

Selective Fair Scheduling over Fading Channels

Apostolos Destounis*, Georgios S. Paschos*, David Gesbert†

* Mathematical and Algorithmic Sciences Lab, France Research Center, Huawei Technologies Co. Ltd.

† EURECOM

emails: * {apostolos.destounis, georgios.paschos}@huawei.com, † gesbert@eurecom.com

Abstract—Imposing fairness in resource allocation incurs a loss of system throughput, known as the Price of Fairness (PoF). In wireless scheduling, PoF increases when serving users with very poor channel quality because the scheduler wastes resources trying to be fair. This paper proposes a novel resource allocation framework to rigorously address this issue. We introduce *selective fairness*: being fair only to selected users, and improving PoF by momentarily blocking the rest. We study the associated admission control problem of finding the user selection that minimizes PoF subject to selective fairness, and show that this combinatorial problem can be solved efficiently if the feasibility set satisfies a condition; in our model it suffices that the wireless channels are stochastically dominated. Using selective fairness, we formulate the PoF minimization subject to an SLA, which ensures that an *ergodic subscriber* is served frequently enough. In this context, we propose an online policy that combines the *Drift-Plus-Penalty* technique with *Gradient-Based Scheduling* experts, and we prove it achieves the optimal PoF . Simulations show that our intelligent blocking outperforms by 40% in throughput the baseline approach which satisfies the SLA by blocking low-SNR users without considering the overall PoF minimization.

I. INTRODUCTION

Throughput efficiency and fairness is a well-explored fundamental tradeoff in wireless communications [1]. The status quo is to use *opportunistic schedulers* to exploit the fading peaks and strike a balance between throughput and fairness [1], [2]. In this paper, we examine the less investigated Price of Fairness (PoF) [3], which captures the negative impact of fairness on system throughput. PoF depends on channel quality differences; for example, PoF will increase steeply when a base station attempts to serve users with *unreasonably* poor channel quality, for instance users trying to access spectrum from strongly shadowed areas (e.g. building basements, tunnels), remote areas (cell-edge users), or operating with ill-functioning or sub-standard RF equipments. Our simulations show a 10%-40% throughput degradation when adding 1-10 users with SNR 20dB less. Since service quality is anyway low for such users, the baseline approach is to block users when their SNR drops below a threshold. However, this approach is inaccurate since it disregards the channel condition of the other users. *We propose a novel framework for user blocking, where the above threshold is adjusted over time to optimize the PoF subject to fairness and admission probability. For the same admission probability, our scheme yields a 40% total throughput gain over the baseline, unraveling significant room for optimization when blocking cell-edge users.*

We start with a K -user scheduling problem over a wireless fading channel. The service must be fair, but our key idea is that it is allowed temporarily to exclude some users from service. In this context, we introduce a novel fairness objective

called *selective fairness*: a subset of users $S \subseteq \{1, \dots, K\}$ is fairly treated, while the remaining users S^c receive no service. We consider the minimization of PoF (equivalent to system throughput maximization) subject to selective fairness for *at least* s_{\min} users. Solving this problem is non-trivial because the throughput contribution of a user to a fairness-constrained system is very complicated. Mathematically, the problem is of combinatorial nature, potentially involving the solution of $O(2^K)$ large convex programs. We show, however, that if the system satisfies the *subspace monotonicity* property, then the problem can be solved efficiently. We further prove that the subspace monotonicity property is satisfied when the fading channels are *stochastically ordered*, a practical case of interest. Then we propose an online policy, referred to as *Selective-GBS*, which is based on $O(K)$ number of experts, i.e., online simulated policies that provide insight for good scheduling decisions. The throughput vector obtained by Selective-GBS is shown to converge to the optimal solution of the above PoF minimization.

The initial theoretical framework assumes we know how many users to block, which is impractical. It is more reasonable to regulate the blocking according to a probabilistic service guarantee over multiple scheduling problems. The system will block users when their channel quality happens to be very poor, but also ensure that they are served most of the times they attempt to access the service. In the second part of the paper we extend selective fairness to a stochastic setting, where the blocking is controlled by a virtual queue evolving across scheduling problems. Combining the queue with Selective-GBS, we design an online policy, referred to as *Online Selective Fair (OSF)* scheduler, that maximizes system performance while satisfying the probabilistic guarantee and being α -fair to selected users. Contrary to the fixed-threshold baseline, OSF excels at identifying the exact impact of fairness on system throughput, and it exploits the available blocking “budget” to dynamically decide which users will be treated in a fair manner.

A. Related work

The concept of opportunistic scheduling dates back to 1995 [4]. The *Gradient-Based Scheduler* (GBS) was proposed and analyzed early in the 2000s, cf. [5]–[7]. It has been shown to provide a stochastic approximation of the optimal solution of the Network Utility Maximization problem [1], [5], [6], while it can also provide good short-term fairness performance by using discounting factors when averaging [2], [8]. For these reasons, and also for its great simplicity, GBS has become the

de facto scheduling policy in 3G base stations [9]. To capture frequency-time resource blocks in LTE systems, GBS was later extended to a multichannel version [7], [10], [11], keeping the original properties. Prior work has shown how to extend GBS to handle systems with many antennas, which is necessary for the 4G and emerging 5G wireless networks [12]. As of today GBS remains the prominent practical scheduler, therefore we restrict our approach to be backward compatible with GBS.

It is anecdotally known that most operators tune the GBS to achieve *proportional fairness*, a tradeoff between maximizing total throughput and providing equal throughput shares to all users [1], [13]. When some users have much lower channel quality, maintaining proportional fairness results in a reduction of system throughput because the base station is forced to assign a great number of resources to them with small throughput return. Prior work has analyzed the phenomenon of *PoF* in general resource allocation problems [3], [14]. In this paper, we are interested to judiciously exclude some users from service in order to minimize *PoF*. To our knowledge, there exist no prior work studying the optimization of *PoF*. We find that in systems with 5% of users having low channel quality, *PoF* optimization can improve total throughput by 40% over the existing state of art which simply blocks low-SNR users below a threshold. This gain is attributed to the fact that the optimal set of users to block at each realization varies from case to case, and it can not be determined by a simple predefined SNR threshold for blocking.

