

Reciprocity and Efficiency in Peer Exchange of Wireless Nodes through Convex Optimization

Fernando Paganini, *Fellow, IEEE*, Martín Zubeldía, *Student Member, IEEE* and Andrés Ferragut

Abstract—This paper considers the allocation of exchange rates in a network of wireless nodes which engage in peer-to-peer dissemination. Here, in addition to the desirable throughput efficiency, it is important to ensure a level of rate reciprocity between peers, an issue that has been studied before only for wired networks. For the wireless substrate efficiency and reciprocity may be in conflict, due to the non-uniform link capacities of different peering choices, and the possible interference between them. We use convex optimization to formulate a relevant tradeoff, measuring reciprocity through a Kullback-Leibler divergence between sent and received rates. We propose decentralized methods which involve peer-to-peer interactions, and are shown to converge to the corresponding tradeoff point. Illustrative simulations are provided.

Index Terms—Wireless networks, peer-to-peer, reciprocity, convex optimization.

1 INTRODUCTION

The efficient dissemination of a large content file among a set of network nodes requires the orchestration of multiple transfers, within the constraints of the communication substrate. Since there is typically no central network planner, efficiency must be pursued by decentralized algorithms. Among them, the peer-to-peer (P2P) approach in which nodes exchange pieces of the file of common interest has emerged with great strength. By turning client nodes into servers, a service capacity is deployed that scales with content demand.

A fundamental premise of such cooperative dissemination is that a certain level of *reciprocity* is enforced: peers will participate as servers inasmuch as they receive as clients. For this reason, the most successful P2P implementations (e.g. BitTorrent [7]) include heuristic algorithms to incentivize reciprocation. Formal studies have also tackled this question, with crisp results in the case of a *proportional* criterion for reciprocity first investigated in [24]. Dynamic iterative algorithms based on this principle can be proved to converge to the desired fairness [23], and lead to alternate protocol implementations [15]. Further analysis of this idealized proportional fairness criterion, and practical approximations through neighbor selection, are provided in [25].

All of the above references focus on the main application domain of P2P, which is the global Internet, with most peers sitting behind wired access links; for tractability it is assumed that these are the only bottlenecks, and the remainder of the network is idealized. Under this assumption the overall service capacity is fixed, the only decision is how to allocate it between client peers. There can still be tradeoffs

with performance as recognized in [9], when peer arrivals and departures are considered. However the *efficiency* of the network substrate at any time is agnostic to the peering choice, so from this standpoint reciprocity is not a burden.

In this paper we are interested in the reciprocity question for a network of *wireless* nodes which engage in mutual communication under the peer-to-peer philosophy. Our motivation is academic, but we believe of interest due to the established potential of this dissemination method. There has been work on epidemic approaches to content dissemination (e.g. [22]), where wireless networks motivate a time-varying network graph, but where capacity is ignored: exchanges become a one-shot interaction between nodes. Here we consider the opposite situation of long-term dissemination over a network that has fixed topology, but where capacity is constrained, depending on the peering choices and on the management of interference. There is, of course, ample literature on resource allocation in wireless networks from the network utility maximization perspective (NUM, see e.g. [6], [16]); here efficiency and fairness between *sending* flows is considered. This is, however, different from the send/receive reciprocity sought here.

The paper is organized as follows. In Section 2 we set up a general framework for studying efficiency and a convex measure of reciprocity, namely the Kullback-Leibler divergence (e.g., [3]) between sent and received rates. Two versions, global and peerwise reciprocity, are considered, the latter more amenable to decentralized optimization. In the wired network case they can be optimized with methods from [25], as reviewed in Section 3. For a wireless substrate, a first difference addressed in Section 4 is that outgoing links from each peer will have a non-uniform capacity according to the destination. This introduces an efficiency/reciprocity tradeoff, formalized in terms of a convex optimization problem. We develop a solution based on Gauss-Seidel updates that involve decentralized peerwise information, which is proved to converge, and tested in simulation examples.

In Section 5 we tackle the additional issue of link inter-

-
- F. Paganini and A. Ferragut are with the Mathematical Analysis in Telecommunications and Energy (MATE) Research Group, Universidad ORT Uruguay. E-mail: {paganini,ferragut}@ort.edu.uy Research supported in part by AFOSR-US under Grant FA9550-15-1-0183.
 - M. Zubeldia is with LIDS, MIT, Cambridge, MA, USA. Email: zubeldia@mit.edu

ference, which makes decentralization far more challenging. Through duality our convex optimization can be decoupled between the reciprocity and medium access components, but similar to NUM problems the latter involves maximum-weight scheduling. Given its complexity, we adapt to our problem the Markov approximation strategy of [5], [12]; this leads to a primal-dual dynamics that is decentralized and is shown to converge to the optimal tradeoff. We also present simulation examples of this procedure. Conclusions are given in Section 6, and some proofs deferred to Appendices. A partial version of our results was presented in [18].

2 EFFICIENCY AND RECIPROCITY IN P2P SHARING

We consider a fixed population of N peers who engage in bilateral exchange of information¹. Their connectivity is specified by a set \mathcal{L} of allowable links, or alternatively through an adjacency matrix $A = (a_{ij})$, where $a_{ij} = 1$ when i can send data to j , and $a_{ij} = 0$ otherwise; we set $a_{ii} = 0$. A is assumed *symmetric*, and with no rows of zeros (no disconnected peers).

Define a *resource sharing matrix* $Z \in \mathbb{R}_+^{N \times N}$ in which z_{ij} represents the offered throughput from peer i to peer j . Z is constrained by connectivity, i.e. it must satisfy

$$z_{ij} \geq 0, \quad z_{ij} = 0 \text{ if } a_{ij} = 0. \quad (1)$$

Z will also be subject to bandwidth constraints, of a nature depending on the network substrate: postponing this question, we refer to a set \mathcal{Z} of feasible resource sharing matrices.

The aggregate sent and received rates per peer are

$$s_i(Z) = \sum_j z_{ij} \quad \forall i; \quad r_j(Z) = \sum_i z_{ij} \quad \forall j. \quad (2)$$

In matrix form, we can write: $Z\mathbf{1} = s$, $\mathbf{1}^T Z = r^T$, where $s, r, \mathbf{1} \in \mathbb{R}^N$ are interpreted as column vectors, the latter being the vector of ones, and T denotes transpose.

We now proceed to specify the desirable objectives on the resource sharing matrix $Z \in \mathcal{Z}$. A first natural objective is the total rate of the exchange, obtained by

$$R(Z) = \sum_{i,j} z_{ij} = \sum_i s_i = \sum_j r_j. \quad (3)$$

Maximizing this quantity would lead to the most *efficient* use of the network in the sense of total throughput. This can also be fit in the language of network utility maximization (NUM) [6], [14] by considering a linear utility in either sent or received rates.²

This efficiency objective is however insufficient for peer-to-peer networks; it may happen that the maximum rate allocation leaves a peer receiving no data, which is not compatible with a bidirectional exchange. Instead, to provide proper incentives for cooperation between peers it is desirable to have approximate parity between the rates a peer sends and receives. See [9], [24] for further justification of this fairness criterion.

1. In BitTorrent [7] parlance we focus only on the *leechers*, who both send and receive data.

2. A strictly concave utility function of rate is often preferred, which gives stronger properties to the optimization. We will not pursue this here.

This motivates our introduction of a quantitative measures of discrepancy between the upload and download rate vectors. A first choice is the Kullback-Leibler (KL) divergence or relative entropy (see [3], [8])

$$D(s||r) := \sum_j s_j \log \left(\frac{s_j}{r_j} \right), \quad (4)$$

a jointly convex function of both vectors. Since s and r have the same sum ($R(Z)$ in (3)), their KL divergence is always non-negative [8] and only zero if $s = r$, i.e. if every peer receives as much throughput as it provides to the network.

Therefore, an allocation matrix Z with small $D(s||r)$ achieves an approximate level of reciprocity, *globally* between each peer and the rest of the swarm.

An alternative, *peerwise* notion of reciprocity is to require that peers i and j share equal amounts of bandwidth, i.e. that Z is symmetric. Approximate peerwise reciprocity can be measured by the KL divergence between *matrices*

$$D(Z||Z^T) := \sum_{i,j:a_{ij}=1} z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right). \quad (5)$$

These dual notions of reciprocity are closely related to the (global and peerwise) proportional fairness notions in [24] for wired networks; we elaborate on these connections below.

The following Lemma formalizes the intuitive fact that the latter notion of reciprocity is more restrictive than the first:

Lemma 1. $D(Z||Z^T) \geq D(s||r)$ for any allocation Z , and equality holds if and only if

$$\frac{z_{ij}}{z_{ji}} = \frac{s_i}{r_i} \quad \text{for any } i, j \text{ with } a_{ij} = 1.$$

Proof. The result follows from the log-sum inequality (see [8]), which for each fixed i implies

$$\sum_j z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right) \geq \left(\sum_j z_{ij} \right) \log \left(\frac{\sum_j z_{ij}}{\sum_j z_{ji}} \right) = s_i \log \left(\frac{s_i}{r_i} \right).$$

Adding over i gives the desired bound. Conditions for equality also follow from those in [8]. \square

An alternative expression for the peerwise reciprocity measure is obtained by grouping terms (i, j) and (j, i) in (5), writing $D(Z||Z^T) := \sum_{\{i>j:a_{ij}=1\}} d(z_{ij}, z_{ji})$, where

$$\begin{aligned} d(x, y) &:= x \log \left(\frac{x}{y} \right) + y \log \left(\frac{y}{x} \right) \\ &= D \left(\left[\begin{array}{c} x \\ y \end{array} \right] \middle\| \left[\begin{array}{c} y \\ x \end{array} \right] \right). \end{aligned} \quad (6)$$

The function $d(x, y)$ is well-defined and non-negative when $x > 0, y > 0$, zero only if $x = y$. At the boundary of \mathbb{R}_+^2 it has the following behavior:

- When $x \rightarrow 0+$ and $y > 0$, $d(x, y) \rightarrow +\infty$. We can adopt the extended function definition $d(0, y) = +\infty$ if $y > 0$. Similarly, $d(x, 0) = +\infty$ if $x > 0$.
- If both $x \rightarrow 0, y \rightarrow 0$, following a straight line with slope $m > 0$, $d(x, mx) \rightarrow 0$. But more unbalanced

approaches $(x, y(x))$ to $(0, 0)$ can yield different limits, so as a function of two variables there is no limit. We adopt the convention

$$d(0, 0) = 0, \quad (7)$$

but emphasize the lack of continuity at this boundary point.

