

Tie-breaking Can Maximize Fairness without Sacrificing Throughput in D2D-assisted Networks

Vincenzo Mancuso¹, Arash Asadi^{1,2}, and Peter Jacko³

¹IMDEA Networks Institute, ²University Carlos III of Madrid, ³Lancaster University
Email: vincenzo.mancuso@imdea.org, arash.asadi@imdea.org, and p.jacko@lancaster.ac.uk

Abstract—Opportunistic schedulers such as MaxRate and Proportional Fair are known for trading off between throughput and fairness of users in cellular networks. In this paper, we propose a novel solution that integrates opportunistic scheduling design principles and cooperative D2D communication capabilities in order to maximize fairness without sacrificing throughput. Specifically, we develop a mathematical approach and design a smart tie-breaking scheme which maximizes the fairness achieved by the MaxRate scheduler. However, our approach could be applied to improve fairness of any scheduler. In addition, we show that users that cooperatively form D2D clusters benefit from both higher throughput and fairness. Our scheduling scheme is simple to implement, scales linearly with the number of clusters, and is able to double the throughput of Equal Time schedulers and to outperform by 20% or more Proportional Fair schedulers, while providing a user fairness index comparable to or better than Proportional Fair.

Index Terms—Cellular networks; Opportunistic scheduling; D2D communications; Tie-breaking.

I. INTRODUCTION

It has been shown that the cellular throughput can be dramatically improved by using opportunistic schedulers such as MaxRate [1] and Proportional Fair [2]. However, opportunistic schedulers proposed for cellular networks face a trade off between throughput and fairness when it comes to prioritizing the users based on their channel qualities [3]. Hence, with the existing cellular architectures, opportunistic schedulers cannot achieve maximum throughput and fairness at the same time, unless all users experience the very same channel quality [4].

In contrast, in this work, we show how to evolve the cellular architecture by leveraging device-to-device (D2D) communications to achieve *maximum throughput and maximum fairness*. In particular, we explore the possible gain from having D2D-assisted connections within *clusters* of mobile users, as shown in Fig. 1, where each cluster is treated by the base station as a regular mobile user in a cell. We propose to change the normal cellular operation (e.g., in LTE-A) as follows: at each scheduling frame, the scheduled mobile user is responsible for the traffic of its *entire* cluster, i.e., it acts as *cluster head*. The cellular traffic managed by the cluster head is then immediately exchanged within the cluster via D2D communications on a secondary wireless interface (e.g., an 802.11-based protocol like WiFi-Direct or WiGig). Note that the cluster head is, in principle, *opportunistically* different at any frame, thereby achieving maximum throughput.

In turn, fairness is achieved in the following way. The scheduler selects some of the connections for transmission,




Fig. 1. Example scenario: three clusters in a base station, with five, ten and fifteen mobile users, respectively.

however, there are situations in which two or more schedulable connections are in *tie*, i.e., they can be scheduled with the same transmission rate. These ties are usually ignored or broken randomly [5]. In contrast, we show that a smart *tie-breaking* strategy allows to compensate for the channel quality differences randomly experienced by data connections that are active in the cellular network.

So far, nobody has investigated the possibility of enhancing fairness of opportunistic schedulers by utilizing tie-breaking methods. Thus, we are the first to explore tie-breaking schemes for improving the fairness achieved with a MaxRate scheduler, though other opportunistic schedulers could also be enhanced with our approach. We select the MaxRate scheduler since it maximizes system throughput, and we show that improved fairness levels can be achieved without paying any throughput cost, i.e., *we study how to break ties in order to maximize connection fairness while maintaining the maximum cell throughput*. However, our approach could be applied to improve fairness of any scheduler.

We propose a novel solution—and design an algorithm—that integrates *opportunistic scheduling design principles* and *cooperative D2D communication capabilities* to maximize both fairness and throughput in cellular networks. Our main contributions can be summarized as follows:

- 1) we develop a mathematical framework for the evaluation of fairness in an architecture based on cooperative D2D clusters and opportunistic scheduling in cellular networks;
- 2) we formulate a problem and show how to optimize tie-breaking based on the knowledge of average channel qualities;
- 3) we analytically derive *MaxRate-MaxFair*, an optimal tie-breaking scheduler for the case of two connections;

- 4) based on the analysis, we design PIKe, a novel tie-breaking heuristic to provide improved fairness under MaxRate scheduling with multiple connections;
- 5) we evaluate the performance of the proposed solution through an extensive numerical simulation study.

The remainder of this paper is organized as follows. In Section II, we review the literature on works relevant to our proposal. Our system model is presented in Section III. Section IV analytically tackles fairness issues which are due to MaxRate scheduling, and discusses a few schemes to improve the level of fairness in the system without reducing the total throughput. For the simple case of a system with two active data connections, we analytically design *MaxRate-MaxFair*, a scheme that achieves the highest possible fairness and maximum throughput. Inspired by the results of the two-connection case, in Section V, we propose PIKe, a heuristic for solving the generic case of N connections to be scheduled. In Section VI, we use numerical simulations to evaluate the throughput and the fairness that can be achieved with our proposal in a variety of realistic scenarios. PIKe can achieve 50% higher throughput than conventional networks while providing a user fairness index close to one. Finally, we summarize and conclude the paper in Section VII.

II. RELATED WORK

D2D communications. D2D communications include all technologies that allow direct communication among users without involvement of an infrastructure or an access point [6]. D2D communications may be utilized in a variety of scenarios such as cooperative communications, packet forwarding, and relaying. In [6] and [7], the authors explore some applications of D2D communications in cellular networks such as P2P, multiplayer gaming, and multicast transmissions. Doppler *et al.* [8] explore D2D communications establishment and management in LTE-A networks. They show that D2D communications using the very same cellular resources (*inbound relay*) increase the throughput by up to 65%. The authors of [9] propose to introduce *fixed ad hoc relay stations*, operating on licensed or unlicensed bands, to enable P2P communications within a cell or adjacent cells. Their approach reduces the call blocking probability by routing new calls to less congested cells. The authors of [10] propose to use D2D communications by forming clusters among mobile users with single antenna to emulate a MIMO device. Yu *et al.* [11] propose D2D communications in cellular networks for local traffic handling. In that work, D2D transmissions are meant to handle communications among two mobile devices, however users do not help each other to relay traffic to the base station. Also, all transmissions occur over the same interface as cellular communications, and D2D resources are allocated by the base station.