Our work is related to the literature of *admission control in stochastic networks* and *call admission control in cellular networks*, however there are important differences. Works in stochastic networks focus mainly on admitting fractions of elastic traffic cf. [15], [16], while in our case we admit a *number* of users. Regarding the combinatorial call admission control for cellular networks, the most relevant work to ours is [17], where the authors discuss admission control and blocking rates with opportunistic scheduling and evaluate the performance of simple mechanisms. In this paper, we derive a low-complexity policy that explicitly maximizes the cell spectral efficiency subject to fairness and a blocking constraint. More broadly, a differentiating factor of our work from existing literature is the consideration of special fairness constraints.

II. SYSTEM MODEL

A. Wireless downlink

We consider a *wireless downlink* with one transmitting base station and K receiving users. Time is slotted $t = 1, 2, \dots$, and at each time slot, the base station can transmit data to *one of the users* with the ultimate goal to optimize the time-average data transmissions. An extension to simultaneous service of multiple users is possible via [7], [10], [11], but we avoid it in the interest of presentation clarity.

If user k is scheduled at slot t , data is sent to this user at a transmission rate $R_k(t) \in \mathcal{R}$, where $\mathcal{R} = \{r_1, r_2, \dots, r_L\}$ is a finite set of possible transmission rates. The vector $\mathbf{R}(t) \in \mathcal{R}^K$ is random, i.i.d. over time, and its randomness is attributed to the wireless channel fading, thus it is independent of the past choices of the base station. The realization of $\mathbf{R}(t)$ is provided to the base station just before the scheduling decision is made,




Fig. 1: Illustration of feasible set \mathcal{X} for a 2-user downlink.

as it is customary in contemporary systems.

Let $I_k^\pi(t) \in \{0, 1\}$ denote the scheduling decision at time t regarding user k under scheduling policy π , where $I_k^\pi(t)$ is 1 if user k is scheduled, and 0 otherwise. A policy that schedules only one user at each slot, i.e., satisfies $\sum_k I_k^\pi(t) \leq 1, \forall t$, is called *feasible*, and we denote with Π the set of all feasible policies. In our model, the base station always has available data for each user,¹ therefore the instantaneous rate of data received by user k during slot t is $\mu_k^\pi(t) = R_k(t)$ if $I_k^\pi(t) = 1$ and zero otherwise, the accumulated user throughput at t

$$\bar{x}_k^\pi(t) = \frac{\sum_{\tau=1}^t \mu_k^\pi(\tau)}{t},$$

and the *user k throughput* is $x_k = \liminf_{t \rightarrow \infty} \bar{x}_k^\pi(t)$. The vector of user throughputs, denoted with \mathbf{x} , is our key performance metric.

Definition 1 (Feasible throughputs). The set of feasible throughputs \mathcal{X} is the set of all throughput vectors \mathbf{x} that can be achieved by any policy in Π .

Let $p_r = \mathbb{P}(\mathbf{R}(t) = \mathbf{r}), \mathbf{r} \in \mathcal{R}^K$ denote the probability distribution of channel rate vectors. By considering all slots with $\mathbf{R}(t) = \mathbf{r}$ and scheduling user k with probability ϕ_{kr} , taking liminf leads to the *convex* set \mathcal{X} :

$$\mathcal{X} = \left\{ \mathbf{x} \geq \mathbf{0} \left| \begin{array}{l} x_k = \sum_{\mathbf{r} \in \mathcal{R}^K} \phi_{kr} p_r r_k \\ 0 \leq \phi_{kr} \leq 1 \\ \sum_k \phi_{kr} = 1, \forall \mathbf{r} \in \mathcal{R}^K \end{array} \right. \right\}. \quad (1)$$

Theorem 1 (Feasible throughputs [18]). The set of feasible throughputs \mathcal{X} is given by set (1).

B. Efficiency and fairness

When allocating resource to multiple users, the goal is to drive the system to a throughput $\mathbf{x} \in \mathcal{X}$ that balances two conflicting objectives (i) high efficiency, and (ii) fairness. Consider the 2-user example of figure 1, where \mathcal{X} is the shown gray area and the marginal user throughputs satisfy $X_1^{\max} > X_2^{\max}$. Point A corresponds to the maximum sum throughput—it is the point in \mathcal{X} that maximizes $x_1 + x_2$. An arising issue with point A however, is that user 2 receives zero throughput, which is unfair. The fairest point is C, which ensures that the users receive the maximum possible equal throughputs. In this case however, the total system throughput is significantly reduced. In practice, engineers desire a tradeoff between the two extremes, A and C. Point B, known as *proportional fairness* provides such a tradeoff. Next we formalize the fairness notions of interest.

Definition 2 (Fairness objectives). Let \mathcal{X} be a convex set of feasible throughputs.

¹If we replace GBS with a *max-weight*-type policy, it is possible to generalize our work to stochastic arrivals using the framework in [18].




Fig. 2: System throughput sensitivity ($1 - PoF$) of alpha-fairness to users with poor channel quality.

- A throughput vector $\mathbf{x} \in \mathcal{X}$ is called *max-sum-throughput* if for any vector $\mathbf{y} \in \mathcal{X}$ it holds:

$$\sum_{k=1}^K x_k \geq \sum_{k=1}^K y_k.$$

- A throughput vector $\mathbf{x} \in \mathcal{X}$ is called *max-min fair* if for any vector $\mathbf{y} \in \mathcal{X}$ it holds:

$$y_i > x_i \Rightarrow \exists j : y_j < x_j \leq x_i.$$

- A throughput vector $\mathbf{x} \in \mathcal{X}$ is called *proportionally fair* if for any vector $\mathbf{y} \in \mathcal{X}$ it holds:

$$\sum_k \frac{y_k - x_k}{x_k} \leq 0.$$

Since \mathcal{X} in (1) is a closed set, a max-sum-throughput vector always exists, but it may not be unique. Since \mathcal{X} is convex, it contains a unique max-min fair vector [19], and because the proportionally fair vector is the optimal solution to maximizing the sum of logarithms (i.e. a strictly convex function) over \mathcal{X} , it exists and it is unique.

