



HAL
open science

Insensitive bandwidth sharing in data networks

Thomas Bonald, Alexandre Proutière

► **To cite this version:**

Thomas Bonald, Alexandre Proutière. Insensitive bandwidth sharing in data networks. Queueing Systems, 2003. hal-01267144

HAL Id: hal-01267144

<https://hal.science/hal-01267144v1>

Submitted on 4 Feb 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Insensitive bandwidth sharing in data networks*

T. Bonald and A. Proutière

France Telecom R&D

38-40, rue du Général Leclerc

92794 Issy-les-Moulineaux Cedex 9, France

{Thomas.Bonald,Alexandre.Proutiere}@francetelecom.com

December 16, 2002

Abstract

We represent a data network as a set of links shared by a dynamic number of competing flows. These flows are generated within sessions and correspond to the transfer of a random volume of data on a pre-defined network route. The evolution of the stochastic process describing the number of flows on all routes, which determines the performance of the data transfers, depends on how link capacity is allocated between competing flows.

We use some key properties of Whittle queueing networks to characterize the class of allocations which are insensitive in the sense that the stationary distribution of this stochastic process does not depend on any traffic characteristics (session structure, data volume distribution) except the traffic intensity on each route. We show in particular that this insensitivity property does not hold in general for well-known allocations such as max-min fairness or proportional fairness. These results are illustrated by several examples on a number of network topologies.

1 Introduction

The majority of traffic in current data networks is elastic. Traffic is composed of flows transporting digital documents of one form or another and the rate of these flows adjusts with respect to the congestion level, typically under the control of TCP [13]. In practice, the rate of a flow does not only depend on the capacity of links on its path and the number of competing flows in progress, but on many other parameters including the version of TCP used by the sources and the scheduling and buffer management schemes implemented in network nodes. However, to gain insight into the performance of data networks it is useful to make some simplifying assumptions with regard to the way bandwidth is shared.

1.1 Utility-based allocations

The way bandwidth is shared in current data networks or should be shared in future data networks has been the subject of considerable recent research. Max-min fairness, where the rate of individual flows is made as equal as possible [3], has long been stated as an ideal objective. Kelly, in particular, has questioned this accepted wisdom and suggests bandwidth should be allocated in order to maximize some overall utility, where the utility of an individual flow is a function of its rate only. The notion of proportional fairness was introduced and shown to be realized by a certain distributed congestion control

*to appear in Queueing Systems

algorithm [16]. In line with this approach, a number of studies have identified the utility function that corresponds to the allocation realized by TCP or some idealized version of TCP [15, 17, 18, 19, 21, 23]. It turns out that max-min fairness may also be viewed as a utility-based allocation with a particular limit definition of per flow utility [23].

Network utility in the above mentioned studies is determined in a *static* scenario, i.e., with a fixed number of permanent flows. Flows do not last indefinitely in practice. Each flow corresponds to the transfer of a finite volume of data (referred to as the flow size) and ceases when the transfer is completed. The evolution of the number of flows in progress clearly depends on the way new flows are generated, their sizes, and the way bandwidth is shared between competing flows. In particular, the fact that an allocation is optimal in the sense of some utility function in a static scenario does not necessarily imply that this allocation is optimal in a dynamic scenario. The allocation may well lead to a steady state where overall utility is low. For instance, maximizing the mean flow rate in some network topologies may lead to instability under the usual traffic conditions in the sense that the number of flows increases indefinitely [4]. In this case, allocating link capacities to maximize the mean flow rate in a static scenario in fact minimizes the mean flow rate in a dynamic scenario since the latter is zero in steady state. This example illustrates the fact that bandwidth sharing objectives cannot reasonably be defined without taking flow-level dynamics into account.

The study of flow-level dynamics in data networks proves difficult in general, even for the simplest network topologies [11]. Prior to the present work, explicit performance results were only available for proportional fairness in homogeneous “lines” and “grids” [4]. We prove here that, except for proportional fairness in so-called homogeneous “hypercubes” (the multi-dimensional generalization of lines and grids), utility-based allocations are *sensitive* in the sense that the steady state distribution depends on detailed traffic characteristics. This notably explains why the analysis of flow-level dynamics is so hard for these allocations. It also suggests that these allocations are unlikely to be optimal in any sense that is independent of the detailed traffic characteristics such as the flow arrival process and the flow size distribution.

1.2 Insensitive allocations

The above observation leads to the following question: is it possible to define an allocation which is *insensitive* in the sense that the steady state distribution does not depend on any traffic characteristics except the traffic intensity on each network path? Such an insensitivity property is the key to simple and robust performance results. Network provisioning rules can then be developed based on traffic intensity forecasts only, independently of the complex traffic structure which is continually evolving as new applications emerge. The practical value of insensitivity is best illustrated by the enduring success of Erlang’s loss formula in telephone networks [10]. This formula gives the proportion of calls that are blocked as a simple function of capacity and traffic intensity, independently of the distribution of call durations. The only required assumption is that calls arrive as a Poisson process, which is verified in practice as calls are generated independently by a large number of users. This insensitivity property explains why Erlang’s formula is still used for dimensioning current telephone networks, despite the fact that telephone traffic characteristics have changed considerably since Erlang’s publication in 1917.

The first insensitivity result for elastic traffic was given in [20] for a single bottleneck whose capacity is fairly shared between flows in progress. The underlying model is the processor-sharing queue. In particular, assuming Poisson arrivals of flows with i.i.d. sizes, the distribution of the number of flows in progress in steady state is insensitive to the flow size distribution. In fact, flows do not arrive as a Poisson process in data networks. Flows form part of sessions, each session being composed of a succession of flows separated by an interval of inactivity generally referred to as a “think-time”. A typical example

is the succession of Web pages downloaded by a user in a period of continuous activity. The resulting flow arrival process may be strongly correlated, depending on the number of flows in a session and the distribution of successive flow sizes and think-time durations [6, 24]. It turns out that the steady state distribution is in fact insensitive to this correlation, provided we assume that *sessions* arrive as a Poisson process. This was proved in [2, 6] using key properties of Kelly queueing networks [14]. The Poisson assumption is reasonable when sessions are generated by a large number of users and has indeed been verified in practice [24]. This insensitivity result still holds when flow rates are all limited by a common fixed constraint referred to as the access rate [2, 6]. The corresponding model is a symmetric queue [14] for which the same arguments indeed apply.

The objective of the present paper is to extend the insensitivity result to any network topology and any access rate constraints (not necessarily the same for all flows). Using key properties of Whittle queueing networks [5, 25], we characterize the class of insensitive allocations and derive explicit results which determine their performance. These allocations differ in general from utility-based allocations. They could be used as bandwidth sharing objectives to be realized by future packet-level mechanisms. However, it is also expected that the performance of these allocations is close to that of the allocations realized by existing packet-level mechanisms such as the congestion control algorithms of TCP. The derived formulas and the resulting insight could then be used to define engineering guidelines for data networks equivalent to those developed for the telephone networks over the years since the discovery of Erlang’s formula.

1.3 Outline

In the next section we describe the considered flow-level model. In Section 3 we characterize the class of insensitive allocations and present key properties satisfied by these allocations. It is demonstrated in Section 4 that this class does not contain utility-based allocations, with the notable exception of proportional fairness in homogeneous “hypercubes”. These results are illustrated on a number of network topologies in Section 5. Section 6 concludes the paper.

2 Flow-level modeling of data networks

In this section, we introduce a generic flow-level model of data networks. We then show how this model can be represented by a processor-sharing queueing network with state-dependent service capacities, with virtually any traffic characteristics.

2.1 Network model

We represent a data network as a set of links $\mathcal{L} = \{1, \dots, L\}$ where each link $l \in \mathcal{L}$ has a capacity $C_l > 0$. A random number of flows compete for access to these links. Each flow is characterized by a volume of information to be transferred (referred to as the flow size) on a route consisting of a set of links. The flows are “elastic” in the sense that their duration depends on their rate which varies as new flows begin and others cease. Specifically, a flow of size s arriving at time t_{start} on route r is completed at time t_{end} given by:

$$s = \int_{t_{\text{start}}}^{t_{\text{end}}} c(t) dt,$$

where $c(t)$ denotes the flow rate at time t , i.e., the capacity allocated to this flow on *each* link of route r at time t , $t_{\text{start}} \leq t \leq t_{\text{end}}$. This rate is limited by the capacity C_l of each link $l \in r$ that is shared between

all flows in progress on route r and on other routes containing link l . It may additionally be constrained by a fixed maximum limit representing external constraints such as the user's access line.

Capacity allocation. We consider K classes of flow in this data network. Each class k is characterized by a route r_k consisting of a non-empty set of links and a per-flow rate limit $a_k > 0$ we refer to as the "access rate". We adopt the convention that either $a_k < \min_{l \in r_k} C_l$, in which case the access rate is limiting, or $a_k = \infty$. We denote by $x = (x_1, \dots, x_K)$ the network state, where x_k is the number of flows of class k in progress. It is assumed that the total capacity ϕ_k allocated to flows of class k is equally shared between these flows and depends on the network state x only. The allocation must satisfy the capacity constraints:

$$\sum_{k:l \in r_k} \phi_k(x) \leq C_l, \quad l = 1, \dots, L, \quad \text{and} \quad \phi_k(x) \leq x_k a_k, \quad k = 1, \dots, K. \quad (1)$$

The allocation is said to be *Pareto-efficient* if for any state x and any class k such that $x_k > 0$, there exists a saturated link l on route r_k or the rate of each flow of class k is maximum, i.e.,

$$\exists l \in r_k, \quad \sum_{k':l \in r_{k'}} \phi_{k'}(x) = C_l \quad \text{or} \quad \phi_k(x) = a_k x_k. \quad (2)$$

Traffic conditions. The evolution of the network state x does not only depend on the way capacity is allocated between flows in progress but on traffic characteristics, i.e., on the way new flows are generated and on the distribution of their size. The traffic characteristics considered in this paper are quite general and described in detail in §2.3-2.4. It is sufficient at this stage to assume that the marked point process of flow arrivals of each class, with marks corresponding to the flow sizes, is stationary and ergodic. Denote by ρ_k the traffic intensity of class k . This corresponds to the mean volume of information offered by flows of class k per unit of time. We refer to the usual traffic conditions as the inequalities:

$$\sum_{k:l \in r_k} \rho_k < C_l, \quad l \in \mathcal{L}. \quad (3)$$

It is worth noting that the Pareto-efficiency of an allocation is not sufficient to ensure network stability under the usual traffic conditions. Specifically, there are Pareto-efficient allocations for which the total number of flows tends to infinity from any initial state although the inequalities (3) are satisfied [4]. The issue of defining stability conditions is still largely open except for some specific allocations and under restrictive assumptions on traffic characteristics [4, 9].

