

Interference-aware Proportional Fairness for Multi-Rate Wireless Networks

Douglas M. Blough, Giovanni Resta, Paolo Santi

Abstract—In this paper, we consider how proportional fairness in wireless networks is impacted by spatial reuse and the interference it produces. We observe that, in scenarios where spatial reuse is possible (e.g., in high-density WLAN environments), the classic notion of time-based proportional fairness can be severely impacted: some users might experience very large interference penalties while other users might get larger bandwidth proportions than what they would have received with time-based proportional fairness and no spatial reuse. To account for this, we introduce the concept of interference-aware STDMA time-based proportional fairness (*i*-STPF), and compare it to ordinary STDMA time-based proportional fairness (STPF). We present an *ε*-*i*-STPF scheduling algorithm, and prove that it approximates the time-based fair bandwidth allocation (up to a small positive constant ϵ), while providing an aggregate throughput that is within a constant factor from optimal. We also present a heuristic *i*-STPF scheduling algorithm and compare it through simulation to a similar heuristic STPF scheduler, and to an interference-aware, rate-based scheduler. The results show that the *i*-STPF scheduler: *i*) achieves excellent aggregate throughput about 35% higher than rate-based throughput; and *ii*) maintains a close approximation to time-based fairness without interference.

I. INTRODUCTION

Whenever resources are allocated to users with different capabilities in a system or network, there is an inherent trade-off between maximizing the overall performance of the system and allocating the resources fairly among the users. Typically, there is an attempt to maximize overall performance subject to the constraint that each user should receive at least some minimum allocation. Different ways to define the constraint give rise to different notions of fairness [7]. In networks, proportional fairness is often used, where allocations are dependent (typically, inversely proportional) on the users' resource requirements. Resources taken from a single heavy resource user can potentially benefit numerous low-resource users, thereby boosting the overall performance significantly while still guaranteeing minimum allocations to all users.

In the specific case of multi-rate wireless networks, where links with different data rates coexist, two notions of proportional fairness have been defined. In *rate-based proportional fairness* (a.k.a. throughput fairness), each user is given an equal number of communication opportunities for transmitting an equal amount of data at each opportunity. It is well-known that, when applied to IEEE 802.11 networks, rate-based fairness causes a “performance anomaly”, according to which the throughput of all users is degraded to match the throughput of the user with the lowest quality channel [3]. To solve this anomaly, the notion of *time-based proportional fairness*¹ has

been introduced, according to which users are provided equal amounts of time to use the channel, and their bandwidths then depend on the overall number of users and their link data rates. Thus, the allocated bandwidth share is inversely proportional to the resource requirement (the air time for a fixed amount of data) and the approach is proportionally fair. Time-based proportional fairness is now widely accepted as the most appropriate notion of proportional fairness for multi-rate wireless networks [2], [14], [20].

In this paper, we are interested in studying how proportional fairness in multi-rate wireless networks is impacted by *interference*. In practice, two wireless links separated by a sufficient distance can communicate simultaneously. This is referred to as *spatial reuse* of the channel and it is beneficial whenever the aggregate rate of the two links (factoring in their mutual interference) exceeds the average rate of the two links communicating separately. However, it is not at all clear what fairness property is maintained, if any, in this situation. Since interference impacts are typically not symmetric, one link's data rate might drop substantially while the other's might be hardly impacted. A severe interference penalty on one user, which is in fact caused by the resource management system itself in a scheduled (STDMA) environment, cannot be justified by any reasonable fairness model known to us.

In this paper, we introduce a notion of *interference-aware proportional fairness* that can be applied to any multi-rate wireless network operating in the same wireless channel, under the constraint that all flows are single hop. The notion of interference-awareness can be applied to extend both rate-based and time-based fairness, leading to the notion of Interference-aware, STDMA Time-based Proportional Fairness (*i*-STPF) and Interference-aware, STDMA Rate-based Proportional Fairness (*i*-SRPF), respectively. The following question then naturally arises: are the throughput benefits provided by time-based fairness vs. rate-based fairness still present when interference is incorporated in the definition of fairness? A major contribution of this paper is investigating this question, disclosing that *time-based fairness is superior to rate-based fairness also in presence of interference*. More specifically, the results of our evaluations quantify this throughput benefit in approximately 35%.

Another major contribution of this paper is introducing an interference-aware, time-based fair algorithm called *ε*-*i*-STPF, which is shown to approximate the time-based fair bandwidth allocation (up to a small positive constant ϵ), while providing an aggregate throughput that is within a *constant factor* from optimal. We also present a heuristic *i*-STPF scheduling algorithm, called GREEDYINTIMEFAIR

D. Blough's research was supported in part by the National Science Foundation under Grants CNS-1017248 and CNS-1319455.

¹In the rest of this paper we use the term “rate-based fairness” as a shorthand of “rate-based proportional fairness”. Similarly for time- fairness.

and compare it through simulation to a similar heuristic STPF scheduler, GREEDYTIMEFAIR, which disregards interference when allocating transmission time to users. The results show that GREEDYINTTIMEFAIR achieves an aggregate throughput that is only about 4–8% lower than GREEDYTIMEFAIR, but with much better time-based proportional fairness: GREEDY-INTTIMEFAIR only deviates from time-based proportional fairness without spatial reuse by about 3%, while GREEDY-TIMEFAIR’s deviation is as high as 40%.

While the notion of interference-aware fairness introduced in this paper can be applied to any multi-rate wireless network with single hop flows, the presentation of the main results and simulation experiments are tailored to a typical interference environment with multiple interfering WLANs. This situation is common today where high-density WLAN deployments abound in urban areas and enterprises [13].

II. RELATED WORK

Fair scheduling in wireless networks has been extensively studied in the past. We do not mention here work on cellular networks (see, e.g., [6] and references therein), where the problem is to choose the best mobile user in a cell at each time slot, since spatial reuse, and, hence, interference, is not considered.

A line of research related, but different, to our work is the one in which proportionally fair AP/user allocation is investigated. In [8], the authors consider this problem in a multi-rate WLAN setting where multiple APs coexist and form a network. However, adjacent APs are assumed to be assigned orthogonal channels, and, hence, interference is not considered. In [9], the authors consider the same problem in a setting where APs operate on the same channel, and interference is modeled by the SINR interference model. In our network model, we assume that user/AP association is already taken care of (e.g., by one of the mentioned approaches), while the problem we face is building a fair STDMA schedule for the given user/AP association.