Opportunistic scheduling. The notion of opportunistic scheduling was introduced by Knopp and Humblet in [1] by proposing the *MaxRate* scheduler. MaxRate exploits multiuser diversities in wireless channel by scheduling the user with the highest transmission rate in each frame. The greedy behavior

of MaxRate leads to unfairness among users with heterogeneous channels. Therefore, MaxRate is rarely used, whereas it is commonly referred to as an upper bound for the achievable throughput. Furthermore, MaxRate does not define a strategy to break scheduling ties, and common implementations use uniformly random tie-breaking [5].

Some of the proposed opportunistic schedulers can provide certain notion of fairness at the expense of reducing throughput. For instance, *Proportional Fair* (PF) represents the state-of-the-art opportunistic scheduling with fairness constraints. PF gives priority to users with relative good channel quality and who have received less throughput in the past, i.e., PF uses the r.v. $R_n(t)/\mu_n(t)$ as scheduling priority at time t , where $R_n(t)$ is the rate achievable in the current frame, and $\mu_n(t)$ is the average throughput received by user n and computed via a low pass filter. The PF scheduler, patented by Qualcomm for High Data Rate (HDR) system 1xEV-DO [2], was shown to maximize the aggregate logarithmic throughput, not the cumulative throughput [12]. While the network capacity is fully achievable in case of homogeneous users [13], this is not the case in presence of multi-class flows [4]. Other opportunistic schedulers, e.g., *MaxWeight* [14] and *Exponential rule* [15], take scheduling decisions based on a metric that combines queue sizes and transmission rates. These schedulers do not achieve the maximum possible throughput.

Tie-breaking. Advanced tie-breaking can be used to improve performance in network and computer systems. In [5], Neely proposed to break ties in favor of the user with the longest queue. The authors of [16] proposed a modification of the MaxWeight scheduler, to serve the user with highest product of head-of-line packet's waiting time and actual transmission rate, with ties broken in favor of older packets. In [17], the authors proposed the oldest-first tie-breaking rule in which the scheduler breaks ties based on the life-time of the flows (i.e., packet waiting time in the queue). For flow-level models, in which users randomly arrive and depart, the authors of [4] and [18] proposed prioritization of flows with higher departure probability. Indeed, [19] proved that opportunistic schedulers fail to achieve fluid-optimality unless the tie-breaking prioritizes such flows. The majority of the existing literature overlooks the existence of scheduling ties, whereas the few authors that take ties into consideration either opt for uniformly random tie-breaking [20] or leverage a secondary constraint to break ties [5], [16], [17], [21], [22].

All these works have one common assumption, i.e., ties are expected to happen in rare events and their impact on system performance is marginal. In contrast, this paper focuses on scheduling ties as an important part of the scheduling decision, which allows for achieving high user fairness with aggressive opportunistic scheduling schemes. Moreover, note that the clustering mechanism we adopt in this work to boost throughput also generates frequent tie opportunities (see results in Section VI-A). Therefore, the study of tie-breaking with respect to fairness consequences becomes very relevant. As we remark, to date, there is not work exploiting D2D com-

munications and tie-breaking in opportunistic scheduling to improve fairness while achieving maximum throughput, while the feasibility of opportunistic scheduling of D2D clusters has been recently demonstrated in [23].

III. SYSTEM MODEL

In this section we present our D2D-based opportunistic scheduling system that can be leveraged to maximize both throughput and fairness in a cellular system.

A. Connections

We consider a cellular network with a set \mathcal{N} of N persistent connections of users to a base station, using dedicated wireless channels. In the following, we focus on downlink communications, though the model is applicable also to uplink communications. The base station operates in a synchronous time-slotted way, and its task is to schedule connections for transmission in every *frame* $t = 0, 1, \dots$. We also assume that the base station has a queue for storing packets to be delivered to each connection, and queues are never empty (fully backlogged assumption), so that we can evaluate the behavior of the system in *saturation*, i.e., under the worst scheduling operational conditions.

Transmission channels are *heterogeneous*, i.e., *not* satisfying the i.i.d. assumption common to many works on cellular networks. The channel is characterized by stationary Rayleigh fading. Therefore, the SNR can be described as a random process $C(t)$ with mean γ and Cumulative Distribution Function (CDF) given by $F(z) = 1 - e^{-\frac{z}{\gamma}}$, $z \geq 0$, $\gamma \in \Gamma$, $n \in \mathcal{N}$, where, for the sake of tractability, we have introduced Γ as the discrete set of average values for the SNR of transmissions.

We assume that the information available at the base station corresponds to the steady-state distribution of SNR. Note that, for practical systems, in which measured channel conditions form a discrete set, existing patents [24], [25] propose to keep track of historical observations of SNR values to provide an estimate of the steady-state distribution of SNR.

We assume that the selection of Modulation and Coding Scheme (MCS) is perfect (i.e., transmissions are affected by negligible error rate), so that we ignore retransmissions. Each possible MCS $k \in \{1, \dots, K\}$ is selected by connection n with probability $p_{n,k}$, and achieves transmission rate r_k , which only depends on the instantaneous SNR value. Eventually, we consider a system with no power control, which is realistic for downlink transmission schemes in LTE, although this assumption can be easily relaxed.

B. Scheduling of Clusters of Users

In the previous discussion we had implicitly associated every connection with a single user. However, mobile users may form clusters using cooperative D2D communications, in which only one of the users, namely the *cluster head*, connects in a given frame to the base station and relays traffic for the other users. A cluster is formally defined as follows.

Definition 1. (Cluster) *A cluster is a group of mobile users that can communicate with each other using a data rate more*

advantageous than with the cellular base station. Only one cluster member, namely the cluster head, is allowed to receive data from the base station. The downlink traffic received at the cluster head can belong to any of the cluster members, and it is distributed in the cluster using D2D techniques.

The scheduling algorithm is MaxRate, i.e., the cluster that contains the user with the highest MCS is scheduled. We propose to operate clusters in opportunistic way: (i) the cluster head can change on a per-frame basis, as it is opportunistically selected as the cluster member with the highest current MCS; (ii) an entire cluster is scheduled as an individual user whose MCS is the highest among members; (iii) the cluster head relays the downlink packets to the final destination on a secondary wireless interface, using D2D for intra-cluster communications. Therefore, in this work, a *connection* n can be a user or, more in general, a cluster composed by $m_n \geq 1$ mobiles, and its instantaneous SNR is the *highest* SNR among the mobile users composing the cluster. In particular, the probability $p_{n,k}$ that a scheduled connection n receives data encoded according to the k -th MCS can be computed based on the cluster members' SNR CDF and the MCS thresholds used in the cellular system (see Appendix A).

