(Near) Optimal Resource-Competitive Broadcast with Jamming

[Extended Abstract]

Seth Gilbert*
Dept. of Computer Science
National Univ. of Singapore
Singapore
seth.gilbert@comp.nus.edu.sg

Valerie King
Dept. of Computer Science
University of Victoria
Victoria, BC, Canada
val@cs.uvic.ca

Seth Pettie †
Dept. of Electrical Engineering
and Computer Science
University of Michigan
Ann Arbor, MI, USA
pettie@umich.edu

Ely Porat
Dept. of Computer Science
Bar-Ilan University
Ramat Gan, Israel
porately@cs.biu.ac.il

Jared Saia‡
Dept. of Computer Science
University of New Mexico
Albuquerque, NM, USA
saia@cs.unm.edu

Maxwell Young‡
Dept. of Computing
Drexel University
Philadelphia, PA, USA
myoung@cs.drexel.edu

ABSTRACT

We consider the problem of broadcasting a message from a sender to $n \geq 1$ receivers in a time-slotted, single-hop, wireless network with a single communication channel. Sending and listening dominate the energy usage of small wireless devices and this is abstracted as a unit cost per time slot. A jamming adversary exists who can disrupt the channel at unit cost per time slot, and aims to prevent the transmission of the message. Let T be the number of slots jammed by the adversary. Our goal is to design algorithms whose cost is *resource-competitive*, that is, whose per-device cost is a function, preferably $o(T)$, of the adversary's cost. Devices must work with limited knowledge. The values n, T, and the adversary's jamming strategy are unknown.

For 1-to-1 communication, we provide an algorithm with an expected cost of $O(\sqrt{T\ln(1/\epsilon)} + \ln(1/\epsilon))$, which succeeds with probability at least $1-\epsilon$ for any tunable parameter $\epsilon > 0$. For 1-to-n broadcast, we provide a very different algorithm that succeeds with high probability and yields an expected cost per device of $O(\sqrt{T/n}\log^4 T + \log^6 n)$. Therefore, the bigger the system, the better advantage achieved over the adversary!

We complement our upper bounds with *tight or nearly tight lower bounds*. We prove that any 1-to-1 communication algorithm with constant probability of success has expected cost $\Omega(\sqrt{T})$. For 1-to-n broadcast we show that some node has cost $\Omega(\sqrt{T/n})$. Finally, we consider a more powerful adversary that can spoof messages from the receiver, rather than just jam the channel. We prove that any 1-to-1 communication algorithm in this model has expected cost $\Omega(T^{\varphi-1})$, where $\varphi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. This matches an earlier upper bound of King, Saia, and Young.

Categories and Subject Descriptors

C.2.1 [**Computer Communications Networks**]: Network Architecture and Design—*Wireless communication*; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Routing and layout*

Keywords

Distributed algorithms; theory; resource competitive; attack resistance; jamming; wireless sensor networks

1. INTRODUCTION

The nature of the wireless medium allows malicious devices to threaten the availability of a network by disrupting communication. Such *jamming attacks* have been demonstrated empirically [4, 7] and they challenge the security of sensor networks, future wireless technologies such as the Michigan Micro-Mote [24], SPECKNET [38], and amorphous computing [1, 12].

Wireless devices (*nodes*) are typically battery-powered and once this energy supply is exhausted, replacement may be impossible. How can energy-starved devices defend themselves against a powerful jamming adversary? Gilbert et al. [20] formalized a model of *resource-competitiveness* to deal with problems of this nature. Both the nodes *and* the adversary possess a common resource (energy, in this case) and the goal is to design algorithms whose resource consumption grows asymptotically slower than that of the adversary's. The model is summarized below.

1.1 Resource-Competitiveness

An implicit assumption in many models is that malicious (bad) nodes incur zero cost for attacking. However, this premise is false since attacking requires the expenditure of network resources such

*This research is supported in part by the Agency for Science, Technology and Research (A*STAR), Singapore, under SERC Grant 1224104049.

†This research is supported by NSF grants CCF-1217338 and CNS-1318294 and a grant from the US-Israel Binational Science Foundation. Part of this work was performed at the MADALGO center at Aarhus University, supported by the Danish National Research Foundation grant no. DNRF84.

‡This research is supported by NSF grant CNS-1318294.

SPAA'14, June 23–25, 2014, Prague, Czech Republic.

http://dx.doi.org/10.1145/2612669.2612679

as bandwidth, computation, or energy. In wireless networks populated by battery-powered devices, an algorithm's performance can be measured by the relative energy costs inflicted upon both the good and bad nodes. If the costs to the latter are disproportionately high, then sustained attacks are not feasible since the bad nodes will rapidly deplete their onboard energy supply; the bad nodes are effectively *bankrupted*.

We now formally define what it means for a distributed algorithm $\mathcal{A}$ to be *resource competitive*; this was recently proposed in [20]. Assume a system of n nodes where node v is classified as either *good*, if its actions are prescribed by $\mathcal{A}$, or *bad* if otherwise. Define G as the set of good nodes and F as the set of bad nodes. We assume that the bad nodes may collude and coordinate their attacks; to this end, we assume they are controlled by a single adversary.

Let $\mathcal{C}(i)$ denote the energy expenditure incurred by node i over an execution of algorithm $\mathcal{A}$. If node i is good, then $\mathcal{C}(i)$ is node i's cost for executing the actions prescribed by $\mathcal{A}$. Otherwise, node i is bad and $\mathcal{C}(i)$ is node i's cost for pursuing an arbitrary strategy. Let $T = \sum_{j \in F} \mathcal{C}(j)$ be the total cost to the adversary. That is, T is what the adversary spends in trying to disrupt the network, and we assume that T is *unknown* to the good nodes. Let ρ be a function of T, and possibly other parameters like n; call this the *cost function*. Let τ be a function of any variables except T; call this the *efficiency function*. Algorithm $\mathcal{A}$ is *resource competitive* if it guarantees $\max_{i \in G}\{C(i)\} = O(\rho + \tau)$.

The cost function ρ intuitively captures the relative performance between good nodes and the adversary when $T > 0$. Clearly, *a small ρ is desirable* and, in many cases, we can achieve a function ρ that is *asymptotically* smaller than T (i.e. $o(T)$). However, when $T = 0$, the efficiency function τ captures the unavoidable cost to attain a goal even in the absence of attack. Efficiency in the absence of an attack is important since an algorithm that is costly even when $T = 0$ is undesirable; therefore, τ should also be small (i.e. $O(1)$ or, for large systems of n nodes, $O(\text{polylog } n)$). It is useful to make this separation between the cases $T > 0$ and $T = 0$ explicit via defining these two functions.

1.2 Network Model

Time is divided into discrete *slots* and a node incurs a cost of 1 for sending or listening per slot. If not sending/listening, a node is assumed to be in the energy-efficient sleep state, which has zero cost. This aligns with the operational costs of current devices, which are dominated by transceiver (*radio*) usage [31].

The adversary represents all bad nodes (who may collude and coordinate their attacks if they wish) and pursues an arbitrary jamming strategy. Her energy budget is finite but *unknown* to the (good) nodes. The adversary is *adaptive*: she knows the actions of all nodes in previous time slots and uses this information to inform future attacks. While we consider malicious attacks, in practice the adversary may also represent an abstraction for noise due to collisions, fading effects, or other non-malicious interference.

An *ℓ-uniform adversary* may partition n nodes into at most $1 \leq \ell \leq n$ sets, each of which experiences a different jamming schedule (see [34]). For 1-to-1 communication, we consider a powerful 2-uniform adversary and assume that both devices can be authenticated, i.e., the adversary cannot spoof messages from either device. For 1-to-n communication, we consider a 1-uniform adversary and only assume that the message m can be authenticated. This is a partially-authenticated model where only a *single public key* (the original sender's key) is known. For more on authentication in practice, see [22, 39]. Therefore, the adversary cannot modify m without this being detected and ignored, and we omit further discussion of this in our analysis. Critically, authentication does not imply the existence of shared secrets between nodes.

When two or more messages are sent in the same slot, a message collision occurs and a good node who is listening receives only noise. In practice, noise is detected via *clear channel assessment (CCA)* [33]. When a node hears noise, it cannot tell whether this noise is the result of jamming or due to legitimate messages colliding. A slot is *clear* if it contains neither noise nor any message.

The adversary is assumed to know our protocols except for any random bits generated in the current slot. We adopt the standard assumption that each node can generate independent random bits. The use of randomness in wireless sensor networks is common in the literature (for example [6, 26, 34–36]) and underlies the analysis of frequency hopping spread-spectrum techniques (see [25, 28]) and standard backoff protocols (see [8] and references therein). In practice, the research community has developed functionality along these lines in [17, 37]. Furthermore, without any randomness, an adversary can easily force a cost of $T + 1$ since sending and listening will be deterministic.

1.3 Main Results

We address fundamental communication problems and provide algorithms that are optimally or near optimally resource competitive. Our algorithms are randomized and each has some small probability of failure. We say an event holds *with high probability* (or w.h.p.) if its probability is at least $1 - \frac{1}{\max\{n^c, T^c\}}$ for some tunable constant $c > 0$. Throughout, n is unknown to the good nodes. Let T be the number of slots the adversary jams (this is also unknown to the good nodes). Our results are as follows.

THEOREM 1. *Assume a 2-uniform adaptive adversary. Let $\epsilon > 0$ be a (small) tunable parameter and assume that both Alice (the sender of m) and Bob (the receiver) can be authenticated. There exists an algorithm for 1-to-1 communication with the following guarantees.*

- *Bob receives m with probability at least $1 - \epsilon$.*
- *Alice and Bob incur an expected cost of $O(\sqrt{T \ln(1/\epsilon)} + \ln(1/\epsilon))$.*
- *Alice and Bob terminate within an expected $O(T)$ slots, which is asymptotically optimal.*

Therefore, when Bob can be authenticated, we improve on the (Las Vegas) algorithm of [23] with expected cost $O(T^{\varphi-1} + 1) = O(T^{0.62} + 1)$, where φ is the golden ratio. By combining both algorithms one can achieve expected cost $O(\min\{\sqrt{T \log(1/\epsilon)} + \log(1/\epsilon), T^{\varphi-1} + 1\})$, that is, one with no dependence on ϵ when $T = 0$. We show that Theorem 1 is asymptotically optimal for constant error rate ϵ.

