Towards Jamming-Resistant and Competitive Medium Access in the SINR Model

Andrea Richa, Jin Zhang
Computer Science and Engineering, SCIDSE
Arizona State University
Tempe, AZ 85287, USA
{aricha,jzhang82}@asu.edu

Christian Scheideler
Department of Computer Science
University of Paderborn
D-33102 Paderborn, Germany
scheideler@upb.de

Stefan Schmid
Deutsche Telekom Laboratories & TU Berlin
D-10587 Berlin, Germany
stefan@net.t-labs.tu-berlin.de

ABSTRACT

The efficient coordination of medium access is arguably one of the most relevant applications of distributed computing. Recently, progress has been made in the design of robust medium access (MAC) protocols that guarantee a competitive throughput against a powerful jammer which can block the medium an arbitrary constant fraction $(1-\varepsilon)$ of the time. These MAC protocols exploit the remaining ε-fraction optimally in the sense that a significant part is used for successful transmissions. However, so far these throughput guarantees only hold for rather simplistic interference models such as Unit Disk Graphs.

This paper reports on our first insights on the design of a robust medium access protocol SinrMac for the more realistic physical interference model which takes into account the signal to interference plus noise ratio (SINR) at the receiver. This model is more difficult, as there is no longer an objective distinction of idling and busy time periods which can be used to dynamically adjust the wireless nodes' backoff periods. We discuss an approach that introduces individual "idle/busy thresholds" which are adapted dynamically and, unlike the multiplicative backoff periods, in an *additive* manner. We find that a reasonable convergence speed (and throughput) can be achieved if there exists some meaningful upper bound $\hat{\tau}$ on the noise level in the network; surprisingly, however, our first simulation results indicate that adaptive changes of the idly/busy thresholds do not yield a better throughput than static thresholds set to $\hat{\tau}$.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems

General Terms

Algorithms, Design

Keywords

Distributed Control, Medium Access, Wireless, SINR, Interference, Jamming, Throughput

1. INTRODUCTION

Medium access is a central challenge in wireless computing. In addition to the complication that the participants (or *nodes*) of a wireless network may gather in an ad-hoc fashion, join and leave arbitrarily over time, or even be mobile, communication may be disrupted by external interference from co-existing networks, microwaves, or even jammers. Despite the topic's apparent relevance, researchers still do not well understand how to design efficient MAC protocols that guarantee a provably high throughput.

Recently, Awerbuch et al. presented first distributed algorithms that provide performance guarantees against a powerful adversary who can jam the medium an arbitrary constant $(1-\varepsilon)$-fraction of any time window of size T: more formally, for some given $T \in \mathbb{N}$ and $0 < \varepsilon \leq 1$, their $(T, 1-\varepsilon)$-*bounded* adversary can jam at most $(1-\varepsilon)w$ of the time steps, for any time window of size $w \geq T$. This adversary is even allowed to be adaptive in the sense that it has complete knowledge of the protocol's execution history. While the first results applied to a single-hop network only [1] (where it is also possible to elect a leader in a self-stabilizing manner [5], and a modified protocol is even competitive against reactive jamming [4]), the findings were subsequently generalized to *Unit Disk Graphs (UDG)* [3].

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

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

S3'11, September 19, 2011, Las Vegas, Nevada, USA.

2. MODEL

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

$$SINR = \frac{d(u,v)^{-\alpha}}{\mathcal{N} + \sum_{w \in S} d(w,v)^{-\alpha}} \geq \beta_1,$$

where $\mathcal{N}$ captures the (e.g., thermal) noise, S is the set of nodes with concurrent transmissions, and β_1 is the *SINR threshold*.

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

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

We assume that in addition to the nodes there is an *adversary*: the idea is that our conservative definition of adversary subsumes many different forms of intentional and unintentional interference. Concretely, like in [1], we want to allow the adversary to know the protocol and its entire history and to use this knowledge in order to *jam* the wireless channel at will at any round (i.e, the adversary is *adaptive*). However, unlike in previous works [1], the adversary is not bounded over time in the sense that it can only jam a subset of the time periods, but *with respect to energy*: for each time period of length T, the adversary has a certain energy budget to disrupt communications. Rather than assuming some jammer locations in the Euclidean plane from which it can transmit at different energy levels, we propose a model where the jammer has a certain budget B_v *for each wireless node* $v \in V$. Henceforth, we assume that this budget is the same for every node and we will simply refer to it by B. Such a jammer is called a (B,T)-*bounded adversary*: in every time interval of size $w \geq T$, the adversary can add $B \cdot w/T$ to the noise level $\mathcal{N}$ of each node.