As detailed in Section VI-A, scheduling clusters instead of users not only brings advantages in terms of system throughput, but also in terms of fairness. However, we will mainly focus on inter-cluster fairness, since the actual per-user fairness depends on the way resources are shared within a cluster. Accordingly, when we refer to throughput and fairness, we assume that cluster resources are divided equally between cluster members. Note that throughput unfairness due to heterogeneous channel qualities within the same cluster is smoothed by the adopted cooperative D2D communications. However, the particular mechanism to manage intra-cluster fairness is left out of the scope of this manuscript. Here, we rather focus on studying possible gains in throughput and in fairness among *connections*.

IV. MAXIMAL FAIRNESS WITH MAXRATE SCHEDULING

Maximum throughput is achieved in our setting by using MaxRate, which in each frame transmits data to a connection with the highest instantaneous SNR, i.e., the process of selecting connection $A(t)$ must satisfy, for every t , that $A(t) = n$ implies that the transmission rate of connection n is higher than or equal to the one of any other connection in \mathcal{N} . Therefore, by definition, MaxRate is throughput-optimal, and so is our proposal.

The objective here is to study when it is possible, and how, to achieve *perfect* fairness (i.e., equal throughputs) given that the scheduler is MaxRate. We focus on fairness in the sense of equalizing the expected time-average throughput across connections, independently of their average channel quality. This maximizes the widely adopted Jain's fairness index. The only degree of freedom that the system offers to play with fairness consists in the occurrence of *ties* in the scheduling scheme, which can be frequent in systems using only few

discrete MCS values. Note also that adopting clusters with opportunistically selected cluster heads allows to increase the probability to transmit with high MCS, which in turn increases the probability of tie between connections. This effect can be exploited by designing a smart *tie-breaking rule* to employ if at least two connections compete for being scheduled with the same MCS. Formally, we use the following definitions.

Definition 2. (Best set \mathcal{M} and best MCS k) $\mathcal{M}(t, k(t))$ is the set of connections that can be scheduled with the k -th MCS in frame t , and k is the best MCS that can be used in the system in frame t , according to the SNR of the connections. We will use \mathcal{M} as short for $\mathcal{M}(t, k(t))$.

Definition 3. (Tie) A tie occurs when, in a given frame t , two or more connections can be scheduled by MaxRate with the same MCS k , which is the best possible MCS in the system at that scheduling epoch, that is: $|\mathcal{M}(t, k(t))| > 1$.

Definition 4. (Tie-breaking) A tie-breaking scheme is a procedure to select exactly one connection $i \in \mathcal{M}(t, k(t))$ to be scheduled when a tie occurs at time t .

In what follows, we first examine the multiple-connection case, finding it intractable to solve exactly. Then we study the two-connection case in detail, for which we give complete answers and which serves to develop fundamental intuition. In Section V, inspired by the results achieved for two connections, we design a heuristic for the multi-connection case.

In many of the arguments below we will rely on the fact that the expected long-term fairness (throughput distribution over an indefinitely long interval) is equivalent to the expected one-slot fairness (average per-slot throughput distribution), due to the stationarity channel assumption we made earlier.

A. Analysis of the Multiple-Connection Case

Let us denote by $Q_{n,k} := \sum_{l=1}^{k-1} p_{n,l}$ the probability that connection $n \in \mathcal{N}$ has an MCS strictly worse than k . Note that $Q_{n,1} = 0$. Let $h_n \in \{0, 1\}$ denote whether the current MCS of connection n is higher than or equal to the current MCS of any other connection ($h_n = 1$) or not ($h_n = 0$). Then, vector $\mathbf{h} := (h_n)_{n \in \mathcal{N}}$ identifies with elements 1 precisely those connections that are allowed to transmit under MaxRate. Note that, by definition, $\mathbf{h} = \mathbf{0}$ cannot happen.

We can now define

$$R^{\mathbf{h}} := \sum_{k=1}^K \left(r_k \prod_{n: h_n=1} p_{n,k} \prod_{n: h_n=0} Q_{n,k} \right), \quad (1)$$

which represents the expected (both one-slot and time-average) cumulative throughput of the system in situation \mathbf{h} (i.e., when exactly the connections specified by \mathbf{h} are in a tie situation). $R^{\mathbf{h}}$ is therefore the ‘‘tie throughput’’ to be shared between connections for which $h_n = 1$.

The number of 1’s in vector \mathbf{h} is denoted by $L_0(\mathbf{h})$ (so-called *zero ‘‘norm’’*). If $L_0(\mathbf{h}) = 1$, then there is a single connection with highest MCS, so it will be scheduled. If

$L_0(\mathbf{h}) \geq 2$, then there are several connections in a tie, and the scheduler must decide who to serve. Without loss of generality this can be done randomly. Then we need to define $L_0(\mathbf{h})$ parameters $0 \leq \alpha_n^{\mathbf{h}} \leq 1$ for every connection n such that $h_n = 1$, denoting the probability of serving connection n in situation \mathbf{h} .

Let us denote by $\mathcal{H}_1 := \{\mathbf{h} : L_0(\mathbf{h}) \geq 1\}$ the set of all vectors $\mathbf{h} \neq \mathbf{0}$. Then the aggregate throughput of the system under MaxRate scheduler is equal to $\sum_{\mathbf{h} \in \mathcal{H}_1} R^{\mathbf{h}}$, and the perfectly fair share is thus $R^* := \sum_{\mathbf{h} \in \mathcal{H}_1} R^{\mathbf{h}} / N$. We further denote by $\mathcal{H}_2 := \{\mathbf{h} : L_0(\mathbf{h}) \geq 2\}$ the set of all vectors representing ties (of at least two connections). Then, achieving perfect fairness means that the following equalities hold:

$$R^{e_n} + \sum_{\mathbf{h} \in \mathcal{H}_2: h_n=1} \alpha_n^{\mathbf{h}} R^{\mathbf{h}} = R^*, \quad \forall n \in \mathcal{N}; \quad (2)$$

$$\sum_{n \in \mathcal{N}} \alpha_n^{\mathbf{h}} = 1, \quad \forall \mathbf{h} \in \mathcal{H}_2; \quad (3)$$

where e_n is the unit vector with 1 at n -th position, and 0’s otherwise, representing the situation when connection n is the unique connection achieving the highest transmission rate.