THEOREM 2. *Consider any 1-to-1 communication algorithm in which Alice sends a message to Bob with probability $1 - \epsilon$ for any constant $\epsilon > 0$. Let A and B be Alice's and Bob's costs, respectively. A 1-uniform adaptive adversary can force $\mathrm{E}(A) \cdot \mathrm{E}(B) > (1 - O(\epsilon))T$. In particular, $\max\{E(A), E(B)\} = \Omega(\sqrt{T})$.*

A natural problem is to communicate a message m from a source node to all n nodes in the system. While a cost of roughly $O(\sqrt{T})$ (in expectation) can be obtained via an extension of Theorem 1, we achieve a much more powerful result.

THEOREM 3. *Assume a 1-uniform adaptive adversary and assume m can be authenticated. There exists an algorithm for 1-to-n communication with the following guarantees:*

- *The cost to each node is* $O\left(\sqrt{\frac{T}{n}} \cdot \log^4 T + \log^6 n\right)$ *w.h.p.*
- *All nodes terminate in* $O(T + n\log^2 n)$ *time slots, w.h.p. This latency is optimal as a function of* T.

Therefore, the expected resource costs incurred by good nodes decrease as n grows! Define a *fair* algorithm to be one where all nodes have the same expected cost. To within a polylogarithmic factor, the cost function in Theorem 3 is asymptotically optimal.

THEOREM 4. *Assume a* 1*-uniform adaptive adversary. Any fair algorithm that achieves* 1*-to-*n *communication with constant probability of failure imposes a cost of* $\Omega(\sqrt{T/n})$ *per node.*

Theorem 1 holds when messages from Alice and Bob can be authenticated, and Theorem 3 holds when only m can be authenticated. We prove that giving the adversary the power to spoof messages from Bob actually changes the asymptotic complexity of 1-1 communication. The bound in Theorem 5 matches an algorithm of King, Saia, and Young [23].

THEOREM 5. *Consider a 1-to-1 communication protocol such that Alice sends a message to Bob with constant probability of failure, given a 2-uniform adaptive adversary who can spoof messages from Bob. In any such protocol, the expected cost to either Alice or Bob is* $\Omega(T^{\varphi-1})$ *where* $\varphi = \frac{1+\sqrt{5}}{2}$ *is the golden ratio.*

1.4 Related Work

King et al. [23] provide a Las Vegas resource-competitive algorithm for 1-to-1 communication with an expected cost of $O(T^{\varphi-1}+1) = O(T^{0.62}+1)$ where $\varphi = \frac{\sqrt{5}+1}{2}$ is the golden ratio. The accompanying 1-to-n broadcast algorithm in [23] requires that $\log n$ is *known* and a cost of roughly $T^{\varphi-1}\log n$; therefore, the performance of this algorithm *worsens as* n *increases*. Gilbert and Young [21] give a Monte Carlo 1-to-n partial broadcast algorithm with a better cost ratio. However, the result critically depends on knowing n (not just $\log n$) and still allows the adversary to prevent a small, but constant, fraction of the nodes from receiving the broadcast. Our results address these previous shortcomings by obtaining a cost ratio that improves as n increases without having any information about n, and informing all nodes with high probability.

Much of the work on mitigating jamming attacks focuses on heuristics [3, 5, 10, 13, 25, 28, 32, 40]. We only summarize those with worst-case guarantees. Gilbert et al. [19] derive bounds on the duration for which communication can be disrupted between two devices using deterministic protocols. Pelc and Peleg [30] examine an adversary who jams randomly. Koo et al. [9] address a jamming adversary whose energy budget is known. An interesting series of results by Awerbuch et al. [6] and Richa et al. [34–36] address an adversary whose jamming is bounded within any sufficiently large time window; Dams et al. [11] employs distributed-learning algorithms to overcome this same type of windowed-jamming adversary. Alistarh et al. [2] demonstrate non-cryptographic authentication given a jamming adversary. Ogierman et al. [29] study medium access with adversarial jamming under the signal-to-interference-plus-noise ratio (SINR) model. In the case of multiple channels, Dolev et al. [14, 15] and Gilbert et al. [18], and Emek and Wattenhofer [16] examine communication problems when the adversary cannot jam all channels simultaneously, while Meier et al. [26] examine the problem of node discovery.

These previous results provide valuable solutions to challenging attack models, however, many also require nodes to incur significant costs, either due to sending or listening, relative to the adversary, and this aspect may pose problems in the energy-constrained wireless networks (a few of these results are incomparable given certain model assumption). By taking a resource-competitive approach, we can show that whatever costs are incurred by the nodes are exceeded (asymptotically) by the costs to the adversary.

1.5 Outline

In Sections 2 and 3, we present resource-competitive algorithms for 1-to-1 and 1-to-n broadcast and in Section 4 we give tight or nearly tight lower bounds for these problems.

2. 1-TO-1 COMMUNICATION

Figure 1 provides the pseudocode for 1-to-1 BROADCAST with the canonical players Alice and Bob as sender and receiver, respectively. Let $\epsilon > 0$ be a tunable parameter set prior to execution. The algorithm proceeds in epochs indexed by $i \geq 11 + \lg\ln(8/\epsilon)$, each consisting of a send phase and a nack (negative acknowledgement) phase, each lasting 2^i time slots.[1] We will classify each phase based on the fraction of slots jammed by the adversary.

DEFINITION 1. **(q-Blocking)** *The adversary* q*-blocks a phase if it jams at least a* q *fraction of the slots, for* $0 \leq q \leq 1$. *A repetition that is not* q*-blocked is* q*-unblocked.*

The rationale for the send phase is clear. According to a birthday paradox argument, if Alice sends in $\Theta(\sqrt{2^i \ln(1/\epsilon)})$ random slots and Bob listens in $\Theta(\sqrt{2^i \ln(1/\epsilon)})$ random slots, then, in the absence of jamming, Alice will transmit m to Bob with probability $1-\epsilon$ and Bob will halt. To stop transmission of m the adversary must jam Bob for at least a constant fraction of the slots. However, because the adversary is 2-uniform Alice cannot tell if Bob was jammed and therefore does not know if m was transmitted. If Bob has yet to receive m, he sends a nack message back to Alice using the same protocol, which, in the absence of jamming, Alice will correctly receive with probability $1-\epsilon$. If Alice does not receive a nack (and was not heavily jammed) she assumes Bob received m and already halted; therefore she halts as well. How does Bob know to halt in the event that Alice prematurely halted? If, in a subsequent epoch, Bob hears little jamming and yet does not receive m, he assumes Alice has halted prematurely and halts. In a similar fashion Alice will halt in the nack phase only if she hears little jamming and does not receive a nack.

Adaptive adversaries are difficult to reason about because their choices can be subtly informed by nodes' past behavior. Lemma 1 allows us to focus on a restricted class of adversarial strategies.

LEMMA 1. *Without loss of generality, in any phase all of the un-jammed slots precede all the jammed slots.*

PROOF. Within one phase of epoch i, the behavior of nodes in each slot (whether they send or listen) is independent of their past behavior. The adversary can gain no information from the nodes by jamming, nor can it influence their future behavior by jamming. Thus, an adversary that chooses to jam slot k after leaving slots 1 through $k-1$ unjammed is equivalent to an adversary that leaves slot k unjammed as well but commits to jamming the last time slot 2^i. In this way all jamming can, without loss of generality, be postponed to a contiguous interval at the end of the phase. □

Lemma 1 shows that we can assume the adversary observes the behavior of all nodes up until a certain point and then *jams for the remainder of the phase*, independent of the nodes' behavior. The point at which such jamming begins is clearly a choice that needs to be made online.

[1] We use $\lg x$ to refer to $\log_2 x$.

1-to-1 BROADCAST for epoch $i \geq 11 + \lg\ln(8/\epsilon)$

- Send Phase — For each of the 2^i slots, do:
  - With probability $\sqrt{\frac{\ln(8/\epsilon)}{2^{i-1}}}$, Alice sends m.
  - With probability $\sqrt{\frac{\ln(8/\epsilon)}{2^{i-1}}}$, Bob listens.
- If Bob receives m or hears less than $\frac{\sqrt{2^{i-1}\ln(8/\epsilon)}}{4}$ noisy slots without hearing m, he terminates.
- Nack Phase — For each of the 2^i slots, do:
  - With probability $\sqrt{\frac{\ln(8/\epsilon)}{2^{i-1}}}$, Bob sends `nack`.
  - With probability $\sqrt{\frac{\ln(8/\epsilon)}{2^{i-1}}}$, Alice listens.
- If Alice hears less than $\frac{\sqrt{2^{i-1}\ln(8/\epsilon)}}{4}$ noisy slots without hearing a `nack`, she terminates.

Figure 1: Pseudocode for epoch i of 1-to-1 BROADCAST.

PROOF. (Theorem 1) The 1-to-1 BROADCAST algorithm can fail in several ways: Alice or Bob can exceed their energy budgets, as a function of the adversaries cost T; Alice can halt prematurely, before Bob receives m; and Bob can halt prematurely, falsely thinking that Alice has already halted prematurely. We first show that with probability $1 - o(\epsilon)$, Alice's and Bob's costs never exceed twice their expectations, then address halting. Let $i = s + \lg\ln(8/\epsilon)$ be the epoch index, where $s \geq 11$, and let $p_i = \sqrt{\ln(8/\epsilon)/2^{i-1}}$ be the sending/listening probability in epoch i.

