

Discounted-Rate Utility Maximization (DRUM): A Framework for Delay-Sensitive Fair Resource Allocation

Atila Eryilmaz

Department of Electrical and Computer Engineering
The Ohio State University
Email: eryilmaz.2@osu.edu

Irem Koprulu

Department of Electrical and Computer Engineering
The Ohio State University
Email: irem.koprulu@gmail.com

Abstract—We introduce a new optimization framework, built over a *discounted-rate* metric, that captures the sensitivities of wireless users to time-variations in their fairness measure of rate allocations. The resulting, so-called, Discounted-Rate Utility Maximization (DRUM) formulation not only accommodates traditional long-term and less-explored instant fairness concepts in its extremes, but also encompasses all intermediate degrees of sensitivity to fluctuations in the users’ rate allocations.

After introducing the versatile DRUM formulation, we fully characterize its solution in the *instantly-fair* and *long-term-fair* extremes for the general class of ω -weighted α -fair utility functions. These solutions reveal the non-trivial impact of fading channel statistics and the utility function parameters on the rate allocations, even under perfectly symmetric network conditions. In particular, we demonstrate that the rate allocations lie between the *maximum* and the *harmonic mean* of the fading-channel rates.

To achieve rates in-between these extremes, we also address the general solution of DRUM by proposing a novel low-complexity dynamic rate allocation algorithm that does not require the knowledge of the channel statistics. This algorithm is proven to achieve the optimal performance of the *instantly-fair* and *long-term-fair* solutions as the discount parameter approaches its lower and upper limits, respectively. We also study the fairness and rate allocation characteristics of our algorithm for intermediate values of the discount parameter in a Rayleigh-Fading environment.

I. INTRODUCTION

Fair allocation of shared resources in a multi-user communication system has been a topic of great interest and activity over decades. The core objective of these efforts has been to obtain comprehensive models and systematic means for sharing heterogeneous resources amongst multiple users so that each user is satisfied (according to varying concepts of “satisfaction”) with its relative share *in the long-run*. This paper releases this implicit assumption on the long-term measure of fairness in order to accommodate varying degrees of *short-term* sensitivities of users to their received share.

Initial breakthroughs in the well-founded formulation and systematic solution of the fair resource allocation problem has been made in the seminal works of [1], [2] (see [3] for more references). These works formulated the problem of fair resource allocation in *wired* communication networks, with Internet in mind, through a *static utility optimization problem* subject to link capacity constraints. Through the use of dual optimization methods, they have also developed decentralized rate allocation (also called, congestion control) and scheduling strategies that converge to the optimal fair allocation. The unifying work of [4] expanded the reach of this approach by

introducing a common class of, so called, ω -weighted α -fair utility functions (cf. (6)) that encapsulates all important fairness concepts, including *proportional fairness*, *MaxMin fairness*, *sum-rate fairness* (cf. Def. 2) within a common formulation.

This comprehensive mathematical optimization framework for fair allocation is later embraced by the wireless networking research community in the fruitful development of efficient and fair allocation of *wireless* network resources (see, for example, [5], [6], [7], [8], [9]). These, and many subsequent works have unified previously-designed queue-length-based *backpressure* routing and *max-weight* scheduling strategies (developed in the seminal works of [10], [11]) with the fair rate allocation framework under the common umbrella of *Network Utility Maximization* (NUM) [12], also called *Cross-Layer Design* [13] or *Stochastic Network Optimization* [14] in different contexts.

These works have revealed that queue-lengths in the previous works can be viewed as Lagrange multipliers of related constraints of an associated static utility maximization problem, and vice versa. This revelation has led to many interesting follow-up efforts, which are still very active today, that utilize this connection to develop new adaptive policies by introducing different queue-lengths and virtual queue-lengths that accommodate new constraints of interest, such as reliability constraints, delay constraints, deadline constraints, etc.

However, throughout this development the measure of fairness in the utility function has remained to be based on *long-term average rates*. While this may be an acceptable measure in a wired communication network whereby the resources are essentially static, it is no longer the case in mobile wireless networks whereby resources are time-varying and possibly highly heterogeneous. Since the mobile users will be subject to time-varying channel conditions, their measure of fairness will typically depend on their rates in the short-run rather than the long-run. With increasingly mobile and increasingly delay-sensitive applications in the horizon, therefore, it is essential that the fair resource allocation framework in wireless networks accommodates varying delay-sensitivities into its formulation and solution. In this work, we address this need as follows:

- We introduce (cf. Section II) a Discounted-Rate Utility Maximization (DRUM) formulation that uses β -discounted user rates. DRUM not only encapsulates NUM as a special case but also extends it significantly by allowing different degrees of sensitivity to time-variations in the user rates.

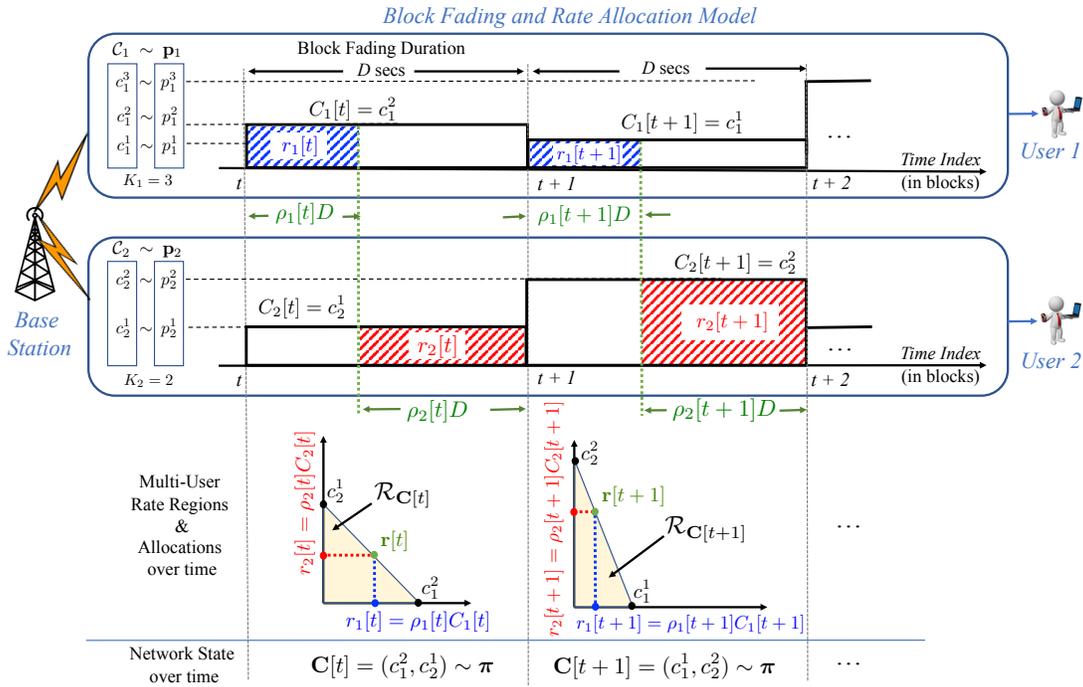


Fig. 1. Channel fading and rate allocation model for a two-user setting: In each fading-block t of duration D seconds, the maximum achievable rate $C_n[t]$ for each user $n \in \{1, 2\}$ is randomly realized from its corresponding set $C_n \triangleq \{c_n^1, \dots, c_n^{K_n}\}$ of rates according to a distribution $\mathbf{p}_n \triangleq \{p_n^1, \dots, p_n^{K_n}\}$. This determines the achievable rate region $\mathcal{R}_{C[t]}$ for the block, from which individual rates $\mathbf{r}[t] = (r_1[t], r_2[t])$ can be allocated to individual users. These rates are achieved through a *time-sharing* strategy whereby $\rho_n[t]$ determines the fraction of the block allocated to user n , as illustrated in slots t and $t+1$.

