Talk is Cheap(er): Mitigating DoS and Byzantine Attacks in Sensor Networks

David R. Cheriton School of Computer Science

Technical Report CS-2010-14

Valerie King[†], Jared Saia[‡], Maxwell Young[*]
[*]Department of Computer Science, University of Victoria, BC, Canada, val@cs.uvic.ca
[†]Department of Computer Science, University of New Mexico, NM, USA, saia@cs.unm.edu
[‡]David R. Cheriton School of Computer Science, University of Waterloo, ON, Canada, m22young@cs.uwaterloo.ca

"The power and prestige of Byzantium were founded above all on its gold..."
-*History of the Byzantine State* by George Ostrogorsky

***Abstract*—Sensor networks are extremely vulnerable to denial-of-service (DoS) attacks due to their shared communication medium and the constrained power supply of the devices. We examine the scenario where one player attempts to send a message m to another player in a single-channel network. Assume an adversary that jams for T instances where T is *unknown* to either player and where cost is measured in terms of energy expenditure. We give a protocol guaranteeing delivery of m while the ratio of either player's expected cost to the adversary's cost is $O(T^{\varphi-2}) \approx O(T^{-0.382})$ where $\varphi = \frac{1+\sqrt{5}}{2}$ is the golden ratio.**

Our result implies that to prevent communication of m, the adversary incurs an asymptotically higher cost than the expected cost incurred by either correct player. This property holds even when the receiver suffers a Byzantine fault. We extend our analysis to multiple receivers and later, to a multi-hop setting where both senders and receivers can suffer Byzantine faults. Our work pertains to a wide variety of adversaries that are energy constrained. Notably, we can tolerate an *adaptive* adversary who launches attacks using knowledge of the past actions of all players. Moreover, in networks with a sufficient amount of communication traffic, we can tolerate a *reactive* adversary who may detect a transmission in the current time slot and then decide to jam. Finally, we apply our result to the fundamental network communication problem of reliable broadcast which deals with conveying a message from one player to all other players in the network. In a popular sensor network grid model, we give a reliable broadcast protocol that, in comparison to previous work in the grid, offers improved resilience to DoS attacks.

I. Introduction

In addition to traditional network security challenges, the shared communication medium of sensor networks renders them vulnerable to jamming attacks [1]. A jamming attack occurs when an attacker transmits concurrently with another (possibly legitimate) transmission such that communication is disrupted within the area of interference. Consequently, this type of behavior represents a simple denial-of-service (DoS) attack that threatens the availability of sensor networks [2]. While this can be viewed as a competition for the communication medium, the *de facto* resource being consumed is energy. In the energy-critical domain of sensor networks, such attacks can rapidly deplete the onboard energy supply of a device.

Wireless network cards typically offer states such as *off*, *sleep*, *receive* and *transmit*. While the sleep state requires negligible power, remarkably the cost of the transmit and receive states are roughly equivalent. With the recent Telos motes the radio transmit and receive costs are 38mW and 35mW, respectively, while the sleep state cost is 15μW [3]. Therefore, the cost of the transmit/receive state exceeds that of the sleep state by a factor greater than 2000. Similar relationships hold for the MICAz and MICA2 motes [4], [5]. Therefore, two considerations inform our protocol design: (1) minimizing the time spent outside the sleep state improves energy-efficiency and (2) we must account for the costly receive state.

In this work, we focus on mitigation by making DoS attacks more expensive for an adversary to launch relative to the cost incurred by correct players. Critical to our investigation is the following 2-Player Scenario: Player A (a sender) wishes to send a message m to Player B (a receiver). However, there exists a jamming adversary who aims to prevent transmission of m. The adversary may also control the receiver if it suffers a Byzantine fault. This necessitates a delicate balancing of cost between the sender and a correct receiver. The sender must not be manipulated into depleting its energy supply by a Byzantine receiver via repeated retransmission requests for m. Conversely, a correct receiver should not unfairly bear the brunt of the cost. Of course, it is not known *a priori* whether a receiver is correct or Byzantine. Under this challenging scenario, we give a communication protocol which guarantees that the ratio of either correct player's expected cost to the adversary's cost is $O(T^{\varphi-2})$ where φ is the golden ratio.

Therefore, the cost incurred by the adversary exceeds the expected cost of a correct player. In the energy-scarce domain of sensor networks, this greatly undermines the efficacy of DoS attacks. We generalize our result to a multi-player setting and apply this to the problem of reliable broadcast in a multi-hop setting where a bounded number of senders and receivers suffer Byzantine faults. In the popular grid model (see [6]–[12]) we obtain a reliable broadcast protocol with increased

resilience to DoS attacks over previous results in the grid.

II. Our Model and Assumptions

Network Model: Various aspects of theoretical models for sensor networks have been critiqued [13]. Consequently, our assumptions are grounded in the empirical literature. There is a single channel, and in our single-hop settings (Sections IV&V), we assume a time division multiple access (TDMA)-like medium access control (MAC) protocol; that is, a time-slotted network. For example, the well-known LEACH [14] protocol for sensor networks is TDMA-based. Secure synchronization has been demonstrated for the MICA2 devices [15]. Only for ease of exposition is a *global* broadcast schedule assumed in our final multi-hop scenario (Section VI); however, this is avoidable if nodes maintain multiple schedules as with S-MAC [16]. Even then, global scheduling (and the resulting energy-efficiency advantages) has been clearly demonstrated by experimental work in [17].

The adversary or a player is said to be *active* in a time slot (or just a 'slot') s if it is either in the send (transmit) or receive state during s. As with the Telos and MICA models, each device possesses a single transceiver and only one state is possible per slot. The cost to a player for sending and receiving (or detecting a collision) for one slot is normalized to $\alpha_{cor} = 1$, but our protocols can be easily modified to address the small differences between the send/receive state costs that exist in practice. The adversary is active for T slots at a cost of $\alpha_{adv} = 1$ per slot and T is unknown to either player.

Each player in the receive state can detect whether a collision has occurred in a slot. Clear channel assessment (which subsumes carrier sensing) is a common feature on devices [18] and considered practical under the IEEE 802.11 standard [19]. Collisions are only detectable by the receiver [2]. When a collision occurs, we assume the transmission for that slot is lost. The absence of channel activity cannot be forged; however, the adversary may *forge a collision*. That is, even if no correct players are transmitting, adversarial nodes may transmit simultaneously (or replay interference) in order to cause receiving players to detect a collision; forging costs 1.

Energy Competitiveness: We introduce the notion of *energy competitiveness*. Let A and B respectively denote the expected number of slots for which a correct Player A and correct Player B are active before successfully terminating a protocol for the 2-Player Scenario. Let T be the total number of slots for which the adversary is active. A protocol is $O(T^{-c})$*-energy competitive* if for any $T > 0$, $\frac{A}{T} = O(T^{-c})$ and $\frac{B}{T} = O(T^{-c})$ for a constant $c < 1$. We refer to A/T, B/T as *energy ratios*.

Why is energy competitiveness useful? First, equal asymptotic costs (in terms of energy) apply to each player which is useful in homogenous sensor networks where devices have similar energy constraints. Second, central to a DoS attack is the notion that an adversary can force a player to incur a higher cost; call this the *unfairness property*. If this property holds for all T, DoS attacks are especially devastating. To see this, consider when players incur 5x the cost of the adversary under a DoS attack. For a T of moderate size, a player's absolute cost is tolerable. Throughput is reduced, but the players may still communicate several times prior to exhausting their energy supplies. In contrast, for large T, nodes may have insufficient energy to communicate even once (i.e. zero throughput).

If an energy-competitive protocol is used, then the unfairness property fails for sufficiently large T since the energy ratio is $O(T^{-c})$. At this point, the correct players enjoy the advantage. Since our analysis is asymptotic, the players may incur a higher cost than the adversary for limited T. However, as discussed above, this is acceptable so long as this absolute cost is tolerable and we show this is true for our protocols.

Las Vegas Property: The protocol for the 2-Player Scenario, and the subsequent multi-player extension with n players, is Las Vegas. That is, if the players have sufficient energy to execute the protocol until the termination condition, then m is successfully sent from the sender to the correct receiver(s) with probability 1; there is no probability or error.

The Las Vegas property is valuable in multi-hop sensor networks for the following reason. Let n be the number of devices within transmitting distance of a device, and let N be the total number of devices in the network. Monte Carlo communication protocols that succeed with high probability in n are possible. However, typically, $n \ll N$ and messages will traverse multiple hops; consider $\Omega(N)$ hops. *Then even if the failure probability for a single hop is exponentially small in n, the probability that some hop fails can still be very large in terms of N*. For example, if N is exponential in n and the failure probability for each hop is $O(n^{-c})$ for some constant $c > 0$, or even $O(2^{-n})$, then communication fails along the chain with at least constant probability.

Alternatively, we might achieve protocols that succeed with high probability in N. For small networks this is reasonable. However, in large networks, N may not be known *a priori*. Furthermore, achieving a high probability guarantee in N typically involves $\Omega(\log N)$ operations which, for large N, may be too costly. Therefore, by devising Las Vegas protocols, we avoid assumptions that are problematic given that $n \ll N$.

