What is optimal scheduling in wireless networks?*

(Invited Paper)

Anna Pantelidou and Anthony Ephremides
Institute for Systems Research
Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742
{apantel, etony}@umd.edu

ABSTRACT

We consider a wireless network consisting of multiple transmitters with multicast traffic destined for a set of receivers. We are interested in the problem of joint scheduling and rate control under two performance objectives; the objective of maximizing the total sum throughput of the network and of being proportionally fair with respect to the received rate at each receiver. We first consider static wireless networks, and then extend our analysis for the more general and more realistic case of time-varying networks. We finally verify our analytical results through a set of simulations.

1. INTRODUCTION

Medium access control methods have a great impact on the performance of wireless systems where the limited resources need to be shared in an efficient and effective manner among multiple contending users. The presence of medium access control becomes essential to mitigate the interference when the number of users in the network increases. Clearly, the decisions regarding how to restrict access to the channel are tightly coupled with the characteristics of the physical layer, in particular with the corresponding powers and rates at which the contending nodes operate.

In this paper, we consider a single-hop network of multiple transmitters, each multicasting traffic destined for a set of receivers. Each transmitter is associated with a multicast session and the receivers of various sessions are allowed to overlap. We are interested in the problem of jointly scheduling the transmitters and controlling their rates under two criteria, the criteria of total throughput and proportional fairness. We incorporate the physical layer into the scheduling decisions through the *physical interference* model which asserts that a transmission is successful if the ratio of the received signal power over the total interference and thermal noise powers at a receiver exceeds a certain threshold (SINR criterion).

The problem of joint scheduling and rate control in the context of wireless systems has been studied extensively in the literature under a variety of different settings ([1], [2], [3], [4], [5], [6], [7], etc.).

*This work was supported under the MURI grant W911NF-08-1-0238.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
WICON'08 November 17-19, 2008, Maui, Hawaii, USA
Copyright 2008 ICST 978-963-9799-36-3

In [1] the problem of joint scheduling and rate control in static wireless networks under unicast traffic is considered. The optimal solution for the problem of maximizing the total (sum) throughput with and without a minimum rate requirement for every transmitter is characterized. It is further shown that when the transmission powers are large a pure Time Division Multiple Access (TDMA) scheme, where a single node transmits at any given time, is optimal with respect to maximizing the total throughput in the presence of minimum rate constraints. In addition, the problem of obtaining a max-min fair and proportionally fair rate allocation is formulated in terms of a linear and a non-linear program respectively. However, in [1] the optimal solution is not characterized in neither formulation.

In a different work [2] we consider the problem of joint scheduling and rate control for unicast traffic in static wireless networks under the objective of proportional fairness. We assume a restricted problem that allows to employ only a subset of the total possible rate control and scheduling actions. Specifically, we restrict our attention to rates that can be achieved by scheduling the transmitters one at a time, in a pure TDMA fashion, and also rates that can be achieved by letting all the transmitters to operate concurrently. We further characterize the optimal proportionally fair probability distribution over this restricted set of rate control and scheduling decisions which specifies the probability with which each decision is chosen. Thus we show that contingent upon the channel conditions and the amount of interference that each transmitter causes to the others it may be preferable to allow all the transmitters to operate simultaneously for a certain fraction of the time. This is due to the fact that although the individual rate of any transmitter when it is activated in a pure TDMA fashion is higher than the corresponding rate under concurrent operation, the latter may still be preferable when the transmitters do not interfere much with each other and hence can be assumed to operate almost independently. The work in [2] aims to exactly characterize this trade-off in the scheduling, namely to characterize when it is better, with respect to the objective of proportional fairness to allow more concurrent activations transmitting at lower rates or to schedule fewer activations at higher rates.

However, although [1] and [2] consider static wireless networks in practice the wireless links are hardly ever static due to fading, mobility and other effects. Towards this end, a large body of research studied the problem of scheduling under time-varying wireless networks. In particular, a model that has been extensively studied is the downlink channel of a base station, transmitting unicast data traffic to a set of mobile terminals under the assumption that the base station can serve a single terminal at any given time in a TDMA manner. One such scheduling algorithm is the proportional fair sharing scheduler (PFS) introduced by Qualcomm for its High

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991

Data Rate (HDR) system. The PFS selects a single terminal for transmission at any given time; the one that maximizes the ratio of a user's instantaneous rate to the average rate it has received so far. Therefore, those terminals that received comparably lower average data rates until the current decision instant are more likely to be selected in the optimal solution. The PFS is shown to be optimal with respect to the objective of proportional fairness ([3], [4], [5]).

In a subsequent work [6] a time-varying wireless network is considered in which multiple transmitters are allowed to operate simultaneously in order to send their unicast traffic to their respective destinations. The channel process is assumed to be stationary and ergodic, but otherwise arbitrary. Under this generalized framework a rate allocation policy is introduced that is optimal with respect to maximizing a strictly concave, increasing utility function. In spite of the generality of the formulation, the model in [6] does not capture the objective of proportional fairness. In a different work [7] a broader class of utility functions is accounted for and the problem of optimal rate allocation for a switch serving a set of queues is considered. However, in [7] the state process of the switch, which defines its service rate, is constrained to alternate among states according to a stationary and ergodic, finite state Markov Chain.

Although these works consider the problem of rate control for utility maximization under unicast traffic, a large amount of traffic in networks is comprised of multicast data. In [8] a base station with multicast traffic destined towards various groups of receivers is considered. It is assumed that at any given time only a single multicast group can be selected for transmission and that all the terminals in the group must receive at the same rate. Hence, the decision that needs to be made at any given time is which *unique* multicast group to serve and at which *rate* under two objectives; first the objective is to be proportionally fair with respect to the total rate of each multicast group, and in the sequel, with respect to the overall rate of each terminal when it is a member of various multicast groups.

In this paper, we generalize prior work in that we consider multicast traffic under both static and time-varying environments where concurrent activations of the transmitters are allowed. Specifically, in the first part of this paper we consider static wireless networks under multicast traffic. We first obtain the optimal rate control and scheduling policy to maximize the total throughput of the network. Next, similarly to [2], we restrict our attention to two simple, yet non-trivial, schemes, namely concurrent activation of all the transmitters as well as a TDMA scheme that activates a single multicast transmitter at any given time. We obtain a proportionally fair probability distribution over this reduced set of rate control and scheduling decisions. Our results generalize [2] in two aspects; (i) we consider multicast traffic, and (ii) we employ a weaker set of assumptions. Moreover, since unicast is a special case of multicast traffic our model can be used to obtain the results in [2]. Unlike [1] our objective is to not only formulate the problem of proportional fairness as a convex optimization problem, but also to explicitly characterize the optimal solution and thus be able to deduce under which channel conditions, reflecting to corresponding rates, one scheme is better than the other.

