

Network utility maximization for overcoming inefficiency in multirate wireless networks.

Andrés Ferragut
 Universidad ORT &
 Universidad de la República
 Montevideo, Uruguay

José García
 Universidad ORT
 Montevideo, Uruguay

Fernando Paganini
 Universidad ORT
 Montevideo, Uruguay

Abstract—Wireless local area networks, in particular the ones based in the IEEE 802.11 standard, are nowadays used around the world to deliver Internet access, and are becoming increasingly prevalent. One of the crucial performance issues that these networks introduce is the possibility of having multiple transmission rates in the physical layer.

In this paper, we use the Network Utility Maximization framework to characterize the cross-layer interaction between the transport protocols such as TCP and the underlying MAC level rate adaptation. We describe the resource allocation imposed by current wireless networks in this framework, and characterize its equilibrium. Moreover, we propose alternative resource allocations that overcome the inefficiencies found in current protocols, and show simple mechanisms to impose more efficient equilibria in single cell scenarios. We also present simulations of these mechanisms in action, and discuss further generalizations to more complex networks.

I. INTRODUCTION

The Network Utility Maximization (NUM) framework originating in the work of Kelly [1], has been extensively applied in recent years for cross-layer optimization in wireless networks, (see [2], [3] and references therein). We can roughly classify this literature in terms of the multiple access technology considered. For the case of a *scheduled* MAC, the cross-layer optimization can be tackled through a dual decomposition approach, although the scheduling component is difficult. In the case of a random MAC, much of the effort has gone into modelling collisions; while this leads to non-convexities (see [2]), recent results [4], [5] have shown how to cast the problem in appropriate variables to approach the performance of the scheduled case.

We argue, however, that when considering wireless LAN technologies that are now prevalent in the world, the emphasis on collisions is misplaced. In recent versions of the IEEE 802.11 [6] wireless LAN standard, the loss of performance due to collisions is not as important as one might expect, due to two main reasons: on one hand, most of the traffic is downlink, and the Access Point (AP) does not collide with itself; secondly, with the 802.11 default parameters, collision probabilities and resolution time are low when compared with data transmission, whenever the number of stations is not too large. These conclusions follow from detailed models of the IEEE 802.11 Distributed Coordination Function (DCF), pioneered by [7]. Other aspects of the technology have a much larger impact, in particular the adaptation of physical

layer modulation rates (ARF [8]), as studied recently in [9]. Understanding cross-layer issues in this context is the main focus of this paper.

We begin in Section II by analyzing the effective rates the TCP layer can achieve when operating above a wireless MAC layer, specifically 802.11. This provides us with a characterization of the data rates offered by the MAC layer to the TCP layer, which is the basis for the rest of the analysis.

Secondly, in Section III we analyze the joint effect of using TCP above the multiple rates offered by the MAC layer. A first contribution of this paper is to pose this problem in the NUM framework, by taking into account the TCP behavior and the losses incurred in the AP buffer, in a single cell scenario. In our model, we exhibit the inefficiency that results from this arrangement, where users with high PHY rates are severely penalized, with minimal benefit for the slow users. This generalizes previous results ([9]) to fully take into account the TCP behavior.

In a second step in the analysis, we interpret the behavior in terms of a NUM problem with a natural constraint, but where the user utilities are inversely scaled with capacity. Based on this model, in Section IV we find ways of removing this bias, in particular we find a suitable price scaling that achieves global convergence to the unbiased NUM problem. We also explore ways to implement price scaling at a packet level, and we develop a new algorithm based on an appropriate Active Queue Management scheme. In Section V we present extensions of the analysis to a general network, involving possibly wired and wireless links. In Section VI we exhibit the performance of the proposed algorithms through simulations. Conclusions are given in Section VII.

II. IMPACT OF OVERHEADS IN TCP OVER 802.11.

The purpose of this section is to quantify the impact of protocol overheads when using TCP over 802.11. Using 802.11 implies that, when transmitting a packet of length L and due to the use of the Distributed Coordination Function, the Access Point (AP) must comply with a series of backoff and waiting times, as well as headers included by the PHY layer.

This means that the MAC layer offers a service to the upper layer consisting of a transmission rate $C_i \leq PHY_i$. This has

been analyzed before [7], [9], [10] and we will recall and extend this analysis here.

The time it takes to send this packet has a fixed component given by

$$T_i^0 := DIFS + H + \frac{L}{PHY_i} + SIFS + MAC_ACK_i, \quad (1)$$

that includes the time in the air and all overheads, plus a random number of time slots $K\sigma$, where $K \sim U\{0, \dots, CW\}$. In Table I we show typical values of these parameters for 802.11g.