Cryptographic Authentication: Several results show how cryptographic authentication can be implemented in sensor networks [1], [20]–[23]. Therefore, it is important to consider the impact of cryptographic authentication and we assume it here. However, the adversary may capture and subvert a limited number of players; these players are said to suffer a Byzantine fault and are controlled by the adversary. Such attacks are well-known to the research community [1], [2]. However, subverting nodes is assumed to be a costly operation; for example, tamper-resistant approaches may be applied to sensor devices and techniques for evading physical attacks are known (see [1] and references therein). Therefore, we assume the number of such instances is limited. In the 2-Player Scenario and the n-player extension, the Multi-player Scenario, we assume only receivers can suffer Byzantine faults since communication is doomed with a Byzantine sender. Furthermore, energy-competitive communication is non-trivial even with authentication due to jamming and the

forging of collisions; this is discussed in Section IV-A.

Types of Adversaries: The adversary has full knowledge of past actions by correct players. This allows for *adaptive attacks* whereby the adversary may alter its behavior based on observations it has collected over time. Furthermore, the adversary is *reactive*: in any slot, the adversary may detect a transmission and then jam this transmission. The effectiveness of reactive jamming has been shown experimentally [24].

Our adversary captures the worst-case disruption of transmissions due to non-malicious failures such as software errors or accidental deviations from a global schedule. Our adversary also models challenging malicious behavior. Various jamming strategies feature prominently in the literature such as constant, reactive, and random jamming [1], [24]. Our energy-competitive protocols apply to all such adversaries with a bounded energy supply. The is particularly pertinent to sensor networks where devices are typically battery powered.

A. Our Main Contributions

Throughout, let $\varphi = (1+\sqrt{5})/2$ denote the golden ratio. Our three main analytical contributions are given below.

Theorem 1. COMPETITIVE COMMUNICATION *has the following properties:*

- *If the adversary is never active ($T = 0$), then each player is active for $O(1)$ slots in expectation.*
- *Otherwise, if the adversary is active for $T > 0$ slots, then* COMPETITIVE COMMUNICATION *is $O(T^{\varphi-2})$-energy competitive where $\varphi - 2 \approx -0.382$.*
- *If both players are correct and execute* COMPETITIVE COMMUNICATION *until their respective termination conditions are met, then transmission of m is guaranteed.*

For the more general single-hop setting with a sender and n receivers, we provide the protocol MULTI-PLAYER COMPETITIVE COMMUNICATION (MPCC) and obtain the following:

Theorem 2. MPCC *has the following properties:*

- *If all receivers are correct, then the expected cost to the sender and each receiver is $O(\ln^2 n)$.*
- *If corrupt receivers are active for a total of $T > 0$, then* MPCC *is $O(\max\{\ln^2 n, T^{\varphi-1}\}/T)$-energy competitive.*
- MPCC *terminates within $O(T^{\varphi/(\varphi-1)})$ slots assuming $\ln n = O(T)$. Each correct player with sufficient energy terminates successfully at this point.*

Finally, we address reliable broadcast in the multi-hop sensor network grid model where devices are referred to as nodes. By leveraging results in [7], we obtain the following result.

Theorem 3. *Assume $t < (r/2)(2r+1)$ nodes in any $(2r+1) \times (2r+1)$ square centered about a correct node p in the grid are Byzantine and used by the adversary to disrupt p's communications for $\beta \leq B_0$ slots. Let $C = \{$Nodes q at $(x,y)| -r \leq x \leq r \ \wedge \ y \geq 0\}$ be a corridor of nodes in this network. There exists a protocol with the following properties:*

- *If each correct node p knows B_0 and has sufficient energy to execute the protocol, then reliable broadcast along the corridor C is guaranteed.*
- *If $\beta = O(r^2 \ln^{2/(2-\varphi)} r)$, each correct node incurs a cost (not an energy ratio) of $O(r^4 \ln^{2/(2-\varphi)} r)$ active slots.*
- *If $\beta = \omega(r^2 \ln^{2/(2-\varphi)} r)$, the energy ratio is $o(1)$ and decreases as $O\left(\frac{r^{2(2-\varphi)} \cdot \ln^2 r}{\beta^{2-\varphi}}\right) \approx O\left(\frac{r^{0.764} \cdot \ln^2 r}{\beta^{0.382}}\right)$.*

While this last result applies to a single corridor, it is sufficient to prove reliable broadcast since the network can be covered piece-wise by such corridors (see Section VI). We note that switching from the sleep state to an active state incurs a cost. However, in our protocols, the number of state switches is limited by the number of active slots and, therefore, our asymptotic analysis holds. We also examine the constants in our asymptotic analysis and investigate practical values for α_{cor} and α_{adv} in Section IV-C. Finally, we believe our notion of energy competitiveness, and its application to reactive adversaries, is novel and relevant to designing adversarial fault-tolerant algorithms for sensor networks.

III. RELATED WORK

Jamming Adversaries: Several works demonstrate that wireless devices are vulnerable to adversarial jamming [24]–[27]. Defenses include spread spectrum techniques, frequency or channel hopping, and mapping with rerouting (see [2], [28]–[30] and references therein).

There are a number of theoretical models for adversarial jamming. Gilbert *et al.* [31] examine communication between two players with collision detection in a time-slotted network against an adversary who interferes with an unknown number of transmissions. Cryptographic authentication is not assumed. The adversary cannot forge the absence of channel activity; we also use this property. The authors derive bounds on (1) the *jamming gain*, which is defined as the amount of energy used to prevent communication relative to the amount of energy used by continuous jamming [32], and (2) *disruption-free complexity*, which measures how long the adversary may disrupt a protocol without broadcasting. We note that while jamming gain measures the relative energy cost of two adversarial strategies, it does not explicitly address the cost to the adversary relative to the correct players as we do here. Pelc and Peleg [33] examine an adversary that randomly corrupts messages; we do not require the adversary to behave randomly. Awerbuch *et al.* [34] give a jamming-resistant MAC protocol in a single-hop network with an adaptive, but non-reactive, adversary. Richa *et al* [35] significantly extend this work to multi-hop networks. Dolev *et al.* [36] address a variant of the gossiping problem when multiple channels are jammed. Gilbert *et al.* [37] derive bounds on the time required for information exchange when a reactive adversary jams multiple channels. Meier *et al.* [38] examine the delay introduced by a jamming adversary for the problem of node discovery, again in a multi-channel setting. Dolev *et al.* [39] address secure communication using multiple channels with a non-reactive adversary. In these previous works, the impact of receive state costs is not explicitly addressed and there is no notion of energy competitiveness.

Reliable Broadcast: Reliable broadcast has been extensively studied in the grid model [6]–[10], [40]–[42]. Energy-efficient reliable broadcast is achieved by King *et al.* [10], [40] but the authors do not consider a jamming adversary. With a reactive jamming adversary, Bhandhari [43] show reliable broadcast is possible when the amount of jamming is bounded and known *a priori*; however, correct nodes must expend considerably more energy than the adversary. Progress toward fewer broadcasts is made by Bertier *et al.* [11]; however, each node must spend a significant amount of time in the costly listening state. Alistarh *et al.* [12] assume collision detection and achieve non-cryptographic authenticated reliable broadcast. They apply their result to the grid with a jamming adversary; however, nodes must again spend a significant time in the receive state.

IV. Our Cost Competitive Protocol

When a node is not active, it is assumed to be in the energy-efficient sleep state. We discuss how (1) randomness foils an adaptive adversary and (2) energy competitiveness is possible against a reactive adversary given sufficient network traffic.

Tolerating an Adaptive Adversary: A simple but critical feature of our protocols is: the probability that a player is active in one slot is independent from the probability that the player is active in another slot. Therefore, knowing that a player was active for k slots in the past conveys no information about future activity. Believing otherwise is the trap of the well-known"Gambler's Fallacy" [44].

In contrast, we could have specified that a player be active for a fixed k number of slots chosen probabilistically in each epoch. However, a adaptive adversary knows past information. Therefore, within an epoch, if a player is active for $k' < k$ slots, the adversary knows that the player will be active for $k - k'$ more slots in the remainder of the epoch. This information can allow the adversary to jam more effectively. Instead, by having a player be active independently in each slot, knowledge of the past cannot help the adversary.

Tolerating a Reactive Adversary: If the adversary knows, free of cost, when m is sent then energy competitiveness is impossible. However, a reactive adversary must detect transmissions. Such detection is referred to as *clear channel assessment* (CCA) [18]. Consider two cases with respect to CCA. In the first case, the adversary wishes to detect channel activity, but does not care about the traffic content. Detection is performed via the radio chip using the *received signal strength indicator* (RSSI) [45]. If the RSSI value is below the clear channel threshold, then the channel is assumed to be clear [46]. RSSI incurs a cost on the order of 10^{-6} W – roughly three orders of magnitude smaller than the send or receive costs. Therefore, an adversary can detect channel activity at essentially zero-cost; *however, the trade-off is that this detection is indiscriminate.*

In the second case, the adversary desires knowledge of the traffic content, such as source/destination information. Now the adversary must enter the receive state. Although the time spent in the receive state may vary, the cost is substantial being on the order of 10^{-3} W. Denote this cost by ρ_{adv}.

