Toward Optimal Utilization of Shared Random Access Channels

Joseph (Seffi) Naor[*]
Computer Science Department
Technion - I.I.T.
Haifa 32000, Israel
naor@cs.technion.ac.il

Danny Raz
Computer Science Department
Technion - I.I.T.
Haifa 32000, Israel
danny@cs.technion.ac.il

Gabriel Scalosub
Department of Computer Science
University of Toronto
Toronto, ON, Canada
scalosub@cs.toronto.edu

***Abstract*—We consider a multipacket reception channel shared by several communication applications. This is the case, for example, in a single radio mesh network where neighboring cells use the same radio channel. In such scenarios, unlike the common multiple access model, several transmissions may succeed simultaneously, depending on the actual locations of the sending and receiving stations, and thus channel utilization may be greater than 1.**

Our goal is to derive a decentralized access control mechanism that maximizes the channel utilization, while taking into account fairness among the different users. We focus on a simple case where each user can adjust a single parameter that determines its transmission probability in any time slot, and develop such a protocol for the general problem, where users are distributed arbitrarily, based on strong motivation which is derived from analytical bounds for homogeneous interferences. We further show, using extensive simulations, that this protocol achieves a high utilization of radio resources compared to any other protocol (not necessarily based on a simple parameter), while maintaining fairness between all users.

I. INTRODUCTION

Wireless networks often involve a joint usage of common communication channels, in a multiple access environment. In most of the models describing such settings, simultaneous transmission by more than one station results in a collision which causes all transmissions at that time to fail. Medium access control (MAC) protocols based on carrier sensing with collision detection (CSMA/CD) or collision avoidance (CSMA/CA), are used in such scenarios in order to deal with collisions, in an attempt to maximize the system's throughput. However, in many current wireless networks, such as mesh WiFi networks, or 802.15 clusters, simultaneous usage of the same wireless channel is possible, i.e., there could be several successful simultaneous transmissions using the same channel, at the same geographical proximity. Consider for example the settings described in Figure 1, where we outline two WiFi stations, A, B, and their transmission ranges. Assume each station has a client which is supposed to receive a transmission from that station. If the clients of A and B are a and b respectively, then simultaneous transmissions will cause a collision at client a, while b can receive the message from B. This is due to the fact that b is much closer to B than to A, and thus the signal from station B is much stronger compared to the signal of station A. Client a, on the other hand, is approximately at equal distance from A and B and thus the interference is very strong with respect to the signal. If, however, the clients of A and B are a' and b respectively, then both simultaneous transmissions will succeed, since they do not collide at either of the receiving ends.

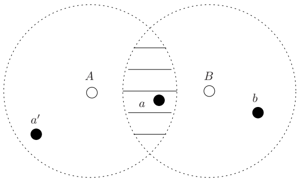

Fig. 1. Outline of two stations, A, B and their transmissions ranges.

Indeed, the above observations may seem trivial in the sense that there is no reason to refrain from having both stations transmit in the latter case. However, if we assume that the locations of the clients are being drawn from some (possibly unknown) distribution, then the stations have *no knowledge* of where exactly their target resides, and therefore the stations cannot know which of the scenarios is the situation at hand. In spite of this lack of knowledge, stations must still decide whether or not to transmit at any given time slot. This is the main difficulty we try to address in this work. Our approach is based on incorporating any available prior knowledge of the interferences between the various stations into the model, and the goal is to develop a theory as well as protocols for analyzing and solving this difficulty.

One can consider different approaches for this problem. One way is to assume a powerful centralized unit that has full knowledge and control over all stations. At each time slot, this centralized unit decides which of the stations should transmit and which should remain silent. While powerful, it is very hard to implement such protocols and a distributed approach is much more scalable and useful in practice. In this paper we concentrate on a family of simple distributed protocols, where each station i chooses a single parameter R_i and at each time slot it transmits with probability R_i. We show that this simple approach can achieve a very good utilization of the channel if the right parameter is chosen.

One problem that arises in such a distributed setting is fairness. There may be different ways to consider fairness.

---

[*] Research supported in part by ISF Grant 1366/07 and US-Israel BSF Grant 2002276.

One way is to require that all stations have an equal share of the radio resources. Clearly, such a requirement may result in a very low utilization of the network, especially for arbitrarily deployed networks. We therefore require a weaker notion of fairness, namely, making sure that no station is starved.

The model we consider is general, and can be applied to many wireless environments, such as the ones emerging in wireless mesh networks, ultra wideband (UWB) environments such as the ones appearing in wireless personal area networks (WPAN), and other ad-hoc networking environments, where carrier sensing approaches such as CSMA/CD and CSMA/CA may not be applicable. Our results show that taking spatial considerations into account, even if these considerations are somewhat noisy, may significantly increase the network's throughput and efficiency. We now elaborate a little on some of the environments that might benefit from our approach:

a) Wireless mesh networks: Wireless mesh networks are among the most promising access techniques, allowing robust connectivity at a relatively low cost. Such networks consist of access points which both provide connectivity to wireless clients, as well as acting as mesh routers and forward traffic to and from the wired network. Single radio wireless mesh networks are currently the most commonly deployed commercial solution [1]. In these networks, a single radio channel is used for the communications between the backhaul nodes, as well as the communications between the wireless clients, such as laptops, or handheld devices, and the access points. By contrast, in dual-radio (or multi-radio) networks different radio channels are used for each of the latter types of communication. In such settings, our model captures, for example, high load downstream traffic to the various clients of the access points, in a single-radio environment. In recent years a lot of research was done on various methods of improving MAC protocols for such networks. Our model provides an alternative view of medium access control in wireless mesh environments, and we note that our results may also be applied to multi-radio wireless mesh networks, by applying our proposed protocol to each channel independently.

b) Ultra wideband networks: There has been extensive research in recent years trying to design good MAC protocols for UWB networks. Our suggested approach is especially appealing for UWB environments, which pose a major difficulty in utilizing CSMA-based protocols for channel access. This difficulty is based primarily on the fact that channel assessment is difficult due to low power emissions, as well as due to the fact that UWB uses pulse position modulation, which is practically carrierless [2].

