

A Queueing Analysis of Max-Min Fairness, Proportional Fairness and Balanced Fairness

T. Bonald[†], L. Massoulié^{*}, A. Proutière[†], J. Virtamo[‡]

[†]France Telecom
{thomas.bonald, alexandre.proutiere}@francetelecom.com

^{*}Microsoft Research
lmassoul@microsoft.com

[‡]Helsinki University of Technology
jorma.virtamo@hut.fi

Submission to the special issue of Queueing Systems:
Queueing Models for Fair Resource Sharing
December 15, 2005

Abstract

We compare the performance of three usual allocations (max-min fairness, proportional fairness and balanced fairness) in a communication network whose resources are shared by a random number of data flows. The model consists of a network of processor-sharing queues. The vector of service rates, which is constrained by some compact, convex capacity set representing the network resources, is a function of the number of customers in each queue. This function determines the way network resources are allocated. We show that this model is representative of a rich class of wired and wireless networks. We give in this general framework the stability condition of max-min fairness, proportional fairness and balanced fairness and compare their performance on a number of toy networks.

1 Introduction

Following the seminal work of Kelly et al. [22], considerable recent research effort has been devoted to the issue of optimally allocating resources of wired and wireless communication networks [21, 26, 28, 32, 33, 34, 38]. As a general rule a resource allocation is said to be optimal if it maximizes the overall “utility” of a given set of users, where the utility of each user is some concave function of the bit rate allocated to that user. The dynamic nature of traffic is not taken into account in this framework. In real networks, data flows do not last forever but arrive at random times and leave the network once the corresponding digital document has been transferred. This results in a random, dynamic set of active users that impacts the bit rate that can be allocated to each of them. Conversely, the bit rate allocated to each user determines how long that user will stay active and thus impacts the evolution of the set of active users. The study of resource allocation cannot therefore be decoupled from that of the stochastic process describing the set of active users. This is best illustrated by an example.

Consider a wireless system where a base station transmits data to mobiles that experience constant radio conditions that can be classified as either “good” or “bad”. Those mobiles that experience bad radio conditions require more radio resources (e.g., bandwidth, transmission power) than the other mobiles to get the same bit rate. Thus the overall bit rate is maximized by allocating *all* resources to those mobiles having a good radio channel. Such an allocation is not “fair”, however. It would be much more “fair” to maximize the minimum bit rates (so-called max-min fairness). But most resources are then consumed by those mobiles having a bad radio channel. We see that there is a trade-off between “efficiency” and “fairness”, which is captured by the above mentioned utility maximizing allocations. A well-known example is proportional fairness, obtained with a logarithmic utility function [22].

Now let us account for the fact that the number of active mobiles is dynamic. It is then unclear what an “optimal” resource allocation is. Maximizing the overall bit rate in each state starves those mobiles that experience bad radio conditions. Such mobiles, which are served only in the absence of mobiles with good radio conditions, stay longer in the system. This results in a steady state where most active mobiles have bad radio conditions. Conversely, max-min fairness gives relatively few resources to those mobiles with good radio conditions, which results in an unnecessarily high number of such mobiles in steady state. We see that the performance of an allocation depends critically on the steady state it leads to. The apparent trade-off between “efficiency” and “fairness” vanishes.

It may be argued that network resources should rather be allocated so as to minimize the mean duration of data transfers. Such an “optimal” allocation may depend on detailed traffic characteristics like the distribution of data transfer volumes, however. A more practically interesting objective is to allocate resources in such a way that the steady state does *not* depend on these traffic characteristics, making performance robust with respect to user behaviour. Balanced fairness is such an “insensitive” allocation [8, 9].

The objective of the present paper is to describe a class of models that allows us to compare the performance of various resource allocations in a dynamic setting with a randomly varying number of ongoing flows. Specifically, the model consists of a network of processor-sharing queues where each queue corresponds to a particular flow class. The service rate of each queue, which represents the overall bit rate allocated to the corresponding flow class, depends on the entire network state. The vector of service rates is constrained by some compact, convex set representing the capacity of the considered communication network. Such a queueing network is representative of a rich class of wired and wireless networks, as illustrated by the various toy examples presented in the following. We use these simple, reference examples as a basis to compare the performance of max-min fairness, proportional fairness and balanced fairness. A complete study involving more realistic scenarios is certainly needed to validate the preliminary conclusions drawn in the present paper.

An overview of related work is given in the next section (see also [40] for a survey in the context of wired networks). Section 3 describes the model while section 4 gives the definition and key properties of various resource allocations. Stability issues are addressed in section 5. The following two sections are devoted to examples of wired and wireless networks, respectively. Section 8 concludes the paper.

2 Related work

The modeling of communication networks as networks of processor-sharing queues started with the analysis of radio systems by Telatar and Gallager [44] and Stamatelos and Koukoulidis [42]. Based on the observation that TCP shares bandwidth in an approximately fair way, Heyman et al. [20] and Massoulié and Roberts [31] applied similar models to wired IP networks. Practical dimensioning rules were developed on this basis by Berger and Kogan [4]. Ben Fredj et al. observed the insensitivity of the results to detailed traffic characteristics like the structure of user sessions [3]. Many papers proposed modified models that account more precisely for the way bandwidth is shared by TCP, see [25] for example.

Stability issues were first addressed by Tassiulas and Ephremides [43] in their seminal work on ad-hoc networks. A class of resource allocations that maximize the stability region of some wireless networks were considered in related papers by Armony and Bambos [1] and Bambos and Michailidis [2]. De Veciana et al. [45], Bonald and Massoulié [7] and Ye [46] proved the stability of fixed networks under various allocations like max-min fairness and proportional fairness. In the present paper, we extend these results to any network whose resources are constrained by some compact, convex capacity set (see proposition 1 and theorem 2).

The first analytical performance result for networks with multiple resources was derived by Massoulié and Roberts [31]. They observed that under proportional fairness, the Markov process describing the number of ongoing flows on each route of a homogeneous linear network is reversible. This result was extended to homogeneous grids by Bonald and Massoulié [7] and further generalized to homogeneous hypercubes by Kelly (unpublished communication; see [8]). Bonald and Proutière [8] derived similar results for any network topology under balanced fairness and proved the insensitivity of these results in a similar way as in [3]. They also proved that utility maximizing allocations lead to sensitive results except for proportional fairness in homogeneous hypercubes. This explains why so few analytical results exist for such allocations. The known results reduce to a mean-field approximation of max-min fairness in a homogeneous star network by Fayolle et al. [18] and a fluid limit analysis of so-called α -fair allocations by Kelly and Williams [23]. We prove in this paper that max-min fairness is also sensitive in the general context of a compact, convex capacity set except in the trivial case where the network reduces to a set of independent links (see theorem 1). We also find a new class of capacity sets for which proportional fairness is insensitive (see §4.4). This class is quite restrictive, however, and proportional fairness turns out to be sensitive in most cases of practical interest.

The performance of balanced fairness has been studied in a number of papers. Some bounds and recursive formulas were derived in [9, 11, 12, 13]. The notion of balanced fairness, initially introduced for wired networks with single-path routing, was applied to networks with traffic splitting by Leino and Virtamo [27] and to ad-hoc networks by Penttinen and Virtamo [37]. These results are included in the unified framework of a compact, convex capacity set introduced in the present paper.

While network performance is generally evaluated at normal traffic loads, Massoulié and Roberts [30] and Boyer et al. [14] explored the possibility of an overload and recognized the need for admission control. The related issue of load balancing has recently been addressed by Bonald et al. [6]. The corresponding model is a queueing network with both state-dependent arrival rates and service rates. We do neither consider admission control nor load balancing in the present paper.

Finally, a number of papers are devoted to the integration of data traffic with streaming traffic. Núñez Queija et al. [35] and Delcoigne et al. [17] evaluated the performance of data

traffic in a network where priority is given to streaming traffic. Key et al. [24] proved that the stability of a network with adaptive streaming traffic is not affected by this traffic under α -fair sharing. Bonald and Proutière [10] derived insensitive performance bounds for a similar network under balanced fairness. We here restrict the analysis to data traffic alone.

3 Model

In this section, we first present the model in terms of traffic characteristics and network resources. We then describe the corresponding queueing network and introduce the performance metric used in the rest of the paper.

3.1 Traffic characteristics

We model traffic at *flow* level. Specifically, we ignore the complex interaction of packet level mechanisms (congestion control, scheduling, routing, buffer management) at short-time scales and are only interested in the long-term resource allocation they realize. The flow content is then viewed as a fluid which is transmitted as a continuous stream through the network. The transmission rate changes at flow arrivals and flow departures only. This abstraction is suitable for the performance evaluation of relatively *large* flows (more than 100 packets, say) that typically generate most traffic.

In practice, flows are generated within sessions, each session being composed of a succession of flows separated by intervals of inactivity generally referred to as “think-times”. A typical example is the succession of Web pages downloaded by a user in a period of continuous activity. This may result in a strongly correlated flow arrival process, depending on the number of flows in a session, the distribution of flow sizes and think-time durations and their possible correlation. The *session* arrival process, on the other hand, is typically a Poisson process [36]. This is due to the fact that sessions are generated independently by a large population of users. For a limited user population, sessions may simply be considered as permanent, cf. [4, 20].