COMPETITIVE COMMUNICATION(m)

1: $i \leftarrow 2$
2: $A\_terminate \leftarrow$ `false`
3: **while** ($A\_terminate ==$ `false`) **do**
4: *Epoch 1:*
5: **for** slot $s = 1$ to 2^{ci} **do**
6: • Player A sends m with probability $\frac{2}{2^i}$.
7: **if** (Player B has not received m) **then**
8: • Player B listens with probability $\frac{2}{2^{(c-1)i}}$.
Epoch 2:
9: **for** slot $s = 1$ to 2^i **do**
10: **if** (Player B has not received m) **then**
11: • Player B sends `req`.
12: • Player A listens with probability $\frac{4}{2^i}$.
13: **if** (Player A listens and detects no collision or `req`) **then**
14: • $A\_terminate \leftarrow$ `true`
15: $i \leftarrow i + 1$

Fig. 1. Pseudocode for COMPETITIVE COMMUNICATION.

Conversely, the cost to a correct node, ρ_{cor}, is comparable for small payloads; we assume large content can be broken into small payloads. For a small payload, the total message size is small (assuming small headers/footers [47], [48]) and a correct player also incurs a cost on the order of 10^{-3} W. Therefore, while $\rho_{adv} < \rho_{cor}$, these values are comparable. We can re-normalize the costs in Section II by ρ_{adv} such that the adversary's cost is again $\alpha_{adv} = 1$, while the correct players' cost, α_{cor}, is a constant multiple of ρ_{adv} (see Section IV-C).

Which case matters? If an adversary seeks to disrupt *any* traffic, then the first case applies and the only cost is due to jamming. However, there are many situations where this behavior quickly disables the adversary. For example, consider a terrain with multiple sensor networks where the adversary wishes to target only a select few collecting critical data. Or there is a single network that executes several distributed applications and the adversary wishes to interfere with a critical few. If the adversary jams indiscriminately, it will quickly exhaust its energy supply in disrupting mostly non-critical transmissions. The more challenging scenario occurs when the adversary first identifies transmissions and then jams only those it is targeting; this is the second case. Using the re-normalization above, our asymptotic analysis for energy competitiveness holds; this is investigated further in Section IV-C.

Finally, we note that the conclusion of our argument aligns with claims put forth in empirical results on reactive jamming; that is, such behavior does not necessarily result in a more energy-efficient attack because the adversary must still be listening to the channel for broadcasts prior to committing itself to their disruption [24].

A. Protocol Overview

Figure IV gives the pseudocode for our 2-PLAYER SCENARIO protocol called COMPETITIVE COMMUNICATION. Each round $i \geq 2$ consists of 2 epochs and c is a constant to be determined later. We summarize a round i of the protocol:

• *Epoch 1:* This epoch consists of 2^{ci} slots. In each slot j, Player A sends m with probability $\frac{2}{2^i}$ for an expected total of $2^{(c-1)i+1}$ slots. In each slot, Player B listens with probability $\frac{2}{2^{(c-1)i}}$ for an expected total of 2^{i+1} slots.

• *Epoch 2:* This epoch consists of 2^i slots. If Player B has not received m, then Player B sends a request for retransmission, `req`, for all 2^i slots. Player A listens in each slot with probability $4/2^i$ for an expected total 4 slots.

Termination Conditions: Termination conditions are important because the adversary cannot be allowed to keep either player executing in perpetuity while simultaneously forcing them to incur a *higher* cost. Player B terminates the protocol upon receiving m. Due to authentication, messages cannot be spoofed or modified; therefore, this termination condition suffices. Player A terminates if it listens to a slot in Epoch 2 with neither a collision nor a `req`. Since the adversary cannot forge the absence of channel activity, this condition suffices. In other words, Player A continues into the next round if (1) Player A listens to zero slots or (2) all slots listened to by Player A in Epoch 2 contain a collision or a `req`. We highlight the two situations where this condition is met:

• *Situation I:* Player B is correct and has not received m.

• *Situation II:* Player B is Byzantine and sends `req`s, *or* Player B is correct and terminated and the adversary forges collisions to trick Player A into thinking a valid `req` was jammed.

Situation I occurs prior to the successful delivery of m possibly due to jamming attacks in Epoch 1. Situation II addresses an attack that can be employed in Epoch 2 after the delivery of m. The adversary can deplete Player A's energy by making it appear as if Player B repeatedly did not receive m and is requesting a retransmission. This attack affects Player A only. Also, if Player B is correct, the attack is only effective once m is received. If a correct Player B has not received m, a `req` will be issued anyway and the attack accomplishes nothing.

B. Analyzing Cost-Competitiveness

Define a *jamming-1* round as a round where the adversary jams at least half of the slots in Epoch 1; otherwise, it is a non-jamming-1 round. Similarly, a *jamming-2* round is a round where the adversary jams or forges collisions in at least half the slots in Epoch 2. Regarding the constant c: clearly $c > 1$ or Line 8 of our protocol is nonsensical. Also, if $c \geq 2$, then the expected cost to Player A is at least as much as the cost to Player B. In such a case, the cost to the adversary for a Situation II attack is less than the expected cost to Player A since a Byzantine receiver can sleep through Epoch 1. As discussed above, we must avoid this in Situation II. Therefore, we have $1 < c < 2$. Throughout, assume ceilings on the number of active slots of a player if it is not an integer.

Lemma 1. *Consider a round of* COMPETITIVE COMMUNICATION *that is not jamming-1 and where neither player has terminated. The probability that Player B does not receive the message from Player A is less than e^{-2}.*

Proof: Let $s = 2^{ci}$ be the number of slots in Epoch 1. Let p_A be the probability that Player A sends in a particular slot. Let p_B be the probability that Player B listens in a particular slot. Let $X_j = 1$ if the message is not delivered from Player A to Player B in the j^{th} slot. Then $Pr[$ m is not delivered in Epoch 1$]$=$Pr[X_1X_2\cdots X_s = 1]$=$Pr[X_s = 1 \mid X_1X_2\cdots X_{s-1} = 1] \cdot \prod_{i=1}^{s-1} Pr[X_i]$. Let $q_j = 1$ if the adversary does not jam in slot j; otherwise, let $q_j = 0$. The value of q_j can be selected arbitrarily by the adversary. Then $Pr[X_1X_2\cdots X_s = 1] = 1 - p_Ap_Bq_i$ and substituting for each conditional probability, we have $Pr[X_1X_2\cdots X_s] = (1 - p_Ap_Bq_1)\cdots(1 - p_Ap_Bq_s) = \prod_{j=1}^{s}(1 - p_Ap_Bq_j) \leq e^{-p_Ap_B\sum_{j=1}^{s} q_j} < e^{-2}$ since $p_Ap_B\sum_{j=1}^{s} q_j > (2/2^i)(2/2^{(c-1)i})(s/2) = (2/2^i)(2/2^{(c-1)i})(2^{ci}/2) = 2$ since the round is not jamming-1 and so the adversary jams less than $s/2$ slots. ■

We next consider the simplest case where there is never a jamming-1 round and there are no Situation II attacks:

Lemma 2. *Assume that Player B is correct and there are no jamming-1 rounds or Situation II attacks. If the adversary is active for $T > 0$ slots, then the cost ratio is $O(1/T)$. Otherwise, $T = 0$, the expected cost to each player is $O(1)$.*

Proof: Using Lemma 1, the expected cost to Player A is at most $\sum_{i=2}^{\infty} e^{-2(i-2)} \cdot (2 \cdot 2^{(c-1)i} + 4) \leq \sum_{i=2}^{\infty}(e^{5-i} + 4 \cdot e^{-2(i-1)}) = O(1)$. Similarly, the expected cost to Player B is at most $\sum_{i=2}^{\infty} e^{-2(i-2)} \cdot (2^{i+1} + 2^i) \leq \sum_{i=2}^{\infty}(e^{5-i} + e^{4-i})$ $= O(1)$. The result follows for $T > 0$. ■

We next consider the cost to Player A prior to the successful transmission of m. Note that this does not necessarily yield the cost to Player A over the entire protocol:

Lemma 3. *Assume there is at least one jamming-1 round in round i. Then, prior to successful delivery of m, the expected cost to Player A is $O(2^{(c-1)i})$ and, therefore, the cost ratio of Player A before the message is delivered is $O(T^{\frac{c-1}{c}}/T)$.*

Proof: Let i be the last jamming-1 round. Then the adversary has been active for at least $2^{ci}/2$ slots. Using Lemma 1, the expected cost to Player A prior to m being delivered is $O(2^{(c-1)i}) + \sum_{k=1}^{\infty} e^{-2(k-1)} \cdot (2 \cdot 2^{(c-1)(i+k)} + 4) = O(2^{(c-1)i})$ since $c < 2$. Therefore, the energy ratio is $O(2^{(c-1)i}/2^{ci})$ and, letting $T = 2^{ci}$, this is $O(T^{(c-1)/c})$. ■

Now consider when Situation II attacks may occur:

Lemma 4. *Assume that Player B has received m by round i and that round i is non-jamming-2. Then the probability that Player A retransmits m in round $i + 1$ is less than e^{-2}.*