A connection between fairness and convex optimization is rigorously established by the problem of *Network Utility Maximization* (NUM) [16]:

$$\max_{\mathbf{x} \in \mathcal{X}} \sum_{k=1}^K g_\alpha(x_k), \quad (2)$$

where we have used the α -fair function

$$g_\alpha(x) = \begin{cases} \frac{x^{1-\alpha}}{1-\alpha}, & \alpha \in [0, 1) \cup (1, \infty) \\ \log x, & \alpha = 1. \end{cases}$$

Problem (2) is useful since by tuning the value of α we obtain different fair vectors as solutions to the optimization problem:

- 1) choosing $\alpha = 0$ yields max-sum-throughput,
- 2) choosing $\alpha = 1$ yields proportional fairness [13],
- 3) choosing $\alpha \rightarrow \infty$ yields max-min fairness [20].

Consider the *Gradient-Based Scheduling* (GBS) policy: schedule the user that maximizes $R_k(t)g'_\alpha(\bar{x}_k^{GBS}(t))$, where $g'_\alpha(x) = x^{-\alpha}$, and recall that $R_k(t)$ is the instantaneous transmission rate of user k and $\bar{x}_k^{GBS}(t)$ is the accumulated throughput of user k ; clearly $GBS \in \Pi$. Let $\mathbf{x}^*(\alpha)$ be a solution of (2), prior work [5], [6] has shown that $\bar{x}_k^{GBS}(t) \xrightarrow{\text{a.s.}} x_k^*(\alpha)$, $\forall k$. Hence, we can use GBS and tune the value of α to operate the system at any desirable point. There is anecdotal evidence that 3G and LTE base stations use GBS with $\alpha \approx 1$.

III. OPTIMIZING PRICE OF FAIRNESS

Since $\mathbf{x}^*(\alpha)$ is the solution of (2) for some α , it follows that $\mathbf{x}^*(0)$ denotes a max-sum-throughput vector. We define

the price of fairness similar to [3]:

$$PoF \triangleq \frac{\sum_k x_k^*(0) - \sum_k x_k^*(\alpha)}{\sum_k x_k^*(0)}, \quad (3)$$

which depicts how expensive it is to offer α -fairness in terms of loss of sum throughput. Note that $1 - PoF$ shows the fraction of maximum total throughput achieved by the α -fair point.

In wireless systems it is common for some users to have very poor signal reception. Such users can have a negative effect in PoF , which we showcase with a simulation example. In a downlink system with Rayleigh fading, we simulate GBS scheduling for $K = 10$ users with mean channel rate 1, while adding progressively 1..10 users with mean channel rate 20dB less. We plot $1 - PoF$ for different values of α in figure 2. Proportional fairness experiences a significant efficiency drop of almost 10% for 1 weak user, and 40% for 10.

Our idea is to economize fairness by excluding weak users from service. However, selecting users is challenging because we do not know *a priori* their *long-term* contributions to the throughput of a fairness-constrained system.

A. Selective fairness

As a first step we introduce a novel fairness metric, called *selective fairness*. We partition the set of users to two sets, $S \cup S^c = \{1, \dots, K\}$. For set S we guarantee α -fairness, while the users in the complementary set S^c are not served at all.² Our goal will be to select the set S carefully in order to decrease the price of fairness.

To concretely define selective fairness we need some technical tools. Define the subspace $\mathcal{X}(S) \subseteq \mathcal{X}$, where all non-selected users S^c must have zero throughputs:

$$\mathcal{X}(S) \triangleq \{\mathbf{x} \in \mathcal{X} \mid x_k = 0, \forall k \in S^c\}.$$

Also, we need an $|S|$ -dimensional representation of the vectors in $\mathcal{X}(S)$. Let \mathbf{e}_i be a K -dimensional column vector of zeros with the exception of element i which is one, e.g. for $K = 3$ we have $\mathbf{e}_2 = (0, 1, 0)^T$. Note that $(\mathbf{e}_i)_{i \in \mathcal{K}}$ is a basis of \mathbb{R}^K . Then define the $K \times |S|$ *dimensionality reduction matrix*

$$\mathbf{D}_S \triangleq (\mathbf{e}_k)_{k \in S}.$$

For the example of 3 users, and $S = \{1, 3\}$, we have

$$\mathbf{D}_{\{1,3\}} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}.$$

If we multiply an element of $\mathcal{X}(S)$ or \mathcal{X} with \mathbf{D}_S , we can remove the dimensions that correspond to S^c . Last, consider the set with the $|S|$ -dimensional representations of $\mathcal{X}(S)$:

$$\Gamma(S) \triangleq \{\mathbf{u} \mid \mathbf{u} = \mathbf{x}\mathbf{D}_S, \mathbf{x} \in \mathcal{X}(S)\}.$$

Definition 3 (Selective fairness). A vector $\mathbf{x} \in \mathcal{X}$ is called (S, α) -selective fair if

- 1) $\mathbf{x} \in \mathcal{X}(S)$,
- 2) and $\mathbf{x}\mathbf{D}_S$ is α -fair in the set $\Gamma(S)$.

The definition posits that the selected users in S will be allocated α -fair throughputs in the subspace $\Gamma(S)$, while the rest users receive zero throughput.

²The concept of selective fairness can be generalized to larger user partitions and different fairness objectives per subset.




Fig. 3: 3-user feasible throughputs and the eight possible $(S, 1)$ -selective fair points.

Fix a subset $S \subseteq \{1, \dots, K\}$, and consider the conditions:

$$x_k = 0, \quad \forall k \in S^c, \quad (4)$$

$$\mathbf{x} \in \arg \max_{\mathbf{u} \in \mathcal{X}} \sum_{k \in S} g_\alpha(u_k). \quad (5)$$

Theorem 2. The vector \mathbf{x} is (S, α) -selective fair if and only if it satisfies (4)-(5).

Proof. Definition 3-(1) is equivalent to (4). We also establish the equivalence of definition 3-(2) to (5):

$$\mathbf{x} \in \arg \max_{\mathbf{u} \in \mathcal{X}} \sum_{k \in S} g_\alpha(u_k) \stackrel{(4) \text{ or } \mathbf{x} \in \mathcal{X}(S)}{\Leftrightarrow}$$

$$\mathbf{x} \in \arg \max_{\mathbf{u} \in \mathcal{X}(S)} \sum_{k \in S} g_\alpha(u_k) \Leftrightarrow$$

$$\mathbf{x} D_S \in \arg \max_{\mathbf{u} \in \Gamma(S)} \sum_{i=1}^{|S|} g_\alpha(u_i),$$

where the last follows by a dimension permutation. \square

An implication of the theorem is that the (S, α) -selective fair point is the limit point of GBS policy if we preclude users in S^c from scheduling.