An additional comment is the symmetry of $d(x, y)$ in its two variables, which highlights the fact that our peerwise reciprocity measure is symmetric: $D(Z||Z^T) = D(Z^T||Z)$.

Having defined the key notation, we state the main question of interest for our paper: under the physical constraints of each specific network scenario, what is a suitable tradeoff between efficiency and reciprocity, and whether such allocation can be found through decentralized peer interactions.

3 WIRED NETWORKS WITH UPLOAD CONSTRAINT

In this section we provide background on prior results for this question in the case of a wired P2P network, under the usual assumption that the only bottleneck is the overall upload bandwidth μ_i from each peer i . In this case we can characterize the allowable resource sharing matrices as

$$\mathcal{Z} = \left\{ Z \in \mathbb{R}_+^{N \times N} \text{ satisfying (1), } \sum_j z_{ij} = \mu_i \ \forall i \right\}. \quad (8)$$

Here the vector s of total sending rates is fixed at $\mu = (\mu_i)$, and the overall transfer rate is $R(Z) = \sum_i \mu_i$ for all $Z \in \mathcal{Z}$: all allowable allocations are equally efficient.³

Therefore in this case there is no tradeoff, the remaining objective of reciprocity, in its global version, can be stated in terms of the following convex optimization:

Problem 1. *Given a connectivity matrix A and upload bandwidths $\mu = (\mu_i)$, find $Z \in \mathcal{Z}$ defined by (8) that minimizes $D(\mu||r(Z))$.*

We note again that if $r = \mu$ is feasible within \mathcal{Z} , it will be optimal; otherwise we are seeking a certain kind of approximation. An equivalent formulation (since μ is fixed) is

$$\max_Z \sum_j \mu_j \log(r_j(Z)), \text{ subject to } Z \in \mathcal{Z}.$$

In this version it can be interpreted as an instance of (weighted) *proportional fairness*, extensively studied in Internet resource allocation [14]. Here, we choose each node's weight as its own contribution to the network.

The set of solutions to Problem 1 can be characterized by Lagrangian duality; we recall some results in [25], based on a stream of related literature [20], [23]. All solutions Z^* to Problem 1 correspond to a unique vector $r^* = r(Z^*)$, characterized by a unique set of multipliers or prices $\lambda_i^* > 0$, $i = 1, \dots, N$, such that:

- $r_i^* = \lambda_i^* \mu_i$ for every peer. So λ_i^* defines the proportional reciprocity the peer receives from the network.

3. One could, instead, define \mathcal{Z} by an inequality constraint, but this deliberate inefficiency would serve no purpose, and will not be pursued.

- $z_{ji}^* > 0$ only for $j \in \arg \min\{\lambda_j^* : a_{ji} = 1\}$; furthermore, in this case $\lambda_i^* = [\lambda_j^*]^{-1}$. So at optimality a peer can only receive/send rate to another of *inverse* price.

Also in [25] is a detailed study of a prominent decentralized algorithm for reciprocity, proposed in [15], [23], [24]:

$$z_{ij}(t+1) = \mu_i \frac{z_{ji}(t)}{r_i(t)}. \quad (9)$$

In this *proportional reciprocity* scheme, peer i allocates to peer j the fraction of its bandwidth μ_i equal to the proportion of bandwidth received from peer j in the previous step. In matrix form we write $Z(t+1) = \mathcal{R}[Z(t)]$ by introducing the *reciprocity mapping*

$$\mathcal{R}[Z] := \text{diag}(\mu_i/r_i(Z)) \cdot Z^T.$$

Algorithm (9) is closely related to the so-called Sinkhorn procedure for matrix row and column renormalization [20]. A summary of its main properties is:

- Any solution Z^* of Problem 1 is a fixed point of \mathcal{R}^2 , square of the reciprocity mapping. Furthermore $Z^+ := \mathcal{R}[Z^*]$ is also a solution of Problem 1.
- In general, Z^+ need not be equal to Z^* , i.e. Z^* need not be a fixed point of the map \mathcal{R} itself. However the point $\tilde{Z} = \frac{Z^* + Z^+}{2}$ is another optimum and a fixed point of \mathcal{R} .
- In the special case where $r = \mu$ is feasible, there is always a symmetric optimal allocation.

The most important fact is the following convergence result.

Theorem 2 ([23], [25]). *Given an initial condition $Z(0) \in \mathcal{Z}$ with $z_{ij}(0) > 0$ whenever $a_{ij} = 1$, the sequence generated by (9) satisfies $\lim_{k \rightarrow \infty} Z(2k) = Z^*$, $\lim_{k \rightarrow \infty} Z(2k+1) = Z^+$, where both Z^* and Z^+ are optimal points of Problem 1. Furthermore, $r(Z(t))$ converges to the optimal rate vector r^* .*

Thus, provided initially all exchange options are explored, the even and odd subsequences converge to (possibly different) optimal allocations, and the fairness objective is achieved.

We finish the section by highlighting an additional fact: for any fixed point \tilde{Z} of \mathcal{R} , we have

$$\tilde{z}_{ij} = \frac{\mu_i}{\tilde{r}_i} \tilde{z}_{ji} \implies \frac{\tilde{z}_{ij}}{\tilde{z}_{ji}} = \frac{\mu_i}{\tilde{r}_i} \ \forall j,$$

the condition for equality in Lemma 1. We conclude that $D(\tilde{Z}||\tilde{Z}^T) = D(\mu||\tilde{r})$. Since as mentioned before there is always a fixed point of \mathcal{R} among the optima of Problem 1, we have the following consequence:

Corollary 3. *The minimum of $D(Z||Z^T)$ under $Z \in \mathcal{Z}$ defined by (8) has the same value as Problem 1, and a subset of its solutions.*

In other words, even if our objective is the looser *global* reciprocity each peer receives from the network, in this wired network case one can impose with no penalty the more stringent *peerwise* reciprocity.

4 WIRELESS NETWORKS: MULTIPLE RATES

We now move to consider a wireless network substrate, in which peers occupy certain spatial locations, connected by wireless channels. There are at least two differences between this situation and the wired case:

- (i) Wireless channels often adapt their rate to physical layer parameters such as signal-to-noise ratio, itself affected by distance. As a result, the sending rate will no longer be agnostic to the choice of receiving peer.
- (ii) Wireless links may interfere with each other. Therefore not all peers may transmit at the same time.

Remark 1. Another characteristic, which could partially offset the efficiency loss of point (ii), is that wireless nodes may broadcast, sending pieces at the same time to many receiving peers. This feature is not, however, easy to combine with the peer-to-peer philosophy of bilateral exchange. In particular, the choice of transmit rate, the piece to send, and the enforcement of reciprocity, all come into question under one-to-many transmissions. We will not pursue this direction further in this paper, leaving it open for future research.

In this section we focus on issue (i), postponing to the next section the consideration of interference. So for now we assume all peers have separate transmission channels, which they can allocate independently. Given two peers i and j , let μ_{ij} denote the maximum rate at which peer i can transmit to j , if it were to choose *only* this destination.

By time-sharing between destinations the peer can achieve the sending rates $z_{ij} = p_{ij}\mu_{ij}$ where $\sum_j p_{ij} = 1$. Here p_{ij} is the proportion of time devoted by peer i to neighbor j , again we assume no inefficient idle time. Therefore the achievable rate allocations are characterized by decoupled constraints

$$\mathcal{Z}_i := \left\{ Z_i = (z_{ij})_{j=1}^N, z_{ij} = 0 \text{ if } a_{ij} = 0, \sum_j \frac{z_{ij}}{\mu_{ij}} = 1 \right\} \quad (10)$$

on the rows of Z . The set \mathcal{Z} of achievable matrices is assimilated with the Cartesian product of the \mathcal{Z}_i ,

$$\mathcal{Z} = \left\{ Z \in \mathbb{R}_+^{N \times N} : Z_i \in \mathcal{Z}_i, i = 1, \dots, N \right\}. \quad (11)$$

Note that this is a generalization of (8), which corresponds to the special case $\mu_{ij} = \mu_i$ for all j , where the channel from peer i has the same quality for all destinations.

In general the matrix $M = (\mu_{ij})$ need not be symmetric: differences in peer channel qualities (e.g. transmission power) may cause $\mu_{ij} \neq \mu_{ji}$; indeed asymmetry was already present in the wired scenario.

We now look at a motivating example.

Example 1. Consider a wireless network with 3 peers which are all neighbors, and the matrix of maximum rates

$$M = \begin{bmatrix} 0 & 2 & 2 \\ 2 & 0 & 1 \\ 2 & 1 & 0 \end{bmatrix}.$$

Here maximum rates are symmetric, but not always uniform among outgoing links. We specify the time-sharing matrix P and the resulting rate allocation Z :

$$P = \begin{bmatrix} 0 & p & 1-p \\ q & 0 & 1-q \\ v & 1-v & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 0 & 2p & 2-2p \\ 2q & 0 & 1-q \\ 2v & 1-v & 0 \end{bmatrix}. \quad (12)$$

The set \mathcal{Z} of allowable allocations corresponds to all above matrices Z where p, q, v vary in the interval $[0, 1]$.