Proof: Let $s = 2^i$ be the number of slots in Epoch 2 and let $p = 4/2^i$ be the probability that Player A listens in a slot. For slot j, define X_j such that $X_j = 1$ if Player A does not terminate. Then $Pr[$ Player A retransmits m in round $i + 1] = Pr[X_1X_2\cdots X_s = 1]$. Let $q_j = 1$ if the adversary does not jam in slot j; otherwise, let $q_j = 0$. The q_j values are determined arbitrarily by the adversary. Since Player A terminates if and only if it listens and does not detect any activity, then $Pr[X_j = 1] = (1 - pq_j)$. Therefore, $Pr[X_1X_2\cdots X_s = 1] \leq e^{-p\sum_{j=1}^{s} q_j} < e^{-2}$. ■

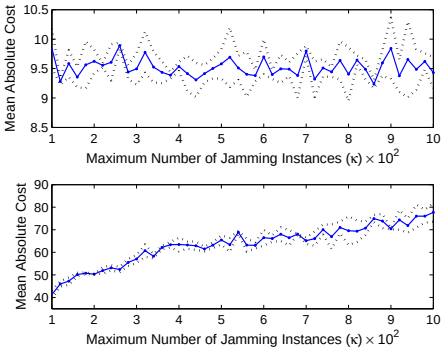

Fig. 2. Mean absolute cost (maximum of either player) for a random jammer with $p_j = 0.5$ (top) and 0.9 (bottom). Dotted lines signify 95% confidence intervals.

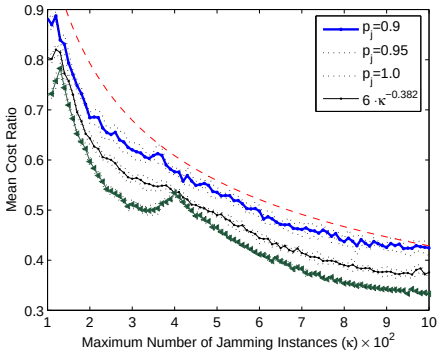

Fig. 3. Mean energy ratio (maximum of either player) for a random jammer with $p_j = 0.9$, 0.95 and 1.0. Dashed line is $6 \cdot \kappa^{\varphi-2}$. Dotted lines signify 95% confidence intervals.

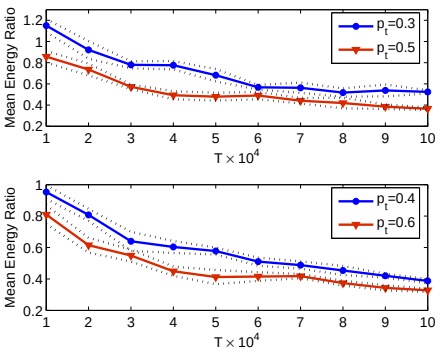

Fig. 4. Mean energy ratio (maximum of either player) for a reactive jammer with $p_t = 0.3, 0.4, 0.5$ and 0.6, separated for clarity. Dotted lines signify 95% confidence intervals.

Lemma 5. *Assume there is at least one jamming-1 round. The energy ratio of Player A is $O(T^{c-1}/T)$ and the energy ratio for a correct Player B is $O(T^{\frac{1}{c}}/T)$.*

Proof: Let i be the last round which is jamming-1. Let $j \geq i$ be the last round which is jamming-2; if no such jamming-2 round exists, then assume $j = 0$. Either way, the cost to the adversary is $\Omega(2^{ci} + 2^j)$.

By Lemma 3, the expected cost to Player A prior to successfully transmitting m is $O(2^{(c-1)i})$. Using Lemma 4, the expected cost to Player A prior to terminating is $O(2^{(c-1)i}) + O(2^{(c-1)j}) + \sum_{k=1}^{\infty} e^{-2(k-1)} \cdot (2 \cdot 2^{(c-1)(j+k)} + 4) = O(2^{(c-1)j})$ since $c < 2$. Therefore, the total expected cost to Player A is $O(2^{(c-1)i} + 2^{(c-1)j})$. The energy ratio for Player A is highest when $i = 0$ which yields $O(T^{(c-1)}/T)$. Using Lemma 3, the overall energy ratio for Player A is $\max\{O(T^{(c-1)/c}/T), O(T^{c-1}/T)\} = O(T^{c-1}/T)$.

Finally, assume Player B is correct. Since the last jamming-1 round occurred in round i, Player B is active for $O(2^i)$ slots. Using Lemma 1, Player B's expected cost prior to receiving m is $O(2^i) + \sum_{k=1}^{\infty} e^{-2(k-1)} \cdot (2 \cdot 2^{i+k} + 2^{i+k}) = O(2^i)$. Since Player B terminates upon receiving m, the energy ratio is maximized when $j = 0$ which yields $O(T^{1/c}/T)$. ■

We now give the proof for Theorem 1 stated in Section II-A:

Proof of Theorem 1: We first examine the worst case energy ratios for both players over all values for i and j:

Player A: First, Lemma 3 addresses the case where no Situation II attacks occur, and here the energy ratio is $O(T^{(c-1)/c}/T)$. Second, Lemma 5 addresses the case where Situation II attacks occur. Here the cost to the adversary is $\Omega(2^{ci} + 2^j)$ while the expected cost to Player A is $O(2^{(c-1)i} + 2^{(c-1)j})$. The worst energy ratio arises with $i = 0$ with an energy ratio of $O(T^{c-1}/T)$.

Player B: For Player B, the expected cost prior to the successful transmission of m is $O(2^i)$. Overall, the cost to the adversary is $T = \Omega(2^{ci} + 2^j)$. The energy ratio is maximized when $j = 0$, which yields an energy ratio of $O(T^{1/c}/T)$.

Therefore, the exponents of interest which control the energy ratios are $(c-1)/c$, $c-1$, and $1/c$ and the adversary knows c before picking i and j. The value of c that should be chosen must minimize $\max\{(c-1)/c, c-1, 1/c, \}$. Since $1 < c < 2$, we have $1/c > (c-1)/c$. Therefore, we solve for c in $c - 1 = 1/c$, this gives $c = (1 + \sqrt{5})/2$ which is the golden ratio.

If $T \geq 1$, then by setting $c = \varphi$ in Lemma 2 and the above argument, the energy ratio is $\max\{O(1/T), O(T^{\varphi-2})\} = O(T^{\varphi-2}) \approx O(T^{-0.382})$. If $T = 0$, then by plugging $c = \varphi$ into Lemma 2, the expected cost to each player is $O(1)$. ■

Discussion: Note that there is an $O(1)$ up-front cost per execution of the protocol when there are no jamming-1 or jamming-2 attacks. This is the price for communication in the presence of a powerful adversary, even if that adversary is not always active. In exchange, an adversary incurs a penalty that increases commensurate with the amount of disruption it causes. Therefore, the unfairness property does not hold.

Finally, might players share a secret schedule? This would reduce the active costs in Theorem 1 where neither player knows if the other is active with certainty. Unfortunately, such a schedule becomes known to the adversary if the receiver suffers a Byzantine fault which invalidates any improved analysis and allows Player A to be manipulated.

C. Numerical Results

We further evaluate COMPETITIVE COMMUNICATION since it is a building block for our next two protocols. Our aims are modest, we: (1) show small constants in our asymptotic analysis and (2) use back-of-the-envelope calculations to estimate the behavior for pessimistic α_{cor} and α_{adv} values. The impact of radio irregularity, weather and terrain should be evaluated with a prototype; however, this is outside the scope of our paper. This is left to future work and we note that such detrimental factors affect other DoS-resistant proposals, and there exist many results on mitigating their effect.

Parameters: The Telos mote is powered by two AA batteries used in series which yields roughly 18000 J assuming a combined 1.8V cut-off voltage (see [3]). The send/receive costs are 35 mW/38 mW; however, we use the total operational cost of 41 mW [3]. Let $|m|$ denote the packet (message) size. Headers and footers of 2 bytes each are possible (see [47], [48]). With $|m| = 8$ and 12 bytes this gives a respective payload of 4 and 8 bytes. The Telos' CC2420 radio has a specification of 250

kbps, but in practice the data rate is typically less. Therefore, we overestimate the time for sending/receiving at 3 ms which accommodates our message sizes and is a reasonable TDMA slot size. With these parameters, a player can be active for $M = 18000\,\text{J}/(41\,\text{mW}\times 3\,\text{ms}) \approx 1.4\times 10^8$ slots. A MATLAB implementation of our protocol is used to investigate (1) and (2) with random and reactive jammers.

Random Jammers: We start by considering the absolute costs for each player when jamming occurs with probability $p_j = 0.5$ and 0.9 in each slot of an epoch. The adversary is able to jam for at most κ instances. Each trial terminates upon successful communication, or when κ is exhausted, and a run consists of 50 trials. Each point in Figure 2 represents the mean of 3 such runs and the maximum of Player A and Player B is plotted. For $p_j = 0.5$, the absolute cost is negligible. For $p_j = 0.9$, the cost is slow growing and still very small relative to M. This is representative for $p_j \leq 0.9$ (omitted for space) and we observe that the absolute costs in our analysis are tolerable up to this point. For larger values of p_j, we consider energy ratios. Figure 3 depicts the energy ratio for $p_j = 0.9, 0.95$ and 1.0 (constant jammer). Clearly, the energy ratio favors the players. Finally, a loose upper bound of $6\kappa^{\varphi-2}$ is plotted which implies a small constant in our asymptotic analysis.