User performance. Users perceive performance essentially through the mean time necessary to transfer a document. In the following, we evaluate performance in terms of *throughput*, defined as the ratio of the mean flow size to the mean flow duration in steady state. Assuming network stability and applying Little's formula [1], the throughput of flows of any class k is related to the expected number of flows of class k in steady state through the relationship:

$$\gamma_k = \frac{\rho_k}{E[x_k]}. \quad (4)$$

In the simplest case where the network reduces to a single link of capacity C and a single class of traffic intensity ρ without limiting access rate, the corresponding model is the processor-sharing queue (provided the allocation is Pareto-efficient). For the general traffic characteristics described in §2.3-2.4 and under

the usual traffic condition $\rho < C$, the number of flows has a geometric distribution of mean ρ/C in steady state [2]. Thus the flow throughput γ is simply given by:

$$\gamma = C - \rho. \tag{5}$$

2.2 A processor-sharing queueing network

We now introduce an open queueing network of processor-sharing nodes with state-dependent service capacities. We show in §2.3-2.4 that this queueing network can represent the data network described in §2.1 with virtually any traffic characteristics (arbitrary flow size distribution, correlated arrivals of flows within sessions, etc). The queueing network consists of N processor-sharing nodes with state-dependent capacities, that is, the capacity (or speed) ψ_i of node i depends on the state $y = (y_1, \dots, y_N)$, where y_i is the number of customers in node i . We only assume that $\psi_i(y) > 0$ if and only if $y_i > 0$. Exogenous arrivals at node i form a Poisson process of rate ν_i . The services required at node i are exponential i.i.d. of mean $1/\mu_i$. After service completion at node i , a customer moves to node j with probability p_{ij} and leaves the network with probability $p_i \equiv 1 - \sum_j p_{ij}$. We assume that each customer eventually leaves the network, so that the effective arrival rate λ_i at node i is uniquely defined by the equations:

$$\lambda_i = \nu_i + \sum_j \lambda_j p_{ji}, \quad i = 1, \dots, N.$$

We denote by $\rho_i = \lambda_i/\mu_i$ the traffic intensity at node i . The stochastic process $Y = \{Y_t, t \geq 0\}$ that describes the evolution of the number of customers at each node is an irreducible Markov process.

2.3 Poisson flow arrivals

Consider the data network of §2.1 where flows of each class arrive as an independent Poisson process. This may be represented by the above considered processor-sharing queueing network, where each customer corresponds to an ongoing flow in case of exponential flow size distributions, or to a phase of an ongoing flow in case of phase-type flow size distributions.

Exponential flow size distribution. If flows have exponential i.i.d. sizes, the corresponding processor-sharing queueing network has $N = K$ nodes and no routing, i.e., $p_{ij} = 0$ for all nodes i, j . The service capacity ψ_i of node i represents the total capacity ϕ_i allocated to flows of class i , which is equally shared between these flows. A simple example is given in Figure 1.

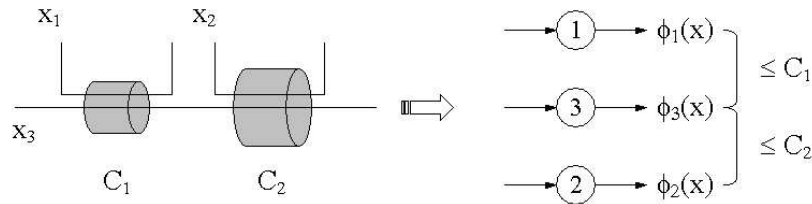


Figure 1: A data network represented as a processor-sharing queueing network

Phase-type flow size distribution. Measurements of the size of flows in real data networks show that their distribution is not exponential but typically much more variable [8]. The considered queueing network allows phase-type distributions, which are known to form a dense subset within the set of all distributions of nonnegative random variables. A phase-type distribution for flows of class k can be represented simply by a set of consecutive nodes $S_k \subset \{1, \dots, N\}$ such that $\nu_i > 0$ for the first node $i \in S_k$ only, and for any node $i \in S_k$, $p_{ij} = 0$ for all nodes j except for $j = i + 1$, if $i + 1 \in S_k$ (refer to Figure 2). As each node $i \in S_k$ represents a phase of flows of the same class k and the capacity $\phi_k(x)$ allocated to flows of class k is fairly shared between these flows, we have

$$\psi_i(y) = \frac{y_i}{x_k} \phi_k(x), \quad \text{with} \quad x_k = \sum_{i \in S_k} y_i.$$

The traffic intensity of flows of class k is given by:

$$\rho_k = \sum_{i \in S_k} \rho_i.$$

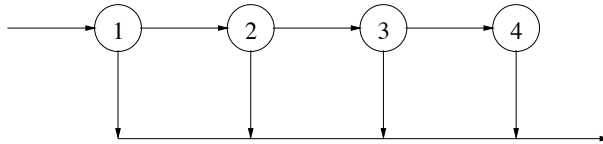


Figure 2: A 4-phase distribution of flow sizes

2.4 Poisson session arrivals

As mentioned in Section 1, flows do not arrive as independent Poisson processes in data networks. They are typically generated within sessions, each session being composed of a succession of flows separated by an interval of inactivity which we call “think-time”. Again, the considered processor-sharing network is sufficiently general to account for this complex structure of traffic, provided sessions arrive as a Poisson process and think-time durations do not depend on the network state (as opposed to flow durations).

Exponential flow size and think-time duration distributions. We first consider the case where successive flow sizes and think-time durations are all exponentially distributed. Think-times can simply be represented by infinite-server nodes, i.e., those nodes i in the set $S_0 \subset \{1, \dots, N\}$ such that:

$$\psi_i(y) = y_i. \tag{6}$$

We still denote by $S_k \subset \{1, \dots, N\}$ the set of nodes representing flows of class k , i.e., such that:

$$\psi_i(y) = \frac{y_i}{x_k} \phi_k(x), \quad x_k = \sum_{i \in S_k} y_i. \tag{7}$$

A session can then be represented as a random walk of a customer in an alternating series of nodes in the sets S_k , $k \neq 0$, and in the set S_0 . That is, for any node $i \notin S_0$, we have $p_{ij} = 0$ for all nodes $j \notin S_0$, and for any node $i \in S_0$, we have $p_{ij} = 0$ for all nodes $j \in S_0$. We assume without loss of generality that

$\nu_i = 0$ and $p_i = 0$ for all nodes $i \in S_0$, which means that a session necessarily starts and ends with a flow (and not a think-time). Again, the traffic intensity of flows of class k is simply given by:

$$\rho_k = \sum_{i \in S_k} \varrho_i. \quad (8)$$

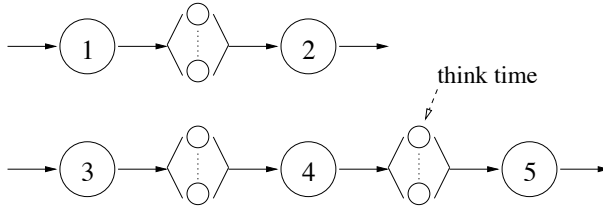


Figure 3: Example of two types of session, composed of two and three flows, respectively

It is worth noting that the distribution of the number of flows per session may be perfectly general. Successive flow sizes and think-time durations may also be correlated. Figure 3 gives an example of two types of session, composed of two and three flows, respectively. The mean flow sizes may well be higher for the first type of session for instance. In fact, arbitrary correlations between successive flow sizes and think-time durations may be represented by considering as many session types as necessary and introducing phase-type distributions.

Phase-type flow size and think-time duration distributions. As in §2.3, assume now that each node represents a phase of a flow or a think-time (and not the flow or the think-time itself). We still denote by S_0 the set of nodes representing think-times, satisfying (6), and S_k the set of nodes representing flows of class k , satisfying (7). Assume without loss of generality that successive phases of the same flow or think-time consist of consecutive nodes. A session with phase-type distributions of flow sizes and think-time durations can be represented as a random walk such that any visit of a customer to a node $i \in S_k$, $k \neq 0$, can be followed by a visit to the node $i + 1 \in S_k$ if this node corresponds to a new phase of the same flow, or a visit to a node $j \in S_0$ representing the first phase of a think-time. Similarly, any visit to a node $i \in S_0$ can be followed by a visit to the node $i + 1 \in S_0$ if this node corresponds to a new phase of the same think-time, or a visit to a node $j \notin S_0$ representing the first phase of a new flow of the same session. The exogenous Poisson processes correspond to new sessions. Specifically, those nodes $i \notin S_0$ such that $\nu_i > 0$ correspond to the first phase of the first flow of a session. The traffic intensity of flows of class k is still given by (8).

3 Insensitive allocations

We now characterize those capacity allocations for which the steady state is insensitive to the above described traffic characteristics. Specifically, we prove in Theorem 1 that the insensitivity property is equivalent to three milder forms of insensitivity, which all imply the balance property. We then give key properties of these allocations and introduce the notion of “balanced fairness”.

3.1 Balance property

Let e_k be the unit vector with 1 in component k and 0 elsewhere, for $k = 1, \dots, K$.

Definition 1 (Balance property) *The capacities ϕ_1, \dots, ϕ_K are said to be balanced if:*

$$\phi_k(x)\phi_{k'}(x - e_k) = \phi_{k'}(x)\phi_k(x - e_{k'}), \quad \forall k, k' \forall x \text{ such that } x_k > 0 \text{ and } x_{k'} > 0.$$

Let $\langle x, x - e_{k_1}, \dots, x - e_{k_1} - \dots - e_{k_{n-1}}, 0 \rangle$ be a direct path from state x to state 0, i.e., a path of length n where $n \equiv \sum_k x_k$ is the number of flows in state x . The balance property implies that the expression

$$\Phi(x) = \frac{1}{\phi_{k_1}(x)\phi_{k_2}(x - e_{k_1}) \dots \phi_{k_n}(x - e_{k_1} - \dots - e_{k_{n-1}})}, \quad (9)$$

is independent of the considered direct path. In particular, the capacities are uniquely characterized by the positive function Φ , referred to as the balance function:

$$\phi_k(x) = \frac{\Phi(x - e_k)}{\Phi(x)}, \quad \forall k, \forall x \text{ such that } x_k > 0. \quad (10)$$

Conversely, if there exists a positive function Φ such that the capacities satisfy (10), it can be easily verified that these capacities are balanced. We say that the capacities are balanced by Φ .

Remark 1 *The balance property may be interpreted as the fact that the relative change in the capacity allocated to class k when a flow of class k' is removed is the same as the relative change in the capacity allocated to class k' when a flow of class k is removed, i.e.,*

$$\frac{\phi_k(x - e_{k'})}{\phi_k(x)} = \frac{\phi_{k'}(x - e_k)}{\phi_{k'}(x)}, \quad \forall x \text{ such that } x_k > 0 \text{ and } x_{k'} > 0.$$

3.2 Sufficient condition for insensitivity

Consider an allocation for which the balance property holds. The processor-sharing queueing network introduced in §2.2 can represent virtually any traffic characteristics, provided session arrivals form independent Poisson processes. In view of (6) and (7), it may be readily verified that the corresponding service capacities ψ_1, \dots, ψ_N are balanced by the function Ψ defined by:

$$\Psi(y) = \prod_{i \in S_0} \frac{1}{y_i!} \times \Phi(x) \prod_{k=1}^K \binom{x_k}{y_i, i \in S_k}.$$

The processor-sharing network is then a so-called Whittle network [25]. An invariant measure χ of the corresponding Markov process Y is simply given by¹:

$$\chi(y) = \Psi(y) \prod_{i=1}^N \varrho_i^{y_i}. \quad (11)$$

Summing this expression over all states corresponding to x_k flows of class k , we obtain in view of (8):

$$\varphi(x) \equiv \sum_{y: \sum_{i \in S_k} y_i = x_k} \chi(y) = \prod_{i \in S_0} e^{\varrho_i} \times \Phi(x) \prod_{k=1}^K \rho_k^{x_k} \quad (12)$$

Thus the invariant measures of the number of flows of each class are insensitive to *any* traffic characteristics (flow size distribution, distribution of the number of flows per session, correlation between successive flow sizes and think-time durations, etc) except the traffic intensities ρ_1, \dots, ρ_K . This is actually a direct consequence of the well-known insensitivity of Whittle networks [25]. We conclude that the balance property indeed implies insensitivity.