In [15], [21], the authors consider the wireless link scheduling problem subject to different fairness criteria, and, in [21], also considering energy efficiency criteria. Interference between links is modeled by the SINR interference model. The problem is formulated as a linear problem for max-min fairness, while it becomes a non-linear optimization problem if proportional fairness is considered. Unfortunately, non-linear optimization problems are hard to solve for even a moderate number of links. On the contrary, all scheduling algorithms presented in this paper have polynomial time complexity, and can be used to build schedules in large networks.

In [19], the authors consider the problem of achieving proportionally fair allocation of end-to-end bandwidth in presence of spatial reuse. They present a proportionally fair distributed allocation algorithm, and model interference between links by the SINR model. However, they only consider single-rate links in the network model. Thus, the approach of [19] cannot be applied in popular multi-rate WLANs.

The work that is closest in spirit to ours is [20], where the authors consider two notions of fairness (rate-based fairness and time-based fairness, as defined in the Introduction), and

show that time-based fairness provides considerably higher throughput than rate-based fairness in multi-rate WLANs. However, in [20], the authors consider a single AP WLAN, i.e., a typical TDMA setting where spatial reuse, and, hence, interference, is not considered.

To our knowledge, no prior work has investigated the effect of interference on time-based proportional fairness in an STDMA setting with multi-rate links.

III. NETWORK AND INTERFERENCE MODELS

A. Problem Setting

We consider a scenario in which one-hop wireless networks are densely deployed over a region. In this scenario, the areas served by different access points (APs) can overlap and interference between nodes in neighboring regions must be considered. This scenario covers several practical deployment types, including but not limited to the following: 1) institutional 802.11 WLANs where different basic service set (BSS) areas overlap (often referred to as the overlapping BSS problem); 2) multiple independent WLANs deployed in the same general area, for example a commercial area or an apartment complex; and 3) multiple femtocells deployed by different users within a cellular dead area. Note that, in general, the problem of interference between neighboring areas cannot be solved solely through the use of multiple channels. For example, in IEEE 802.11b/g, there are only three orthogonal channels, which are not enough to eliminate interference in dense deployments.

Let AP_1, \dots, AP_k denote the co-located, interfering APs. For each AP_i , there are n_i users, denoted $u_{1,i}, \dots, u_{n_i,i}$. In total, there are $n = \sum_i n_i$ users in the network. In the following, we use the term *transmitter node* (*node* for short) to refer to either an AP or a user that is acting as transmitter.

The policy used for allocating users to APs is outside the scope of this paper: in what follows, we assume that user/AP allocation is predetermined. Users are assumed to be stationary for a period of time in between movements. The durations of stationary periods are assumed to be relatively long compared to a scheduling period. This scenario is consistent with a typical home/office environment, or with an urban scenario where APs are deployed in establishments such as coffee houses, bars, and restaurants.

Each user $u_{j,i}$ is characterized by a data rate $dr_{j,i}$ experienced on the link to/from AP_i in *absence of interference*. This *interference-free data rate* depends on several factors such as distance to the AP, radio propagation environment, etc. Given the assumed semi-stationary setting, in what follows we assume that user interference-free data rates remain fixed throughout a scheduling period.

User data rates *in presence of interference* from other nodes are computed according to the *graded SINR model*, originally proposed in [10], [11] and formally defined in [17]. Without loss of generality, consider a user $u_{j,1}$ connected to AP_1 . According to this model, the data rate experienced by $u_{j,1}$ when nodes in the set $\mathcal{T} = \{v_2, \dots, v_k\}$ are transmitting simultaneously is given by

$$dr_{j,1}(\mathcal{T}) = f(\text{SINR}(j, 1, \mathcal{T})) ,$$




Fig. 1. Rate function in the graded SINR interference model.

where f is a non-decreasing rate function, and $SINR(j, 1, \mathcal{T})$ is the SINR experienced by user $u_{j,1}$ when nodes in set \mathcal{T} are transmitting.² Note that nodes in \mathcal{T} can be either users (for uplink communications) or APs (for downlink communications). When the set \mathcal{T} is empty, the data rate experienced by user $u_{j,1}$ equals the data rate in absence of interference, i.e., $dr_{j,1}(\emptyset) = dr_{j,1}$. The interference-inclusive data rate for user $u_{j,i}$ connected to an arbitrary AP_i is denoted by $dr_{j,i}(\mathcal{T})$, for a transmitting set \mathcal{T} , and is defined similarly.

While the notion of interference-aware proportional fairness introduced in the following is independent of the specific rate function used, in this paper we will consider two specific rate functions.

The first function, denoted f_s , is a staircase function, with discrete data rate values depending on the experienced SINR value as reported in Table I [5].

Data rate	Min SINR (dB)	Data rate	Min SINR (dB)
6 Mbps	6	24 Mbps	17
9 Mbps	8	36 Mbps	19
12 Mbps	9	48 Mbps	24
18 Mbps	11	54 Mbps	25

TABLE I
MINIMUM SINR VALUES FOR 802.11A/G DATA RATES.

The staircase function f_s , while adherent to practical WLAN settings, is not apt to algebraic manipulation. For simplifying mathematical derivations, when deriving the theoretical performance bound provided by the ϵ_i -STPF scheduling algorithm introduced in Section VI we used another data rate function. This function, denoted f and depicted in Figure 1, is defined as follows: it is 0 if the SINR (expressed in dB) is below a minimum value $\beta_0 = 0$; it is equal to dr_{max} if the SINR is above a maximum value β_{max} ; and it is an increasing function of the SINR in the $[\beta_0, \beta_{max}]$ interval. This is consistent with practical wireless networks in which the data rate is 0 when signal quality is too low, and cannot exceed a certain maximum data rate even with excellent signal quality. For definiteness and in accordance with [17], f is assumed to be increasing linearly between 0 and dr_{max} in the $[\beta_0, \beta_{max}]$ interval, which is consistent with the classic Shannon's information rate in the $[\beta_0, \beta_{max}]$ interval. More specifically, we have $f(x) = \frac{dr_{max}}{\beta_{max}} \cdot x$ for $x \in [0, \beta_{max}]$.³ The values of β_{max} and dr_{max} depend on the technology at hand (e.g., $dr_{max} = 54Mbps$ in IEEE 802.11g, with a typical β_{max} value of 25dB [5]).