Our goal is to design a *symmetric local-control* MAC protocol (i.e., there is no central authority controlling the nodes, and all the nodes are executing the same protocol) that has a "competitive" throughput against any (B,T)-bounded adversary in any multi-hop network that can be modeled by SINR. Intuitively, we want to call a MAC protocol competitive if the number of successful message receptions at the nodes is a "large" fraction of the messages that would have been received if the adversarial contributions to the noise $\mathcal{N}$ are subtracted in the SINR formula for the corresponding time steps.

3. THE MAC PROTOCOL

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

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

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

Our solution, the SinrMac protocol, is formally described in Algorithm 1. The algorithm is essentially interpreting any noise floor smaller than τ_v as an idle channel and increases

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

the sending probabilities accordingly; if on the other hand the noise is relatively high, the sending probabilities are reduced, but only after T_v rounds where the channel was not idle.

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

Algorithm 1 SinrMac

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

4. FIRST RESULTS

Although intuitively, adapting τ_v seems to be crucial to accurately react to the channel states and converge to a good throughput, our first experiments indicate that static τ_v values (fixed at the maximal possible reception power) are often better. In the following, we report on our preliminary simulation study to evaluate the performance of our protocol in terms of throughput and as a function of the network size. We define throughput as the number of messages successfully received in the whole network per round. In our network, nodes are distributed on a square grid (i.e., the number of nodes is $n = a \times a$ for some parameter a), and we allow the simplistic adversary to evenly allocate its jamming budget ($B = 200$, no other noise) among T time steps, i.e., $\mathcal{N} = B/T$ per round. The transmission power for all nodes is set to 4, the SINR ratio is $\beta_1 = 0.5$, and $T = 50$. We set $c = 0.1$, and consider $\hat{p} = \{1/24, 1/2\}$.

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

Figure 1 show an exemplary dependency of the throughput on the different τ_v schemes when $\hat{p} := 1/24$. We see that an adaptive τ_v value approach performs worse than a strategy fixing τ_v at a high level: $\tau_v = 40$ lets p_v stay close to $\hat{p}$, which is still fine since $1/24$ is relatively small. In other experiments, we find that if $\hat{p} = 1/2$, fixing τ_v at a lower level, i.e., $\tau_v = 4$, gives the best throughput, since in this case $\hat{p}$ is much higher, and hence there are more collisions and busy time periods. Here, being able to identify the busy channels and decrease access probabilities accordingly is crucial for the protocol to achieve a good throughput.

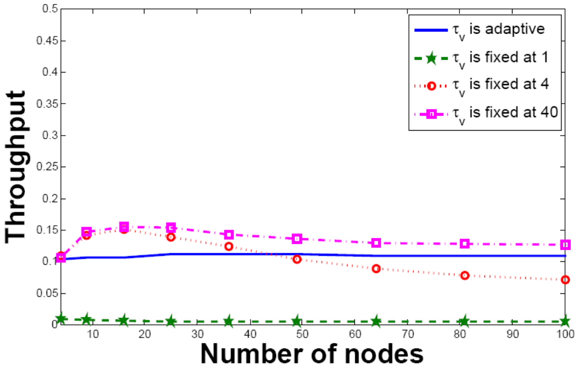

Figure 1: Normalized throughput as a function of the network size and under different τ_v adaption schemes. The result is averaged over 5 runs.

5. CONCLUSION

This paper described a preliminary MAC protocol for the SINR model under jamming activities. Surprisingly, we found that our adaptive idle/busy threshold adaption strategy often performs worse than a static strategy. In our future work, we plan to rigorously evaluate different adapting schemes for τ_v, and study our algorithm under more sophisticated and worst-case adversaries, not only empirically but hopefully also by deriving performance proofs. Obviously, in this process, changes to the protocol presented here may be required.

6. REFERENCES

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

[2] D. Blough, C. Canali, G. Resta, and P. Santi. On the impact of far-away interference on evaluations of wireless multihop networks. In *Proc. MSWIM*, 2009.

[3] A. Richa, C. Scheideler, S. Schmid, and J. Zhang. A jamming-resistant MAC protocol for multi-hop wireless networks. In *Proc. DISC*, 2010.

[4] A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Competitive and fair medium access despite reactive jamming. In *Proc. ICDCS*, 2011.

[5] A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Self-stabilizing leader election for single-hop wireless networks despite jamming. In *Proc. MobiHoc*, 2011.