Costs: Let X be the actual cost of Alice (or Bob) in epoch i. By linearity of expectation $\mathrm{E}(X) \leq p_i(2 \cdot 2^i) = \sqrt{\ln(8/\epsilon)/2^{i-1}} \cdot 2^{i+1} = \sqrt{2^{i+3}\ln(8/\epsilon)}$. By a Chernoff bound the probability that X exceeds twice its expectation is less than $\exp(-\mathrm{E}(X)/3)$. The probability that Alice's cost in *any* epoch exceeds twice its expectation is at most $\sum_{i \geq 11+\lg\ln(8/\epsilon)} \exp\left(-\sqrt{2^{i+3}\ln(8/\epsilon)}/3\right) < \sum_{s \geq 11} (\epsilon/8)^{2^{(s+3)/2}/3} = o(\epsilon)$.

Note that the expected cost to Alice and Bob in phases $11 + \lg\ln(8/\epsilon)$ through i is $O(\sqrt{2^i \ln(1/\epsilon)} + \ln(1/\epsilon))$. If the adversary jams $T = \Omega(2^i)$ slots in phase i then Alice and Bob have spent $O(\sqrt{T\ln(1/\epsilon)} + \ln(1/\epsilon))$, that is, neither will exceed their energy budget through epoch i. For Alice or Bob to exceed their energy budget the adversary must, at the very least, get one or both parties to continue to epoch $i+1$ while jamming $o(2^i)$ slots. By Lemma 1, we can assume the adversary jams a suffix of the 2^i slots, though the moment when she begins jamming can be chosen adaptively, by observing the behavior of Alice and Bob.

Send Phase — Alice is still running: We need to prove two claims. First, any adversarial strategy that stops Alice from transmitting m to Bob with probability $1 - O(\epsilon)$ must jam $\Omega(2^i)$ slots. Second, if the adversary does, in fact, jam $\Omega(2^i)$ slots Bob will not halt (correctly) and proceed to epoch $i + 1$, with high probability. For the first claim, the probability that Alice fails to transmit the message to Bob in the first unjammed $2^i/2$ slots is $(1-p_i^2)^{2^{i-1}} = \epsilon/8$. Thus, any adversarial strategy that stops the transmission of m with probability greater than $\epsilon/8$ must be committed to jamming at least half the slots. Turning to the second claim, if the adversary decides to jam the last $2^i/2$ slots, the expected number of jammed slots heard by Bob is at least $\sqrt{2^{i-1}\ln(8/\epsilon)}$ and the probability he hears less than $\sqrt{2^{i-1}\ln(8/\epsilon)}/4$ (possibly halting prematurely) is, by a Chernoff bound, less than $\exp(-(3/4)^2\sqrt{2^{i-1}\ln(8/\epsilon)}/2) < (\epsilon/8)^{2^{(s-5)/2}} < \epsilon/8$.

Send Phase — Alice has halted prematurely: It is no longer possible for the message m to be sent so the correct behavior is for Bob to halt. To prevent this the adversary must cause Bob to hear a large number of jammed slots. If the adversary jams less than $2^i/16$ slots then the expected number heard by Bob is less than $\sqrt{2^{i-1}\ln(8/\epsilon)}/8$ and, by a Chernoff bound, the probability that Bob hears less than $\sqrt{2^{i-1}\ln(8/\epsilon)}/4$ is $\exp(-\sqrt{2^{i-1}\ln(8/\epsilon)}/24)$ $= (\epsilon/8)^{2^{(s-1)/2}/24} < \epsilon/8$.

Nack Phase: The analysis of the nack phase is identical. If Bob is still running, the adversary cannot stop Alice from receiving a nack with probability greater than $\epsilon/8$ without jamming $2^i/2$ slots. If the adversary does jam Alice for least $2^i/2$ slots then Alice will hear a sufficient number to continue to epoch $i+1$, with probability $1 - (\epsilon/8)^{2^{(s-5)/2}}$. Finally, if Bob has halted (after receiving m or prematurely), then Alice will halt with probability at least $1 - \epsilon/8$ unless the adversary jams $2^i/16$ slots.

To sum up, if the adversary wants to prevent Bob from receiving m, or prevent Alice from receiving a nack, or prevent the the second party to halt after the first has halted, it must $(1/16)$-block one of the phases. If the adversary does not $(1/16)$-block one of the phases, the probability of any type of failure in this epoch is at most $2(\epsilon/8 + (\epsilon/8)^{2^{(s-5)/2}} + (\epsilon/8)^{2^{(s-1)/2}/24}) < \epsilon/2$, where $s = i - \lg\ln(8/\epsilon) \geq 11$. So long as epoch i is the *last* epoch that is at least $(1/16)$-blocked by the adversary, the expected cost to Alice or Bob after epoch i is at most $\sum_{j \geq i+1} \mathrm{E}(\text{cost in epoch } j) \cdot \Pr(\text{still running in epoch } j)$, which is $\sum_{j \geq i+1} p_j 2^{j+1} \cdot (\epsilon/2)^{j-(i+1)}$ $= O(\sqrt{2^i \ln(1/\epsilon)}) = O(\sqrt{T\ln(1/\epsilon)})$. A similar calculation gives the total latency as $O(2^i) = O(T)$ which is asymptotically optimal since the adversary can always force T latency. □

3. 1-TO-n COMMUNICATION

The pseudocode for *epoch* i of 1-to-n BROADCAST is provided in Figure 2. Each epoch consists of $b\,i^2$ *repetitions* each consisting of 2^i slots. The *status* t_u of node u is initially `informed` if u is the sender and `uninformed` otherwise. The variable S_u is reset to 16 at the beginning of each epoch and is non-decreasing throughout the epoch. The parameters $b > 0$ and $d > 0$ are sufficiently large constants and the *first* epoch i is some sufficiently large constant.

While the pseudocode is simple, the design decisions and mechanisms that lead to correctness are intricate. We take some time to provide a discussion of these decisions now.

3.1 A Tour of the Algorithm

The variable S_u controls the probability that u is sending or listening in a given slot. We want S_u to be sufficiently high so that the message is quickly disseminated, but not so high that u expends too much energy. The right bound for an epoch-i repetition, for $i > \log n$, is about $\sqrt{2^i/n}$.[2] However, we do not assume u knows n, even approximately, so it cannot jump straight to the ideal S_u value. We implicitly determine an estimate of n by the following strategy. In each slot of a repetition every u sends with probability $S_u/2^i$; if u is `informed`, then it sends the message and if it is `uninformed` then it sends noise. The purpose of sending noise is to let all nodes gauge how large n is relative to 2^i (assuming no

[2]We want the number of `informed` nodes to increase geometrically in each repetition. An n-party version of the birthday paradox shows this is possible if each `informed` node sends in $\Theta(\sqrt{2^i/n})$ random slots and `uninformed` nodes listen in $\Theta(\sqrt{2^i/n})$ random slots. However, as we will show, the required analysis is far more involved than a birthday paradox argument.

1-to-n BROADCAST for epoch i with node u

- $S_u \leftarrow 16$
- Repeat $b \cdot i^2$ times:
  - For each of the 2^i slots:
    - If $t_u \in \{\texttt{informed}, \texttt{helper}\}$, then send m with probability $\frac{S_u}{2^i}$
    - If $t_u = \texttt{uninformed}$, then send noise with probability $\frac{S_u}{2^i}$
    - Listen with probability $\frac{S_u\, di^3}{2^i}$
  - Let C_u be the number of clear slots heard and $C'_u = \max\{0, C_u - \frac{1}{2}S_u di^3\}$
    - $S_u \leftarrow S_u \cdot 2^{C'_u/(S_u di^4)}$
  - Execute at most one of the following Cases (in order):
    1. If $S_u > 360 \cdot 2^{i/2}$, then terminate
    2. $t_u = \texttt{uninformed}$: If m is heard, then $t_u \leftarrow \texttt{informed}$
    3. $t_u = \texttt{informed}$: If m is heard more than $\frac{di^3}{200}$ times, then $t_u \leftarrow \texttt{helper}$ and $n_u \leftarrow \frac{2^i}{(S_u)^2}$
    4. $t_u = \texttt{helper}$: If $S_u \geq 360\sqrt{\frac{2^i}{n_u}}$, then terminate

Figure 2: Pseudocode for epoch i of 1-to-n BROADCAST

jamming). If a node u hears a sufficient number of clear slots it increases S_u. Note that u expects to listen in $S_u di^3$ slots; if all are clear then S_u will increase by a roughly $2^{1/(2i)}$ factor at the end of the repetition, which is quite small. There are two reasons we need the $\{S_u\}$-values to increase slowly. First, we need to spend about $\log n < i$ repetitions when $S_u \approx \sqrt{2^i/n}$ in order to quickly disseminate the message, so it is important that we do not increase S_u too aggressively and overshoot the ideal value. Second, in order for all nodes to have roughly the same cost, it is important that S_u/S_w be bounded for any two nodes u and w. By increasing S_u and S_w tentatively, we can bound the divergence S_u/S_w over all bi^2 repetitions. (Of course, the adversary can artificially keep S_u low by jamming a large fraction of the slots. To push the nodes into epoch $i \gg \log n$ it will need to jam about $T = \Omega(i^2 2^i)$ slots.)

It is clearly a bad idea for nodes to halt as soon as they receive the message. In order to distribute the costs effectively, `informed` nodes must stay around to help further disseminate the message. The question is *how long* should they keep running and under what circumstances? A natural halting criterion is *stop when u has heard the message a sufficient number of times*, say poly(i). By a Chernoff bound, a node u that halts can deduce that all nodes have heard the message at least once, w.h.p. This idea does not lead naturally to an algorithm with cost $\tilde{O}(\sqrt{T/n})$. The adversary can jam at a rate that will cause roughly half the nodes to hear messages beyond the halting threshold, leaving the other half to continue running the protocol. To get the remaining nodes to hear the message a sufficient number of times they must up their sending rates (the $\{S_u\}$ values) by a constant factor. The adversary can jam at a rate to cause half the nodes to halt again, necessitating the remaining nodes to up their sending rates, and so on. The last node running will therefore spend about $\tilde{O}(\sqrt{T/n} + \sqrt{T/(n/2)} + \sqrt{T/(n/4)} + \cdots) = \tilde{O}(\sqrt{T})$. That is, an algorithm employing this strategy does not benefit from having a large number of nodes.