The total rate is $\sum_{i,j} z_{ij} = 4+q+v$, so efficiency is no longer agnostic to the peering choice: the set of efficient allocations is

$$\mathcal{Z}^{\text{eff}} = \left\{ Z = \begin{bmatrix} 0 & 2p & 2-2p \\ 2 & 0 & 0 \\ 2 & 0 & 0 \end{bmatrix}, p \in [0, 1] \right\}.$$

Also note that there are no symmetric matrices in \mathcal{Z}^{eff} , even though M is symmetric.

We now look at reciprocity, computing the vectors

$$s(Z) = \begin{bmatrix} 2 \\ 1+q \\ 1+v \end{bmatrix}, \quad r(Z) = \begin{bmatrix} 2(q+v) \\ 2p+1-v \\ 3-2p-q \end{bmatrix}.$$

It is easily checked that in this case $s = r$ is feasible (thus minimizing $D(s||r)$), achieved for $p = \frac{1}{2}$ and $q+v = 1$. Therefore global reciprocity is reached by the allocations in

$$\mathcal{Z}^{\text{rec}} = \left\{ Z = \begin{bmatrix} 0 & 1 & 1 \\ 2q & 0 & 1-q \\ 2(1-q) & q & 0 \end{bmatrix}, q \in [0, 1] \right\}.$$

Among them, one matrix (for $q = \frac{1}{2}$) is symmetric, thus achieving peerwise reciprocity ($D(Z||Z^T) = 0$).

The main observation is that $\mathcal{Z}^{\text{eff}} \cap \mathcal{Z}^{\text{rec}} = \emptyset$, one cannot satisfy both objectives simultaneously.

4.1 Trading off efficiency with global reciprocity

The example shows that there is a tradeoff between efficiency and reciprocity in wireless P2P settings. This suggests managing the tradeoff through a combined cost that contemplates both factors, such as

$$J(Z) = D(s||r) - \alpha R(Z) = \sum_i s_i \left[\log \left(\frac{s_i}{r_i} \right) - \alpha \right]. \quad (13)$$

Here the parameter $\alpha > 0$ weighs the importance assigned to efficiency. Note that $J(Z)$ is a convex function of Z , so its minimization over \mathcal{Z} is a convex optimization problem.

Problem 2. Given A and M , find $Z \in \mathcal{Z}$ defined by (11) that minimizes $J(Z)$.

Example 2 (Continuation of Example 1). We minimize the cost $J(Z)$ over matrices $Z(p, q, v)$ as in (12). We argue that it suffices to confine our search to $p = \frac{1}{2}$, $q = v$. This is because for any point (p, q, v) we can find another point $(p', q', v') = (1-p, v, q)$ with the same efficiency and reciprocity: $R(Z') = R(Z)$, and $D(s'||r') = D(s||r)$, in fact s', r' coincide with s, r modulo a permutation of the last two components. Therefore $J(Z') = J(Z)$. Invoking convexity of J , it can be no larger at the point $(Z + Z')/2$, which corresponds to $p = \frac{1}{2}$, $q = v$.

We thus consider the scalar valued function in $q \in [0, 1]$:

$$\begin{aligned} f(q) &= J(Z(1/2, q, q)) \\ &= 2 \log \left(\frac{1}{2q} \right) + 2(1+q) \log \left(\frac{1+q}{2-q} \right) - \alpha(4+2q). \end{aligned}$$

Minimizing $f(q)$ does not yield a closed form solution, but we find that the optimal q^* satisfies

$$\begin{cases} \frac{1}{2} < q^* < 1 & \text{if } 0 < \alpha < 2 + \log(2); \\ q^* = 1 & \text{if } \alpha \geq 2 + \log(2). \end{cases}$$

Thus if the weight α is large we just get the optimal efficiency solution. For moderate values of α we have a compromise with reciprocity which, however, always yields

$$Z(q^*) = \begin{bmatrix} 0 & 1 & 1 \\ 2q^* & 0 & 1 - q^* \\ 2q^* & 1 - q^* & 0 \end{bmatrix}$$

which is non-symmetric and with $s^* \neq r^*$ (since $q^* > \frac{1}{2}$).

Returning to the general case, we may attempt to solve Problem 2 through duality, which was a powerful method in the wired network situation. Writing the Lagrangian

$$\begin{aligned} L(Z, \lambda) &= J(Z) + \sum_i \lambda_i \left(\sum_j \frac{z_{ij}}{\mu_{ij}} - 1 \right) \\ &= \sum_i \left[s_i \log \left(\frac{s_i}{r_i} \right) - \lambda_i \right] + \sum_{i,j} z_{ij} \left[\frac{\lambda_i}{\mu_{ij}} - \alpha \right], \end{aligned}$$

leads after some analysis to the saddle point condition:

$$\lambda_i^* = \max_{\{j: a_{ij}=1\}} \left[-\log \left(\frac{s_i^*}{r_i^*} \right) + \alpha - 1 + \frac{s_j^*}{r_j^*} \right] \mu_{ij}.$$

In comparison to the conditions reviewed in Section 3 for the wired case ($\mu_{ij} = \mu_i$) we do not have here a clean interpretation for the optimal multipliers as reciprocity factors. And the preceding coupled transcendental equation does not suggest an immediate path for decentralization.

4.2 Trading off efficiency and peerwise reciprocity – best response algorithm

The decentralization objective motivates us to consider an alternative convex optimization problem:

Problem 3. Given A and M , find $Z \in \mathcal{Z}$ defined by (11) that minimizes the cost

$$\begin{aligned} E(Z) &= D(Z||Z^T) - \alpha R(Z) \\ &= \sum_{i,j: a_{ij}=1} z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha \right]. \end{aligned} \quad (14)$$

In the wired case of Section 3, it follows from Corollary 3 that the minimum of $E(Z)$ coincides with that of $J(Z)$ in (13) (in that case the throughput term is constant). In the wireless situation this is no longer true. Still, this alternative of trading off peerwise reciprocity with efficiency is a valid option to achieve our tradeoff in a decentralized way.

Grouping the cost in Problem 3 as

$$E(Z) = \sum_i \left\{ \sum_j z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha \right] \right\};$$

suggests the decentralized algorithm where peer i responds to the received rates z_{ji} by optimizing the term in braces over its decision variables z_{ij} :

Problem 4 (Best response iterate). For each i , and given z_{ji} for all j , solve

$$\min_{Z_i \in \mathcal{Z}_i} \sum_j z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha \right], \quad (15)$$

with \mathcal{Z}_i in (10).

To solve this problem for fixed $Z_i^T = (z_{ji})$ (i -th row of Z^T), construct the Lagrangian with multiplier λ_i for the constraint:

$$L_i(Z_i, Z_i^T, \lambda_i) = \sum_j z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha + \frac{\lambda_i}{\mu_{ij}} \right] - \lambda_i.$$

To minimize over Z_i for fixed λ_i (and Z_i^T), we impose

$$\frac{\partial L_i}{\partial z_{ij}} = \log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha + \frac{\lambda_i}{\mu_{ij}} + 1 = 0, \quad (16)$$

whose solution gives the reciprocity rule

$$z_{ij} = z_{ji} e^{\alpha - 1 - \frac{\lambda_i}{\mu_{ij}}}. \quad (17)$$

The value of λ_i can be found by imposing the constraint (10).

Remark 2. An important observation is that in the wired case ($\mu_{ij} = \mu_i \forall j$) we obtain in (17) $z_{ij} = \kappa_i z_{ji}$ for all j , namely a proportional allocation of upload rates as a function of rates received. After imposing the constraint we find $\kappa_i = \frac{\mu_i}{r_i}$ and therefore this solution is precisely the proportional reciprocity iteration (9), assuming we carry out these updates simultaneously for all rows (peers). We have thus re-interpreted this algorithm as the best response iteration in Problem 4.

Motivated by its good properties in the wired case, we test this best response generalization.

Example 3. We simulate in Matlab the best-response iteration for a 3-node network with maximal rate matrix

$$M = \begin{bmatrix} 0 & 3 & 2 \\ 3 & 0 & 1 \\ 2 & 1 & 0 \end{bmatrix}.$$

Figure 1 shows the trajectories of the cost $E(Z)$ in (14) when $Z(t)$ is updated using the best-response approach; the Gauss Seidel alternative will be described later. We see that the best response algorithm is unable to reach the optimum of Problem 3. This fact is not surprising, because even in the wired case the iteration (9) need not converge to optimality in $D(Z||Z^T)$, it can oscillate between two suboptimal points.

The convergence property that (9) did have was reaching the optimal global reciprocity $D(s||r)$. This suggests looking here at its counterpart, the cost $J(Z)$ in (13). However, in this wireless case the best-response iteration does not perform well either, as shown in Figure 2.

Given the failure to optimize when all peers perform simultaneous updates, one might consider the alternative of solving Problem 4 one peer at-a-time; we tested this alternative as well (see [18]) but it does not reach the optimum either. The proper way to carry out one at-a-time updates is now discussed.




Figure 1. Evolution of $E(Z)$ for different algorithms.




Figure 2. Evolution of $J(Z)$ for different algorithms.

4.3 Gauss-Seidel algorithm and its convergence

The Gauss-Seidel approach for solving a convex optimization problem where the variables are split in N components, consists of N successive updates, in each of which the global cost is optimized over one group of variables, with the others fixed. In our case, peers would in turns update their rows Z_i of the allocation matrix, which would naturally happen if they are not synchronized and use a common update interval.

