

Flow-level Stability of Utility-Based Allocations for Non-Convex Rate Regions

T. Bonald and A. Proutière

France Telecom R&D / Ecole Normale Supérieure, Paris

Abstract—We investigate the stability of utility-maximizing allocations in networks with arbitrary rate regions. We consider a dynamic setting where users randomly generate data flows according to some exogenous traffic processes. Network stability is then defined as the ergodicity of the process describing the number of active flows. When the rate region is convex, the stability region is known to coincide with the rate region, independently of the considered utility function. We show that for non-convex rate regions, the choice of the utility function is crucial to ensure maximum stability. The results are illustrated on the simple case of a wireless network consisting of two interacting base stations.

Index Terms—Resource allocation, maximum stability.

I. INTRODUCTION

Designing fair and efficient resource allocations is a central issue in data networks. Since the seminal work of Kelly [8], optimization approaches have been extensively used to this end: an allocation is said to be optimal if it maximizes the overall “utility” of the set of active data flows.

Formally, consider a network whose resources are shared by a set of data flows. We identify each flow through its “class”, which defines the resources it requires for the transfer of its packets. There is an arbitrary set of K flow classes. Denote by $x = (x_1, \dots, x_K)$ the number of active flows of each class and by $\phi = (\phi_1, \dots, \phi_K)$ the total long-term throughput of each class, i.e., the throughput of each class- k flow is ϕ_k/x_k . Now given some increasing and strictly concave function U representing the utility of a flow as a function of its throughput, an allocation maximizes the overall network utility in state x if ϕ solves the optimization problem:

$$\max \sum_k x_k U(\phi_k/x_k), \quad \text{subject to } \phi \in \mathcal{R}, \quad (1)$$

where \mathcal{R} is some compact subset of \mathbb{R}_+^K representing the rate region, that is the set of all vectors of achievable throughputs. The allocation is unique if the rate region \mathcal{R} is convex. Usual allocations are the α -fair allocations [14] based on the utility functions $U_\alpha(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha)$ for all $\alpha > 0$, $\alpha \neq 1$, and $U_1(\cdot) = \log(\cdot)$. The parameter α measures the degree of fairness of the allocation [15]: the total throughput tends to be maximized when $\alpha \rightarrow 0$ but the allocation is then very unfair; $\alpha = 1$ gives proportional fairness [8]; $\alpha = 2$ corresponds to the minimum potential delay allocation [12]; the limiting case $\alpha \rightarrow \infty$ leads to the most fair allocation, namely max-min fairness [3].

The main advantage of optimization approaches is that standard decomposition techniques lead to congestion control


algorithms that realize the corresponding allocations in a decentralized way [8],[10],[14]. Most studies in this field of research consist in analyzing the performance of these algorithms *at packet level*, assuming a fixed number of active flows. In practice, data flows do not last for ever but arrive at random times and leave the network once the last packet of the corresponding document has been received. This results in a dynamic number of active flows. A key requirement for an allocation is to ensure maximum stability *at flow level* in the sense that the number of active flows remains finite for the highest traffic intensity.

In the following, we assume that class- k flows are generated according to a Poisson process of intensity λ_k and have i.i.d. exponential sizes of mean σ_k bits¹; the corresponding traffic intensity is $\rho_k = \lambda_k \sigma_k$ bit/s. We are interested in the stability region of utility-based allocations, that is the largest set of vectors $\rho = (\rho_1, \dots, \rho_K)$ such that the network is stable. This issue has so far been addressed for convex rate regions only. The stability region then coincides with the rate region, independently of the considered utility function (see below). We here consider an arbitrary rate region (possibly discrete) and show that the stability region is generally larger than the rate region and depends on the considered allocation. Specifically, we prove that:

- the maximum stability region is the smallest convex, coordinate convex set containing \mathcal{R} ;
- the stability region of an α -fair allocation depends on α : it is maximum when $\alpha \rightarrow 0$ and minimum when $\alpha \rightarrow \infty$. In particular, the most fair allocations cannot guarantee maximum stability.

We illustrate the results on the simple case of a wireless network consisting of two interacting base stations, whose rate region is naturally non-convex.

Notation: For any $\mathcal{S} \subset \mathbb{R}^K$, we denote by $\check{\mathcal{S}}$ the largest open subset of \mathcal{S} and by $\bar{\mathcal{S}}$ the smallest convex, coordinate-convex set containing \mathcal{S} (see below for a discrete set \mathcal{S}).