Our solution is somewhat counterintuitive. When a node u hears the message a sufficient number of times in one repetition in epoch j ($dj^3/200$) it becomes a `helper` and, assuming S_u is about the ideal value $\sqrt{2^j/n}$, estimates n by $n_u = 2^j/(S_u)^2$. It continues to act exactly like an `informed` node, except that when S_u climbs to $360\sqrt{2^i/n_u}$, in a subsequent epoch $i \geq j$, it halts. We prove that when S_u reaches this threshold, all other nodes have `helper` status w.h.p. An important feature of this approach is that once `helper` nodes begin halting, the ability of other nodes to halt is not affected. Note that it is hearing *silence* that causes S_u to grow from 16 to $360\sqrt{2^i/n_u}$ in one epoch, and silence is *free*. To prevent `helpers` from halting the adversary is forced to jam a constant fraction of the slots.

There is one last issue related to halting. With some tiny but non-zero probability, all but one node will become a `helper` and halt. The remaining node will never become a `helper`, and therefore never reach the halting condition described above, i.e., its expected cost would be *infinite*. Therefore, we need an alternative halting condition (Case 1 in the pseudocode) to force these exceptionally unlucky nodes to halt and preserve our cost function of $\tilde{O}(\sqrt{T/n})$.

3.2 Preliminaries

Before we begin our main analysis, we state some technical lemmas regarding the version of Chernoff bounds used here. We also prove some preliminary results that are used later on.

Standard Chernoff Bounds: We review well-established Chernoff bounds that we employ in this work:

THEOREM 6. *([27]) Let $X_1, \ldots, X_n$ be independent trials such that $\Pr(X_i) = p$ and let $X = \sum_{i=1}^n X_i$. For any $\delta > 0$,*

$$\Pr(X > (1+\delta)\,\mathrm{E}[X]) \leq \left[\frac{e^{\delta}}{(1+\delta)^{(1+\delta)}}\right]^{\mathrm{E}[X]}$$

$$\Pr(X < (1-\delta)\,\mathrm{E}[X]) \leq \left[\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right]^{\mathrm{E}[X]}.$$

We only use the following corollaries of Theorem 6.

COROLLARY 1. *([27]) Let $X_1, \ldots, X_n$ be independent trials such that $\Pr(X_i) = p$ and let $X = \sum_{i=1}^n X_i$. For any δ, where $0 < \delta < 1$,*

$$\Pr(X > (1+\delta)\,\mathrm{E}[X]) \leq e^{-\delta^2\,\mathrm{E}[X]/3}$$

$$\Pr(X < (1-\delta)\,\mathrm{E}[X]) \leq e^{-\delta^2\,\mathrm{E}[X]/2}$$

Furthermore: $\Pr(|X - \mathrm{E}[X]| > \sqrt{3\,\mathrm{E}[X]\ln(1/\epsilon)}) < 2\epsilon$

The last bound of Corollary 1 is obtained from the first two by setting $\delta = \sqrt{3\ln(1/\epsilon)/\,\mathrm{E}[X]}$, if $\delta < 1$, and from Theorem 6 if $\delta \geq 1$. In our application we normally set $\epsilon < 1/\mathrm{poly}(n)$. The following inequality is well known:

FACT 1. *$1 - y \geq e^{-2y}$ for any $0 \leq y \leq 1/2$.*

Useful Properties of Our Algorithm: Let A be the set of all non-terminated nodes that currently know m at some point in epoch i. Let V be the set of all nodes that have not terminated. Define $S_A = \sum_{u \in A} \frac{S_u}{2^i}$ and $S_V = \sum_{u \in V} \frac{S_u}{2^i}$. Define p_m to be the probability that exactly one node sends m while all others stay silent, and let p_c be the probability that an unjammed slot is clear. We have the following bounds on p_m and p_c.

LEMMA 2. *Throughout 1-to-n* BROADCAST, $S_A \cdot e^{-2S_V} \leq p_m \leq eS_A \cdot e^{-S_V}$ *and* $e^{-2S_V} \leq p_c \leq e^{-S_V}$.

PROOF. The probability that some node $u \in A$ transmits m while $V - \{u\}$ stay silent is $\sum_{u \in A} \left(\frac{S_u}{2^i} \cdot \prod_{v \in V - \{u\}} (1 - \frac{S_v}{2^i}) \right)$. By Fact 1 this is at least $\sum_{u \in A} \frac{S_u}{2^i} \cdot e^{-2S_V} = S_A e^{-2S_V}$ and at most $\sum_{u \in A} \frac{S_u}{2^i} \cdot e^{-(S_V - S_u/2^i)} < eS_A e^{-S_V}$. The probability of an unjammed slot being clear is $p_c = \prod_{v \in V} (1 - \frac{S_v}{2^i})$, which is at most e^{-S_V} and at least e^{-2S_V}, by Fact 1. □

We establish properties about S_u and t_u that later allow us to make guarantees about epochs $i > \lg n$.

LEMMA 3. *With high probability, every node u has $S_u = 16$ for all epochs $i \leq \lg n$.*

PROOF. The probability that a slot is clear is at most $p_c < e^{-S_V} = e^{-\sum_v S_v/2^i} \leq e^{-16n/2^i}$. When $i \leq \frac{1}{2} \lg n$ $p_c < e^{-16\sqrt{n}}$ and the probability that there are *any* clear slots is superpolynomially small. For any $i \leq \lg n$ we have $p_c \leq e^{-16}$. The expected number of clear slots heard by u is at most $S_u di^3/e^{16}$ but it must hear at least $S_u di^3/2$ to increase S_u. By a Chernoff bound this happens with probability $\exp(-\Omega(i^3)) = \exp(-\Omega(\lg^3 n))$. □

LEMMA 4. *In epochs $i \leq \lg n$, w.h.p. no node becomes a* `helper` *or terminates, i.e., each node is either* `uninformed` *or* `informed`.

PROOF. By Lemma 3, when $i \leq \lg n$, $S_u = 16$ for all u. The probability m is successfully sent in a slot is $p_m \leq eS_A \cdot e^{-S_V}$ by Lemma 2, which is at most $e|A|(16/2^i)e^{-16n/2^i} \leq 16e^{-15}$. A node u will hear less than $(16e^{-15}) \cdot S_u di^3 = (16)^2 e^{-15} di^3$ transmissions of m in expectation but must hear $di^3/200$ to become a helper. By a Chernoff bound, w.h.p. this does not happen. □

3.3 Informing All Nodes

In this section, we analyze how nodes adjust their sending and listening probabilities, and the rate at which m is disseminated. We begin by addressing the divergence between $\{S_u\}$ values.

LEMMA 5. *Consider any epoch $i > \lg n$. With probability $1 - \exp(-\Omega(i))$, we have $S_u/S_v \leq 2$ throughout the epoch, for any two nodes u and v.*

PROOF. We begin with an informal argument. Suppose that in one repetition a $q \leq 1$ fraction of the slots are clear. We expect u to hear $C_u = qdi^3 S_u$ clear slots, so $C'_u = \max\{0, (q-1/2)\} \cdot di^3 S_u$. Then, S_u updates to $S_u \cdot 2^{\frac{C'_u}{S_u di^4}} = S_u \cdot 2^{\frac{\max\{0,(q-1/2)\}}{i}}$. However, C_u may not be close to its expectation, due to both random chance *and* the adversary's choice of when to begin jamming. We will show that for all u, regardless of the adversary's choice, $C_u = qdi^3 S_u \pm O(\sqrt{S_u di^4})$ with probability $1 - \exp(-\Omega(i))$. That is, S_u will drift from its ideal value by a factor of $2^{O(1/\sqrt{S_u di^4})}$. Over bi^2 repetitions the total drift of S_u will be bounded by $\sqrt{2}$, hence S_u/S_v will be bounded by 2.

Fix a repetition j in epoch i, a node u, and a specific time slot when the adversary begins jamming; by Lemma 1 we can assume the adversary jams for an interval at the end of the repetition. Let q be the actual fraction of clear slots and C_u be the number observed by u. By linearity of expectation we have $\mathrm{E}[C_u] = qdi^3 S_u \leq di^3 S_u$. Fix a constant c and let $R_u = \sqrt{c \cdot di^4 S_u}$, i.e., R_u is at least as large as $\sqrt{\mathrm{E}[C_u] \cdot ci}$ since $q \leq 1$. By a Chernoff bound the probability that $|\mathrm{E}[C_u] - C_u| > R_u$ is less than $2e^{-ci/3}$. By the union bound it follows that for each repetition j (bi^2 values), node u (n values), and each jamming time (2^i values) both C_u and C'_u are within $\sqrt{c \cdot di^4 S_u}$ of their expectations with probability $1 - bi^2 n 2^{i+1} e^{-ci/3} > 1 - e^{-(c/3 - O(1))i}$, since $i > \lg n$.

At the end of repetition j, node u sets $S_u = S_u \cdot 2^{C'_u/(di^4 S_u)}$, which w.h.p. is $S_u \cdot 2^{(\mathrm{E}[C'_u] \pm R_u)/di^4 S_u} = S_u \cdot 2^{\max\{0,(q-1/2)\}/i} \cdot 2^{\pm\sqrt{c/(di^4 S_u)}}$, where q is common to all nodes u in repetition j. Note that since $S_u \geq 16$, the error factor is never more than $2^{\pm\sqrt{c/(16di^4)}}$. Over bi^2 repetitions the accumulated error factor is, with high probability, at most $2^{\pm\sqrt{c/(16di^4)} bi^2} = 2^{\pm\sqrt{cb^2/(16d)}}$. For $d > 4cb^2/16$ sufficiently large the error factor is between $1/\sqrt{2}$ and $\sqrt{2}$ and S_u/S_v bounded by 2, for any u and v. □

Lemma 5 makes a claim about all epochs greater than $\log n$ whereas several lemmas stated later are concerned with epochs beyond $\log n + 8$. We are not concerned about whether m is quickly disseminated in those 8 epochs, so long as other properties relating to correctness are maintained. For example, Lemma 6 states that we never simultaneously have both `uninformed` and `helper` nodes, w.h.p.