Note, however, that they must take into account *all* cost terms affected by their decision, which includes others beyond those in (15). To address this, write the cost in Problem 3 as

$$E(Z) = E_i(Z_i, Z_i^T) + E_{-i}(Z_{-i}),$$

where

$$\begin{aligned} E_i(Z_i, Z_i^T) &:= \sum_{j: a_{ij}=1} \left[z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right) + z_{ji} \log \left(\frac{z_{ji}}{z_{ij}} \right) - \alpha z_{ij} \right] \\ &= D(Z_i \| Z_i^T) + D(Z_i^T \| Z_i) - \alpha \mathbf{1}^T Z_i, \end{aligned} \quad (18)$$

and the term E_{-i} depends *only* on the allocations of other peers, denoted collectively by Z_{-i} .

One step in the Gauss Seidel algorithm is given by:

Problem 5. For fixed i , and given Z_i^T minimize $E_i(Z_i, Z_i^T)$ over $Z_i \in \mathcal{Z}_i$.

This step requires the same information as the best-response version, namely the rates z_{ji} received from other peers, which is the basic assumption of any reciprocity scheme. Both can be computed numerically using convex optimization techniques.

Note that given Z_i^T , the function $E_i(Z_i, Z_i^T)$ is *strictly* convex in Z_i , thus Problem 5 has a unique solution Z_i . A study through duality leads to the saddle point condition

$$\frac{\partial \tilde{L}_i}{\partial z_{ij}} = \log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha + \frac{\lambda_i}{\mu_{ij}} + 1 - \frac{z_{ji}}{z_{ij}} = 0, \quad (19)$$

which is a modest variant of (16).

Clearly, the Gauss-Seidel iteration will compute a sequence $Z(t)$ with monotonically decreasing values of $E(Z(t))$; will it reach optimality? In the remaining trajectory of Figure 1 we show a simulation for Example 3, which indeed exhibits convergence to the optimum; this behavior is robust to initial conditions. The following statement claims this behavior is quite general.

Theorem 4. Consider the Gauss-Seidel iteration where $Z(t+1)$ is generated from $Z(t)$ by N successive steps of Problem 5, one row at a time. Let the initial condition $Z(0) \in \mathcal{Z}$ satisfy $z_{ij}(0) > 0$ whenever $a_{ij} = 1$. Then any limit point Z^* of $Z(t)$ is a global optimum of Problem 3.

A proof is given in Appendix A, by adapting the argument from a standard convergence result for Gauss-Seidel algorithms in [2]. That result applies to an optimization over Cartesian product, which is the case here as the constraints (10) are decoupled for each of the rows of Z . The cost is required to be continuously differentiable and convex, and strictly convex in each component of the Cartesian product when the others are held constant. These conditions hold here as well, except at the boundaries $z_{ij} = 0$, $z_{ji} = 0$, where our cost is not well-defined or tends to infinity.

By requiring that $Z(0)$ uses initially all available peering options, the initial problem is well defined and it is not hard to see that this property will be preserved, $z_{ij}(t) > 0$ whenever $a_{ij} = 1$. However in the limit the boundary can be approached, yielding $z_{ij}^* = 0$ (and necessarily, $z_{ji}^* = 0$) for a pair of neighbors; indeed this may be the optimal allocation, avoiding the use of a very inefficient link. We illustrate this situation with an Example.

Example 4. Consider a line network of 4 nodes, with maximal rate and resource sharing matrices M, Z below ($p, q \in [0, 1]$):

$$M = \begin{bmatrix} 0 & 2 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 2 & 0 \end{bmatrix}; \quad Z = \begin{bmatrix} 0 & 2 & 0 & 0 \\ 2p & 0 & 1-p & 0 \\ 0 & 1-q & 0 & 2q \\ 0 & 0 & 2 & 0 \end{bmatrix},$$

A simple analysis reveals that for any $\alpha > 0$ the optimal cost is achieved at $p = q = 1$, i.e. it is optimal not to use the links between peers 2 and 3. Starting our Gauss-Seidel algorithm from a random initial condition that uses all links, we observe convergence. Indeed after 5 rounds of updates the matrix is already using values of p and q of the order of $1 - 10^{-4}$, indicating the undesirable link is being turned off.

Extending the proof in [2] to accommodate this kind of phenomenon is non-trivial, see Appendix A.

Remark 3. We are not claiming that the alternative cost $J(Z)$ could be also minimized by this algorithm, indeed Figure 2 shows a small but nonzero gap, reflecting a difference between $D(Z||Z^T)$ and $D(s||r)$ at the convergence point. In the wireless case, the equality conditions in Lemma 1 may not hold at the optimum of Problem 3.

5 EFFICIENCY/RECIPROCITY TRADEOFF UNDER INTERFERENCE

We now consider a second characteristic of wireless local area networks: the use of a shared medium gives rise to *interference*, preempting certain links from being activated simultaneously. A typical case is when a common wireless channel is used in the network, so nodes interfere with all others within their range. Thus, medium access must be controlled either by centralized scheduling or by decentralized random access control.

The optimization of medium access in a wireless network, and its interaction with other layers such as routing or congestion control has been the subject of a large amount of research, in particular within the setting of network utility maximization, see [6], [12], [16] and references therein. In this section we will extend the methodology to pursue the efficiency/reciprocity tradeoff of interest.

A standard method of treating interference is to identify the *independent sets* of links, which can transmit simultaneously with no interference. Consistently with our method of analysis, we will represent each of these independent sets through a binary matrix $X \in \{0, 1\}^{N \times N}$ that identifies the active links. All allowable independent sets are specified by the class of configurations

$$\mathcal{X} = \{X^1, \dots, X^k, \dots, X^K\} \subset \{0, 1\}^{N \times N}.$$

At any time instant, only one of such configurations can be active. We remark that:

- Every $X \in \mathcal{X}$ has the hard zeros of the neighborhood structure, $x_{ij} = 0$ whenever $a_{ij} = 0$.
- Independent sets in \mathcal{X} need not be *maximal*: allowing the temporary use of inefficient configurations is often the only way to obtain tractable decentralized solutions.

As in the previous section we introduce time-sharing among configurations to achieve a richer mix of transmission rates. Let π_k denote the probability (fraction of time) configuration X^k is active. The possible time-shares are given by the unit simplex in \mathbb{R}^K ,

$$\Pi := \left\{ \pi \in \mathbb{R}^K : \pi_k \geq 0, \sum_{k=1}^K \pi_k = 1 \right\}.$$

Also for each $\pi \in \Pi$ denote by

$$P(\pi) = \sum_k \pi_k \cdot X^k. \quad (20)$$

the matrix of fractions of time p_{ij} each link is activated. Let \mathcal{P} denote the set of all possible such P matrices over $\pi \in \Pi$. \mathcal{P} is the convex hull of \mathcal{X} , and is thus convex.

Finally, we recall the matrix M of maximum transmission rates per link μ_{ij} ; the effective transmission rate at the link is $z_{ij} = p_{ij}\mu_{ij}$ as before. In matrix form $Z = M \circ P$, where \circ denotes the componentwise (Hadamard) product; the resulting set \mathcal{Z} of possible allocation matrices is also convex.

Remark 4. We can recast the situation of Section 4 in this more general context. In that case the only interfering links are those outgoing from the same peer i , who can only talk to one other peer at once. So we can choose a set \mathcal{X} of matrices with a single "1" per row, of the structure defined by A .⁴ The corresponding set \mathcal{P} are the row-stochastic matrices of structure A , and the set \mathcal{Z} coincides with the one in (11).

Within this new set \mathcal{Z} of allowable file-sharing matrices we now study tradeoffs between efficiency and reciprocity, through a reformulation of Problem 3:

Problem 6. Given a connectivity matrix A , a matrix of link bandwidths M , a family of independent sets \mathcal{X} , and the function $E(Z)$ defined in (14), solve

$$\begin{aligned} \min_{Z \in \mathbb{R}_+^{n \times n}} E(Z) \\ \text{subject to } Z = M \circ P, \\ P \in \mathcal{P} = \text{co}(\mathcal{X}). \end{aligned}$$

This is a convex optimization problem; below, we will seek distributed solution methods. We first discuss some properties of optimal allocations.

5.1 Symmetry in allocations

The class \mathcal{X} of independent set matrices is said to be *symmetric* if $X \in \mathcal{X} \implies X^T \in \mathcal{X}$; i.e. reversing all peer transmissions does not introduce interference. Note that:

- This does not mean the matrices X themselves are symmetric; links (i, j) and (j, i) will in many situations not be active at once.
- The structure of Section 4 is *not* symmetric. For instance in Example 1 we can have links 1-3 and 2-3 active at once, but not the other way round.
- If \mathcal{X} is symmetric (closed under transposition), then so is its convex hull \mathcal{P} .

Remark 5. An important case of a symmetric class \mathcal{X} occurs in the 802.11 (WiFi) standard when the request-to-send/clear-to-send option is activated. Here links (i, j) and (j, i) cannot be on at once. Now, activating a link (i, j) requires a free medium (measured by carrier-sense), and a bidirectional (RTS/CTS) handshake between nodes, which establishes that both the forward and backward links are free of interference. In that case, if $x_{ij} = 1, x_{ji} = 0$ is allowed, so is $x_{ij} = 0, x_{ji} = 1$, with the rest unchanged. This is a stronger condition than symmetry of \mathcal{X} .

The following result concerns optimal allocations for symmetric interference sets under the (restrictive) condition of symmetric maximal rates.

⁴ These are all maximal independent sets; in that situation where interference is within the node itself, there is no need to allow inefficient configurations.

Proposition 5. If \mathcal{X} is symmetric, and $M = (\mu_{ij}) = M^T$, then \mathcal{Z} is a symmetric set, and the optimum of Problem 6 is achieved at a symmetric matrix ($Z = Z^T$).