However, wireless systems are time varying in general. Hence we proceed by considering time-varying networks. We obtain an optimal policy that maximizes the total throughput of the network. We further introduce a systematic algorithm, that is "channel-aware" and "opportunistic" and achieves proportional fairness by taking greedy decisions at any given time as in [9]. We extend the results of [3], [4], [5], [6] and [7] by considering the problem of proportionally fair rate allocation for multicast traffic. Further we extend [3], [4] and [5] in that our model allows for concurrent activations. Our results are also different from [6] and [7] by assuming a relaxed set of assumptions on the channel process. Since a multicast transmission is a general case of unicast our results are valid for unicast transmissions as well. We also extend the results of [8] in that we allow multiple multicast transmissions to be scheduled simultaneously.

The rest of this paper is organized as follows. In Section 2 we present the network model we consider. In Section 3 we discuss the problem of total throughput maximization and proportional fairness in static networks. Next, in Section 4 we extend our discussions to time-varying networks under a broader class of utility functions, which captures the objectives of total throughput and proportional fairness. In Section 5 we verify our analytical results through a set of simulations. Finally, we conclude the paper in Section 6.

2. MODEL

We consider a single-hop wireless network of T multicast transmitters and D receivers. The time is slotted. We denote by $\mathcal{T}$ and $\mathcal{D}$ the sets of transmitters and receivers in the network respectively. Each transmitter $k \in \mathcal{T}$ wishes to multicast at a common rate (single rate multicast) to a set of receivers $\mathcal{D}(k) \subseteq \mathcal{D}$. We call the pair $(k, \mathcal{D}(k))$ a *multicast session*. Our model is general enough to account for the special cases of unicast and broadcast traffic as well, i.e., when the cardinality $|\mathcal{D}(k)|$ of the set $\mathcal{D}(k)$ is equal to 1 and D respectively. We assume that a receiver $d \in \mathcal{D}$ can be a member of more than one multicast session, i.e., for any two multicast transmitters $j, k \in \mathcal{T}$, it is possible that $\mathcal{D}(k) \cap \mathcal{D}(j) \neq \emptyset$. For example, in Figure 1 receiver 1 is an intended receiver for both transmitters 1 and 2. In this work we do not consider bursty traffic but assume, instead, that each transmitter is saturated and always has enough data to send. Moreover, at each time slot n each transmitter k transmits at a power level $P_n(k)$. We assume that $P_n(k)$ takes two possible values, namely $P_k^{\max}$ when transmitter k operates at its maximum power and 0 when it remains silent. We denote by $\mathbf{P}_n$ the K-dimensional vector of transmission powers at time slot n, i.e., $\mathbf{P}_n = (P_n(k),\ k \in \mathcal{T})$. We also denote by $N(d)$ the noise power level at receiver $d \in \mathcal{D}$. The network model we consider is depicted in Figure 1. Note that although we restrict our attention to a single-hop network, our model can be used to address and

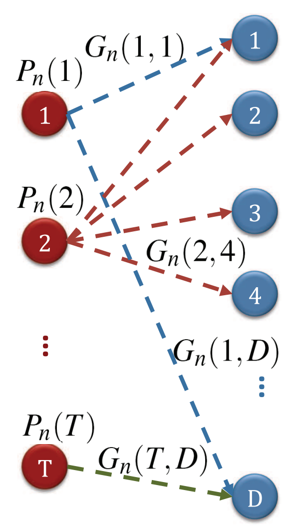

Figure 1: A network of T multicast transmitters and D receivers.

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991

illuminate the scheduling and rate control problem in full-fledged multi-hop networks under fixed routing.

Consider a channel process $\{\mathbf{S}_n\}_{n=0}^{\infty}$ that specifies the changes in the channel conditions between every transmitter and receiver in the network. The time-variability of the channel process is due to node mobility, channel fading, shadowing, simple path loss, etc. The process $\{\mathbf{S}_n\}_{n=0}^{\infty}$ is stationary and although it can vary from one slot to another, it is assumed to be fixed during the duration of a time slot. For each time slot n, the channel state $\mathbf{S}_n = \{\mathbf{G}_n(i,j), \ \ i \in \mathcal{T}, \ j \in \mathcal{D}\}$ specifies the overall channel effect $\mathbf{G}_n(i,j)$ between each transmitter i and receiver j. The process $\{\mathbf{S}_n\}_{n=0}^{\infty}$ takes values from a continuous set $\mathcal{S}$ which has probability density function $f_{\mathcal{S}}(\cdot)$. In the first part of this paper we consider static channel processes, i.e., $\mathbf{S}_n = \mathbf{s} \in \mathcal{S}$, for all $n = 1, 2, \ldots$. In the second part we generalize our discussions to address the more general, and more realistic, case of time-varying channel processes.

Regardless of the definition of the channel process, we define the outcome of a transmission based on the *physical interference* model. Specifically, we say that a transmission from a transmitter to its intended receiver is successful if the Signal to Interference plus Noise Ratio (SINR) at the receiver exceeds a specified threshold. Although approximate in general, this model accounts for the fact that in a wireless network *all* the concurrently transmitting nodes may interfere and cause a transmission to fail, depending on the channel conditions and the respective powers of the concurrently transmitting nodes. The SINR threshold depends on various communication-related parameters, such as the transmission rate, the modulation type, the coding technique, the target bit error probability, etc. In this work, we only examine the dependence of the SINR threshold on the rate of the transmission and assume that the rest of the parameters remain constant.

When at some time slot n a transmitter k *multicasts* at a *common* rate r to all its intended receivers in the set $\mathcal{D}(k)$, then the SINR at each receiver $d \in \mathcal{D}(k)$ must exceed the corresponding threshold, i.e.,

$$\frac{P_n(k)G_n(k,d)}{N(d) + \sum_{j \in \mathcal{T}, j \neq k} P_n(j)G_n(j,d)} \geq \gamma_{n,d}(r), \ \forall d \in \mathcal{D}(k). \quad (1)$$

We consider that each receiver $d \in \mathcal{D}$ has multi-packet reception (MPR) capabilities, in other words at any given time it may receive successfully from multiple transmitters as long as the corresponding SINR from each one of them exceeds the required threshold. Thus, it is possible for two, or more, multicast transmitters with overlapping receiving nodes to concurrently transmit successfully.