It is also important to consider $\alpha_{cor} > \alpha_{adv}$. There is no consensus on the size of a jamming packet, $|m_j|$, required to disrupt a transmission. In [24], $|m_j| = 20$ bytes using MICA2 motes. However, smaller jamming packets have examined in more general wireless local area networks. In [25], $|m_j| \approx 3$ bytes and, in [26], $|m_j|$ is just a few bits! In comparison to an adversary that transmits $|m_j| = 2, 3$ or 5 bits, the players must handle a $|m| = 96$ bit (12 byte) message. This implies a factor discrepancy of $96/5 \approx 20$, $96/3 \approx 30$, $96/2 \approx 50$, respectively. Consequently, we consider $\alpha_{adv} = 1$ but $\alpha_{cor} = 20, 30$ and 50. Extrapolating using $\alpha_{cor} \cdot 6 \cdot \kappa^{\varphi-2}$, we note that the energy ratios for $\alpha_{cor} = 30$ and 50 drop below 1 for κ larger than 2.8×10^5, 8×10^5 and 3.1×10^6, respectively. These values are still small relative to M. For example, with $|m_j| = 5$ bits, the adversary must allow successful delivery of the first packet or face spending $M - 2.8 \times 10^5 \approx 10^8$ time slots suffering a disadvantageous energy ratio. By the same argument, hundreds of communications can occur before the players exhaust their energy supply.

Reactive Jammers: As discussed above, our protocol is resistant to indiscriminate jamming. Now consider a reactive adversary who can listen to 2 bytes of a header in a $|m| = 8$ byte packet and then decide whether to jam at zero cost. Therefore, $\alpha_{adv} = 1$ while $\alpha_{cor} = 4$. In each slot, channel traffic occurs with probability p_t from an outside source. By our discussion in Section IV, the adversary does not enter the costly receive state unless it detects (at zero cost) traffic on the channel via RSSI; otherwise, the adversary sleeps. Each trial consists of the adversary jamming m until its supply T is exhausted. For each value of T, a run of 10 trials are performed and each point in Figure 4 represents the mean of 3 runs. For values up to $T = 1 \times 10^5$ and $p_t = 0.3, 0.4, 0.5$ and 0.6 the players quickly achieve an advantageous energy ratio.

Discussion: These preliminary results suggest that our protocol is resistant to (i) efficient random jammers and (ii) reactive jammers when sufficient non-critical traffic is present. As with any DoS attack, throughput suffers. However, in situations where the successful transmission of a moderate number of packets is critical, these back-of-the-envelope calculations are an encouraging first-approximation of performance.

V. Multiple Receivers & Bounded Delivery Time

Multi-Player Competitive Communication, abbreviated MPCC, handles the scenario where a single sender a sends m to a set of n receivers $R_a = \{b_1, ..., b_n\}$. Furthermore, we bound the *latency*: the number of slots prior to successful termination. For (an unknown) T slots, we achieve $O(T^{\varphi/(\varphi-1)})$ latency given $\ln n = O(T)$; we discuss this later.

Our protocol is intuitive: execute Competitive Communication where a is Player A and each correct receiver acts as Player B. The probabilities for sending and receiving are modified and there are two more epochs (2 & 4) where players act deterministically. Again, a correct receiver terminates upon receiving m while the sender a terminates upon listening to a slot without receiving `req` or detecting a collision. Note that `req` messages can collide in Epochs 3 & 4; this is correct and the sender will retransmit. Figure V gives the pseudocode for MPCC. If the adversary jams, then *none of the correct receivers receive m in that slot*. A round is again defined as jamming-1 if at least $1/2$ of the slots in Epoch 1 are jammed while a round is jamming-2 if at least $1/2$ the slots in Epoch 2 are jammed or forged. Due to space constraints, several proofs are omitted below (see our technical report [49]).

Lemma 6. *Consider a non-jamming-1 round. The probability that at least one correct receiver does not receive the message from the sender is less than $1/n^2$.*

Proof: Let s be the number of slots in Epoch 1 of round i. Let $p_a = 3\ln n/2^i$ be the probability that the sender transmits in a particular slot. Let $p_b = 2/2^{(\varphi-1)i}$ ($p_b = 1$ for $i \leq 2$) be the probability that a correct receiver b_u listens in a particular slot. Let $X_j = 1$ if the message is not transmitted from sender a to receiver b_u in the j^{th} slot. Then Pr[m is not successfully transmitted to the b_u during Epoch 1]=$Pr[X_1X_2\cdots X_s = 1]$=$Pr(X_s = 1 \mid X_1X_2\cdots X_{s-1} = 1)$. Let $q_j = 1$ if the adversary does not jam given $X_1X_2\cdots X_{i-1}$; otherwise, let $q_j = 0$. The value of q_j can be selected arbitrarily by the adversary. Then $Pr[X_i = 1 \mid X_1\cdots X_{i-1} = 1] = 1 - p_A p_r q_i = 1 - 6\ln n/2^{\varphi i}$. Then we have $Pr[X_1X_2\cdots X_s] = (1 - p_a p_b q_1)\cdots(1 - p_a p_b q_s) \leq \prod_{j=1}^{s}(1 - p_a p_b q_j) \leq e^{-p_a p_b \sum_{j=1}^{s} q_j} < 1/n^4$ since $p_a p_b \sum q_j > 4\ln n$. Taking a union bound, the probability that at least one correct receiver has not received m is at most n^{-2}. ∎

Lemma 7. *Assume that by round i all correct receivers have heard the message m. Assume that round i is non-jamming-2. Then the probability that the sender retransmits the message in round $i+1$ is less than $1/n^2$.*

MULTI-PLAYER COMPETITIVE COMMUNICATION(m, a, R_a)

$i \leftarrow \lceil \ln(2\ln n)/(\varphi-1)\ln 2 \rceil$
$a\_terminate \leftarrow$ `false`
while ($a\_terminate ==$ `false`) **do**
  Epoch 1:
  for slot 1 to $2^{\varphi i}$ **do**
    - a sends m with probability $\frac{3\ln n}{2^i}$.
    - Each b_u that has not received m listens with probability $\frac{2}{2^{(\varphi-1)i}}$.

  Epoch 2:
  for slot 1 to $2^{(\varphi-1)i+1}$ **do**
    - a sends m.
    - Each b_u that has not received m listens.

  Epoch 3:
  for slot 1 to 2^i **do**
    - Each b_u that has not received m sends `req`.
    - a listens with probability $\frac{2\ln n}{2^i}$.
    - **if** (a listens but detects no collision or `req`) **then**
      - $a\_terminate \leftarrow$ `true`

  Epoch 4:
  for slot 1 to $2^{(\varphi-1)i+1}$ **do**
    - a listens.
    - **if** (a detects neither a collision nor `req`) **then**
      - $a\_terminate \leftarrow$ `true`
    - Each b_u that has not received m sends `req`.

  $i \leftarrow i+1$

Fig. 5. Pseudocode for MULTI-PLAYER COMPETITIVE COMMUNICATION.

Proof: This is computed similarly to the proof of Lemma 6. Let s be the number of slots in Epoch 2 and let $p = 4\ln n/2^i$ be the probability that the sender listens in a slot. For slot j, define X_j such that $X_j = 1$ if the sender does not terminate. Then $Pr[$ The sender retransmits m in round $i+1] = Pr[X_1X_2\cdots X_s = 1]$. Let $q_j = 1$ if the adversary does not jam given $X_1X_2\cdots X_{i-1}$; otherwise, let $q_j = 0$. The q_j values are determined arbitrarily by the faulty nodes in collusion. Since the sender terminates if and only if it listens and does not detect any activity, then $Pr[X_j = 1] = (1-pq_j)$. Therefore, $Pr[X_1X_2\cdots X_s = 1] \leq e^{-p\sum_{j=1}^{s} q_j} < n^{-2}$. ■

Lemma 8. *Assume all receivers are correct and there are no jamming-1 or jamming-2 rounds. Then the energy ratio of any correct player (sender or receiver) is* $O(\ln^2 n/\max\{1,T\})$.

Proof: Let $d = \frac{\ln(2\ln n)}{(\varphi-1)\ln 2}$. Using Lemma 6, the expected cost to the sender is at most $\sum_{i=d}^{\infty} n^{-2(i-d)} \cdot (2^{(\varphi-1)i} \cdot 3\ln n + 2^{(\varphi-1)i+1} + 2\ln n + 2^{(\varphi-1)i+1}) = O(\ln^2 n)$. Similarly, Therefore, Lemma 6, the expected cost to Player B is at most $\sum_{i=d}^{\infty} n^{-2(i-d)} \cdot (2^{i+1} + 2^{(\varphi-1)i+1} + 2^i + 2^{(\varphi-1)i+1}) = O(\ln^2 n)$. The result follows for an adversary that is active for $T > 0$ slots. ■

Lemma 9. *Assume there is at least one jamming-1 round. The energy ratio of the sender and any correct receiver is* $O(\max\{\ln^2 n/T, T^{\varphi-1}/T\})$.