¹This measure may be of infinite sum, in which case the Markov process Y is transient or null recurrent.

3.3 Necessary condition for insensitivity

A key result of the present paper is that the converse is also true: an allocation for which the invariant measures of the number of flows of each class are insensitive to any traffic characteristics except the traffic intensities ρ_1, \dots, ρ_K is balanced. In fact, each of the following milder forms of insensitivity implies the balance property:

- (I1) **Insensitivity to the flow size distribution:** For Poisson flow arrivals and i.i.d. flow sizes, the invariant measures of the process describing the number of flows of each class remain unchanged when for any class, the exponential distribution of flow sizes is replaced by any phase-type distribution with the same mean.
- (I2) **Insensitivity to the flow arrival process:** For exponential i.i.d. flow sizes, the invariant measures of the process describing the number of flows of each class remain unchanged when for any class, the Poisson flow arrivals are replaced by Poisson *session* arrivals with the same flow arrival rate.
- (I3) **Time-scale insensitivity:** For Poisson flow arrivals and exponential i.i.d. flow sizes, the invariant measures of the process describing the number of flows of each class remain unchanged when for any class, flow inter-arrival times and flow sizes are multiplied by the same constant.

Theorem 1 *Any allocation that satisfies one of the properties (I1), (I2), (I3) is balanced.*

The proof of Theorem 1, given in Appendix A, directly follows from the necessary condition for insensitivity in processor-sharing networks proved in [5]. In view of §3.2, all three insensitivity properties above are equivalent.

3.4 Properties of insensitive allocations

In view of previous results, there exists a continuum of insensitive allocations, each characterized by a positive function Φ according to (10). In view of the capacity constraints (1), this function must satisfy the following inequalities in any state x :

$$\sum_{k:l \in r_k} \frac{\Phi(x - e_k)}{\Phi(x)} \leq C_l, \quad l = 1, \dots, L, \quad \text{and} \quad \frac{\Phi(x - e_k)}{\Phi(x)} \leq x_k a_k, \quad k = 1, \dots, K. \quad (13)$$

From (12), the invariant measures of the number of flows of each class are insensitive to any traffic characteristics except the traffic intensities ρ_1, \dots, ρ_K , and proportional to:

$$\varphi(x) = \Phi(x) \prod_{k=1}^K \rho_k^{x_k}. \quad (14)$$

This corresponds to an invariant measure of the Markov process $X = \{X_t, t \geq 0\}$ describing the evolution of the number of flows of each class for Poisson flow arrivals and exponential i.i.d. flow sizes.

Stability. The Markov process X is positive recurrent if and only if:

$$\sum_x \varphi(x) < \infty, \quad (15)$$

in which case it has the stationary distribution:

$$\pi(x) \equiv \lim_{t \rightarrow \infty} \Pr(X_t = x) = \frac{\varphi(x)}{\sum_{x'} \varphi(x')}. \quad (16)$$

Proposition 1 *The traffic conditions (3) are necessary for the Markov process X to be positive recurrent.*

Proof. Assume that $\sum_{k:l \in r_k} \rho_k > C_l$ for some link l . In view of the capacity constraints (13), we have:

$$\Phi(x) \geq \frac{1}{C_l} \sum_{k:l \in r_k} \Phi(x - e_k).$$

Let \mathcal{X} be the set of states x such that $x_k = 0$ for all k such that $l \notin r_k$. It follows from the previous inequality that for any state $x \in \mathcal{X}$ with $n = \sum_k x_k$ flows, the expression $\Phi(x)C_l^n$ is larger than the number of paths of length n from state x to state 0, i.e.,

$$\Phi(x) \geq \binom{n}{x_k, l \in r_k} \frac{1}{C_l^n}.$$

Thus the invariant measure φ given by (14) satisfies:

$$\sum_{x \in \mathcal{X}} \varphi(x) \geq \sum_{n=0}^{\infty} \left(\frac{\sum_{k:l \in r_k} \rho_k}{C_l} \right)^n = \infty.$$

The Markov process X is not positive recurrent. □

We prove in §3.5 that for the insensitive allocation referred to as “balanced fairness”, the stability condition (15) holds *if and only if* the usual traffic conditions (3) are satisfied.

Performance. Assume that stability condition (15) holds. The mean number of flows of class k in steady state is then given by:

$$E[x_k] = \sum_x x_k \pi(x).$$

The throughput of flows of class k then follows from (4). It is worth noting that flow throughput does not depend on any flow characteristics (position in the session, flow size, etc) except the class of the flow. This follows from the fact that, whatever the considered subclass of flows of a given class k (e.g., the second flows of class k of a given type of session, or those flows of class k that have a given size distribution), which can be represented simply by a subset of nodes S_k , the flow throughput of this subclass is equal to the flow throughput of class k :

Proposition 2 *For any processor-sharing network described in §2.2 satisfying (6), (7), (8), we have:*

$$\gamma_k \equiv \frac{\sum_{i \in S_k} \varrho_i}{\sum_{i \in S_k} E[y_i]} = \frac{\sum_{i \in I_k} \varrho_i}{\sum_{i \in I_k} E[y_i]}, \quad \forall I_k \subset S_k, I_k \neq \emptyset.$$

Proof. It follows from (11) and (12) that for all $i \in S_k$:

$$E[y_i] = \frac{\sum_y y_i \Psi(y) \prod_{i'} \varrho_{i'}^{y_{i'}}}{\sum_y \Psi(y) \prod_{i'} \varrho_{i'}^{y_{i'}}} = \varrho_i \frac{\sum_x x_k \Phi(x) \rho_k^{x_k-1} \prod_{k' \neq k} \rho_{k'}^{x_{k'}}}{\sum_x \Phi(x) \prod_{k'} \rho_{k'}^{x_{k'}}} = \frac{\varrho_i}{\rho_k} E[x_k].$$

It is then sufficient to sum this equality over all $i \in I_k$. □

Remark 2 *It follows from Proposition 2 that the mean duration of a flow is proportional to its size. It is sufficient to consider as many nodes as necessary to approximate a deterministic flow size distribution (cf. §2.3) and to apply Proposition 2 to this subset of nodes. Equivalently, the property follows from the fact that the mean sojourn time of a customer in any node of a Whittle network is proportional to its service time [5, 22].*

3.5 Balanced fairness: the most efficient insensitive allocation

Most insensitive allocations are inefficient in the sense that link capacities are not fully allocated. In fact, there is a unique insensitive allocation for which in any state $x \neq 0$, a network link is saturated or a flow rate limit constraint is attained. In view of (13), the corresponding balance function Φ is recursively defined by $\Phi(0) = 1$ and:

$$\forall x \neq 0, \quad \Phi(x) = \max \left(\max_l \left\{ \frac{1}{C_l} \sum_{k: l \in r_k, x_k > 0} \Phi(x - e_k) \right\}, \max_{k: x_k > 0} \left\{ \frac{1}{a_k x_k} \Phi(x - e_k) \right\} \right). \quad (17)$$

We refer to this allocation as “balanced fairness”. Note that the recursive expression (17) provides a way to evaluate numerically the corresponding balance function for any network topology. In Section 5, we give a number of examples where the balance function has a closed-form expression.

Remark 3 *Any insensitive and Pareto-efficient allocation necessarily coincides with balanced fairness. In particular, if balanced fairness is not Pareto-efficient on a given network, this implies that there is no insensitive and Pareto-efficient allocation for this network (cf. the example of hypercycles in §5.3).*

In view of (14), an invariant measure of the number of flows of each class is given by:

$$\varphi(x) = \Phi(x) \prod_{k=1}^K \rho_k^{x_k}. \quad (18)$$

A key property of balanced fairness is that the stability condition (15) holds *if and only if* the usual traffic conditions (3) are satisfied. The proof of Theorem 2 is given in Appendix B.

Theorem 2 *The stability condition (15) holds for balanced fairness if and only if the usual traffic conditions (3) are satisfied.*

4 Utility-based allocations

In view of Theorem 1, the insensitive allocations are those for which the balance property holds. In this section, we prove that utility-based allocations do not satisfy the balance property, except for proportional fairness in homogeneous “hypercubes”.

4.1 Fair allocations

As mentioned in Section 1, most allocations considered so far in the literature are based on the notion of *utility*. Assume the utility of a flow is an increasing and strictly concave function U of its rate. A unique allocation is then defined by maximizing the overall utility:

$$\sum_{k=1}^K x_k U \left(\frac{\phi_k(x)}{x_k} \right), \quad (19)$$

under the capacity constraints (1). We say that these allocations are “fair” in the sense that the utility function U is the same for all classes of flow. In particular, it follows from the Kuhn-Tucker Theorem that the capacity allocated to flows that share the same “constraining” links (in the sense that the corresponding Lagrange multipliers are positive) is fairly shared between these flows.

The allocation associated with the log utility function is known as proportional fairness [14]. Another example is the range of allocations associated with the power functions $U = (\cdot)^\alpha$, where the parameter α , $\alpha < 1$, $\alpha \neq 0$, captures the trade-off between efficiency (in terms of overall allocated capacity $\sum_{k=1}^K \phi_k(x)$) and fairness. Specifically, the allocation maximizes the overall capacity when $\alpha \rightarrow 1$ and tends to max-min fairness when $\alpha \rightarrow -\infty$ [23]. For convenience, we also refer to max-min fairness as a utility-based allocation.

Remark 4 *Utility-based allocations are Pareto-efficient. In view of Remark 3, a utility-based allocation which is insensitive coincides with balanced fairness.*

Proportional fairness has been shown to be insensitive to the flow size distribution in homogeneous “lines” and “grids”, for Poisson flow arrivals and i.i.d. sizes [4]. In view of Theorem 1 and Remark 4, it thus corresponds to balanced fairness on these network topologies. More generally, proportional fairness coincides with balanced fairness in homogeneous “hypercubes”, the multi-dimensional extension of “lines” and “grids” (see Proposition 4 below). A key result of this paper is that this is actually the only network topology for which there exists an insensitive utility-based allocation. The proof of Theorem 3 below is given in Appendix C. The considered networks are implicitly assumed to be connected in the sense that for any non-empty subset \mathcal{L}' of \mathcal{L} , $\mathcal{L}' \neq \mathcal{L}$, there exists a route containing a link in \mathcal{L}' and a link in $\mathcal{L} \setminus \mathcal{L}'$. In addition, all links are assumed to be limiting in the sense that removing any link changes the set of feasible allocations defined by the capacity constraints (1). A network is said to be homogeneous if all (limiting) links have the same capacity.