It is easy to calculate that there are $2^N - 1$ vectors belonging to \mathcal{H}_1 and $2^N - N - 1$ vectors belonging to \mathcal{H}_2 . Therefore, there are $2^N - 1$ equations (out of which one is redundant), while having $N2^{N-1} - N$ unknowns, which is significantly more (except for $N = 2$).

We can reformulate the above problem formally as a linear programming (LP) problem:

$$\max 0 \quad (4)$$

$$\sum_{\mathbf{h} \in \mathcal{H}_2: h_n=1} \alpha_n^{\mathbf{h}} R^{\mathbf{h}} = R^* - R^{e_n}, \quad \forall n \in \mathcal{N}; \quad (5)$$

$$\sum_{n \in \mathcal{N}: h_n=1} \alpha_n^{\mathbf{h}} = 1, \quad \forall \mathbf{h} \in \mathcal{H}_2; \quad (6)$$

$$\alpha_n^{\mathbf{h}} \geq 0, \quad \forall \mathbf{h} \in \mathcal{H}_2 \text{ and } \forall n \in \mathcal{N} \text{ such that } h_n = 1. \quad (7)$$

The constant objective ($\max 0$) indicates that we are in fact interested in finding whether there is a feasible solution satisfying all the constraints. Because of the non-negativity of $\alpha_n^{\mathbf{h}}$ (expressed in (7)), feasibility of (5) necessarily requires having $0 \leq R^* - R^{e_n}$ for all n . Moreover, it is easy to see that every $\mathbf{h} \in \mathcal{H}_2$ gives one necessary condition (by adding up (5) for all n such that $h_n = 1$, and simplifying using (6) for all vectors $\mathbf{g} \leq \mathbf{h}$ and using (7) for the remaining unknowns). For instance, a two-connection tie $\mathbf{h} = e_{n,m}$ gives

$$R^{e_{n,m}} \leq (R^* - R^{e_n}) + (R^* - R^{e_m}). \quad (8)$$

This means that the throughput associated to ties of connections n and m cannot be higher than what n and m need to reach perfect fairness.

B. Analysis of the Two-Connection Case

If there are many connections, the size of the problem becomes too large to be solvable at milliseconds scale in

a real base station. Therefore, we analyze the case of two connections with the aim to get more insight into the problem in order to design a heuristic. The two-connection case can be analyzed and solved without the need of using numerical methods. To make the notation more intuitive, we denote the following quantities:

$$R^{(1)} := R^{e_1} = \sum_{k=2}^K r_k p_{1,k} Q_{2,k}, \quad (9)$$

$$R^{(2)} := R^{e_2} = \sum_{k=2}^K r_k p_{2,k} Q_{1,k}, \quad (10)$$

$$R^{(X)} := R^{e_{1,2}} = \sum_{k=1}^K r_k p_{1,k} p_{2,k}, \quad (11)$$

which represent the expected (both one-slot and time-average) transmission rates in the following three cases: Eq. (9) expresses the rate of connection 1 when it has an MCS strictly better than connection 2; Eq. (10) is for connection 2 having an MCS strictly better than connection 1; and Eq. (11) is for the case of tie. Note that the aggregate throughput of the system under MaxRate scheduler is equal to $R^{(1)} + R^{(2)} + R^{(X)}$.

In the following proposition we give a sufficient and necessary condition for a scheduler that achieves both maximal throughput and perfect fairness (i.e., all connections achieve the same throughput).

Proposition 1. *The MaxRate scheduler can achieve both one-slot and time-average fairness if and only if*

$$\left| R^{(1)} - R^{(2)} \right| \leq R^{(X)}. \quad (12)$$

The proof is given in Appendix B. Here, it is worth to discuss when such a condition might hold. Indeed, there are some intuitive sufficient conditions stated next, which are independent of the transmission rates r_k .

Proposition 2. *The MaxRate scheduler can achieve both one-slot and time-average fairness if any of the following conditions hold:*

- 1) $p_{1,k} = p_{2,k}$ for all $k \geq 2$ (i.e., the channels of the two connections are statistically equal);
- 2) $|p_{1,k} Q_{2,k} - p_{2,k} Q_{1,k}| \leq p_{1,k} p_{2,k}$ for all $k \geq 2$;
- 3) $p_{1,k} \geq p_{2,k}$ for all $k \geq 2$ and $p_{2,K} \geq 1/2$.

The proof uses arguments similar to the ones used in the proof of Proposition 1 and is thereby omitted for brevity. Moreover, there may be weaker conditions which make it likely that perfect fairness is achievable. For instance, if one of the following conditions holds, perfect fairness is achievable:

- 1) $p_{1,k} p_{2,k}$ is large enough for all k large enough;
- 2) $|p_{1,k} Q_{2,k} - p_{2,k} Q_{1,k}|$ small enough for all k large enough;
- 3) probabilities $p_{n,k}$ for all k large enough are approximately equal for the two connections;
- 4) the expression $p_{1,k} Q_{2,k} - p_{2,k} Q_{1,k}$ often changes sign as k grows.

Finally, taking into account that transmission rates r_k grow somewhat exponentially with k (see Table I in Appendix A), it is much more important that the two connections are statistically similar in the upper MCS range rather than in the lower MCS range.

Let us define now the *MaxRate-MaxFair scheduler* for two connections, by introducing a bias in the tie-breaking rule of the MaxRate scheduler as described in the following definition.

Definition 5. (MaxRate-MaxFair scheduler) *In case a tie occurs under MaxRate scheduling, serve connection 1 with probability $\alpha^{(X)}$ and serve connection 2 with probability $1 - \alpha^{(X)}$, where*

$$\alpha^{(X)} := \frac{1}{2} + \frac{R^{(2)} - R^{(1)}}{2R^{(X)}}. \quad (13)$$

Moreover, if $\alpha^{(X)} \notin [0, 1]$, then it is not a proper probability value, thus we cut such values off:

$$\begin{cases} \text{if } \alpha^{(X)} < 0 \text{ then } \alpha^{(X)} := 0; \\ \text{if } \alpha^{(X)} > 1 \text{ then } \alpha^{(X)} := 1. \end{cases} \quad (14)$$

The following proposition establishes when $\alpha^{(X)}$ is a proper probability value, so that no cut-off is needed. The proof is immediate, therefore we omit it.