²In this paper, we do not consider transmission power control. Thus, a given transmitter node always uses the same power and the SINR is determined once the set of simultaneous transmitters is known.

³Note that the assumption that $\beta_0 = 0$ is made only to simplify mathematical derivations in the following. Up to simple algebraic manipulations, the results presented in this paper are valid for any β_0 such that $0 < \beta_0 < \beta_{max}$.

B. Practical Issues with Problem Setting

Our primary goal in this paper is to provide a conceptual framework for illustrating the problem with traditional fairness metrics when considering interference and for demonstrating that alternative fairness criteria which account for interference are possible and efficiently achievable. Thus, we are not primarily concerned with providing a complete solution and resolving all practical issues with implementing such a solution in a specific network setting. Nevertheless, since this preliminary study uses the case of multiple overlapping WLANs as the driving example, we provide a brief discussion of how interference-aware proportional fairness might be applied in that scenario.

To achieve interference-aware fairness of the type discussed herein requires cooperation/coordination among multiple APs that are operating on the same channel and producing sufficient interference to degrade performance. In enterprise settings, this coordination could come from a centralized controller. Such enterprise WLAN controllers have become popular in recent years and several companies rely on this model for a substantial portion of their business, e.g. Meru Networks [12]. Such a centralized controller could compute a schedule that it distributes to all APs and also synchronize the APs for the purposes of carrying out the scheduled transmissions. Distributed AP coordination is also feasible, e.g. Aerohive Networks HiveOS operating system [4] allows multiple APs to form a hive for the purposes of coordinating management of network resources in a distributed fashion. Such functionality would allow a set of APs to agree upon a schedule and provide loose synchronization for supporting execution of the schedule, perhaps facilitated by multi-packet link transmissions with a single block ACK to amortize any synchronization overhead.

There is also the question of what information is needed for interference-aware scheduling to be carried out. As will be seen later when scheduling algorithms are discussed, it is necessary to know the rates on the links that are being scheduled in both the interference-free case and with different possible combinations of interferers. One possibility would be for APs to carry out interference measurements similar to those reported in [10] to obtain this information. Such measurements would have to be repeated periodically to deal with user mobility⁴, environmental changes, etc., and the networks would have to operate in a no-spatial-reuse mode during measurement periods. Another possibility would be to start off by using simple interference models to estimate rates and then have APs measure actual rates that occur as different link combinations are used. The measured rates could then replace the estimated values for those specific combinations in the next scheduling period. Gradually, over time, more measured and fewer estimated rate values would be used and the performance achieved by the approach should be comparable to what would be achieved with perfectly accurate rate information.

In terms of the complexities involved with the approach, the target scenario is likely to involve tens of APs and tens

⁴Note that in typical indoor WLAN settings, many users are stationary for significant periods of time and even when walking, velocity is quite low.

of clients per AP. With numbers in that range, the number of measurements that would have to be carried out and the amount of information to be communicated to the scheduler would be quite small and the computation time for the scheduling algorithms discussed later would be negligible.

IV. INTERFERENCE-AWARE PROPORTIONAL FAIRNESS

Our focus is on building a proportionally fair schedule for high-density wireless network deployments. In particular, we are interested in achieving time-based proportional fairness [20], which has been shown in [2], [14], [20] to provide substantial throughput benefits with respect to rate-based fairness in multi-rate WLANs.

It is important to observe that the original definition of time-based proportional fairness refers to a scenario in which a single AP is present in the network, and the problem is scheduling transmissions of the users associated with that specific AP, with only a single user active in each transmission slot. This scenario corresponds to a typical TDMA setting, where a single user is scheduled for transmission in each slot and interference does not occur. In the scenario considered in this paper, though, there are multiple co-existing APs operating on the same channel, and the problem of how to deal with concurrent transmissions comes into play.

One obvious way of dealing with multiple co-existing APs is to use TDMA across all APs, such that, at each slot, a single active user is selected in the network, and only the transmission between the selected user and the respective AP takes place. While this solution allows a straightforward generalization of the notion of time-based proportional fairness to the multiple AP setting, the obtained aggregate throughput is likely to be low. In fact, it is well known in the literature that spatial-TDMA (STDMA), in which *multiple* transmissions take place in the same transmission slot subject to interference constraints, provides higher throughput than TDMA.

Based on the above, we need to generalize the notion of time-based proportional fairness to the STDMA setting. We now introduce two definitions of time-based proportional fairness in STDMA networks. The first definition is an immediate extension of the corresponding fairness notion in a TDMA network. Formally, time-based proportional fairness is achieved in a STDMA network when the following condition is satisfied:

STDMA time-based proportional fairness (STPF): Let $T = \{T_1, \dots, T_t\}$ be the scheduling period composed of t transmission slots of equal duration⁵, $L_h = \{l_1, \dots, l_{n_h}\}$ be the set of links scheduled in slot T_h , $\mathcal{T}_h = \{v_1, \dots, v_{n_h}\}$ be the set of transmitters of those links, and $\mathcal{U}_h = \{u_1, \dots, u_{n_h}\}$ be the set of users associated with those links. STPF is achieved whenever each user in the network appears in T exactly once. Formally,

$$\forall u_{j,i}, \exists \mathcal{U}_h \text{ such that } u_{j,i} \in \mathcal{U}_h \text{ and } u_{j,i} \notin \mathcal{U}_k, \quad k \neq h.$$

The above notion of STPF gives an equal share of the channel occupancy time to each user in the network, so it is

⁵To simplify presentation, in the following we assume slot duration is normalized to 1, so that the amount of data transmitted in a slot on a link equals its data rate.