Parameter	Value
Slot time σ	$9\mu s$
$SIFS$	$10\mu s$
$DIFS$	$28\mu s$
PLCP Header H	$24\mu s$
PHY_i	6Mbps ... 54Mbps
CW_{min}	15 slots
MAC_ACK	$24\mu s$

TABLE I
IEEE 802.11G PARAMETERS

We are interested in the average rate obtained by a station to study the upper layer effective rate. Observing that each packet is treated independently, the transmission times of successive packets form a renewal process, and the renewal reward theorem [11] tells us that in the long range the average rate is:

$$C_i^0 = \frac{L}{EK\sigma + T_i^0} = \frac{L}{\frac{CW_{min}}{2}\sigma + T_i^0}, \quad (2)$$

where we substituted K for its mean. We also took $CW = CW_{min}$ since we are modeling downlink traffic from the AP, which does not collide with itself. We also assume the appropriate PHY_i has been used so that one can neglect packet transmission errors. The denominator of the preceding expression (mean total time) is denoted by T_i .

In Table II we show the corresponding MAC level rates C_i^0 for the different PHY rates allowed in 802.11g with parameters as in Table I. Note the impact of overheads in the highest modulation rates.

When TCP connections are taken into account, another overhead must be considered: the TCP ACK packet. These packets were designed to have low impact on the reverse path, by having a length of 40 bytes. However, due to the overheads added by the MAC layer, the TCP ACK becomes non negligible, in particular at high modulation speeds. We assume that one TCP ACK is sent in the uplink direction for every TCP packet sent downlink. We will also assume that collision probabilities are low between downlink packets and the TCP ACKs. Under these assumptions, the TCP ACK packet introduces another overhead time in the system. The effective data rate then becomes:

$$C_i = \frac{L}{T_i + TCP_ACK_i} \quad (3)$$

where TCP_ACK_i is the average time to transmit a TCP_ACK packet and is given by:

$$TCP_ACK_i := DIFS + H + \frac{L_{ack}}{PHY_i} + SIFS + MAC_ACK_i, \quad (4)$$

where L_{ack} is typically 40 bytes. These effective data rates C_i are also shown in Table II. Note the strong impact of the TCP ACKs in the performance of the protocol, particularly at high modulation rates.

PHY rates	MAC rates (C_i^0)	Eff. data rate (C_i)
54	31.9	22.4
48	29.7	21.2
36	24.6	18.5
24	18.4	14.6
18	14.6	12.1
12	10.4	8.9
6	5.57	5.08

TABLE II
MAC RATES FOR THE CORRESPONDING PHY RATES OF 802.11G IN MBPS. $L = 1500$ BYTES.

In the following, we shall not address the impact of overheads and consider them given, since they are included in the standards. We can thus, for the purpose of modelling, concentrate in the C_i 's, which are the effective data rates at which packets from a TCP connection are served. In the following section we will model the behavior of the TCP connections above this MAC layer.

III. TCP RESOURCE ALLOCATION IN A MULTIRATE WIRELESS ENVIRONMENT

Let us consider first a single cell scenario where N users indexed by $i = 1, \dots, N$ are downloading data from a single AP.

Assume that each station establishes a downlink TCP connection, with sending rate x_i . Packets of these connections will be queued in the interface queue of the AP, before being put into the shared medium to reach their destination. We want to calculate the effective output rate y_i attained by each user. Assume that the Head of Line (HOL) probability for a packet of user i is proportional to the input rates x_i . Then, through a renewal reward argument similar to [9] we have:

$$y_i = \frac{p_{HOL,i}L}{\sum_j p_{HOL,j}L/C_j} = \frac{x_i}{\sum_j x_j/C_j} \quad (5)$$

which is similar to eq. (5) of [9], but without the collision terms.

To complete the loop, we must now model the TCP behavior, that determines the input rates x_i . Recall (c.f. [12]) that TCP Reno-like algorithms can be modelled as a primal controller:

$$\dot{x}_i = k(x_i)(U_i'(x_i) - p_i)$$

where $U(x)$ is an increasing and concave Utility function, p_i is the link loss rate (interpreted as a price) and $k(x_i) > 0$ a scaling factor. In the analysis, we will restrict ourselves

to the α -fair family of utility functions introduced in [13], which verify $U'(x) = Kx^{-\alpha}$ with $\alpha > 0$ a parameter which determines the compromise between efficiency and fairness of the allocation.

We can model the link loss rate as:

$$p_i = \left(\frac{x_i - y_i}{x_i} \right)^+ = \left(1 - \frac{1}{\sum_j x_j / C_j} \right)^+ = p$$

which is simply the proportion of packets that exceed the current service rate, where $(\cdot)^+ = \max(\cdot, 0)$ as usual. With this model the packet loss rate of each flow is the same and the complete dynamics follow:

$$\dot{x}_i = k(x_i)(U'_i(x_i) - p), \quad (6a)$$

$$p = \left(1 - \frac{1}{\sum_j x_j / C_j} \right)^+. \quad (6b)$$

We would like to characterize the equilibrium of this dynamics in terms of a NUM problem. For this purpose, consider the following function:

$$\Phi(x) = \sum_i \frac{x_i}{C_i} - 1 - \log \left(\sum_i \frac{x_i}{C_i} \right),$$

whenever $\sum_i \frac{x_i}{C_i} > 1$ and 0 otherwise.