Proposition 3. *Condition (12) is equivalent to $\alpha^{(X)} \in [0, 1]$ as defined in (13).*

Using MaxRate-MaxFair, connection 1 receives $R^{(1)} + \alpha^{(X)} R^{(X)}$, while the throughput of connection 2 is $R^{(2)} + (1 - \alpha^{(X)}) R^{(X)}$. Such a throughput distribution is the fairest possible, and the aggregate is maximum, as stated in the following proposition, which is the main result of this section and validates the MaxRate-MaxFair name of the above scheduler.

Proposition 4. *If (12) holds, then the MaxRate-MaxFair scheduler achieves maximum throughput, and both one-slot and time-average fairness. If (12) does not hold, then the MaxRate-MaxFair scheduler achieves maximum throughput, and the difference between individual throughputs is the minimum achievable with tie-breaking schemes.*

The proof is presented in Appendix C. According to Proposition 4, when condition (12) does not hold, the scheduler still achieves maximum throughput, but will not be perfectly fair anymore. Nevertheless, the difference between individual throughputs will be the minimum possible, and may significantly outperform randomized tie-breaking. Note that, from the proof of Proposition 4, it follows that, using the Jain's fairness index as metric, the MaxRate-MaxFair scheduler achieves the smallest possible distance from the perfectly fair throughput distribution. In fact, the Jain's fairness index is maximized when differences are minimized. The result is formalized in the following corollary.

Corollary 1. *MaxRate-MaxFair scheduler achieves the highest possible Jain's fairness index achievable by means of any tie-breaking scheme.*

V. MAXIMAL FAIRNESS WITH MAXRATE SCHEDULING OF MULTIPLE CONNECTIONS

Having developed fundamental intuition based on exact results for the case of two connections, we now set out to design a tie-breaking scheme for MaxRate schedulers for a generic number of connections N , achieving better fairness than with randomized tie-breaking. Note that, as discussed in Section IV-A, extending the analytical approach from the two-connection case would require $N2^{N-1} - N$ tie-breaking parameters for all possible ties of $2, 3, \dots, N$ connections, which grows too fast to be implementable (2, 6, 28, 75, ...). Instead, we focus on scalable solutions that rely on at most N tie-breaking parameters.

A. WRR Tie-breaking

We design a scheduler in which tie-breaking is resolved as if Weighted Round Robin (WRR) was implemented (only for ties). Thus, we assume that for each connection n there is a non-negative parameter α_n used as follows.

Definition 6. (WRR Tie-breaking) *If, at a given scheduling epoch t , \mathcal{M} is the set of connections that are currently tied in the highest MCS (see Definition 2), then the probability (or, the average fraction of time) that connection $n \in \mathcal{M}$ is served in such situations is as follows:*

$$\alpha_n / \sum_{m \in \mathcal{M}} \alpha_m, \quad \text{where } \alpha_m \geq 0 \quad \forall m \in \mathcal{M}. \quad (15)$$

The expected throughput of connection n under MaxRate with WRR tie-breaking is then:

$$\sum_{k=1}^K r_k \sum_{\mathcal{M} \ni n} \left[\frac{\alpha_n}{\sum_{m \in \mathcal{M}} \alpha_m} \prod_{m \in \mathcal{M}} p_{m,k} \prod_{m \notin \mathcal{M}} Q_{m,k} \right], \quad (16)$$

where $\mathcal{M} \ni n$ denotes any set of connections that includes n .

Using (16) it is, however, intractable to obtain values of α_n which would equalize the expected individual throughput of all connections, since it leads to a large system of non-linear equations. Hence, in the following we propose a heuristic.

B. The PIKe Heuristic

To compute α_n for each connection n , we propose a heuristic that follows the scheme used for $\alpha^{(X)}$ in a two-connection system (see Eq. (13)) and that scales linearly with the number of connections. With our heuristic, connection n competes with the rest of connections, which we globally indicate as *virtual connection* “ $-n$ ”, and we assume that the target of connection n is to achieve a portion $1/N$ of the cell throughput. We name our heuristic Priority Keying (PIKe), since it essentially maps channel qualities onto connection priorities in order to break ties.

Following the ideas from the two-connection case, we first define $Q_{-n,k}$ as the CDF for the best MCS of all the

connections except connection n , $Q_{-n,k} := \prod_{m=1, m \neq n}^N Q_{m,k}$,

and $p_{-n,k}$ as the probability that at least one of the connections (except connection n) is in MCS k and no other connection is in a better MCS, formally, $p_{-n,k} = Q_{-n,k+1} - Q_{-n,k}$. Then, (9) to (11) can be rewritten as follows, for $n = 1 \dots N$:

$$R^{(n)} := \sum_{k=2}^K r_k p_{n,k} Q_{-n,k}; \quad (17)$$

$$R^{(-n)} := \sum_{k=2}^K r_k p_{-n,k} Q_{n,k}; \quad (18)$$

$$R_n^{(X)} := \sum_{k=1}^K r_k p_{n,k} p_{-n,k}. \quad (19)$$

The total cellular throughput of MaxRate is $R_{Tot} = R^{(n)} + R^{(-n)} + R_n^{(X)}$, which is the same for all values of n , as it can be easily verified.