apparently fair. However, in STDMA multiple users are scheduled in the same slot. This means that, during its transmission slot, user $u_{j,i}$ does not experience its interference-free data rate $dr_{j,i}$, but it experiences a *degraded* data rate $dr_{j,i}(\mathcal{T}_h - \{v_{j,i}\})$, where h is the index of the slot in which user $u_{j,i}$ is scheduled for transmission and $v_{j,i}$ is the transmitter associated with $u_{j,i}$'s link. Thus, the actual portion of bandwidth that $u_{j,i}$ receives is not proportional to $dr_{j,i}$, but to the interference degraded data rate $dr_{j,i}(\mathcal{T}_h - \{v_{j,i}\})$. Some users can, in fact, experience an excessive interference penalty, where their bandwidth share is reduced substantially based on the network scheduling algorithm, rather than being based on the inherent conditions of their links. It seems then reasonable to define a notion of time-based proportional fairness, which is aimed at giving users a share of the available bandwidth that is *proportional to the interference-free data rate* $dr_{j,i}$, and not to the degraded data rate $dr_{j,i}(\mathcal{T}_h - \{v_{j,i}\})$ as done in STPF. This leads us to the following notion:

Interference-aware STDMA time-based proportional fairness (*i*-STPF): Assume the same definitions as in STPF. active in multiple slots. Let $d_{j,i} = c \cdot dr_{j,i}$ be the *virtual demand* associated with user $u_{j,i}$, where $c > 0$ is an arbitrary constant, and let $d_{j,i}^h$ be the amount of virtual demand satisfied in transmission slot T_h in which $u_{j,i}$ is active, where $d_{j,i}^h = c \cdot dr_{j,i}(\mathcal{T}_h - \{v_{j,i}\})$. *i*-STPF is achieved when the virtual demand of each user is satisfied at the end of the scheduling period T . Formally:

$$\forall u_{j,i}, \sum_{h: u_{j,i} \in \mathcal{U}_h} d_{j,i}^h = d_{j,i}.$$

It is immediate to see that under *i*-STPF, at the end of the scheduling period each user gets a share of the bandwidth that is equal to its virtual demand divided by the total virtual demand, i.e. bandwidth is allocated fairly based on the interference-free data rates.

In practice, achieving exact *i*-STPF might not be possible. In fact, only a small number of data rates can be used, e.g., in IEEE 802.11g, and guaranteeing that the satisfied demand at the end of the scheduling period equals the virtual demand for each user might be impossible. For this reason, we introduce the notion of ϵ -approximate *i*-STPF:

ϵ -approximate *i*-STPF (ϵi -STPF): Assume the same definitions as in *i*-STPF; ϵi -STPF is achieved when the virtual demand of each user is approximately satisfied at the end of the scheduling period T . Formally:

$$\forall u_{j,i}, \sum_{h: u_{j,i} \in \mathcal{U}_h} d_{j,i}^h \in [d_{j,i}(1 - \epsilon), d_{j,i}(1 + \epsilon)].$$

The following example clarifies the different notions of fairness considered in this paper. In Figure 2, we have a network composed of 4 co-existing APs, each with a single associated user. In case of TDMA, each transmission slot accommodates a single user. With rate-based proportional fairness (RPF), each user gets 25% of the bandwidth allocation. This can be accomplished by giving each of the 54 Mbps users two time slots, the 12 Mbps user 9 time slots, and the 2 Mbps user 54 time slots to comprise one scheduling period. The aggregate throughput in this case is $(\frac{4}{67} \times 54) + (\frac{9}{67} \times 12) + (\frac{54}{67} \times 2) \approx$

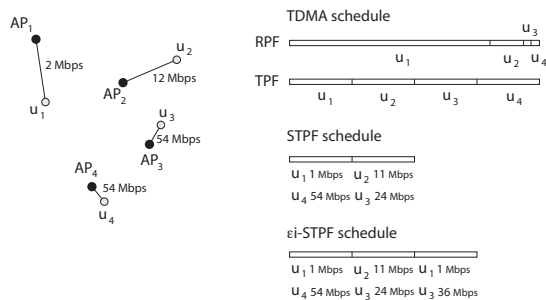


Fig. 2. Example of the different fair scheduling policies. We assume a downlink scenario. The interference-free data rates are reported as labels on the links. Link data rates in presence of interference are reported in the respective schedules.

6.4 *Mbps*. With time-based proportional fairness (TPF), each user is assigned an equal share of transmission time. The corresponding aggregate throughput is:

$$\frac{2 + 12 + 54 + 54}{4} = \frac{122}{4} = 30.5 \text{ Mbps} .$$

The resulting fair bandwidth allocation for each user is:

$$u_1 = \frac{2}{122} = 1.64\%, \quad u_2 = \frac{12}{122} = 9.84\% ,$$

$$u_3 = u_4 = \frac{54}{122} = 44.26\% .$$

In case of STDMA, we have two possible notions of fairness. According to STPF, each user must get an equal share of transmission time, possibly with a reduced data rate due to interference with a concurrent transmission. The resulting schedule, reported in Figure 2, is composed of two equal-length slots, accommodating two transmissions each. Although the data rates on the single links are reduced in three out of four links, spatial reuse substantially increases aggregate throughput w.r.t. to the TDMA scenario. In fact, the aggregate throughput is:

$$\frac{1 + 54 + 11 + 24}{2} = \frac{90}{2} = 45 \text{ Mbps} ,$$

which is about 50% higher than the TDMA throughput with TPF. However, the bandwidth allocation resulting from STPF scheduling is far from the time-based fair allocation. In fact, the portions of bandwidth allocated to users are as follows:

$$u_1 = \frac{1}{90} = 1.11\%, \quad u_2 = \frac{11}{90} = 12.22\% ,$$

$$u_3 = \frac{24}{90} = 26.67\%, \quad u_4 = \frac{54}{90} = 60.00\% .$$

Notice in particular that users u_3 and u_4 , which have the same interference-free data rate and hence should receive the same portion of bandwidth according to time-based proportional fairness, receive instead very different amounts of bandwidth, with user u_4 unfairly getting 60% of the available bandwidth, and user u_3 receiving only 26.67% of the bandwidth.

This unfairness caused by wireless interference is removed using ϵi -STPF scheduling. A virtual demand equal to the interference-free data rate is assigned to each user, and the

schedule is built with the purpose of satisfying virtual demands, up to $\epsilon = \frac{1}{8}$. The resulting schedule is reported in Figure 2. The schedule results in an aggregate throughput of:

$$\frac{2 + 11 + 60 + 54}{3} = \frac{127}{3} = 42.3 \text{ Mbps} ,$$

which is only slightly lower than the aggregate throughput with STPF scheduling, and about 40% higher than the TDMA throughput with TPF. The resulting bandwidth allocation is however much fairer in a time-based sense, approximating i -STPF up to $\epsilon = \frac{1}{8}$. The portions of bandwidth allocated to users with ϵi -STPF scheduling are as follows, and are very close to the time-based fair allocation:

$$u_1 = \frac{2}{127} = 1.57\%, \quad u_2 = \frac{11}{127} = 8.66\% ,$$

$$u_3 = \frac{60}{127} = 47.24\%, \quad u_4 = \frac{54}{127} = 42.53\% .$$

The salient features of the considered proportional fair scheduling policies are summarized in Table II.