It turns out that the performance of “fair” allocations like max-min fairness and proportional fairness is not highly sensitive to these detailed traffic characteristics. This insensitivity property is even exact for balanced fairness [8]. Thus in the rest of the paper, we assume that flows arrive as a Poisson process and have i.i.d. exponential sizes. The study of the degree of sensitivity of max-min fairness and proportional fairness is beyond the scope of the present paper.

3.2 Network resources

Each flow is characterized by the network resources it requires. A typical example in a wired network is the set of links on the path from source to destination (see section 6). We consider an arbitrary set of N flow classes. All flows within the same class have the same resource requirements. Class- i flows arrive as a Poisson process of intensity $\lambda_i > 0$ and have i.i.d. exponential sizes of mean $\sigma_i > 0$ (in bits). Let $\rho_i = \lambda_i \sigma_i$ be the *traffic intensity* of class- i flows (in bit/s). This is the traffic volume generated by class- i flows per unit of time. We denote by ρ the vector of traffic intensities (ρ_1, \dots, ρ_N) .

Let x_i be the number of class- i flows in progress. We refer to the vector $x = (x_1, \dots, x_N)$ as the *network state*. The total bit rate allocated to class- i flows in state x is equally shared by these flows and denoted by $\phi_i(x)$. By convention, we let $\phi_i(x) = 0$ if $x_i = 0$. We refer to the vector $\phi(x) = (\phi_1(x), \dots, \phi_N(x))$ as the *resource allocation* in state x . The resource allocation

is constrained by some compact, convex, coordinate convex¹ set \mathcal{C} referred to as the *capacity set*. A number of examples of capacity sets are given in sections 6 and 7. We denote by $\mathcal{C}(x)$ the set of vectors of \mathcal{C} whose i -th component is equal to 0 if $x_i = 0$. Note that $\phi(x) \in \mathcal{C}(x)$ for all states x .

3.3 A queueing network

The considered model may be viewed as a network of N processor-sharing queues where each queue corresponds to a flow *class*. Customers arrive at queue i as a Poisson process of intensity λ_i and have i.i.d. exponential service requirements of mean σ_i . The service rate $\phi_i(x)$ of queue i depends on the entire network state x . This queueing network is a Whittle network [41] if and only if the resource allocation is *balanced* in the sense that for any pair of classes i, j and any state x such that $x_i > 0$ and $x_j > 0$,

$$\phi_i(x)\phi_j(x - e_i) = \phi_i(x - e_j)\phi_j(x), \quad (1)$$

where e_i is the N -dimensional vector with 1 in component i and 0 elsewhere. This balance property is equivalent to the reversibility of the Markov process X that describes the evolution of the network state, with transition rates:

$$q(x, x + e_i) = \lambda_i \quad \text{and} \quad q(x, x - e_i) = \frac{\phi_i(x)}{\sigma_i}.$$

The invariant measure of X is then explicit:

$$\pi(x) = \Phi(x)\rho_1^{x_1} \dots \rho_N^{x_N}. \quad (2)$$

The function Φ , generally referred to as the *balance function*, is defined by $\Phi(0) = 1$ and

$$\forall x \neq 0, \quad \Phi(x) = (\phi_{i_1}(e_{i_1})\phi_{i_2}(e_{i_1} + e_{i_2}) \dots \phi_{i_n}(e_{i_1} + e_{i_2} + \dots + e_{i_n}))^{-1}, \quad (3)$$

where $0, e_{i_1}, e_{i_1} + e_{i_2}, \dots, e_{i_1} + e_{i_2} + \dots + e_{i_n} \equiv x$ denotes any path from state 0 to state x and $n = \sum_{i=1}^N x_i$. Note that this definition does not depend on the considered path in view of (1). If the invariant measure π has a finite sum, the Markov process X is positive recurrent and its stationary distribution follows by normalisation.

When the balance property holds, the queueing network is *insensitive* in the sense that the stationary distribution of X does not depend on the distribution of service requirements beyond the mean. It may in fact be shown that this stationary distribution is then insensitive to *all* traffic characteristics described in §3.1 except for the traffic intensity [8]. Moreover, the balance property is a necessary *and* sufficient condition for insensitivity. For a non-balanced allocation, the stationary distribution is sensitive to *all* traffic characteristics and generally intractable. This is the case of most networks under max-min fairness and proportional fairness (see section 4).

3.4 Performance metric

Assume that the network is *stable* in the sense that the Markov process X is positive recurrent. A number of performance metrics could then be considered to evaluate the quality-of-service

¹A set $A \subset \mathbb{R}_+^N$ is said to be coordinate convex if $b \in A$ implies $a \in A$ for all $a \in \mathbb{R}_+^N$ such that $a \leq b$ component-wise.

experienced by the data transfers. We focus on the *mean flow duration* for each class. By Little’s law, the mean duration of class- i flows is given by:

$$\frac{\mathbb{E}[X_i]}{\lambda_i},$$

where X_i denotes the number of class- i flows in steady state. A convenient, equivalent performance metric is the *flow throughput*, defined as the ratio of the mean flow size to the mean flow duration. In view of the above expression, the class- i flow throughput γ_i is the ratio of the class- i traffic intensity $\rho_i = \lambda_i \sigma_i$ to the mean number of class- i flows:

$$\gamma_i = \frac{\rho_i}{\mathbb{E}[X_i]}. \quad (4)$$

4 Resource allocation

This section is devoted to the definition and the properties of max-min fairness, proportional fairness and balanced fairness. We first introduce two key properties of resource allocations.

4.1 Pareto efficiency

A resource allocation is said to be *Pareto efficient* if all resources are “consumed” in the sense that the bit rate allocated to one flow cannot be increased without decreasing the bit rate allocated to another flow. Formally, the allocation ϕ is said to be Pareto efficient if, for any network state x , any vector $\varphi \in \mathcal{C}(x)$ such that $\varphi \geq \phi(x)$ component-wise is equal to $\phi(x)$.

Note that Pareto efficiency implies that $\phi(x)$ belongs to the boundary of the capacity set in any state x . Both properties are in fact equivalent if the boundary of the capacity set does not contain any segment parallel to one of the class- i axes, $i = 1, \dots, N$.

4.2 Homotheticity

A resource allocation is said to be *homothetic* if changing the resources available for a given class by some factor changes the bit rate allocated to this class by the same factor. This may be formalized as follows. For any $a \in \mathbb{R}_+^N$, denote by ϕ^a the allocation associated with the capacity set $a \times \mathcal{C}$ defined as the set of vectors $a \times \varphi$, $\varphi \in \mathcal{C}$, where $a \times \varphi$ denotes the vector whose i -th component is equal to the product $a_i \varphi_i$. The allocation ϕ is said to be homothetic if:

$$\forall a \in \mathbb{R}_+^N, \quad \phi^a = a \times \phi.$$

As shown in sections 6 and 7, the homotheticity property simplifies the analysis of some networks whose capacity set is the homothetic version of that of some other networks.

4.3 Max-min fairness

The notion of max-min fairness is well known in political science [39]. It was introduced as a design objective for communication networks by Bertsekas and Gallager [5]. The principle of max-min fairness is to allocate network resources in such a way that the bit rate of a flow cannot be increased without decreasing the bit rate of a flow having a smaller bit rate. Max-min fairness is uniquely defined by the following *water-filling* procedure:

1. start from a bit rate equal to zero for all flows;
2. increase the bit rate of all flows at the same speed until the bit rate of some flows is constrained by the capacity set; freeze the bit rate of these flows;
3. apply step 2 repeatedly to non-frozen flows until the bit rate of all flows is constrained by the capacity set.

Note in particular that max-min fairness is Pareto efficient. If the boundary of the capacity set does not contain any segment parallel to one of the class- i axes, $i = 1, \dots, N$, the water-filling procedure stops after the first instance of step 2 and the allocation is simply given by the intersection of the line of direction x with the boundary of the capacity set.

Max-min fairness is non-homothetic (cf. the example of §7.1) and non-balanced except if the network reduces to a set of independent links, as stated in the following theorem proved in Appendix A.

Theorem 1 *Max-min fairness is balanced if and only if for some $L \geq 1$, the network reduces to a set of L independent links in the sense that there exists a partition I_1, \dots, I_L of the set of classes $\{1, \dots, N\}$ and some positive constants c_1, \dots, c_L such that:*

$$\mathcal{C} = \left\{ \varphi : \sum_{i \in I_1} \varphi_i \leq c_1, \dots, \sum_{i \in I_L} \varphi_i \leq c_L \right\}.$$

4.4 Proportional fairness

Proportional fairness was introduced by Kelly et al. [22]. This is the most commonly cited utility maximizing allocation. For any network state x , $\phi(x)$ is defined as the unique vector φ that maximizes the quantity

$$\sum_{i=1}^N x_i \log(\varphi_i)$$

under the constraint $\varphi \in \mathcal{C}(x)$. Like any utility maximizing allocation, proportional fairness is Pareto efficient. It may easily be verified, on the other hand, that this is the only utility maximizing allocation that is homothetic.