Proposition 5. *The MaxRate scheduler for a system with $N \geq 2$ connections can achieve both one-slot and time-average fairness if and only if*

$$\left| R^{(n)} - R^{(-n)} \right| \leq R^{(X)}, \quad \forall n \in \mathcal{N}. \quad (20)$$

The proof of the Proposition 5 derives from the proof of Proposition 1. The value of the tie-breaking probability α_n is computed from the following equation:

$$R^{(n)} + \alpha_n R_n^{(X)} = \frac{1}{N} R_{Tot}. \quad (21)$$

The resulting value of α_n is then:

$$\alpha_n = \frac{1}{N} + \frac{\sum_{k=1}^K r_k [p_{-n,k} Q_{n,k} - (N-1) p_{n,k} Q_{-n,k}]}{N \sum_{k=1}^K p_{n,k} p_{-n,k} r_k}. \quad (22)$$

However, when the expected transmission rate of connection n is strictly higher than the average fair individual share (i.e., $R^{(n)} > \frac{1}{N} R_{Tot}$), then α_n is negative, which is not acceptable for the WRR scheduling scheme. This corresponds to a situation in which the resources in ties are not enough to equalize the connection throughputs without loss of throughput maximality. Therefore, we propose the following transformation which preserves the order of α_n , i.e., preserves the *priority list* among connections, without wasting throughput:

$$\min_{m \in \mathcal{N}} \alpha_m < 0 \Rightarrow \alpha_n := \alpha_n - \min_{m \in \mathcal{N}} \alpha_m, \quad \forall n \in \mathcal{N}. \quad (23)$$

In the next section we will quantify the level of fairness achieved in the system with our heuristic. Although we will use static clustering scenarios to illustrate the advantages of PIKe, we remark that our methodology and findings apply to clusters whose composition varies in time, as some users may turn on/off their devices or migrate to another cluster or another cell. In particular, we note that in presence of *flows* with limited duration, probabilities $p_{n,k}$ described in Session III can be adapted to represent the steady-state probabilities of flow n given the arrival and flow-size distributions.

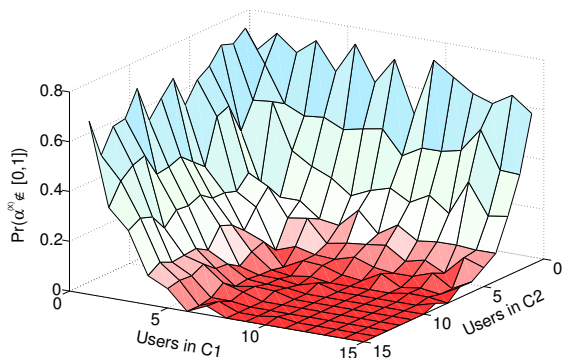


Fig. 2. Probability that perfect fairness cannot be achieved.

VI. EVALUATION

In this section we validate numerically our proposal. We use the Jain’s fairness index to compare the effectiveness of PIKe as compared to commonly deployed Equal Time (ET) and Proportional Fair (PF) schedulers with unclustered mobile users. Additionally, we compare PIKe to a MaxRate scheduler applied to connections (i.e., clustered users) in which tie-breaking is operated uniformly at random (denoted as MR). As for the throughput, we normalize it to cell capacity. Since our work does not investigate intra-cluster mechanisms, we assume that the total connection throughput can be equally shared among clustered members within the connection.

For the sake of tractability, in what follows we assume that mobile users belong to one of three predefined SNR *classes*, which correspond to *poor*, *average*, and *good* average SNR, i.e., set Γ contains three elements only. The designated average SNR for different classes are chosen in a manner that the mean achievable rates for *poor*, *average*, and *good* users are 20%, 50%, and 80% of the maximum transmission rate achievable in the system, respectively. Therefore, with the MCS values reported in Table I in Appendix A and the assumed Rayleigh fading model, the average SNR values for *poor*, *average* and *good* users are $\gamma_m = 7\text{ dB}$, 16 dB , 23 dB , respectively.

A. Impact of Cluster Size on Tie-breaking Opportunities

We start by evaluating the importance of clustering for fair tie-breaking operations. Under MaxRate scheduling, both user-based and cluster-based scheduling are throughput-optimal. However, the advantage of clustering is twofold: (i) opportunistic cluster head selection yields higher channel qualities used to transmit in the system, which contributes to equalize transmission rates among connections (note that efficiency curves increase logarithmically according to Shannon’s results); and (ii) clustering offers the possibility to re-distribute the throughput among cluster members, thus yielding potentially higher fairness. We exemplify these effects by considering the specific case of two connections. In this case, the exact analysis of Subsection IV-B applies, and PIKe implements exactly the *MaxRate-MaxFair* tie-breaking of Definition 5.

Fig. 2, which is the result of 20,000 random instances for two connections with random number and quality of clustered users, is in line with our intuition on the effect of clustering. The figure shows the probability that $\alpha^{(X)}$ is




Fig. 3. Throughput comparison (average and standard deviation) of connections with 1 to 10 clustered users, with uniform quality distribution.

outside the acceptable range $[0, 1]$, when the perfect fairness can not be achieved. The figure reveals that perfect fairness can be achieved almost surely when both connections consist of more than 5 clustered users. The probability of *fairness non-achievability* radically increases as the cluster size drops below 5 users, since the average connection qualities can be very unbalanced and yield large $|R^{(1)} - R^{(2)}|$. In contrast, when clusters are large enough (i.e., with more 5 – 10 members), the fact that each cluster head exhibits the highest SNR in its cluster makes the probability to use the best MCS close to 1, and thus $|R^{(1)} - R^{(2)}|$ approximates 0, while increasing the probability of ties. Therefore, (12) is met with high probability.

B. Throughput of Connections

The clustering gain versus the conventional cellular architecture is illustrated here with simple numerical calculations.

Fig. 3 shows the difference in throughput achieved by ET, PF¹ and PIKe as a function of the number of connections in the network. Results are averaged over 2,000 random instances, and user qualities are uniformly distributed among *poor*, *average*, and *good*. Cluster sizes are chosen at random, ranging uniformly from 1 to 10 members. Of course, the throughput of ET and PF only depends on the number of mobile users and their channel qualities, but we keep using the number of connections as reference. Interestingly, PIKe doubles the throughput of ET and outperforms PF by more than 20%. Most importantly, PIKe can nearly achieve 100% of the achievable throughput.

C. Fairness between Connections

Finally, we validate our proposed PIKe heuristic for maximum rate with maximum fairness by comparing its fairness performance figures with the ones obtained under legacy ET and PF schedulers, and under MR. Note that, according to the common understanding, PF should have much higher fairness than MaxRate-based approaches [26], while we show that the opposite is true with PIKe. For ET and PF, we numerically simulate the scheduling of single users, then we sum up the throughput of users according to which cluster/connection they belong to. Note that PIKe and MR use both MaxRate, hence they achieve the same throughput.

¹PF results are obtained by simulating a scheduling process in which the average user throughput is computed with an autoregressive filter with exponential time constant equal to 1,000 frames. However results computed with time constant in the range 50 to 5,000 do not significantly differ.




Fig. 4. Fairness comparison (connections comprising 1 to 10 clustered users).