Proof. $Z^T = (M \circ P)^T = M^T \circ P^T = M \circ P^T$, so symmetry of \mathcal{Z} follows from that of \mathcal{P} . In addition we observe that $E(Z) = E(Z^T)$, since $R(Z) = R(Z^T)$ and

$$D(Z||Z^T) = \sum_{i>j} z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right) + z_{ji} \log \left(\frac{z_{ji}}{z_{ij}} \right) = D(Z^T||Z).$$

Then if Z is an optimal allocation, so is Z^T and by convexity the symmetric matrix $\frac{1}{2}(Z + Z^T)$ must also be optimal. \square

5.2 Dual decomposition for the optimal tradeoff

As a first attempt towards decomposing Problem 6 between the set of peer nodes, introduce the Lagrangian

$$L_0(Z, \pi, \Lambda) := E(Z) + \langle \Lambda, Z \rangle - \langle \Lambda, M \circ P(\pi) \rangle,$$

where Λ is a matrix multiplier with the same structure A , and $\langle \Lambda, Z \rangle = \sum_{ij} \lambda_{ij} z_{ij}$ is the standard matrix inner product. The last term above can be rewritten as

$$\begin{aligned} \langle \Lambda, M \circ P \rangle &= \sum_{ij} \lambda_{ij} \mu_{ij} \overbrace{\sum_k \pi_k X_{ij}^k}^{p_{ij}} \\ &= \sum_k \pi_k \underbrace{\sum_{ij} \lambda_{ij} \mu_{ij} X_{ij}^k}_{w_k(\Lambda)}, \end{aligned}$$

where $w_k(\Lambda)$ is the *weight* of the schedule $X^k \in \mathcal{X}$, obtained by adding the weights $\lambda_{ij} \mu_{ij}$ of the active links in the independent set. The vector $w(\Lambda) \in \mathbb{R}^K$ collects the weights of all independent schedules.

The dual function corresponding to Problem 6 is found by minimizing $L_0(Z, \pi, \Lambda)$ over the primal variables Z and π . This optimization decomposes between the two primal variables:

- 1) Minimization over $Z \in \mathbb{R}_+^{N \times N}$, with structure A , of

$$E(Z) + \langle \Lambda, Z \rangle = \sum_{i,j} z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha + \lambda_{ij} \right].$$

This problem is very similar to the one considered in Section 4, and could be tackled with similar means. Indeed, for fixed Λ a Gauss-Seidel iteration would provide a decentralized solution.

- 2) Maximization over $\pi \in \Pi$ of

$$\sum_k \pi_k w_k(\Lambda),$$

which amounts to the classical max-weight-scheduling [21]: concentrate the probability mass on the schedule $X^k \in \mathcal{X}$ with maximum weight w_k ; if there are weight ties between schedules, any convex combination thereof is optimal. This second problem is of combinatoric complexity, given the growth of the number of independent sets as a function of link nodes. This is the main limitation of the dual decomposition strategy.

Remark 6. Even if the complexity issue were not present, another undesirable feature of max-weight scheduling for this problem are oscillations around optimality. In the typical case where links (i, j) and (j, i) interfere, no single independent set can achieve reciprocity, so optimality requires allocating positive fractions of time p_{ij}, p_{ji} to both links. This can only be enabled by equalizing link weights, $\lambda_{ij} \mu_{ij} = \lambda_{ji} \mu_{ji}$. As the multiplier Λ varies close to its optimum, this tie is broken and max-weight schedules “chatter” very quickly. This phenomenon is common when using duality in non-strictly convex problems.

To avoid the above difficulties we turn to an approximate allocation, following an approach proposed in [12] for network utility maximization.

5.3 Markov approximation to optimal scheduling

In [12], Jiang and Walrand propose a distributed method based on random access to a wireless medium, which computes an approximation to the max-weight schedule. Other related references around the same time are [17], [19]. A clear and general presentation of this approach is given by Chen et al. in [5], which we largely follow here.

Given a vector of weights for each independent set, the idea is to select the following mix of schedules:

$$\hat{\pi}_k^\tau(\Lambda) := \frac{e^{\frac{w_k(\Lambda)}{\tau}}}{\sum_{l=1}^K e^{\frac{w_l(\Lambda)}{\tau}}}, \quad X^k \in \mathcal{X}, \quad (21)$$

where $\tau > 0$ is a “temperature” parameter. As $\tau \rightarrow 0+$, the above distribution $\hat{\pi}^\tau$ becomes concentrated on the schedules of maximum weight.

The interest of such probability distribution is that it can be computed in a decentralized way through a stochastic process, which goes by the names of Gibbs sampler [4] or Glauber dynamics [13]; when combined with a gradual “cooling” of τ it is called Simulated Annealing [11].

Specifically, construct a continuous-time Markov chain over the configuration space \mathcal{X} , with transition rates

$$q(X, X + e_{ij}) = e^{\frac{\lambda_{ij} \mu_{ij}}{\tau}} \mathbb{1}_{\{X + e_{ij} \in \mathcal{X}\}}, \quad (22a)$$

$$q(X, X - e_{ij}) = \mathbb{1}_{\{X - e_{ij} \in \mathcal{X}\}}. \quad (22b)$$

Here e_{ij} denotes the matrix with a single ‘1’ in entry (i, j) , so transitions only add one new link or turn off an existing one. Note that the exponent in (22a) amounts to the difference in the weights of the schedules X and $X + e_{ij}$, from where it follows that: $\hat{\pi}^\tau$ in (21) satisfies the detailed balance equations⁵

$$\hat{\pi}^\tau(X) q(X, X + e_{ij}) = \hat{\pi}^\tau(X + e_{ij}) q(X + e_{ij}, X);$$

this implies the chain is time-reversible and its steady-state distribution is $\hat{\pi}^\tau(\Lambda)$.

The above stochastic dynamics admit a fully decentralized implementation in network links, provided that the wireless network is endowed with carrier-sense multiple access (CSMA) to determine which link activations keep the

⁵ In these balance equations we temporarily change notation, dropping the dependence on Λ and writing $\hat{\pi}^\tau(X)$ instead of $\hat{\pi}_k^\tau$ for the schedule $X^k = X$.

schedule within the non-interfering class \mathcal{X} .⁶ In that case, each link can regulate the aggressiveness with which it tries to seize the medium in (22a), as a function of the current weight.

Remark 7. For this to be a viable resource allocation strategy implies that competition for the medium occurs at a much faster time-scale than the adaptation of the weights Λ . Fortunately, this is a natural situation in practical wireless networks, where for instance 802.11 medium access works at the millisecond scale, whereas the peer-to-peer reciprocity dynamics will typically occur at the scale of seconds; for more discussion see our simulations below.

Remark 8. In [26] (see also a summary in [18]), a Gibbs sampler was used for the entire resource allocation, building a Markov chain with transitions that reward efficiency and penalize an integrated measure of the reciprocity discrepancy. Here we confine the Markov dynamics to the scheduling component for given multipliers. The slower-scale resource allocation will be optimized below using convex techniques.

What is the relationship of this approximate schedule with the original optimization? As explained in [5], [12], it amounts to adding an entropy term to the weight maximization objective, of the form

$$H(\pi) := \sum_k \pi_k \log \left(\frac{1}{\pi_k} \right), \quad (23)$$

and solving for

$$\begin{aligned} & \max_{\pi \in \Pi} \langle \Lambda, M \circ P(\pi) \rangle + \tau H(\pi) \\ & = \max_{\pi \in \Pi} \sum_k \pi_k [w_k(\Lambda) - \tau \log(\pi_k)]. \end{aligned} \quad (24)$$

The above problem amounts to computing the Fenchel conjugate of the convex negative entropy function $-H(\pi)$ over the probability simplex, evaluated at the vector of weights $w(\Lambda)$. It is easily checked (see [3]) that it has the solution (21), and that the maximum in (24) is the log-sum-exp function of the weights:

$$\varphi^\tau(\Lambda) := \tau \log \left(\sum_k e^{\frac{w_k(\Lambda)}{\tau}} \right). \quad (25)$$

This is a convex function, a property that can be sharpened in our situation.

Lemma 6. If \mathcal{X} contains at least one schedule other than the $\{e_{ij}\}$, then $\varphi^\tau(\Lambda)$ is strictly convex in $\Lambda \in \mathbb{R}_+^{n \times n}$.

A proof is given in Appendix B. Note that the hypothesis holds in all cases except the one where all links interfere with each other *and* we do not allow the null independent set $X = 0$. Since in particular in the Markov chain approach the latter is always a possibility, the hypothesis is not restrictive.

⁶ Strictly speaking, CSMA may still give collisions due to the hidden node problem; this effect is considered negligible in [12]. It can also be avoided by enabling the request-to-send/clear-to-send option in 802.11.

Lemma 7. The gradient of $\varphi^\tau(\Lambda)$ (expressed as a matrix) is given by:

$$\frac{\partial \varphi^\tau}{\partial \Lambda} = \left[\frac{\partial \varphi^\tau}{\partial \lambda_{ij}} \right]_{i,j=1}^{N \times N} = M \circ P(\hat{\pi}^\tau(\Lambda)). \quad (26)$$

Proof. We compute the partial derivatives of φ^τ in (25) as follows. Fix a link (i, j) , then

$$\begin{aligned} \frac{\partial \varphi^\tau}{\partial \lambda_{ij}} &= \tau \frac{1}{\sum_k e^{\frac{w_k(\Lambda)}{\tau}}} \sum_k e^{\frac{w_k(\Lambda)}{\tau}} \frac{\partial w_k}{\partial \lambda_{ij}} \frac{1}{\tau} \\ &= \sum_k \hat{\pi}_k^\tau(\Lambda) \mu_{ij} X_{ij}^k \\ &= \mu_{ij} p_{ij}(\hat{\pi}^\tau(\Lambda)). \end{aligned}$$

□

5.4 A Primal - Dual approach to the optimal tradeoff

Armed with the above machinery, we are now ready to formulate a distributed method for a set of wireless nodes under interference to agree on (an approximation to) the optimal efficiency-reciprocity tradeoff. The approach parallels what was done in [5] for network utility maximization. One technical difference is that our primal cost $E(Z)$ is not strictly convex as defined, in contrast to utility functions which are assumed strictly concave.