B. PoF minimization with selective fairness

In figure 2 we saw that when we block users with poor channel quality, the *PoF* decreases. It is therefore natural to ask the question *if we were allowed to block all but s_{\min} users in order to minimize PoF, which users would we block?* Next, we pursue a (S, α) -selective fair point that minimizes *PoF* (and thus maximizes system efficiency) subject to serving at least s_{\min} users in a fair manner.

$$T(S^*) = \max_{S \subseteq \{1, \dots, K\}} \sum_k x_k \quad (6)$$

$$\text{s.t. } \mathbf{x} \in \arg \max_{\mathbf{u} \in \mathcal{X}} \sum_{k \in S} g_\alpha(u_k) \quad (7)$$

$$x_k = 0, \quad k \in S^c \quad (8)$$

$$|S| \geq s_{\min}, \quad (9)$$

where, (7)-(8) ensure that the solution vector is (S, α) -selective fair, the constraint (9) ensures that at least s_{\min} users are served, and the objective (6) aims to achieve the maximum sum throughput denoted as $T(S^*)$ (from (3) this is equivalent to minimizing *PoF*). Problem (6)-(9) is an admission control problem with fairness constraints.

We give a pictorial example of *PoF* minimization. Consider a 3-user wireless downlink with feasible throughputs shown in figure 3. The system operates with proportional fairness




Fig. 4: Subspace comparison for the example of figure 3 and $|S| = 2$.

($\alpha = 1$) and must serve at least $s_{\min} = 2$ users. Figure 3 shows all possible selective fair points, where the points $(\{k\}, 1)$, $k = 1, 2, 3, \emptyset$ (indicated as dots with white interior) are infeasible due to constraint (9). Among the feasible $(S, 1)$ -selective fair points, the optimization selects the point with the maximum total throughput, which in this case is $(\{1, 2\}, 1)$.

Corollary 1. Suppose a system is operated with a policy that solves (6), and denote the total achieved throughput by $T_{\mathcal{K}}$ when serving users \mathcal{K} . If $|\mathcal{K}| \geq s_{\min}$, then:

$$T_{\mathcal{K}} \leq T_{\mathcal{K}'}, \quad \forall \mathcal{K} \subseteq \mathcal{K}'.$$

Hence, the optimization (6) is useful because it ensures that the system performance does not drop even when adding users with low average SNRs.

However *PoF* minimization with selective fairness (6) is in general a combinatorial Mixed Integer Convex Program: to determine the solution we potentially need to inspect exponential to K user subsets, and for each subset solve a difficult NUM problem. The NUM problem is difficult due to the possibly very large number of fading states, and explicit solutions for Rayleigh fading are only known for the case of max-min fairness [21]. Nevertheless, we present next a condition which is sufficient to break the combinatorial structure, and simplifies the solution of *PoF* minimization.

C. Subspace monotonicity

Consider a permutation of user indices $\sigma(\cdot)$ and the arising $|S|$ -dimensional subspace $\Gamma^\sigma(S)$, which is the same as $\Gamma(S)$ but with dimensions permuted by σ . We can now compare the subspaces $\Gamma(S_1), \Gamma(S_2)$ for two user sets with same cardinality $|S_1| = |S_2| = |S|$. We say that the subspace of S_1 dominates that of S_2 if there exists permutation σ_2 such that

$$\Gamma(S_1) \supseteq \Gamma^{\sigma_2}(S_2).$$

Figure 4 showcases a comparison of subspaces with $|S| = 2$ on the example of figure 3.

If a subspace dominates another, then it also contains a better selective fair vector; let $T(S)$ be the total throughput of the (S, α) -fair vector under a feasible S . Suppose there exists permutation σ_2 such that $\Gamma(S_1) \supseteq \Gamma^{\sigma_2}(S_2)$. Then $T(S_1) \geq T(S_2)$. Next we generalize this observation.

Definition 4 (Subspace monotonicity). For feasible throughputs \mathcal{X} with dimension K , we say that the set \mathcal{X} satisfies the subspace monotonicity if there exists an ordering of users $(\sigma(1), \sigma(2), \dots, \sigma(K))$ and permutations σ_S , such that:

$$\Gamma(\{\sigma(1), \dots, \sigma(|S|)\}) \supseteq \Gamma^{\sigma_S}(S), \quad \forall S \subseteq \mathcal{K}.$$

Subspace monotonicity expresses the condition that if the users are ordered according to $\sigma(\cdot)$, the subspace of first

$|S|$ users, denoted by $\Gamma(\{\sigma(1), \dots, \sigma(|S|)\})$, dominates the subspace of any other subset with same cardinality.

Corollary 2. Suppose that the set \mathcal{X} satisfies the subspace monotonicity. Then an optimal solution of (6) is of the form $(x_{\sigma(1)}^*, x_{\sigma(2)}^*, \dots, x_{\sigma(j)}^*, 0, \dots, 0)$, where $s_{\min} \leq j \leq K$.

Therefore, if the subspace monotonicity is satisfied, then problem (6) can be solved by a polynomial number of calls to an oracle that solves the NUM problem. Hence, we are motivated to ask: *when is subspace monotonicity satisfied?*

For the interested user, subspace monotonicity is trivially satisfied for a resource allocation problem constrained on a simplex ($\sum_k x_k \leq 1$), where users can be ordered with marginal utilities. Next, we will show that it also applies to our problem under a specific condition for the channels. Recall that $R_k(t)$ is the user k instantaneous channel rate, which is random independently distributed across users and time. In the remaining of the paper we make the following assumption.

Assumption 1 (Stochastic dominance). The channels are stochastically dominated: there exists a permutation $\sigma(\cdot)$ of user indices, such that $R_{\sigma(1)}(t) \geq_{\text{st}} \dots \geq_{\text{st}} R_{\sigma(K)}(t)$, where $R_{\sigma(i)}(t) \geq_{\text{st}} R_{\sigma(j)}(t)$ means

$$P(R_{\sigma(i)}(t) > x) \geq P(R_{\sigma(j)}(t) > x), \quad \forall x, t.$$

Assumption 1 is mild and holds for many practical cases of interest, such as identically distributed fading channels with different means.