Lemma 1: Φ is a convex function of x .

Proof: $\Phi(x)$ can be written as $\Phi(x) = g(f(x))$ where $f(x) = \sum_i x_i / C_i$ is a linear function of x and $g(u) = (u - 1 - \log(u))\mathbf{1}_{\{u > 1\}}$, where $\mathbf{1}$ represents the indicator function. It is easy to see that, for $u > 0$, $g'(u) = \max\{0, 1 - 1/u\}$ which is nonnegative and increasing function of u . Therefore, g is increasing and convex and thus Φ is convex [14]. ■

Consider now the following convex optimization problem:

Problem 1:

$$\max_x \sum_i \frac{1}{C_i} U_i(x_i) - \Phi(x) \quad (7)$$

We have the following:

Theorem 2: The equilibrium of the dynamics (6) is the unique optimum of Problem 1.

Let $V(x)$ denote the objective function in equation (7). The proof of the Theorem depends on the following lemma for $V(x)$, proved in the Appendix:

Lemma 3: The upper level sets of V , $\{x : V(x) \geq \gamma\}$ are compact and cover the positive orthant as $\gamma \rightarrow -\infty$.

Proof of Theorem 2: By Lemma 3 and the concavity of the objective function, there is a unique optimum for Problem 1, and it must satisfy the optimality conditions, namely:

$$\frac{1}{C_i} U'_i(x_i) - \frac{\partial}{\partial x_i} \Phi(x) = 0 \quad \forall i$$

By substituting Φ we have:

$$\frac{1}{C_i} \left[U'_i(x_i) - \left(1 - \frac{1}{\sum_j x_j / C_j} \right)^+ \right] = 0$$

Identifying the last term as $p = p(x)$ in (6), the optimality conditions become:

$$U'_i(x_i) - p = 0 \quad \forall i$$

which is the equilibrium condition of (6). ■

Theorem 4: The equilibrium of the dynamics (6) is globally asymptotically stable.

Proof: Consider $V(x)$ as a Lyapunov function of the system. Differentiating along trajectories:

$$\dot{V} = \nabla V \cdot \dot{x} = \sum_i \frac{k(x_i)}{C_i} (U'_i(x_i) - p)^2 \geq 0$$

So V is increasing along the trajectories. Moreover, $\dot{V} = 0$ only when $x = x^*$ the solution of Problem 1. Therefore, the upper level sets of V are forward invariant and since by Lemma 3, they cover the positive orthant, this equilibrium is globally asymptotically stable. ■

The function Φ in Problem 1 plays the role of a penalty function, and thus Problem 1 can be seen as an approximation of the following problem:

Problem 2 (Modified Network Problem):

$$\max_x \sum_i \frac{1}{C_i} U_i(x_i)$$

subject to the constraint:

$$\sum_i \frac{x_i}{C_i} \leq 1$$

which is the equivalent to the Network problem of [1], with two variants. The first one is that the constraint is rewritten in terms of x_i / C_i , the ‘‘time proportion’’ the shared medium is used by connection i . The sum of the allocated time proportions must be less than one and this is a natural constraint¹.

The second main difference with [1] is the scaling factor C_i for the user utility. This implies a negative bias to users with higher rates, since they would have less weight in the net utility. This bias leads to a radical equalization of rates to approximately the one of the slowest station, a well known effect in 802.11 environments (c.f. [9]).

Remark 1: In the case where all users share a common utility function (e.g. equal RTT TCP/Reno connections), the solution of problem Problem 2 reduces to

$$x_i^* = \frac{1}{\sum_j 1/C_j},$$

which is the harmonic mean of the data rates. This is in accordance with [9] where this rate is obtained as an upper bound on the realistic rate permanent connections experiment, and collisions are considered. As compared with [9], in our result the behavior of TCP is fully taken into account. As we mentioned before, we are disregarding collisions in the downlink traffic, and focusing on the issue of the multiple

¹This set can also be interpreted as the convex hull of the rates obtained by assigning all capacity to each of the nodes, which is equivalent to models used in scheduled networks [3].

rates. Moreover, note that this result is independent of the TCP flavor in use, so the MAC layer is actually determining the allocation.