LEMMA 6. *In a repetition of epoch $i > \lg n$, if any node becomes a* `helper` *node then no nodes remain* `uninformed`*, with probability $1 - \exp(-\Omega(i^3))$, regardless of the adversary's jamming strategy.*

PROOF. Recall that p_m is the probability that any given unjammed time slot contains a message. Suppose, for the purpose of analysis, that the adversary commits to leaving a q fraction of the slots unjammed. The expected number of messages heard by u is $L_u = qp_m S_u di^3$. If $L_u < di^3/400$ then by a Chernoff bound the probability that u hears $2 \cdot L_u$ messages (meeting the threshold to become a `helper`) is $\exp(-\Omega(i^3))$. If $L_u \geq di^3/400$ the probability that u hears zero messages (and stays uninformed) is $\exp(-\Omega(i^3))$ for d large enough, also by a Chernoff bound. By Lemma 5, S_u and S_w differ by a factor of at most 2, so the analysis above applies equally well to all nodes. By a union bound over all n nodes and all 2^i jamming fractions q, the probability that one `helper` node and one `uninformed` node exist after the repetition is $n \cdot 2^i \cdot \exp(-\Omega(i^3)) = \exp(-\Omega(i^3))$. □

LEMMA 7. *Call a repetition in epoch i successful if each node u increases S_u by a factor at least $2^{1/(10i)}$. Consider a repetition in epoch $i \geq \lg n + 8$ where $S_V < \sqrt{\frac{n}{2^i}}$. The probability that the adversary can prevent the repetition from being successful without $\frac{1}{10}$-blocking it is $\exp(-\Omega(i^3))$.*

PROOF. By Lemma 1 the adversary can be assumed to jam an interval of slots at the end of the repetition. The adversary wants to prevent nodes from hearing clear slots, so his best strategy is to jam the maximum $2^i/10$ slots allowed in an unblocked repetition, leaving $9 \cdot 2^i/10$ unjammed. By Lemma 2, the expected number of clear slots witnessed by u is at least $p_c \cdot \frac{9}{10} S_u di^3 \geq e^{-2S_V} \cdot \frac{9}{10} \cdot S_u di^3 \geq \frac{9 \cdot S_u di^3}{10 \cdot e^{1/8}} > 0.79 \cdot S_u di^3$. (The lower bound on i and upper bound on S_V implies $2S_V \leq 1/8$.) By a Chernoff bound the probability that u witnesses less than $0.6 S_u di^3$ clear slots is $\exp(-\Omega(i^3))$. If $C_u \geq 0.6 S_u di^3$ then $C'_u \geq 0.1 S_u di^3$ and S_u is increased by a $2^{1/(10i)}$ factor. By a union bound over all $n < 2^i$ nodes, the probability that each u grows S_u by this factor is $1 - \exp(-\Omega(i^3))$. □

LEMMA 8. *Consider any epoch $i \geq \lg n + 8$ and threshold $h > 0$. If $S_V \leq h$ before one repetition and $S_V \geq 4h$ later in the epoch then at least $3i$ of the intervening repetitions were $\frac{1}{2}$-unblocked, with probability $1 - \exp(-\Omega(i))$.*

PROOF. The proof the Lemma 5 shows that the divergence in S_u-values from their expectations (due to random chance and any adversarial jamming strategy) is by a factor between $\frac{1}{\sqrt{2}}$ and $\sqrt{2}$ over the entire epoch, with probability $1 - \exp(-\Omega(i))$. That is, of the factor 4 increase in S_V (from h to $4h$), at most $\sqrt{2}$ is due to the samples $\{C_u\}$ deviating from their expectations. We can therefore assume without loss of generality that C_u always matches its expectation and analyze the number of $\frac{1}{2}$-unblocked repetitions that cause S_V to grow by a $2\sqrt{2}$ factor.

If the adversary q'-blocks a repetition then the fraction q of clear slots must be at most $1 - q'$. If $q \leq 1/2$ then S_u (and S_V) does not increase. If $q > 1/2$ (implying the repetition is $\frac{1}{2}$-unblocked) then S_u (and S_V) increases by a $2^{(q-1/2)/i} \leq 2^{1/(2i)}$ factor, since $q \leq 1$. Thus, to achieve a $2\sqrt{2} = 2^{3/2}$ factor increase in S_V we need $(3/2)/(1/(2i)) = 3i$ repetitions that are $\frac{1}{2}$-unblocked. □

To summarize, Lemmas 7 and 8 imply that after approximately $10\,i\,\lg(\sqrt{2^i/n}) < 5\,i^2$ $\frac{1}{10}$-unblocked repetitions, w.h.p. $S_V \geq \sqrt{\frac{n}{2^i}}$. If, later in the epoch, $S_V > 4\sqrt{\frac{n}{2^i}}$ then we have witnessed $3i$ or more $\frac{1}{2}$-unblocked repetitions. We set $b \geq 10$ so there are ample repetitions for S_V to grow sufficiently large. We will now prove that in those $3i$ $\frac{1}{2}$-unblocked repetitions, all nodes will become `informed` and attain `helper` status.

LEMMA 9. *Consider an epoch $i \geq \lg n + 8$ before any nodes have achieved `helper` status. With probability $1 - \exp(-\Omega(i))$ the adversary cannot prevent S_V from exceeding $4\sqrt{n/2^i}$ and cannot prevent all nodes from attaining `helper` status, without $\frac{1}{10}$-blocking a constant fraction of the repetitions.*

PROOF. By Lemma 7 the adversary cannot prevent S_V from growing to $\sqrt{\frac{n}{2^i}}$ without $\frac{1}{10}$-blocking a constant fraction of the repetitions. Moreover, it cannot prevent S_V from then growing to $4\sqrt{\frac{n}{2^i}}$ without $\frac{1}{2}$-blocking a constant fraction of the remaining repetitions. By Lemma 8 there are at least $3i$ $\frac{1}{2}$-unblocked repetitions while $\sqrt{\frac{n}{2^i}} \leq S_V \leq 4\sqrt{\frac{n}{2^i}}$.

To prevent dissemination of m, the adversary's optimum strategy is to jam as much as possible; namely, $2^i/2 - 1$ slots in a $\frac{1}{2}$-unblocked repetition, and all 2^i slots in a $\frac{1}{2}$-blocked repetition. By Lemma 1, we can assume that the adversary commits to jamming an interval of slots at the end of a repetition. We allow the adversary to decide whether to $\frac{1}{2}$-block the repetition *after* all nodes have committed to which slots they will send and listen. If they can accomplish a task in the first $2^i/2$ slots, then to stop them the adversary is forced to $\frac{1}{2}$-block the repetition. Therefore, we analyze the probability of accomplishing a task within $2^i/2$ slots.

So long as $S_V \leq 4\sqrt{\frac{n}{2^i}}$ we will have $S_V \leq 1/4$ since $i \geq \lg n + 8$. Recall that A is the set of `informed` nodes at the beginning of some repetition, and $S_A = \sum_{u \in A} S_u/2^i$. The probability that an unjammed slot contains m is p_m, which is at least $S_A/e^{2S_V} \geq S_A/e^{1/2}$. By Lemma 5, if $S_V \geq \sqrt{\frac{n}{2^i}}$ then for every node u, $S_u \geq \frac{1}{2}\sqrt{\frac{2^i}{n}}$. Therefore, $p_m \geq S_A/e^{1/2} \geq \frac{|A|}{e^{1/2}2\sqrt{2^i n}}$.

Let X_m be the number of slots containing m and I_m be the number of nodes that hear m. We will show that conditioned on X_m being a constant fraction of its expectation, I_m will be at least a constant fraction of its expectation (about $i^3|A|$) with probability $1 - \exp(-\Omega(i^3))$. When $|A|$ is very small, however, the probability that X_m is too small may not be completely negligible.

By linearity of expectation we have $\mathrm{E}[X_m] \geq p_m \cdot \frac{1}{2}2^i \geq \frac{|A|\sqrt{2^i}}{4e^{1/2}\sqrt{n}} \geq \frac{|A|\cdot\sqrt{2^i}}{6.6\sqrt{n}}$. By a Chernoff bound, the probability that X_m is less than $\mathrm{E}[X_m]/3$ is at most $\exp(-(2/3)^2\,\mathrm{E}[X_m]/2) = \exp(-(2/9)\,\mathrm{E}[X_m])$. Call a repetition good if this holds and bad otherwise. We will proceed under the assumption that the repetition is good, i.e., $X_m \geq \mathrm{E}[X_m]/3 > \frac{|A|\cdot\sqrt{2^i}}{19.8\sqrt{n}}$. The probability that a node u hears m in this repetition is at least $1 - \left(1 - \frac{di^3 S_u}{2^i}\right)^{X_m}$

$$\geq 1 - e^{-\frac{di^3 S_u}{2^i} \cdot \frac{|A|\cdot\sqrt{2^i}}{19.8\sqrt{n}}} \quad \text{By the lower bound on } X_m$$

$$\geq 1 - e^{-\frac{di^3}{2\sqrt{2^i \cdot n}} \cdot \frac{|A|\cdot\sqrt{2^i}}{19.8\sqrt{n}}} \quad \text{By the lower bound on } S_u$$

$$\geq \begin{cases} \frac{di^3|A|}{79.2n} & \text{If the exponent is} \geq -1 \text{ and Fact 1} \\ 1 - 1/e & \text{Otherwise} \end{cases}$$

If we are in the second case in the last inequality then $\mathrm{E}[I_m] \geq n(1 - 1/e)$ by linearity of expectation and, by a Chernoff bound, I_m is at least $n/4$ with probability $1 - \exp(-\Omega(n))$. If we are in the first case then $\mathrm{E}[I_m] \geq di^3|A|/79.2$ and $I_m \geq i^3|A|$ with probability $1 - \exp(-\Omega(i^3))$, again by a Chernoff bound assuming that $d > 79.2$. Thus, after $\lg(n/4)/\lg(i^3)$ good repetitions, the number of `informed`/`helper` nodes will be at least $n/4$ with high probability. We need to bound the number of bad repetitions, where $X_m < \mathrm{E}[X_m]/3$, and show that when $|A| \geq n/4$, all nodes become helpers with high probability. The second claim is easy to establish. If $|A| \geq n/4$ then the expected number of messages heard by a node u is at least $X_m \cdot di^3 S_u/2^i \geq di^3|A|/39.6n > di^3/160$. The probability that u hears at least $di^3/200$ messages (exceeding the threshold to become a `helper`) is, by a Chernoff bound, $1 - \exp(-\Omega(i^3))$.