c) Wireless personal area networks: In the case of the relatively new IEEE 802.15.4 standard (Zigbee) for WPAN, the system's performance depends heavily upon maintaining synchronization in the network using beacon frames. The behavior of the proposed synchronization scheme is in itself a multiple access environment. However, as of yet, the standard does not fully specify the protocols to be used in this respect [3]. We believe that the protocol we suggest in Section III can be used to address these challenges, due to its simplicity and robustness.

d) Cellular networks: Medium access in cellular networks usually does not use CSMA-type protocols due to their inefficiencies when using licensed band. In such settings, especially for initial uplink access, random access protocols, e.g., variants of slotted Aloha, are used and were shown to provide good performance. These results have also been extended to random channel selection in OFDMA wireless networks [4]. Our work provides a general framework for designing random access protocols in such settings.

In the remainder of this section we give a formal definition of our network model, present a summary of our results, and discuss related work. In subsequent sections we present a high-level motivation for our protocol, followed by a formal presentation of the protocol. We then present a simulation study of its performance, and the performance of several other simple protocols for non-homogenous interferences for environments which cannot support carrier-sensing. We complete our simulation study by discussing some of the practical aspects of implementation, and evaluating the robustness of our approach to inaccurate parameters estimation. Finally, in Section V we present some conclusions and open questions.

A. Model

We consider a network consisting of $n \geq 2$ stations using a common wireless medium. It is assumed that time is divided into fixed length slots, and transmission of one packet takes a single slot. We further focus on the decisions made by the stations in each time slot, as well as on the overall system performance in every time slot. We focus our attention on *random access* disciplines employed by the various stations. For every station i, we let $R_i \in [0,1]$ denote the *probability* that station i transmits. Due to interferences, the probability of a *successful* transmission also depends upon the transmission of other stations. Given a probabilities-vector $\overline{R} = (R_1, \ldots, R_n) \in [0,1]^n$, we define the *success probability* of station i's transmission as:

$$r_i(\overline{R}) = R_i \cdot \Pi_{j \neq i}(1 - \alpha_{i,j} R_j),$$

where for every $1 \leq i, j \leq n$, $\alpha_{i,j} \in [0,1]$ is a fixed network-dependent parameter denoting the interference inflicted on i upon simultaneous transmission of both i and j.

One way to think about the $\alpha_{i,j}$s is considering them as the probability that a transmission by j will interfere with a transmission of i. Then, in order for i's transmission to succeed, transmissions of other stations should not interfere. Clearly, if for all i, j, $\alpha_{i,j} = 0$, i.e., there are no interferences, then the best protocol would simply have all stations transmitting with probability 1, which implies an optimal use of resources, both per-station and globally. On the other hand, if for all i, j, $\alpha_{i,j} = 1$, then our model coincides with the classic collision channel used for analyzing multiple access environments, where in any case of simultaneous transmissions a collision occurs, resulting in the failure of all the transmissions. For some more concrete intuition as to the interferences represented by the values of $\alpha_{i,j}$, consider for example the situation depicted in Figure 1. Assume stations A and B each has a client chosen in its transmission range uniformly at random, such that the choices of the clients of the two stations are independent. It follows that the probability of a station's client being in the marked area of the intersection of the two transmission areas is proportional to the quotient of the area of the intersection and the station's overall transmission area. By letting $\alpha_{i,j}$ be this quotient, the above network scenario is captured precisely by our model.

Interferences are called *homogeneous* if there exists some $\alpha \in (0, 1)$ such that for all i, j, $\alpha_{i,j} = \alpha$.[1] If this is not the case, then we refer to the interferences as *non-homogeneous*. Note that in most practical cases, interferences are expected to be non-homogeneous. Given a station i, its *neighbors* are defined to be the set of all stations j such that $\alpha_{i,j} > 0$.

Given a probabilities-vector $\overline{R} = (R_1, \ldots, R_n)$, the *throughput* of the system $\varphi(\overline{R})$ is considered to be the overall use of resources in the system given probabilities-vector $\overline{R}$,

$$\varphi(\overline{R}) = \sum_{i=1}^{n} r_i(\overline{R}) = \sum_{i=1}^{n} R_i \prod_{j \neq i} (1 - \alpha_{i,j} R_j).$$

$\varphi(\overline{R})$ should be interpreted as the expected number of successful transmissions in a given time slot. Note that a-priori, $\varphi(\overline{R})$ can take any value between 0 and n, where the former is its value e.g. in case where $\alpha_{i,j} = 1$ for all i, j, and $R_i = 1$ for all i, and the latter is its value where $\alpha_{i,j} = 0$ for all i, j, and $R_i = 1$ for all i. In what follows we refer to a probabilities-vector $\overline{R}$ as *uniform*, if $R_i = R_j$ for all i, j.

We refer to the above setting as the *spatial interferences multiple access (SIMA) environment*.

B. Our Results

We present a model that enables capturing the effect of interferences in shared wireless networks. Motivated by analytical results for homogeneous interferences (which are less probable to appear in real life scenarios, see [5]), we devise a new, simple distributed random-access protocol for the general settings of non-homogeneous interferences.

As a test case to exhibit the benefits of our protocol, we consider the setting arising in single-radio mesh networks, focusing on the client downlink level, in traffic-intensive scenarios, without resorting to carrier-sensing. We perform an extensive simulation study of the new protocol's performance, as well as the performance of several other protocols for the problem, on random architectures of wireless mesh networks of various magnitudes.

Our new protocol is shown to obtain higher throughput than all other protocols for the problem. This is despite the fact that, as explained above, it only uses a single variable to determine each station's transmission probability in each slot, and thus it is very simple, and its operation is fully distributed. Moreover, it turns out that our proposed protocol is very robust; it performs well both under low load and high load conditions, and it obtains high throughput regardless of the number of stations, which in particular renders it the best choice for networks with a varying number of stations.

