑ. In this case, w is guaranteed to receive a message, so according to the *Sade* protocol, it lowers p_w by $(1+\gamma)$.

The upper bound on $\mathbb{E}[p'_S]$ certainly also holds 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 $\neg E_2$. We distinguish between three cases.

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

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

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.

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

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

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

$$((1/a)x_0 + (1/b)x_1 + (1/c)x_2)^2 (a^2 x_0 + b^2 x_1 + c^2 x_2) \leq 1 \tag{3}$$

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

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

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

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

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

Combining the claims, Equation (3) 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) 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 5 (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,

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

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

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

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

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

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