By this analysis, the probability of a bad repetition is $\Pr[X < \frac{\mathrm{E}[X_m]}{3}] < \exp(-(\frac{2}{9})\,\mathrm{E}[X_m]) < \exp(-(\frac{2}{9})\frac{|A|\cdot\sqrt{2^i}}{6.6\sqrt{n}})$, which is $\exp(-\Omega(i^3))$ for $|A| \geq i^3$. However, for all $|A| \geq 1$, it is at most $\exp(-(\frac{2}{9}) \cdot \frac{16}{6.6}) < 0.59$ since $\sqrt{2^i/n} \geq 16$. From the adversary's perspective, there is no need to block bad repetitions, only good ones. Moreover, after the first $\frac{1}{2}$-unblocked good repetition $|A|$ will be i^3 and we only need $\lg(n/4)/\lg(i^3)$ more repetitions for all nodes to attain `helper` status, with probability $1 - \exp(-\Omega(i^3))$. By Lemma 8 there is room for $3i$ $\frac{1}{2}$-unblocked repetitions until $S_V \geq 4\sqrt{n/2^i}$. With probability $1 - \exp(-\Omega(i))$, we do not have $3i - \frac{\lg(n/4)}{\lg(i^3)}$ bad repetitions before the first good one. □

3.4 Terminating the Algorithm

In this section, we show that w.h.p. all `helper` nodes terminate and derive the cost and efficiency functions.

LEMMA 10. *If node u assumes `helper` status in epoch $i > \lg n$, then w.h.p. $S_u \geq \frac{2^{i/2}}{45\cdot\sqrt{n}}$.*

PROOF. Let u be an `informed` node that changes its status to `helper`. Can $S_u < \frac{2^{i/2}}{45\sqrt{n}}$ when u receives enough messages to assume `helper` status? If so, then by Lemma 5 we know that each other node w has $S_w \leq 2 \cdot \frac{2^{i/2}}{45\sqrt{n}}$ which implies that $S_V < \frac{1}{2^i}(n \cdot \frac{2\cdot 2^{i/2}}{45\sqrt{n}}) \leq \frac{2\cdot\sqrt{n}}{45\cdot 2^{i/2}}$. Then, the expected number of messages that u hears is then at most $p_m \cdot S_u\, di^3 \leq \frac{eS_A S_u di^3}{e^{S_V}} \leq eS_A S_u\, di^3 \leq eS_V S_u\, di^3 < e \cdot \frac{1}{2^i} \cdot (n \cdot \frac{2\cdot 2^{i/2}}{45\sqrt{n}})(\frac{2^{i/2}}{45\sqrt{n}})di^3 < \frac{di^3}{370}$. By Chernoff bounds and a union bound, it follows that w.h.p. u would not receive more than $\frac{di^3}{200}$ messages and therefore would not change its status to `helper`. □

LEMMA 11. *Consider an epoch $i > \lg n$. If $S_u \geq \frac{360 \cdot 2^{i/2}}{\sqrt{n_u}}$, then w.h.p. all nodes have* `helper` *status.*

PROOF. $S_u = \frac{2^{j/2}}{\sqrt{n_u}}$ when u changes its status to `helper` in some epoch j. By Lemma 10, $S_u \geq \frac{2^{j/2}}{45\sqrt{n}}$ and so it follows that $\sqrt{n_u} \leq 45\sqrt{n}$. In some epoch $i \geq j$, by assumption $S_u \geq \frac{360 \cdot 2^{i/2}}{\sqrt{n_u}} \geq \frac{8 \cdot 2^{i/2}}{\sqrt{n}}$. By Lemma 5, this implies that $S_V > \frac{4\sqrt{n}}{2^{i/2}}$. By Lemma 9, this implies that all nodes have already set their status to `helper`. □

LEMMA 12. *Consider an epoch $i > \lg n$ and let u have status* `helper`*. W.h.p. the adversary cannot prevent S_u from exceeding $\frac{360 \cdot 2^{i/2}}{\sqrt{n_u}}$ without $\frac{1}{10}$-blocking a constant fraction of the repetitions in the epoch.*

PROOF. By Lemma 1, the adversary's jammed slots can be allocated to the end of the repetition. The expected number of clear slots that a node u hears in a repetition where the adversary is not blocking is at least $\frac{9\, p_c\, S_u\, di^3}{10} \geq \frac{9\, p_c\, S_u\, di^3}{10 e^{2 S_V}} \geq 0.12\, S_u\, di^3$. By Chernoff bounds and a union bound, the probablity that any node u hears less than $0.1\, S_u\, di^3$ clear slots is $e^{-\Omega(i)}$. S_u increases by at least a $2^{0.1/i}$ factor in each such repetition, and over at least $bi^2 - 1$ such repetitions, S_u will increase to $2^{0.1\, b\, i} \geq 2^i$ since $b \geq 10$. Therefore, w.h.p. S_u will exceed $\frac{360 \cdot 2^{i/2}}{\sqrt{n_u}}$ in this epoch. □

To summarize, Lemma 11 implies that a `helper` node u can safely terminate when $S_u \geq \frac{360 \cdot 2^{i/2}}{\sqrt{n_u}}$ because w.h.p., all remaining nodes will also have their status set to `helper`. Lemma 12 guarantees w.h.p. that $S_u \geq \frac{360 \cdot 2^{i/2}}{\sqrt{n_u}}$ will be achieved for every `helper` node. We can now prove the main claims of Theorem 3.

PROOF. First, we analyze the cost of the algorithm when the adversary never $\frac{1}{10}$-blocks a constant fraction of the repetitions in an epoch. For epochs $i \leq \lg n$, Lemma 3 guarantees w.h.p. that $S_u = 16$ so in these epochs the expected cost per node is $O(i^5) = O(\log^5 n)$ (a cost of $O(i^3)$ per repetition for $O(i^2)$ repetitions) regardless of the adversary's behavior. Furthermore, in epochs $i \leq \lg n$ Lemma 4 guarantees that w.h.p. all nodes have status `uninformed` or `informed`. Given this property, Lemma 9 shows that w.h.p. all nodes assume `helper` status by the end of epoch $i = 8 + \lg n$ at the latest. By Lemmas 11 and 12, w.h.p. all `helper` nodes terminate in the next epoch at the latest. At this point, Lemmas 9 implies w.h.p. $S_V \leq \frac{4\sqrt{n}}{2^{i/2}}$ and so, by Lemma 5, w.h.p that each node u has $S_u = O(\sqrt{2^i/n}) = O(1)$. The total cost is therefore $O(\lg^6 n)$ since we have a cost of $O(\lg^5 n)$ per epoch for $\lg n + O(1)$ epochs.

Now consider the case where the adversary does $1/10$-block a constant fraction of the repetitions in some epoch; let ℓ be the last such epoch. The cost to the adversary is $T = \Omega(2^\ell \ell^2)$. The expected cost to a node u, as a function of its final S_u value, is $O(S_u d\, \ell^5) = O(\sqrt{\frac{2^\ell}{n}} \ell^5) = O(\sqrt{\frac{T/\ell^2}{n}} \ell^5) = O(\sqrt{T/n} \lg^4 T)$.

We now turn to the latency. If the adversary never $1/10$-blocks a constant fraction of the repetitions in an epoch, the algorithm will terminate in epoch $\ell \leq \lg n + 8$ with high probability. The latency will be $O(\ell^2 2^\ell) = O(n \log^2 n)$. Let ℓ be the last epoch that the adversary $1/10$-blocks a constant fraction of the repetitions. With high probability all nodes will halt in epoch $\ell + 1$, so the latency will be $O(\ell^2 2^\ell) = O(T)$ which is asymptotically optimal in T given that the adversary can force at least T latency by jamming.

All events analyzed so far hold with high probability, yet a node u may be unlucky and never achieve `informed` or `helper` status before all other nodes have become helpers and halted. Without an alternative halting criterion the expected cost per node would be infinite. Under normal circumstances a node u will become `informed`, then become a `helper` in some epoch j, setting $n_u = 2^j/(S_u)^2$, then halt in an epoch $i \geq j$ when $S_u \geq 360\sqrt{2^i/n_u}$, which is at most $360 \cdot 2^{i/2}$. If any node u finds that $S_u > 360 \cdot 2^{i/2}$, it knows that some correctness/efficiency property has already been violated and can therefore halt. Moreover, to keep u from detecting such a violation, the adversary is forced to 1/10-block most repetitions in an epoch. □

4. LOWER BOUNDS

We present lower bounds for the 1-to-1 and 1-to-n communication problems. Throughout, we abuse terminology somewhat by referring to T as the adversary's *budget* which may be a fixed upper bound on what the adversary actually spends (perhaps the budget is dictated by the amount of energy supplied by the adversary's battery). However, an adversary who has a fixed budget (known or unknown to the nodes) is certainly no stronger than our original adversary, and thus the lower bounds we derive will also hold against our original adversary. Our lower bounds are also strong in that Theorems 2 and 4 assume only a (weak) 1-uniform adaptive adversary that can jam but not send any meaningful messages. Theorem 5 shows that the 1-to-1 communication problem is very sensitive to the power of the adversary. In particular, it is provably more expensive to perform 1-to-1 communication against a 2-uniform adversary that can broadcast messages indistinguishable from Bob's. In [23], the authors give a 1-to-1 communication algorithm in this model with cost $O(T^{\varphi-1})$, where $\varphi = \frac{1+\sqrt{5}}{2}$. Theorem 5 shows that this is asymptotically optimal.