Start by formulating the approximate tradeoff through the addition of the entropy term to Problem 6:

Problem 7. Given a connectivity matrix A , a matrix of link bandwidths M , a family of independent sets \mathcal{X} , and the function $E(Z)$ defined in (14), solve

$$\begin{aligned} & \min E(Z) - \tau H(\pi) \\ & \text{subject to } Z = M \circ P(\pi), \\ & \quad Z \in \mathbb{R}_+^{n \times n}, \quad \pi \in \Pi, \end{aligned}$$

where $P(\pi)$, $H(\pi)$ are given respectively by (20),(23).

The Lagrangian for this modified problem is

$$L(Z, \pi, \Lambda) = E(Z) + \langle \Lambda, Z \rangle - [\langle \Lambda, M \circ P(\pi) \rangle + \tau H(\pi)], \quad (27)$$

for which we seek a saddle point: minimum in (Z, π) , maximum in Λ . Again the primal variables appear suitably decomposed, however we will not optimize over both at the same time. Rather, consistently with the faster time-scale of the medium access, we will assume that for fixed Z and Λ the stochastic approach described above is used to optimize the second part of (27), resulting in the mean schedule (21), and the reduced Lagrangian

$$\begin{aligned} \hat{L}(Z, \Lambda) &:= \min_{\pi \in \Pi} L(Z, \pi, \Lambda) \\ &= L(Z, \hat{\pi}^\tau(\Lambda), \Lambda) \\ &= E(Z) + \langle \Lambda, Z \rangle - \varphi^\tau(\Lambda), \end{aligned} \quad (28)$$

with $\hat{\pi}^\tau$ from (21) and $\varphi^\tau(\Lambda)$ from (25). The following statement is proved in Appendix B:

Proposition 8. If (Z^*, π^*, Λ^*) is a saddle point of $L(Z, \pi, \Lambda)$, then (Z^*, Λ^*) is a saddle point of $\hat{L}(Z, \Lambda)$. Conversely, if (Z^*, Λ^*) is a saddle point of $\hat{L}(Z, \Lambda)$, then with $\pi^* = \hat{\pi}(\Lambda^*)$ from (21), (Z^*, π^*, Λ^*) is a saddle point of $L(Z, \pi, \Lambda)$.

We have thus reduced the optimization problem to finding (at the slower scale of rate adaptation) a saddle point of $\hat{L}(Z, \Lambda)$, which is convex in Z and strictly concave in Λ due to Lemma 6. To achieve this purpose in a decentralized way, we turn to a primal-dual dynamics, expressed below in continuous time for some fixed parameters $\beta > 0, \gamma > 0$.

$$\dot{Z} = -\beta \frac{\partial \hat{L}}{\partial Z} = -\beta \left[\Lambda + \frac{\partial E}{\partial Z} \right], \quad (29a)$$

$$\dot{\Lambda} = \gamma \frac{\partial \hat{L}}{\partial \Lambda} = \gamma \left[Z - \frac{\partial \varphi^\tau}{\partial \Lambda} \right] = \gamma [Z - M \circ P(\hat{\pi}^\tau(\Lambda))], \quad (29b)$$

where we have invoked (26). Writing these equations componentwise over the matrices, and recalling the definition of $E(Z)$ we have:

$$\dot{z}_{ij} = -\beta \left[\lambda_{ij} + \log \left(\frac{z_{ij}}{z_{ji}} \right) - \frac{z_{ji}}{z_{ij}} + 1 - \alpha \right]; \quad (30a)$$

$$\dot{\lambda}_{ij} = \gamma [z_{ij} - \mu_{ij} p_{ij}], \quad (30b)$$

where p_{ij} is the steady-state link utilization fraction obtained from the stochastic dynamics in the medium access (MAC) layer, which occurs at a faster time-scale. This means that the right-hand side of (30b) is proportional to the excess rate sent by the upper layers to the MAC queue, so λ_{ij} is proportional to the occupation of this queue, an attractive feature for a decentralized implementation.

In regard to (30a), as in previous sections the only information required to implement this gradient rule are the current reciprocity fractions $\frac{z_{ij}}{z_{ji}}$ of the exchange of node i with other peers j , something that can also be measured. The primal-dual dynamics is thus decentralized.

Proposition 9. *Solutions $(Z(t), \Lambda(t))$ of (29a-29b) converge asymptotically to a saddle point of $\hat{L}(Z, \Lambda)$. Thus, combined with the mapping $\pi(t) = \hat{\pi}^\tau(\Lambda(t))$ from (21), the trajectory $(Z(t), \pi(t))$ converges to a solution of Problem 7.*

This convergence follows from Lyapunov methods of a similar nature of those in [1], [10]. For completeness it is provided in Appendix B. Note that we do not require strict convexity of the primal cost, our damping is provided by strict convexity of Lemma 6.

5.5 Simulation Example

To illustrate the behavior of the above procedure we implemented in Matlab a discrete approximation to (29). Here it is convenient to choose $\gamma \gg \beta$, which helps convergence due to the stronger convexity in the variable Λ . In practical terms, it gives a higher priority to congestion (balance of the MAC queue) than achieving the correct efficiency/reciprocity tradeoff. For the MAC layer, we implemented a stochastic simulation of the Markov chain (22).

To set time-scales in somewhat realistic terms, we considered the following. For a WiFi network of ~ 100 Mbps total throughput, with packets of $\sim 10^4$ bits, the medium access rate is in the order of 10^4 packets (cycles of medium access) per second. We can let the Markov chain access the medium 100 times with fixed weights λ_{ij} , and still allow for update

steps every 0.01 seconds of the primal-dual dynamics. For these steps we choose $\gamma = 1, \beta = 0.01$; in this way prices react at the scale of 0.01 sec, whereas rates z_{ij} vary at the scale of 1 sec, consistent with a reasonable schedule of reciprocity (P2P network clients [7] commonly perform updates every 10 sec).

We tested these algorithms in the 4 node network of Figure 3. Nodes can exchange bidirectionally with their immediate neighbors, node 1 having a faster transmission rate than the rest. The corresponding matrix of maximum exchange rates is

$$M = \begin{bmatrix} 0 & 2 & 0 & 2 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix}.$$




Figure 3. Test network.

Assuming links interfere when they share a node among their source or destination, the independent sets are: the empty set, 8 single link schedules, and 8 two-link schedules which activate opposite sides of the square, in all possible directions. The temperature parameter was set at $\tau = 1$. The tradeoff parameter was set at $\alpha = 1$.







Figure 4. Evolution of rates z_{ij} .

Figure 5. Rates z_{ij} with idealized MAC layer.Figure 7. Multipliers λ_{ij} with idealized MAC layer.Figure 6. Evolution of multipliers λ_{ij} .

Figures 4 and 6 show the time trajectories of peer rates and the corresponding multipliers. For comparison purposes we also plot in Figures 5 and 7 a simulation of the primal-dual dynamics where the MAC layer is idealized, replacing the link probabilities by $P(\hat{\pi}^\tau)$ where $\hat{\pi}^\tau$ is given by (21). We observe a close match between the mean behavior of the stochastic dynamics and the deterministic counterpart: a damped oscillatory transient typical of this kind of primal-dual dynamics. The multipliers have a noisier behavior, consistent with their more aggressive update schedule.

Three values of link rates emerge in steady state: the largest for the faster links of node 1; an intermediate rate for nodes 2, 4 attempting to reciprocate to node 1; the lowest value for communications of nodes 2, 3, 4 amongst themselves. So we see a network that successfully negotiates the efficiency-reciprocity tradeoff.

6 CONCLUSION

In this paper we have investigated the distinct objectives of throughput efficiency and peer reciprocity which characterize peer-to-peer exchange, successively incorporating features of a wireless substrate: multiple-rate physical layers and interference. In the first case, we have developed a decentralized reciprocity mechanism that provably optimizes a tradeoff between overall throughput and peerwise KL divergence. In the interference case, the optimal tradeoff can be decomposed into the natural layers of rate allocation and link scheduling. The latter is difficult to solve, but a Markov approximation method was developed along the lines of [5], [12], leading to a primal-dual dynamics that converges to an approximation of the optimal tradeoff and can be implemented in a decentralized manner.

APPENDIX A. PROOF OF THEOREM 4

The sequence $Z(t)$ obtained from the Gauss-Seidel iteration belongs to a compact set \mathcal{Z} given in (11); consider a subsequence $Z(t_k)$ that converges as $k \rightarrow \infty$ to a limit point $Z^* \in \mathcal{Z}$. By definition $z_{ij}^* \geq 0$ for all links with $a_{ij} = 1$.

The cost $E(Z(t))$ is monotonically decreasing and lower bounded, so it has a limit. However since the cost includes the terms

$$d(z_{ij}, z_{ji}) = z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right) + z_{ji} \log \left(\frac{z_{ji}}{z_{ij}} \right)$$

which are not continuous at zero, it is important to isolate the links, if any, where the limit $z_{ij}^* = 0$. A first observation is that if $z_{ij}^* = 0$, then necessarily $z_{ji}^* = 0$. Otherwise $d(z_{ij}(t_k), z_{ji}(t_k)) \rightarrow \infty$ due to its second term, which implies $E(Z(t_k)) \rightarrow \infty$, contradicting the fact that it is decreasing.