Example 1: To see the biasing effect of Problem 2 consider the following example. Assume 3 users are downloading data from a single AP, and they have common utilities $U(x) = -\frac{1}{\tau^2 x}$ which model the TCP/Reno response, while their MAC layer rates are $C_i = 10$. In this case $x_i^* = 3.333$ for all three users. If user 3 for instance changes its radio conditions to $C_3 = 1$, the new allocation results:

$$x_1^* = x_2^* = x_3^* = 0.8333.$$

Note that the fastest destinations are heavily penalized due to the user 3 inefficiency².

Remark 2: We can analyze the system also from a connection-level perspective. If connections of rate C_i are not permanent but they arrive as a Poisson Process of intensity λ_i and have exponentially distributed workloads, in the case where all the utilities are the same, the system becomes a Discriminatory Processor Sharing queue [15]. In particular, jobs with effective rate C_i have weight $1/C_i$. This has been noted before in [10], [16]. The results of this section therefore generalize these previous results to the case where the utility functions are not equal for each connection. Note that in this case the queue is not a DPS queue anymore.

The previous remarks and examples suggest that ways to remove the bias in Problem 2 must be explored. In the following sections we propose alternative resource allocations for the single cell networks.

IV. A MORE EFFICIENT RESOURCE ALLOCATION FOR A SINGLE CELL

Based on the preceding discussion, it would seem natural to remove the bias in Problem 2 by considering:

Problem 3 (Wireless Network Problem):

$$\max_x \sum_i U_i(x_i)$$

subject to the constraint:

$$\sum_i \frac{x_i}{C_i} \leq 1,$$

Observe that the constraint here is similar to the ones discussed in [3], where the schedulable region of the system is defined in terms of the convex hull of the transmission rates. The purpose of this section is to analyze how to achieve the solution of Problem 3 without resorting to a complicated scheduling mechanism in the AP.

To see the difference with Problem 2, consider the following interesting property:

Proposition 5: If $U(x) = K \log(x)$ for all connections, then the equilibrium of Problem 3 is $x_i^* = C_i/n$. In particular, the allocated rate for user i depends only on its own effective rate and the total number of users.

²The result from Problem 1 would be $x_1^* = x_2^* = x_3^* = 0.89$, which shows that the barrier function approximation is very close.

This shows that imposing proportional fairness between users protects the fastest users from the lower rate ones. For instance, in Example 1, high rate users would be unaffected by the change in C_3 . When TCP Reno is in use, we do not have such protection but the situation is nevertheless better, as the following example shows.

Example 2: Assume we have the same situation of Example 1. When all three users have $C_i = 10$, the equilibrium of Problem 3 is the same as before, $x_i^* = 3.33$. When user 3 changes its radio conditions and $C_i = 1$, the equilibrium changes to $x_1^* = x_2^* = 1.93$, $x_3^* = 0.61$. We see then that in this case we can increase total network throughput by $\approx 80\%$ with respect to Example 1, and fastest users are not as heavily penalized.

We would like to drive the network to the equilibrium of Problem 3. For this purpose, consider the Lagrangian:

$$\mathcal{L}(x, p) = \sum_i U_i(x_i) - p \left(\sum_i \frac{x_i}{C_i} - 1 \right).$$

A simple primal-dual gradient algorithm to solve this optimization problem is:

$$\dot{x}_i = k(x_i) \left(U'(x_i) - \frac{p}{C_i} \right), \quad (8a)$$

$$\dot{p} = \left(\sum_i \frac{x_i}{C_i} - 1 \right)_p^+, \quad (8b)$$

where $k(x_i) > 0$ as before and $(\cdot)_p^+$ is the usual positive projection. It is well known [17], [18] that the trajectories of the dynamics given by (8) converge globally to the optimum of Problem 3.

In the wired case, dual algorithms have interpreted the price variable as the queueing delay [19], [20]. This is also the case here in this modified version. By integrating \dot{p} we see that p tracks the amount of time the shared medium is not capable of coping with the demands, and thus accumulating as delay in the queue.

More formally, let b_i denote the amount of data of connection i in the buffer (assume it is non empty). Then

$$\dot{b}_i = x_i - y_i$$

and the delay d is given by:

$$d = \sum_i \frac{b_i}{C_i}.$$

Therefore, recalling equation (5) we have:

$$\dot{d} = \sum_i \frac{\dot{b}_i}{C_i} = \sum_i \frac{x_i}{C_i} - 1.$$

Observe further that when all capacities are equal $C_i = C$ we recover the delay based model of [19], [20].

From equations (8) we see that in order to appropriately solve Problem 3, we need to scale the price to which the user reacts by the effective rate C_i . This makes sense since connections with higher rates use the medium more efficiently, and

thus should be charged less whenever this resource is scarce. Note however that this poses problems on implementation, because it prevents from using directly the queueing delay as the price. Moreover, the user endpoint has to be notified of the correct MAC level rate, which is infeasible.