- We fully characterize (cf. Section III) the solutions of DRUM for the two extremes of so-called *long-term fair* and *instantly fair* allocations for the general class of ω -weighted α -fair utility functions, and also explicitly describe the impact of channel statistics on these allocations.
- We develop (cf. Section IV) a novel *Dynamic-DRUM (D-DRUM) Algorithm* for solving the DRUM problem not only at the extremes of $\beta \in \{0, 1\}$ but for the general case of any $\beta \in [0, 1)$. We then prove that D-DRUM achieves the optimal instantly-fair and long-term-fair allocations at the extremes of $\beta = 0$ and $\beta \uparrow 1$, and show that it stably spans rates in between these extremes with intermediate choices of β .
- We present (cf. Section V) comprehensive investigations of the performance characteristics of the Dynamic DRUM Algorithm that go beyond average rates. In particular, we present several interesting observations on the *distribution of, correlation between* individual user rates with varying β .

Throughout the paper, we consistently denote: random variables with upper-case letters, e.g., X ; realizations with lower-case letters, e.g., x ; sets with script letters, e.g., \mathcal{X} ; vectors with boldface letters, e.g., \mathbf{X} or \mathbf{x} ; component-wise inequalities with \succ and \succeq ; and the non-negative real numbers with \mathbb{R}_+ .

II. SYSTEM MODEL AND DRUM PROBLEM FORMULATION

A. Channel and Network Model

We consider the service of N users over wireless fading channels which are block-fading in time (see Fig. 1). For each user, the channel state is constant over blocks of duration of D seconds that are synchronized among all users so that the

network operates in a common slotted time t aligned with these blocks. While the results of this work hold when the channel states are dependent across users and dependent but ergodic over time, for simplicity of exposition, we will assume that channel states change independently from one block to the next.

User $n \in \{1, \dots, N\}$ has a set of K_n possible non-negative-valued channel states $C_n \triangleq \{c_n^1, \dots, c_n^{K_n}\}$ which are observed with corresponding probabilities $\mathbf{p}_n \triangleq \{p_n^1, \dots, p_n^{K_n}\}$. In particular, the *channel state* of user n at time t is a random variable $C_n[t] \in C_n$ with distribution \mathbf{p}_n , i.e., $\Pr(C_n[t] = c_n^k) = p_n^k$, and indicates the maximum achievable rate for that user in that block. Without loss of generality, we assume that $\mathbb{E}[C_n] > 0$, since otherwise we can omit such a user from the network.

The *network state* at time t is given by the realizations of all user states at that time, i.e., $\mathbf{C}[t] \triangleq (C_1[t], \dots, C_N[t]) \in \mathcal{C} \triangleq \prod_{n=1}^N C_n$ where $\prod \cdot$ indicates the Cartesian product of sets. Hence, a given network state $\mathbf{c} \triangleq (c_1, \dots, c_N) \in \mathcal{C}$ is observed with p.m.f. $\boldsymbol{\pi} \triangleq (\pi_{\mathbf{c}})_{\mathbf{c}}$, where $\pi_{\mathbf{c}} \triangleq \Pr(\mathbf{C}[t] = \mathbf{c})$.

B. Instantaneous and Average Achievable Rate Regions

The channel state of a user indicates the highest rate that can be achieved during the given block duration, *if the channel is used by that user alone*. If multiple users are utilizing the channel simultaneously, interference prevents either user from achieving their highest rate. In that case, we consider a time-sharing in which users access the channel at disjoint sub-intervals of the block duration. While time-sharing is a commonly used strategy and allows us to get closed-form solutions to various optimization problems in the sequel, we note that

it is also possible to accommodate more sophisticated multi-access rate regions with efficient allocation strategies as in [15]. If user n utilizes the channel during a $\rho_n[t]$ fraction of the block t , it can achieve a maximum rate of $R_n[t] = \rho_n[t]C_n[t]$. Accordingly, each vector $\boldsymbol{\rho}[t] \triangleq (\rho_n[t])_n$ in the set of fractions $\Psi \triangleq \{\boldsymbol{\rho} \succeq \mathbf{0} : \sum_{n=1}^N \rho_n \leq 1\}$ corresponds achieving the point $(\rho_1[t]C_1[t], \dots, \rho_N[t]C_N[t])$ in the multi-user rate region. Thus, for each network state $\mathbf{C}[t] = \mathbf{c} = (c_n)_{n=1}^N \in \mathcal{C}$, the *instantaneous achievable rate region* at time t is:

$$\mathcal{R}_{\mathbf{C}[t]} = \mathcal{R}_{\mathbf{c}} \triangleq \text{Conv}(\{\mathbf{0}, c_1\mathbf{e}_1, \dots, c_N\mathbf{e}_N\}) \subset \mathbb{R}_+^N, \quad (1)$$

where $\text{Conv}(\cdot)$ indicates the convex hull of a set, and \mathbf{e}_n is the n^{th} standard unit vector. The lower part of Fig. 1 illustrates the time-varying rate region $\mathcal{R}_{\mathbf{C}[t]}$ and the concept of time-sharing for a two-user setting. The *average achievable rate region* $\overline{\mathcal{R}}$, also called the *capacity region*, is obtained by averaging the achievable rate region over all possible network states $\mathbf{c} \in \mathcal{C}$:

$$\overline{\mathcal{R}} = \sum_{\mathbf{c} \in \mathcal{C}} \pi_{\mathbf{c}} \cdot \mathcal{R}_{\mathbf{c}}, \quad (2)$$

where the summation is *set addition*.

The following example illustrates the concepts we have introduced so far in a key two-state setting that will be revisited in later sections to demonstrate important characteristics of our new framework as we develop them.

Example 1 (Two-User and Two-State Channel Setting). We consider a two user scenario with independent and identically distributed (i.i.d.) maximum achievable rates $\{C_n[t]\}_{n=1,2}$ taking values from a two-state set $\mathcal{C} = \{c^H, c^L\}$ with $c^H > c^L \geq 0$, and $p \triangleq \mathbb{P}(C_n[t] = c^H) = 1 - \mathbb{P}(C_n[t] = c^L) \in (0, 1)$.