We also consider practical aspects of implementing our protocol in a fully distributed manner. We conducted a simulation study investigating the effect of errors in the estimation of the interference parameters between the various stations. Our results show that our protocol is robust with respect to these errors, as long as the errors are not larger than $\sim$ 40%. From an implementation point of view, the fact that our protocol is independent of the protocols applied in other levels of the network, facilitates the seamless operation of varied wireless devices connected to a unified wireless network, thus rendering it useful in various other multiple access environments.

C. Previous Work

There has been extensive work over the past decades, considering various approaches to medium access control. These approaches include random access protocols (e.g., Aloha and its variants), multipacket reception (MPR) models extending basic random access models, and game theoretical models for medium access in wireless networks, to name but a few.

Random access protocols, starting with the celebrated Aloha protocol and its variants, have been the foundation for designing cutting edge medium access protocols in a multitude of settings. Much of the work in recent years considered protocols which are based on carrier sensing, where the two most predominant paradigms are based on collision detection (CSMA/CD) and collision avoidance (CSMA/CA). In our work we consider non-carrier-sensing environments. These approaches are prevalent in designing MAC protocols for UWB environments, which are receiving much interest in recent years (see [2] for a survey, and specifically the difficulties of using CS-based approaches). Furthermore, even in more commonly deployed wireless networks (e.g., based on IEEE 802.11) carrier sensing might actually cause throughput degradation, especially in extreme high-load conditions, where performing carrier sensing is too time consuming [6]. Sensor networks are another fundamental networking environment which benefits immensely from avoiding the use of carrier sensing. In such networks one of the most scarce resources is power, and designing good medium access protocols which avoid the need to be constantly alert in performing carrier sensing (e.g., during backoff periods) might help provide better performance in such settings [7].

Issues involving selfish behavior of agents in multiple access environments have received much attention in recent years. In particular there has been a lot of work done on understanding systems' stability, throughput, and convergence to equilibrium, under stochastic assumptions on packet generation at the various stations [8]–[11]. Some of these works also considered models for interferences, rate control, and multipacket reception [12]–[17]. Our work differs substantially from the above body of research due to the fact that we assume that stations are always backlogged, thus modeling high-load conditions, which might be well beyond the stability region of classical protocols. Our main focus is on modeling and exploiting the *spatial diversity* in such environments, and providing a unified model as well as protocols which aim at maximizing the overall throughput, while guaranteeing fairness in channel

---

[1] The case where interferences are all zero, or all one, is trivial (see [5]).

access (taking into account neighboring interfering stations).

Our work is greatly motivated by the results appearing in [5], which considers a simplified game theoretical version of our model where all interferences are homogeneous, and the stations contend for channel access. They show that by using a penalization scheme, the selfish stations can be coerced into employing a randomized access control strategy whose throughput is very close to the optimal throughput obtained by a centralized scheduler.

Recent work has considered the impact of standard backoff mechanisms (such as CSMA/CA) in networks with spatially distributed nodes [18]. In this work it is shown that the classical protocols lead to substantial unfairness in channel access (due to the capture effect), and eventually lead to significant throughput reduction. The authors propose improved backoff mechanisms, in a multipacket reception (MPR) setting. Our work provides an appealing alternative to such backoff mechanisms, while maintaining both fairness in channel access, as well as improved throughput. Another recent advocation of the importance of MPR appears in [19]. The authors consider the problem of scalability in ad-hoc networks, and show that MPR combined with many-to-many communication, eventually leads to better capacity. Our work provides an alternative view and model for MPR, and we show that this view is useful in designing good protocols which provide higher throughput.

There has been much interest recently in distributed protocols enabling high throughput by avoiding collisions [20], [21]. We believe however that collisions are an inherent feature of the wireless medium, and good distributed protocols for dealing with these issues should consider the fundamental cause for collisions, which we consider to be inter-station interferences. Better understanding the role of interferences on collisions, lies in the core of designing good medium access protocols, and the solutions we propose in this paper provide an alternative view of such a quest, as well as some first steps along the path suggested by our model.

II. THE SCENARIO OF HOMOGENEOUS INTERFERENCES

In this section we consider a special case of our model, where interferences are homogeneous, i.e., there exists some network parameter $\alpha \in (0,1)$ such that for all stations i, j, $\alpha_{i,j} = \alpha$. Although this case does not seem to naturally arise in real life scenarios, it serves as a starting point in our attempt to design good protocols for the general case.

As described in the introduction, we focus our attention on single-parameter protocols, where every station i merely chooses its transmission probability R_i. Our goal is to design a protocol, or assign a transmission probability to every station, attempting at obtaining two objectives simultaneously. The first objective is maximizing the overall system's throughput, $\varphi(\overline{R})$, given the probabilities-vector $\overline{R}$. The second objective is obtaining fairness in channel access among the various stations. Clearly, a necessary condition for fairness is maintaining that for every station i, $R_i > 0$, which implies that every station has some non-zero probability of accessing the channel, and therefore it also has a non-zero probability (albeit small) of successfully transmitting a packet. However, when restricting our attention to the case of homogeneous interferences, fairness can be characterized by a much stronger necessary condition, namely, maintaining that for every two stations i, j, we have $R_i = R_j$. This implies that all stations have *the same* probabilities of both channel access, and successfully transmitting a packet. This condition is implied by the fact that in the case of homogeneous interferences, all stations are symmetric. Using the terms defined earlier, we focus our attention on uniform probabilities-vectors.

The following theorem provides a characterization of the optimal uniform probabilities-vector in this case.

Theorem 2.1: For the case of homogeneous interferences, where there is an $\alpha \in (0,1)$ such that $\alpha_{i,j} = \alpha$ for every i, j , the uniform probabilities-vector $\overline{R}$ which maximizes the overall throughput is defined by $R_i = \min\left\{\frac{1}{\alpha n}, 1\right\}$ for all i.