Proof: Let i be the last round which is jamming-1 and let j be the last round which is jamming-2, $j \geq i$. Then the cost to the adversary is $\Omega(2^{\varphi i} + 2^j)$.

Sender: The energy ratio is highest for the sender when $i = 0$ so let $T = \Omega(2^j)$. In this case, using Lemma 7, the expected cost to the sender prior to successfully terminating is $O(\ln^2 n) + \sum_{k=1}^{\infty} n^{-2(k-1)} \cdot O(2^{(\varphi-1)(j+k)} + \ln n) = O(2^{(\varphi-1)j} + \ln^2 n)$. Therefore, the energy ratio for sender is $O(\max\{\ln^2 n, T^{\varphi-1}\}/T)$.

Receivers: For each correct receiver, the energy ratio to T is highest when $j = 0$ (a correct receiver is not active in the rounds indexed by j), so let $T = \Omega(2^{\varphi i})$. In the worst case, all rounds up to i have been jamming-1, in which case the expected cost to each correct receiver up to the end of round $i+1$ is $O(\ln^2 n + 2^i)$. Therefore, the energy ratio for each receiver is $O(\max\{\ln^2 n, T^{\varphi-1}\}/T)$ noting that $1/\varphi = \varphi - 1$. ■

Lemma 10. MPCC *terminates within* $O(\log T)$ *rounds or, equivalently,* $O(T^{\varphi/(\varphi-1)})$ *slots assuming* $\ln n = O(T)$.

Proof: If the adversary is not active for all slots in Epoch 2, then all correct receivers obtain m. Once all correct receivers terminate, the adversary must be active in all slots in Epoch 4 to prevent Player A from terminating. Therefore, prior to successful termination of all correct players (including the sender), the adversary is active for at least $2^{(\varphi-1)i+1}$ slots per round i in Epochs 2 & 4. For $d = \frac{\ln(2\ln n)}{(\varphi-1)\ln 2}$, we seek the number of rounds ρ such that $\sum_{i=d}^{\rho} 2^{(\varphi-1)i} \geq T$ which yields $\rho \leq \varphi \cdot \lg(T + 5\ln n)$. Each round i has at most $4 \cdot 2^{\varphi \cdot i+1}$ slots so ρ rounds equal at most $2^{\varphi+3} \cdot (T + 5\ln n)^{\varphi/(\varphi-1)}$ slots. ■

Theorem 2 of Section II-A follows from Lemmas 8, 9 and 10.

Discussion: The value n is the number of devices within the broadcast range of the sender. Therefore, for large networks, we expect n to be very small relative to the total number of network devices N. For a determined adversary, we expect $T \gg \ln^2 n$ since $T > n$ i.e. the number of transmissions a device can perform will likely exceed the number of neighboring devices. In this case, the energy ratio approaches $O(T^{\varphi-1}/T)$.

Finally, for any protocol, the optimal latency $\ell_{OPT} = \Omega(T)$ since the adversary can jam constantly for T slots. In contrast, assuming $\ln n = O(T)$, by Lemma 10 our protocol completes within $O(T^{\varphi/\varphi-1})$ slots. Noting $\varphi/(\varphi-1) - 1 = \varphi$, we have:

Corollary 1. MPCC *is within a* $O(T^{\varphi})$*-factor of* ℓ_{OPT}.

VI. DOS-RESISTANT RELIABLE BROADCAST IN THE GRID

Reliable broadcast addresses the problem of propagating m from a dealer to the rest of the network. The problem has been extensively studied in the grid model [6]–[12], [43]. Reliable broadcast is still possible when t Byzantine nodes can each jam at most n_c transmissions [43]. Unfortunately, the protocol of [43], and the subsequent refinement by [11], requires that *correct nodes possess much more energy than the Byzantine nodes*. In particular, while the send state costs are improved in [11], both [11], [43] allow the adversary to force a correct node to *listen* for $\Omega(t \cdot n_c)$ slots. In contrast, each Byzantine node is active for n_c. This $\Omega(t)$-factor discrepancy affords the

adversary a tremendous advantage and our main result is a protocol that mitigates this advantage.

The Grid Model: Each node $p(x, y)$ is situated at (x, y) in a grid. The dealer d is located at $(0, 0)$ and seeks to propagate m to all correct nodes in the network. When a node p sends a message, all awake nodes within L_∞ distance r (i.e. the $(2r + 1) \times (2r + 1)$ square centered about p) receive the message; this *neighborhood* is denoted by $N(p)$. Analogous results hold for the Euclidean metric (see [7]). There are $t < (r/2)(2r + 1)$ Byzantine nodes in any broadcast neighborhood. For any correct node p, the adversary can use its t Byzantine nodes in $N(p)$ to jam for up to $B_0 = t \cdot n_c$ slots total. There is a global broadcast schedule that assigns each node a slot for broadcasting; a specification is unimportant here (see [9] for an example). A complete iteration of the broadcast schedule is called a *cycle* and a node sends at least once per cycle.

Receive State Costs & Multi-Hop Networks: In this multi-hop scenario, the amount of jamming in a neighborhood (but not the total for adversary in the network) is bounded by B_0 and known. This is required in [11], [43] and a similar assumption is made in [34], [35]. A bound seems necessary so that correct nodes know when to wake up as m propagates outward. Our protocol uses this bound to synchronize sending/receiving. Alternatively, nodes may perpetually listen for transmissions. As discussed in Section III, this is adopted by other reliable broadcast protocols; however, the receive state costs are problematic. Therefore, we assume the existence of B_0 which is interpreted as the number of times a Byzantine node can deviate from the global broadcast schedule within some time frame before being identified and subjected to various defensive techniques (see [2]). Staying under B_0 in each time frame allows the adversary to attack throughout the lifetime of the network and we assume that B_0 is large.

Overview: Our protocol starts at slot 0 and synchronizes the timing of nodes for sending and listening. While this synchronization is not mathematically challenging, a full description yields an unreadable protocol. *For ease of exposition, our treatment addresses each node q in $C = \{q(x, y) | -r \leq x \leq r \wedge y \geq 0\}$; that is, a corridor of width $2r+1$ moving up from d.* Traversing the x-coordinates is nearly identical and the grid can be covered piecewise by these two types of corridors.

Node q issues a COMMIT(q, m) message if q has committed to m. Node q_2 sends HEARD(q_2, q_1, m) if q_2 has received COMMIT(q_1, m). As in [7], p commits to m when, through Steps 3 & 4, it receives $t + 1$ COMMIT(q, m) or HEARD(q_2, q_1, m) from node-disjoint paths all lying within a single $(2r + 1) \times (2r + 1)$ area. Due to space constraints, we omit this in the pseudocode and refer the reader to [7].

For each node p, $R_p = N(p) \cap C$. We say a node p initiates MPCC(m, p, R_p) in the context of the global broadcast schedule. This means the slots indexed by s in Epoch 1 of MPCC are those slots assigned to p by the global schedule. In Epoch 2, we deal with sending by nodes in R_p slightly differently. In each cycle, there is one extra slot assigned to each node. The slots indexed by s in Epoch 2 of MPCC are those slots used by nodes in R_p; a node in R_p can send `req`, and p is listening for a collision/`req` then. Figure 6 gives our pseudocode where $L = 2^{\varphi+3} \cdot (B_0 + 5 \ln n)^{\varphi/(\varphi-1)}$ given Lemma 10.

DoS-Resistant Reliable Broadcast

1: Starting in round 1, and ending no later than round L, node d executes MPCC(m, d, R_d) and each node $i \in R_d$ commits to the first value it receives from d.

The following step is executed by all nodes:

2: Starting in round $2yL$, and ending no later than round $(2y + 1)L - 1$, node $p(x, y)$ initiates MPCC(COMMIT$(p, m), p, R_p)$.

The following steps are executed by each node excluding those nodes in $N(d)$:

3: **for** $i = 0$ to $r - 1$ **do**

4: Starting in round $2(y - r + i)L$, and ending no later than round $2(y - r + i)L + L - 1$, node $p(x, y)$ listens for COMMIT messages by initiating MPCC(COMMIT$(q, m), q(x', y'), R_q)$ with each node in row $y' = y - r + i$ in C and where $p \in R_q$.

5: Starting in round $2(y - r + i)L + L$, and ending no later than round $2(y - r + i)L + 2L - 1$, node $p(x, y)$ listens for HEARD messages by initiating MPCC(HEARD$(q_2, q_1, m), q_2, R_{q_2})$ with each node $q_2 \in B'_p$ in row $y + i$ and where $p \in R_{q_2}$.

6: Starting in round $2(y-r)L+L$, and ending no later than round $2(y-r)L+2L-1$, node q_2 sends a HEARD message by initiating MPCC(HEARD$(q_2, q_1, m), q_2, R_{q_2})$ where q_1, q_2 are sister nodes.

Fig. 6. Pseudocode for DoS-Resistant Reliable Broadcast.