Lemma 1. Consider a K -user wireless downlink where Assumption 1 holds. Then the feasible throughputs \mathcal{X} satisfy the subspace monotonicity.

Missing proofs due to limited space are provided in the report [22].

D. Minimizing PoF with online experts

We propose an efficient online scheme to minimize PoF (6). Building on the subspace monotonicity property, our scheme is shown to achieve optimal performance using a small number of GBS experts. First we define the notion of *Gradient-Based Scheduler (GBS) expert* for set $S \subseteq \mathcal{K}$. GBS(S) uses the real system observations $\mathbf{R}(t)$ to simulate the (S, α) -selective fair performance.

GBS(S) expert:

Initialize: $\bar{x}_k^S(0) = 0, \forall k \in S$.

Iterate: At t , observe $\mathbf{R}(t)$ and do:

- 1) *Scheduling:* Choose user k^* randomly from the set

$$\arg \max_{k \in S} \frac{R_k(t)}{(\bar{x}_k^S(t))^\alpha}$$

- 2) *Throughput update:*

$$\bar{x}_k^S(t+1) = \frac{t}{t+1} \bar{x}_k^S(t) + \frac{1}{t+1} R_k(t) \mathbb{1}_{(k=k^*)}$$

Observe that the GBS(S) expert converges to $T(S)$ and hence ensures α -fairness for the user set S . In particular, GBS(S) for $\alpha = 1$ mimicks the opportunistic scheduler of LTE for $S \subseteq \mathcal{K}$ users. Therefore, the S expert tracks the would-be system evolution if the users in S^c were blocked.

To minimize PoF , we present the Selective-GBS policy, which at t restricts opportunistic scheduling to the subset $S^*(t)$, provided by expert with the highest throughput. We note

that the scheduling decision of Selective-GBS is ultimately made according to $\bar{x}_k^{sel}(t)$, which may be different from what the corresponding expert sees, $\bar{x}_k^{S^*}(t)$.

Selective-GBS:

Input: s_{\min} , permutation of user indices $\sigma(\cdot)$.

Initialize: $\bar{x}_k^{sel}(0) = 0, \forall k = 1, \dots, K$. Consider the stronger user set per cardinality:

$$S_i = \{\sigma(1), \sigma(2), \dots, \sigma(i)\}, \forall i \in \{1, 2, \dots, K\} \quad (10)$$

Initialize all GBS(S) experts for $S = S_{s_{\min}}, \dots, S_K$.

Iterate: At t , observe $\mathbf{R}(t)$ and do:

- 1) *Expert update:* Throughput update $\bar{x}_k^S(t)$ for each of the experts GBS($S_{s_{\min}}$), GBS($S_{s_{\min}+1}$), \dots , GBS(S_K).
- 2) *Expert selection:* Choose the expert with highest total accumulated throughput

$$S^*(t) \in \arg \max_{\{S_i\}_{i \geq s_{\min}}} \sum_{k \in S_i} \bar{x}_k^{S_i}(t)$$

- 3) *Scheduling:* Choose user k^* randomly from the set

$$\arg \max_{k \in S^*(t)} \frac{R_k(t)}{(\bar{x}_k^{sel}(t))^\alpha} \quad (11)$$

- 4) *Throughput update:*

$$\bar{x}_k^{sel}(t+1) = \frac{t}{t+1} \bar{x}_k^{sel}(t) + \frac{1}{t+1} R_k(t) \mathbb{1}_{(k=k^*)}$$

Theorem 3. Assume stochastically dominated channels and permutation $\sigma(\cdot)$ such that corollary 2 holds. Then the Selective-GBS converges almost surely to $T(S^*)$:

$$\mathbb{P} \left\{ \lim_{t \rightarrow \infty} \sum_k \bar{x}_k^{sel}(t) = T(S^*) \right\} = 1.$$

The Selective-GBS is an online policy that does not require knowledge of fading statistics and can adapt to system changes. Additionally, it requires only $O(K)$ number of GBS(S) experts, and hence a) selecting the best expert at each slot, and b) storing expert data in memory, is efficient.

As input, Selective-GBS requires the minimum number of accepted users s_{\min} and the permutation $\sigma(\cdot)$ that orders the users in decreasing average SNR. The latter can be kept updated using frequent measurements. Determining s_{\min} is however complicated in practice; we should avoid blocking certain subscribers forever. The next section proposes our main result that deals with this issue: a novel stochastic framework which allows to *determine the optimal set of blocked users at each scheduling realization while providing blocking frequency guarantees.*

IV. STOCHASTIC SELECTIVE FAIRNESS

In wireless systems the actual number of users and their average channel quality varies from base station to base station and from time to time, hence it is impossible to operate selective fairness with a predetermined s_{\min} . Instead, we require a scheme that can adapt the number of selected users in an online fashion in order to a) improve system efficiency, and b) ensure Quality of Service.

A. Novel SLA

We follow a stochastic approach; we propose a novel Service Level Agreement (SLA), such that an arbitrary user

is guaranteed to be selected with probability $1 - \epsilon$, where ϵ is tunable. As opposed to the snapshot scheduling of section III, we consider a sequence of scheduling problem instances z_1, z_2, \dots , where each z_n is a realization of the random variable \mathbf{Z} describing the spatial distribution of the users, that is the set of active users $\mathcal{K}(\mathbf{Z})$ and their average SNRs due to slow fading, denoted by \overline{SNR}_k for user $k \in \mathcal{K}(\mathbf{Z})$. Then let $y_k(\mathbf{Z}) = 1$ indicate the event that user k is selected at a random realization \mathbf{Z} , we require the probabilistic constraint:

$$\mathbb{P}_{\mathbf{Z}} (y_k(\mathbf{Z}) = 1 | k \in \mathbf{Z}) \geq 1 - \epsilon, \forall k.$$

The users are considered statistically equivalent, hence the constraint intuitively captures the behavior of an ergodic subscriber, i.e., one who buys a service with the associated SLA, and then uses the service over a long time horizon each time from possibly different locations. In this context, the constraint implies that such a subscriber will be rarely blocked.




Fig. 5: Scheduling over multiple realizations.