It is well-known that the transmission rate is an increasing function of the SINR threshold (See e.g., [10].). Hence, it follows that when the transmission rate is lowered, the corresponding value of the threshold decreases, and thus more transmissions can jointly satisfy the condition of (1). On the other hand, an increase in the transmission rate increases the SINR threshold, which restricts the number of transmitters that can concurrently access the channel successfully. By varying the transmission rates, and hence the thresholds, we can obtain two extreme cases: (i) a pure TDMA scheme, where only a single transmitter can successfully transmit at any given time, and (ii) a scheme where all transmitters can successfully multicast concurrently. In general for a network of T multicast transmitters there exist $2^T - 1$ possible ways to schedule them[1]. We call the possible ways at which the transmitters can be activated as *actions*. Let us denote by $\mathcal{A}$ the set of all possible actions with $|\mathcal{A}| = 2^T - 1$. It is not immediate which subset of the transmitters must operate at any given time and at which rate, i.e., a scheduling and rate control action needs to be determined. The optimal action to be selected depends on several factors such as the adopted performance objective, the channel and network conditions, etc.

We proceed to define the set of feasible transmission rates that can be achieved through all possible rate control and scheduling actions. Let $\mathcal{R}_n(\mathbf{s})$ be the feasible rate region at time slot n when the state is $\mathbf{S}_n = \mathbf{s} \in \mathcal{S}$. Based on the above, $\mathcal{R}_n(\mathbf{s})$ is defined as

$$\mathcal{R}_n(\mathbf{s}) = \Bigg\{ \mathbf{r} = (r_1, \ldots, r_T) : \forall k \in \mathcal{T}, \ \forall \, d \in \mathcal{D}(k), \\ \frac{P_n(k)G_n(k,d)}{N(d) + \sum_{j \in \mathcal{T}, j \neq k} P_n(j)G_n(j,d)} \geq \gamma_{n,d}(r_k) \Bigg| P_n(k) \in \{0, P_k^{\max}\} \Bigg\}. \quad (2)$$

It is easy to observe that for any given power vector $\mathbf{P}_n$, there exist multiple rate vectors in $\mathcal{R}_n(\mathbf{s})$ that satisfy the SINR criterion. Without loss of generality, we are interested only in those vectors in $\mathcal{R}_n(\mathbf{s})$ that correspond to the *maximum* transmission rates for which the SINR criterion is jointly satisfied at every receiver. This is reasonable since as all the transmitters are saturated they will always transmit at the highest rate possible whenever activated. We denote this new region by $\tilde{\mathcal{R}}_n(\mathbf{s})$. Under this assumption, the set $\tilde{\mathcal{R}}_n(\mathbf{s})$ is a discrete set that contains $2^{|\mathcal{A}|}$ rate vectors, which can be obtained by activating all possible subsets of the T transmitters. Each such rate vector $\mathbf{r} \in \tilde{\mathcal{R}}_n(\mathbf{s})$ corresponds to an achievable rate under a scheduling action $j \in \mathcal{A}$.

In the rest of the paper we obtain medium access control schemes that take the characteristics of the physical layer into account. Our objective is to maximize the total throughput of the network and address fairness by employing the criterion of proportional fairness under multicast traffic for both static and time-varying networks. We define by throughput the overall traffic that reaches all the receivers of a multicast session. Thus for any two multicast transmitters that transmit at equal rates, the transmitter that has a higher number of receivers is assumed to contribute more with respect to throughput.

3. STATIC NETWORKS

In this section we consider networks that are time-invariant. We assume that the channel effect between every transmitter and every receiver is due to pure path loss which is fixed in time. Hence, at every time slot n the channel state $\mathbf{S}_n$ is given by $\mathbf{S}_n = \{\mathbf{G}(i,j), \ \ i \in \mathcal{T}, \ j \in \mathcal{D}\}$, where $0 < G(i,j) \leq 1$ is the path loss between transmitter i and receiver j. As we mentioned previously, every rate vector in the rate region $\tilde{\mathcal{R}}_n(\mathbf{s})$ corresponds to an action $j \in \mathcal{A}$. Due to the time-invariability of the network with a little abuse of notation we refer to the feasible rate region as $\tilde{\mathcal{R}}$. We further index each rate vector in $\tilde{\mathcal{R}}$ by the action $j \in \mathcal{A}$ that achieves it. Thus, we denote by r_k^j the rate at which transmitter $k \in \mathcal{T}$ transmits to all of the receivers $\mathcal{D}(k)$ in its multicast session, under action $j \in \mathcal{A}$.

Since we consider single-rate multicast, the rate of transmitter k is equal to the rate of every receiver d in its multicast group ($d \in \mathcal{D}(k)$). Thus, we can characterize the rate of each receiver $d \in \mathcal{D}(k)$ through the transmission rate of its corresponding transmitter.

In Section 3.1 we obtain a scheduling and rate control policy that maximizes the total (sum) throughput of the network. Since maximizing the total throughput of the network can be unfair to certain transmitters that face a poor channel, in Section 3.2 we consider the commonly used criterion of proportional fairness. We further ob-

[1]We subtract the scheduling action that corresponds to all transmitters being silent since given the fact that their queues are always saturated at least one transmitter will be activated at any given time.

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991

tain a rate control and scheduling policy that is proportionally fair with respect to the received rate.

Let $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_{|\mathcal{A}|})$ denote a probability distribution over the set of all possible rate control and scheduling actions in $\mathcal{A}$[2]. We then define the *effective rate* of transmitter $k \in \mathcal{T}$ to be

$$r_k = \sum_{j \in \mathcal{A}} r_k^j \pi_j.$$

3.1 Total throughput

In this work we measure throughput in terms of the total received rate of all receivers in a multicast session. Under the assumption that the network is static, the maximization problem can be formulated as follows:

$$\max_{\boldsymbol{\pi}} \sum_{k \in \mathcal{T}} |\mathcal{D}(k)| \sum_{j \in \mathcal{A}} r_k^j \pi_j \tag{3}$$

s.t.

$$\pi_j \geq 0, \qquad j \in \mathcal{A}, \tag{4}$$

$$\sum_{j \in \mathcal{A}} \pi_j = 1. \tag{5}$$

Note that the problem described by (3)-(5) is a linear program. Its solution is given by the following proposition.