Theorem 3 *Consider a network for which a utility-based allocation is insensitive:*

- *If there is at least one limiting access rate, the network reduces to a single link and a single class.*
- *If there is no limiting access rate, the network is a homogeneous hypercube and the allocation is that realized by proportional fairness.*

Note that, as proportional fairness and max-min fairness differ in homogeneous hypercubes, max-min fairness is sensitive for *any* network topology not reduced to a single link. If the network reduces to a single link, all flows must have the same access rate. This corresponds to the case of a single bottleneck with a common access rate considered in [2, 6]. All utility-based allocations coincide with balanced fairness in this case. In the presence of different access rates, all utility-based allocations are sensitive. Bandwidth must be shared according to balanced fairness to preserve insensitivity (cf. §5.4).

4.2 Discriminatory allocations

Assume now that a class of flows is not determined by an access rate a_k and a route r_k only, but also by a fixed weight $w_k > 0$. As the function $w_k U(\cdot/w_k)$ is increasing and strictly concave, a unique allocation is defined by maximizing the overall utility:

$$\sum_{k=1}^K x_k w_k U \left(\frac{\phi_k(x)}{x_k w_k} \right), \quad (20)$$

under the capacity constraints (1). We refer to such an allocation as “discriminatory” as soon as at least two classes have different weights. The allocation associated with the log utility function U is known as weighted proportional fairness [16]. It can be shown in a very similar way to [23] that the allocation associated with the power function $U = (\cdot)^\alpha$ tends to weighted max-min fairness [21] when $\alpha \rightarrow -\infty$. More generally, it follows from the Kuhn-Tucker Theorem that the capacity allocated to flows that share the same “constraining” links (in the sense that the corresponding Lagrange multipliers are positive) is shared in proportion to their weights. For a single link in the absence of access rate, the corresponding model is the so-called discriminatory processor-sharing queue [12].

Proposition 3 *Discriminatory allocations are sensitive.*

Proof. Consider a link l that belongs to at least two routes, say routes r_1, r_2 , such that the corresponding classes have different weights, $w_1 \neq w_2$. Denote by \mathcal{K} the set of classes k such that route r_k contains link l . As link l is limiting, there exists a state x , with $x_k = 0$ for all $k \notin \mathcal{K}$ and $x_k > C_l/a_k$ for all $k \in \mathcal{K}$, such that:

$$\forall k \in \mathcal{K}, \quad \frac{x_k w_k}{\sum_{k' \in \mathcal{K}} x_{k'} w_{k'}} C_l < \min_{l' \in r_k, l' \neq l} C_{l'}.$$

For any utility function U , it follows from the Kuhn-Tucker Theorem that link l is the only saturated link in state x and:

$$\forall k \in \mathcal{K}, \quad \phi_k(x) = \frac{x_k w_k}{\sum_{k' \in \mathcal{K}} x_{k'} w_{k'}} C_l.$$

By choosing x_1 and x_2 sufficiently large, we can assume that this expression still holds in states $x - e_1$ and $x - e_2$. In particular,

$$\frac{\phi_1(x)}{\phi_1(x - e_2)} = 1 - \frac{w_2}{\sum_{k \in \mathcal{K}} x_k w_k} \neq 1 - \frac{w_1}{\sum_{k \in \mathcal{K}} x_k w_k} = \frac{\phi_2(x)}{\phi_2(x - e_1)}$$

The balance property is violated. The proof then follows from Theorem 1. \square

Discriminatory allocations are advocated as a means for realizing service differentiation [7]. Users pay more for a higher weight and expect a better quality of service. Proposition 3 suggests, however, that the degree of service differentiation is difficult to control since sensitive to detailed traffic characteristics. It turns out that these allocations are approximately insensitive provided the ratio $\max_k w_k / \min_k w_k$ is not too high [4]. But the service differentiation is then not significant. In case of priorities, i.e., when the ratio $\max_k w_k / \min_k w_k$ tends to infinity, the sensitivity is strong and the service differentiation difficult to control [4]. This calls into question the use of discriminatory allocations as an efficient means for realizing service differentiation in data networks.

5 Application to specific network topologies

In this section, we give the allocation realized by balanced fairness on a number of network topologies and compare the resulting performance with that of max-min fairness. Specifically, the balance function (17) associated with balanced fairness is evaluated explicitly when it has a closed-form expression. Note that the stationary distribution of the number of flows of each class follows immediately in view of (14)-(16) and Theorem 2. The flow throughput is then given by (4). As max-min fairness is sensitive, the corresponding performance results are obtained by simulation with Poisson flow arrivals and i.i.d. exponential flow sizes of unit mean, unless otherwise specified. Note that the corresponding Markov process describing the number of flows of each class is positive recurrent under the usual traffic conditions [9]. Unless otherwise specified, we assume that there is no limiting access rate.

5.1 Lines, grids and hypercubes

Lines and grids shown in Figure 4 have already been considered in [4]. These are actually hypercubes of dimension $n = 2$, according to the following definition.

Definition 2 (Hypercube) *A hypercube of dimension n is a network of n sets of route, referred to as directions, such that the set of links is the set of intersections of n routes of different directions.*

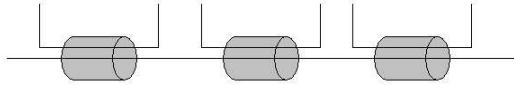


Figure 4: Examples of a line (left) and a grid (right)

Note that a hypercube of dimension $n = 1$ just consists of a single link. Consider a hypercube of dimension $n \geq 2$ with unit capacity links. Denoting by $\mathcal{D}_1, \dots, \mathcal{D}_n$ the corresponding directions, we get:

$$\Phi(x) = \left(\sum_{r_k \in \mathcal{D}_1} x_k, \dots, \sum_{r_k \in \mathcal{D}_n} x_k \right).$$

It may indeed be easily be verified that this function satisfies (17) in any state x .

Proposition 4 *Proportional fairness coincides with balanced fairness in homogeneous hypercubes.*

Proof. Let $r_1 \in \mathcal{D}_1, \dots, r_n \in \mathcal{D}_n$. It can be easily verified that for any Pareto-efficient allocation, all links are saturated and all routes that belong to the same direction receive the same bandwidth, so that $\sum_{i=1}^n \phi_i(x) = 1$ for unit capacity links. The bandwidth allocation realized by proportional fairness maximizes the function:

$$\sum_{i=1}^n \sum_{k \in \mathcal{D}_i} x_k \log \left(\frac{\phi_i(x)}{x_k} \right).$$

We obtain for any directions i, j :

$$\sum_{k \in \mathcal{D}_i} \frac{x_k}{\phi_i(x)} = \sum_{k \in \mathcal{D}_j} \frac{x_k}{\phi_j(x)},$$

so that

$$\phi_i(x) = \frac{\sum_{k \in \mathcal{D}_i} x_k}{\sum_k x_k}.$$

The result then follows from (10). □

Figure 5 compares the performance results obtained with balanced fairness and max-min fairness for a line of $L = 5$ unit capacity links. The traffic intensity is the same on each route. The “short” route refers to any single link route, while the “long” route refers to the 5-link route. Note that balanced fairness leads to better flow throughput than max-min fairness on the short route and worse flow throughput on the long route.

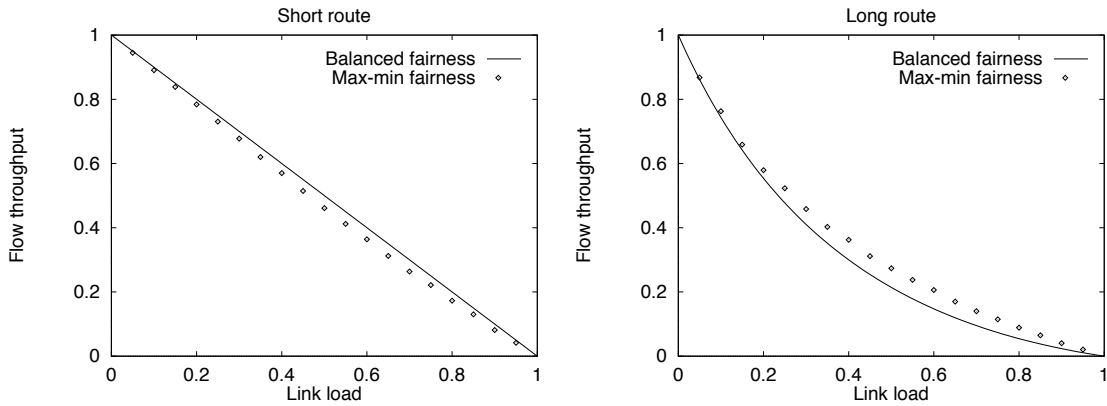


Figure 5: Comparison of balanced fairness and max-min fairness on a 5-link line

Deriving the explicit bandwidth sharing realized by balanced fairness in non-homogeneous hypercubes is difficult in general. An exception is the line with unequal link capacities. For simplicity, we assume without loss of generality that the minimum link capacity is equal to one. Denoting by x_0 the number of flows on the route that contains all links and by x_l the number of flows on the route that contains link l only, we have for any x such that $x_l > 0$ for some link l :

$$\Phi(x) = \sum_{y: \sum_{l: x_l > 0} y_l \leq x_0} \prod_{l: x_l > 0} \binom{x_l + y_l - 1}{y_l} \frac{1}{C_l^{x_l + y_l}}.$$

5.2 Trees

The network topology we refer to as a “tree” and illustrate in Figure 6 is practically interesting as it may represent an access network that typically consists of several multiplexing stages.

Definition 3 (Trees) *A tree is a network of K routes and $L = K + 1$ links such that a single link, referred to as the trunk, belongs to all routes, and all other links, referred to as the branches, belong to a single route.*

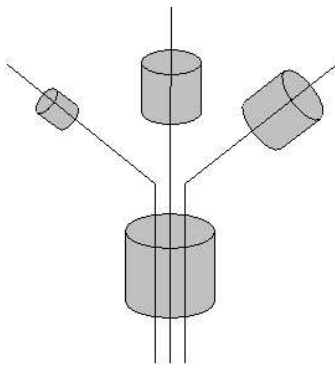


Figure 6: A 3-branch tree

Let $r_k = \{0, k\}$, where link 0 refer to the trunk and link k to a branch, $k = 1, \dots, K$. We assume without loss of generality that $C_0 = 1$, $C_k \leq 1$ for all branches k , $C_k = 1$ for at most one branch k , and $\sum_{k=1}^K C_k > 1$. We have:

$$\Phi(x) = \prod_{k=1}^K \frac{1}{C_k^{x_k}}$$

for all states x for which the sum of the capacity of “active” branches (those branches k such that $x_k > 0$) is less than 1, and:

$$\Phi(x) = \sum_{z: z \leq x, I(z) \neq \emptyset} \left(\frac{\sum_k (x_k - z_k) - 1}{\sum_{k \notin I(z)} (x_k - z_k)} \right) \left(\frac{\sum_{k \notin I(z)} (x_k - z_k)}{(x_k - z_k), k \notin I(z)} \right) \left(\frac{\sum_{k \in I(z)} (x_k - z_k)}{(x_k - z_k), k \in I(z)} \right) \prod_{k=1}^K \frac{1}{C_k^{z_k}}$$

otherwise, where z is a K -dimensional vector of integers, $z \leq x$ means $z_k \leq x_k$ for all k , $I(z)$ is the empty set if $C.1(z) \equiv \sum_{k: z_k > 0} C_k > 1$ and the set of indices k such that $z_k = 0$, $x_k > 0$ and $C.1(z) + C_k > 1$ otherwise. For a 2-branch tree, this last expression reduces to:

$$\Phi(x) = \sum_{z_1 \leq x_1} \binom{x_1 - z_1 + x_2 - 1}{x_1 - z_1} \frac{1}{C_1^{z_1}} + \sum_{z_2 \leq x_2} \binom{x_1 + x_2 - z_2 - 1}{x_2 - z_2} \frac{1}{C_2^{z_2}}.$$

In Figure 7, we compare the performance of balanced fairness and max-min fairness on a 2-branch tree. Note that all utility-based allocations coincide with max-min fairness on trees (see Lemma 3 in Appendix C). For max-min fairness, two sets of simulation results are given, depending on the ratio r of the mean flow size of class 1 to the mean flow size of class 2. We know from Theorem 1 and Theorem 3 that the performance of max-min fairness is sensitive to this parameter (property (I3) does not hold). In a homogeneous case ($C_1 = C_2 = 0.6$, $\rho_1 = \rho_2$, left plot), both allocations lead to very close performance and max-min fairness is approximately insensitive. In a non-homogeneous case ($C_1 = 1, C_2 = 0.5$, $\rho_1 = 0.01$, right plot), the performance of balanced fairness and max-min fairness may differ significantly due to the strong sensitivity of max-min fairness.