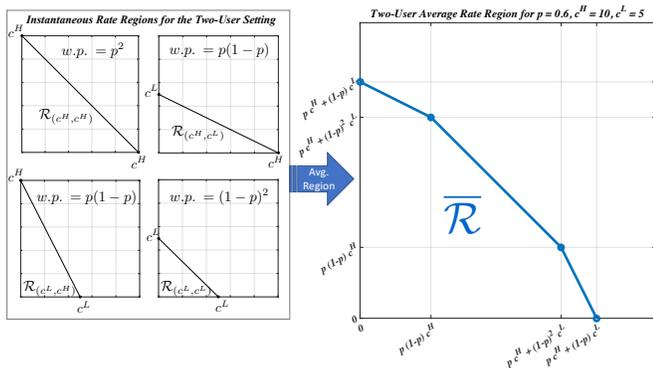


Fig. 2. Average rate region $\overline{\mathcal{R}}$ for a two-user network with i.i.d. maximum achievable rates $\{C_n[t]\}_n$ taking values from $\mathcal{C} = \{c^H, c^L\}$ with $c^H > c^L \geq 0$, and $p \triangleq P(C_n[t] = c^H) \in [0, 1]$.

Thus, the network state $\mathbf{C}[t]$ has four possible realizations $\mathcal{C} = \{(c^H, c^H), (c^H, c^L), (c^L, c^H), (c^L, c^L)\}$ with probabilities $p^2, p(1-p), p(1-p)$ and $(1-p)^2$, respectively. For each realization of the network state $\mathbf{C}[t]$, the instantaneous achievable rate region $\mathcal{R}_{\mathbf{C}[t]}$ at time t is given by a corresponding triangular area illustrated on the left side of Fig. 2. Then, the average achievable rate region $\overline{\mathcal{R}}$ is obtained from (2) by averaging these four triangular regions, resulting in a symmetric pentagon that is illustrated on the right side of Fig. 2 for the particular set of values $p = 0.6$, $c^H = 10$ and $c^L = 5$.

C. Discounted-Rate Utility Maximization (DRUM) Problem

Def. 1 (β -Discounted Rate). For a given $\beta \in [0, 1]$, we define the β -discounted rate of user n at time $t \geq 0$ as

$$R_n^{(\beta)}[t] \triangleq \frac{\sum_{\tau=-T_s+1}^t \beta^{t-\tau} R_n[\tau]}{\sum_{\tau=-T_s+1}^t \beta^{t-\tau}}, \quad t > -T_s, \quad (3)$$

where $T_s \in \{0, 1, \dots\}$ denotes the starting time of the system operation. Also, define the *accumulated¹ β -discounted rate* $\overline{R}_n^{(\beta)}[t]$ at time t as the limit of $R_n^{(\beta)}[t]$ with $T_s \rightarrow \infty$, i.e.,

$$\overline{R}_n^{(\beta)}[t] \triangleq \begin{cases} R_n[t], & \text{if } \beta = 0 \\ (1-\beta) \sum_{\tau=-\infty}^t \beta^{t-\tau} R_n[\tau], & \text{if } \beta \in (0, 1) \\ \liminf_{T_s \rightarrow \infty} \frac{1}{t+T_s} \sum_{\tau=-T_s+1}^t R_n[\tau], & \text{if } \beta = 1 \end{cases} \quad (4)$$

Finally, let $\overline{R}_n^{(\beta)}[t] \triangleq \mathbf{E}[\overline{R}_n^{(\beta)}[t]] \in \mathbb{R}_+$ be the *mean accumulated β -discounted rate* for user n at time t .

This definition provides an effective means of capturing varying degrees of delay sensitivity to service rates within a common framework. The numerator in (3) adds the rates from t backwards by decreasing the weights by a factor of β each time the time index is decremented. As such, by choosing different values of $\beta \in [0, 1]$, this operation allows us to emphasize or de-emphasize the past values of $R_n[\cdot]$ in the measure of rate at time t . The denominator in (3) serves as a normalization factor to balance the effect of β in the numerator.

We can see the versatility of this discounted-rate metric from (4): when $\beta = 0$ the accumulated β -discounted rate $\overline{R}_n^{(0)}[t]$ reduces to the *instantaneous rate* $R_n[t]$ achieved in the same slot; and when $\beta = 1$ the β -discounted rate $\overline{R}_n^{(1)}[t]$ becomes the *long-term time-average rate* received since the beginning of time. In-between these extremes, the parameter $\beta \in (0, 1)$ allows us to put different emphasis on the rates received in the recent and remote past.

This versatile metric lies at the center of a new utility maximization framework that we introduce next, which can accommodate delay sensitivities in the fair allocation. Using appropriate choices of the discount parameter $\beta \in [0, 1]$, this new formulation will be able to capture the sensitivity of users to delayed service within a unified framework, whereby $\beta = 0$ corresponds to extreme delay-sensitivity, while $\beta \uparrow 1$ approaches the delay-insensitive case.

Def. 2 (Discounted-Rate Utility Maximization (DRUM)). Suppose that the system starts at time $-T_s \leq 0$. Given a set of

¹We note that, even though we use the qualification *accumulated* for this metric due to its *discounted time-averaging* nature, $\{\overline{R}_n^{(\beta)}[t]\}_t$ can in general be a sequence of random variables that need not converge to a single value.

user utility functions $\{U_n(x)\}_{n=1}^N$, the Discounted-Rate Utility Maximization (DRUM) problem is given by:

$$\begin{aligned} \max \quad & \sum_{n=1}^N \bar{u}_n(\{R_n^{(\beta)}[t]\}_t) \\ \text{s.t.} \quad & \mathbf{R}[t] \in \mathcal{R}_{\mathbf{C}[t]}, \quad t > -T_s, \end{aligned} \quad (5)$$

with $\bar{u}_n(\{R_n^{(\beta)}[t]\}_t) \triangleq \liminf_{T \rightarrow \infty} \frac{1}{T + T_s} \sum_{t=-T_s+1}^T \mathbf{E} [U_n(R_n^{(\beta)}[t])]$ measuring the *average utility* achieved by user n over time with respect to its β -discounted rate performance.

In particular, we will study DRUM for the wide class of ω -weighted α -fair utility functions defined as:

$$U_n^{[\alpha]}(x) \triangleq \begin{cases} \omega_n \frac{x^{1-\alpha}}{1-\alpha}, & \alpha \geq 0, \alpha \neq 1 \\ \omega_n \log(x) & \alpha = 1, \end{cases}, \quad x > 0, \quad (6)$$

for a given weight vector $\boldsymbol{\omega} \triangleq (\omega_n)_n \succ \mathbf{0}$, i.e., $\omega_n > 0$ for all n . Depending on the value of α these utility functions span a wide range fairness from *sum-rate* (when $\alpha = 0$) to *proportional* (when $\alpha = 1$) to *MinMax-fair* (as $\alpha \uparrow \infty$) allocations [3], [4].