PROPOSITION 1. *The optimal policy that solves* (3)-(5) *with the objective to maximize the total throughput of the network assigns a non-zero probability to those actions $j \in \mathcal{A}$ that maximize the total sum throughput T_j, where*

$$T_j := \sum_{k \in \mathcal{T}} |\mathcal{D}(k)| r_k^j.$$

If there exists a unique such action the optimal policy will choose it with probability 1. *Otherwise, in the presence of more than one maximizers, the optimal policy will arbitrarily time-share among them.*

It is clear that optimizing the total throughput of the network leads to an efficient utilization of the network resources since we maximize the total rate that the network can support. Nevertheless, similarly to the unicast case [1], it can lead to serious unfairness among the transmitters with poor channel conditions by prohibiting them from accessing the channel. In the next section we consider the utility of proportional fairness [11] which has been widely used as a performance metric in wireless networks since it provides a good compromise between efficiency and fairness [12].

3.2 Proportional fairness

In this section we focus on the objective of proportional fairness. Specifically a feasible rate vector $\mathbf{r} = (r_k, k \in \mathcal{T})$ is said to be *proportionally fair* if for any other feasible rate vector $\mathbf{r}' = (r'_k, k \in \mathcal{T})$ it is true that

$$\sum_{k \in \mathcal{T}} \frac{r'_k - r_k}{r_k} \leq 0.$$

It was shown in [11] that the objective of proportional fairness is equivalent to maximizing the sum of the logarithms of user rates over all the feasible rate vectors.

Recall that r_k^j is the transmission rate of transmitter k under action j. Hence, obtaining the proportionally fair rate can be given as an optimization problem with respect to the probability distribution over which each action is chosen. This is given next:

$$\max_{\boldsymbol{\pi}} \sum_{k \in \mathcal{T}} |\mathcal{D}(k)| \log \left(\sum_{j \in \mathcal{A}} \pi_j r_k^j \right)$$

s.t.

$$\pi_j \geq 0, \quad \forall j \in \mathcal{A},$$

$$\sum_{j \in \mathcal{A}} \pi_j = 1.$$

Let the vector $\boldsymbol{\pi}^\star = (\boldsymbol{\pi}_1^\star, \ldots, \boldsymbol{\pi}_{|\mathcal{A}|}^\star)$ be the vector of probabilities with which each action is chosen in order to achieve proportional fairness. The fact that the number of possible actions is exponentially increasing in the number of multicast transmitters renders an analytical computation of each action probability infeasible. Instead, in what follows we are going to focus our attention on a restricted, yet non-trivial, problem in which we consider actions given by two extreme threshold selections. In particular we are going to decrease the SINR threshold values to the maximum such values that allow all the multicast transmitters to operate concurrently. We will call this approach as “Lowering the Rates” or “LR” in short. It will correspond to “Action 0”. Further we consider T possible actions (“Action 1” to “Action T”) that are obtained by increasing the SINR thresholds so that only a single transmitter can access the channel at any given time. We will call this approach as “Scheduling” or “SCH”. Hence, each action $k = 1, \ldots, T$ corresponds to transmitter k being active in an interference free manner. Thus, the two schemes “LR” and “SCH” yield a total of $T + 1$ actions, as shown in Figure 2. Under the “LR” scheme although all transmitters operate concurrently their individual rates may be very low due to the effects of interference. On the other hand, under the “SCH” scheme although a single transmitter is active at any given time its rate is higher than its corresponding rate under the scheme “LR”, at the cost of accessing the channel for a smaller fraction of the time due to the time-sharing.

Next, we find the optimal proportionally fair probability distribution over the aforementioned restricted set of actions by solving the following problem:

$$\max_{\boldsymbol{\pi}} \sum_{k \in \mathcal{T}} |\mathcal{D}(k)| \log(\pi_0 r_k^0 + \pi_k r_k^k) \tag{6}$$

s.t.

$$\pi_j \geq 0, \quad \forall j \in \{0, 1, \ldots, T\}, \tag{7}$$

$$\sum_{j=0}^{T} \pi_j = 1. \tag{8}$$

Before we characterize the optimal policy solving (6)-(8), we provide some useful definitions. Let $\mathcal{J}$ be a subset of the set $\mathcal{T}$, such that for every $j \in \mathcal{J}$ it is true that $\pi_j > 0$. Also let the complement $\mathcal{J}^c$ of the set $\mathcal{J}$, i.e., $\mathcal{J}^c = \mathcal{T} \setminus \mathcal{J}$, be the set such that for every $i \in \mathcal{J}^c$ it follows that $\pi_i = 0$. Given these definitions, the optimal policy is defined as follows.

PROPOSITION 2. *The optimal proportionally fair policy that solves* (6)-(8) *has the following properties:*

1. If

$$\sum_{k \in \mathcal{T}} \frac{r_k^0}{r_k^k} \leq 1,$$

[2]We assume that this probability distribution exists by requiring ergodicity on the selection of the different actions.

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991

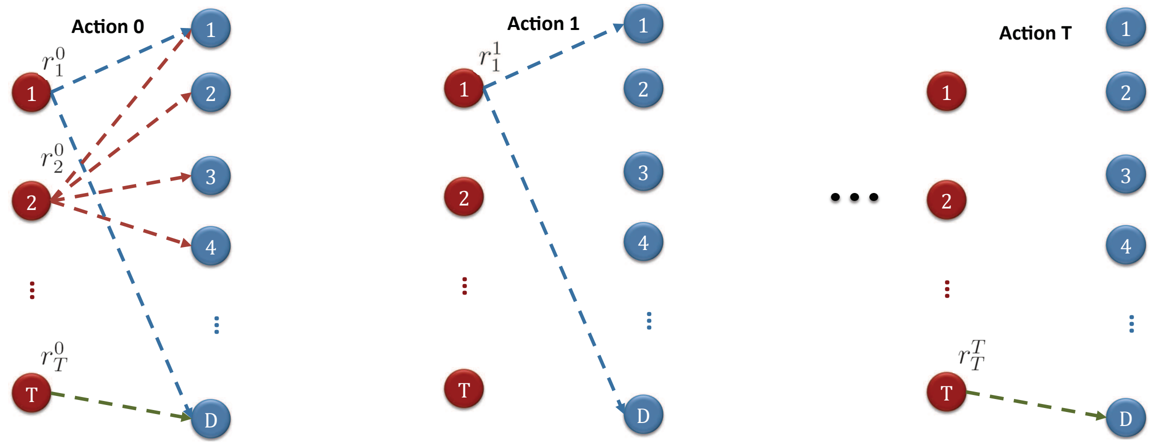

Figure 2: The $T+1$ possible configurations obtained by either scheduling T transmitters one by one or by allowing all the transmitters to transmit simultaneously. The rate of transmitter k under configuration j is denoted by r_k^j.

then each multicast transmitter $k \in \mathcal{T}$ is scheduled to transmit individually with probability

$$\pi_k^\star = \frac{|\mathcal{D}(k)|}{\sum_{j\in\mathcal{T}} |\mathcal{D}(j)|}, \quad \forall k \in \mathcal{T},$$

and the probability of concurrent operation satisfies $\pi_0^\star = 0$.