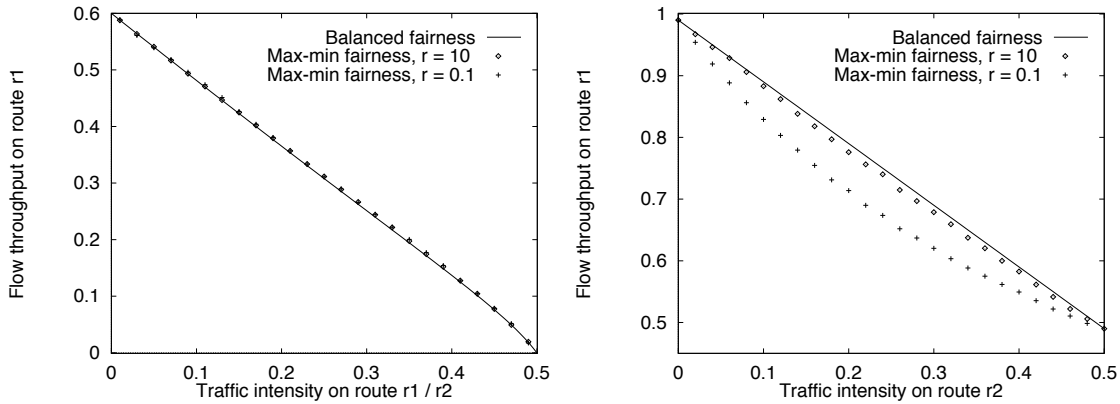


Figure 7: Comparison of balanced fairness and max-min fairness on a 2-branch tree (left plot: homogeneous case, right plot: non-homogeneous case)

5.3 Triangles, pyramids and hypercycles

As illustrated Figure 8, a triangle is composed of 3 links and 3 routes, each route containing 2 links. Similarly, a pyramid, for which we give a 2-D representation in Figure 8, is composed of 4 links and 4 routes, each route containing 3 links. More generally, we define hypercycles as follows:

Definition 4 (Hypercycles) *A hypercycle is a network of L links and L routes such that the set of links is the set of intersections of $L - 1$ routes.*

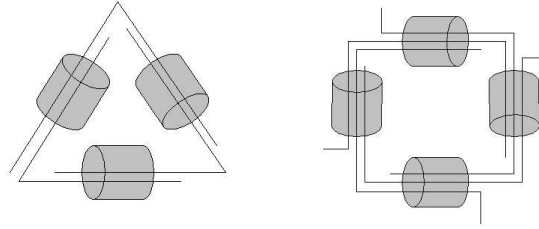


Figure 8: The triangle and the pyramid

Even for the simplest hypercycles, it proves extremely difficult to evaluate explicitly the balance function associated with balanced fairness. This is notably due to the fact that balanced fairness is not Pareto-efficient on these networks (see Lemma 8 in Appendix C). The inefficiency of balanced fairness in homogeneous hypercycles is illustrated in Figure 9, which gives the fraction of wasted bandwidth for the triangle when $x_3 = 10$ and for the pyramid when $x_3 = x_4 = 10$. The wasted bandwidth is here defined as the maximum bandwidth that can be added to a route under the capacity constraints (1). By definition, the wasted bandwidth is equal to zero for a Pareto-efficient allocation. It can be verified that for any integer $n \geq 1$, $\phi_1(x) = (n + 1)/(3n + 1)$ and $\phi_2(x) = \phi_3(x) = 1/2$ in state $x = (1, n, n)$ for the triangle with unit capacity links, so that the wasted bandwidth can be as high as $1/6$ for this network topology, while $\phi_1(x) = (2n^2 + 3n + 1)/(11n^2 + 6n + 1)$ and $\phi_2(x) = \phi_3(x) = \phi_4(x) = 1/3$ in state $x = (1, n, n, n)$ for the pyramid with unit capacity links, so that the wasted bandwidth can be as high as $5/33$ for this network topology.

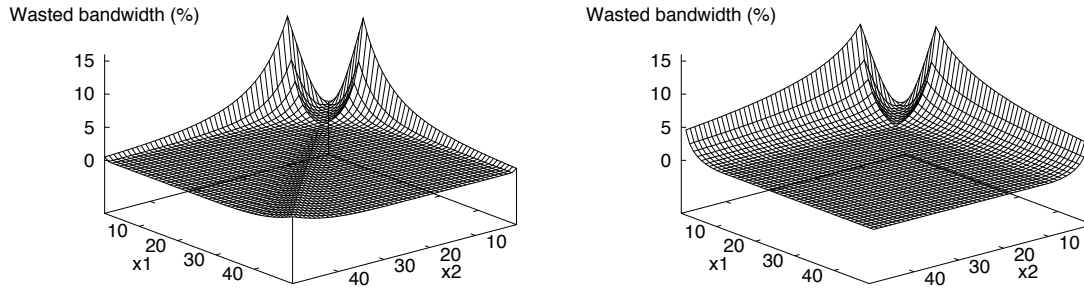


Figure 9: Inefficiency of balanced fairness on the triangle (left) and on the pyramid (right)

Figure 10 presents a performance comparison of balanced fairness and max-min fairness for the triangle with unit capacity links, when the three routes have the same traffic intensity and when the traffic intensity of one route is fixed at 0.05 (the plot corresponds to the flow throughput on this route, say r_1). In both cases, we observe that balanced fairness and max-min fairness give very similar results. This suggests that the inefficiency of balanced fairness does not have a strong impact on performance. We also observe that while the performance of balanced fairness is monotonic with respect to the traffic intensity, this is not the case of max-min fairness (cf. right plot). This can be explained by the fact that when the load of link $r_2 \cap r_3$ tends to 1, the traffic generated on each route r_2, r_3 is “smoothed” and looks like a fluid of constant rate 0.5. As max-min fairness is Pareto-efficient, the flow throughput on route r_1 tends to that of a single link of capacity $C = 0.5$ and traffic intensity $\rho = 0.05$, that is $C - \rho = 0.45$ in view of (5). This is not the case of balanced fairness which is inefficient on this network topology.

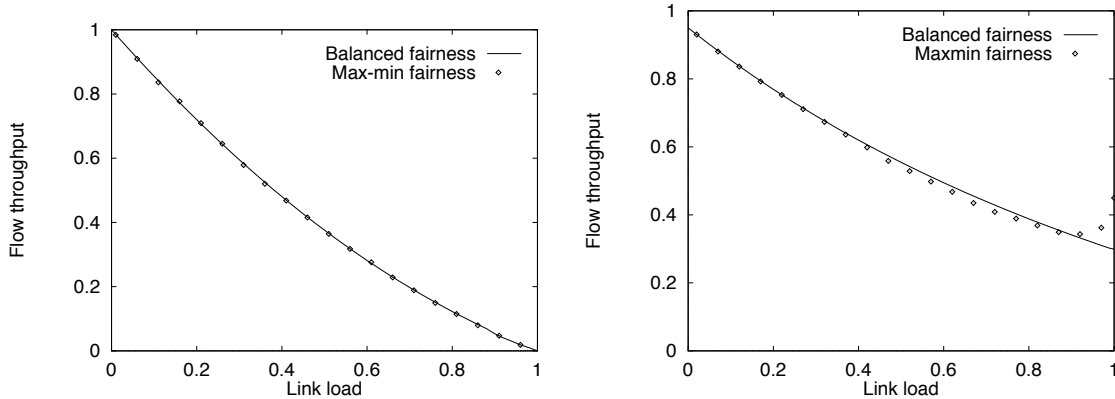


Figure 10: Comparison of balanced fairness and max-min fairness on the triangle (left plot: same traffic intensity on each route, right plot: traffic intensity of one route fixed at 0.05)

5.4 A single link with different access rates

Finally, we consider a single unit capacity link with different access rates $a_1, \dots, a_K < 1$. We have:

$$\Phi(x) = \prod_{k=1}^K \frac{1}{a_k^{x_k} x_k!}, \text{ if } a \cdot x \leq 1,$$

and

$$\Phi(x) = \sum_{z: z \leq x, I(z) \neq \emptyset} \left(\frac{\sum_k (x_k - z_k) - 1}{\sum_{i \notin I(z)} (x_k - z_k)} \right) \binom{\sum_{k \notin I(z)} (x_k - z_k)}{(x_k - z_k), k \notin I(z)} \binom{\sum_{k \in I(z)} (x_k - z_k)}{(x_k - z_k), k \in I(z)} \prod_{k=1}^K \frac{1}{a_k^{z_k} z_k!},$$

otherwise, where z is a K -dimensional vector of integers, $z \leq x$ means $z_k \leq x_k$ for all k , $I(z)$ is the empty set if $a \cdot z \equiv \sum_k z_k a_k > 1$, and the set of indices k such that $x_k > 0$ and $a \cdot z + a_k > 1$ otherwise. Figure 11 compares the performance of balanced fairness and max-min fairness for two access rates $a_1 = 0.2$ and $a_2 = 0.5$. The traffic intensity is the same for each class of flows.

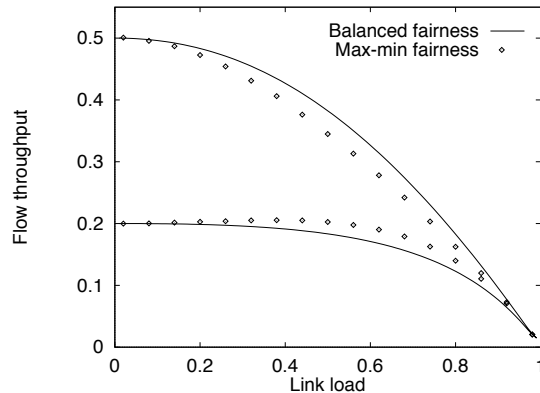


Figure 11: Comparison of balanced fairness and max-min fairness for two distinct access rates

6 Conclusion

Most allocations considered so far in the literature are based on the notion of utility, which is evaluated on the basis of the rate allocated to each flow in a static scenario. We suggest that it is more appropriate when defining bandwidth sharing objectives in data networks to study flow-level dynamics. In particular, we argue that a useful and highly desirable property is that capacity allocation should lead to performance that is *insensitive* to detailed traffic characteristics. Like Erlang’s loss formula for telephone networks, sharing network resources in an insensitive way would allow the development of simple future proof engineering rules for data networks.