Proportional fairness is non-balanced except for some specific capacity sets like those of so-called homogeneous hypercubes [8], sets of vectors φ such that $\varphi_1^\alpha + \dots + \varphi_N^\alpha \leq 1$, for some $\alpha \geq 1$, and combined, homothetic versions of these sets. An example with $N = 3$ classes is the set of vectors φ such that $\varphi_1^\alpha + \varphi_2^\alpha \leq 1$ and $\varphi_1^\alpha + \varphi_3^\alpha \leq 1$, for which we get in any state $x \neq 0$:

$$\phi_1(x) = \left(\frac{x_1}{x_1 + x_2 + x_3} \right)^{\frac{1}{\alpha}},$$

$$\phi_2(x) = \left(\frac{x_2 + x_3}{x_1 + x_2 + x_3} \right)^{\frac{1}{\alpha}} \quad \text{if } x_2 > 0, \quad \phi_3(x) = \left(\frac{x_2 + x_3}{x_1 + x_2 + x_3} \right)^{\frac{1}{\alpha}} \quad \text{if } x_3 > 0.$$

It may easily be verified that the balance property indeed holds in this case.

4.5 Balanced fairness

The notion of balanced fairness was introduced by Bonald and Proutière as a means to approximately evaluate the performance of “fair” allocations like max-min fairness and proportional fairness in wired networks [8]. Balanced fairness is defined as the unique *balanced* allocation (cf. §3.3) such that $\phi(x)$ belongs to the boundary of the capacity set in any state $x \neq 0$. Let Φ be the corresponding balance function. In view of (3), we have for any state $x \neq 0$:

$$\phi_i(x) = \frac{\Phi(x - e_i)}{\Phi(x)}, \quad i = 1, \dots, N. \quad (5)$$

Thus $\Phi(x)$ is recursively defined as the minimum positive constant β such that the vector $(\Phi(x - e_1), \dots, \Phi(x - e_N))/\beta$ belongs to \mathcal{C} , with the convention $\Phi(0) = 1$ and $\Phi(x) = 0$ for all $x \notin \mathbb{Z}_+^N$. If the capacity set is a polytope, which means that $\mathcal{C} = \{\varphi : \varphi A \leq C\}$ for some $L \geq 1$, an $N \times L$ -dimensional non-negative matrix A and a L -dimensional positive vector C , this recursion takes the following simple form:

$$\Phi(x) = \min_{l=1, \dots, L} \frac{1}{C_l} \sum_{i=1}^N \Phi(x - e_i) A_{il}.$$

Explicit expressions for the flow throughput can be derived from this recursion [11, 12]. It turns out that in most examples considered in sections 6 and 7, the capacity set is a polytope.

Note that balanced fairness is not necessarily Pareto efficient [8]. If balanced fairness is not Pareto efficient for a particular capacity set \mathcal{C} , no balanced allocation is Pareto efficient on \mathcal{C} . Conversely, if balanced fairness is Pareto efficient for a capacity set \mathcal{C} , no other balanced allocation is Pareto efficient on \mathcal{C} . In particular, if a Pareto efficient allocation (like max-min fairness or proportional fairness) is balanced for a capacity set \mathcal{C} , it necessarily coincides with balanced fairness on \mathcal{C} .

Finally, balanced fairness is homothetic. This follows from the fact that for any vector $a \in \mathbb{R}_+^N$, the allocation defined by $\phi^a = a \times \phi$ is balanced and belongs to the boundary of the capacity set. The corresponding balance function Φ^a is given by:

$$\Phi^a(x) = \Phi(x) \frac{1}{a_1^{x_1}} \dots \frac{1}{a_N^{x_N}}. \quad (6)$$

5 Stability issues

We give in this section the stability condition of max-min fairness, proportional fairness and balanced fairness.

5.1 Necessary condition

We first show that for *any* allocation, a necessary stability condition is that the traffic intensity vector belongs to the capacity set.

Proposition 1 *For any allocation, a necessary stability condition is $\rho \in \mathcal{C}$.*

Proof. If $\rho \notin \mathcal{C}$, there exists a half-space \mathcal{H} that contains \mathcal{C} but not ρ (by the convexity of \mathcal{C}) and whose boundary intersects each axis (by the compactness and the coordinate convexity of \mathcal{C}). We denote by δ_i the intersection of the boundary of \mathcal{H} with the class- i axis. Since

$$\mathcal{H} = \left\{ \varphi : \sum_{i=1}^N \frac{\varphi_i}{\delta_i} \leq 1 \right\},$$

the network associated with the capacity set $\mathcal{H} \cap \mathbb{R}_+^N$ corresponds to a multi-class single server queue with unit speed, class- i arrival rate λ_i and class- i mean service time σ_i/δ_i . The load of this queue is equal to:

$$\sum_{i=1}^N \frac{\rho_i}{\delta_i}.$$

Since $\rho \notin \mathcal{H}$, we have:

$$\sum_{i=1}^N \frac{\rho_i}{\delta_i} > 1,$$

so that the queue is unstable. Since \mathcal{H} contains \mathcal{C} , the network is unstable. \square

5.2 Sufficient condition

The above necessary stability condition is also a sufficient stability condition up to the critical case where the traffic intensity vector belongs to the boundary of the capacity set. Let $\check{\mathcal{C}}$ be the interior of the capacity set. The following result is proved in Appendix B.

Theorem 2 *For max-min fairness, proportional fairness and balanced fairness, a sufficient stability condition is $\rho \in \check{\mathcal{C}}$.*

6 Wired networks

We represent a wired network as a set of L links and K routes where each route k is defined as a subset r_k of the set of links $\{1, \dots, L\}$. Let $C_l > 0$ be the capacity of link l . We denote by C the vector (C_1, \dots, C_L) . We refer to the *incidence matrix* as the $K \times L$ -dimensional matrix A whose k, l -entry is equal to 1 if $l \in r_k$ and 0 otherwise.

6.1 Single-path routing

In case of single-path routing, each flow is assigned a particular route in the network so that $N = K$. The corresponding capacity set is given by:

$$\mathcal{C} = \{\varphi : \varphi A \leq C\}.$$

Consider for instance the 3-branch symmetric tree of figure 1. We have in this case:

$$C = (2, 1, 1, 1) \quad \text{and} \quad A = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}.$$

The corresponding capacity set is shown in figure 1.




Figure 1: A symmetric tree network and its capacity set.

The flow throughput (4) depends on the allocation and on the vector of traffic intensities. In all numerical applications in the paper, we take the same traffic intensities for all classes. For balanced fairness, we obtain applying the results of [12]:

$$\gamma_1 = \gamma_2 = \gamma_3 = \frac{(2 - \varrho)(3 - \varrho)(6 + \varrho)}{(4 - \varrho)(9 + \varrho)},$$

where ϱ denotes the total traffic intensity $\rho_1 + \rho_2 + \rho_3$. Note that the stability condition is $\varrho < 2$. Figure 2 gives the flow throughput as a function of the traffic intensity ϱ . In this and subsequent figures, the performance of max-min fairness and proportional fairness is obtained by simulation (except when analytical results are available) with the same mean flow size for all classes. Each simulation point corresponds to the flow throughput measured over 1,000,000 events (arrival or departure) after a warm-up period of 100,000 events. The performance of balanced fairness, on the other hand, is always obtained analytically. We observe from figure 2 that the performance of max-min fairness and proportional fairness cannot be distinguished from that of balanced fairness.




Figure 2: Flow throughput for the network of figure 1.

Now consider the 2-branch tree of figure 3, defined by:

$$C = (2, 1) \quad \text{and} \quad A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$




Figure 3: An asymmetric tree network and its capacity set.

The corresponding capacity set is shown in figure 3.

For balanced fairness, the flow throughputs are given by [12]:

$$\gamma_1 = 2 - \rho, \quad \gamma_2 = \frac{(2 - \rho)(4 - \rho)}{8 - \rho},$$

where ρ denotes the total traffic intensity $\rho_1 + \rho_2$. The stability condition is still $\rho < 2$. As shown in figure 4, the three allocations again give very similar results.




Figure 4: Flow throughput for the network of figure 3.

Note that the model can be used to represent *multicast* flows as well. In the network of figure 3 for instance, the class-2 flows may correspond to multicast flows from one source (located between the two links) to two destinations (one per link). Similarly for the network of figure 1, a multicast flow whose packets are replicated on each of the three branches after their transmission on the root link may be represented by an additional route that contains all network links. In all cases, the bit rate of the multicast flow is determined by the most constraining link of the multicast tree.

6.2 Multi-path routing

Now assume that each flow class is not assigned a single route but a *set* of routes. Specifically, each class i chooses at any time *one* route in a subset s_i of the set of routes $\{1, \dots, K\}$. Let \mathcal{R}

be the set of $N \times K$ stochastic matrices² such that on each row i , the i, k -entries are equal to 0 for all k except one in the set s_i . Each matrix $R \in \mathcal{R}$ corresponds to a particular route choice. The routing scheme, referred to as *multi-path routing*, is assumed to operate at a much higher frequency than the typical flow duration (cf. §3.1). The capacity set is then the *convex hull* of the capacity sets associated with the routing matrices $R \in \mathcal{R}$:

$$\mathcal{C} = \text{convex hull of } \{\varphi : \exists R \in \mathcal{R}, \varphi R A \leq C\}.$$

Consider the example of figure 5 with two unit capacity links. Class 2 chooses either one link or the other. We have in this case:

$$C = (1, 1), \quad A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad R = \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{pmatrix}.$$

The corresponding capacity set is shown in figure 5. It coincides with the capacity set of the tree network of figure 1, so that both networks are in fact equivalent. The flow throughput is the same for the three classes and given by figure 2 as a function of the overall traffic intensity.