¹These technical assumptions allow us to describe the network state x as a Markov process; the results remain valid for more realistic traffic assumptions but require a much more complex proof, cf. [11].

II. CONVEX RATE REGIONS

We first recall stability results for convex rate regions². A typical example is the rate region of a wireline network. The class of a flow is then defined by the set of links on its route through the network. The rate region \mathcal{R} is the polytope of vectors ϕ such that $\phi A \leq C$, where A is the routing matrix ($A_{kl} = 1$ if class- k flows go through link l , $A_{kl} = 0$ otherwise) and C is the vector of link capacities. Many other examples leading to convex rate regions are presented in [6].

The flow-level stability of utility-based allocations was investigated for wireline networks in [16],[5],[17] and for any convex rate region in [6]. The results are derived under the so-called *time-scale separation* assumption: the congestion control algorithms converge so fast compared to the flow-level dynamics that the allocation ϕ can be assumed to be given by (1) at any time. Recently, the stability analysis was extended to scenarios where the time-scale separation assumption does not hold [9]. The following theorem summarizes the above mentioned stability results in the context of α -fair allocations.

Theorem 1: For any convex rate region, the maximum stability region is equal to the rate region and is achieved by all α -fair allocations.

It is worth noting that the stability region does not depend on the fairness parameter α that characterizes the allocation, provided $\alpha > 0$. When $\alpha = 0$, the allocation (which is not unique) maximizes the total throughput but does not ensure maximum stability [5]. This is due to the fact that those flows that consume less resources are naturally given priority, leading to the starvation of other flows. The stability region is then strictly included in the rate region.

III. MAXIMUM STABILITY FOR ARBITRARY RATE REGIONS

We now consider an arbitrary rate region. Wireless networks where base stations interact through interference typically have non-convex rate regions, see Section V. A first issue is to determine the maximum stability region, that is the set of traffic vectors ρ such that there exists an allocation stabilizing the network.

Theorem 2: For an arbitrary rate region \mathcal{R} , the maximum stability region is $\bar{\mathcal{R}}$.

Proof. If $\rho \notin \bar{\mathcal{R}}$, we proved in [6] that the network is unstable for any allocation. Now assume that ρ belongs to the largest open subset of $\bar{\mathcal{R}}$. An allocation that stabilizes the network is given by the so-called MaxProjection policy [2]: for any state x , ϕ solves the optimization problem

$$\max \langle x, \phi \rangle, \quad \text{subject to } \phi \in \mathcal{R},$$

where $\langle \cdot, \cdot \rangle$ is the usual scalar product on \mathbb{R}^K . \square

The rest of the paper is devoted to the stability analysis of α -fair allocations. Due to space limitations, we consider the

case of $K = 2$ classes only. This simple scenario is sufficient to illustrate the sensitivity of the stability region to the fairness parameter α . The case of $K = 3$ or more classes is briefly discussed in Section VI.

IV. A NETWORK WITH TWO CLASSES

A. A discrete rate region

We first consider the case where the rate region \mathcal{R} consists of a finite set of n points A^1, A^2, \dots, A^n in \mathbb{R}_+^2 , for some $n \geq 2$. Since the rate region is assumed to be coordinate-convex, we can order these n points from left to right and from top to bottom in the sense that for all $i = 1, \dots, n-1$,

$$A_1^i < A_1^{i+1} \text{ and } A_2^i > A_2^{i+1}. \quad (2)$$

This is illustrated in Figure 1(a) for a discrete rate region of $n = 4$ points.




Fig. 1. (a) A discrete rate region (b) The corresponding cones associated with some α -fair allocation (note that A^2 is not scheduled in this example).

In view of (1), the α -fair allocation schedules in state $x = (x_1, x_2)$ a point A^j that maximizes³:

$$x_1^\alpha \frac{(A_1^j)^{1-\alpha}}{1-\alpha} + x_2^\alpha \frac{(A_2^j)^{1-\alpha}}{1-\alpha}, \quad j = 1, \dots, n. \quad (3)$$

Define for all $i \neq j$:

$$s_{\alpha}^{i,j} = \left(\frac{(A_2^j)^{1-\alpha} - (A_2^i)^{1-\alpha}}{(A_1^i)^{1-\alpha} - (A_1^j)^{1-\alpha}} \right)^{1/\alpha}. \quad (4)$$