Proof: Let $\overline{R}$ be any uniform probabilities-vector. Denote by x the probability of transmission of every station in $\overline{R}$. By the definition of the system's throughput, we have

$$\varphi(\overline{R}) = \sum_{i=1}^{n} r_i(\overline{R}) = \sum_{i=1}^{n} R_i \prod_{j \neq i}(1-\alpha_{i,j}R_j) = nx(1-\alpha x)^{n-1}.$$

By taking derivatives with respect to x, we obtain that the optimal value of x is $x_0 = \frac{1}{\alpha n}$, thus completing our proof. ■

We refer to the probability-vector implied by Theorem 2.1 as the *optimal uniform protocol*. It might be interesting to try and evaluate the performance of the optimal uniform protocol, compared to an optimal (possibly non-uniform) probabilities-vector, which maximizes the overall throughput, and need not necessarily maintain fairness. Although this task is difficult to do for the general case, in the special case of homogeneous interferences some results obtained in [5] provide tight bounds on the performance of an optimal probabilities-vector for the case where interferences are homogeneous.

Theorem 2.2 ([5]): Given n stations, and any $k \in \{1, \ldots, n-1\}$, if for all stations i, j, $\alpha_{i,j} = \alpha$ and $\alpha \in \left[\frac{1}{k+1}, \frac{1}{k}\right)$, then any probabilities-vector $\overline{R}_k^*$ where exactly k stations choose to transmit (i.e., choose to transmit with probability 1), and exactly $n-k$ stations choose to remain idle (i.e., choose to transmit with probability 0), attains optimal throughput $\varphi(\overline{R}_k^*) = k(1-\alpha)^{k-1}$.

Figure 2 depicts the throughput obtained by the probabilities-vectors $\overline{R}_k^*$ as a function of α, for $k = 1, \ldots, 5$. The value of probability-vector $\overline{R}_k^*$ is denoted by $v_k = v_k(\alpha)$. The optimal throughput as a function of α is given by the maximal curve among all plots.

We note that the above theorem implies a simple centralized algorithm for scheduling the transmissions of the participating stations in the case where interferences are homogeneous. This algorithm, which uses a greedy approach, works as follows: Initially, no stations are chosen. Assume the algorithm has chosen some $0 \leq m \leq n$ stations to transmit. This implies that the current system's performance has value $v_m(\alpha)$. If $m < n$, and $v_{m+1}(\alpha) \geq v_m(\alpha)$, then the algorithm picks some unchosen station and adds it to the schedule. Otherwise, it terminates.

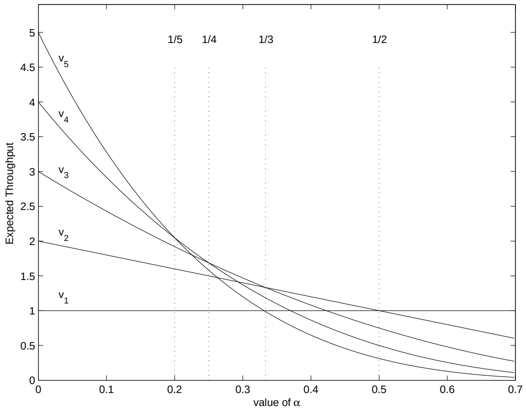

Fig. 2. Different values of $v_k = v_k(\alpha)$, up to $k = 5$.

Theorem 2.2 proves that this simple centralized strategy is guaranteed to obtain the optimal throughput possible.

Although the above algorithm is optimal, it has several drawbacks. The first obvious drawback is the fact that it relies on a central entity to manage transmissions, thus rendering such an approach inapplicable in a distributed environment. A second drawback is the fact that the above algorithm fails to satisfy even the our weakest concept of fairness, since some stations are assigned zero transmission probability. We note that in the very restrictive settings of homogeneous interferences, this can be overcome by keeping counters as to the number of times a station participates in the schedule.

The above results enable us to provide an explicit quantification of how far is the throughput obtained by an optimal uniform protocol, which guarantees fairness, from the optimal throughput possible by a centralized scheduler. The following simple lemma shows that the ratio between these two values is at most $e \sim 2.718$.

Lemma 2.3: Given homogeneous interference parameter α, let $\overline{R}$ denote the optimal uniform protocol for α, and let $\overline{R}^*$ denote an optimal probabilities-vector for α. It follows that $\varphi(\overline{R}^*) \leq e \cdot \varphi(\overline{R})$.

Proof: By Theorem 2.2 we have

$$\varphi(\overline{R}^*) = k(1-\alpha)^{k-1} \leq \frac{1}{\alpha},$$

for $\alpha \in \left[\frac{1}{k+1}, \frac{1}{k}\right)$. On the other hand, by Theorem 2.1 we have

$$\varphi(\overline{R}) = \frac{1}{\alpha}(1-1/n)^{n-1} \geq \frac{1}{e\alpha},$$

which completes the proof. ■

We emphasize that although this guarantee is the same as the guarantee given by the slotted-Aloha protocol and its variants, the model considered here is inherently different. The assumption in the model underlying Aloha-like protocols is that any two simultaneous transmissions result in the loss of all transmissions in that time slot, while in our model it is perfectly likely to have non-zero throughput even in the case of multiple simultaneous transmissions.

The above results will serve as a starting point in our attempt to design a single-parameter random access protocol which provides fairness to the various stations, for the more practical scenario where interferences are non-homogeneous.

III. PRACTICAL PROTOCOLS FOR NON-HOMOGENEOUS INTERFERENCES

In this section, we build upon the results provided in the previous section, and present a simple distributed protocol for the more realistic setting of SIMA where interferences are non-homogeneous, i.e., for every two stations i, j we have an interference parameter $\alpha_{i,j}$, reflecting the interference of j with respect to i. We first discuss the intuition underlying our protocol, and then turn to present the details of the protocol and discuss some implementation issues.