2. *If*

$$\sum_{k\in\mathcal{T}} \frac{r_k^0}{r_k^k} > 1,$$

then the optimal policy is a threshold policy with threshold function $R(\mathcal{J})$, where

$$R(\mathcal{J}) = \frac{1 - \sum_{j\in\mathcal{J}} r_j^0/r_j^j}{\sum_{m\in\mathcal{J}^c} |D(m)|}. \tag{9}$$

Specifically:

(a) *A multicast transmitter $j \in \mathcal{T}$ is scheduled to transmit individually with probability $\pi_j^\star > 0$ (i.e., $j \in \mathcal{J}$) given by*

$$\pi_j^\star = \frac{|\mathcal{D}(j)| - \sum_{i\in\mathcal{J}^c} |\mathcal{D}(i)| \frac{r_j^0/r_j^j}{1-\sum_{j\in\mathcal{J}} r_j^0/r_j^j}}{\sum_{k\in\mathcal{T}} |\mathcal{D}(k)|}, \tag{10}$$

if and only if

$$\frac{r_j^0}{|\mathcal{D}(j)| r_j^j} < R(\mathcal{J}). \tag{11}$$

(b) *All transmitters operate concurrently with probability $\pi_0^\star$ given by*

$$\pi_0^\star = \frac{\sum_{m\,\in\mathcal{J}^c} |\mathcal{D}(m)|}{\left(\sum_{k\in\mathcal{T}} |\mathcal{D}(k)|\right)\left(1 - \sum_{j\in\mathcal{J}} r_j^0/r_j^j\right)}. \tag{12}$$

Proposition 2 characterizes the optimal solution based on the threshold function $R(\mathcal{J})$ which itself is a function of the set $\mathcal{J}$. Hence, in order to completely characterize the optimal policy we need to specify the set $\mathcal{J}$. Note that since the optimal policy is of threshold type, the cardinality $|\mathcal{J}|$ of set $\mathcal{J}$ is sufficient to specify set $\mathcal{J}$. Let $R(j)$ denote $\{R(\mathcal{J}) : |\mathcal{J}| = j\}$. Further, let us reorder the multicast sessions with respect to their corresponding values of the ratios $r_j^0/|\mathcal{D}(j)|r_j^j$, $j \in \mathcal{T}$ values in increasing order, i.e.,

$$\frac{\tilde{r}_1^0}{|\tilde{\mathcal{D}}(1)|\tilde{r}_1^1} \le \frac{\tilde{r}_2^0}{|\tilde{\mathcal{D}}(2)|\tilde{r}_2^2} \le \ldots \le \frac{\tilde{r}_T^0}{|\tilde{\mathcal{D}}(T)|\tilde{r}_T^T},$$

where $\tilde{r}_j^0$, $\tilde{r}_j^j$, and $\tilde{\mathcal{D}}(j)$ denote the rates r_j^0, r_j^j, and $\mathcal{D}(j)$ respectively under the new ordering. We will make use of the following property of the threshold function $R(j)$ to obtain the cardinality of set $\mathcal{J}$.

Lemma 1. *The threshold function $R(j)$ defined in Proposition 2 satisfies the following:*

$$R(j-1) \le R(j), \quad \textit{if and only if} \;\; j \in \mathcal{J}.$$

Hence, $R(j)$ is increasing for all $j \in \mathcal{J}$ and decreasing for all $j \in \mathcal{J}^c$. Using this fact, the cardinality of $\mathcal{J}$ can be found as below.

Proposition 3. *The cardinality of set $\mathcal{J}$ under the optimal policy specified in Proposition 2 is given by the following:*

$$|\mathcal{J}| = \arg\max_{\ell\in\{0,1,\ldots,T\}} \frac{1 - \sum_{j=1}^{\ell} \tilde{r}_j^0/\tilde{r}_j^j}{\sum_{m=\ell+1}^{T} |\tilde{\mathcal{D}}(m)|}. \tag{13}$$

From Proposition 2, we observe that the quantity $\sum_{k\in\mathcal{T}} \frac{r_k^0}{r_k^k}$ is, in a sense, a criterion to determine the level of interference in the network. For instance, the condition $\sum_{k\in\mathcal{T}} \frac{r_k^0}{r_k^k} < 1$ can be seen as a criterion for high levels of interference in the network since it can be translated to a situation in which the rates under concurrent operation are much lower than the corresponding rates under individual operation. Hence, when $\sum_{k\in\mathcal{T}} \frac{r_k^0}{r_k^k} \le 1$, concurrent transmissions are avoided ($\pi_0 = 0$) and the optimal scheme is to activate the transmitters one at a time in a TDMA fashion. Otherwise, the case where it is true that $\sum_{k\in\mathcal{T}} \frac{r_k^0}{r_k^k} > 1$ can be seen as a scenario of low interference, as the individual rates under concurrent operation are comparable to the rates achieved by a TDMA scheme. Hence, the optimal policy assigns a positive probability to the "LR" scheme

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991

and thus allows concurrent transmissions. Furthermore, the transmitters that are further selected to be activated individually (i.e., those that are chosen to be in the set $\mathcal{J}$) are the most disadvantaged multicast transmitters, i.e., whose ratios $r_j^0/|\mathcal{D}(j)|r_j^j$ are the lowest due to the effects of interference. This is ensured by ordering the transmitters with respect the ratios $r_j^0/|\mathcal{D}(j)|r_j^j$ and assigning the transmitters with the $|\mathcal{J}|$ lowest values to set $\mathcal{J}$ through (13).