B. Problem formulation

In this section we formalize a *PoF* minimization over multiple scheduling realizations. Let $\mathcal{T}(n) = \{t_n, \dots, t_{n+1}-1\}$ be the set of time slots spanned by the n -th realization of the spatial distribution, \mathbf{Z}_n . Then, given the realization of the spatial distribution in $\mathcal{T}(n)$, z_n , we solve a selective-fair scheduling problem where we need to decide which users to admit in the system for these slots; let this decision be denoted by $\mathbf{y}(z_n)$. We need the following assumption for time scale separation:

Assumption 2 (Time scale separation). The changes in spatial realizations occur in a slower time scale than the fast fading. In particular, the duration of a realization $\mathcal{T}(n)$ is long enough for the empirical throughput averages to converge to the α -fair rates, that is

$$\left| \frac{1}{|\mathcal{T}(n)|} \sum_{t=t_n}^{t_{n+1}-1} \mu_k(t) - T_k^*(z_n, \mathbf{y}(z_n)) \right| \leq \delta,$$

for some small $\delta > 0$, where $T_k^*(z_n, \mathbf{y}(z_n))$ is the optimal selective α -fair rate vector for spatial realization z_n , given the admission decision $\mathbf{y}(z_n)$.

The above assumption implies that user arrivals/departures and slow fading fluctuations occur slower than the convergence of the scheduler, which in practice takes a few seconds.

Let the admission policy $\mathbf{y}(\cdot)$ be a function from spatial realizations to deciding if a user will be admitted or not. Our general problem then is to find the admission policy to solve the following stochastic optimization:

$$\max_{\mathbf{y}} \mathbb{E}_{\mathbf{Z}} \left\{ \sum_k T_k^*(\mathbf{Z}, \mathbf{y}(\mathbf{Z})) \right\} \quad (12)$$

$$\text{s.t. } \mathbb{P}_{\mathbf{Z}} (y_k(\mathbf{Z}) = 1 | k \in \mathbf{Z}) \geq 1 - \epsilon, \forall k, \quad (13)$$

where the constraint (13) ensures the satisfaction of the SLA of each subscriber, and the maximal value of (12) represents the best way the system can use the available admission budget to minimize the *PoF*. Note that this minimization is very

complicated since it jointly considers i) the actual number of active users at each realization, ii) the random user locations and corresponding average SNRs, and iii) the cost of fairness.

We proceed to reformulate the problem as an optimization of time-averages. First, it will be useful to transform the per-user SLA constraints into an equivalent constraint for the total number of admitted users:

Lemma 2. Assume the spatial distribution of all users is the same. Then, (13) is equivalent to

$$\mathbb{E}_{\mathbf{Z}} \left\{ \sum_{k \in \mathcal{K}(\mathbf{Z})} y_k(\mathbf{Z}) \right\} \geq (1 - \epsilon) \mathbb{E}_{\mathbf{Z}} \{ |\mathcal{K}(\mathbf{Z})| \}$$

Then, we use the *law of large numbers* to express the objective function as a time average over scheduling realizations:

$$\mathbb{E}_{\mathbf{Z}} \left\{ \sum_k T_k^*(\mathbf{Z}, \mathbf{y}(\mathbf{Z})) \right\} = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left\{ \sum_{n=0}^{N-1} \sum_{k \in \mathbf{z}_n} T_k^*(z_n, \mathbf{y}(z_n)) \right\},$$

where the expectation is taken with respect to possible random control (if the control is deterministic it can be eliminated). Finally, constraint (13) can be rewritten as a time average of indicator functions:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left\{ \sum_{n=0}^{N-1} \mathbb{1}_{(k \in \mathbf{z}_n)} y_k^n(\mathbf{z}_n) \right\} \geq (1 - \epsilon) \mathbb{E}_{\mathbf{Z}} \{ |\mathcal{K}(\mathbf{Z})| \}. \quad (14)$$

The problem to solve now has become:

$$\max_{\{\mathbf{y}^n\}_{n=0}^{N-1}} \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left\{ \sum_{n=0}^{N-1} \sum_{k \in \mathbf{z}_n} T_k^*(z_n, \mathbf{y}^n(z_n)) \right\}. \quad (15)$$

V. ONLINE SELECTIVE FAIR POLICIES

The form of (15) motivates a stochastic optimization approach [18], where we can use a virtual queue to ensure the satisfaction of the time-average constraint.

Define a *virtual queue* $Q(n)$, which evolves at the same time scale as the spatial process as follows:

$$Q(n+1) = [Q(n) + A(n) - D(n)]^+, \quad (16)$$

where $A(n)$ is defined as

$$A(n) = \begin{cases} |\mathcal{K}(\mathbf{z}_n)|, & \text{w.p. } 1 - \epsilon \\ 0, & \text{w.p. } \epsilon, \end{cases} \quad (17)$$

and $D(n) = D(\mathbf{z}_n, \mathbf{y}^n(\mathbf{z}_n)) = \sum_{k \in \mathcal{K}(\mathbf{z}_n)} y_k^n(\mathbf{z}_n)$ is the number of users admitted at the n -th spatial realization. Note that $A(n)$ is a random variable with mean $(1 - \epsilon) \mathbb{E}_{\mathbf{Z}} \{ |\mathcal{K}(\mathbf{Z})| \}$. The idea then is that, if the queue is stable, the mean of its *service* $D(n)$ will be larger than the one of its *arrivals* $A(n)$, and the equivalent SLA constraint (14) will be met. Therefore, our policy will strive to stabilize $Q(t)$. We observe that the virtual queue can be seen as a counter to track/learn the Lagrange multiplier corresponding to the SLA constraint.

A. Known selective fair throughputs

We first deal with the case where given a realization z_n , the corresponding selective-fair sum throughputs $\sum_k T_k^*(z_n, \mathbf{y})$ are known for each \mathbf{y} .

Drift Plus Penalty (DPP):

Initialize: Fix parameter $V > 0$.

Iterate over scheduling realizations: At the beginning of t_n :

1) *User selection:*

$$\mathbf{y}^n = \arg \max_{\mathbf{y}} \left[\sum_k T_k^*(\mathbf{Z}_n, \mathbf{y}) + \frac{Q(n)}{V} \sum_{k \in \mathcal{K}(\mathbf{Z}_n)} y_k \right]. \quad (18)$$

2) *Virtual queue update:* Set $D(n) = \sum_k y_k^n$, draw a random variable $A(n)$ and update the queue as in (16).