One may first consider the characteristics of a centralized approach to maximizing the throughput in an environment where interferences are non-homogeneous. This problem is essentially that of finding a probabilities-vector maximizing the overall system throughput. We note that if we limit our attention to merely finding a subset of the stations to schedule simultaneously, then the problem reduces to that of finding a maximum size independent set in an undirected graph, which is NP-hard even to approximate. It thus follows that unlike the case where interferences are homogenous, we do not expect to find a centralized algorithm obtaining the optimal throughput in the more general case.

We therefore build upon the intuition underlying the analysis presented in the previous section for homogenous interferences. We consider algorithms that aim at both providing fair channel access to the various stations, as well as aiming at maximizing the overall system's throughput. We note that unlike the strong necessary condition for fairness which was described for the case of homogeneous interferences, for the general case fairness should take into account the amount of interferences sensed, and inflicted, by a station. We therefore restrict our attention to the weaker necessary condition, which maintains that every station has some non-zero probability of transmitting. This ensures that every station has some positive probability of transmitting a packet, and no station is starved.

As might be evident, the goals of both maximizing throughput as well as maintaining fairness, or merely non-starvation, are somewhat conflicting goals. Consider, for example, one station x which interferes with a set A of n stations, such that no two stations in A interfere with each other. If $n \geq 2$ then any policy aiming at maximizing the overall throughput would always schedule the stations of A, and never schedule x. This serves as a simple example motivating our focus on protocols where every station has a non-zero probability of transmitting.

In the following subsection we present our random access protocol, which is based upon the above intuitions and design criteria. We later present a simulation study of our protocol applied to the setting of single-radio downlink mesh networks, where transmitting stations are access points which transmit to clients in their area, and must do so by sharing a single channel. We present several other prominent protocols for the problem, neither of which resorts to carrier-sensing, and compare these protocols with our newly suggested protocol.

A. *Protocol* INTERFERENCESRAND

For the purpose of illustration, consider a wireless network where multiple stations use a single radio channel. A transmission of some station A will fail due to a simultaneous transmission of another station B, if A's message is sent to a client that is also affected by B's transmission (see Figure 1). Note that this reasoning applies also if one of the receivers is itself a transmitting station; in this case if this station does not transmit it receives the transmission if it is not in range of another station that transmits, and if it does transmit, it is well within its own range so it will not receive the transmission.

In order to present our protocol and the intuition underlying it, it is convenient to consider the area covered by a station as a planar disc, and consider clients as distributed uniformly at random in this area. We note that this is not a necessity, and we use this terminology only for the sake of illustration. In such settings a transmission of station A to a client a will fail due to station B's transmission with probability proportional to the area of the intersection of the discs (the lined area in Figure 1), divided by the area of the disc of station A.[2]

Our new protocol INTERFERENCESRAND is motivated by the results appearing in the previous section; if the interferences were homogeneous, then by Theorem 2.1 the optimal uniform protocol is having every station transmit with probability $\frac{1}{\alpha n} = \frac{1}{\sum_{j=1}^{n} \alpha}$. In the case of non-homogeneous interferences, the above protocol suggests employing the following strategy: Every station i transmits with probability

$$R_i = \frac{1}{1 + \sum_{j \neq i} \alpha_{i,j}}.$$

The above choice implies, for example, that an isolated station (i.e., a station i such that $\alpha_{i,j} = 0$ for all $j \neq i$) transmits with probability 1. On the other end we might have a standard collision channel where we have $\alpha_{i,j} = 1$ for every $i \neq j$, in which case the above protocol reduces to the basic case of slotted Aloha. Protocol INTERFERENCESRAND described in Protocol 1 gives a formal definition of our protocol for the general (non-homogeneous) setting.

Protocol 1 INTERFERENCESRAND(station i)

```
Initial Setup:
for all stations $j$ that are neighbors of $i$ do
    estimate $\alpha_{i,j}$
end for
After Setup:
set $\tau_i = \frac{1}{1+\sum_{j \neq i} \alpha_{i,j}}$
at each time slot, transmit with probability $\tau_i$
```

There are several implementation issues and other practical aspects, which arise in the context of the above protocol. In particular, one of the most important aspects to consider is the estimation of the values of $\alpha_{i,j}$. We discuss these issues in Section IV-C.

B. *Additional Protocols*

In order to evaluate the performance of our new protocol INTERFERENCESRAND, we compare its performance with the performance of several broadly used protocols for multiple access environments. Specifically, we consider the following additional protocols:

1) CLUSTERIZE: This is a clustering protocol, whose variant is used, e.g., in IEEE 802.15.4 (Zigbee). Stations are divided into clusters, and in every cluster, a TDM discipline is used to determine the transmission schedule. The division into clusters is done by greedily assigning stations to clusters. A yet-unassigned station i is chosen at random and is defined to be the cluster-head. Its cluster is defined by all yet-unassigned stations j such that $\alpha_{i,j} > 0$, which are then assigned to i's cluster. This generates a *local* clustering scheme, and note that it is possible for two stations i, j corresponding to different clusters to have $\alpha_{i,j} > 0$. In each cluster stations transmit one after the other, in a Round-Robin fashion.
2) SQRTRAND: Every station transmits with probability proportional to the square root of the number of its neighbors. I.e., station i transmits with probability $\tau_i = \frac{1}{\sqrt{|\{j|\alpha_{i,j}>0\}|}}$.
3) INTERSECTRAND: Every station transmits with probability proportional to the number of its neighbors. I.e., station i transmits with probability $\tau_i = \frac{1}{|\{j|\alpha_{i,j}>0\}|}$.
4) GREEDY: All stations transmit simultaneously.
5) HALFRAND: Each station transmits with probability 0.5.

IV. SIMULATION STUDY

A. *Simulation Description*

In this section we present a simulation study comparing the performance of the various protocols presented in the previous sections in the setting of a synchronous single radio wireless mesh network, where access points are always backlogged, for the scenario of downlink traffic from the access points to randomly chosen clients in their coverage area. Our simulation study of the performance of the above protocols consists of several components. We consider access points as uniform devices, with uniform unit-size discs transmission range, centered at the access point's location. For any two access points i and j, we let the interference parameter $\alpha_{i,j}$ be the ratio between the intersection of both access points' coverage areas, and the overall coverage area of station i. As mentioned before, since in our settings the coverage area of all stations is the same, this implies that the interferences matrix defined by these parameters is symmetric (see Figure 1).