Theorem 2. *Consider a 1-to-1 communication protocol in which Alice sends a message to Bob with probability $1 - \epsilon$ for any constant $\epsilon > 0$. Let A and B be Alice and Bob's empirical costs, and T be the budget of an adaptive, 1-uniform adversary. The adversary can force $\mathrm{E}(A)\,\mathrm{E}(B) > (1 - O(\epsilon))T$. In particular, $\max\{E(A), E(B)\} = \Omega(\sqrt{T})$.*

PROOF. In general, Alice and Bob do not know the adversary's budget T, nor do they know the behavior of the other party. However, for the lower bound proof we assume that T is common knowledge, and that after each time slot both Alice and Bob know the action taken by the other. In particular, when Alice sends and Bob listens, both parties know the message is sent and can halt immediately. These assumptions only make our lower bound stronger.

The adversary commits to the following strategy. Let a_i and b_i be the probability of sending/listening chosen by Alice and Bob just before the ith time slot. The adversary jams if and only if it has not already jammed T slots and $a_i b_i > 1/T$. Alice and Bob can clearly pursue one of two strategies: (i) force the adversary to exhaust her budget in the first T slots by setting $(a_i)_{i \leq T}$ and $(b_i)_{i \leq T}$ sufficiently high, then send the message in the next round by setting $a_{T+1} = b_{T+1} = 1$. Alternatively, they can (ii) choose (a_i) and (b_i) sufficiently low such that the adversary never jams. Furthermore, no mixture of strategies (i) and (ii) can have a strictly lower expected cost than pursuing the best of (i) and (ii). We analyze strategy (ii) then note that the analysis extends to strategy (i).

The rest of the proof is organized as follows. We show that (I) rather than charging 1 unit for sending/listening and 0 for sleeping, it suffices to consider a fractional cost model, (II) without loss of generality, Alice and Bob choose the infinite vectors (a_i) and (b_i) obliviously, (III) it is advantageous to always maximize $a_i b_i$ and to set all coordinates of (a_i) equal and all coordinates of (b_i) equal.

(I) In slot i Alice chooses to send/listen with probability a_i. Rather than charge her 1 if she does, in fact, send/listen we charge

her a_i regardless. By linearity of expectation, the expected cost to Alice and Bob in this fractional model is exactly their cost in the 0/1 cost model. We now need to argue that all non-oblivious algorithms can be made oblivious in the fractional model without increasing their expected costs.

(II) In general an adaptive algorithm for (Alice,Bob) is an infinite decision tree. A node at depth i is labeled with a pair (a_i, b_i) where $a_i b_i \leq 1/T$; it has four children labeled with the empirical behavior of Alice and Bob in the ith time step, i.e., whether they send/listen or sleep. One child is a leaf (Alice sends, Bob listens, and the algorithm halts) whereas the other three children are identified with behaviors at round $i+1$. Since, by (I), Alice's/Bob's cost for the ith round are *independent of the empirical behavior* of Alice/Bob, if they do not terminate after round i they are free to follow *any* of the other three children. Without loss of generality, we can assume they commit to taking the *best* child; the one minimizing $\mathrm{E}(A)\,\mathrm{E}(B)$. In this way all branching can be eliminated without degrading expected costs. In other words, Alice and Bob can be assumed to commit to vectors (a_i) and (b_i) in advance.

(III) Suppose that Bob commits to (b_i). We first show that without loss of generality, Alice's best response is to choose (a_i) such that each $a_i = 0$ or a_i is maximum, that is, such that $a_i b_i = 1/T$. Consider some vector (a_i) that violates the claim, where $0 < a_1 < 1/(b_1 T)$. Let z be the expected cost to Alice in all slots $i \geq 2$, conditioned on not halting after slot 1. Her expected cost over all slots $i \geq 1$ is $a_1 + (1 - a_1 b_1)z = z + a_1(1 - b_1 z)$, which can always be minimized by setting $a_1 = 0$ (if $b_1 z \leq 1$) or $a_1 = 1/(b_1 T)$ (if $b_1 z \geq 1$). The same argument applies to Bob as well. We can clearly ignore any slots in which $a_i = b_i = 0$, so without loss of generality $a_i b_i = 1/T$, for all i. Since Alice and Bob need only succeed with some constant probability we can assume that they halt (unsuccessfully) after some fixed step $t = \Omega(T)$. The probability of failure is $(1 - 1/T)^t < e^{-t/T}$. Define $\hat{a} = (\prod_{i \leq t} a_i)^{1/t}$ and $\hat{b} = (\prod_{i \leq t} b_i)^{1/t}$ to be the geometric means of their probability vectors. Since $(\hat{a}, \hat{a}, \cdots)$ and $(\hat{b}, \hat{b}, \ldots)$ are also valid vectors (that is, they do not induce the adversary to jam any slots) since $\hat{a} \cdot \hat{b} = (\prod_{i \leq t} a_i b_i)^{1/t} = (1/T^t)^{1/t} = 1/T$. Let $p_i = (1 - 1/T)^{i-1}$ be the probability that the algorithm is still running at time step i. Then:

$$\begin{aligned}
\mathrm{E}(A) \cdot \mathrm{E}(B) &= \left(\sum_{i \leq t} a_i p_i\right) \cdot \left(\sum_{i \leq t} b_i p_i\right) \\
&= \sum_{i,j \leq t} a_i b_j \left(1 - \frac{1}{T}\right)^{i+j-2} \\
&\geq \sum_{i,j \leq t} \hat{a}\hat{b} \left(1 - \frac{1}{T}\right)^{i+j-2} \\
&\qquad \text{since the geometric mean} < \text{arithmetic mean} \\
&= \hat{a}\hat{b}\left((1 - O(\exp(-t/T)))T\right)^2 \\
&= \Omega(T) \qquad\qquad \text{since } \hat{a}\hat{b} = 1/T
\end{aligned}$$

The $O(\exp(-t/T))$ reflects the fact that the sums are truncated at t. Thus, the limit of $\mathrm{E}(A)\,\mathrm{E}(B)$ is exactly T as the probability of failure goes to zero.

Suppose Alice and Bob pursue strategy (i) and exhaust the adversary's budget. By the same argument $a_i b_i$ should be infinitesimally larger than $1/T$ to trigger the adversary to jam. Furthermore, $\mathrm{E}(A)\,\mathrm{E}(B)$ is optimized when all a_i (and all b_i) are equal, which implies that $\sum_{i \leq T} a_i > T^{1-\delta}$ and $\sum_{i \leq T} b_i > T^{\delta}$, for some $\delta > 0$. Finally, since our argument makes no assumption on the k-uniformity of the adversary, we assume the weakest adversary (to obtain the strongest lower bound) which is 1-uniform. □

Theorem 4. *Assume a 1-uniform adaptive adversary. Any fair algorithm that achieves 1-to-n communication with probability at least $1/2$ imposes an expected cost per node of $\Omega(\sqrt{T/n})$.*

PROOF. Assume a *fair* algorithm $\mathcal{A}$ (Section 1.3) that achieves 1-to-n broadcast with constant probability at least $\epsilon > 0$ and has an expected cost $g(T)$ for the sender and each receiver. We now design a new algorithm $\mathcal{A}'$ for two players, Alice and Bob, as follows. Bob will simulate those actions taken by the n receivers under $\mathcal{A}$. We must be careful as Bob cannot send and listen simultaneously. This is not an issue for communications between just the n receivers, all of which are simulated by Bob. However, we must address communication between the sender and any of the n receivers. We allocate a pair of slots in $\mathcal{A}'$ for each slot in $\mathcal{A}$. If there was (a probability of) both sending and listening in a slot by receivers under $\mathcal{A}$, then Bob sends in the first slot and listens in the second slot of the corresponding pair. Alice will simulate actions taken by the sender, duplicating the action in each pair of slots.