Theorem 4. The DPP satisfies the SLAs, and yields sum throughput within $O(1/V)$ of the maximum in (15).

Proof: Let $\mathbf{y}^*(z_n)$, $n = 1, 2, \dots$ denote the optimal solution of (12)-(13), i.e. a (possibly randomized) function from a scheduling realization to a user admission decision; existence of optimal stationary randomized policies is established in [23]. To prove optimality of DPP, we will compare to the performance of $\mathbf{y}^*(z_n)$ with respect to the Lyapunov drift. Indeed, we first define the drift of policy \mathbf{y} :

$$\Delta(\mathbf{y}, \mathbf{Z}, Q) = \sum_{k \in \mathcal{K}(\mathbf{Z})} (VT_k^*(\mathbf{Z}, \mathbf{y}) + Qy_k).$$

Recall that \mathbf{y}^n denotes the decision of DPP at realization n , and observe that it is designed to minimize the quantity $\Delta(\cdot)$. Then, for a given scheduling realization z_n and given value of the virtual queue $Q(n)$, we can show that

$$\begin{aligned} & \mathbb{E}\{Q^2(n+1) - Q^2(n)\} - V\mathbb{E}\left\{\sum_k T_k^*(z_n, \mathbf{y}^n(z_n))\right\} \\ & \leq B + Q(n)\mathbb{E}\{A(z_n)\} - \mathbb{E}\{\Delta(\mathbf{y}^n, z_n, Q(n))\} \\ & \leq B + Q(n)\mathbb{E}\{A(z_n)\} - \mathbb{E}\{\Delta(\mathbf{y}^*, z_n, Q(n))\}, \end{aligned}$$

where B is a constant that depends only the parameters of the system. Additionally, from optimality of \mathbf{y}^* we have

$$\begin{aligned} & \mathbb{E}\{Q^2(n+1) - Q^2(n)\} - V\mathbb{E}\left\{\sum_k T_k^*(z_n, \mathbf{y}^n(z_n))\right\} \\ & \leq B - V\mathbb{E}\left\{\sum_{k \in \mathcal{K}(\mathbf{Z})} T_k^*(z_n, \mathbf{y}^*(z_n))\right\}, \end{aligned}$$

from which we can deduce that (i) $Q(n)$ is (mean rate) stable therefore the SLA constraints hold and (ii) the value of the average throughput less than the optimal by at most B/V , finishing the proof. ■

B. Unknown selective fair throughputs

In practice determining $\sum_k T_k^*(z_n, \mathbf{y})$ for a set of users \mathbf{y} knowing their channel statistics is very challenging, with the exception of proportional and max-min fairness, where approximations exist [21], and the case of max throughput where the solution is to always admit all users. Therefore, we cannot always directly apply the DPP policy. In this section we propose a combination of Selective-GBS (which requires s_{\min}) and DPP (which requires $\sum_k T_k^*(z_n, \mathbf{y})$), to design a policy which progressively learns the best user set and thus does not need the information of neither s_{\min} , nor $\sum_k T_k^*(z_n, \mathbf{y})$. The only information needed is the ordering of the users according to their SNRs, which is found by measurements.

Online Selective Fair (OSF) Scheduler:

Initialize: Fix parameter $V > 0$, virtual queue service $S^*(0) = S$.

Iterate t_n (over scheduling realizations):

- 1) *Update virtual queue:* Generate the random variable $A(n)$, as per (17), set $D(n) = |S^*(t_n)|$, i.e. the cardinality of the subset chosen in the last slot of last realization, and update the queue: $Q(n+1) = [Q(n) - D(n) + A(n)]^+$.
- 2) *User ordering:* Consider current user realization z_n , and $K_n = |\mathcal{K}(z_n)|$. Then permute users with $\sigma(\cdot)$ such that $\overline{SNR}_{\sigma(1)} \geq \overline{SNR}_{\sigma(2)} \geq \dots \geq \overline{SNR}_{\sigma(K_n)}$.
- 3) *Initialize next realization:* Reset $t = 0$, construct sets S_i as per (10), set $\bar{x}_k(0) = 0, \forall k$ and $\bar{x}_k^{S_i}(0) = 0, \forall k \in S_i$ for each expert i , and run the next realization using the following policy:

Iterate t :

- 1) *Expert update:* Execute one step of each expert $\text{GBS}(S_i), i \in \{1, 2, \dots, K_n\}$ using $\mathbf{R}(t)$. Let $\bar{x}^{S_i}(t+1)$ be the vector of accumulated throughputs of expert i .
- 2) *Expert selection:* Choose the expert that maximizes:

$$S^*(t) = \arg \max_{\{S_i\}_i} \left[\sum_{k \in S_i} \bar{x}_k^{S_i}(t+1) + |S_i| \frac{Q(n)}{V} \right]$$

- 3) *Scheduling:* Choose user $k^*(t)$ randomly from the set:

$$\arg \max_{k \in S^*(t)} \frac{R_k(t)}{(\bar{x}_k(t))^\alpha}$$

- 4) *Throughput update:*

$$\bar{x}_k(t+1) = \frac{t}{t+1} \bar{x}_k(t) + \frac{1}{t+1} R_k(t) \mathbb{1}_{(k=k^*(t))}$$

Corollary 3 (Optimality of OSF). Let assumptions of i) time-scale separation, ii) statistically identical users, and iii) stochastically dominated channels, hold. Theorem 3 implies that at each scheduling realization the experts converge to the corresponding selective fair throughputs, hence OSF chooses the action that maximizes (18), and by Theorem 4 *OSF satisfies the SLA constraints and yields sum throughput within $O(1/V)$ of the maximum in (15).*

A few remarks about OSF: searching for the best expert and storing expert data is $O(K)$, hence OSF has low-complexity. Also, to determine $D(n)$, we estimated the number of users admitted in last realization using the $|S|$ of the best expert at the last slot. This is ensured to be accurate when the experts have converged. In our simulations, this method remains accurate even when the convergence is not completely reached.