Note that the optimal proportionally fair probability distribution for the case of unicast traffic $(\pi_0^{\mathrm{u}\star}, \ldots, \pi_T^{\mathrm{u}\star})$ under the restricted set of actions follows directly from our formulation by setting the cardinality of the set $\mathcal{D}(k)$ for every transmitter equal to one, i.e., $|\mathcal{D}(k)| = 1$ for every $k \in \mathcal{T}$. Then the solution of the unicast case is given by

COROLLARY 1. *The optimal proportionally fair policy for unicast traffic has the following properties:*

1. *If*
$$\sum_{k\in\mathcal{T}} \frac{r_k^0}{r_k^k} \leq 1,$$
then each transmitter $k \in \mathcal{T}$ is scheduled to transmit individually with probability
$$\pi_k^{\mathrm{u}\star} = \frac{1}{T}, \quad \forall k \in \mathcal{T},$$
and the probability of concurrent operation is zero , i.e., $\pi_0^{\mathrm{u}\star} = 0$.

2. *If*
$$\sum_{k\in\mathcal{T}} \frac{r_k^0}{r_k^k} > 1,$$
then the optimal policy is a threshold policy with threshold function $R(\mathcal{J})$, where
$$R(\mathcal{J}) = \frac{1 - \sum_{j\in\mathcal{J}} r_j^0/r_j^j}{T - |\mathcal{J}|}. \tag{14}$$
Specifically,
   (a) *A transmitter $j \in \mathcal{T}$ is scheduled to transmit individually with probability $\pi_j^{\mathrm{u}\star} > 0$ (i.e., $j \in \mathcal{J}$) given by*
$$\pi_j^{\mathrm{u}\star} = \frac{1}{T} - \sum_{i\in\mathcal{J}^c} \frac{r_j^0/r_j^j}{T\left(1 - \sum_{j\in\mathcal{J}} r_j^0/r_j^j\right)}, \tag{15}$$
   if and only if
$$\frac{r_j^0}{r_j^j} < R(\mathcal{J}). \tag{16}$$
   (b) *All transmitters operate concurrently with probability $\pi_0^{\mathrm{u}\star}$ given by*
$$\pi_0^{\mathrm{u}\star} = \frac{T - |\mathcal{J}|}{T\left(1 - \sum_{j\in\mathcal{J}} r_j^0/r_j^j\right)}. \tag{17}$$

COROLLARY 2. *The cardinality of set $\mathcal{J}$ under the optimal policy specified in Corollary 1 is given by the following:*
$$|\mathcal{J}| = \underset{\ell\in\{0,1,\ldots,T\}}{\arg\max} \frac{1 - \sum_{j=1}^{\ell} \tilde{r}_j^0/\tilde{r}_j^j}{T - \ell}. \tag{18}$$

Note also that Corollaries 1 and 2 extend [2] where we had assumed that for every unicast transmitter $j \in \mathcal{T}$ the rates under individual operation r_j^j are all equal to each other.

4. TIME-VARYING NETWORKS

The wireless channel is time-varying in nature and fluctuates with time due to several reasons such as the effects of user mobility, fading, shadowing, etc. In this section we generalize our previous discussions to capture this time-variability. For the purpose of this paper we consider the utility function of α-fairness [13] defined as

$$U^\alpha(\theta) = \begin{cases} \log(\theta) & \text{if } \alpha = 1 \\ (1-\alpha)^{-1}\theta^{1-\alpha} & \text{otherwise.} \end{cases} \tag{19}$$

Although our results, presented next, hold for arbitrary utility functions that are strictly concave and increasing in the received rate of a user our choice of $\alpha-$fairness stems from the fact that it can capture some well known fairness criteria. In particular, for $\alpha = 1$ it corresponds to the criterion of proportional fairness, for $\alpha \to \infty$ to max-min fairness, and for $\alpha = 0$ it yields the total (sum) throughput objective. Thus, employing the utility of α-fairness suffices to extend the criteria we considered for static networks, namely total throughput and proportional fairness, to time-varying ones.

Let $R_{n,k}^{\boldsymbol{\pi}}$ denote the multicast rate assignment of the kth transmitter under some policy $\boldsymbol{\pi}$ at time slot n. Then the time average of the rate of transmitter k is given by

$$\theta_{n,k}^{\boldsymbol{\pi}} = \frac{1}{n}\sum_{\nu=1}^{n} R_{\nu,k}^{\boldsymbol{\pi}}. \tag{20}$$

We dub this time average as the “effective rate” of the kth transmitter.

The effective rate described in (20) can also be written in a recursive form as

$$\theta_{n+1,k}^{\boldsymbol{\pi}} = \theta_{n,k}^{\boldsymbol{\pi}} + \epsilon_n Y_{n,k}^{\boldsymbol{\pi}}, \tag{21}$$

where

$$\epsilon_n = \frac{1}{n+1}, \tag{22}$$

$$Y_{n,k}^{\boldsymbol{\pi}} = R_{n+1,k}^{\boldsymbol{\pi}} - \theta_{n,k}^{\boldsymbol{\pi}}. \tag{23}$$

We denote by $\mathbf{R}_n^{\boldsymbol{\pi}} = (R_{n,k}^{\boldsymbol{\pi}}, k \in \mathcal{T})$ and $\boldsymbol{\theta}_n^{\boldsymbol{\pi}} = (\theta_{n,k}^{\boldsymbol{\pi}}, k \in \mathcal{T})$ the T-dimensional vectors of the current rates at time slot n and the effective rates up to time slot n of each transmitter $k \in \mathcal{T}$ respectively. Our objective is to solve the following convex optimization problem:

$$\max_{\boldsymbol{\theta}\in\bar{\mathcal{R}}} \sum_{k=1}^{T} |\mathcal{D}(k)| U_k^\alpha(\theta_k), \tag{24}$$

where $U_k^\alpha(\theta_k)$ is given by (19) and $\bar{\mathcal{R}}$ is

$$\bar{\mathcal{R}} = \Big\{ \mathbf{r} = (r_1, \ldots, r_T) : \exists\, \mathbf{r}(\mathbf{s}) \in \mathcal{R}_n(\mathbf{s}),\ \forall \mathbf{s} \in \mathcal{S} \\ \text{s.t. } \mathbf{r} = \int_{\mathcal{S}} \mathbf{r}(\mathbf{s}) f_{\mathcal{S}}(\mathbf{s}) d\mathbf{s} \Big\}. \tag{25}$$

In the sequel, we present an optimal, centralized policy, which at each slot takes joint scheduling and rate control decisions with the objective to solve (24). At every slot n the policy has knowledge of both the current channel conditions at slot n and the effective rates of each transmitter until the end of slot $n-1$. Consider the following set of rate vectors

$$\mathcal{M}(\boldsymbol{\theta}, \mathbf{s}) = \left\{ \tilde{\mathbf{r}} \ : \ \tilde{\mathbf{r}} = \underset{\mathbf{r}\in\tilde{\mathcal{R}}_n(\mathbf{s})}{\arg\max} \left\{ \nabla_{\boldsymbol{\theta}} U_k^\alpha(\boldsymbol{\theta})^T \mathbf{r} \right\} \right\}, \tag{26}$$