III. CHARACTERIZATION AND COMPARISON OF OPTIMAL INSTANTLY-FAIR & LONG-TERM-FAIR DRUM SOLUTIONS

In Def. 2, we have introduced our discounted-rate utility maximization (DRUM) problem for a set of utility function parameters: $\alpha \geq 0$, $\boldsymbol{\omega} \succ \mathbf{0}$; and discount parameter: $\beta \in [0, 1]$. In this section, our goal is to characterize the complete solution of DRUM, for all $\alpha \geq 0$ and $\boldsymbol{\omega} \succ \mathbf{0}$, under the two tractable cases of $\beta = 0$ and $\beta = 1$, corresponding, respectively, to the *instantly* and *long-term* ω -weighted α -fair rate allocations.

The investigation in this section will reveal (cf. Propositions 1 and 2) how the statistics of channel state processes factor into the solution of DRUM for different choices of α , $\boldsymbol{\omega}$, and β parameters, even under perfectly symmetric utility functions and channel conditions (cf. Corollary 1). These results will act as benchmarks in the next section when we design and investigate a dynamic policy that will work for all $\alpha, \boldsymbol{\omega}, \beta$.

First, we characterize the *instantly-fair*, i.e., $\beta = 0$, DRUM solution for all allowed values of $(\boldsymbol{\omega}, \alpha)$ parameters of $\{U_n^{[\alpha]}\}_n$.

Proposition 1 (Complete Instantly-Fair Solution of DRUM).

► For $\alpha = 0$, an² optimal instantly-fair rate allocation $\dot{\mathbf{R}}^{(0)}[t] \triangleq (\dot{R}_n^{(0)}[t])_n$ when $\mathbf{C}[t] = \mathbf{c} = (c_n)_n$ is given by:

$$\dot{R}_n^{(0)}[t] = \dot{R}_n^{(0)}(\mathbf{c}) = \frac{c_n \mathcal{I} \left(n \in \arg \max_{m \in \{1, \dots, N\}} \{\omega_m c_m\} \right)}{\left| \arg \max_{m \in \{1, \dots, N\}} \{\omega_m c_m\} \right|}, \quad (7)$$

for each n . In the last expression, $\mathcal{I}(A)$ denotes the indicator function of event A , and $|\mathcal{A}|$ is the cardinality of set \mathcal{A} .

²We note that in both of these cases of $\beta = 0$ and $\beta = 1$, when $\alpha = 0$ there are possibly multiple solutions to DRUM, including the one we provide.

► For $\alpha > 0$, the optimal instantly-fair rate allocation $\dot{\mathbf{R}}^{(0)}[t] \triangleq (\dot{R}_n^{(0)}[t])_n$ for any $\mathbf{C}[t] = \mathbf{c} \succ \mathbf{0}$ is given by³:

$$\dot{R}_n^{(0)}[t] = \dot{R}_n^{(0)}(\mathbf{c}) = \frac{(\omega_n c_n)^{1/\alpha}}{\sum_{m=1}^N \frac{(\omega_m c_m)^{1/\alpha}}{c_m}}, \quad \forall n, \quad (8)$$

Consequently, the accumulated rates $\dot{\mathbf{R}}^{(0)}[t] \triangleq (\dot{R}_n^{(0)}[t])_n$ (cf. (4)) achieved under the instantly-fair allocation are given by $\dot{\mathbf{R}}^{(0)}[t] = \dot{R}_n^{(0)}(\mathbf{C}[t])$ with the right-hand-side given by (7) and (8) when $\alpha = 0$ and $\alpha > 0$, respectively.

Proposition 1 (see [16] for the proof) explicitly provides the optimal instantly-fair DRUM solution for a given $(\boldsymbol{\omega}, \alpha)$ pair as a function of the network state random vector $\mathbf{C}[t]$. As such, the optimal rate allocation $\dot{\mathbf{R}}^{(0)}(\mathbf{C}[t])$ is a *random vector* that is governed by the distribution $\boldsymbol{\pi} = (\pi_{\mathbf{c}})_{\mathbf{c}}$ of \mathbf{C} and does not converge to a constant as t diverges. This comes from the highly delay-sensitive nature of the instantly-fair allocation due to the choice $\beta = 0$. In contrast, we will see in the next proposition that in the long-term-fair case when $\beta = 1$, the optimal allocation will be a constant vector. The above solution of the instantly-fair allocation also captures the impact of $\boldsymbol{\omega}$ and α parameters on the DRUM solution. We shall comment further on this impact in contrast with the long-term case after we provide the characteristics of the long-term ω -weighted α -fair DRUM solution, i.e., when $\beta = 1$, in the following proposition.

Proposition 2 (Complete Long-Term-Fair Solution of DRUM).

For any $\alpha \geq 0$, the optimal long-term-fair rates $\dot{\mathbf{R}}^{(1)}[t] \triangleq (\dot{R}_n^{(1)}[t])_n$ converge to a constant average accumulated rate vector $\dot{\mathbf{r}}^{(1)}$ (cf. Def. 1) that solves:

$$\dot{\mathbf{r}}^{(1)} \in \arg \max_{\mathbf{r} \triangleq (r_n)_n \in \bar{\mathcal{R}}} \sum_{n=1}^N U_n^{[\alpha]}(r_n), \quad (9)$$

where $U_n^{[\alpha]}(\cdot)$ is the ω -weighted α -fair function given in (6), and $\bar{\mathcal{R}}$ is the average achievable rate region defined in (2). The solution of (9) exists for $\alpha \geq 0$ and is unique for $\alpha > 0$. Given the solution of $\dot{\mathbf{r}}^{(1)}$ from (9), we can now characterize the optimal long-term-fair rate allocation $\dot{\mathbf{R}}^{(1)}[t]$ at each time t for different choices of α .

► For $\alpha = 0$, an³ optimal long-term-fair rate allocation $\dot{\mathbf{R}}^{(1)}[t] \triangleq (\dot{R}_n^{(1)}[t])_n$ when $\mathbf{C}[t] = \mathbf{c} = (c_n)_n$ is given by the same allocation (7) as in the instantly-fair case.

► For $\alpha > 0$, the optimal long-term-fair rate allocation $\dot{\mathbf{R}}^{(1)}[t] = \dot{\mathbf{R}}^{(1)}(\mathbf{c}) = \rho_n^*(\mathbf{c}) c_n$ when $\mathbf{C}[t] = \mathbf{c} = (c_n)_n$ is found by solving for $\{\rho^*(\mathbf{c})\}_{\mathbf{c} \in \mathcal{C}}$ that satisfies, for each n :

$$\sum_{\mathbf{c} \in \mathcal{C}} \pi_{\mathbf{c}} \rho_n^*(\mathbf{c}) c_n \mathcal{I} \left(n \in \arg \max_{m \in \{1, \dots, N\}} \frac{\omega_m c_m}{\left(\dot{r}_m^{(1)} \right)^\alpha} \right) = \dot{r}_n^{(1)}, \quad (10)$$

where $(\rho_n^*(\mathbf{c}))_n \in \Psi \triangleq \{\boldsymbol{\rho} \succeq \mathbf{0} : \sum_{n=1}^N \rho_n \leq 1\}$, $\forall \mathbf{c} \in \mathcal{C}$.