Fig. 5. Fairness comparison (connections comprising 5 to 10 clustered users).

Fig. 4 reports fairness indexes for the different scheduling schemes for systems with 2 to 6 connections, each formed by 1 to 10 members. Fairness indexes are shown in terms of *box-and-whiskers* plots, reporting minimum and maximum values recorded over the set of simulations performed, the 25th and 75th percentiles of their values, and the average (the solid dashes in the boxes reported in the figure). Interestingly, the level of fairness achieved by our proposed scheme is very high, and PIKe achieves at least as much fairness as PF. Recalling that the throughput achieved by PF is much lower than the one achieved by PIKe, this result is very encouraging.

Even more interestingly, simulations accounting for connections with at least 5 clustered users reveal that PIKe can outperform ET and PF in terms of fairness, as depicted in Fig. 5, where the boxes delimited by the 25th and 75th percentiles are very close to 1 for the PIKe scheme. Specifically, Figs. 4 and 5 also show that PIKe reduces the distance from perfect fairness (i.e., from 1) by 50% with respect to MR.

To conclude, we remark that PIKe is beneficial for both throughput and fairness, which means that it would allow *better worst-case performance* in comparison to ET and PF. Indeed, Fig. 6 shows that the minimum throughput received by a cluster member in the system, using PIKe, is much higher than the one achieved with ET (by a factor ~ 1.5 or more), and PF (by a factor ~ 2).

VII. CONCLUSIONS

In this paper, we have shown how tie-breaking can lead to maximal fairness in cellular networks without sacrificing any throughput. Although scheduling ties are usually ignored and uniformly random tie-breaking is an accepted practice, we have shown that this practice is rather inefficient. Our simulations indicate that there is a great potential for fairness and throughput enhancements in customized tie-breaking as




Fig. 6. Minimum throughput attained by a cluster member (with connections comprising 5 to 10 clustered users).

soon as tie probability becomes relevant, as achieved, e.g., with D2D and opportunistic clustering strategies. We have rigorously formulated a tie-breaking scheme that achieves maximal fairness *and* maximal throughput. The complexity of the optimal approach is not tractable for more than two connections. Inspired by the optimal solution for the case of 2 connections, we have designed PIKe, a heuristic that scales linearly with the number of connections. Our results confirm that PIKe achieves almost perfect fairness and maximal throughput and largely benefits of D2D-assisted opportunistic clustering schemes.

REFERENCES

- [1] R. Knopp and P. Humblet, "Information capacity and power control in single-cell multiuser communications," in *Proceedings of IEEE ICC*, vol. 1, Jun. 1995, pp. 331–335.
- [2] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi, "CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users," *IEEE Communications Magazine*, vol. 38, no. 7, pp. 70–77, 2000.
- [3] X. Liu, E. K. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks," *Computer Networks*, vol. 41, no. 4, pp. 451–474, 2003.
- [4] U. Ayesta, M. Eraisquin, and P. Jacko, "A modeling framework for optimizing the flow-level scheduling with time-varying channels," *Performance Evaluation*, vol. 67, pp. 1014–1029, 2010.
- [5] M. Neely, "Order optimal delay for opportunistic scheduling in multi-user wireless uplinks and downlinks," *IEEE/ACM Transactions on Networking*, vol. 16, no. 5, pp. 1188–1199, 2008.
- [6] L. Lei, Z. Zhong, C. Lin, and X. Shen, "Operator controlled device-to-device communications in LTE-advanced networks," *IEEE Wireless Communications*, vol. 19, no. 3, pp. 96–104, 2012.
- [7] F. H. Fitzek, M. Katz, and Q. Zhang, "Cellular controlled short-range communication for cooperative P2P networking," *Wireless Personal Communications*, vol. 48, no. 1, pp. 141–155, 2009.
- [8] K. Doppler, M. Rinne, C. Wijting, C. Ribeiro, and K. Hugl, "Device-to-device communication as an underlay to LTE-advanced networks," *IEEE Communications Magazine*, vol. 47, no. 12, pp. 42–49, 2009.
- [9] H. Wu, C. Qiao, S. De, and O. Tonguz, "Integrated cellular and ad hoc relaying systems: iCAR," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 10, pp. 2105–2115, 2001.
- [10] A. Sendonaris, E. Erkip, and B. Aazhang, "User cooperation diversity. Part I. System description," *IEEE Transactions on Communications*, vol. 51, no. 11, pp. 1927–1938, 2003.
- [11] C.-H. Yu, K. Doppler, C. B. Ribeiro, and O. Tirkkonen, "Resource sharing optimization for device-to-device communication underlying cellular networks," *IEEE Transactions on Wireless Communications*, vol. 10, no. 8, pp. 2752–2763, 2011.
- [12] H. Kushner and P. Whiting, "Convergence of proportional-fair sharing algorithms under general conditions," *IEEE Transactions on Wireless Communications*, vol. 3, pp. 1250–1259, 2004.
- [13] S. Borst, "User-level performance of channel-aware scheduling algorithms in wireless data networks," *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, pp. 636–647, 2005.