We choose a mesh network configuration by randomly and uniformly placing access points in a bounded 2-dimensional plane. Such a configuration fully defines the interferences matrix between all stations. For every access point i defined by its unit disc location, we choose a client's location uniformly at random within the access point's transmission range, and designate this location as the target of the transmission. Figure 3 shows the outline of 2 different configurations with

---

[2] If all stations have identical transmission ranges (e.g., all stations use the same devices and power levels), then the interferences are symmetric, i.e., for every i, j, $\alpha_{i,j} = \alpha_{j,i}$.

$n = 8$ access points, each with a transmission target (client) in its range. Every access point decides whether or not to transmit to its designated client, according to the protocol used. A transmission of access point i to its target client is considered successful if and only if access point i indeed transmits, and every other access point j such that i's target client is in j's range, does not transmit. Reconsidering Figure 3, if a transmission target lays in the intersection of the transmission ranges of its access point and another access point, then simultaneous transmission by both of these access points will cause the transmission to fail at the target client.

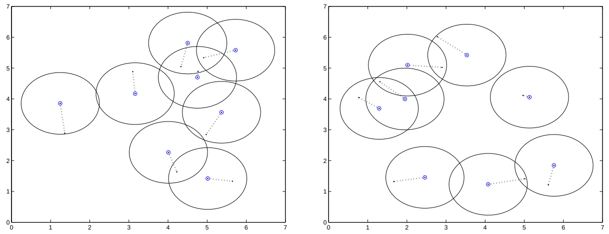

(a)

(b)

Fig. 3. Examples of different stations configurations, each with a transmission target ($n = 8$).

We measure the *throughput* of a specific protocol for a specific choice of targets by the various access points, as the ratio between the number of successful transmissions, and the overall number of access points.

Our simulations consist of checking 10 different configurations for values of n - the number of access points - ranging from 25 to 1500, over a $40{\times}40$ domain. For each configuration we conducted 20 rounds of choosing a transmission target, and verifying the successful transmissions.

B. Simulation Results

Figure 4 shows the average throughput results for protocol INTERFERENCESRAND, as well as for the additional 5 protocols considered in our study. Figure 5 supplies a high-resolution view of our results for the case where the number of stations is over 500, which appear in the marked rectangle in Figure 4. Note that a large number of stations represents high load and high interference since the portion of intersecting areas increases. Figure 6 shows the average standard deviation of the various protocols.

As can be seen, our new protocol INTERFERENCESRAND, and the GREEDY protocol, demonstrate the best performance for moderate values of n. In this range, the overall area covered by the access points is less than the overall area of the domain. In particular, the mesh does not provide network access to the overall area. For higher loads, the performance of INTERFERENCESRAND is better than the performance of all other protocols. This transition has been noted in several other smaller scale simulations we have conducted. It appears that the transition point is roughly the number of stations for which the overall area of the discs (n times π) equals the area of the domain. In other words, as the sum of transmission areas increases beyond the area of the domain, the policy implemented by our new protocol, INTERFERENCESRAND, of being aggressive in proportion with the actual amount of interference one suffers from one's neighbors, is the best strategy. Combining this with the previous observation, our simulation shows that if we are to design a wireless mesh network which completely covers a given area, our new protocol, INTERFERENCESRAND, yields the best performance. On the other hand, for a relatively small number of stations, it appears that being greedy and constantly transmitting is the best strategy. This is probably due to the relatively small average value of $\alpha_{i,j}$, when the domain area is much larger than the number of stations (which is proportional to the area covered by them). Unlike the other protocols, which either exhibit good performance for a small number of stations, or for a large number of stations, the INTERFERENCESRAND protocol performs well for both high and low load conditions. In this sense the protocol is very robust, and is clearly the best choice in cases where the system load, and coverage capabilities, varies. More specifically, our protocol obtains 10-20% better throughput than all protocols save the GREEDY protocol for the case where the load is relatively low, and 10-20% better throughput than the GREEDY protocol for high levels of load. As can be seen in Figure 6, all protocols (except HALFRAND for small values of n) have very small standard deviation, which provides some confidence as to their performance for varying number of stations.

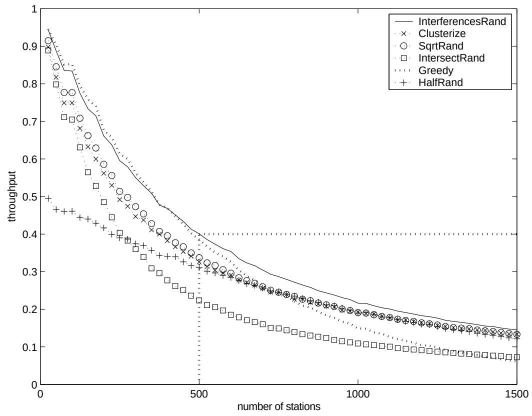

Fig. 4. Comparison of the throughput of all 6 protocols.

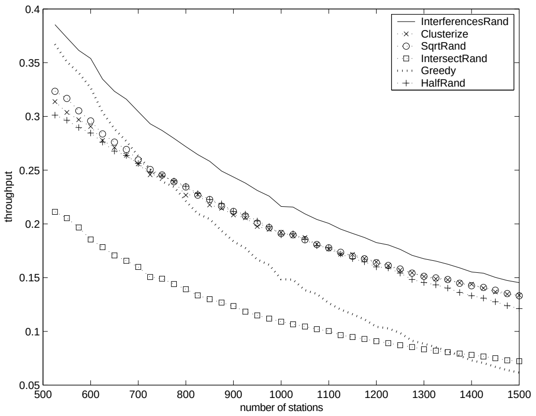

Fig. 5. High resolution view of results for high-loads.