³Note that in this case, if any $c_n = 0$, then the utility function value of that user becomes $-\infty$, hence that scenario is excluded.

Proposition 2 (see [16] for the proof) characterizes the optimal long-term-fair DRUM solution implicitly through the solutions of (9) and (10) for a given (ω, α) pair. These equations can in general be difficult to solve and require knowledge of the network statistics π . While we defer the solution of these equations through simple updates to the next section, next we consider an important special case of the network scenario that allow us to explicitly solve them and provide an interesting comparison between the accumulated rate performances of instantly-fair and long-term-fair optimal allocations.

Corollary 1 (Comparison of Instantly and Long-Term-Fair Allocations under Symmetry). *Suppose that: (i) each channel $C_n[t] \in \mathcal{C} \triangleq \{c^1, \dots, c^K\}$ is independently and identically distributed⁴ (i.i.d.) according to a common distribution $\mathbf{p} \triangleq (p^1, \dots, p^K)$ with $\mathbb{P}(C_n[t] = 0) = 0$; and (ii) the $\{U_n^{[\alpha]}(\cdot)\}_n$ weights are $\omega_n = 1, \forall n$. Then, we can write the average accumulated discounted rate $\tilde{r}_n^{(\beta)}[t] \triangleq \mathbf{E} \left[\tilde{R}_n^{(\beta)}[t] \right]$ (cf. Def. 1) for the instantly and long-term-fair optimal allocations as follows: for each slot t ,*

$$\tilde{r}_n^{(\alpha)}[t] = \begin{cases} \frac{1}{N} \mathbf{E} \left[\max_{m \in \{1, \dots, N\}} C_m[t] \right], & \text{if } \alpha = 0 \\ \mathbf{E} \left[\frac{(C_n[t])^{\frac{1}{\alpha}}}{\sum_{m=1}^N (C_m[t])^{\frac{1}{\alpha}-1}} \right], & \text{if } \alpha > 0. \end{cases} \quad (11)$$

$$\tilde{r}_n^{(1)}[t] = \frac{1}{N} \mathbf{E} \left[\max_{m \in \{1, \dots, N\}} C_m[t] \right], \quad \forall \alpha \geq 0. \quad (12)$$

It is insightful to rewrite (11) for proportionally-fair ($\alpha = 1$) and MaxMin-fair ($\alpha \uparrow \infty$) cases:

$$\tilde{r}_n^{(\alpha)}[t] = \begin{cases} \frac{1}{N} \mathbf{E} [C_n[t]], & \text{if } \alpha = 1 \\ \mathbf{E} \left[\left(\sum_{m=1}^N \frac{1}{C_m[t]} \right)^{-1} \right], & \text{as } \alpha \uparrow \infty. \end{cases} \quad (13)$$

Corollary 1 reveals the non-trivial impact of the fairness parameter α on the average accumulated rates under the instantly and long-term fair allocations. In particular, (12) reveals that the long-term-fair optimal allocation $\tilde{\mathbf{R}}^{(1)}(\mathbf{c})$ is insensitive to the choice of $\alpha \geq 0$ value under these symmetric conditions, and equally shares the service between users that achieve the maximum achievable rate at the given network state \mathbf{c} . In contrast, the optimal instantly-fair allocation (11) is sensitive to α values: when $\alpha = 0$, i.e., when the utility functions are linear, it performs the same allocation as the optimal long-term-fair allocation; and when $\alpha > 0$ it becomes increasingly sensitive to the network state distribution as α increases.

The above corollary also reveals the interesting impact of the network state $\mathbf{C}[t]$ statistics on the achieved average accumulated rates. In particular, (12) shows that the average accumulated rate of long-term-fair allocation is governed by the *maximum* of all achievable rates $\{C_m[t]\}_m$ for all $\alpha \geq 0$. In contrast, (11) and (13) reveal that the average accumulated

rate of instantly-fair allocation varies from their *maximum* in the *sum-rate-optimal* case ($\alpha = 0$), to their *arithmetic mean* in the *proportionally-fair* case ($\alpha = 1$), to their *harmonic mean* in the *MaxMin-fair* case ($\alpha \uparrow \infty$). This relationship is further demonstrated in the following example.

Example 2 (DRUM Solutions for the Two-User and Two-State Setting). *For the same setting as in Example 1, the long-term-fair solution of DRUM (cf. Proposition 2) allocates all resources to the user who sees a good channel (if both users see a good channel, they share equally) regardless of α due to symmetry. This scheme achieves a high average rate for both users at the boundary of the average rate region $\bar{\mathcal{R}}$ as depicted in Fig. 3. This high rate, however, comes at the risk of possible high delay since it is insensitive to intermittent and bursty services. If a user's channel remains in a bad state for an extended duration, as it will in many realistic wireless settings, that user will receive low rates for that duration, resulting in delays in its average service experience.*

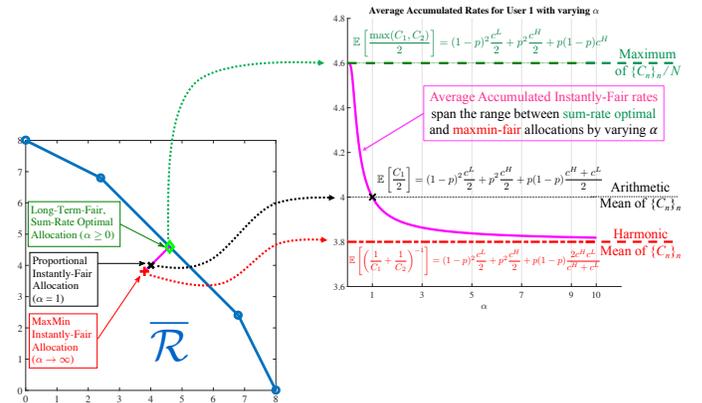


Fig. 3. Average accumulated instantly-fair rates achieved by user 1 in the two-user, two-state setting of Example 1 with varying $\alpha \geq 0$. As α goes from 0 to ∞ , the rate ranges between the sum-rate-optimal and MaxMin-fair allocations.

The instantly-fair solution of DRUM (cf. Proposition 1) achieves the same rates as the long-term-fair case for $\alpha = 0$. When $\alpha > 0$, however, it allocates positive rates to both users at every block, ensuring a steady distribution of the resources. Moreover, as α increases the rates of both users will become more similar, and equal to the average harmonic mean in the limit. Achieving fairness at each block, however, comes at the cost of lower rate. Fig. 3 shows how the achieved rate decreases with increasing α . The figure also shows that the range of instantly-fair rates lies between the maximum of the channel rates and the harmonic mean of the channel rates. As such, these statistics emerge as important measures in characterizing the largest cost of long-term-fair versus instant-fair in terms of the statistics of the channel rates.