A. The Multirate RED algorithm

In order to drive the system to the optimum of Problem 3, we propose to use a simple Active Queue Management policy which we call the Multirate RED algorithm (MRED).

Instead of using queueing delay as the price, we propose to use as a proxy the buffer length b , and to generate the price, the AP discards packets randomly with probability p_i proportional to $\frac{b}{C_i}$ for connection i . This gives a linear Random Early Detection (RED) algorithm, but with probabilities related to the effective data rates. Note that less packets will be dropped for connections with higher MAC rates. Moreover, this mechanism can be implemented in the AP resorting only to local information, such as destination address and current rate for this destination.

The closed loop dynamics for the proposed system is:

$$\dot{x}_i = k(x_i) (U'(x_i) - \kappa b / C_i), \quad (9a)$$

$$\dot{b} = \left(\sum_i x_i - y_i \right)_b^+ = \left(\sum_i y_i \right) \left(\sum_i \frac{x_i}{C_i} - 1 \right)_b^+ \quad (9b)$$

where $\kappa > 0$ is the proportionality constant of RED. These equations are similar to (8). In particular, in equilibrium, the x_i and $p = \kappa b$ will satisfy the KKT conditions of Problem 3. Stability results for these equations are harder to obtain, in Section VI we explore its behavior by simulation.

V. EXTENSION TO GENERAL NETWORK TOPOLOGY

In this section, we will discuss how to extend the previous analysis to the case where multirate wireless links and wired ones are present in the system. The main purpose of this section is to find a common price for the different type of network capacity constraints to which users should react in order to allocate resources properly. We identify three classes of capacity constraints:

- The classical wired constraints, $\sum_i x_i \leq c$.
- The wireless multirate constraints, $\sum_i \frac{x_i}{c_i} \leq 1$.
- The contention between cells constraints: $\sum_j \alpha_j \leq 1$, where α_j is the proportion of time each contending node is using the shared medium.

Consider then a network of J nodes, $j = 1, \dots, J$, connected by links $l = 1, \dots, L$. These links can be wired or wireless, and have an effective transmission rate c_l . In the case of wired links, c_l is the link capacity. In the wireless case, it is the effective data rate. Let $i = 1, \dots, n$ represent the connections, with rate x_i , and R the classical routing matrix: $R_{li} = 1$ if connection i traverses link l and 0 otherwise.

To represent the contention inherent to the wireless medium, we group the links l in contention sets: two links belong to the same contention set if they cannot be transmitting simultaneously and define $G_{kl} = 1$ if link l belongs to

contention set k and 0 otherwise; we call G the contention matrix.

The network capacity constraints can then be written as $Hx \leq \mathbf{1}$ where H is given by $H = GC^{-1}R$, with G, R previously defined, $C = \text{diag}(c_l)$ and $\mathbf{1}$ a column vector of ones.

To see how this framework enables us to model different situations, consider the following examples:

Example 3 (Wired network): If all links are wired, the contention matrix G is the Identity matrix. By taking R and C as before we recover the classical wired constraints $\sum_{i \in l} x_i \leq c_l$.

Example 4 (Single wireless cell): If there is only one wireless AP with N users in the cell, we can take R as the identity matrix, C as the wireless effective capacities and $G = \mathbf{1}^T$ (there is only one contention region where all links participate). We then recover the constraints discussed in Section IV.

Example 5 (Wireless Distribution System): To see a more complete example, consider the network composed of wired and wireless links shown in Figure 1. This topology appears in outdoor wireless distribution scenarios. We can model the capacity constraints of this network with the above framework by taking:

$$G = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}$$

$$C = \text{diag}(c, c_{AP_1}, c_{AP_2}, c_1, c_2, c_3, c_4)$$

$$R = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$




Fig. 1. Topology of a mixed wired-wireless distribution system with 4 end-users.

We can now pose the general Network Utility Maximization problem, which is:

Problem 4 (Wired-Wireless Network Problem):

$$\max_x \sum_i U_i(x_i)$$

subject to:

$$Hx \leq \mathbf{1}.$$

The previous problem seeks an optimal allocation within the natural constraints of the network, expressed in terms of allocated time. These constraints are equivalent to the ones used in the scheduling literature [3]. However, modelling the constraints as time proportions leads to a natural way to control the input rates of the sources without explicit scheduling for each destination in the nodes. something that is difficult to implement and requires significant message passing between them. In our approach, we assume that the buffers are served as FIFO and we control the input rates to the network to obtain a suitable time sharing.

Consider now the Lagrangian of Problem 4:

$$L(x, p) = \sum_i U_i(x_i) - p^T (Hx - \mathbf{1}) \quad (10)$$

where $p = (p_1, \dots, p_K)^T$ is the vector of prices. We see therefore that we have one price for each contention graph.