Let us then classify the links in two classes $\mathcal{L} = \mathcal{L}^+ \cup \mathcal{L}^0$, those with positive and zero limits, denoting the respective incidence matrices by A^+ , A^0 , both symmetric, with $A = A^+ + A^0$. In general, for matrix $Z \in \mathcal{Z}$ we denote by Z^+ , Z^0 its projections over the respective link sets, i.e. the

Hadamard products $Z^+ = Z \circ A$, $Z^0 = Z \circ A^0$. We can also break down the cost (14) in two terms as

$$E(Z) = E^+(Z^+) + E^0(Z^0), \quad (31)$$

$$E^0(Z^0) := \sum_{i,j: a_{ij}^0=1} z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha \right], \quad (32)$$

$$E^+(Z^+) := \sum_{i,j: a_{ij}^+=1} z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha \right]. \quad (33)$$

Focusing on the subsequence t_k , we note that

$$\begin{aligned} \liminf_{k \rightarrow \infty} E(Z(t_k)) &= \liminf_{k \rightarrow \infty} [E^+(Z^+(t_k)) + E^0(Z^0(t_k))] \\ &= E^+(Z^*) + \liminf_{k \rightarrow \infty} E^0(Z(t_k)) \\ &\geq E^+(Z^*). \end{aligned} \quad (34)$$

The first step uses (31), the second follows since $Z^+(t_k) \rightarrow Z^*$, a point where $E^+(\cdot)$ is continuous. The last inequality is due to the fact that the efficiency terms in (32) vanish and the divergence part is non-negative.

The Gauss-Seidel iteration proceeds in rounds of length N , in which the rows are successively updated. Given $Z(t)$, denote by $Y^l(t)$ the matrix where the first l rows have been replaced by the corresponding ones for $Z(t+1)$. Note that cost is reduced in each step,

$$E(Z(t)) \geq E(Y^1(t)) \geq \dots \geq E(Y^l(t)) \geq \dots \geq E(Z(t+1)).$$

Consider now a modified Gauss-Seidel round $Y^{l,+}(t_k)$, $l = 1, \dots, N-1$ starting at the projection $Z^+(t_k)$, and optimizing only over the links in \mathcal{L}^+ with the portion $E^+(\cdot)$ of the cost.

Lemma 10. $\lim_{k \rightarrow \infty} E^+(Y^{l,+}(t_k)) = E^+(Z^*)$.

Proof of Lemma. We first note that from the monotonicity of the cost sequence

$$E^+(Z^+(t_k)) \geq E^+(Y^{1,+}(t_k)) \geq \dots \geq E^+(Y^{N-1,+}(t_k)),$$

we have the upper bound

$$\limsup_{k \rightarrow \infty} E^+(Y^{l,+}(t_k)) \leq E^+(Z^*), \quad l = 1, \dots, N-1. \quad (35)$$

To establish a lower bound, consider the first iterate $Y^{1,+}(t_k)$, and perturb it to a matrix $\tilde{Y}^1(t_k)$, as follows: include in the links $(1, j) \in \mathcal{L}^0$, if any, the term $\tilde{y}_{1j}(t_k) = z_{j1}(t_k)$ so that the corresponding divergence term disappears, and the contribution to the efficiency is of the order of

$$\delta_k := \max_{(1,j) \in \mathcal{L}^0} z_{j1}(t_k),$$

which goes to zero as $k \rightarrow \infty$.

To ensure the first row in $\tilde{Y}^1(t_k)$ stays within the feasibility constraint Z_i in (10) requires suitable reductions in some of the links $(1, j) \in \mathcal{L}^+$; however since these occur around a point $y_{1j}^{1,+} > 0$ where the cost is smooth, their contribution to the overall cost is also of the order of δ_k . This means that

$$E(\tilde{Y}^1(t_k)) = E^+(Y^{1,+}(t_k)) + O(\delta_k).$$

As constructed, $\tilde{Y}^1(t_k)$ is a feasible point in the first step of the *original* Gauss-Seidel iteration starting from $Z(t_k)$, which had optimum $Y^1(t_k)$. Hence

$$\begin{aligned} E^+(Y^{1,+})(t_k) + O(\delta_k) &\geq E(Y^1(t_k)) \\ &\geq E(Z(t_k+1)) \\ &\geq E(Z(t_{k+1})). \end{aligned}$$

Taking \liminf in the preceding expression as $k \rightarrow \infty$ leads to

$$\liminf_{k \rightarrow \infty} E^+(Y^{l,+}(t_k)) \geq \liminf_{k \rightarrow \infty} E(Z(t_{k+1})) \geq E^+(Z^*),$$

where we invoked (34). This, together with (35), establishes the lemma for $l = 1$. An analogous procedure can be used for the remaining (finite number of) Gauss-Seidel iterates $Y^{l,+}(t_k)$. \square

We can are now in similar conditions to those in the proof of Prop 3.8, Section 3.3.5 of Bertsekas and Tsitsiklis [2]. We have a sequence $Z^+(t_k)$ convergent to Z^* , a point where the reduced cost $E^+(\cdot)$ is continuous, and such that the resulting round of Gauss-Seidel iterates $Y^{l,+}(t_k)$ have the same limiting cost $E^+(Z^*)$.⁷

The proof in [2] can now be emulated directly with the above sequences, we provide only a sketch. One must show first that the change $Y^{1,+}(t_k) - Z^+(t_k)$ in the first G-S update goes to zero with k , by a contradiction argument that relies on the common asymptotic cost $E^+(Z^*)$, and the strict convexity of the function E^+ as a function of the first row. From here it follows that the common limit Z^* must be minimizing in the first component, when the others are held fixed, which gives a gradient condition

$$\langle \nabla_1 E^+(Z^*), Z_1 - Z_1^* \rangle \geq 0 \quad \forall Z_1 \in \mathcal{Z}_1.$$

Repeating the procedure in the following G-S updates provides analogous conditions for other components of the gradient $\nabla E^+(Z^*)$, implying Z^* must be an optimum of $E^+(\cdot)$. Under our convention (7) the term $E^0(Z^*) = 0$, so the overall cost is minimized.

APPENDIX B. PROOFS OF SECTION 5

Proof of Lemma 6

Proof. Without loss of generality set $\tau = 1$. The Hessian of the log-sum-exp function $f(w) = \log(\sum_k e^{w_k})$ is computed in [3] and shown to verify $\nabla^2 f \geq 0$; it follows also that its null-space $\ker[\nabla^2 f]$ has dimension 1, spanned by $\mathbf{1}$, the vector of all ones.

The function $\varphi(\Lambda)$ is the composition of f with the linear function $w(\Lambda)$. For convenience represent the latter as $w = B \text{vec}(\Lambda)$, where $\text{vec}(\Lambda)$ contains the λ_{ij} in vector format, and B is an incidence matrix between links and independent sets, multiplied by the diagonal matrix of the (μ_{ij}) . In this format we have

$$\nabla^2 \varphi = B^T \nabla^2 f B \geq 0.$$

⁷ A subtle point: we are *not* claiming that the G-S iteration from $Z^+(t_k)$, continued beyond the first round, is consistent with the next point $Z^+(t_{k+1})$ down the line. The proof of [2] invokes this fact *only* to argue that the intermediate G-S iterates have the same limiting cost. Here we have established this specifically in Lemma 10.

For a vector v to be in $\ker(\nabla^2\varphi)$ requires $Bv \in \ker[\nabla^2 f]$ and hence $Bv = c\mathbf{1}$, $c \in \mathbb{R}$. Now since \mathcal{X} contains the individual link schedules e_{ij} , then we must have $\mu_{ij}v_{ij} = c$ for every link. But by hypothesis there is at least one more independent set with a number of links $l \neq 1$. That schedule would give an entry lc in Bv , therefore it must be that $c = 0$, hence $v = 0$. We conclude $\nabla^2\varphi$ is positive definite and hence φ is strictly convex. \square

Proof of Proposition 8

Proof. For the direct implication, start with a saddle point (Z^*, π^*, Λ^*) of L . In particular π^* is the minimizer over this variable π , so by (28) we have $\hat{L}(Z^*, \Lambda^*) = L(Z^*, \pi^*, \Lambda^*)$.

Also, starting with $L(Z^*, \pi^*, \Lambda^*) \leq L(Z, \pi, \Lambda^*)$ and minimizing over π with fixed Z yields

$$\hat{L}(Z^*, \Lambda^*) \leq \hat{L}(Z, \Lambda^*) \quad \forall Z.$$

Allowing now variations in Λ we can write

$$\begin{aligned} \hat{L}(Z^*, \Lambda^*) &= L(Z^*, \pi^*, \Lambda^*) \geq L(Z^*, \pi^*, \Lambda) \\ &\geq \min_{\pi \in \Pi} L(Z^*, \pi, \Lambda) = \hat{L}(Z, \Lambda), \end{aligned}$$

establishing the saddle point condition in \hat{L} .