IV. DYNAMIC ALGORITHM FOR GENERAL SOLUTION OF DRUM AND ITS OPTIMALITY GUARANTEES

In section III, we have presented the optimal DRUM solution for instantly-fair and long-term-fair rate allocations which correspond to the extreme values of the rate-discount parameter $\beta \in [0, 1]$. In both cases the solution requires the full knowledge

⁴Note that until here, we allowed channel states to be correlated across users.

of the channel statistics π . It is possible to develop simple dynamic algorithms that do not require the knowledge of π , which can be asymptotically long-term-fair-optimal in the limit as time goes to infinity. In particular, queue-length-based policies that achieve this goal have been developed in several works (e.g., [6], [7], [8]). However, we are interested not only in solving the instantly-fair and long-term fair allocations but in solving the problem in the range between these extremes, which calls for a new design. In this section, we address this need by introducing a dynamic algorithm that provides a solution to the DRUM problem for values of $\alpha \geq 0$ and $\beta \in [0, 1)$, and does not require the knowledge of the channel statistics π . Our algorithm is not restricted to symmetric channel conditions and weights, but applies to the most general setup.

Def. 3 (Dynamic-DRUM (D-DRUM) Algorithm). The Dynamic-DRUM algorithm starts at an arbitrary time $t = -T_s$ with the β -discounted rate initialized to $R_n^{(\beta)}[-T_s] = 0$ for all $n \in \{1, \dots, N\}$. Then, at time slot $t > -T_s$, given $\mathbf{R}^{(\beta)}[t-1] = (R_n^{(\beta)}[t-1])_n$ and the network state $\mathbf{C}[t]$, the algorithm allocates $\mathbf{R}^{(\beta)}[t]$ as:

$$\mathbf{R}[t] \in \arg \max_{\mathbf{r} \in \mathcal{R}_{\mathbf{C}[t]}} \sum_{n=1}^N U_n \left(\beta R_n^{(\beta)}[t-1] + (1-\beta)r_n \right) \quad (14)$$

and uses this $\mathbf{R}[t]$ to update $\mathbf{R}^{(\beta)}[t]$ from $\mathbf{R}^{(\beta)}[t-1]$ as:

$$\mathbf{R}^{(\beta)}[t] = \beta \mathbf{R}^{(\beta)}[t-1] + (1-\beta)\mathbf{R}[t] \quad (15)$$

For the particular class of ω -weighted α -fair utility functions $\{U_n^{[\alpha]}(x)\}_n$ defined in (6), the solution to the maximization problem in (14) is given by:

$$R_n[t] = \left(\left(\frac{(1-\beta)^{1-\alpha} c_n \omega_n}{\lambda^*} \right)^{\frac{1}{\alpha}} - \frac{\beta \cdot R_n^{(\beta)}[t-1]}{1-\beta} \right)^+ \quad (16)$$

where $(x)^+ \triangleq \max\{0, x\}$ and $\lambda^* > 0$ satisfies

$$\sum_{n=1}^N \frac{1}{c_n} \left(\left(\frac{(1-\beta)^{1-\alpha} c_n \omega_n}{\lambda^*} \right)^{\frac{1}{\alpha}} - \frac{\beta \cdot R_n^{(\beta)}[t-1]}{1-\beta} \right)^+ = 1 \quad (17)$$

Note that (16) and (17) correspond to a waterfilling solution, which can be solved exactly with $O(N \log(N))$ computational-complexity: At each time t , the algorithm determines the number of users which can be allocated a nonzero rate without violating the total capacity constraint in (17). Finding this number requires sorting the N users based on their maximum λ value that makes $R_n[t]$ zero in (16). The $O(N \log(N))$ complexity follows from noting that this sorting operation governs the computational complexity of the D-DRUM algorithm.

We note that the update rule (15) bears a similarity to the update rule in the proportionally-fair rate allocation policy introduced in [17], but differs from it in two important aspects. First, the rate allocation decision in (14) differs from that in [17]. Second, our policy is not specific to the proportionally-fair case but accommodates all α -fair utility functions, and integrates the discount parameter β into its allocation. In this respect, D-DRUM can be viewed as an alternate to the proportionally-fair scheduler in [17] when $\alpha = 1$ and $\beta \uparrow 1$,

and also a generalization in accommodating all other α -fairness metrics. In the following theorem, we establish this fact that the rate allocation achieved by our algorithm achieves the optimal rate allocations for both $\beta = 0$ and as $\beta \uparrow 1$.

Theorem 1 (Instant and Long-Term Optimality of D-DRUM). *D-DRUM achieves the instantly-fair optimal allocation when $\beta = 0$ and converges to the long-term-fair optimal allocation as $\beta \uparrow 1$. Specifically, the discounted rates $\mathbf{R}^{(\beta)}[t]$ achieved under D-DRUM satisfies:*

$$\mathbf{R}^{(0)}[t] = \dot{\mathbf{R}}^{(0)}[t], \quad \forall t > -T_s, \quad (18)$$

where $\dot{\mathbf{R}}^{(0)}[t]$ is defined in Proposition 1, and

$$\lim_{\beta \uparrow 1} \lim_{t \rightarrow \infty} \mathbf{R}^{(\beta)}[t] = \dot{\mathbf{R}}^{(1)}, \quad \text{w.p. 1}, \quad (19)$$

where $\dot{\mathbf{R}}^{(1)}$ is defined in Proposition 2.

While the proof of optimality for instantly-fair allocation is immediate, the proof of optimality for long-term-fair allocation requires a more subtle upper and lower-bounding to establish (see [16] for the complete proof). With this theorem, D-DRUM yields a new means for solving the long-term fair allocation problem that is attacked in several earlier works through dual methods (e.g., [6], [7], [8]). Yet, D-DRUM also provides the versatility to emphasize short-term fairness experience through the choice of β parameter. Next, we revisit the two-user and two-state scenario from Examples 1 and 2 to demonstrate that our D-DRUM algorithm attains the long-term and instantly-fair allocations for each $\alpha \geq 0$ as its β parameter approaches 0 and 1, and spans all rates in-between by varying $\beta \in (0, 1)$.

Example 3 (D-DRUM Algorithm Performance for the Symmetric Two-User and Two-State Setting). *For the same setting as in Example 1, the average rate performance of the D-DRUM algorithm with varying β is depicted in Fig. 4.*

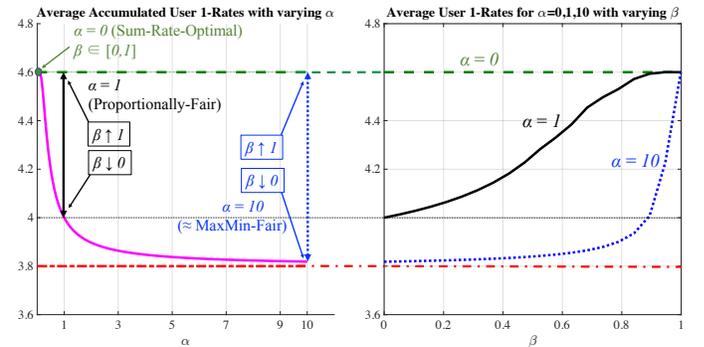


Fig. 4. Average accumulated β -discounted rates achieved by user 1 in the two-user, two-state setting of Example 1 with varying $\beta \in [0, 1)$ for $\alpha \in \{0, 1, 10\}$. For each α , as β goes from 0 to 1, the average accumulated β -discounted rate increases from its instantly-fair value to the long-term-fair value as shown in the plot on the right.