Proof of Theorem 3: In [7], it is shown that each node $p(x, y)$ can obtain m by majority filtering on messages from $2t + 1$ node-disjoint paths contained within a single $(2r+1)\times(2r+1)$ area since at least $t + 1$ will be m. Our correctness proof is similar; however, we argue along a corridor and show that nodes in the y^{th} row can commit to m by slot $2yL - 1$.

Base Case: Each node in $N(d)$ commits to the correct message m immediately upon hearing it directly from the dealer by round L. Therefore, clearly, every node $p(x, y) \in N(d)$ commits by round $2yL - 1$.

Induction Hypothesis: Let $-r \leq a \leq r$. If each correct node $p'(x', y') \in N(a, b)$ commits to m by round $2y'L - 1$, then each correct node $p(x, y) \in N(a, b+1) - N(a, b)$ commits to m in round $2yL - 1$.

Induction Step: We now show $2t + 1$ connectedness within a single neighborhood and we argue simultaneously about the time required for p to hear messages along these disjoint paths. The node $p(x, y)$ lies in $N(a, b + 1) - N(a, b)$ and can be considered to have location $(a - r + z, b + r + 1)$ where $0 \leq z \leq r$ (the case for $r + 1 \leq z \leq 2r$ follows by symmetry). We demonstrate that there exist $r(2r + 1)$ node-disjoint paths $P_1, ..., P_{r(2r+1)}$ all lying within the same neighborhood and that the synchronization prescribed by our protocol is correct:

One-Hop Paths: the set of nodes $A_p = \{q(u,v) \mid (a-r) \leq u \leq (a+z)$ and $(b+1) \leq v \leq (b+r)\}$ lie in $N(a,b)$ and are neighbors of p. Therefore, there are $r(r+z+1)$ paths of the form $q \to p$ where $q \in A_p$.

By their position relative to $p(x,y)$, each correct node $q(u,v) \in A_p$ is such that $v = y - r + c$ for some fixed $c \in \{0,...,r-1\}$. Therefore, by the induction hypothesis, q commits to m by round $2(y-r+c)L - 1$. By the protocol, $q(u,v)$ sends COMMIT messages using MPCC in round $2vL = 2(y-r+c)L$ until round $2(v+1)L - 1 = (2(y-r+c)+1)L - 1$ at the latest. By the protocol, $p(x,y)$ listens for COMMIT messages from q starting in round $2(y-r+c)L$ until round $(2(y-r+c)+1)L - 1$ at the latest. Therefore, p and q are synchronized in the execution of MPCC and p will receive q's message by round $(2(y-r+c)+1)L - 1 = (2(b+c+1)+1)L - 1$ at the latest. Since this occurs for all nodes in A_p, node p has received all COMMIT messages from A_p by round $(2(y-1)+1)L - 1 = (2(b+r)+1)L - 1 \leq (2(b+r+1)+1)L - 1 = 2yL - 1$.

Two-Hop Paths: consider the sets $B_p = \{q(u,v) \mid (a+z+1) \leq u \leq (a+r)$ and $(b+1) \leq v \leq (b+r)\}$ and $B'_p = \{q'(u',v') \mid (a+z+1-r) \leq u' \leq (a)$ and $(b+r+1) \leq v' \leq (b+2r)\}$. The nodes in B_p lie in $N(a,b)$ while the nodes in B'_p lie in $N(p)$. Moreover, the set B'_p is obtained by shifting left by r units and up by r units. Recall that there is a one-to-one mapping between the nodes in B_p and the nodes in B'_p; these are sister nodes. There are $r(r-z)$ paths of the form $q \to q' \to p$.

Consider a correct node $q(u,v) \in B_p$ and its sister node $q'(u',v') \in B'_p$ where $v' = v + r$ by definition. Again, given the location of $q(u,v)$ relative to $p(x,y)$, we have $v = y - r + c$ for some fixed $c \in \{0,...,r-1\}$. By the induction hypothesis, q commits to m by round $2vL - 1$. Then by DOS-RESISTANT RELIABLE BROADCAST, q sends a COMMIT message using MPCC in round $2vL = 2(y-r+c)L$ until round $2(v+1)L - 1 = (2(y-r+c)+1)L - 1$ at the latest. By DOS-RESISTANT RELIABLE BROADCAST, $q'(u',v')$ receives COMMIT messages from q using MPCC starting in round $2(v'-r+c)L = 2vL = 2(y-r+c)L$ and ending no later than round $2(v'-r+c+1)L - 1 = 2(v+1)L - 1 = (2(y-r+c)+1)L - 1$. Therefore, q and q' are synchronized in the execution of MPCC and q' will receive q's message by round $(2(y-r+c)+1)L - 1 \leq 2yL - 1$ at the latest.

By the above, each node $q'(u',v') \in B'_p$ can start sending a HEARD message using MPCC in round $2(v'-r)L + L$ and ending no later than round $2(v'-r)L + 2L - 1$. Starting in round $2(y-r+c)L + L$, node $p(x,y)$ uses MPCC to listen for a HEARD message from $q'(u',v')$ where $v' = y+c$. Therefore, p is listening to q' starting in $2(y-r+c)L + L = 2(v'-r)L + L$ and ending no later than $2(v'-r)L + 2L - 1$; p and q' are synchronized. Therefore, p receives all HEARD messages by round $2(v'-r)L + 2L - 1$ when $v' = y + r - 1$; that is, by round $2(y-1)L + 2L - 1 = 2yL - 1$.

Therefore, a total of $r(r+z+1) + r(r-z) = r(2r+1)$ node-disjoint paths from $N(a,b)$ to $PN(a,b)$ exist, all lying in in a single neighborhood $N(a, b+r+1)$. For an adversary corrupting $t < (r/2)(2r+1)$ nodes, a correct node can majority filter to obtain m. Furthermore, we have shown that any $p(x,y) \in N(a,b+1)$ executes MPCC $r(2r+1) = O(r^2)$ times in order to receives all COMMIT and HEARD messages by round $2yL - 1$. Therefore, p can commit to the correct message by round $2yL - 1$; this concludes the induction.

Cost Analysis: In our protocol, a correct node p is involved in $k \leq r(2r+1) + 1$ executions of MPCC. For each such execution, let τ_i be the number of slots for which the adversary is active where $i = 1, ..., k$. Denote the adversary's total active time by $\beta = \sum_{i=1}^{k} \tau_i < B_0$. We examine two cases:

Case 1: Assume the adversary is active for a total of $\beta = \sum_{i=1}^{k} \tau_i = O(r^2 \ln^{2/(2-\varphi)} r)$ slots over all k executions of MPCC by p. Note that $\beta = 0$ is possible. For each execution, node p incurs a cost of $O(\max\{\ln^2 r, \tau_i^{\varphi-1}\})$ by Theorem 2. Therefore, over $k = O(r^2)$ executions, clearly the total cost (not the energy ratio) incurred by p is $O(r^4 \ln^{2/(2-\varphi)} r)$.

Case 2: Otherwise, assume that the adversary jams for at least one slot. Using Theorem 2, the energy ratio is:

$$O\left(\frac{\sum_{i=1}^{k} \max\{\ln^2 r, \tau_i^{\varphi-1}\}}{\sum_{i=1}^{k} \tau_i}\right) = O\left(\frac{r^2 \ln^2 r}{\beta} + \frac{(r^2)^{(2-\varphi)} \cdot \ln^2 r}{\beta^{2-\varphi}}\right)$$

by noting $(\sum_i \tau_i^{\varphi-1})/(\sum_i \tau_i) \leq k^{2-\varphi}(\sum_i \tau_i)^{\varphi-2}$ since $\tau_i^{\varphi-1}$ is concave and $(\sum_i \tau_i^{\varphi-1})/k \leq (\frac{1}{k}\sum_i \tau_i)^{\varphi-1}$ by the corollary of Jensen's inequality that applies to concave functions. ■

Discussion: We make two points regarding our protocol. First, for $\beta = O(r^2 \ln^{2/(2-\varphi)} r)$, the absolute cost to p is $O(r^4 \ln^{2/(2-\varphi)} r)$. But for $\beta = \omega(r^2 \ln^{2/(2-\varphi)} r)$, the energy ratio is dominated by $\approx O(r^{0.764} \ln^2 r/\beta^{0.382})$ and, at this point, it favors p; the adversary no longer has the advantage.

Our second point concerns the upper bound $t = (r/2)(2r+1) - 1$ for which reliable broadcast is feasible under DoS attacks [43]; recall that correct nodes require more energy. However, if nodes are subverted, it seems reasonable that Byzantine and correct nodes will each possess roughly equal energy. Assume that each node can be active for n_c slots. Consider the particular situation where the adversary targets p by having one of its nodes jam for n_c instances. This exhausts p's energy and *reliable broadcast fails*. This attack is suggested in [43]. In contrast, with our protocol, p can expect to avoid being disabled in this situation if $t = o(n_c^{2-\varphi}/(r^{2(2-\varphi)} \ln^2 r))$ for sufficiently large n_c. We note that p's survival is not guaranteed and, if $t = \Omega(n_c^{2-\varphi}/(r^{2(2-\varphi)} \ln^2 r))$, the adversary can expect to disable p by using enough of its nodes. However, this is an improvement over previous results (which cannot tolerate this attack) and it illustrates the importance of accounting for receive state costs when considering bounds on t.