Figure 5: A symmetric network with multi-path routing and its capacity set.

Consider now the example of figure 6 with two links of different capacities:

$$C = (2, 1), \quad A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad R = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}.$$

The corresponding capacity set, shown in figure 6, does not coincide with any wired network with static routes. This is in fact a homothetic version of the capacity set of figure 3.

Using the homotheticity property (6) and the results of [12], we get for balanced fairness:

$$\gamma_1 = \frac{8 - 3\rho}{4}, \quad \gamma_2 = \frac{(8 - 3\rho)(4 - \rho)(8 - \rho)}{2(64 - 24\rho + \rho^2)},$$

where ρ denotes the total traffic intensity $\rho_1 + \rho_2$. The stability condition is $\rho < 8/3$ in this case. The results, shown in figure 7, are similar for the three allocations except for proportional fairness that leads to a slightly higher flow throughput for class 2 at moderate load.

²A matrix is said to be stochastic if its entries are non-negative and the sum of the entries on each row is equal to 1.




Figure 6: An asymmetric network with multi-path routing and its capacity set.




Figure 7: Flow throughput for the network of figure 6.

6.3 Traffic splitting

Suppose now that each class i can use *all* routes in the set s_i at the same time. Let \mathcal{S} be the set of $N \times K$ stochastic matrices such that on each row i , the i, k -entries are equal to 0 for all k except those in the set s_i . Each matrix $S \in \mathcal{S}$ corresponds to a particular traffic splitting scheme. The capacity set is:

$$\mathcal{C} = \{\varphi : \exists S \in \mathcal{S}, \varphi SA \leq C\}.$$

Consider the example of figure 8, which corresponds to the network of figure 5 but with traffic splitting instead of multi-path routing. We have in this case:

$$C = (1, 1), \quad A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad S = \begin{pmatrix} 1 & 0 \\ \alpha_1 & \alpha_2 \\ 0 & 1 \end{pmatrix} \quad \text{with} \quad \alpha_1 + \alpha_2 = 1.$$

The corresponding capacity set, shown in figure 8, coincides with that of the tree network of figure 1 where the capacity of one of the three branches is equal to 2.

We get for balanced fairness:

$$\gamma_1 = \gamma_3 = \frac{2(2 - \varrho)(3 - \varrho)}{12 - 5\varrho}, \quad \gamma_2 = 2 - \varrho,$$

where ϱ denotes the total traffic intensity $\rho_1 + \rho_2 + \rho_3$. The stability condition is $\varrho < 2$. The results are given in figure 9.




Figure 8: A wired network with traffic splitting and its capacity set.




Figure 9: Flow throughput for the network of figure 8.

7 Wireless networks

Consider a wireless network with a pre-determined set of *transmission profiles* to be used in a sequential way. Each transmission profile corresponds to a specific allocation of the downlink radio resources used to transmit data from the base stations to the mobiles. We denote by M the number of transmission profiles and by C the $M \times N$ -dimensional matrix whose m, i -entry gives the bit rate allocated to class- i flows in transmission profile m . Let \mathcal{T} be the set of M -dimensional non-negative row vectors summing to 1. Each element $\tau \in \mathcal{T}$ defines a particular *schedule* in the sense that the component m of τ gives the fraction of time the transmission profile m is used. As mentioned in §3.1, the scheduling is assumed to operate at a much higher frequency than the typical flow duration, so that the capacity set is given by:

$$\mathcal{C} = \{\varphi : \exists \tau \in \mathcal{T}, \varphi \leq \tau C\}.$$

7.1 A single base station

We first consider the case of a single base station that serves each flow one at a time. There are N transmission profiles. The base station transmits to a class- i flow in transmission profile

$i, i = 1, \dots, N$. Thus C is an N -dimensional diagonal matrix and the capacity set \mathcal{C} is the homothetic version of \mathcal{T} , the set of possible schedules. In view of the homotheticity property, proportional fairness and balanced fairness allocate the transmission time equally among the flows. The corresponding model is a processor sharing queue. Max-min fairness, on the other hand, is not homothetic. It equalizes the bit rates so that the fraction of transmission time allocated to each class- i flow is inversely proportional to C_i , the i -th diagonal element of the matrix C . The corresponding model is a *discriminatory* processor sharing queue.




Figure 10: A single base station with two classes (time-sharing).

Consider the simple case of two classes, as shown in figure 10 when $C_1 = 3$ and $C_2 = 1$. The model is basically that considered in section 1 where mobiles experience either “good” or “bad” radio conditions. For proportional fairness and balanced fairness, the flow throughputs are given by:

$$\gamma_1 = C_1(1 - \varrho) \quad \text{and} \quad \gamma_2 = C_2(1 - \varrho),$$

where ϱ is equal to $\rho_1/C_1 + \rho_2/C_2$. The stability condition is $\varrho < 1$.

For max-min fairness, we obtain using the results of Fayolle et al. [18]:

$$\gamma_1 = C_1 \frac{(1 - \varrho)(2 - \varrho)}{2 - \varrho - \varrho \frac{C_1}{C_2} \frac{C_2 - C_1}{C_2 + C_1}} \quad \text{and} \quad \gamma_2 = C_2 \frac{(1 - \varrho)(2 - \varrho)}{2 - \varrho - \varrho \frac{C_2}{C_1} \frac{C_1 - C_2}{C_1 + C_2}}.$$

The results are given in figure 11 as a function of the overall traffic intensity $\rho_1 + \rho_2$ when $C_1 = 3$ and $C_2 = 1$. Observe that max-min fairness leads to much worse performance for class 1 since the corresponding flows are allocated three times less transmission time than class-2 flows. The corresponding gain for class-2 flows is negligible because these flows generate most of the load.




Figure 11: Flow throughput for the network of figure 10.

It is interesting to compare these results with those obtained for the *maximum* capacity set, as predicted by information theory. We assume that the radio channel can be represented by a Gaussian channel. The system then corresponds to the Gaussian broadcast channel [16], whose capacity set is given for two classes by:

$$\varphi_1 = W \log_2 \left(1 + \frac{P_1}{N_1} \right), \quad \varphi_2 = W \log_2 \left(1 + \frac{P_2}{N_2 + P_1} \right), \quad P_1 + P_2 \leq P,$$

where W is the bandwidth, P is the total transmit power, N_1 is the noise power for class-1 mobiles and $N_2 > N_1$ is the noise power for class-2 mobiles.




Figure 12: A single base station with two classes (broadcast channel).

Unlike the above considered time sharing scheme, the base station transmits *simultaneously* to both classes. Data is jointly coded so as to cancel the interference due to class-2 data at the reception of class-1 data. We take $W = 1$, $P/N_1 = 8.5$ dB and $P/N_2 = 0$ dB so that the maximum bit rates of class 1 and class 2 are approximately equal to 3 and 1 as in the example of figure 10. The capacity set is given by figure 12. Note that, unlike in the previous examples, this is not a polytope. The corresponding flow throughputs are shown in figure 13. Observe that the stability condition is less stringent than with a time sharing scheme.

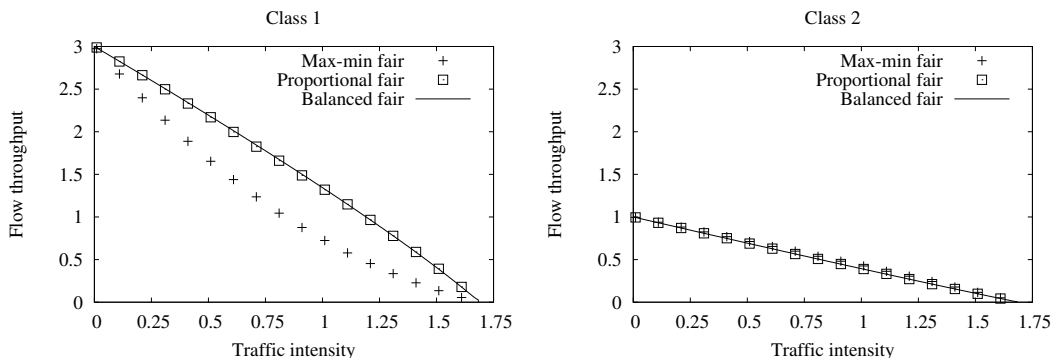


Figure 13: Flow throughput for the network of figure 12.

7.2 Soft hand-over

We now show how the model may be used to represent a soft hand-over mechanism that allows a mobile to receive data from several base stations. This is in fact the analog of the multi-path routing and traffic splitting schemes introduced for wired networks in §6.2-6.3, depending on whether or not the mobile is able to receive data from several base stations at the same time.




Figure 14: Soft hand-over with alternating (top) vs. simultaneous transmission (bottom).