VI. NUMERICAL RESULTS

In this section we evaluate the performance of OSF against policies that block users with average SNR below a threshold, and against the case where no admission control is done and all users are served in a fair manner. We consider a system with 100 subscribers, each requesting a probabilistic SLA of being served 95% of the time. The system evolves as per figure 5. At each scheduling realization, a subscriber is active with probability 0.1 and is placed uniformly at random in a cell. The parameters are chosen such that a user at




Fig. 6: Simulation results for blocking probability 5%: (a) Plot legends. (b) Empirical admission probabilities for $\alpha = 1$. (c) Empirical average system throughput for $\alpha = 1$. (d) Price of Fairness for different levels of fairness (values of α).

the edge of the cell has an average SNR of -5dB . Each scheduling realization lasts for 3000 time slots, and at each slot the channel realizations are drawn from a Rayleigh fading distribution.

The results are presented in figure 6. Figures 6b-6c show the case $\alpha = 1$; the OSF policy achieves the SLA (6b) and outperforms the best naive threshold admission control that also achieves the SLA by approximately 40% in terms of total throughput (6c). In addition, it outperforms the policy that does no admission control by a wider margin. Figure 6d compares PoF for different values of α . In low values of α , where fairness constraints are loose, all policies have similar performance. As α increases, having a good user admission strategy makes fairness cheaper; the PoF of OSF increases slower than the other two policies. *We attribute the gains of our approach to opportunistic minimization of PoF by appropriate dynamic user blocking over multiple scheduling realizations.*

We also experimented with more stringent values of the SLA. For $\alpha = 1$ the gain in system throughput with respect to the best threshold policy that satisfies the SLA is 10% and 2% for corresponding selection probabilities 99% and 99.9%. Decreasing gain for more stringent SLA is to be expected, since for selection probability 1 all the policies will yield the same PoF —they are all forced to accept all users at all times. However, we mention that the probability of good coverage of LTE is measured in drive tests to be 90-95% [24], which allows a fairly large amount of blocking.

VII. CONCLUSIONS

In this paper, we introduced *selective fairness*, the idea of providing fairness to some users and blocking the others, in order to mitigate the adverse effect of users with very poor channel quality. Extending this concept to the stochastic setting, we derived an online policy that maximizes the system's throughput subject to being selectively fair and satisfying an SLA on the user blocking probability. Our intelligent blocking outperforms by 40% naive approaches that simply block low-SNR users. It is worth noting that our results can be easily extended to multiple SLA classes using one queue for each class, as well as systems with multiple cells using an operator-wide global queue. Analyzing the transient behavior of the Selective-GBS with experts and extending our results to cases where users have different spatial distributions are very interesting topics for further study.

REFERENCES

- [1] H. J. Kushner and P. A. Whiting, "Convergence of proportional-fair sharing algorithms under general conditions," *IEEE Transactions on Wireless Communications*, pp. 1250–1259, 2004.
- [2] P. Viswanath, D. N. C. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, pp. 1277–1294, 2002.
- [3] D. Bertsimas, V. F. Farias, and N. Trichakis, "The price of fairness," *Oper. Res.*, pp. 17–31, 2011.
- [4] R. Knopp and P. A. Humblet, "Information capacity and power control in single-cell multiuser communications," in *IEEE ICC*, 1995.
- [5] R. Agrawal and V. Subramanian, "Optimality of certain channel aware scheduling policies," in *Allerton*, 2002.
- [6] A. L. Stolyar, "On the asymptotic optimality of the gradient scheduling algorithm for multiuser throughput allocation," *Oper. Res.*, pp. 12–25, 2005.
- [7] J. Huang, V. G. Subramanian, R. Agrawal, and R. A. Berry, "Downlink scheduling and resource allocation for OFDM systems," *IEEE Transactions on Wireless Communications*, pp. 288–296, 2009.
- [8] A. Eryilmaz and I. Koprulu, "Discounted-rate utility maximization (DRUM): A framework for delay-sensitive fair resource allocation," in *WiOpt*, 2017.
- [9] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushyana, and S. Viterbi, "CDMA/HDR: a bandwidth efficient high speed wireless data service for nomadic users," *IEEE Communications Magazine*, pp. 70–77, 2000.
- [10] A. Leith, M. S. Alouini, D. I. Kim, X. Shen, and Z. Wu, "Flexible proportional-rate scheduling for OFDMA system," *IEEE Transactions on Mobile Computing*, pp. 1907–1919, 2013.
- [11] F. Capozzi, G. Piro, L. A. Grieco, G. Boggia, and P. Camarda, "Downlink Packet Scheduling in LTE Cellular Networks: Key Design Issues and a Survey," *IEEE Communications Surveys Tutorials*, pp. 678–700, 2013.
- [12] D. Bethanabhotla, G. Caire, and M. J. Neely, "WiFlix: Adaptive video streaming in massive MU-MIMO wireless networks," *IEEE Transactions on Wireless Communications*, pp. 4088–4103, 2016.
- [13] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, pp. 237–252, 1998.
- [14] D. Bertsimas, V. F. Farias, and N. Trichakis, "On the efficiency-fairness trade-off," *Manage. Sci.*, pp. 2234–2250, 2012.
- [15] X. Xiang, C. Lin, X. Chen, and X. Shen, "Toward optimal admission control and resource allocation for LTE-A femtocell uplink," *IEEE Transactions on Vehicular Technology*, pp. 3247–3261, 2015.
- [16] M. J. Neely, E. Modiano, and C. P. Li, "Fairness and optimal stochastic control for heterogeneous networks," *IEEE/ACM Transactions on Networking*, pp. 396–409, 2008.
- [17] T. Bonald and A. Proutière, "Wireless downlink data channels: User performance and cell dimensioning," in *ACM MobiCom*, 2003.
- [18] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Resource allocation and cross-layer control in wireless networks," *Found. Trends Netw.*, pp. 1–144, 2006.
- [19] B. Radunovic and J. Y. L. Boudec, "A unified framework for max-min and min-max fairness with applications," *IEEE/ACM Transactions on Networking*, pp. 1073–1083, 2007.
- [20] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions on Networking*, pp. 556–567, 2000.
- [21] R. Combes, Z. Altman, and E. Altman, "On the use of packet scheduling in self-optimization processes: Application to coverage-capacity optimization," in *WiOpt*, 2010.
- [22] A. Destounis, G. Paschos, and D. Gesbert, "Selective fair scheduling over fading channels," in *arXiv.org*, 2017.
- [23] M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan and Claypool, 2010.
- [24] J. Furuskog, K. Werner, M. Riback, and B. Hagerman, "Field trials of LTE with 44 MIMO," *Ericsson review*, 2010.