- [14] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 466–478, 1993.
- [15] S. Shakkottai, R. Srikant, and A. L. Stolyar, "Pathwise optimality of the exponential scheduling rule for wireless channels," in *Advances in Applied Probability*, 2004, pp. 1–30.
- [16] B. Sadiq and G. de Veciana, "Throughput optimality of delay-driven Maxweight scheduler for a wireless system with flow dynamics." in *Proceedings of Annual Allerton Conference on Communication, Control and Computing*, 2009.
- [17] S. Liu, L. Ying, and R. Srikant, "Throughput-optimal opportunistic scheduling in the presence of flow-level dynamics," *IEEE/ACM Transactions on Networking*, vol. 19, no. 4, pp. 1057–1070, 2011.
- [18] P. Jacko, "Value of information in optimal flow-level scheduling of users with Markovian time-varying channels," *Performance Evaluation*, vol. 68, no. 11, pp. 1022–1036, 2011.
- [19] U. Ayesta, M. Erausquin, M. Jonckheere, and I. M. Verloop, "Scheduling in a random environment: Stability and asymptotic optimality," *IEEE/ACM Transactions on Networking*, vol. 21, no. 1, pp. 258–271, Feb. 2013.
- [20] W. Wang, V. Subramanian, and D. Guo, "Low complexity scheduling algorithms for wireless networks with full duplex state exchange," in *Information Sciences and Systems (CISS), 2014 48th Annual Conference on*. IEEE, 2014, pp. 1–6.
- [21] B. Ji, G. R. Gupta, X. Lin, and N. B. Shroff, "Performance of low-complexity greedy scheduling policies in multi-channel wireless networks: Optimal throughput and near-optimal delay," in *Proceedings of IEEE INFOCOM*. IEEE, 2013, pp. 2589–2597.
- [22] S. Bodas, S. Shakkottai, L. Ying, and R. Srikant, "Scheduling in multi-channel wireless networks: Rate function optimality in the small-buffer regime," *IEEE Transactions on Information Theory*, vol. 60, no. 2, pp. 1101–1125, 2014.
- [23] A. Asadi, V. Mancuso, and R. Gupta, "An SDR-based Experimental Study of Outband D2D Communications," in *to appear in Proceedings of IEEE INFOCOM [online at <https://www.dropbox.com/s/f4uu17qslaqkpi/Infocom-D2D-Shared.pdf?dl=0>]*, Apr. 2016.
- [24] E. F. Chaponniere, P. J. Black, J. M. Holtzman, and D. N. C. Tse, "Transmitter directed code division multiple access system using path diversity to equitably maximize throughput," US Patent US 6,449,490 B1, 2002.
- [25] U. Ayesta and P. Jacko, "Method for selecting a transmission channel within a time division multiple access (TDMA) communications system," EU Patent EP2 384 076 B1, 2013.
- [26] R. Kwan, C. Leung, and J. Zhang, "Proportional fair multiuser scheduling in LTE," *IEEE Signal Processing Letters*, vol. 16, no. 6, pp. 461–464, 2009.
- [27] S. Sesia, I. Toufik, and M. Baker, *LTE-the UMTS long term evolution: from theory to practice*. Wiley, 2011.

APPENDIX

A. Transmission Rates and Cluster MCS Selection

The instantaneous achievable rate of connection n at slot t , $R_n(t) = r_k$, depends on the adopted MCS $k = 1, 2, \dots, K$. We assume that the actual MCS for connection n at slot t is selected as a function of the instantaneous SNR $C_n(t)$, i.e.:

$$\begin{aligned} R(t) = r_k &\iff MCS_n(t) = k \iff C_n(t) \in [c_k; c_{k+1}[, \\ 0 &= c_1 < c_2 < \dots < c_K < c_{K+1} = \infty, \\ 0 &= r_1 < r_2 < \dots < r_K. \end{aligned}$$

A connection n is a cluster composed by m_n mobiles, and its instantaneous SNR is the *highest* SNR among the users composing the cluster. Hence, the probability $p_{n,k}$ that a scheduled connection n receives data encoded with MCS k can be computed based on the SNR CDF and the MCS thresholds:

$$p_{n,k} = \int_{c_k}^{c_{k+1}} dF_n(z), \quad (24)$$

where $F_n(z)$ is the CDF of the maximum of m_n random variables representing the SNR values of each of the m_n mobiles forming cluster n :

$$F_n(z) = \prod_{j=1}^{m_n} \left(1 - e^{-\frac{z}{\gamma_j}}\right), \quad z \geq 0, \quad n \in \mathcal{C}. \quad (25)$$

Table I shows a list of possible modulation and coding schemes for LTE-like networks [27], their coding rate, and the SNR threshold (in dB) that has to be reached to achieve a negligible error rate. The table also contains the net transmission rate, in bits per symbol, achieved with each MCS. The Implementation Margin (IM) in Table I is a value that represents the noise due to non-ideal receiver. In our simulation, MCS thresholds c_k include both SNR and IM.

TABLE I
MODULATION AND CODING SCHEMES AND THEIR THRESHOLDS

Modulation	Coding Rate	SNR (dB)	IM (dB)	SNR+IM (dB)	Bits per symbol
QPSK	1/8	-5.1	2.5	-2.6	0.25
	1/5	-2.9		-0.4	0.4
	1/4	-1.7		0.8	0.5
	1/3	-1		1.5	0.67
	1/2	2		4.5	1
	2/3	4.3		6.8	1.3
	3/4	5.5		8.0	1.5
16QAM	4/5	6.2	3	8.7	1.6
	1/2	7.9		10.9	2
	2/3	11.3		14.3	2.66
	3/4	12.2		15.2	3
64QAM	4/5	12.8	4	15.8	3.2
	2/3	15.3		19.3	4
	3/4	17.5		21.5	4.5
	4/5	18.6		22.6	4.8

B. Proof of Proposition 1

Proof: Observe first that $R^{(n)}$ is the minimum achievable transmission rate for connection $n = \{1, 2\}$ under the MaxRate scheduler, since this connection must be served if it has an MCS strictly better than the other connection. Moreover, $R^{(n)} + R^{(X)}$ is the maximum achievable transmission rate for connection $n = \{1, 2\}$ with a MaxRate scheduler. Then, fairness cannot be achieved if either $R^{(1)} > R^{(2)} + R^{(X)}$ or $R^{(2)} > R^{(1)} + R^{(X)}$, which is equivalent to $|R^{(1)} - R^{(2)}| > R^{(X)}$. In contrast, fairness can be achieved if (12) holds, as shown in (4). Therefore, the equivalence holds. ■

C. Proof of Proposition 4

Proof: Consider the MaxRate scheduler with randomized tie-breaking with bias $\alpha \in [0, 1]$ for connection 1. Then, the expected time-average and one-slot individual throughputs are $R^{(1)} + \alpha R^{(X)}$ and $R^{(2)} + (1 - \alpha)R^{(X)}$, respectively. It is straightforward to verify that, if (12) holds, plugging $\alpha^{(X)}$ for α the throughput of each connection is equal to $(R^{(1)} + R^{(2)} + R^{(X)})/2$.

If (12) does not hold, then suppose that $R^{(1)} > R^{(2)} + R^{(X)}$ (case $R^{(2)} > R^{(1)} + R^{(X)}$ is analogous). The difference in the individual throughputs is $R^{(1)} + \alpha R^{(X)} - (R^{(2)} + (1 - \alpha)R^{(X)}) = R^{(1)} - R^{(2)} - R^{(X)} + 2\alpha R^{(X)}$, which is minimized if $\alpha = 0$. Indeed $\alpha^{(X)} = 0$, because of the cut-off of a negative value given by (13). ■