Consider the simple example of figure 14 with two base stations and three classes. The bit rate of class-1 and class-3 flows when scheduled is equal to 3 and the bit rate of class-2 flows when scheduled is equal to 1. If the mobile is not able to receive data from the two base stations at the same time, there are three transmission profiles and we get:

$$C = \begin{pmatrix} 3 & 0 & 3 \\ 3 & 1 & 0 \\ 0 & 1 & 3 \end{pmatrix}.$$

The capacity set, shown in figure 14 (top graph), is the homothetic version of the capacity set of a 3-branch symmetric tree network (cf. figure 1). For balanced fairness, it follows from the homotheticity property (6) and the results of [12] that:

$$\gamma_1 = \gamma_3 = \frac{9 - \varrho}{3} \left(\frac{9 - \varrho}{18 - 5\varrho} + \frac{9(3 - \varrho)}{54 - 9\varrho - \varrho^2} \right)^{-1}, \quad \gamma_2 = \frac{3 - \varrho}{9} \left(\frac{3 - \varrho}{18 - 5\varrho} + \frac{9 - \varrho}{54 - 9\varrho - \varrho^2} \right)^{-1},$$

where ϱ denotes the total traffic intensity $\rho_1 + \rho_2 + \rho_3$. The stability condition is $\varrho < 3$. The corresponding flow throughputs are given in figure 15 (top graphs).

Now if the mobile is able to receive data from both base stations at the same time, there is an additional transmission profile and we get:

$$C = \begin{pmatrix} 3 & 0 & 3 \\ 3 & 1 & 0 \\ 0 & 1 & 3 \\ 0 & 2 & 0 \end{pmatrix}.$$

The capacity set, shown in figure 14 (bottom graph), is the homothetic version of the capacity set of a 3-branch asymmetric tree network. For balanced fairness, we obtain:

$$\gamma_1 = \gamma_3 = \frac{(18 - 5\varrho)(18 - 2\varrho)}{3(36 - 7\varrho)}, \quad \gamma_2 = \frac{18 - 5\varrho}{9},$$




Figure 15: Flow throughput for the wireless network of figure 14 with alternating (top) vs. simultaneous transmission (bottom).

where ρ denotes the total traffic intensity $\rho_1 + \rho_2 + \rho_3$. The stability condition is $\rho < 18/5$. The flow throughputs are given in figure 15 (bottom graphs). Observe that the three allocations differ quite significantly for classes 1 and 3 at moderate load.

7.3 Interference avoidance

In practice, the downlink capacity of a base station is mainly limited by the interference due to other base stations. To limit the impact of interference, one may coordinate the transmission of the base stations. Consider the example of figure 16 with two base stations, one class per base station and three transmission profiles. Two profiles correspond to cases where one base station is on and the other is off, one profile corresponds to the case where both base stations are on. We have in this case:

$$C = \begin{pmatrix} 3 & 0 \\ 0 & 3 \\ 2 & 2 \end{pmatrix}.$$

The corresponding flow throughputs are given in figure 17. Note that the three allocations have very similar performance.

7.4 Ad-hoc networks

Finally, we consider a wireless network without base stations. Such an ad-hoc network consists of a set of L node-to-node wireless links and K routes where each route k is defined as a subset




Figure 16: Two base stations with interference avoidance.




Figure 17: Flow throughput for the network of figure 16.

r_k of the set of links $\{1, \dots, L\}$. We denote by A the incidence matrix as for wired networks. We assume as above that an arbitrary number M of transmission profiles are used in a sequential way. We denote by C the $M \times L$ -dimensional matrix whose m, l -entry gives the capacity of link l in transmission profile m . Nodes can store and relay the packets of any flow. Again, it is assumed that the scheduling operates at a much higher frequency than the typical flow duration. The considered ad-hoc network may then be viewed as a wired network with a varying topology, as shown in the examples below.

We first consider single-path routing where each flow is assigned a particular route in the network. We have $N = K$ in this case. The capacity set is that of a wired network whose link capacities are determined by the schedule:

$$\mathcal{C} = \{\varphi : \exists \tau \in \mathcal{T}, \quad \varphi A \leq \tau C\}.$$

Consider for instance the example of figure 18 with two links and two routes. We assume nodes cannot transmit and receive at the same time, which results in two transmission profiles. We have:

$$A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}.$$




Figure 18: Ad-hoc network with two links and two routes.

In this case, the capacity set coincides with that of a wireless network with a single base station (cf. figure 10).

Another example is given in figure 19. The network consists of three links and three routes. There are again two transmission profiles. We have in this case:

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} 2 & 0 & 2 \\ 0 & 2 & 0 \end{pmatrix}.$$

The capacity set is the homothetic version of that of a 2-link linear wired network. For balanced fairness, it follows from the homotheticity property (6) and the results of [8] that:

$$\gamma_1 = \gamma_3 = 2 - \varrho, \quad \gamma_2 = \frac{(2 - \varrho)(3 - \varrho)}{6 - \varrho},$$

where ϱ denotes the total traffic intensity $\rho_1 + \rho_2 + \rho_3$. The stability condition is $\varrho < 2$. The flow throughputs are given in figure 20.

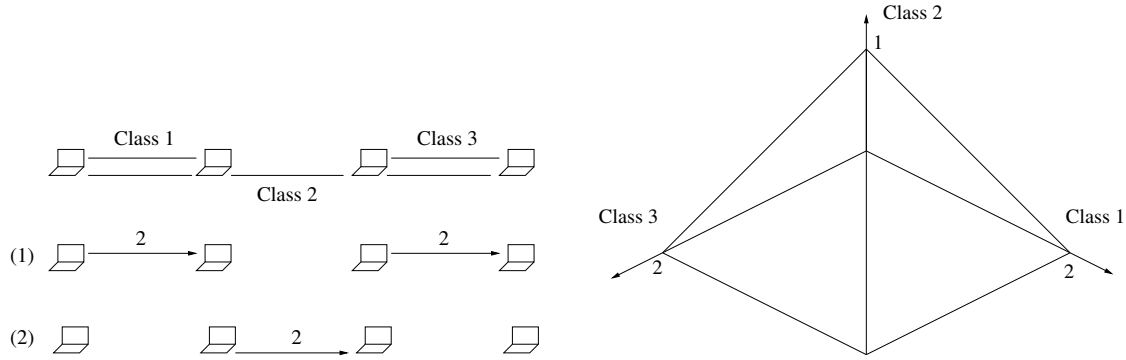


Figure 19: Ad-hoc network with three links and three routes.

8 Conclusion

The issue of resource allocation cannot be addressed without taking the random, dynamic nature of traffic into account. By modeling traffic at flow level, we have shown that a rich class of communication networks can be represented by a queueing network with state-dependent service




Figure 20: Flow throughput for the ad-hoc network of figure 19.

rates. The vector of service rates, which is constrained by some compact, convex capacity set, determines how network resources are allocated.

We proved that the stability condition of the network under max-min fairness, proportional fairness and balanced fairness is that the vector of traffic intensities belongs to the interior of the capacity set. We also proved that max-min fairness does not satisfy the balance property except in the trivial case where the network reduces to a set of independent links. This explains why so few analytical results exist for this allocation and confirms that one should not expect such results. Proportional fairness is not balanced either except in some specific cases (cf. §4.4). Balanced fairness, on the other hand, is balanced by construction and thus leads to a tractable queueing network. No practical implementation of this allocation exists, however. Balanced fairness should rather be considered as a mathematical tool useful for the performance evaluation of more practical allocations like max-min fairness and proportional fairness.

The numerical results of sections 6 and 7 show that balanced fairness indeed constitutes a good approximation of proportional fairness in most cases. Balanced fairness underestimates the performance of proportional fairness in some cases (cf. figures 7, 15), but remains a good approximation at high loads in all cases. This is confirmed by the preliminary results of Massoulié suggesting that the allocation $\phi(x)$ of proportional fairness and balanced fairness coincide for large states x [29].

Max-min fairness, on the other hand, may lead to much worse performance than balanced fairness and proportional fairness (cf. figures 11, 13, 15, 20). This is due to the fact that max-min fairness gives an absolute priority to flows with small bit rates. In wireless networks, this results in an inefficient allocation where flows that experience bad radio conditions consume most radio resources. Proportional fairness and balanced fairness, on the other hand, are *homothetic* so that the allocated network resources do not depend on the radio conditions (cf. the example of figure 10). These allocations are much more efficient and robust than max-min fairness in such networks with heterogeneous flow classes.

A natural extension of the present work is the case where the capacity set is a function of the network state. Such a model is representative of a wireless network where the base stations interact through interference depending on their activity state for instance. If the base stations do not cooperate as in the example of figure 16, the model does not belong to the framework described in the present paper. The stability condition is unknown in general and the way resources should be allocated remains an open issue.

Acknowledgements

We thank the reviewers for their very useful comments.

Appendix

A Proof of theorem 1

Consider a capacity set \mathcal{C} for which max-min fairness has the balance property. We denote by C_i the intersection of the boundary of the capacity set \mathcal{C} with the class- i axis. We say that two classes i, j are equivalent if $C_i = C_j$ and $\mathcal{C}(e_i + e_j)$ is the set of vectors $\varphi = \varphi_i e_i + \varphi_j e_j$ such that $\varphi_i + \varphi_j \leq C_i$. Note that this relationship is reflexive and symmetric. The following lemma shows that it is transitive. It thus defines an *equivalence* relationship.