Acknowledgements: We gratefully thank Srinivasan Keshav and James Horey for many valuable discussions.

REFERENCES

[1] J. P. Walters, Z. Liang, W. Shi, and V. Chaudhary, *Security in Distributed, Grid, Mobile, and Pervasive Computing. Chapter 17: Wireless Sensor Network Security: A Survey*. Auerbach Publications, 2007.

[2] A. D. Wood and J. A. Stankovic, "Denial of Service in Sensor Networks," *Computer*, vol. 35, no. 10, pp. 54–62, 2002.

[3] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling Ultra-Low Power Wireless Research," in *IPSN*, 2005.

[4] Crossbow, MICAz Wireless Measurement System. http:// courses.ece. ubc.ca/ 494/files/ MICAz_Datasheet. pdf.

[5] ——, Mica2 Wireless Measurement System. https:// www.eol.ucar.edu/ rtf/facilities/ isa/internal/ CrossBow/ DataSheets/ mica2.pdf.

[6] V. Bhandari and N. H. Vaidya, "On Reliable Broadcast in a Radio Network," in *PODC*, 2005, pp. 138–147.

[7] ——, "On Reliable Broadcast in a Radio Network: A Simplified Characterization," CSL, UIUC, Tech. Rep., May 2005.

[8] ——, "Reliable Broadcast in Wireless Networks with Probabilistic Failures," in *INFOCOM*, 2007, pp. 715–723.

[9] C.-Y. Koo, "Broadcast in Radio Networks Tolerating Byzantine Adversarial Behavior," in *PODC*, 2004, pp. 275–282.

[10] Valerie King and Cynthia Phillips and Jared Saia and Maxwell Young, "Sleeping on the Job: Energy-Efficient and Robust Broadcast for Radio Networks," *Accepted to Algorithmica*, 2010.

[11] M. Bertier, A.-M. Kermarrec, and G. Tan, "Message-Efficient Byzantine Fault-Tolerant Broadcast in a Multi-Hop Wireless Sensor Network," in *ICDCS*, 2010, pp. 408–417.

[12] D. Alistarh, S. Gilbert, R. Guerraoui, Z. Milosevic, and C. Newport, "Securing Your Every Bit: Reliable Broadcast in Byzantine Wireless Networks," in *SPAA*, 2010, pp. 50–59.

[13] B. Raman and K. Chebrolu, "Censor Networks: A Critique of "Sensor Networks" from a Systems Perspective," *Computer Communication Review*, vol. 38, pp. 75–78, 2008.

[14] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-Efficient Communication Protocol for Wireless Microsensor Networks," in *HICSS*, 2000, pp. 3005–3014.

[15] S. Ganeriwal, C. Pöpper, S. Čapkun, and M. B. Srivastava, "Secure Time Synchronization in Sensor Networks," *ACM Transactions on Information and System Security*, vol. 11, no. 23, 2008.

[16] W. Ye, J. Heidemann, and D. Estrin, "An Energy-Efficient MAC Protocol for Wireless Sensor Networks," in *INFOCOM*, 2002, pp. 1567–1576.

[17] Y. Li, W. Ye, and J. Heidemann, "Energy and Latency Control in Low Duty Cycle MAC Protocols," in *WCNC*, 2005, pp. 676–682.

[18] I. Ramachandran and S. Roy, "Clear Channel Assessment in Energy-Constrained Wideband Wireless Networks," *IEEE Wireless Communications*, vol. 14, no. 3, pp. 70–78, 2007.

[19] J. Deng, P. K. Varshney, and Z. J. Haas, "A New Backoff Algorithm for the IEEE 802.11 Distributed Coordination Function," in *CNDS*, 2004, pp. 215–225.

[20] D. Liu and P. Ning, "Multi-Level μTESLA: Broadcast Authentication for Distributed Sensor Networks," *ACM Transactions in Embedded Computing Systems*, vol. 3, pp. 800–836, 2004.

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

[22] Y. W. Law, J. Doumen, and P. Hartel, "Survey and Benchmark of Block Ciphers for Wireless Sensor Networks," *ACM Transactions on Sensor Networks*, vol. 2, no. 1, pp. 65–93, 2006.

[23] C. Karlof and N. Sastry and D. Wagner, "TinySec: A Link Layer Security Architecture for Wireless Sensor Networks," in *SenSys*, 2004, pp.162-175.

[24] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks," in *MobiHoc*, 2005, pp. 46–57.

[25] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa, "On the Performance of IEEE 802.11 Under Jamming," in *INFOCOM*, 2008, pp. 1265–1273.

[26] G. Lin and G. Noubir, "On Link Layer Denial of Service in Data Wireless LANs," *Wireless Communications & Mobile Computing*, vol. 5, no. 3, pp. 273–284, 2005.

[27] N. Aschenbruck, E. Gerhards-Padilla, and P. Martini, "Simulative Evaluation of Adaptive Jamming Detection in Wireless Multi-hop Networks," in *ICDCS Workshops*, 2010, pp. 213–220.

[28] W. Xu, K. Ma, W. Trappe, and Y. Zhang, "Jamming Sensor Networks: Attack and Defense Strategies," IEEE *Networks*, vol. 20, no. 3, pp. 41–47, 2006.

[29] X. Liu, G. Noubir, R. Sundaram, and S. Tan, "SPREAD: Foiling Smart Jammers Using Multi-Layer Agility," in *INFOCOM*,2007, pp. 2536–2540.

[30] V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein, "Using Channel Hopping to Increase 802.11 Resilience to Jamming Attacks," in *INFOCOM*, 2007, pp. 2526–2530.

[31] S. Gilbert, R. Guerraoui, and C. C. Newport, "Of Malicious Motes and Suspicious Sensors: On the Efficiency of Malicious Interference in Wireless Networks," in *OPODIS*, 2006, pp. 215–229.

[32] T. Brown, J. James, and A. 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*, 2006, pp. 120–130.

[33] A. Pelc and D. Peleg, "Feasibility and Complexity of Broadcasting with Random Transmission Failures," in *PODC*, 2005, pp. 334–341.

[34] B. Awerbuch, A. Richa, and C. Scheideler, "A Jamming-Resistant MAC Protocol for Single-Hop Wireless Networks," in*PODC*, 2008, pp. 45–54.

[35] A. Richa, C. Scheideler, S. Schmid, and J. Zhang, "A Jamming-Resistant MAC Protocol for Multi-Hop Wireless Networks," in *DISC*, 2010.

[36] 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 *DISC*, 2007, pp. 208–222.

[37] S. Gilbert, R. Guerraoui, D. Kowalski, and C. Newport, "Interference-resilient information exchange," in *INFOCOM*, 2009, pp. 2249–2257.

[38] D. Meier, Y. A. Pignolet, S. Schmid, and R. Wattenhofer, "Speed Dating Despite Jammers," in *DCOSS*, 2009, pp. 1–14.

[39] S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport, "Secure communication over radio channels," in *PODC*, 2008, pp. 105–114.

[40] V. King, C. Phillips, J. Saia, and M. Young, "Sleeping on the Job: Energy-Efficient and Robust Broadcast for Radio Networks," in *PODC*, 2008, pp. 243–252.

[41] M. Bertier, A.-M. Kermarrec, and G. Tan, "Brief announcement: Reliable broadcast tolerating byzantine faults in a message-bounded radio network," in *Proceedings of the 22^{nd} International Symposium on Distributed Computing (DISC)*, 2008, pp. 516–517.

[42] V. Vaikuntanathan, "Brief announcement: Broadcast in Radio Networks in the Presence of Byzantine Adversaries," in *PODC*, 2005.

[43] V. Bhandhari, J. Katz, C.-Y. Koo, and N. Vaidya, "Reliable Broadcast in Radio Networks: The Bounded Collision Case," in *PODC*, 2006, pp. 258 – 264.

[44] J. Sundali and R. Croson, "Biases in Casino Betting: The Hot Hand and the Gamblers Fallacy," *Judgment and Decision Making*, vol. 1, no. 1, pp. 1–12, 2006.

[45] K. Srinivasan and P. Levis, "RSSI is Under Appreciated," in*EmNets*, 2006.

[46] J. Bardwell, "Converting signal strength percentage to dbm values." [Online]. Available: http:// www.wildpackets.com/ elements/ whitepapers/ Converting_Signal_Strength.pdf

[47] J. Polastre, J. Hill, and D. Culler, "Versatile Low Power Media Access for Wireless Sensor Networks," in *SenSys*, 2004, pp. 95–107.

[48] B. Cody-Kenny, D. Guerin, D. Ennis, R. S. Carbajo, M. Huggard, and C. M. Goldrick, "Performance Evaluation of the 6LoWPAN Protocol on MICAz and TelosB Motes," in *PM2HW2N*, 2009, pp. 25–30.

[49] V. King, J. Saia, and M. Young, "Talk is Cheap(er): Mitigating DoS and Byzantine Attacks in Wireless Sensor Networks," Tech. Rep. CS-2010-14, 2010, http://www.cs.uwaterloo.ca/research/tr/2010/CS-2010-14.pdf.