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991

where $\tilde{\mathcal{R}}_n(\mathbf{s})$ is the discrete set of rate vectors obtained by redefining (2)[3] and $\nabla_{\boldsymbol{\theta}} U_k^{\alpha}(\boldsymbol{\theta})$ is the vector of partial derivatives of the utility function with respect to the average rate vector $\boldsymbol{\theta}$. For the special case of proportional fairness, i.e., when $\alpha = 1$, the gradient vector may not be well-defined for small values of n, when some of the θ_ks are zero. To resolve this problem, (26) can be modified as $\mathcal{M}(\boldsymbol{\theta}, \mathbf{s}) = \left\{\tilde{\mathbf{r}} \ : \ \tilde{\mathbf{r}} = \arg\max_{\mathbf{r} \in \tilde{\mathcal{R}}_n(\mathbf{s})} \sum_{k=1}^{T} |\mathcal{D}(k)| \frac{r_k}{\theta_k + d_k}\right\}$, where d_ks are arbitrarily small positive constants. At every time slot n, when $\mathbf{S}_n = \mathbf{s}$ for some $\mathbf{s} \in \mathcal{S}$ and $\boldsymbol{\theta}_{n-1} = \boldsymbol{\theta}$, the optimal *multicast transmission rate vector* $\mathbf{R}_n(\boldsymbol{\theta}, \mathbf{s})$ is selected arbitrarily from the set described in (26), i.e.,

$$\mathbf{R}_n(\boldsymbol{\theta}, \mathbf{s}) \in \mathcal{M}(\boldsymbol{\theta}, \mathbf{s}). \tag{27}$$

This policy was shown in [9] to be optimal with respect to solving the optimization in (24). In fact the results in [9] are more general and include joint rate and power control decisions.

5. SIMULATION RESULTS

In this section, we analyze the performance of the proposed policies through a set of simulations. Throughout our simulation analysis we consider a *static* single-hop, wireless network with three transmitters and six receivers as shown in Figure 3. Specifically, $\mathcal{D}(1) = \{1, 2, 3\}$, $\mathcal{D}(2) = \{4, 5\}$ and $\mathcal{D}(3) = \{6\}$. The duration of a time slot is assumed to be equal to one second. For simplicity, we set the maximum transmission powers equal for each transmitter, i.e., $P(k) = P$, $k = 1, 2, 3$, where $P = 6.0 * 10^{-5}$ Watts. Further, the power of the additive white Gaussian noise is assumed common at all receivers and equal to $N = 3.34 * 10^{-6}$ Watts.

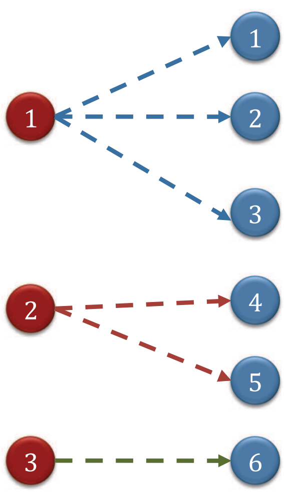

Figure 3: Simulation topology

We focus only on the criterion of proportional fairness. Towards this end, we consider three proportionally fair policies and compare their performance by varying the level of interference at every receiver. Specifically, the first policy we consider is a pure TDMA scheme that allocates the probability with which each transmitter accesses the channel so that its effective rate is proportionally fair. The second policy we consider is a restricted rate control policy that can take upon actions from both the schemes "SCH" and "LR", i.e., that can either activate the three transmitters one at a time or all together. Again, the probability with which each action is selected is such that the effective rate at each receiver is proportionally fair (i.e., the optimal probability distribution solves the problem defined in (6) -(8)). Finally, we consider a general rate control policy given in Section 4, where the policy can select any out of the total $2^3 - 1$ possible rate control and scheduling decisions to achieve proportional fairness. We compare the performance of the above three policies when the channel effects are due to pure path loss (i.e., static network case). Towards this end, we parameterize the path loss matrix $\mathbf{G}$, defining the path losses between the 3 transmitters and the 6 receivers, as follows:

$$\mathbf{G} = \begin{bmatrix} 0.8 & 0.9 & 0.75 & \beta & \beta & \beta \\ \beta & \beta & \beta & 0.85 & 0.9 & \beta \\ \beta & \beta & \beta & \beta & \beta & 0.7 \end{bmatrix},$$

where $\beta \in [0, 1]$ is a parameter we define as the *interference coefficient* since it multiples only the interfering channel coefficients. This parameter reflects the level of interference between each multicast session in the network. For example, when $\beta = 0$, the channels between the three multicast sessions can be seen as three parallel channels that can operate simultaneously without causing any interference to each other. At the other extreme, when $\beta = 1$, the path losses over the interfering channels are equal to 1, and therefore the level of interference at every receiver is very high. Throughout this section, we assume that the data rate $r(\cdot)$ is given by the Shannon formula, i.e., $r(\mathrm{SINR}) = \log_2(1 + \mathrm{SINR})$, under the assumption of unit bandwidth.

In Figures 4, 5 and 6, the proportionally fair rates of each multicast session are plotted as a function of the interference coefficient β for the three policies considered. As expected, when interference levels are low (i.e., $\beta = 0$), both the restricted and the general proportionally fair rate control policies achieve equivalently and much better than a pure TDMA scheme as when interference is negligible the LR scheme is the best choice at all times. On the other hand, when the level of the interference is high (i.e., $\beta = 1$) all three policies are converging to the TDMA schedule as any simultaneous transmission is strictly suboptimal. Naturally, the general rate control policy achieves better than the restricted case as it has more actions available at its disposal. However, we observe that in this specific example, the performance difference between the two policies are not significant, advocating that the restricted rate control policy can be useful at least under certain scenarios. One last point worth mentioning is the following observation. In Figure 6, the effective rate of transmitter 3 under the restricted policy is higher than that of the general policy for a certain range of β, which initially may appear counter-intuitive. But in a second thought, this can be actually expected as the objective considered in these simulations is the proportional fairness as defined in (6). Specifically, although the rate of the unicast source 3 is achieving a higher rate under the restricted policy for certain ranges of β, the effective rates of multicast source 1 and 2 under the general rate control policy are higher than the ones corresponding to the restricted policy under the same range of β. Given that multicast sessions 1 and 2 serve more receivers, it is natural that a policy favoring these two sessions is bound to achieve better with respect to the criterion given in (6).