Lemma 1 *Consider three classes, say 1,2,3, such that classes 1,2 are equivalent and classes 1,3 are equivalent. Then classes 2,3 are equivalent.*

Proof. We let $C_1 = 1$ without any loss of generality. We then have $C_1 = C_2 = C_3 = 1$. By the convexity of the capacity set, $\mathcal{C}(e_2 + e_3)$ contains the set of vectors $\varphi = \varphi_2 e_2 + \varphi_3 e_3$ with $\varphi_2 + \varphi_3 \leq 1$. Since the max-min fair allocation $\phi(x)$ belongs to the boundary of \mathcal{C} for any state x , we deduce that either $\phi_2(e_2 + e_3) \geq 1/2$ or $\phi_3(e_2 + e_3) \geq 1/2$. Using the fact that $\phi_1(e_1 + e_3) = 1/2$ and $\phi_1(e_1 + e_2) = 1/2$, it follows from the balance property applied to state $e \equiv e_1 + e_2 + e_3$ that either $\phi_1(e) \leq \phi_2(e)$ or $\phi_1(e) \leq \phi_3(e)$. Now since $\phi_2(e_1 + e_2) = 1/2$ and $\phi_3(e_1 + e_3) = 1/2$, it follows again from the balance property applied to state e that $\phi_2(e) = \phi_3(e)$. We conclude that $\phi_1(e) \leq \phi_2(e) = \phi_3(e)$.

Now assume that $\phi_1(e) < \phi_2(e)$. In view of the water-filling procedure described in §4.3, this implies that there is a segment from $\phi_1(e)e$ to $\phi(e)$ that belongs to the boundary of \mathcal{C} . The direction of this segment is $e_2 + e_3$. It then follows from the convexity of \mathcal{C} that $\phi_1(e) \geq 1$. Thus $\phi_2(e) > 1$, which contradicts the coordinate convexity of \mathcal{C} . We deduce that $\phi_1(e) = \phi_2(e)$ and $\phi_1(e) = \phi_3(e)$. Since $\phi_1(e_1 + e_3) = 1/2$ and $\phi_1(e_1 + e_2) = 1/2$, it follows from the balance property applied to state e that $\phi_2(e_2 + e_3) = 1/2$ and $\phi_3(e_2 + e_3) = 1/2$. Since $\phi(e_2 + e_3)$ belongs to the boundary of \mathcal{C} , we obtain by the convexity of the capacity set that $\mathcal{C}(e_2 + e_3)$ is the set of vectors $\varphi = \varphi_2 e_2 + \varphi_3 e_3$ such that $\varphi_2 + \varphi_3 \leq 1$. \square

To prove theorem 1, we need the following additional two results.

Lemma 2 *Consider n equivalent classes, say $1, \dots, n$, for some $n \geq 2$. Then $\mathcal{C}(e_1 + \dots + e_n)$ is the set of vectors $\varphi = \varphi_1 e_1 + \dots + \varphi_n e_n$ such that $\varphi_1 + \dots + \varphi_n \leq C_1$.*

Proof. The proof is by induction on n . The property holds for $n = 2$. Assume that it holds for $n - 1$ classes, for some $n \geq 3$. Consider n equivalent classes. We let $C_1 = \dots = C_n = 1$ without any loss of generality. We denote by e the state $e_1 + \dots + e_n$.

By the induction assumption, we have $\phi_1(e - e_i) = \phi_i(e - e_1)$ for all $i = 2, \dots, n$. It then follows from the balance property applied to state e that $\phi_1(e) = \phi_i(e)$ for all $i = 2, \dots, n$. By the convexity of the capacity set, $\mathcal{C}(e)$ contains the set of vectors φ such that $\varphi_1 + \dots + \varphi_n \leq 1$. Since $\phi(e) \equiv \phi_1(e)e$ belongs to the boundary of \mathcal{C} , we deduce that $\phi_1(e) \geq 1/n$. By the convexity of \mathcal{C} , the result will be proved if we show that $\phi_1(e) = 1/n$.

Assume $\phi_1(e) > 1/n$. By the induction assumption, we know that $\phi_i(e + e_1 - e_j) = 1/n$ for all $i, j = 2, \dots, n, i \neq j$. Applying the balance property to state $e + e_1$, we deduce that

$$\phi_2(e + e_1) = \dots = \phi_n(e + e_1).$$

Using again the induction assumption, we know that $\phi_1(e + e_1 - e_2) = 2/n$. Since $\phi_1(e) = \phi_2(e)$, we have $\phi_2(e) > 1/n$ and it follows from the balance property applied to state $e + e_1$ that $\phi_1(e + e_1)/2 < \phi_2(e + e_1)$. Thus:

$$\phi_1(e + e_1)/2 < \phi_2(e + e_1) = \dots = \phi_n(e + e_1).$$

In view of the water-filling procedure, this implies that there is a segment from $\phi_1(e + e_1)(e + e_1)/2$ to $\phi(e + e_1)$ that belongs to the boundary of \mathcal{C} . The direction of this segment is $e_2 + \dots + e_n$. It then follows from the convexity of \mathcal{C} that $\phi_1(e + e_1) \geq 1$. Thus $\phi_2(e + e_1) > 1/2$. By the coordinate convexity of the capacity set, the vector $\phi_1(e + e_1)e_1 + \phi_2(e + e_1)e_2$ belongs to $\mathcal{C}(e_1 + e_2)$. Since classes 1,2 are equivalent, we deduce that $\phi_1(e + e_1) + \phi_2(e + e_1) \leq 1$, a contradiction. \square

Lemma 3 Consider n classes, say $1, \dots, n$, such that no two different classes are equivalent. Then $\mathcal{C}(e_1 + \dots + e_n)$ is the set of vectors $\varphi = \varphi_1 e_1 + \dots + \varphi_n e_n$ such that $\varphi_1 \leq C_1, \dots, \varphi_n \leq C_n$.

Proof. The proof is by induction on n . The property holds for $n = 1$. Assume that it holds for $n - 1$ classes, for some $n \geq 2$. Consider n classes such that no two different classes are equivalent. We denote by e the vector $e_1 + \dots + e_n$ and by C the vector $C_1 e_1 + \dots + C_n e_n$. Note that, by the coordinate convexity of \mathcal{C} , it is sufficient to prove that $C \in \mathcal{C}$ to conclude the proof. We assume without any loss of generality that $\phi_1(e) \leq \phi_2(e) \leq \dots \leq \phi_n(e)$ and distinguish between two cases:

- If $\phi_1(e) < \phi_n(e)$, it follows from the water-filling procedure that there is a segment from $\phi_1(e)e$ to $\phi(e)$ that belongs to the boundary of \mathcal{C} . By the induction assumption, all vectors $\varphi = \varphi_1 e_1 + \dots + \varphi_n e_n$ such that $\varphi_1 = C_1, \varphi_2 \leq C_2, \dots, \varphi_n \leq C_n$ and $\varphi_i = 0$ for some $i = 2, \dots, n$ belong to \mathcal{C} . By the convexity of \mathcal{C} , we deduce that $\phi_1(e) = C_1$. Since $\phi_i(e - e_i) = C_i$ and $\phi_1(e - e_i) = C_1$ for all $i = 2, \dots, n$ by the induction assumption, we deduce from the balance property applied to state e that $\phi_i(e) = C_i$ for all $i = 2, \dots, n$. Thus $C \in \mathcal{C}$.
- If $\phi_1(e) = \phi_n(e)$, we have $\phi_1(e) = \phi_2(e) = \dots = \phi_n(e)$. By the induction assumption, we have $\phi_i(e - e_j) = C_i$ for all $i, j = 1, \dots, n, i \neq j$ so that, applying the balance property to state e , $C_1 = \dots = C_n$. If $\phi_1(e) = C_1$, we have $\phi_i(e) = C_i$ for all $i = 1, \dots, n$ and $C \in \mathcal{C}$. Since $\phi_1(e) \leq C_1$ by the coordinate convexity of \mathcal{C} , we now assume that $\phi_1(e) < C_1$. By the convexity of \mathcal{C} , we have $\phi_2(e) \geq (n - 1)C_1/n$. Note that the inequality is strict for $n = 2$ since classes 1,2 are not equivalent. By the induction assumption, we know that $\phi_1(e + e_1 - e_2) = C_1$. Applying the balance property to state $e + e_1$ we get

$$\phi_1(e + e_1)\phi_2(e) = C_1\phi_2(e + e_1),$$

so that $(n - 1)\phi_1(e + e_1)/n \leq \phi_2(e + e_1)$, with a strict inequality for $n = 2$. Noting that $1/2 < (n - 1)/n$ for $n > 2$, we obtain $\phi_1(e + e_1)/2 < \phi_2(e + e_1)$ for any $n \geq 2$. In view of the water-filling procedure, this implies that there is a segment from $\phi_1(e + e_1)(e + e_1)/2$ to $\phi(e + e_1)$ that belongs to the boundary of \mathcal{C} . Using the same argument as above, we

deduce from the convexity of \mathcal{C} that $\phi_1(e + e_1) = C_1$. In view of the above equality, we get $\phi_2(e) = \phi_2(e + e_1)$. Since $\phi_1(e) < C_1 = \phi_1(e + e_1)$, there is a segment from $\phi(e)$ to $\phi(e + e_1)$ that belongs to the boundary of \mathcal{C} . Again, this implies that $\phi_2(e) = C_2$. Since $\phi_1(e) = \phi_2(e)$ and $C_1 = C_2$, we get $\phi_1(e) = C_1$, a contradiction. \square