By denoting $q = H^T p$, the KKT conditions of Problem 4 are $U'_i(x_i) = q_i$ where q_i is given by:

$$q_i = \sum_{l:i \in l} \sum_{k:l \in k} \frac{p_k}{c_l} \quad (11)$$

Therefore, the connection must react to a price which is the sum of the prices of the contention graphs it traverses, divided by the corresponding link capacities.

Again, to solve Problem 4, we can use a primal-dual algorithm which gives the following dynamics:

$$\dot{x}_i = k(x_i) (U'(x_i) - q_i) \quad (12a)$$

$$\dot{p} = (Hx - \mathbf{1})_p^+, \quad (12b)$$

$$q = H^T p. \quad (12c)$$

These dynamics are globally asymptotically stable [17], [18] and its equilibrium is the solution of Problem 4. In this context, the prices track again the queueing delays at each node.

The remaining issue is whether this prices can be correctly generated and transmitted to the sources via the MRED implementation discussed in Section IV-A. Characterizing the topologies where decentralization is possible is, at the time of writing, ongoing work. However, we will present now two important examples in which this mechanism can be applied.

Example 6 (Mixed wired-wireless access network): A typical configuration for wireless access coverage is to distribute access points in non overlapping channels across the region to cover, and wire them to the Internet access node. This produces the tree topology of Figure 2. There, the APs are connected to a central switch, which is also connected to the router handling the Internet connection. End users are then connected to the

APs via 802.11 for example. In this case, each user traverses three contention graphs, one per link.




Fig. 2. Topology of a mixed wired-wireless access network.

Assuming the link capacities and user distributions shown in Figure 2, the corresponding price for user i is calculated according to equation 11 as:

$$q_i = \frac{p_{access}}{c_{access}} + \frac{p_{dist}}{c_{dist}} + \frac{p_{AP_i}}{c_i}$$

where p_{access} , p_{dist} and p_{AP_i} are the queueing delays suffered by the data packets of user i along the network, scaled by the corresponding link capacities.

By using the Multirate RED algorithm in each link in the network, we can therefore transmit this price to the source, and impose the notion of fairness of Problem 4 by emulating the dynamics of (12). We shall show that this mechanism indeed works in Section VI.

Example 7 (Wireless distribution system): A variation of the above example occurs when the area to cover is large, for instance, large outdoor deployments. In this case, the distribution links that connect each AP with the wired network are replaced by a wireless cell that backhauls all the APs and which is directly connected to the Internet router, as in Figure 1. The APs have 2 radio interfaces: one to connect to the backhaul link and one for local connectivity.

In this case, the corresponding price for user i can again be calculated according to 11 as:

$$q_i = \frac{p_{access}}{c_{access}} + \frac{p_{BH}}{c_{AP_i}} + \frac{p_{AP_i}}{c_i}$$

where now p_{BH} reflects the queueing delay in the backhaul node, and c_{AP_i} is the MAC layer rate at which the AP of user i connect to the backhaul node. Again, by using Multirate RED in this tree topology we can impose the notion of fairness of Problem 4.

VI. IMPLEMENTATION AND SIMULATIONS

As we discussed in section IV, the price to which a TCP connection should react in order to attain the equilibrium of

Problem 3 is the queuing delay. However, this price should be scaled by the effective data rate C_i the connection experiments in each link. Clearly, this is difficult to implement without resorting to scheduling. Moreover, typical TCP connections use loss based congestion control mechanisms, such as TCP Reno. Therefore, we propose to use the MRED algorithm developed in Section IV-A at each node to attain the optimum of Problem 3.

To test the proposal in a real environment, we implemented this algorithm in the Network Simulator `ns-2` [21]. Our implementation is based on the library `dei80211mr` [22]. Two important extensions were made to the library: the existing ARF mechanism was updated to cope with the possibility of a single node having different modulation rates for different destinations, which reflects the real behavior of current APs. The second modification was to implement the Multirate RED (MRED) queue, where the described early packet discard takes place.

Note that the cross-layer information needed for implementation of the mechanism is minimal: whenever a packet for next-hop j is received, it is discarded with probability $p_j = \kappa b / C_j$ where κ acts as a scaling parameter, b is the current queue length, and C_j is the corresponding effective rate for the current modulation rate the AP maintains with destination i (as in Table II). In the case of wired links, the link capacity is used to scale this drop probability. The non-dropped packets are served then on a FIFO basis.

We now present several simulation scenarios to illustrate the behavior of the proposed algorithm.

A. Single-cell scenario

We simulate the topology shown in Figure 3, which consists of a single cell 802.11g scenario in which 3 users are connected with a modulation rate $PHY_i = 54Mbps$, and some time later, a fourth user is added at the lowest possible modulation $PHY_5 = 6Mbps$. All four connections use TCP/Newreno and share equal Round Trip Times (RTTs), then having similar utility functions.