We have characterized the class of insensitive allocations and proved that utility-based allocations do not belong to this class in general. We have identified the most efficient insensitive allocation, referred to as balanced fairness, for which explicit performance results were given for a number of practically interesting network topologies. It remains to find how such an allocation could be realized in a distributed way through the congestion control algorithms implemented in the sources or the scheduling and buffer management schemes implemented in network nodes.

Appendix

A Balance property of insensitive allocations

Proof of Theorem 1. Consider the processor-sharing network introduced in §2.3 representing the data network with Poisson flow arrivals and exponential i.i.d. flow sizes, i.e., with $N = K$ nodes and $\nu_i/\mu_i = \rho_i$ for $i = 1, \dots, N$. We refer to this processor-sharing network as the *initial* network. From [5, Theorem 2], the following insensitivity property (P) implies the balance property:

- (P) The invariant measures of the Markov process describing the number of customers at each node of the initial network remain unchanged when for any node i and for any α_i , $0 < \alpha_i < 1$, the exponential i.i.d. services at node i are replaced by i.i.d. services, exponentially distributed of mean $1/\alpha_i \times 1/\mu_i$ with probability α_i , null with probability $1 - \alpha_i$.

The proof then follows from the fact that each property (I1), (I2), (I3) implies property (P):

(I1) \Rightarrow (P) Consider the initial network where the i.i.d. exponential services are replaced by i.i.d. 2-phase services with the same mean. This new network consists of $\tilde{N} = 2N$ nodes with $\tilde{S}_i = \{i, i + N\}$ for each class i , $i = 1, \dots, N$, exogenous arrival rates $\tilde{\nu}_i = \nu_i$ and $\tilde{\nu}_{i+N} = 0$, routing probability $\tilde{p}_{i,i+N} = \alpha_i$ for some α_i , $0 < \alpha_i < 1$, $\tilde{p}_{ij} = 0$ otherwise, and service rates $\tilde{\mu}_i = \omega \times \mu_i$ and $\tilde{\mu}_{i+N} = \omega/(\omega - 1) \times \alpha_i \mu_i$, for some $\omega > 1$. From property (I1), the invariant measures of the process describing the number of customers at each node i of the initial network are invariant measures of the process describing the number of customers at each pair of nodes $\{i, i + N\}$ of this new network. We deduce (P) by letting ω tend to infinity in the corresponding balance equations.

(I2) \Rightarrow (I3) Consider the initial network where the Poisson flow arrivals are replaced by Poisson session arrivals, each session being composed of a geometrically distributed number of flows. This new network consists of $\tilde{N} = 2N$ nodes with $\tilde{S}_i = \{i\}$ for each class i , $i = 1, \dots, N$, $\tilde{S}_0 = \{N+1, \dots, 2N\}$, exogenous arrival rates $\tilde{\nu}_i = \alpha_i \nu_i$ for some α_i , $0 < \alpha_i < 1$, and $\tilde{\nu}_{i+N} = 0$, routing probability $\tilde{p}_{i,i+N} = 1 - \alpha_i$, $\tilde{p}_{ij} = 0$ otherwise, and service rates $\tilde{\mu}_i = \mu_i$ and $\tilde{\mu}_{i+N} = \omega$, for some $\omega > 0$. From property (I2), the invariant measures of the process describing the number of customers at each node of the initial network are invariant measures of the process describing the number of customers at each node $1, \dots, N$ of this new network. We deduce (I3) by letting ω tend to infinity in the corresponding balance equations.

(I3) \Rightarrow (P) Consider the initial network where the arrival rates and service rates at any node i are multiplied by the same constant α_i , $0 < \alpha_i < 1$. This also corresponds to the initial network with the same arrival rates but where the services at node i are replaced by exponentially distributed services of mean $1/\alpha_i \times 1/\mu_i$ with probability α_i , null services with probability $1 - \alpha_i$.

□

B Stability condition for balanced fairness

To prove Theorem 2, we need the following result.

Lemma 1 Consider any other positive function $\tilde{\Phi}$ such that $\tilde{\Phi}(0) = 1$ and the inequalities (13) are satisfied. We have:

$$\forall x, \quad \tilde{\Phi}(x) \geq \Phi(x).$$

Proof. The proof is by induction on the total number of flows $n = \sum_{k=1}^K x_k$. As $\tilde{\Phi}(0) = \Phi(0) = 1$, the inequality is satisfied for $n = 0$. Now assume it is satisfied for $n = m$, $m \geq 0$. Let x be any state with $n = m + 1$ total number of flows. From (13) and (17), we get:

$$\begin{aligned} \tilde{\Phi}(x) &\geq \max \left(\max_l \left\{ \frac{1}{C_l} \sum_{k:l \in r_k} \tilde{\Phi}(x - e_k) \right\}, \max_{k:x_k > 0} \left\{ \frac{1}{a_k x_k} \tilde{\Phi}(x - e_k) \right\} \right) \\ &\geq \max \left(\max_l \left\{ \frac{1}{C_l} \sum_{k:l \in r_k} \Phi(x - e_k) \right\}, \max_{k:x_k > 0} \left\{ \frac{1}{a_k x_k} \Phi(x - e_k) \right\} \right) = \Phi(x). \end{aligned}$$

□

Proof of Theorem 2. Consider the insensitive allocation characterized by the balance function:

$$\forall x, \quad \tilde{\Phi}(x) = \sum_{z \in \mathcal{Z}} \left\{ \prod_{k=1}^K \frac{1}{z_k!} \left(\frac{1}{a_k} \right)^{z_k} \times \prod_{l=1}^L \binom{\sum_{k:l \in r_k} z_{kl}}{z_{kl}, k : l \in r_k} \left(\frac{1}{C_l} \right)^{\sum_{k:l \in r_k} z_{kl}} \right\},$$

where $z = (z_k, z_{kl}, l \in r_k)_k$ is a vector of integers and

$$\mathcal{Z} = \left\{ z : \forall k, z_k + \sum_{l \in r_k} z_{kl} = x_k \right\}.$$

It may readily be verified that this function satisfies the capacity constraints (13). In fact, this follows from the fact that $\tilde{\Phi}(x)$ is the normalization constant of the following closed Kelly queueing network. The network consists of L processor-sharing nodes $1, \dots, L$ of respective capacities C_1, \dots, C_L and K infinite-server nodes $1, \dots, K$ of respective per-server capacities a_1, \dots, a_K . There are K classes of customer. Customers of class k visit the infinite-server node k and the processor-sharing nodes $l \in r_k$ in a cyclic way, in a fixed but arbitrary order (each of these nodes is visited exactly once in a cycle). Services at each node are exponential i.i.d. of unit mean. The rate at which customers of class k visit each processor-sharing node $l \in r_k$ and the infinite-server node k is equal to $\tilde{\Phi}(x - e_k)/\tilde{\Phi}(x)$, so that the constraints (13) are satisfied. In particular, it follows from Lemma 1 that $\tilde{\Phi}(x) \geq \Phi(x)$ for all x . From (14), an invariant measure for the Markov process \tilde{X} associated with this allocation for Poisson flow arrivals and i.i.d. exponential services is given by:

$$\forall x, \quad \tilde{\varphi}(x) = \tilde{\Phi}(x) \prod_{k=1}^K \rho_k^{x_k}.$$

This corresponds to an invariant measure for the number of customers of each class in the following open Kelly queueing network. The network is the same as the closed queueing network considered above except that customers of class k arrive as a Poisson process of rate ρ_k , visit the infinite-server node k and the processor-sharing nodes $l \in r_k$, in a fixed but arbitrary order, then leave the network. In particular, we know that \tilde{X} is positive recurrent if and only if the usual traffic conditions (3) hold [14]. As $\tilde{\varphi}(x) \geq \varphi(x)$ for all x , where φ is the invariant measure of the Markov process X given by (14), we conclude the proof by summing this inequality over all states x . \square

C Sensitivity of utility-based allocations

In this section, we prove Theorem 3. The definition of trees, hypercubes and hypercycles are given in Section 5. We need the following lemmas.

Lemma 2 *Any non-homogeneous network contains a tree.*

Proof. Consider a non-homogeneous network. Denote by $\mathcal{L}' \subset \mathcal{L}$ the set of links l such that at least one of the routes containing l also contains a link of capacity smaller than C_l . Note that \mathcal{L}' is non-empty. We denote by l' the link of smallest capacity that belongs to \mathcal{L}' , and by \mathcal{K} the set of classes such that all routes r_k , $k \in \mathcal{K}$, contain link l' . If two of these routes, say r_1 and r_2 , cross each other at a link of smaller capacity than $C_{l'}$, consider the link l'' of smallest capacity that belongs to $r_1 \cap r_2$. Since $l'' \notin \mathcal{L}'$, routes r_1 and r_2 coincide in the network restricted to the set of classes \mathcal{K} . We conclude that the network restricted to the set of classes \mathcal{K} is a tree. \square

Lemma 3 *Utility-based allocations are not balanced in trees.*

Proof. Consider a tree composed of a trunk of unit capacity and K branches $1, \dots, K$ of respective capacities $C_1 \geq \dots \geq C_K$, such that:

$$\sum_{k \neq K} C_k \leq 1 < \sum_k C_k. \quad (21)$$

As any tree contains a restricted tree that satisfies (21), it is sufficient to prove that utility-based allocations are not balanced under this assumption. The class of flows k is characterized by the route r_k which contains the trunk and branch k , and a per-flow rate limit a_k . For any state x such that $x_k > 1/a_k$ for all k , the access rates are not limiting. It then follows from the Kuhn-Tucker theorem that the solution of the optimisation problem (19) satisfies:

$$U' \left(\frac{\phi_k(x)}{x_k} \right) = \eta_0 + \eta_k, \quad x_k > 0.$$

where η_0 and η_k are the Lagrange multipliers associated with the capacity constraints of the trunk and branch k , respectively. Let $\mathcal{K}(x) \subset \{1, \dots, K\}$ be the set of non-saturated branches k , i.e., such that $\phi_k(x) < C_k$. We have $\eta_k = 0$ in this case, so that the capacity $\phi_k(x)/x_k$ allocated to each flow of class k is the same for all $k \in \mathcal{K}(x)$. We get:

$$\phi_k(x) = \frac{x_k}{\sum_{k' \in \mathcal{K}(x)} x_{k'}} (1 - \sum_{k' \notin \mathcal{K}(x)} C_{k'}), \quad k \in \mathcal{K}(x).$$

In particular, any utility-based allocations coincides with max-min fairness. We now consider two cases.