The left plot repeats the right plot in Fig. 3 to indicate the span between the instantly-fair and long-term-fair allocations for sum-rate-optimal ($\alpha = 0$ in green), proportionally-fair ($\alpha = 1$ in black), and essentially MaxMin-fair ($\alpha = 10$ in blue) cases. Then, the right plot gives the average rate that our D-DRUM algorithm provides to user 1 as its discount parameter β ranges

from 0 to 1. For each case of $\alpha = \{0, 1, 10\}$, we see that the algorithm achieves the instantly-fair and the long-term-fair allocations in its extremes, and spans all value in-between as β spans $(0, 1)$. We also see by comparing the right-hand plots for $\alpha = 1$ and $\alpha = 10$ that as α increases, the average rate performance becomes less sensitive to low β values and highly sensitive to changes in β around 1.

Example 3 illustrates the performance of D-DRUM for the symmetric setting of i.i.d. channel statistics and equal weights ω in the choice of α -fair utility functions. In the next example, we consider the case of *non-i.i.d.* channels and *non-uniform* utility weights to show that the optimality characteristics of D-DRUM continue to hold more generally.

Example 4 (Dynamic-DRUM Algorithm Performance for an Asymmetric Two-User and Two-State Setting). We consider an asymmetric two user scenario where both users' maximum achievable rates take values from different sets ($\mathcal{C}_1 = \{4, 11\}$ and $\mathcal{C}_2 = \{5, 10\}$) with different probabilities ($\mathbf{p}_1 = (0.6, 0.4)$ and $\mathbf{p}_2 = (0.4, 0.6)$). Fig. 4 depicts the average rate performance of D-DRUM for a fixed $\alpha = 1$ (i.e. instantly proportionally-fair allocation) and a range of β and ω .

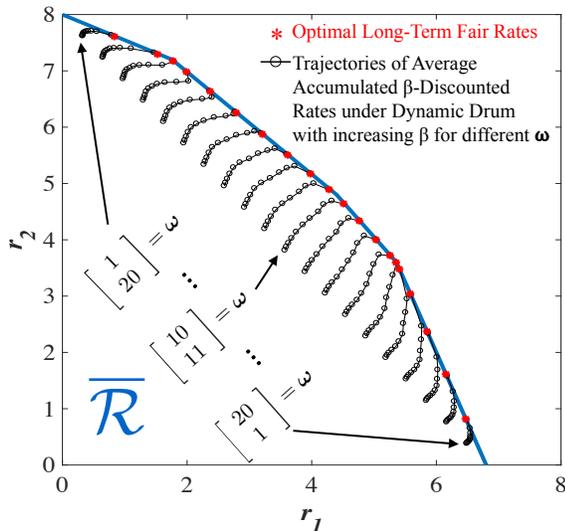


Fig. 5. Trajectories of the average accumulated β -discounted rates achieved by our Dynamic DRUM Algorithm for the asymmetric two-user, two-state setting with varying $\beta \in [0, 1)$ and different ω . For the case of $\alpha = 1$, trajectories for $\omega \in \{(20, 1)^T, (19, 2)^T, \dots, (1, 20)^T\}$ are illustrated as β starts from 0 and as it approaches 1. The rates converge to the long-term optimal rates indicated at the boundary of the asymmetric average achievable rate region $\bar{\mathcal{R}}$.

Each trajectory corresponds to a different weight vector ω while the individual points on each trajectory correspond to a different discount parameter β . For each trajectory, as β starts from 0 and approaches 1, the achieved average rates increase towards the optimal long-term fair allocation on the boundary of the average achievable rate region $\bar{\mathcal{R}}$. The particular limit points depend on how ω weighs one user over the other. For all sets of parameters, the performance of the D-DRUM algorithm approaches the optimal long-term rates as $\beta \uparrow 1$. Here, the optimal long-term rates are independently calculated using a dual queue-length-based approach in [8].

Examples 3 and 4 not only confirm the optimality features of D-DRUM in its extremes (both for symmetric and asymmetric conditions), but also show how the α and β values can jointly impact the average rates provided to the users. These show that D-DRUM is a highly stable and effective strategy for accommodating user sensitivities to short-term service variations.

V. FURTHER OBSERVATIONS ON THE STATISTICAL PERFORMANCE OF THE DYNAMIC-DRUM ALGORITHM

In order to develop a more comprehensive understanding of D-DRUM performance, in this section we extend our previous investigations in two aspects: (1) we consider more realistic continuous-valued fading channels; (2) we study the *distribution* of the user rates under D-DRUM rather than their averages.

In particular, we consider a two-user⁵ i.i.d. Rayleigh fading scenario where scheduled users transmit at a fixed power P over additive white Gaussian noise with power N_0 and $\frac{P}{N_0} = 3dB$. Then, for the channel gain $H_n[t]$ of user n at time t , the maximum achievable rate $C_n[t]$ becomes $\log\left(1 + |H_n[t]|^2 \frac{P}{N_0}\right)$, where $|H_n[t]|^2$ is exponentially distributed.

• **Distribution of Instantaneous User Rates:** Until now, we have focused on the performance of the accumulated β -discounted user rates $\mathbf{R}^{(\beta)}[t]$ as the metric of interest. However, $\mathbf{R}^{(\beta)}[t]$ is only a virtual metric to measure the delay sensitivity of the users, whereas $\mathbf{R}[t]$ is the actual amount of service received. Accordingly, it is important to understand the D-DRUM performance with respect to the statistics of individual user rates $\{\mathbf{R}[t]\}_t$. Fig. 6 illustrates the empirical CDF and pdf of instantaneous rates $R_1[t]$ for user-1 achieved under D-DRUM for varying $\beta \in \{0, 0.5, 0.99\}$ values and $\alpha = 1$.