Policy	High Thr	Time-based Fairness
RPF	no	no
TPF	no	yes
STPF	yes	no
ϵi -STPF	yes	yes

TABLE II
MAIN FEATURES OF THE DIFFERENT FAIR SCHEDULING POLICIES.

V. FAIR SCHEDULING

In this section, we formally define the fair scheduling problems considered in this paper. To simplify notation, we consider in this section a set $L = \{l_1, \dots, l_n\}$ of wireless links to schedule, with corresponding interference-free data rates dr_1, \dots, dr_n . Links are either in downlink (from AP to user) or in uplink (from user to AP) direction. In what follows we make the reasonable assumption that each AP can accommodate only a constant number of users, which implies that the number of APs in the network is $\Theta(n)$, where n is the total number of links to schedule.

For a specific link $l_i \in L$, we denote by t_i and r_i the transmitter and receiver on the link, respectively. Notice that transmitters and receivers of links in L are not necessarily distinct. In fact, APs are typically involved in many communications as either transmitter or receiver. As we shall see, our scheduling algorithms guarantee that both APs and user nodes are involved in a single communication (as either transmitter or receiver) in each transmission slot.

Given a subset $L' \subseteq L$ of active links, $dr_{i,L'}$ denotes the data rate on link l_i when all nodes in L' are simultaneously transmitting, i.e.,

$$dr_{i,L'} = f(\text{SINR}(i, L')) ,$$

where

$$\text{SINR}(i, L') = \frac{P_{ii}}{N + \sum_{j \in L', j \neq i} P_{ji}} ,$$

where P_{ii} is the received power at r_i of the signal transmitted by t_i , and P_{ji} is the received power at r_i of the interfering signal transmitted by t_j .

Before formally introducing the fair scheduling problems, we need to define a traffic model. In accordance with [20], we adopt the *fluid traffic model*, according to which all active flows (corresponding to the n links in our setting) continuously transfer infinite streams of bits. In other words, flows (links) are assumed to be continuously backlogged.

Definition 1 (STPF-scheduling): Given the set of links L , and an arbitrary duration $\tau > 0$ of each transmission slot, find a schedule $\mathcal{S} = \{S_1, \dots, S_T\}$ such that:

- i) no node (either user or AP) is involved in more than one communication in the same transmission slot (primary interference constraint);
- ii) each link appears in \mathcal{S} exactly once, formally:

$$\forall l_i \in L, \exists S_j \in \mathcal{S} \text{ such that } l_i \in S_j \text{ and } l_i \notin S_k, k \neq j$$

- iii) the aggregate amount of data transmitted per unit of time is maximized, formally:

$$\max \left\{ \text{Thr}(\mathcal{S}) = \frac{\sum_i dr_{i,S_{j(i)}}}{\tau T} \right\},$$

where $S_{j(i)}$ represents the set of active links in the slot in which link l_i is scheduled for transmission.

Definition 2 (ϵi -STPF-scheduling): Given a set of links L , each with a virtual demand d_i equal to the interference-free data rate⁶ dr_i on link l_i , which is assumed to be an arbitrary integer > 0 , and an arbitrary duration $\tau > 0$ of each transmission slot, find a schedule $\mathcal{S} = \{S_1, \dots, S_T\}$ such that:

- i) no node (either user or AP) is involved in more than one communication in the same transmission slot (primary interference constraint);
- ii) the virtual demand on each link is exhausted up to ϵ , formally:

$$\forall l_i \in L, \sum_{j: l_i \in S_j} d_{ij} \in [(1 - \epsilon)d_i, (1 + \epsilon)d_i],$$

where $d_{ij} = \tau dr_{i,S_j}$ is the virtual demand satisfied on link l_i when scheduled for transmission in slot S_j .

- iii) the aggregate amount of data transmitted per unit of time is maximized, formally:

$$\max \left\{ \text{Thr}(\mathcal{S}) = \frac{\sum_i \sum_{j: l_i \in S_j} dr_{i,S_j}}{\tau T} \right\}.$$

Thus, the goal of our scheduling algorithms is building an STPF or an ϵi -STPF schedule (condition ii)) with maximum aggregate throughput (condition iii)).

VI. SCHEDULING ALGORITHMS

In this section, we present a fair scheduling algorithm with proven approximation bounds that hold under the assumption of log-distance radio propagation. More specifically, in the following we assume that the radio signal power at distance d from the transmitter is given by P/d^α , where P is the transmission power (the same for all transmitters) and $\alpha > 2$ is the path loss coefficient [16]. Up to straightforward technical details, the presented approximation bounds apply also to more

⁶I.e., we set to 1 the proportionality constant c used to define the virtual link demand.

Algorithm ϵi -STPF:

Input: A set $L = \{l_1, \dots, l_n\}$ of n links
Output: An ϵi -STPF fair schedule S_1, \dots, S_T

1. for each $l_j \in L$, add dr_j copies of l_j in multiset L_M
2. $t = 0$
3. Partition links in L_M into classes L_1, \dots, L_k as defined in (1)
4. for each $L_i \neq \emptyset$, with $1 \leq i \leq k$
5. Partition network deployment region into squares of width $\mu_i \cdot D_{i+1}$
6. 4-color the squares such that no two adjacent squares have the same color
7. for $h = 1, \dots, 4$
8. Select color h
9. repeat
10. For each square A of color h , choose a link $l_j \in L_i$ with receiver in A ; $L_h^i = L_h^i \cup \{l_j\}$
11. $t = t + 1$; $S_t = L_h^i$
12. set duration of slot S_t to $1/f((1 + \eta)^{i-1} \beta_Q)$
13. until all links of L_i in selected squares are scheduled
14. set $T = t$
15. return S_1, \dots, S_T

Fig. 3. The ϵi -STPF scheduling algorithm.

general radio propagation models, such as the model used in [18] which is shown to closely approximate log-normal shadowing. Furthermore, in the following we assume that the interference-free data rates are arbitrary integers, and that the data rate function is function f as depicted in Figure 1.