We show by induction on N that

$$\mathcal{C} = \left\{ \varphi : \sum_{i \in I_1} \varphi_i \leq c_1, \dots, \sum_{i \in I_L} \varphi_i \leq c_L \right\},$$

where I_1, \dots, I_L are the equivalence classes, that form a partition of the set $\{1, \dots, N\}$, and $c_1 = C_{i_1}, \dots, c_L = C_{i_L}$ for any $i_1 \in I_1, \dots, i_L \in I_L$. The property holds for $N = 1$. Assume it holds for $N - 1$ classes, for some $N \geq 2$. We denote by e the vector $e_1 + \dots + e_N$. If $L = N$, the result follows by lemma 3. If $L < N$, there exists an equivalence class that contains at least two elements, say 1,2. By the induction assumption, we have:

$$\mathcal{C}(e - e_1) = \left\{ \varphi : \varphi_1 = 0, \sum_{i \in I_1} \varphi_i \leq c_1, \dots, \sum_{i \in I_L} \varphi_i \leq c_L \right\}$$

and

$$\mathcal{C}(e - e_2) = \left\{ \varphi : \varphi_2 = 0, \sum_{i \in I_1} \varphi_i \leq c_1, \dots, \sum_{i \in I_L} \varphi_i \leq c_L \right\}.$$

By the convexity of \mathcal{C} , we deduce that

$$\left\{ \varphi : \sum_{i \in I_1} \varphi_i \leq c_1, \dots, \sum_{i \in I_L} \varphi_i \leq c_L \right\} \subset \mathcal{C}.$$

The result then follows from lemma 2 which, by the coordinate convexity of \mathcal{C} , implies that:

$$\mathcal{C} \subset \left\{ \varphi : \sum_{i \in I_1} \varphi_i \leq c_1, \dots, \sum_{i \in I_L} \varphi_i \leq c_L \right\}.$$

B Proof of theorem 2

If $\rho \in \check{\mathcal{C}}$, there exists $\varepsilon > 0$ such that $(1 + \varepsilon)\rho \in \mathcal{C}$. For max-min fairness and proportional fairness, we use Foster's criterion to prove that X is positive recurrent [15]. Since the Markov process X has bounded transition rates (by the compactness of the capacity set \mathcal{C}) and a finite number of transitions from any state, we know that it is positive recurrent if for some finite set $A \subset \mathbb{Z}_+^N$, some non-negative function f and some constant $\delta > 0$,

$$\forall x \notin A, \quad \Delta f(x) \stackrel{\text{def}}{=} \sum_{y \neq x} q(x, y)(f(y) - f(x)) \leq -\delta.$$

Max-min fairness. Since $(1 + \varepsilon)\rho \in \mathcal{C}$, it follows from the definition of max-min fairness and the coordinate convexity of the capacity set \mathcal{C} that for all states $x \neq 0$,

$$\phi_i(x) \geq x_i \min_{j=1, \dots, N} \frac{(1 + \varepsilon)\rho_j}{x_j}.$$

Let $\rho_0 \stackrel{\text{def}}{=} \min_{i=1, \dots, N} \rho_i$. Following Fayolle et al. [18], consider the function f defined by:

$$f(x) = \sum_{i=1, \dots, N} \sum_{n=1}^{x_i} \gamma^{\frac{n}{\rho_i}},$$

where γ is some constant such that $\gamma > 1$ and

$$\frac{\gamma^{\frac{1}{\rho_0}}}{1 + \varepsilon} \leq \theta < 1.$$

We have:

$$\Delta f(x) = \sum_{i=1, \dots, N} \left(\lambda_i \gamma^{\frac{x_i+1}{\rho_i}} - \frac{\phi_i(x)}{\sigma_i} \gamma^{\frac{x_i}{\rho_i}} \right) \leq \sum_{i=1, \dots, N} \lambda_i \gamma^{\frac{x_i}{\rho_i}} (1 + \varepsilon) \left(\theta - \frac{x_i}{\rho_i} \min_{j=1, \dots, N} \frac{\rho_j}{x_j} \right).$$

Let $k \stackrel{\text{def}}{=} \arg \min_{j=1, \dots, N} \frac{\rho_j}{x_j}$. We decompose the last sum as follows:

$$\begin{aligned} \sum_{i: \frac{x_i}{\rho_i} > \theta \frac{x_k}{\rho_k}} \lambda_i \gamma^{\frac{x_i}{\rho_i}} (1 + \varepsilon) \left(\theta - \frac{x_i \rho_k}{\rho_i x_k} \right) &\leq \lambda_k \gamma^{\frac{x_k}{\rho_k}} (1 + \varepsilon) (\theta - 1), \\ \sum_{i: \frac{x_i}{\rho_i} \leq \theta \frac{x_k}{\rho_k}} \lambda_i \gamma^{\frac{x_i}{\rho_i}} (1 + \varepsilon) \left(\theta - \frac{x_i \rho_k}{\rho_i x_k} \right) &\leq \gamma^{\theta \frac{x_k}{\rho_k}} (1 + \varepsilon) \theta \sum_{i=1, \dots, N} \lambda_i. \end{aligned}$$

We obtain:

$$\Delta f(x) \leq \gamma^{\theta \frac{x_k}{\rho_k}} \left((1 + \varepsilon)(\theta - 1) \min_{i=1, \dots, N} \lambda_i + \gamma^{(\theta-1) \frac{x_k}{\rho_k}} (1 + \varepsilon) \theta \sum_{i=1, \dots, N} \lambda_i \right).$$

For any $\delta > 0$, there exists a constant C such that

$$(1 + \varepsilon)(\theta - 1) \min_{i=1, \dots, N} \lambda_i + \gamma^{(\theta-1)C} (1 + \varepsilon) \theta \sum_{i=1, \dots, N} \lambda_i \leq -\delta.$$

Since $\gamma > 1$, we deduce $\Delta f(x) \leq -\delta$ for all states $x \notin A$, with $A = \{x \in \mathbb{Z}_+^N : x \leq C\rho\}$.

Proportional fairness. Since $(1 + \varepsilon)\rho \in \mathcal{C}$, it follows from the definition of proportional fairness that for all states x ,

$$\sum_{i=1}^N x_i \log((1 + \varepsilon)\rho_i) \leq \sum_{i=1}^N x_i \log(\phi_i(x)).$$

The function $g : [0, 1] \rightarrow \mathbb{R}$ defined by

$$g(\theta) = \sum_{i=1}^N x_i \log((1 - \theta)(1 + \varepsilon)\rho_i + \theta\phi_i(x))$$

is concave and satisfies $g(0) \leq g(1)$. We deduce $g'(0) \geq 0$, that is

$$\sum_{i=1}^N \frac{x_i}{\rho_i} (\phi_i(x) - (1 + \varepsilon)\rho_i) \geq 0.$$

We obtain:

$$\sum_{i=1}^N \frac{x_i}{\lambda_i} (\lambda_i - \frac{\phi_i(x)}{\sigma_i}) \leq -\varepsilon \sum_{i=1}^N x_i.$$

Let f be the function defined by:

$$f(x) = \sum_{i=1}^N \frac{x_i^2}{2\lambda_i}.$$

Using the fact that for all states $x \in \mathbb{Z}_+^N$,

$$f(x + e_i) = f(x) + \frac{x_i}{\lambda_i} + \frac{1}{2\lambda_i} \quad \text{and} \quad f(x - e_i) = f(x) - \frac{x_i}{\lambda_i} + \frac{1}{2\lambda_i} \quad \text{if } x_i > 0,$$

we get:

$$\begin{aligned} \Delta f(x) &= \sum_{i=1}^N \frac{x_i}{\lambda_i} (\lambda_i - \frac{\phi_i(x)}{\sigma_i}) + \sum_{i=1}^N \frac{1}{2\lambda_i} (\lambda_i + \frac{\phi_i(x)}{\sigma_i}) \\ &\leq -\varepsilon \sum_{i=1}^N x_i + C, \end{aligned}$$

where C is some finite constant. Thus there exists some finite set $A \subset \mathbb{Z}_+^N$ and some constant $\delta > 0$ such that $\Delta f(x) \leq -\delta$ for all states $x \notin A$.

Balanced fairness. We need the following preliminary result.