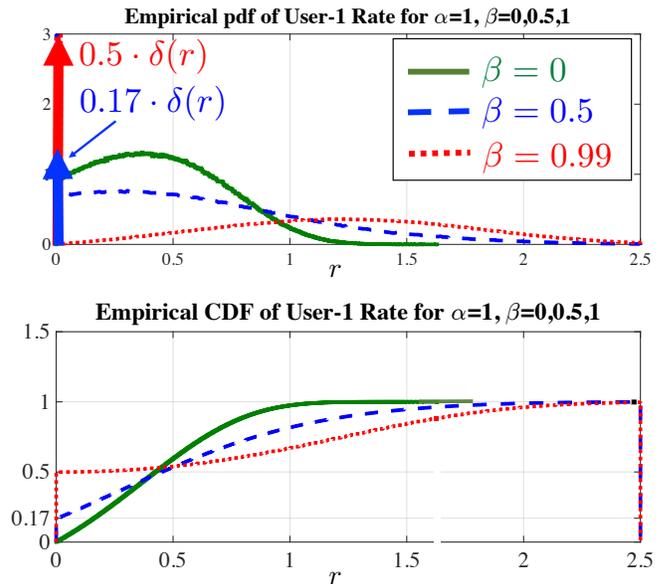


Fig. 6. Distribution (pdf and CDF) of user-1 rates under the Dynamic DRUM algorithm for $\alpha = 1$ and β varying over $\{0, 0.5, 0.99\}$.

When $\beta = 0$, not surprisingly, the distribution of $\{R_1[t]\}_t$ resembles a Rayleigh distribution since the scheduler aims for

⁵See [16] for further results including scenarios with more than two users.

instant-fairness. As β increases towards a mid-level of 0.5, we see that the allocation to the user spreads to a wider range for positive values and a strictly positive probability emerges where the user receives 0 rate. This trend becomes even stronger as β approaches 1, whereby the user receives 0 rate half of the time and receives high rates at other times. This reveals the cost of achieving high long-term rates in opportunistic scheduling, namely, high variability in the instantaneous rates so that good conditions of users can be exploited. These results show that, by adjusting the discount parameter β , the D-DRUM Algorithm provides the means to optimize the tradeoff between achieving high long-term rates and low instantaneous rate variability.

• **Correlation between Individual User Rates:** Another interesting characteristic-of-interest is the relationship between individual user rates that share the common medium. To that end, Fig. 7 depicts the correlation coefficient between the two user rates under D-DRUM for different α values and $\beta \in [0, 1)$.

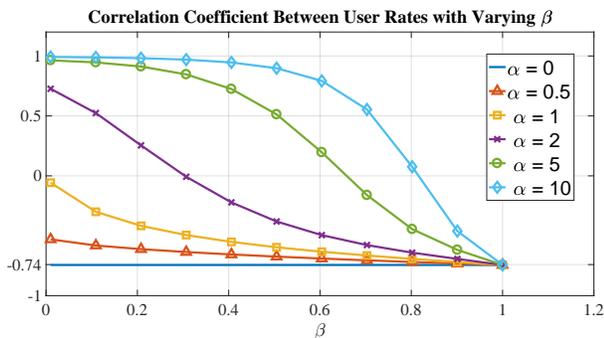


Fig. 7. Correlation coefficient for different α values over β .

For each α , this correlation decreases as β increases, because D-DRUM exploits the time-varying channel capacities more aggressively to achieve higher discounted-rate levels. Accordingly, it increasingly favours the user with high achievable rates while the other tends to receive lower rates, which causes a decrease in the correlation. Also, as α increases, i.e., the fairness requirement gets stronger, the starting correlation level gets higher, reaching 1 in the limit. This is because higher α values enforce increasingly equal instantaneous allocations (under symmetric channel conditions) when β is small. Yet, in all cases, as β increases the averaging affect kicks in (albeit at different rates) to provide sufficient flexibility for D-DRUM to take advantage of the channel variations.

VI. CONCLUSIONS

We introduced a new optimization-based framework for measuring and accommodating the sensitivities of users to time-variations in their received service in a shared wireless fading communication environment. Our framework utilizes discounting of received service rates over time in order to incorporate such sensitivities into a utility maximization formulation. This, so-called Discounted-Rate-Utility-Maximization (DRUM), formulation enables a unified treatment of varying degrees of sensitivities to service rate fluctuations between the extremes of *instantly-fair* (i.e., highly delay-sensitive) and *long-term-fair* (i.e., delay-insensitive) allocations.

Within this new framework, we first characterized the optimal instantly-fair and long-term-fair allocations for the general class of ω -weighted α -fair utility functions. Then, we developed a novel, low complexity Dynamic-DRUM (D-DRUM) Algorithm for the general solution of DRUM for any discount parameter β . We proved that D-DRUM achieves the optimal solutions of the instantly-fair and long-term-fair allocations as β approaches its extremes. We illustrated through extensive numerical investigations that D-DRUM smoothly and stably achieves rate allocations between the instant and long-term optimal extremes by varying β . These investigations also shed light on the temporal dynamics of D-DRUM allocations for individual users and the correlation between the rates of different users.

Based on these results and investigations, we believe that the DRUM framework and the D-DRUM Algorithm show great promise in providing a well-founded means of modeling and managing user sensitivities to short-term fluctuations in a stochastic resource-sharing environment.

REFERENCES

- [1] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, no. 3, March 1998.
- [2] S. H. Low and D. E. Lapsley, "Optimization flow control, I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, pp. 861–875, December 1999.
- [3] R. Srikant, *The Mathematics of Internet Congestion Control*. Boston, MA: Birkhäuser, 2004.
- [4] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. on Networking*, vol. 8, pp. 556–567, Oct. 2000.
- [5] X. Liu, E. Chong, and N. Shroff, "Opportunistic transmission scheduling with resource-sharing constraints in wireless networks," *IEEE JSAC*, vol. 19, no. 10, pp. 2053–2064, October 2001.
- [6] M. Neely, E. Modiano, and C. Li, "Fairness and optimal stochastic control for heterogeneous networks," in *Proc. of IEEE Infocom*, March 2005.
- [7] X. Lin and N. Shroff, "The impact of imperfect scheduling on cross-layer rate control in multihop wireless networks," in *Proc. IEEE Infocom*, Miami, FL, March 2005.
- [8] A. Eryilmaz and R. Srikant, "Fair resource allocation in wireless networks using queue-length based scheduling and congestion control," in *Proceedings of IEEE Infocom*, Miami, FL, 2005.
- [9] A. L. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Systems*, vol. 50, no. 4, pp. 401–457, August 2005.
- [10] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 36, pp. 1936–1948, December 1992.
- [11] —, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. on Inf. Theory*, vol. 39, pp. 466–478, March 1993.
- [12] D. P. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1439–1451, Aug 2006.
- [13] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE JSAC, special issue on Nonlinear Optimization of Communication Systems*, vol. 14, no. 8, August 2006.
- [14] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1–211, 2010.
- [15] A. ParandehGheibi, A. Eryilmaz, A. Ozdaglar, and M. Medard, "Resource allocation for multiple-access channels—an approximate projection approach," *IEEE Trans. on Inf. Theory*, vol. 56, pp. 4417–4437, 2010.
- [16] A. Eryilmaz and I. Koprulu, "Technical report - DRUM: A Framework for Delay-Sensitive Fair Resource Allocation," <https://dl.dropboxusercontent.com/u/37651713/DRUMreport.pdf>.
- [17] P. Viswanath, D. N. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1277–1294, June 2002.