A. ϵi -STPF scheduling

The algorithm for building an ϵi -STPF schedule is derived from the *GradedSINR* algorithm introduced in [17]. Let $dr_{min} > 0$ be the minimum interference-free data rate of the links in L , and let $\beta_Q \geq 1$ be the SNR value such that $f(\beta_Q) = dr_{min}$. Links to be scheduled are partitioned into disjoint classes L_1, \dots, L_k , with links in the same class having similar interference-free data rates and, hence, SNR values. More specifically, link class L_i , with $i = 1, \dots, k$, contain links l_j that satisfy:

$$(1 + \eta)^{i-1} \beta_Q \leq SNR_{l_j} < (1 + \eta)^i \beta_Q, \quad (1)$$

where η is an arbitrary constant such that $1/7 \leq \eta < 1$ and

$$k = \lfloor \log_{1+\eta}(P/\beta_Q N) \rfloor + 1.$$

It is important to observe that k is a constant which does not depend on the number n of links to be scheduled. Furthermore, it is immediate to see that under our working assumption of log-distance radio propagation with path loss exponent $\alpha > 2$, links in the i -th SNR class have lengths that satisfy:

$$D_{i+1} = \left(\frac{P}{(1 + \eta)^i \beta_Q N} \right)^{\frac{1}{\alpha}} < < L_i \leq \left(\frac{P}{(1 + \eta)^{i-1} \beta_Q N} \right)^{\frac{1}{\alpha}} = D_i.$$

The INTIMEFAIR scheduling algorithm is reported in Figure 3 and operates as follows. First, each link $l_j \in L$ is assigned a virtual demand equal to its interference-free data rate dr_j . More specifically, we set the virtual demand for link l_j to dr_j bytes. This is done at step 1 of the algorithm, where dr_j replicas of link l_j (each with virtual demand of 1 byte) are created and added to the multiset L_M of links to be scheduled. Then, links are partitioned into classes according to their SNR values. Each link class is then processed separately, ensuring that each slot accommodates links with similar SNR values. When link class L_i is considered, the network deployment

region is partitioned into square cells whose size depends on i . Then, cells are 4-colored in such a way that no two adjacent cells have the same color. Cells of the same color are then processed separately. When color h is processed, a transmission slot is formed by including, for any h -colored cell A , one link (if existing) with the receiver located in A . This process is repeated until all links in class L_i with receivers in cells with the selected color are scheduled for transmission.

The constant μ_i used to build the cell partitioning when processing links in class L_i is defined as follows [17]:

$$\mu_i = 2 \left(\frac{64(1+\eta)^{i-1} \beta_Q (\alpha - 1)}{\alpha - 2} \right)^{\frac{1}{\alpha}}.$$

We now formally prove that the schedule computed by algorithm INTTIMEFAIR is ϵi -STPF fair, with $1/7 \leq \epsilon \leq 1$.

Lemma 1: Assume that $\frac{1}{7} \leq \eta < 1$ and $\beta_Q \geq 1$. Then, for any link l_h in the multiset L_M , the virtual demand sd_h satisfied on link l_h at the end of the schedule is such that $(1 - \eta) \leq sd_h \leq (1 + \eta)$.

Proof: Due to space limitation, the proofs of all technical results are reported in the complete version of the paper, available as technical report [1]. ■

Lemma 2: Under the assumptions of Lemma 1, the schedule computed by Algorithm INTTIMEFAIR satisfies the primary interference constraint.

Proof: See [1]. ■

Theorem 1: Assume that $\frac{1}{7} \leq \eta < 1$ and $\beta_Q \geq 1$. Then, the schedule computed by Algorithm INTTIMEFAIR is ϵi -STPF fair, with $\epsilon = \eta$.

Proof: See [1]. ■

We now prove that the aggregate amount of data per unit of time transmitted by Algorithm INTTIMEFAIR is within a constant factor from optimum. We first introduce the following definitions, which are a generalization of those introduced in [17].

Definition 3: Given is a set L of links to schedule, and the corresponding link multi-set L_M . The *SNR density* for link class L_i , with $1 \leq i \leq k$, is the maximum number of receivers in a cell of class L_i , where class L_i includes links from the multi-set L_M . The SNR density for class L_i is denoted Δ_i .

Definition 4: Given a set L of links to schedule, the *normalized SNR density* for L , denoted $\Psi(L)$, is defined as

$$\Psi(L) = \max_{1 \leq i \leq k} \left\{ \frac{\Delta_i}{f((1+\eta)^{i-1} \beta_Q)} \right\}.$$

We now prove an upper bound on the length of the schedule computed by Algorithm INTTIMEFAIR.

Theorem 2: The schedule computed by Algorithm INTTIMEFAIR has $O(\Psi(L))$ length.

Proof: See [1]. ■

Corollary 1: The amount of data (bytes) per unit of time transmitted by Algorithm INTTIMEFAIR is

$$\Omega \left(\frac{(1-\eta) \sum_{l_i \in L} dr_i}{\Psi(L)} \right).$$

Proof: See [1]. ■

Theorem 3: The length of the optimal ϵi -STPF schedule is $\Omega(\Psi(L))$.

Proof: See [1]. ■

Corollary 2: The amount of data (bytes) per unit of time transmitted by the optimal algorithm is

$$\Omega \left(\frac{(1+\eta) \sum_{l_i \in L} dr_i}{\Psi(L)} \right).$$

Proof: See [1]. ■

By combining corollaries 1 and 2, we obtain the following theorem, which is the main result of this section.

Theorem 4: The amount of data (bytes) per unit of time transmitted by Algorithm INTTIMEFAIR is within a constant factor from that obtained by the optimal ϵi -STPF fair scheduling algorithm.

It is immediate to see that, if the virtual demands on all links are set to 1, the schedule computed by Algorithm INTTIMEFAIR indeed achieves interference-aware *rate-based* proportional fairness, with a constant approximation bound with respect to optimal. In the next section, we will compare the performance of both versions of the scheduling algorithm (indeed, of the corresponding greedy heuristics), to investigate whether the well known fact that time-based fairness achieves much higher throughput than rate-based fairness holds also in presence of interference.