This expression, that belongs to $\mathbb{R}_+ \cup \{\infty\}$, gives the direction of the switching line between the scheduling points A^i and A^j . Specifically, the inequality:

$$x_1^\alpha \frac{(A_1^i)^{1-\alpha}}{1-\alpha} + x_2^\alpha \frac{(A_2^i)^{1-\alpha}}{1-\alpha} \geq x_1^\alpha \frac{(A_1^j)^{1-\alpha}}{1-\alpha} + x_2^\alpha \frac{(A_2^j)^{1-\alpha}}{1-\alpha}$$

is satisfied in state $x = (x_1, x_2)$ for some $i < j$ if and only if $x_1 \leq s_{\alpha}^{i,j} x_2$. In particular, there exists a set of n cones C^1, \dots, C^n covering \mathbb{R}_+^2 such that point A^i is scheduled when $x \in C^i$. These cones are ordered from left to right (equivalently, from top to bottom) in the sense that a line parallel to the x_1 -axis crosses cones C^1, \dots, C^n in that order when x_1 goes from 0 to ∞ , cf. Figure 1(b).

Let d^i be the direction of the boundary between cones C^i and C^{i+1} , for all $i = 1, \dots, n-1$. This is the tangent of the angle between the line $C^i \cap C^{i+1}$ and the x_2 -axis. We have:

$$d^1 \leq d^2 \leq \dots \leq d^{n-1}.$$

²In all the paper, we assume without loss of generality that the rate region is coordinate-convex (if the rate region contains two distinct points A, B such that $A \leq B$ component-wise, A is never scheduled).

³In this and subsequent expressions, the function $(\cdot)^{1-\alpha}/(1-\alpha)$ is replaced by $\log(\cdot)$ for $\alpha = 1$.

These $n - 1$ directions, that characterize the allocation, are given by the following algorithm:

```

 $i = 1;$ 
while  $i < n$  do
   $l = \arg \min_{j=i+1, \dots, n} s_{\alpha}^{i,j};$ 
  for  $m = i$  to  $l - 1$ ,  $d^m = s_{\alpha}^{i,l};$ 
   $i = l;$ 

```

This algorithm leads to a unique solution in view of the following result. In particular, l may be arbitrarily chosen among those indices j that minimize $s_{\alpha}^{i,j}$.

Proposition 1: The algorithm gives a unique sequence of non-decreasing numbers d^1, \dots, d^{n-1} .

Proof. We first show that the sequence is non-decreasing. If $d^i < d^{i-1}$ for some $i \in \{2, \dots, n-1\}$, there exists some $a < i$ such that $d^{i-1} = s_{\alpha}^{a,i}$ and some $b > i$ such that $d^i = s_{\alpha}^{i,b}$. Since $s_{\alpha}^{i,b} < s_{\alpha}^{a,i}$, it follows easily from (4) that $s_{\alpha}^{a,b} < s_{\alpha}^{a,i}$. But this contradicts the fact that:

$$s_{\alpha}^{a,i} = d^{i-1} = \min_{j=a+1, \dots, n} s_{\alpha}^{a,j}.$$

We now prove that the sequence is unique. If for some i , the set of indices $j > i$ that minimize $s_{\alpha}^{i,j}$ is not unique, let b be the largest such index. We shall prove that:

$$d^i = d^{b-1} \quad \text{and} \quad d^b = \min_{j=b+1, \dots, n} s_{\alpha}^{b,j}. \quad (5)$$

Note that, since the sequence d^1, \dots, d^{n-1} is non-decreasing, this implies that $d^i = d^{i+1} = \dots = d^{b-1}$: the sequence does not depend on the choice of the index $j > i$ that minimizes $s_{\alpha}^{i,j}$. Denote by a this index. If $a = b$, (5) directly follows from the algorithm. Now assume that $a < b$. Since $s_{\alpha}^{i,a} = s_{\alpha}^{i,b}$, it follows from (4) that $s_{\alpha}^{a,b} = s_{\alpha}^{a,i}$. Thus $d^a \leq d^i$. Since the sequence d^1, \dots, d^{n-1} is non-decreasing, we get $d^a = d^i$. In particular,

$$s_{\alpha}^{a,b} = d^a = \min_{j=a+1, \dots, n} s_{\alpha}^{a,j}.$$

Note that b is the largest index $j > a$ that minimizes $s_{\alpha}^{a,j}$. If $s_{\alpha}^{a,c} = s_{\alpha}^{a,b}$ for some $c > b$, it indeed follows from (4) and the equality $s_{\alpha}^{a,c} = s_{\alpha}^{i,c}$ that $s_{\alpha}^{i,c} = s_{\alpha}^{i,a}$, which contradicts the fact that b is the largest index $j > i$ that minimizes $s_{\alpha}^{i,j}$. Now if b is the only index $j > a$ that minimizes $s_{\alpha}^{a,j}$, (5) follows from the algorithm. Otherwise, we apply the same argument as above until b is the only minimizing index. \square