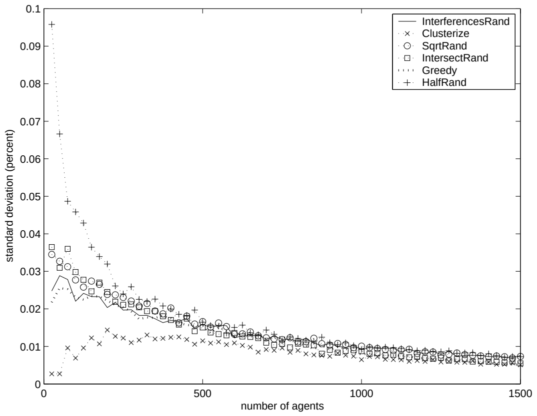

Fig. 6. Average standard deviation percentage of the protocols for different number of stations.

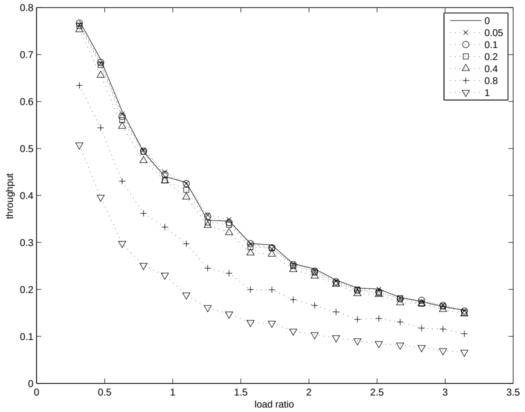

Fig. 7. The performance of protocol INTERFERENCESRAND for varying estimation-errors upper bounds.

C. Practical Aspects and Implementation Issues

An important aspect in the implementation of our new protocol is the issue of synchronization (or lack of synchronization), between the stations and the ability to have unified time slots. Throughout our study, we assume global time slots, which leads to a more tractable analysis. In practice, however, guaranteeing global synchronization is a complex and expensive task. Thus, in practical scenarios, the time slots of the different stations are not synchronized. This should not have a big impact on the performance of the system, since we did not assume collision detection and the average transmission time of all stations is similar.

Another major issue in implementing INTERFERENCESRAND is the estimation of interferences, namely, estimating the values of $\alpha_{i,j}$. It should be noted however that devices in many wireless environments, such as access points in wireless mesh networks, are stationary. It follows that this estimation can be done using common range estimation methods widely used in wireless networks for localization purposes. For example, we may use information obtained from the physical layer and estimate the distance between the stations using methods like received signal strength (RSS), time of arrival (TOA) and time difference of arrivals (TDOA) (see, e.g., [22], [23] and references therein).[3] In order to evaluate the tolerance of our protocol to estimation errors, we conducted an additional simulation study, where we allowed the protocol to err by some percentage in the estimation of its distance to its neighbors (which is therefore reflected in its estimation of the interference parameters). Such estimation errors are common in wireless environments, and in particular in dynamically changing environments. We conducted a simulation study in similar settings as the study presented in Section IV-A above. We considered several upper bounds on the allowed error of the estimation of the values of the interferences parameters, encompassed by the error in the estimation of the distance between a station and its neighbors. For each upper bound r, we considered the case where every station errs in estimating the distance to each of its neighbors by some percentage between 0 and r, chosen uniformly at random, and independently for each neighbor. Errors were allowed to be either overestimates or underestimates of the distance between the two stations, with equal probabilities. This implies that the interference matrix estimated by the stations in such a case need not be symmetric, unlike in Section IV-A.

Figure 7 shows the results comparing several upper bounds on the allowed estimation error, for various network sizes. As mentioned in the previous section, the main parameter governing the performance of all protocols was the system's load, captured by the ratio between the sum of the coverage areas of all stations, and the area of the domain. Motivated by this observation, Figure 7 shows the different performances of our protocol with the different estimation errors upper bounds, as a function of this ratio. Note that for values where the sum of the coverage areas of all stations is larger than the area of the domain, this ratio is greater than 1. As might be expected, our results indeed show that a potentially larger estimation error generally implies a degradation in the protocol's performance. However, our simulation study shows that if the estimation errors do not go beyond ∼40%, the degradation suffered by the protocol is relatively minor.

V. CONCLUSIONS AND OPEN QUESTIONS

We present a generalization of classic multiple access models (such as the one underlying Aloha) by defining a rigorous model for spatial interferences in multiple access environments. Our model takes into account the possibility that several transmissions will succeed simultaneously. This new model captures the fact that collisions are a phenomenon experienced by the receiving ends of the transmissions, and that such collisions depend on the interference level sensed by a receiver from the various simultaneous transmissions.

We use analytical results for the non-realistic case of homogeneous interferences, in order to present a simple distributed protocol for the case where interferences are non-homogeneous. We perform an extensive simulation study in the setting of wireless mesh networks, and compare the performance of our new protocol to several natural and widely used protocols for the problem. Our simulation study shows that

---

[3] Distance estimation suffices for the scenario investigated in our simulation study. In more involved scenarios, additional information may be taken into account in estimating the interference parameters, e.g., prior knowledge of client distribution, geographic data, etc.

our new protocol displays good performance for various levels of load and coverage capabilities, and outperforms all other protocols. It is also shown to be robust to estimation errors of the interferences parameters, and thus can be implemented in a fully distributed manner, even in noisy environments, with minor effects on its performance.

Our work serves as a preliminary attempt to provide a general model for interferences in wireless environments, and to better understand the relation between interferences and collisions. Some of the most interesting research directions which we believe to be of great importance in further understanding the above phenomena include:

a) Fairness vs. throughput: As mentioned earlier (e.g., Section III), there is a tradeoff between the channel access provided to each station, and the overall utilization of the wireless medium. For some network configurations these might be conflicting goals. It is of great importance to try and better understand this tradeoff.

b) Analytic results: Although our results show relatively tight results for homogeneous interferences, obtaining similar results for the non-homogeneous case seems to be a much more challenging task, mostly due to the large number of variables in such settings. A first step in such a direction might be better understanding the performance of an optimal centralized policy for such settings.

c) k-wise interferences: our proposed model only deals with pairwise interferences. In real life, however, more complicated interference relations might occur, which are not captured by our model. It would be interesting to understand the role of such high-tier interferences. A first step in this direction might be considering homogeneous k-wise interferences.

d) Game theoretical modeling: It would be interesting to explore the role of selfishness in a network model where interferences are non homogeneous. This is especially important in cases where various stations might have conflicting goals, and might not all follow a predefined protocol.