VII. SIMULATION EXPERIMENTS

A. Algorithms and metrics

In order to evaluate the performance of interference-aware time-based fairness, we have performed extensive simulation experiments. Algorithm ϵi -STPF, while having provable performance bounds with respect to optimal, performs poorly on networks of bounded size, due to the tight spatial and SNR constraints used to schedule links. This results in very large cells, which induce a sequential schedule in networks of practical size. For this reason, to evaluate interference-aware time-based fairness we introduce a greedy heuristic aimed at producing an ϵi -STPF fair schedule in realistic networks.

The greedy heuristic, which we call GREEDYINTTIMEFAIR (GiTF), works as follows. First, each link is assigned a virtual demand equal to its interference-free data rate. Then, links are scheduled in a greedy fashion: for each fixed-length transmission slot, as many links as possible are added to the slot, subject to the condition that the aggregate throughput is increased when a new link is added, and no currently scheduled link gets a null data rate due to the introduction of the new link. After a slot is created, virtual demand on scheduled links are decreased based on the interference-inclusive data rates. This process is repeated until the virtual demand on all links is satisfied.

For the sake of comparison, we will also consider the following scheduling algorithms:

– **TDMA:** this is a time-based proportional fair TDMA algorithm, where links are scheduled sequentially, and each link is given the same channel access time. TDMA is considered as a baseline for evaluating the throughput benefits provided by STDMA approaches. Furthermore, TDMA provides the

perfectly proportionally fair allocation (in a time-based sense) of bandwidth to users, which is used as a baseline to compute the fairness index used to compare fairness of the various STDMA scheduling algorithms.

– **GreedyIntRateFair** (GiRF): this is a greedy heuristic aimed at building an interference-aware *rate-based* fair schedule. The algorithm works exactly as Algorithm GREEDYINTTIMEFAIR, the only difference being that all the links are assigned the same virtual demand (set to 100). Including GREEDYINTRATEFAIR in the comparison is important to understand whether time-based fairness provides throughput benefits with respect to rate-based fairness also in an STDMA setting.

– **GreedyTimeFair** (GTF): this is a greedy heuristic that builds a time-based fair schedule *without accounting for interference*. In other words, the schedule is greedily built with the purpose of giving each link a transmission opportunity *of the same duration*, while heuristically maximizing the aggregate throughput of the links scheduled in the same slot. Including GREEDYTIMEFAIR in the comparison is important to understand whether accounting for interference in building the schedule is actually necessary to guarantee time-based fairness.

The following metrics are used to evaluate the various scheduling algorithms:

– **aggregate throughput**, which is used to estimate the throughput benefits achieved by the various scheduling approaches;

– **fairness index**, which is used to estimate how close the bandwidth allocation achieved by a scheduling algorithm is to the TDMA allocation. More specifically, denoting by \bar{u}_i the time-based fair fraction of bandwidth allocated to user i by algorithm TDMA, and by u_i^A the fraction of bandwidth allocated to user i by algorithm A , the fairness index is computed as follows:

$$FI = \frac{1}{e^{\frac{1}{n} \cdot \sum_i \left| \ln \frac{\bar{u}_i}{u_i^A} \right|}},$$

where n is the number of users in the network. Notice that $e^{\frac{1}{n} \cdot \sum_i \left| \ln \frac{\bar{u}_i}{u_i^A} \right|}$ represents the average ratio between the fair bandwidth allocation and the actual allocation achieved by algorithm A , and takes value in $[1, \infty]$, with 1 representing perfectly fair allocation and ∞ the maximally unfair allocation. Thus, index FI takes values in $[0, 1]$, with 1 corresponding to a perfectly fair allocation, and $FI \rightarrow 0$ when the allocation becomes less and less fair.

B. Simulation setup

The simulation experiments were performed as follows. A number m of APs is distributed uniformly at random in a square area, with the constraint that APs must be at least $200m$ from each other. For each AP i , a number n_i of users is distributed uniformly at random in a circle of radius $200m$ centered at i , where n_i is chosen uniformly at random in the $[1, 10]$ interval. To compute SNR and SINR values, we use the log-distance radio propagation model with $\alpha = 3.8$, $P = 20dBm$, and $N = -80dBm$, corresponding to a transmission range of about $250m$ at $6Mbps$ without




Fig. 4. Aggregate throughput of scheduling algorithms vs. AP density.

interference. The data rate function used to set link rates as a function of the experienced SINR is function f_s as defined in Section III-A.

A single simulation experiment is performed as follows. After APs and users are deployed as described above, for each $(AP, user)$ pair we randomly select the direction of the link according to probability p_d , which is a simulation parameter set to 0.9 unless otherwise stated. More specifically, the link is set to be a downlink with probability p_d , and it is an uplink otherwise. After link direction is randomly chosen for each $(AP, user)$ pair, a transmission schedule is computed for each of the algorithms mentioned above. Then, link directions are randomly chosen again (with position of AP and users unchanged), and new schedules are computed, and so on, for 100 iterations. At the end of the experiment, the average aggregate throughput and user bandwidth allocation computed across the different schedules is returned as result of a single experiment. The results reported in the following are averaged across 100 random AP/user deployments.

C. Simulation results

We first present results obtained when the deployment area is fixed at $1km^2$, and the number of APs is increased from 5 to 25. Figure 4 reports the aggregate throughput obtained by the different scheduling algorithms, and Figure 5 reports the fairness index of the two time-based fair heuristics. TDMA throughput is independent of AP density, since transmissions are scheduled sequentially and the distribution of clients around their APs does not change with the number of APs. Conversely, the throughput obtained with the STDMA algorithms increases with the number of APs, but saturates at 25 APs. GTF achieves the highest throughput, as much as 3.5 times larger than TDMA throughput. However, as shown in Figure 5, this comes at the expense of fairness: the fairness index is as low as 0.7, meaning that the bandwidth allocation achieved by GTF is as much as 40% away from the time-based fair allocation. GiTF achieves a throughput which is about 5% lower than GTF's, but its fairness index is much higher (close to 0.97), indicating that GiTF bandwidth allocation is only about 3% away from the time-based fair allocation. It should also be noted that time-based fairness shows superior performance to rate-based fairness: GiTF's aggregate throughput is up to 36% higher than GiRF's.