- Assume $C_1 > C_K$. Noting that $\sum_{k \neq 1} C_k < 1$, $\sum_{k \neq 2} C_k \leq 1$, and

$$\frac{1 - \sum_{k \neq 2} C_k}{C_1} < \frac{C_2}{C_1} < \frac{C_2}{1 - \sum_{k \neq 1} C_k},$$

we can choose a state x such that $x_k > 1/a_k + 1$ for all k ,

$$x_1 - 1 < \frac{x_2}{C_2} (1 - \sum_{k \neq 1} C_k) \leq x_1, \quad \frac{x_1}{C_1} (1 - \sum_{k \neq 2} C_k) < x_2 - 1, \quad (22)$$

and

$$x_k > \frac{C_k}{1 - \sum_{k' \neq 1,2} C_{k'}} (x_1 + x_2), \quad k \neq 1, 2. \quad (23)$$

The access rates are not limiting in state x , nor in states $x - e_1$, $x - e_2$. We shall prove that $\phi_1(x)\phi_2(x - e_1) \neq \phi_2(x)\phi_1(x - e_2)$. Let $\bar{C} \equiv 1 - \sum_{k \neq 1,2} C_k$. It easily follows from (22) that:

$$\frac{x_1}{x_1 + x_2} \bar{C} < \frac{x_1}{x_1 + x_2 - 1} \bar{C} < C_1 \quad \text{and} \quad \frac{x_2 - 1}{x_1 + x_2 - 1} \bar{C} < \frac{x_2}{x_1 + x_2} \bar{C} \leq C_2.$$

Using (23), we deduce that $\mathcal{K}(x) = \{1\}$ or $\mathcal{K}(x) = \{1, 2\}$. In addition, as branches $k \neq 1, 2$ are not saturated in state x , they are not saturated in states $x - e_1$, $x - e_2$. It then follows from previous inequalities that $\mathcal{K}(x - e_2) = \{1, 2\}$ and from (22) that $\mathcal{K}(x - e_1) = \{1\}$. Thus, we have $\phi_1(x - e_2) = \bar{C}x_1/(x_1 + x_2 - 1)$ and $\phi_2(x - e_1) = C_2$. If $\mathcal{K}(x) = \{1\}$, we have $\phi_1(x) = 1 - \sum_{k \neq 1} C_k$ and $\phi_2(x) = C_2$. If $\mathcal{K}(x) = \{1, 2\}$, we have $\phi_1(x)/\phi_2(x) = x_1/x_2$. In both cases, the balance property is violated in view of (22).

- Assume $C_1 = C_K$. Consider the state $x = ne + me_1$ where $e = (1, \dots, 1)$, n is an integer such that $n > 1/a_k + 1$ for all k and $m \geq 1$ is the smallest integer such that branch 1 is saturated, i.e., $\phi_1(x) = C_1$ and $\phi_1(x - e_1) < C_1$. Note that the access rates are not limiting in state x , nor in states $x - e_1, x - e_2$. As branch 1 is saturated in state x , it is saturated in state $x - e_2$: $\phi_1(x) = \phi_1(x - e_2) = C_1$. As $\phi_2(x) = \dots = \phi_K(x)$, branches $2, \dots, K$ are not saturated in state x and $\phi_2(x) = (1 - C_1)/(K - 1)$. Similarly, as $\phi_2(x - e_1) = \dots = \phi_K(x - e_1)$ and $\phi_2(x - e_1) \leq \phi_1(x - e_1)$, no branch is not saturated in state $x - e_1$:

$$\phi_1(x - e_1) = \frac{n + m - 1}{Kn + m - 1} < C_1 \quad \text{and} \quad \phi_2(x - e_1) = \frac{n}{Kn + m - 1}.$$

We conclude that $\phi_1(x)\phi_2(x - e_1) > \phi_1(x - e_2)\phi_2(x)$. The capacities are not balanced. □

Lemma 4 *Utility-based allocations are not balanced in homogeneous networks with at least two different access rates.*

Proof. Any network with unit capacity links and at least two different access rates contains a link shared by two classes of flow with different access rates. Consider the restriction of the network to these two classes, say classes 1,2, with $a_1 < a_2$. Let $x = (x_1, 1)$ where x_1 is the smallest integer such that $x_1 a_1 + a_2 > 1$ if $a_2 < 1$, $x_1 = 1$ otherwise. Any utility-based allocation gives $\phi_1(x) = \phi_1(x - e_2) = x_1 a_1$ and $\phi_2(x) = 1 - x_1 a_1$, $\phi_2(x - e_1) = a_2$ if $a_2 < 1$, $\phi_2(x - e_1) = 1$ otherwise. The balance property would imply $x_1 a_1 + a_2 = 1$ if $a_2 < 1$, $a_1 = 1$ otherwise, a contradiction. □

Lemma 5 *Utility-based allocations are not balanced in homogeneous 2-link lines with at least one limiting access rate.*

Proof. Consider a homogeneous line of two unit capacity links. Route r_1 contains both links while routes r_2, r_3 contain a single link. We know from Lemma 4 that utility-based allocations are not balanced in the presence of at least two different access rates. Assume that a utility-based allocation satisfies the balance property in the presence of a common limiting access rate $a < 1$, i.e., there are $K = 3$ classes of flow and $a_1 = a_2 = a_3 = a$. Let $x = (x_1, x_2, 1)$ where $x_1 a \geq 1$ and x_2 is the smallest integer such that $(1 - a)x_2 \geq ax_1$. Link $r_1 \cap r_3$ is not saturated in state x , so that:

$$\phi_1(x) = \frac{x_1}{x_1 + x_2}, \quad \phi_2(x) = \frac{x_2}{x_1 + x_2}, \quad \phi_3(x) = a.$$

As $\phi_2(x - e_3) = \phi_2(x)$, the balance property implies that $\phi_3(x - e_2) = a$. Noting that link $r_1 \cap r_3$ is saturated in state $x - e_2$, we get $\phi_1(x - e_2) = 1 - a$. It then follows from the balance property that:

$$\phi_2(x - e_1) = \phi_1(x - e_2) \frac{\phi_2(x)}{\phi_1(x)} = (1 - a) \frac{x_2}{x_1}.$$

But, as link $r_1 \cap r_3$ is not saturated in state x , it is not saturated in state $x - e_1$ and:

$$\phi_2(x - e_1) = \frac{x_2}{x_1 + x_2 - 1}.$$

This implies $(1 - a)(x_2 - 1) = ax_1$, a contradiction. □

Definition 5 (Incomplete line) *An incomplete line is a network of five routes r_1, r_2, r_3, r_4, r_5 and three links, $r_1 \cap r_2 \cap r_3$, $r_1 \cap r_2 \cap r_4$ and $r_1 \cap r_5$.*

Lemma 6 *Utility-based allocations are not balanced in homogeneous incomplete lines.*

Proof. Consider a homogeneous incomplete line, with three links $r_1 \cap r_2 \cap r_3$, $r_1 \cap r_2 \cap r_4$ and $r_1 \cap r_5$. As this network contains a homogeneous 2-link line, it follows from Lemmas 4 and 5 that utility-based allocations are not balanced in the presence of at least one limiting access rate. We thus consider the case where there is no limiting access rate. If there exists an insensitive utility-based allocation, it must coincide with balanced fairness in view of Remark 4. For $x = (1, 1, 1, 1, 1)$, balanced fairness gives $\phi_1(x) = \phi_5(x) = 1/2$ in view of (17). This is clearly not the allocation realized by max-min fairness. In addition, it follows from the Kuhn-Tucker theorem that the solution of the optimisation problem (19) gives for $x = (1, 1, 1, 1, 1)$,

$$U'(\phi_1(x)) = U'(\phi_3(x)) + U'(\phi_4(x)) + U'(\phi_5(x)).$$

If $\phi_1(x) = \phi_5(x)$, we get $U'(\phi_3(x)) + U'(\phi_4(x)) = 0$, that is $\phi_3(x) = \phi_4(x) = 0$ and $\phi_1(x) = \phi_2(x) = \phi_5(x) = 1/2$. But the allocation $\phi_2(x) = 0$ and $\phi_1(x) = \phi_3(x) = \phi_4(x) = \phi_5(x) = 1/2$ leads to a strictly higher overall utility. \square

Definition 6 (Incomplete square) *An incomplete square is a network of four routes r_1, r_2, r_3, r_4 and three links, $r_1 \cap r_2$, $r_2 \cap r_3$ and $r_3 \cap r_4$.*

Lemma 7 *Utility-based allocations are not balanced in homogeneous incomplete squares.*

Proof. Consider a homogeneous incomplete square, with three links $r_1 \cap r_2$, $r_2 \cap r_3$ and $r_3 \cap r_4$. As this network contains a homogeneous 2-link line, it follows from Lemmas 4 and 5 that utility-based allocations are not balanced in the presence of at least one limiting access rate. We thus consider the case where there is no limiting access rate. If there exists an insensitive utility-based allocation, it must coincide with balanced fairness in view of Remark 4. For $x = (1, 1, 1, 1)$, balanced fairness gives $\phi_1(x) = \phi_2(x) = 2/5$ and $\phi_3(x) = \phi_4(x) = 3/5$ in view of (17), whereas any utility-based allocation gives $\phi_1(x) = \phi_2(x) = \phi_3(x) = \phi_4(x) = 1/2$. \square

Lemma 8 *Utility-based allocations are not balanced in homogeneous hypercycles.*

Proof. Consider a hypercycle of L unit capacity links, $L \geq 3$. In view of Lemma 4, it is sufficient to prove that utility-based allocations are not balanced in the presence of a common access rate a , i.e., there are $K = L$ classes of flow with $a_1 = \dots = a_K = a$. We consider two cases.

- Assume that $a(K - 1) \geq 1$. Note that this includes the case $a = \infty$ where the access rate is not limiting. We prove that the allocation realized by balanced fairness is not Pareto-efficient. This implies that utility-based allocations are not balanced in view of Remark 4. Let $x = (1, 2, \dots, 2)$ and $c = 1/(K - 1)$. As $\phi_2(x) = \dots = \phi_K(x)$, it follows from the capacity constraints (1) that $\phi_K(x) \leq c$. The Pareto-efficiency would imply $\phi_1(x) \geq c$ as $a \geq c$. But, as the network reduces to a single link in state $x - e_1$, we know that $\phi_K(x - e_1) = c$. Thus $\phi_K(x - e_1) \geq \phi_K(x)$ and it follows from the balance property that $\phi_1(x - e_K) \geq \phi_1(x) \geq c$. By symmetry, we have $\phi_1(x - e_2) \geq c$. Now as $\phi_1(x - e_2) = \phi_2(x - e_2)$ and $\phi_3(x - e_2) = \dots = \phi_K(x - e_2)$, it follows from the capacity constraints (1) that $\phi_K(x - e_2) \leq c$. Applying successively the same reasoning, we get $\phi_K(y) \leq c$ in state

$y = (1, \dots, 1, 2)$. Again, the Pareto-efficiency would imply $\phi_1(y) \geq c$. But, as the network reduces to a single link in state $y - e_1$, we know that $\phi_K(y - e_1) = 2/K > c$. Thus $\phi_K(y - e_1) > \phi_K(y)$ and it follows from the balance property that $\phi_1(y - e_K) > \phi_1(y) \geq c$. By symmetry, we get $\phi_k(e) > c$ for all k in state $e \equiv y - e_1 = (1, \dots, 1)$. The capacity constraints (1) are violated.