As proved in Theorem 3 below, the stability region is entirely determined by the set of scheduled points. This set includes the extremal points A^1 and A^n (in view of (3), these points are always scheduled on the x_2 -axis and the x_1 -axis, respectively) and all non-extremal points A^i whose associated cone C^i does not reduce to a line, that is such that $d^{i-1} < d^i$. On the intersection between two cones that do not reduce to a line, any of the two corresponding points may be scheduled (the allocation is not unique).

We say that α -fair allocations are *monotonic cone policies* in the following sense.

Definition 1: An allocation is said to be a monotonic cone policy if there exists a set of non-empty cones C^1, \dots, C^n covering \mathbb{R}_+^2 such that:

- (i) For all $i \neq j$, $\check{C}^i \cap \check{C}^j = \emptyset$.
- (ii) The cones are ordered from left to right as above.
- (iii) The point A^i is scheduled when $x \in \check{C}^i$.
- (iv) The extremal points A^1 and A^n are scheduled on the x_2 -axis and the x_1 -axis, respectively.
- (v) Any of the two points A^i or A^j is scheduled when $x \in C^i \cap C^j$, provided $C^i \neq C^i \cap C^j$ and $C^j \neq C^i \cap C^j$.

The projective cone policies considered by Armony and Bambos [1], [2] are examples of monotonic cone policies in dimension 2. Note that a scheduled point A^i does not necessarily belong to the corresponding cone C^i . It turns out that the stability region of a monotonic cone policy depends on the scheduled points only, and not on the corresponding cones. Specifically, let $\mathcal{S} \subset \mathcal{R}$ be the set of scheduled points, that is the extremal points and all non-extremal points whose associated cone does not reduce to a line. We refer to the *contour* of \mathcal{S} as the broken line joining the points of \mathcal{S} from the left to the right.

Theorem 3: The stability region of a monotonic cone policy is the smallest coordinate-convex set containing the contour of the scheduled points \mathcal{S} .

Proof. Without loss of generality, we assume that all points are scheduled (if some points are not scheduled, it is sufficient to restrict the subsequent analysis to the subset of scheduled points). The proof is by fluid limit techniques [7]. In the fluid limit, the network state is continuous and evolves according to the following differential equations:

$$\begin{aligned} \frac{\partial x}{\partial t} &= \rho - A^i \quad \text{if } x \in \check{C}^i, \\ \frac{\partial x_1}{\partial t} &= \max(\rho_1 - A_1^1, 0), \quad \frac{\partial x_2}{\partial t} = \rho_2 - A_2^1 \quad \text{if } x_1 = 0, \\ \frac{\partial x_1}{\partial t} &= \rho_1 - A_1^n, \quad \frac{\partial x_2}{\partial t} = \max(\rho_2 - A_2^n, 0) \quad \text{if } x_2 = 0, \\ \frac{\partial x_1}{\partial t} &= \rho_1, \quad \frac{\partial x_2}{\partial t} = \rho_2 \quad \text{if } x = 0. \end{aligned}$$

On the intersection between two cones that do not reduce to a line, any of the two corresponding points may be scheduled: the fluid limit does not depend on this choice. We say that the system in fluid limit is stable if it empties in finite time and unstable if it grows linearly after a finite time. This implies the ergodicity and the transience of the original Markov process x , respectively [7], [13].

Let $\delta^i = \rho - A^i$ be the drift vector associated with A^i , for all $i = 1, \dots, n$. We first consider the case where $\delta_1^i < 0$ and $\delta_2^i < 0$ for some $i \in \{1, \dots, n\}$. The set of indices j such that $\delta_1^j < 0$ and $\delta_2^j < 0$ necessarily consists of consecutive indices $\{a, a+1, \dots, b\}$. We then have $\delta_1^j \geq 0$ and $\delta_2^j \leq 0$ for all $j < a$ and $\delta_1^j \leq 0$ and $\delta_2^j \geq 0$ for all $j > b$. Thus the system in fluid limit enters the set of cones C^a, C^{a+1}, \dots, C^b after a finite time and stays there. The system then empties in finite time. This case is illustrated in Figure 2. Similarly, the system is unstable if $\delta_1^i > 0$ and $\delta_2^i > 0$ for some $i \in \{1, \dots, n\}$.