REFERENCES

[1] A. Arjona, C. Westphal, and S. Takala, “Empirical analysis of the single-radio mesh architecture: Performance limitations and challenges,” in *Proceedings of the 2007 international conference on Wireless communications and mobile computing (IWCMC)*, 2007, pp. 541–546.

[2] A. Gupta and P. Mohapatra, “A survey on ultra wide band medium access control schemes,” *Computer Networks*, vol. 51, no. 11, pp. 2976–2993, 2007.

[3] A. Koubâa, M. Alves, M. Attia, and A. V. Nieuwenhuyse, “Collision-free beacon scheduling mechanisms for ieee 802.15.4/zigbee cluster-tree wireless sensor networks,” in *Proceedings of the 7th International Workshop on Applications and Services in Wireless Networks (ASWN)*, 2007.

[4] Y.-J. Choi, S. Park, and S. Bahk, “Multichannel random access in OFDMA wireless networks,” *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 3, pp. 603–613, 2006.

[5] J. Naor, D. Raz, and G. Scalosub, “Homogeneous interference game in wireless networks,” in *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE)*, 2008, pp. 514–521 (See also the detailed Technion technical report CS–2008–11).

[6] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan, “Understanding the real-world performance of carrier sense,” in *Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis (E-WIND)*, 2005, pp. 52–57.

[7] A. Woo and D. E. Culler, “A transmission control scheme for media access in sensor networks,” in *Proceedings of the 7th annual ACM international conference on Mobile computing and networking (MOBICOM)*, 2001, pp. 221–235.

[8] A. B. MacKenzie and S. B. Wicker, “Selfish users in aloha: a game-theoretic approach,” in *Proceedings of the IEEE VTS 53rd Vehicular Technology Conference (VTC)*, Fall 2001, pp. 1354–1357.

[9] Y. Jin and G. Kesidis, “Equilibria of a noncooperative game for heterogeneous users of an aloha network,” *IEEE Communications Letters*, vol. 6, no. 7, pp. 282–284, 2002.

[10] E. Altman, R. El-Azouzi, and T. Jiménez, “Slotted aloha as a game with partial information,” *Computer Networks*, vol. 45, no. 6, pp. 701–713, 2004.

[11] R. T. B. Ma, V. Misra, and D. Rubenstein, “Modeling and analysis of generalized slotted-aloha MAC protocols in cooperative, competitive and adversarial environments,” in *Proceedings of the 26th IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2006, p. 62.

[12] A. B. MacKenzie and S. B. Wicker, “Stability of multipacket slotted aloha with selfish users and perfect information,” in *Proceedings of IEEE INFOCOM 2003, The 22nd Annual Joint Conference of the IEEE Computer and Communications Societies*, 2003, pp. 1583–1590.

[13] M. M. Carvalho and J. J. Garcia-Luna-Aceves, “A scalable model for channel access protocols in multihop ad hoc networks,” in *Proceedings of the 10th annual international conference on Mobile computing and networking (MOBICOM)*, 2004, pp. 330–344.

[14] K.-H. Hui, W.-C. Lau, and O.-C. Yue, “Characterizing and exploiting partial interference in wireless mesh networks,” in *Proceedings of the IEEE International Conference on Communications (ICC)*, 2007, pp. 102–108.

[15] L. Qiu, Y. Zhang, F. Wang, M. K. Han, and R. Mahajan, “A general model of wireless interference,” in *Proceedings of the 13th annual ACM international conference on Mobile computing and networking (MOBICOM)*, 2007, pp. 171–182.

[16] I. Menache and N. Shimkin, “Decentralized rate regulation in random access channels,” in *Proceedings of IEEE INFOCOM 2008, The 27nd Annual Joint Conference of the IEEE Computer and Communications Societies*, 2008, pp. 394–402.

[17] ——, “Efficient rate-constrained nash equilibrium in collision channels with state information,” in *Proceedings of IEEE INFOCOM 2008, The 27nd Annual Joint Conference of the IEEE Computer and Communications Societies*, 2008, pp. 403–411.

[18] W. F. K. Guner D. Celik, Gil Zussman and E. Modiano, “MAC for networks with multipacket reception capability and spatially distributed nodes,” in *Proceedings of IEEE INFOCOM 2008, The 27th Annual Joint Conference of the IEEE Computer and Communications Societies*, 2008, pp. 1436–1444.

[19] J. J. Garcia-Luna-Aceves, H. R. Sadjadpour, and Z. Wang, “Challenges: towards truly scalable ad hoc networks,” in *Proceedings of the 13th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, 2007, pp. 207–214.

[20] E. Modiano, D. Shah, and G. Zussman, “Maximizing throughput in wireless networks via gossiping,” in *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/Performance)*, 2006, pp. 27–38.

[21] A. Brzezinski, G. Zussman, and E. Modiano, “Distributed throughput maximization in wireless mesh networks via pre-partitioning,” *IEEE/ACM Transactions on Networking*, vol. 16, no. 6, pp. 1406–1419, 2008.

[22] N. Patwari, I. Alfred O. Hero, M. Perkins, N. S. Correal, and R. J. O’Dea, “Relative location estimation in wireless sensor networks,” *IEEE Transactions on Signal Processing*, vol. 51, no. 8, pp. 2137–2148, 2003.

[23] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki, “Practical robust localization over large-scale 802.11 wireless networks,” in *Proceedings of the 10th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, 2004, pp. 70–84.