Lemma 4 *Let Φ be the balance function of balanced fairness. The balance function $\tilde{\Phi}$ associated with any other balanced allocation satisfies $\tilde{\Phi}(x) \geq \Phi(x)$ for all states $x \in \mathbb{Z}_+^N$.*

Proof. The proof is by induction on $|x| \stackrel{\text{def}}{=} \sum_{i=1}^N x_i$. Since $\tilde{\Phi}(0) = \Phi(0) = 1$, the inequality is satisfied for $x = 0$. Assume it is satisfied for all states x such that $|x| = n$. Let $x \in \mathbb{Z}_+^N$ be such that $|x| = n + 1$. Since

$$\left(\frac{\tilde{\Phi}(x - e_1)}{\tilde{\Phi}(x)}, \dots, \frac{\tilde{\Phi}(x - e_N)}{\tilde{\Phi}(x)} \right) \in \mathcal{C},$$

it follows from the induction assumption and the coordinate convexity of \mathcal{C} that

$$\left(\frac{\Phi(x - e_1)}{\tilde{\Phi}(x)}, \dots, \frac{\Phi(x - e_N)}{\tilde{\Phi}(x)} \right) \in \mathcal{C}.$$

By the definition of $\Phi(x)$, we deduce that $\tilde{\Phi}(x) \geq \Phi(x)$. □

Since $(1 + \varepsilon)\rho \in \mathcal{C}$, the allocation defined by $\tilde{\phi}(x) = (1 + \varepsilon)\rho$ in all states x is feasible and balanced. The corresponding balance function is given by

$$\tilde{\Phi}(x) = \frac{1}{((1 + \varepsilon)\rho_1)^{x_1}} \cdots \frac{1}{((1 + \varepsilon)\rho_N)^{x_N}}.$$

The result is then a simple consequence of lemma 4:

$$\sum_x \Phi(x)\rho_1^{x_1} \cdots \rho_N^{x_N} \leq \sum_x \tilde{\Phi}(x)\rho_1^{x_1} \cdots \rho_N^{x_N} = \sum_x \frac{1}{(1 + \varepsilon)^{x_1}} \cdots \frac{1}{(1 + \varepsilon)^{x_N}} < \infty.$$

References

- [1] M. Armony and N. Bambos, Queueing dynamics and maximal throughput scheduling in switched processing systems, *Queueing Systems* 44(3) (2003) 209–252.
- [2] N. Bambos and G. Michailidis, Queueing and scheduling in random environments, *Adv. Appl. Prob.* 36 (2004) 293–317.
- [3] S. Ben Fredj, T. Bonald, A. Proutière, G. Régnié and J.W. Roberts, Statistical bandwidth sharing: A study of congestion at flow level, in: *Proc of ACM SIGCOMM*, 2001.
- [4] A.W. Berger, Y. Kogan, Dimensioning bandwidth for elastic traffic in high-speed data networks, *IEEE/ACM Trans. on Networking* 8(5) (2000) 643–654.
- [5] D. Bertsekas and R. Gallager, *Data Networks*, Prentice Hall, 1987.
- [6] T. Bonald, M. Jonckheere and A. Proutière, Insensitive load balancing, in: *Proc. of ACM SIGMETRICS/ Performance*, 2004.
- [7] T. Bonald and L. Massoulié, Impact of fairness on Internet performance, in: *Proc. of ACM SIGMETRICS/ Performance*, 2001.
- [8] T. Bonald and A. Proutière, Insensitive bandwidth sharing in data networks, *Queueing Systems* 44(1) (2003) 69–100.
- [9] T. Bonald and A. Proutière, On performance bounds for balanced fairness, *Performance Evaluation* 55 (2004) 25–50.
- [10] T. Bonald and A. Proutière, On performance bounds for the integration of elastic and adaptive streaming flows, in: *Proc. of ACM SIGMETRICS/ Performance*, 2004.
- [11] T. Bonald, A. Proutière, J. Roberts and J. Virtamo, Computational aspects of balanced fairness, in: *Proc. of ITC 18*, 2003.
- [12] T. Bonald and J. Virtamo, Calculating the flow level performance of balanced fairness in tree networks, *Performance Evaluation* 58 (2004) 1–14.
- [13] T. Bonald and J. Virtamo, A recursive formula for multirate systems with elastic traffic, *IEEE Communications Letters* 9 (2005) 753–755.

- [14] J. Boyer, F. Guillemin, P. Robert and B. Zwart, Heavy tailed M/G/1-PS queues with impatience and admission control in packet networks, in: *Proc. of Infocom*, 2003.
- [15] P. Brémaud, *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*, Springer, New York, 1999.
- [16] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, 1991.
- [17] F. Delcoigne, A. Proutière, G. Régnié, Modeling integration of streaming and data traffic, *Performance Evaluation* 55(3-4) (2004) 185–209.
- [18] G. Fayolle, A. de la Fortelle, J.M. Lasgouttes, L. Massoulié and J.W. Roberts, Best-effort networks: Modeling and performance analysis via large networks asymptotics, in: *Proc. of IEEE INFOCOM*, 2001.
- [19] G. Fayolle, I. Mitrani, and R. Iasnogorodski, Sharing a processor among many job classes, *Journal of the ACM* 27 (1980) 519–532.
- [20] D.P. Heyman, T.V. Lakshman, A.L. Neidhardt, A new method for analysing feedback-based protocols with applications to engineering Web traffic over the Internet, in: *Proc. of ACM SIGMETRICS*, 1997.
- [21] D. Julian, M. Chiang, D. O’Neill and S. Boyd, QoS and fairness constrained convex optimization of resource allocation for wireless cellular and ad hoc networks, in: *Proc. of IEEE INFOCOM*, 2002.
- [22] F.P. Kelly, A. Maulloo and D. Tan, Rate control for communication networks: Shadow prices, proportional fairness and stability, *Journal of the Operat. Res. Society* 49 (1998) 237–252.
- [23] F.P. Kelly and R.J. Williams, Fluid model for a network operating under a fair bandwidth-sharing policy, *Annals of Applied Probability* 14 (2004) 1055–1083.
- [24] P. Key, L. Massoulié, A. Bain and F. Kelly, A network flow model for mixtures of file transfers and streaming traffic, in: *Proc. of ITC 18*, 2003.
- [25] A. Kherani and A. Kumar, Stochastic models for throughput analysis of randomly arriving elastic flows in the Internet, in: *Proc. of IEEE INFOCOM*, 2002.
- [26] S. Kunniyur and R. Srikant, End-to-end congestion control schemes: Utility functions, random losses and ECN marks, in: *Proc. of IEEE INFOCOM*, 2000.
- [27] J. Leino and J. Virtamo, Insensitive traffic splitting in data networks, in: *Proc. of ITC 19*, 2005.
- [28] X. Liu, E.K.P. Chong and N.B. Shroff, A framework for opportunistic scheduling in wireless networks, *Computer Networks* 41(4) (2003) 451–474.
- [29] L. Massoulié, Structural properties of proportional fairness: Stability and insensitivity, submitted, 2005.
- [30] L. Massoulié and J.W. Roberts, Arguments in favour of admission control for TCP flows, in: *Proc. of ITC 16*, 1999.

- [31] L. Massoulié and J.W. Roberts, Bandwidth sharing and admission control for elastic traffic, *Telecommunication Systems* 15 (2000) 185–201.
- [32] L. Massoulié and J.W. Roberts, Bandwidth sharing: Objectives and algorithms, *IEEE/ACM Trans. on Networking* 10(3) (2002) 320–328.
- [33] J. Mo and J. Walrand, Fair end-to-end window-based congestion control, *IEEE/ACM Trans. on Networking* 8(5) (2000) 556–567.
- [34] T. Nandagopal, T. Kim, X. Gao and V. Bharghavan, Achieving MAC layer fairness in wireless packet networks, in: *Proc. of ACM MOBICOM*, 2000.
- [35] R. Núñez Queija, J.L. van den Berg, M.R.H. Mandjes, Performance evaluation of strategies for integration of elastic and stream traffic, in: *Proc. of ITC 16*, 1999.
- [36] V. Paxson and S. Floyd, Difficulties in Simulating the Internet, *IEEE/ACM Trans. on Networking* 9(4) (2001) 392–403.
- [37] A. Penttinen and J. Virtamo, Performance of wireless ad hoc networks under balanced fairness, in: *Proc. of Networking*, 2004.
- [38] B. Radunovic, J.Y. Le Boudec, Rate performance objectives of multihop wireless networks, *IEEE Trans. on Mobile Computing* 3(4) (2004) 334–349.
- [39] J. Rawls, *A Theory of Justice*, Belknap Press, 1971.
- [40] J.W. Roberts, A survey on statistical bandwidth sharing, *Computer Networks* 45 (2004) 319–332.
- [41] R.F. Serfozo, *Introduction to Stochastic Networks*, Springer Verlag, 1999.
- [42] G.M. Stamatelos and V.N. Koukoulidis, Reservation-based bandwidth allocation in a radio ATM network, *IEEE/ACM Trans. on Networking* 5(3) (1997) 420–428.
- [43] L. Tassiulas and A. Ephremides, Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multi-hop radio networks, *IEEE Trans. Aut. Control* 37 (1992) 1936–1938.
- [44] I.E. Telatar and R.G. Gallager, Combining queueing theory with information theory for multiaccess, *IEEE Journal Selected Areas in Comm.* 13 (1995) 963–969.
- [45] G. de Veciana, T.J. Lee and T. Konstantopoulos, Stability and performance analysis of networks supporting elastic services, *IEEE/ACM Trans. on Networking* 9(1) (2001) 2–14.
- [46] H.Q. Ye, Stability of data networks under an optimization-based bandwidth allocation, *IEEE Trans. Aut. Control* 48(7) (2003) 1238–1242.