Fig. 2. Drift vectors and system behavior in fluid limits - A stable scenario.

It now remains to investigate the case where $\delta_1^i \geq 0$, $\delta_2^i \leq 0$ and $\delta_1^{i+1} \leq 0$, $\delta_2^{i+1} \geq 0$ for some $i \in \{1, \dots, n-1\}$. This implies that $\delta_1^j > 0$, $\delta_2^j < 0$ for all $j < i$ and $\delta_1^j < 0$, $\delta_2^j > 0$ for all $j > i+1$, which further implies that the system enters the set $\mathcal{C}^i \cup \mathcal{C}^{i+1}$ after a finite time and stays there. We first consider the case where ρ is strictly below the line (A^i, A^{i+1}) . This case is illustrated in Figure 3(a) and arises when:

$$\frac{-\delta_2^i}{\delta_1^i} > \frac{\delta_2^{i+1}}{-\delta_1^{i+1}}.$$

Let $\theta > 0$ be such that:

$$\frac{-\delta_2^i}{\delta_1^i} > \theta > \frac{\delta_2^{i+1}}{-\delta_1^{i+1}}.$$

One can easily verify that the drift of the Lyapounov function $f(x) = \theta x_1 + x_2$ is lower bounded by a negative constant outside a compact set: the system is stable. Similarly, one proves that the system is unstable when ρ is strictly above the line (A^i, A^{i+1}) , as illustrated in Figure 3(b). \square




Fig. 3. Drift vectors for two scheduled points A^i and A^{i+1} .

Since α -fair allocations are monotone cone policies, it follows from Theorem 3 that the stability region of an α -fair allocation depends on the set of scheduled points only. We give in Figure 4 the stability region of an α -fair allocation when the rate region is that of Figure 1(a) and when the corresponding cones are those of Figure 1(b).

We now study the sensitivity of the stability region to the fairness parameter α . We first show that if the rate region has a convex structure, in the sense that the smallest coordinate-convex set containing its contour is convex, the stability region is maximum and independent of α .

Corollary 1: If \mathcal{R} has a convex structure, the stability region of α -fair allocations is maximum and equal to $\bar{\mathcal{R}}$.

Proof. The convex structure of the rate region implies that $s_{\alpha}^{i,i+1} < s_{\alpha}^{i,j}$ for all $i = 1, \dots, n-1$ and $j > i+1$. It then follows from the algorithm that $d_1 < d_2 < \dots < d_{n-1}$. \square




Fig. 4. Stability region of an α -fair allocation when $\mathcal{S} = \{A^1, A^3, A^4\}$.

We now consider an arbitrary discrete rate region \mathcal{R} . We have the following two key results.

Corollary 2: There exists $\beta \in (0, \infty)$ such that for all $\alpha < \beta$, the stability region of α -fair allocations is maximum and equal to $\bar{\mathcal{R}}$.

Proof. For all $i = 1, \dots, n-1$, there exists $\beta^i \in (0, \infty]$ such that for all $\alpha < \beta^i$, the indice $j > i$ that minimizes $s_{\alpha}^{i,j}$ is that which minimizes the angle between line (A^i, A^j) and the x_1 -axis. Letting $\beta = \min_{i=1, \dots, n} \beta^i$, we conclude that for all $\alpha < \beta$, the stability region of α -fair allocations is the smallest convex, coordinate-convex set containing \mathcal{R} . \square

Corollary 3: There exists $\gamma \geq 0$ such that for all $\alpha > \gamma$, the stability region of α -fair allocations is minimum and equal to the smallest coordinate-convex set containing the contour of the rate region \mathcal{R} .