By construction, $\mathcal{A}'$ solves the 1-to-n broadcast with probability at least that of $\mathcal{A}$; thus, $\mathcal{A}'$ succeeds with probability $\geq 1/2$. Let $E(A)$ denote Alice's expected cost under $\mathcal{A}'$ and note that $E(A) \leq 2 \cdot g(T)$. Let $E(B)$ denote Bob's expected cost under $\mathcal{A}'$ and note that $E(B) \leq n \cdot g(T)$. By Theorem 3, $E(A) \cdot E(B) = \Omega(T)$. Since $E(A) \cdot E(B) \leq 2 \cdot n \cdot g(T)^2$, we must have $g(T) = \Omega(\sqrt{T/n})$. □

Theorem 5. *Consider a 1-to-1 communications protocol such that Alice sends a message to Bob with constant probability of failure, given a 2-uniform adaptive adversary that can spoof messages from Bob. In any such protocol, the expected cost to either Alice or Bob is $\Omega(T^{\varphi - 1})$ where $\varphi = \frac{1+\sqrt{5}}{2}$.*

PROOF. The adversary announces a budget of $\tilde{T}$ and that it will jam Bob (but not Alice, which is possible due to 2-uniformity) if and only if $a_i b_i > 1/\tilde{T}$ and it has not already jammed $\tilde{T}$ slots. It then chooses to either (i) commit to this strategy or (ii) take the place of Bob and *simulate* Bob in scenario (i). In other words, in scenario (ii) the adversary *is* Bob, there is no jamming whatsoever, and there is no correctness criterion for Alice, only a resource-competitive criterion. In scenario (i) the adversary's cost is at most $T = \tilde{T}$ whereas in scenario (ii) its cost is $T = B$, the actual cost incurred by simulating Bob's side of the protocol. Note that since Alice cannot detect when Bob is being jammed, she cannot distinguish scenarios (i) and (ii). Suppose the expected costs incurred by Alice and a (non-adversary) Bob are $O(T^{\alpha})$, where $\alpha < 1$. According to Theorem 2, $\mathrm{E}(A) \cdot \mathrm{E}(B) = \Omega(\tilde{T})$. Let $\delta > 0$ be such that $\mathrm{E}(A) = \Omega(\tilde{T}^{1-\delta})$ and $\mathrm{E}(B) = \Omega(\tilde{T}^{\delta})$. In scenario (ii) (where $T = B$), Alice's expected cost is $\Omega(\tilde{T}^{1-\delta}) = \Omega(T^{(1-\delta)/\delta})$, hence $(1-\delta)/\delta \leq \alpha$. Note that since $\alpha < 1$, it must be that $\delta > 1/2$. In scenario (i) (where $T = \tilde{T}$) Bob's expected cost is $\Omega(T^{\delta})$, hence $\delta \leq \alpha$. Since $(1-\delta)/\delta$ is decreasing in δ, $\max\{(1-\delta)/\delta, \delta\}$ is minimized when $\delta = \frac{1-\delta}{\delta} = \frac{-1+\sqrt{5}}{2} = \varphi - 1$. In other words, either Bob's cost is $\Omega(T^{\alpha})$ in scenario (i) or Alice's cost is $\Omega(T^{\alpha})$ in scenario (ii), where $\alpha = \varphi - 1 > 0.618$. □

5. REFERENCES

[1] Robobees. http://robobees.seas.harvard.edu.

[2] Dan Alistarh, Seth Gilbert, Rachid Guerraoui, Zarko Milosevic, and Calvin Newport. Securing Your Every Bit: Reliable Broadcast in Byzantine Wireless Networks. In *Proceedings of the Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 50–59, 2010.

[3] Ghada Alnifie and Robert Simon. A Multi-Channel Defense Against Jamming Attacks in Wireless Sensor Networks. In *Proceedings of the*

3rd ACM Workshop on QoS and Security for Wireless and Mobile Networks, pages 95–104, 2007.

[4] Nils Aschenbruck, Elmar Gerhards-Padilla, and Peter Martini. Simulative Evaluation of Adaptive Jamming Detection in Wireless Multi-hop Networks. In *International Conference on Distributed Computing Systems Workshops*, pages 213–220, 2010.

[5] Farhana Ashraf, Yih-Chun Hu, and Robin Kravets. Demo: Bankrupting the Jammer. In *Proceedings of the* 9^{th} *International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2011.

[6] Baruch Awerbuch, Andrea Richa, and Christian Scheideler. A Jamming-Resistant MAC Protocol for Single-Hop Wireless Networks. In *Proceedings of the* 27^{th} *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–54, 2008.

[7] Emrah Bayraktaroglu, Christopher King, Xin Liu, Guevara Noubir, Rajmohan Rajaraman, and Bishal Thapa. On the Performance of IEEE 802.11 Under Jamming. In *INFOCOM*, 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 *Proceedings of the Seventeenth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 325–332, 2005.

[9] Vartika Bhandhari, Jonathan Katz, Chiu-Yuen Koo, and Nitin Vaidya. Reliable Broadcast in Radio Networks: The Bounded Collision Case. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 258 – 264, 2006.

[10] Timothy Brown, Jesse James, and Amita Sethi. Jamming and Sensing of Encrypted Wireless Ad Hoc Networks. In *Proceedings of the* 7^{th} *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 120–130, 2006.

[11] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Jamming-Resistant Learning in Wireless Networks. http://arxiv.org/abs/1307.5290, 2013.

[12] Christian Decker, Albert Krohn, Michael Beigl, and Tobias Zimmer. The Particle Computer System. In *Proceedings of the* 4^{th} *International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 443–448, 2005.

[13] Jing Deng, Richard Han, and Shivakant Mishra. Defending Against Path-Based DoS Attacks in Wireless Sensor Networks. In *Proceedings of the 3rd ACM Workshop on Security of Ad Hoc and Sensor Networks*, pages 89–96, 2005.

[14] Shlomi Dolev, Seth Gilbert, Rachid Guerraoui, and Calvin Newport. Gossiping in a Multi-channel Radio Network: An Oblivious Approach to Coping with Malicious Interference. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, pages 208–222, 2007.

[15] Shlomi Dolev, Seth Gilbert, Rachid Guerraoui, and Calvin Newport. Secure Communication over Radio Channels. In *Proceedings of the Symposium on Principles of Distributed Computing (PODC)*, pages 105–114, 2008.

[16] Yuval Emek and Roger Wattenhofer. Frequency Hopping against a Powerful Adversary. In *Proceedings of the* 27^{th} *International Symposium Distributed Computing (DISC)*, pages 329–343, 2013.

[17] Aurélien Francillon and Claude Castelluccia. TinyRNG: A Cryptographic Random Number Generator for Wireless Sensors Network Nodes. In *Proceedings of the* 5^{th} *International Symposium on In Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks and Workshops (WiOpt)*, pages 1–7, 2007.

[18] Seth Gilbert, Rachid Guerraoui, Dariusz Kowalski, and Calvin Newport. Interference-Resilient Information Exchange. In *INFOCOM*, pages 2249–2257, 2009.

[19] Seth Gilbert, Rachid Guerraoui, and Calvin C. Newport. Of Malicious Motes and Suspicious Sensors: On the Efficiency of Malicious Interference in Wireless Networks. In *International Conference On Principles Of Distributed Systems (OPODIS)*, pages 215–229, 2006.

[20] Seth Gilbert, Valerie King, Jared Saia, and Maxwell Young. Resource-Competitive Analysis: A New Perspective on Attack-Resistant Distributed Computing. In *Proceedings of the* 8^{th} *ACM International Workshop on Foundations of Mobile Computing (FOMC)*, 2012.

[21] Seth Gilbert and Maxwell Young. Making Evildoers Pay: Resource-Competitive Broadcast in Sensor Networks. In *Proceedings of the* 31^{th} *Symposium on Principles of Distributed Computing (PODC)*, pages 145–154, 2012.

[22] Chris Karlof, Naveen Sastry, and David Wagner. TinySec: A Link Layer Security Architecture for Wireless Sensor Networks. In *Proceedings of the* 2^{nd} *International Conference on Embedded Networked Sensor Systems*, pages 162–175, 2004.

[23] Valerie King, Jared Saia, and Maxwell Young. Conflict on a Communication Channel. In *Proceedings of the* 30^{th} *Symposium on Principles of Distributed Computing (PODC)*, pages 277–286, 2011.

[24] Yoonmyung Lee, Gyouho Kim, Suyoung Bang, Yejoong Kim, Inhee Lee, Prababl Dutta, Dennis Sylvester, and David Blaauw. A Modular $1mm^3$ Die-Stacked Sensing Platform with Optical Communication and Multi-Modal Energy Harvesting. In *2012 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pages 402–404, 2012.

[25] Xin Liu, Guevara Noubir, Ravi Sundaram, and San Tan. SPREAD: Foiling Smart Jammers Using Multi-Layer Agility. In *INFOCOM*, pages 2536–2540, 2007.

[26] Dominic Meier, Yvonne Anne Pignolet, Stefan Schmid, and Roger Wattenhofer. Speed Dating Despite Jammers. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 1–14, 2009.

[27] Rajeev Motwani and Prbhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[28] Vishnu Navda, Aniruddha Bohra, Samrat Ganguly, and Dan Rubenstein. Using Channel Hopping to Increase 802.11 Resilience to Jamming Attacks. In *INFOCOM*, pages 2526–2530, 2007.

[29] Adrian Ogierman, Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive MAC under Adversarial SINR. http://arxiv.org/abs/1307.7231, 2013.

[30] Andrzej Pelc and David Peleg. Feasibility and Complexity of Broadcasting with Random Transmission Failures. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 334–341, 2005.

[31] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *IPSN*, 2005.

[32] Christina Pöpper, Mario Strasser, and Srdjan Čapkun. Jamming-Resistant Broadcast Communication Without Shared Keys. In *Proceedings of the* 18^{th} *USENIX Security Symposium*, 2009.

[33] Iyappan Ramachandran and Sumt Roy. Clear Channel Assessment in Energy-Constrained Wideband Wireless Networks. *IEEE Wireless Communications*, 14(3):70–78, 2007.

[34] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. A Jamming-Resistant MAC Protocol for Multi-Hop Wireless Networks. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, pages 179–193, 2010.

[35] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive and Fair Medium Access Despite Reactive Jamming. In *Proceedings of the* 31^{st} *International Conference on Distributed Computing Systems (ICDCS)*, pages 507–516, 2011.

[36] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive and Fair Throughput for Co-Existing Networks Under Adversarial Interference. In *Proceedings of the* 31^{st} *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 291–300, 2012.

[37] Deva Seetharam and Sokwoo Rhee. An Efficient Pseudorandom Number Generator for Low-Power Sensor Networks. In 29^{th} *Annual IEEE International Conference on Local Computer Networks*, pages 560–562, 2004.

[38] SPECKNET. `http://www.specknet.org/`.

[39] R. Watro, D. Kong, S. Cuti, C. Gariner, C. Lynn, and P. Kruus. TinyPK: Securing Sensor Networks with Public Key Technology . In *SASN*, pages 59–64, 2004.

[40] Wenyuan Xu, Ke Ma, Wade Trappe, and Yanyong Zhang. Jamming Sensor Networks: Attack and Defense Strategies. IEEE *Networks*, 20(3):41–47, 2006.