For the converse implication, start with a (min-max) saddle point (Z^*, Λ^*) of \hat{L} , and set $\pi^* = \hat{\pi}(\Lambda^*)$. We have

$$\begin{aligned} L(Z^*, \pi^*, \Lambda^*) &= \hat{L}(Z^*, \Lambda^*) \leq \hat{L}(Z, \Lambda^*) \\ &\leq L(Z, \pi, \Lambda^*) \quad \forall \pi \in \Pi. \end{aligned}$$

This shows the minimization part of the saddle condition in L . For the maximization part: note that since $\hat{L}(Z^*, \Lambda)$ is maximized in $\Lambda = \Lambda^*$, and strictly concave, we have the first order condition

$$0 = \frac{\partial \hat{L}}{\partial \Lambda}(Z^*, \Lambda^*) = Z^* - \frac{\partial \varphi^\tau}{\partial \Lambda}(\Lambda^*) = Z^* - M \circ P(\pi^*),$$

where we invoked (26). But then we see that $L(Z^*, \pi^*, \Lambda)$ is independent of Λ and thus maximized in particular at $\Lambda = \Lambda^*$. \square

Proof of Proposition 9

Proof. Consider any saddle point (Z^*, Λ^*) (min in Z , max in Λ) of $\hat{L}(Z, \Lambda)$. Define the Lyapunov function

$$V(Z, \Lambda) = \frac{1}{2\beta} \|Z - Z^*\|_F^2 + \frac{1}{2\gamma} \|\Lambda - \Lambda^*\|_F^2,$$

where $\|\cdot\|_F$ is the Frobenius norm, associated with our matrix inner product. Differentiating along trajectories of (29) gives

$$\begin{aligned} \dot{V} &= \frac{1}{2\beta} \langle Z - Z^*, \dot{Z} \rangle + \frac{1}{2\gamma} \langle \Lambda - \Lambda^*, \dot{\Lambda} \rangle \\ &= \langle Z^* - Z, \frac{\partial \hat{L}}{\partial Z}(Z, \Lambda) \rangle + \langle \Lambda - \Lambda^*, \frac{\partial \hat{L}}{\partial \Lambda}(Z, \Lambda) \rangle \\ &\leq [L(Z^*, \Lambda) - L(Z, \Lambda)] + [L(Z, \Lambda) - L(Z, \Lambda^*)] \quad (36) \\ &= \underbrace{[L(Z^*, \Lambda) - L(Z^*, \Lambda^*)]}_{\leq 0} + \underbrace{[L(Z, \Lambda^*) - L(Z, \Lambda^*)]}_{\leq 0}. \end{aligned} \quad (37)$$

Here the inequalities in (36) are first-order conditions for convexity in Z , concavity in Λ of $\hat{L}(Z, \Lambda)$. The inequalities in (37) follow from the saddle point assumption.

Thus the dynamics is stable in the broad sense. To analyze asymptotic stability, note first that if $\dot{V} = 0$ then all the above inequalities must be equalities. In particular, due to the strict concavity in Λ that follows from Lemma 6, this can only happen when $\Lambda = \Lambda^*$.

If a trajectory $(Z(t), \Lambda(t))$ moves entirely within the set $\{(Z, \Lambda) : \dot{V} = 0\}$, then $\Lambda(t) \equiv \Lambda^*$. But then $\dot{\Lambda} \equiv 0$ so from (29b) we have $Z(t) \equiv M \circ P(\hat{\pi}^\tau(\Lambda^*))$, and the point must be in equilibrium. By construction of the primal-dual dynamics, all equilibria are saddle points.

So the only invariant sets within $\{(Z, \Lambda) : \dot{V} = 0\}$ are equilibrium points; the LaSalle Invariance principle implies that all trajectories converge to one such saddle points, as claimed. \square

REFERENCES

- [1] K. Arrow, I. Hurwicz, and H. Uzawa, *Studies in linear and nonlinear programming*. Stanford: Stanford University Press, 1958.
- [2] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- [3] S. Boyd and L. Vanderberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [4] P. Brémaud, *Markov Chains: Gibbs Fields, Monte Carlo simulation and queues*. New York, NY: Springer, 1999.
- [5] M. Chen, S. C. Liew, Z. Shao, and C. Kai, "Markov approximation for combinatorial network optimization," *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 6301–6327, 2013.
- [6] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," in *Proc. of the IEEE*, vol. 95, Jan 2007, pp. 255–312.
- [7] B. Cohen, "Incentives build robustness in BitTorrent," in *Proc. 1st Workshop on the Economics of Peer-2-Peer Systems*, 2003, pp. 1–5.
- [8] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY: Wiley-Interscience, 1991.
- [9] B. Fan, J. C. Lui, and D.-M. Chiu, "The Design Trade-Offs of BitTorrent-Like File Sharing Protocols," *IEEE/ACM Trans. on Networking*, vol. 17, no. 2, pp. 365–376, 2009.
- [10] D. Feijer and F. Paganini, "Stability of primal-dual gradient dynamics and applications to network optimization," *Automatica*, vol. 46, no. 12, pp. 1974–1981, 2010.
- [11] B. Hajek, "Cooling schedules for optimal annealing," *Mathematics of Operations Research*, vol. 13, no. 2, pp. 311–329, 1988.
- [12] L. Jiang and J. Walrand, "A Distributed CSMA Algorithm for Throughput and Utility Maximization in Wireless Networks," *IEEE/ACM Transactions on Networking*, vol. 18, pp. 960–972, Jun. 2010.
- [13] L. Jiang, M. Leconte, J. Ni, R. Srikant, and J. Walrand, "Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling," *IEEE Transactions on Information Theory*, vol. 58, no. 10, pp. 6541–6555, 2012.
- [14] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 39, pp. 237–252, 1998.
- [15] D. Levin, K. LaCurts, N. Spring, and B. Bhattacharjee, "BitTorrent is an Auction: Analyzing and Improving BitTorrent's Incentives," *ACM SIGCOMM Computer Comm. Review*, vol. 38, pp. 243–254, 2008.
- [16] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1452–1463, Aug. 2006.
- [17] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. V. Poor, "Towards utility-optimal random access without message passing," *Wireless Communications and Mobile Computing*, vol. 10, no. 1, pp. 115–128, 2010.
- [18] F. Paganini, M. Zubeldía, and A. Ferragut, "Trading Off Efficiency and Reciprocity in Wireless Peer-To-Peer File Sharing," in *Proceedings of WiOpt, Mumbai, India*, 2015.

- [19] S. Rajagopalan and D. Shah, "Distributed algorithm and reversible network," in *Conference on Information Sciences and Systems*. IEEE, 2008, pp. 498–502.
- [20] R. Sinkhorn, "Diagonal equivalence to matrices with prescribed row and column sums. II," *Proceedings of the American Mathematical Society*, vol. 45, pp. 195–198, Aug. 1974.
- [21] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *Automatic Control, IEEE Transactions on*, vol. 37, no. 12, pp. 1936–1948, 1992.
- [22] M. Vojnovic and A. Proutiere, "Hop limited flooding over dynamic networks," in *Proc. of the IEEE INFOCOM*, 2011.
- [23] F. Wu and L. Zhang, "Proportional response dynamics leads to market equilibrium," in *Proc. of the 39th ACM Symposium on Theory of Computing (2007)*, 2007, pp. 354–363.
- [24] X. Yang and G. De Veciana, "Fairness, incentives and performance in peer-to-peer networks," *Performance Evaluation*, vol. 63, no. 3, pp. 175–194, 2006.
- [25] M. Zubeldía, A. Ferragut, and F. Paganini, "Neighbor selection for proportional fairness in P2P networks," *Computer Networks*, vol. 83, pp. 249–264, 2015.
- [26] M. Zubeldía, "From resource allocation to neighbor selection in peer-to-peer networks," Master's thesis, Universidad ORT Uruguay, 2014.



Fernando Paganini (M'90–SM'05–F'14) received his Electrical Engineering and Mathematics degrees from Universidad de la República, Montevideo, Uruguay, in 1990, and his M.S. and PhD degrees in Electrical Engineering from the California Institute of Technology, Pasadena, in 1992 and 1996 respectively. His PhD thesis received the 1996 Wilts Prize and the 1996 Clauser Prize at Caltech. From 1996 to 1997 he was a postdoctoral associate at MIT. Between 1997 and 2005 he was on the faculty the Electrical Engineering Department at UCLA, reaching the rank of Associate Professor. Since 2005 he is Professor of Electrical and Telecommunications Engineering at Universidad ORT Uruguay. Dr. Paganini has received the 1995 O. Hugo Schuck Best Paper Award, the 1999 NSF CAREER Award, the 1999 Packard Fellowship, the 2004 George S. Axelby Best Paper Award, and the 2010 Elsevier Scopus Prize. He is a member of the Uruguayan National Academy of Sciences and a Fellow of the IEEE. His research interests are control and networks.

cal Engineering Department at UCLA, reaching the rank of Associate Professor. Since 2005 he is Professor of Electrical and Telecommunications Engineering at Universidad ORT Uruguay. Dr. Paganini has received the 1995 O. Hugo Schuck Best Paper Award, the 1999 NSF CAREER Award, the 1999 Packard Fellowship, the 2004 George S. Axelby Best Paper Award, and the 2010 Elsevier Scopus Prize. He is a member of the Uruguayan National Academy of Sciences and a Fellow of the IEEE. His research interests are control and networks.



Martín Zubeldía received his B.S. degree in Electrical Engineering from Universidad ORT Uruguay in 2012, and his M.Sc. in Telecommunications Engineering from Universidad ORT Uruguay in 2014, working in the Mathematics Applied to Telecommunications (MATE) research group from 2010 to 2014. He is now a Ph.D. student in the Laboratory for Information and Decision Systems (LIDS), at the Massachusetts Institute of Technology.



Andrés Ferragut obtained his PhD degree in Electrical Engineering (Telecommunications) from Universidad de la República, Uruguay (2011). He has held teaching and research positions at the Mathematics and Electrical Engineering Depts. at Univ. de la República (2000 to 2009) and since 2007 he is affiliated with the Mathematics Applied to Telecommunications and Energy (MATE) research group, Universidad ORT Uruguay. He also spent two years (2004–2006) of graduate work at Télécom Paris, France.

In 2012, Dr. Ferragut was awarded the annual Uruguayan Engineering Academy prize for the best doctoral thesis in Electrical Engineering. His research interests are stochastic processes and queueing theory applied to the mathematical modeling of networks.