Proof. For all $i = 1, \dots, n-1$, there exists $\gamma^i \in (0, \infty]$ such that for all $\alpha > \gamma^i$, the indice $j > i$ that minimizes $s_{\alpha}^{i,j}$ is that which minimizes the angle between line $(0, A^j)$ and the x_2 -axis, namely $i+1$. Letting $\gamma = \min_{i=1, \dots, n} \gamma^i$, we conclude that for all $\alpha > \gamma$, all points are scheduled. The stability region is minimum and equal to the smallest convex, coordinate-convex set containing \mathcal{R} . \square

We conclude that the stability region is sensitive to the fairness parameter α : it is maximum when $\alpha \rightarrow 0$ and minimum when $\alpha \rightarrow \infty$. In particular, max-min fairness does not guarantee maximum stability. To the best of our knowledge, this is the first result showing fairness “inefficiency” in a dynamic scenario with a varying number of flows. It differs from the usual efficiency vs. fairness trade-off discussed by many authors in static scenarios, see e.g. [15].

B. A continuous rate region

The above results can be generalized to a continuous rate region. Again, the stability region of an α -fair allocation only depends on the set of scheduled points, which can be obtained by applying the algorithm derived for a discrete rate region to a discretized version of the continuous rate region with a sufficiently small discretizing parameter. We give an example of continuous rate region in the following section.

V. EXAMPLES

We conclude the paper by applying previous results to a simple wireless network consisting of two base stations (BS). We refer to class 1 as those flows served by BS 1 and to class 2 as those flows served by BS 2.

A. Discrete rate region

We first assume the base stations transmit at full power when active, as in CDMA 1xEV-DO systems. We denote by P the power received by each user from each active base station. We assume for simplicity that it is the same for both base stations; this is the case for instance if users are at the same distance of both base stations in an obstacle-free propagation environment, as illustrated by Figure 5.




Fig. 5. A wireless network with two base stations.

The thermal noise power is denoted by N . The bit rate of flows served by a given BS depends on the activity of the other BS, and is assumed to be given by Shannon's formula. Specifically, the bit rate obtained when the other base station is switched on and off is respectively given in bit/s/Hz by:

$$c_{\text{on}} = \log_2 \left(1 + \frac{P}{N + P} \right), \quad c_{\text{off}} = \log_2 \left(1 + \frac{P}{N} \right).$$

Thus the rate region \mathcal{R} contains three points, as illustrated in Figure 6.




Fig. 6. Rate region when the base stations transmit at full power.

The rate region \mathcal{R} has a convex structure if and only if $c_{\text{on}} \geq c_{\text{off}}/2$, that is if

$$1 + \frac{P}{N + P} \geq \sqrt{1 + \frac{P}{N}}.$$

It may be easily verified that this occurs when the signal-to-noise ratio (SNR) P/N is less than the golden number:

$$\Phi = \frac{\sqrt{5} + 1}{2},$$

in which case all α -fair allocations achieve maximum stability in view of Corollary 1.

Now assume that the SNR is higher than the golden number. The stability region of α -fair allocations is maximum if and only if $c_{\text{off}}^{1-\alpha} > 2c_{\text{on}}^{1-\alpha}$, that is if $\alpha < \beta$, with

$$\beta = 1 - \frac{1}{\log_2(c_{\text{off}}) - \log_2(c_{\text{on}})}.$$

If $\alpha > \beta$, all points are scheduled and the stability region is minimum (thus $\gamma = \beta$ in that example). We give in Figure 7 the two corresponding stability regions (in bit/s/Hz) when the SNR is equal to 10 dB.




Fig. 7. The two stability regions when the SNR is equal to 10 dB.

The critical value β is given in Figure 8 as a function of the SNR (higher than $\Phi \approx 2$ dB). It slowly (logarithmically) increases from 0 to 1 when the SNR goes from Φ to ∞ .




Fig. 8. Critical value β as a function of the SNR ($\beta = \infty$ if the SNR is less than $\Phi \approx 2$ dB).

B. Continuous rate region

We now consider the case where the base stations do not necessarily transmit at maximum power. We still denote by P the maximum received power. The rate region \mathcal{R} is the set of vectors ϕ such that:

$$\phi_1 \leq \log_2 \left(1 + \frac{u_1 P}{N + u_2 P} \right),$$

$$\phi_2 \leq \log_2 \left(1 + \frac{u_2 P}{N + u_1 P} \right),$$

for some $u_1, u_2 \in [0, 1]$. The boundary of \mathcal{R} is obtained letting $u_1 = 1$ or $u_2 = 1$, i.e., at least one of the two base stations transmits at full power. The resulting rate region is presented in Figure 9 for different values of the SNR P/N . Note that




Fig. 9. Rate region for different values of the SNR ($-10, -3, 0, 3, 10$ dB, from bottom left to top right).

the rate region does not have a convex structure. It contains the three points of the discrete rate region considered above.