6. CONCLUSIONS

In this paper, we obtained a joint scheduling and rate control policy that at any given time identifies a set of transmitters to access the channel and specifies their respective transmission rates. We considered both static as well as time-varying wireless networks. In the case of static networks we focused on the objectives of total throughput maximization and proportional fairness. When

[3] In fact, our results are valid even if the optimization of (26) is performed over the region $\mathcal{R}_n(\mathbf{s})$ or more complex regions acquired through joint rate and power control as in [9].

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991

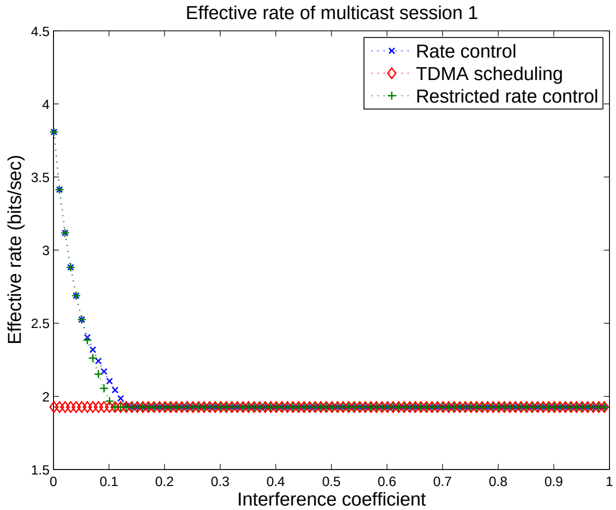

Figure 4: Effective rate of transmitter 1 with increasing β.

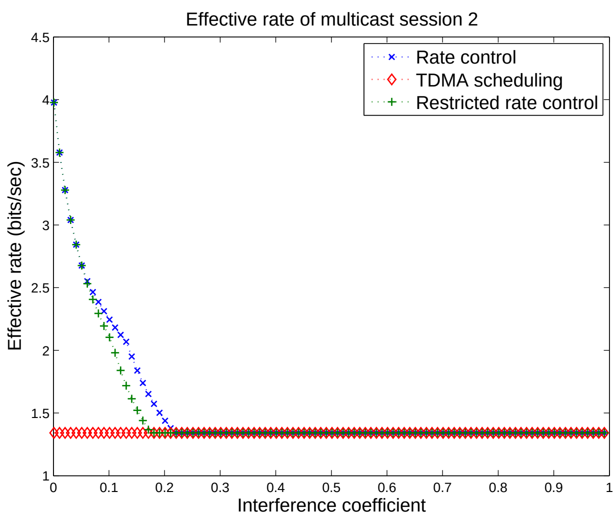

Figure 5: Effective rate of transmitter 2 with increasing β.

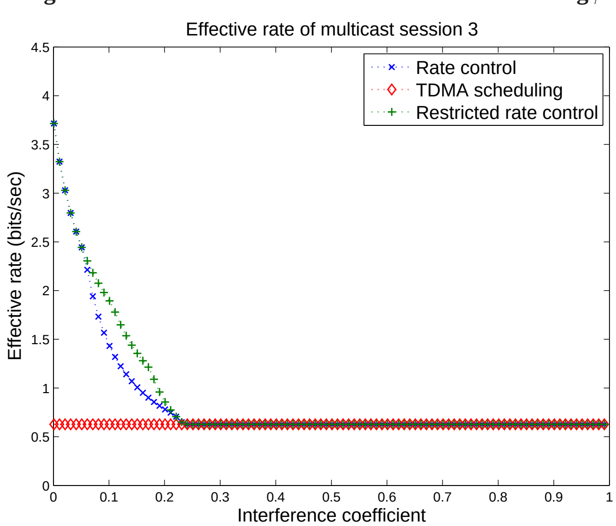

Figure 6: Effective rate of transmitter 3 with increasing β.

the networks are time-varying we considered more general utility functions that can capture the objectives of total throughput, proportional fairness, max-min fairness, etc. Finally, we verified our analytical results through a set of simulations.

7. REFERENCES

[1] R.D. Yates, C. Raman, and N.B. Mandayam. Fair and efficient scheduling of variable rate links via a spectrum server. In *Proceedings of IEEE International Conference on Communications (ICC)*, June 2006.

[2] A. Pantelidou and A. Ephremides. Optimal rate control policies for proportional fairness in wireless networks. In *Proceedings of Conference on Information Sciences and Systems (CISS)*, March 2008.

[3] R. Agrawal, A. Bedekar, R. J. La, and V. Subramanian. Class and channel condition based weighted proportional fair scheduler. In *Proceedings of International Telegraffic Congress (ITC)*, 2001.

[4] H. J. Kushner and P.A. Whiting. Asymptotic properties of proportional-fair sharing algorithms. In *Proceedings of 40th Annual Allerton Conference on Communication, Control, and Computing*, September 2002.

[5] H. J. Kushner and P.A. Whiting. Convergence of proportional-fair sharing algorithms under general conditions. *IEEE Transactions on Wireless Communications*, 3(4), 2004.

[6] R. Agrawal and V. Subramanian. Optimality of certain channel aware scheduling policies. In *Proceedings of 40th Annual Allerton Conference on Communication, Control, and Computing*, September 2002.

[7] A. L. Stolyar. On the asymptotic optimality of the gradient scheduling algorithm for multiuser throughput allocation. *Operations Research*, 53(1), January- February 2005.

[8] H. Won, H. Cai, D. Y. Eun, K. Guo, A. Netravali, I. Rhee, and K. Sabnani. Multicast scheduling in cellular data networks. In *Proceedings of IEEE INFOCOM*, May 2007.

[9] A. Pantelidou and A. Ephremides. A cross-layer view of wireless multicast optimization. In *Proceedings of 46th Annual Allerton Conference on Communication, Control, and Computing*, September 2008.

[10] A. Goldsmith. *Wireless communications*. Cambridge university press, 2005.

[11] F.P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1), 1997.

[12] B. Radunović and J. Le Boudec. Rate performance objectives of multi-hop wireless networks. In *Proceedings of IEEE Infocom*, March 2004.

[13] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5), 2000.

Digital Object Identifier: 10.4108/ICST.WICON2008.4991
http://dx.doi.org/10.4108/ICST.WICON2008.4991