Fig. 3. Topology of a single-cell scenario.

For these modulation rates, the effective data rates according to Table II are $C_i \approx 22.4Mbps$, $i = 1, 2, 3$ and $C_4 \approx 5.1Mbps$. In the first graph of Figure 4, we see that all connections converge to the same throughput, which is approximately $x^* = 3.0Mbps$, the harmonic mean discussed in Remark 1. In the second graph, we show the behavior of the system under the MRED algorithm. In this case the allocation converges approximately to $x_i^* = 4.2Mbps$, $i = 1, 2, 3$ and $x_4^* = 2.1Mbps$, which is the exact solution of Problem 3. Note that the total throughput in the network is increased by more than 20%.




Fig. 4. Comparison between throughputs: without MRED (above), with MRED (below).

B. Wired-wireless scenario

The purpose of this example is to show that Problem 2 captures the behavior of the system when the TCP connections have different RTTs, and thus different utilities, and to show how efficiency can also be improved in this case with the MRED algorithm.

We consider the topology of Figure 5, where two connections with different RTTs share a wireless bottleneck link. In this example, connection 1 has twice the RTT of connection 2, and its station is closer to the AP, having a modulation rate $PHY_1 = 54Mbps$. The second connection has a modulation




Fig. 5. Wired-wireless topology.





Fig. 6. Wired-wireless topology simulation. Above: original allocation. Below: MRED algorithm.

rate of $PHY_2 = 6Mbps$. Both connections use TCP/Newreno, which we model by the utility function $U(x) = -1/(\tau^2 x)$ with τ the connection RTT.

Plugging these values into Problem 2 using the effective data rates of Table II, the allocation results $x_1^* = 2.25Mbps$ and $x_2^* = 4.5Mbps$. In the first graph of Figure 6 we show the results of the corresponding simulation, which shows that indeed the connection throughputs converge approximately to

the values predicted by Problem 2.

By using MRED in the AP we can change the allocation to the one proposed in Problem 3, removing the bias of Problem 2. The resulting allocation is $x_1^* \approx x_2^* \approx 4.1Mbps$. In the second graph of 6 we show the corresponding simulation results. We see that the MRED algorithm approximately drives the system to the new equilibrium. Note that this new equilibrium is $\approx 20\%$ more efficient than the preceding one.

C. Tree topology

In this example, we simulate the topology of Figure 2 with two distribution APs and two users in each AP. The access link capacity is $c_{access} = 20Mbps$ representing a typical access capacity. The distribution links have $c_{dist} = 100Mbps$ and thus are overprovisioned. The wireless cells are identical and have each one two users, with modulation rates $PHY_1 = PHY_3 = 54Mbps$ and $PHY_2 = PHY_4 = 6Mbps$. Each user has a single TCP connection and all connections have equal RTTs.

Plugging these values in Problem 3 gives the following allocation:

$$x_1^* = x_3^* = 6.5Mbps \quad x_2^* = x_4^* = 3.5Mbps$$

By using the MRED algorithm as discussed in Section V, we can drive the system to this allocation. Results are shown in Figure 7, where we see that the throughputs approximately converge to the above equilibrium.




Fig. 7. Throughputs of TCP connections for a wireless access scenario with 4 users. MRED is in use.

VII. CONCLUSIONS

Throughout this paper, we applied the Network Utility Maximization framework to characterize the cross-layer interaction between the TCP transport protocol with an underlying MAC where multiple modulation rates coexist. This situation is present in typical IEEE 802.11 deployment scenarios. We analyzed the impact of overheads in the throughput of TCP connections, and then we described the resource allocation

imposed by current wireless networks in this framework, characterizing its equilibrium through a suitable NUM problem.

In the second part of the paper, we proposed an alternative resource allocation that generalized the fairness and efficiency notions of TCP in wired networks to this context. This new resource allocation overcomes the inefficiencies found in current protocols. We also showed a simple mechanism to impose these more efficient equilibria in single cell scenarios and also showed possible generalizations of this procedure to more complex topologies. Finally we validated the results by simulations that show that the proposed algorithm indeed drives the system to the desired allocations.

In future work, we plan to extend the proposed mechanisms to the new 802.11 additions, where these issues may become more important due to the higher data rates involved, and due to packet aggregation mechanisms.