There are critical values β, γ of α , with $\beta \leq \gamma$, such that the stability region of α -fair allocations is maximum if $\alpha < \beta$ and minimum if $\alpha > \gamma$ (cf. Corollaries 2 and 3). We give in Figure 10 the stability region of α -fair allocations for different values of α . The critical values are $\beta \approx 0.12, \gamma \approx 0.41$ when the SNR is equal to 0 dB and $\beta \approx 0.46, \gamma \approx 0.65$ when the SNR is equal to 10 dB. Note that the stability region decreases continuously from the maximum stability region to the minimum stability region when α goes from β to γ . The critical values β, γ of α are given in Figure 11 with respect to the SNR.




Fig. 10. Stability region for different values of α (a) $\alpha = 0.12, 0.2, 0.41$, (b) $\alpha = 0.46, 0.5, 0.57, 0.61, 0.63, 0.65$, from top-right to bottom-left.




Fig. 11. Critical values β, γ of α as functions of the SNR.

VI. CONCLUSION

The stability region of utility-based allocations is larger than the rate region in non-convex cases. A judicious choice of the utility function ensures maximum stability by restricting the set of scheduled points and thus convexifying the rate region. This is the case of α -fair allocation for a sufficiently small fairness parameter α . Max-min fairness, on the other hand, always schedules all points of the rate region and therefore minimizes the stability region.

We believe these results have strong implications on the choice of the utility function to be used in the design of packet-level algorithms, especially in wireless networks where the interaction of transmitters through interference naturally leads to non-convex rate regions.

Though we considered the case of two classes only, we think the main results extend to the case of three or more classes. That is, the stability region depends on the fairness parameter α , is maximum when $\alpha \rightarrow 0$ and minimum when $\alpha \rightarrow \infty$. The problem is made harder by the fact that the stability region may depend on traffic characteristics like the flow size distribution, as pointed out in [4] for a wireless network where three or more base stations interact through interference.

Acknowledgments: This work was initiated after fruitful discussions with Vivek Mhatre and François Baccelli.

REFERENCES

- [1] M. Armony. Queueing networks with interacting service resources. *PhD thesis, Stanford University*, 1999.
- [2] M. Armony and N. Bambos. Queueing dynamics and maximal throughput scheduling in switched processing systems. *Queueing systems*, 44: 209–252, 2003.
- [3] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, 1992.
- [4] T. Bonald, S. Borst, N. Hegde, A. Proutière. Wireless data performance in multi-cell scenarios. In *Proc. ACM Sigmetrics / Performance*, 2004.
- [5] T. Bonald and L. Massoulié. Impact of fairness on Internet performance. In *Proc. ACM Sigmetrics / Performance*, 2001.
- [6] T. Bonald, L. Massoulié, A. Proutière and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. To appear in *Queueing systems*, 2006.
- [7] J. Dai. On positive harris recurrence of multi-class queueing networks: a unified approach via fluid limit models. *Annals of Applied Probability*, 5: 49–77, 1995.
- [8] F. P. Kelly. Charging and rate control for elastic traffic. *European Trans. on Telecommunications*, 8: 33–37, 1997.
- [9] X. Lin, N. Shroff. The impact of imperfect scheduling on cross-layer rate control in multi-hop wireless networks. In *Proc. IEEE Infocom*, 2005.
- [10] S. Low and D. Lapsley. Rate control in communication networks: Shadow prices, proportional fairness and stability. *IEEE/ACM Trans. Networking*, 7: 861–874, 1999.
- [11] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity. submitted, 2005.
- [12] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *IEEE/ACM Trans. Networking*, 10: 320–328, 2002.
- [13] S.P. Meyn. Transience of multiclass queueing networks via fluid limit models. *Annals of Applied Probability*, 5: 946–957, 1995.
- [14] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Networking*, 8: 556–567, 2000.
- [15] A. Tang, J. Wang, and S. H. Low. Is fair allocation always inefficient? In *Proc. IEEE Infocom*, 2004.
- [16] G. de Veciana, T. J. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting elastic services. *IEEE Trans. Networking*, 1: 2–14, 2001.
- [17] H. Q. Ye. Stability of data networks under an optimization-based bandwidth allocation. *IEEE Trans. Aut. Control*, 48: 1238–1242, 2003.