Figures 6 and 7 report the aggregate throughput and fairness results when the number of APs is increased from 15 to 50, where the deployment area is changed so to keep the AP density fixed to 20 APs per square kilometer. The results




Fig. 5. Fairness index of time-based fair scheduling algorithms vs. AP density.




Fig. 6. Aggregate throughput of scheduling algorithms vs. no. of APs, with fixed density.

show an increasing trend of the aggregate throughput for the STDMA algorithms, while TDMA throughput does not increase with the number of APs due to lack of spatial reuse. The relative performance of the three STDMA algorithms is similar to the case of increasing AP density, with GTF achieving a throughput as much as 6.5 times higher than TDMA, GiTF achieving a throughput about 4–8% lower than GTF, but up to 37% higher than that achieved by GiRF. In terms of fairness, GiTF is substantially superior to GTF, with a bandwidth allocation which is about 3% away from the time-based fair allocation as compared to the 40% difference displayed by GTF. Notice that the fairness index is independent of the number of APs, while it appears to be influenced by AP density (recall Figure 5).

We have also performed a set of experiments in which 20 APs are deployed in a square kilometer, and the traffic mix parameter p_d is changed from 0.95 down to 0.75. The results, not shown due to lack of space, have shown that the traffic mix (ratio of downlink vs. uplink) has negligible effects on the scheduling algorithm performance: both the aggregate throughput and the fairness index are only marginally influenced by the value of p_d , with, e.g., the aggregate throughput varying less than 1% with different values of p_d .

VIII. CONCLUSION

We have illustrated the problems with applying existing fairness concepts to wireless networks with interference caused by spatial reuse. The analysis herein represents a first step toward an interference-aware approach to fairness, where we have tackled the problem for one-hop flows. To generalize our analysis to networks with multi-hop flows is not straightforward but is certainly a worthwhile open problem to consider.

REFERENCES

[1] D. Blough, G. Resta, P. Santi, “Interference-aware Proportional Fairness for Multi-Rate Wireless Networks”, *Tech. Report IIT-11-2013*, 2013.




Fig. 7. Fairness index of time-based fair algorithms vs. no. of APs, with fixed density.

- [2] V. Gamberoza, B. Sadeghi, E.W. Knightly, “End-to-End Performance and Fairness in Multihop Wireless Backhaul Networks”, *Proc. ACM Mobicom*, pp. 287–201, 2004.
- [3] M. Heusse, F. Rousseau, G. Berger-Sabbatel, A. Duda, “Performance anomaly of 802.11b”, *Proc. IEEE INFOCOM*, 2003.
- [4] Aerohive HiveOS Web Site, <http://www.aerohive.com/products/software-management/hiveos>.
- [5] T.-S. Kim, H. Lim, J.C. Hou, “Improving Spatial Reuse Through Tuning Transmit Power, Carrier Sense Threshold, and Data Rate in Multi-Hop Wireless Networks”, *Proc. ACM Mobicom*, pp. 366–377, 2006.
- [6] H.J. Kushner, P.A. Whiting, “Convergence of Proportional-Fair Sharing Algorithms Under General Conditions”, *IEEE Trans. on Wireless Communications*, Vol. 3, n. 4, pp. 1250–1259, 2004.
- [7] T. Lan, D. Kao, M. Chiang, A. Sabharwal, “An Axiomatic Theory of Fairness in Network Resource Allocation”, *Proc. IEEE Infocom*, pp. 1–9, 2010.
- [8] L. Li, M. Pal, Y.R. Yang, “Proportional Fairness in Multi-rate Wireless LANs”, *Proc. IEEE Infocom*, pp. 1004–1012, 2008.
- [9] W. Li, S. Wang, Y. Cui, X. Cheng, R. Xin, M. Al-Rodhaan, A. Al-Dhelaan, “AP Association for Proportional Fairness in Multi-Rate WLANs”, *IEEE Trans. on Networking*, (to appear).
- [10] R. Maheshwari, S. Jain, S.R. Das, “A Measurement Study of Interference Modeling and Scheduling in Low-Power Wireless Networks”, *Proc. 6th ACM SenSys*, pp. 141–154, 2008.
- [11] R. Maheshwari, J. Cao, S.R. Das, “Physical Interference Modeling for Transmission Scheduling on Commodity WiFi Hardware”, *Proc. IEEE Infocom*, pp. 2661–2665, 2009.
- [12] Meru Networks Web Site, <http://www.merunetworks.com/>.
- [13] V. Mhatre, K. Papagiannaki, F. Baccelli, “Interference Mitigation through Power Control in High Density 802.11 WLANs,” *Proceedings of the 26th IEEE International Conference on Computer Communications*, pp. 535–543, 2007.
- [14] B. Radunovic, J. Le-Boudec, “Rate performance objectives of multi-hop wireless networks”, *Proc. IEEE Infocom*, 2004.
- [15] C. Raman, R.D. Yates, N.B. Mandayam, “Scheduling Variable Rate Links via a Spectrum Server”, *Proc. IEEE DySPAN*, pp. 110–118, 2005.
- [16] T.S. Rappaport, *Wireless Communications*, Prentice Hall, Upper Saddle River, NJ, 2002.
- [17] P. Santi, R. Maheshwari, G. Resta, S. Das, D. Blough, “Wireless Link Scheduling under a Graded SINR Interference Model”, *Proc. ACM Workshop on Foundations of Wireless Ad Hoc and Sensor Networking and Computing (FOWANC)*, pp. 3–12, 2009.
- [18] C. Scheideler, A. Richa, P. Santi, “An $O(\log n)$ Dominating Set Protocol for Wireless Ad Hoc Networks under the Physical Interference Model”, *Proc. ACM MobiHoc*, pp. 91–100, 2008.
- [19] P. Soldati, B. Johansson, M. Johansson, “Proportionally Fair Allocation of End-to-End Bandwidth in STDMA Wireless Networks”, *Proc. ACM Mobicom*, pp. 286–297, 2006.
- [20] G. Tan, J. Gutttag, “Time-based Fairness Improves Performance in Multi-Rate WLANs”, *Proc. USENIX Conf.*, pp. 23–36, 2004.
- [21] R.D. Yates, C. Raman, N.B. Mandayam, “Fair and Efficient Scheduling of Variable Rate Links via a Spectrum Server”, *Proc. IEEE ICC*, pp. 5246–5251, 2006.