APPENDIX

Proof of Lemma 3: Denote by $V(x) = \sum_i \frac{1}{C_i} U_i(x_i) - \Phi(x)$ the objective function of Problem 1. We analyze the case $0 < \alpha < 1$, where we have the following bound:

$$\begin{aligned} \sum_i \frac{1}{C_i} U_i(x_i) &= \sum_i \frac{K_i x_i^{1-\alpha}}{C_i (1-\alpha)} \\ &= \sum_i \frac{K_i C_i^{-\alpha}}{1-\alpha} \left(\frac{x_i}{C_i} \right)^{1-\alpha} \\ &\leq K \left(\max_i \frac{x_i}{C_i} \right)^{1-\alpha} \\ &\leq K \left(\sum_i \frac{x_i}{C_i} \right)^{1-\alpha}. \end{aligned}$$

Here $K = \max_i \left\{ \frac{K_i C_i^{-\alpha}}{1-\alpha} \right\}$. Let $y = \sum_i x_i / C_i$ and assume $y > 1$. Using the previous bound we have:

$$V(x) \leq K y^{1-\alpha} - y + 1 + \log(y) = h(y)$$

and since $\alpha > 0$, the function $h(y) \rightarrow -\infty$ when $y \rightarrow \infty$ so for any given γ we can find y_γ such that $g(y) < \gamma$ for all $y > y_\gamma$. Therefore, for all $x \in E_\gamma = \{ \sum_i x_i / C_i > y_\gamma \}$ we have that $V(x) < \gamma$. This proves that the level sets of V , $\{V \geq \gamma\}$, are contained in the complement of the set E_γ , which is compact. Since γ is arbitrary, we also have that $V(x) \rightarrow -\infty$ when $\|x\| \rightarrow \infty$.

The case with $\alpha = 1$ can be proven analogously. For $\alpha > 1$, the utilities are upper bounded and $\Phi(x) \rightarrow \infty$ whenever $\|x\| \rightarrow \infty$ so the result is trivial. ■

ACKNOWLEDGMENTS

This work was partially supported by LATU/Ceibal, CSIC I+D (Universidad de la República), ANII-FCE and AFOSR-US.

REFERENCES

- [1] 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.
- [2] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," in *Proceedings of the IEEE*, vol. 95, no. 1, Jan 2007, pp. 255–312.
- [3] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE Journal on Selected Areas in Communication*, pp. 1452–1463, Aug. 2006.
- [4] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," in *Proceedings of the Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing*, Sep. 2008.
- [5] A. Proutière, Y. Yi, and M. Chiang, "Throughput of random access without message passing," in *Proceedings of the 44th Conference on Information Science and Systems (CISS 08)*, Apr. 2008.
- [6] "IEEE 802.11-2007, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," <http://www.ieee802.org/11/>, Jun. 2007.
- [7] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, Mar. 2000.
- [8] A. Kamerman and L. Monteban, "Wavelan-II: a high performance wireless lan for the unlicensed band," *Bell Labs Technical Journal*, vol. 2, no. 3, pp. 118–133, Aug. 1997.
- [9] A. Kumar, E. Altman, D. Miorandi, and M. Goyal, "New insights from a fixed-point analysis of single cell IEEE 802.11 WLANs," *IEEE/ACM Transactions on Networking*, vol. 15, no. 3, pp. 588–601, Jun. 2007.
- [10] A. Ferragut and F. Paganini, "A connection level model of IEEE 802.11 cells," in *Latin American Networking Conference*, Sep. 2009.
- [11] W. Feller, *An Introduction to Probability Theory and Its Applications*. New York: John Wiley & Sons, 1965.
- [12] R. Srikant, *The Mathematics of Internet Congestion Control*. Boston, MA: Birkhäuser, 2004.
- [13] J. Mo and J. Walrand, "Fair end-to-end window based congestion control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, Oct. 2000.
- [14] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [15] E. Altman, K. Avrachenkov, and U. Ayesta, "A survey on discriminatory processor sharing," *Queueing Systems: Theory and Applications*, vol. 53, no. 1-2, pp. 53–63, Jun. 2006.
- [16] T. Bonald and A. Proutière, "Flow-level stability of utility-based allocation in non-convex rate regions," in *Proceedings of the 40th Annual Conference on Information Systems and Sciences (CISS 2006)*, 2006, pp. 327–332.
- [17] K. Arrow, L. Urwicz, and H. Uzawa, *Studies in Linear and Nonlinear Programming*. Stanford, CA: Stanford University Press, 1958.
- [18] D. Feijer and F. Paganini, "Krasovskii's method in the stability of network control," in *Proceedings of the American Control Conference*, Jun. 2009.
- [19] S. H. Low and D. Lapsley, "Optimization flow control, I: basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, 1999.
- [20] S. H. Low, F. Paganini, and J. C. Doyle, "Internet congestion control," *IEEE Control Systems Magazine*, vol. 22, no. 1, pp. 28–43, Feb. 2002.
- [21] "The Network Simulator NS-2," <http://www.isi.edu/nsnam/ns/>.
- [22] "dei80211mr: a new implementation of 802.11 for ns2," <http://www.dei.unipd.it/wdyn/?IDsezione=5090>.