- Assume that $a(K - 1) < 1$. Let $x = (1, n, m, \dots, m)$ where $m \geq n \geq 1/a$ and n is the smallest integer such that $n(1 - a) \geq ma(K - 2)$. As link $r_2 \cap \dots \cap r_K$ is the only saturated link in states x and $x - e_1$, we have $\phi_2(x) = \phi_2(x - e_1)$ and $\phi_1(x) = a$ for any utility-based allocation. The balance property would imply that $\phi_1(x - e_2) = a$. Now as $\phi_3(x - e_2) = \dots = \phi_K(x - e_2)$, it follows from the capacity constraints (1) that $\phi_3(x - e_2) \leq (1 - a)/(K - 2)$. But, as the network reduces to a single link in state $x - e_2 - e_1$, we have $\phi_3(x - e_2 - e_1) = m/(n - 1 + m(K - 2))$. Thus $\phi_3(x - e_2) < \phi_3(x - e_2 - e_1)$ and the balance property would imply that $\phi_1(x - e_2) < \phi_1(x - e_2 - e_3)$. Hence $\phi_1(x - e_2 - e_3) > a$ and the capacity constraints (1) are violated.

□

Lemma 9 *A homogeneous network which does not contain a hypercycle, an incomplete line nor an incomplete square is a hypercube.*

Proof. The proof is by induction on the number of distinct routes N . The result holds for $N = 1$ and $N = 3$ (note that $N \neq 2$). A homogeneous network of three routes is indeed either a line or a triangle. Assume that the result holds for $N \leq M$, with $M \geq 3$. Consider a homogeneous network of $M + 1$ routes $\mathcal{N}_0 = \{r_0, r_1, \dots, r_M\}$ which does not contain a hypercycle, an incomplete line nor an incomplete square. Denote by \mathcal{L}_0 the set of links of \mathcal{N}_0 . We can choose a route, say r_0 , such that $\mathcal{N} = \mathcal{N}_0 \setminus r_0$ is a connected network. Denote by $\mathcal{L} \subset \mathcal{L}_0$ the set of limiting links of \mathcal{N} . By assumption, \mathcal{N} is a hypercube. We denote by n the dimension of this hypercube and $\mathcal{D}_1, \dots, \mathcal{D}_n$ the corresponding directions.

We first show that we can choose r_0 such that $n \geq 2$. Assume that for any k such that $\mathcal{N}_0 \setminus r_k$ is connected, this network reduces to a single link. If there is a route, say r_1 , such that $\mathcal{N}_0 \setminus r_1$ is not connected, this network has two limiting links, $r_1 \cap r_0$ and $r_1 \cap \dots \cap r_M$. But this implies $r_2 = r_3$, a contradiction. Thus for any k , the network $\mathcal{N}_0 \setminus r_k$ is connected. This network reduces to a single link, which is the intersection of all routes except r_k . We conclude that \mathcal{N}_0 is a hypercycle, a contradiction.

We can thus assume that $n \geq 2$. We first consider the case where a direction of \mathcal{N} contains a single route, say $\mathcal{D}_1 = \{r_1\}$.

- Assume $n = 2$. Then $\mathcal{L}_0 \neq \mathcal{L}$, otherwise \mathcal{N}_0 would be an incomplete line. Any link of $\mathcal{L}_0 \setminus \mathcal{L}$ is the intersection of r_0 and r_1 or r_0 and a single route r_2 of \mathcal{D}_2 . Otherwise, one would obtain a triangle. In the former case, \mathcal{N}_0 is a hypercube of directions \mathcal{D}_1 and $\mathcal{D}_2 \cup r_0$. In the latter case, $\mathcal{L}_0 \setminus \mathcal{L}$ is the set of all intersections $r_0 \cap r_2$, $r_2 \in \mathcal{D}_2$. Otherwise, \mathcal{N}_0 would contain an incomplete square. Thus \mathcal{N}_0 is a hypercube of directions $\mathcal{D}_1 \cup r_0$ and \mathcal{D}_2 .
- Assume $n \geq 3$. Then $\mathcal{N} \setminus r_1$ is a hypercube of dimension $n - 1$. The network $\mathcal{N}_0 \setminus r_1$ is connected. Otherwise, \mathcal{N}_0 would contain an incomplete line. By assumption, $\mathcal{N}_0 \setminus r_1$ is a hypercube, of dimension $n - 1$ or n . If $\mathcal{N}_0 \setminus r_1$ is of dimension n , the corresponding directions are $r_0, \mathcal{D}_2, \dots, \mathcal{D}_K$. Since $r_0 \neq r_1$, there exists $r_2 \in \mathcal{D}_2, \dots, r_n \in \mathcal{D}_n$ such that $r_0 \cap r_2 \cap \dots \cap r_n \neq r_1 \cap r_2 \cap \dots \cap r_n$. We conclude that routes r_0 and r_1 do not cross each other, otherwise they would form a triangle with one of the routes r_2, \dots, r_n . Thus \mathcal{N}_0 is a hypercube of directions $\mathcal{D}_1 \cup r_0, \mathcal{D}_2, \dots, \mathcal{D}_n$. If $\mathcal{N}_0 \setminus r_1$ is of dimension $n - 1$, there exists a direction of \mathcal{N} , say \mathcal{D}_2 , such that the directions of this hypercube are $\mathcal{D}_2 \cup r_0, \mathcal{D}_3, \dots, \mathcal{D}_n$. In particular, all directions of $\mathcal{N}_0 \setminus r_1$ contain at least two routes.

We can thus restrict the analysis to the case where all directions of \mathcal{N} contain at least two routes. We consider two cases.

- Assume $\mathcal{L} \neq \mathcal{L}_0$. Let $l \in \mathcal{L}_0 \setminus \mathcal{L}$. There exists a direction of \mathcal{N} , say \mathcal{D}_1 , such that any route r_1 of this direction does not contain link l . Assume that two routes of the same direction contain link l , say $r_2, r'_2 \in \mathcal{D}_2, r_2 \neq r'_2$. Since r_1 crosses r_2 and r'_2 , these three routes form a triangle, a contradiction. The link l is thus the intersection of r_0 and at most one route of each direction $\mathcal{D}_2, \dots, \mathcal{D}_n$. In addition, the network $\mathcal{N} \setminus r_1$ is a hypercube of dimension n . As the network $\mathcal{N}_0 \setminus r_1$ is connected, it is a hypercube of dimension n or $n + 1$. But link l belongs to n routes of $\mathcal{N}_0 \setminus r_1$. Thus $\mathcal{N}_0 \setminus r_1$ is a hypercube of dimension n . The directions of this hypercube are $\{\mathcal{D}_1 \setminus r_1\} \cup r_0, \mathcal{D}_2, \dots, \mathcal{D}_n$. Since this result holds for any route $r_1 \in \mathcal{D}_1$, \mathcal{N}_0 is a hypercube of dimension n , with directions $\mathcal{D}_1 \cup r_0, \mathcal{D}_2, \dots, \mathcal{D}_n$.
- Assume $\mathcal{L} = \mathcal{L}_0$. Let $l \in r_0$. Assume there exists a link $l' \in \mathcal{L}$ such that $l' \notin r_0$. There exists a direction of \mathcal{N} , say \mathcal{D}_1 , such that $l \in r_1$ and $l' \in r'_1, r_1, r'_1 \in \mathcal{D}_1, r_1 \neq r'_1$. The network $\mathcal{N} \setminus r_1$ is a hypercube of dimension n . As the network $\mathcal{N}_0 \setminus r_1$ is connected, it is a hypercube of dimension n or $n + 1$. But link l' belongs to n routes of $\mathcal{N}_0 \setminus r_1$. Thus $\mathcal{N}_0 \setminus r_1$ is a hypercube of dimension n . Since link l belongs to r_0 and to one route of each direction $\mathcal{D}_2, \dots, \mathcal{D}_n$, the directions of this hypercube are $\{\mathcal{D}_1 \setminus r_1\} \cup r_0, \mathcal{D}_2, \dots, \mathcal{D}_n$. Since r_0 contains links of \mathcal{L} only, these links are $r_1 \cap r_2 \cap \dots \cap r_n$, with $r_2 \in \mathcal{D}_2, \dots, r_n \in \mathcal{D}_n$. Thus $r_0 = r_1$, a contradiction. We conclude that r_0 contains all links of \mathcal{L} , thus \mathcal{N}_0 is a hypercube of dimension $n + 1$, with directions $r_0, \mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n$.

□

Proof of Theorem 3. Consider a network for which a utility-based allocation is insensitive. In view of Theorem 1, the allocation is balanced. It then follows from Lemmas 2-3 that the network is homogeneous, and from Lemmas 6-7-8-9 that the network is a hypercube. In view of Lemma 4, there is at most one limiting access rate. If there is one limiting access rate, the network reduces to a single link in view of Lemma 5 and the fact that any hypercube not reduced to a single link contains a 2-link line. If there is no limiting access rate, the allocation coincides with proportional fairness in view of Proposition 4 and Remark 3. □

References

- [1] F. Baccelli and P. Brémaud, *Palm-Martingale Calculus and Stochastic Recurrences*, Springer-Verlag, 1994.
- [2] 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.
- [3] D. Bertsekas and R. Gallager, *Data Networks*, Prentice Hall, 1987.
- [4] T. Bonald and L. Massoulié, Impact of fairness on Internet performance, in: *Proc. of ACM SIGMETRICS*, 2001.
- [5] T. Bonald and A. Proutière, Insensitivity in processor-sharing networks, *Performance Evaluation* 49 (2002) 193–209.
- [6] T. Bonald, A. Proutière, G. Régnié, J.W. Roberts, Insensitivity results in statistical bandwidth sharing, in: *Proc. of ITC 17th*, 2001.

- [7] D. Clark, Adding service discrimination to the Internet, *Telecommunications Policy* 20 (1996) 169–181.
- [8] M. Crovella and A. Bestavros, Self-similarity in WWW traffic: Evidence and possible cause, in: *Proc. of ACM SIGMETRICS*, 1996.
- [9] G. De Veciana, T.J. Lee and T. Konstantopoulos, “Stability and performance analysis of networks supporting services with rate control - Could the Internet be unstable?”, in: *Proc. of IEEE INFOCOM*, 1999.
- [10] A. K. Erlang, Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges, *Elektroteknikeren* 13 (1917).
- [11] 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.
- [12] G. Fayolle, I. Mitrani and R. Iasnogorodski, Sharing a processor among many classes, *Journal of the ACM* 27 (1980) 519–532.
- [13] V. Jacobson, Congestion avoidance and control, in: *Proc. of ACM SIGCOMM*, 1988.
- [14] F.P. Kelly, *Reversibility and Stochastic Networks*, Wiley, 1979.
- [15] F.P. Kelly, Mathematical modelling of the Internet, in: *Mathematics Unlimited - 2001 and Beyond*, eds B. Engquist and W. Schmid, Springer-Verlag, Berlin, 2001, 685–702.
- [16] F.P. Kelly, A. Maulloo and D. Tan, Rate control for communication networks: Shadow prices, proportional fairness and stability, *Journal of the Operational Research Society* 49 (1998).
- [17] S. Kunniyur and R. Srikant, End-to-end congestion control schemes: Utility functions, random losses and ECN marks, in: *Proc. of IEEE INFOCOM*, 2000.
- [18] R.J. La and V. Anantharam, Charge-sensitive TCP and rate control in the Internet, in: *Proc. of IEEE INFOCOM*, 2000.
- [19] L. Massoulié, Stability of distributed congestion control with heterogeneous feedback delays, *IEEE Transactions on Automatic Control* 47-6 (2002) 895–902.
- [20] L. Massoulié and J.W. Roberts, Bandwidth sharing and admission control for elastic traffic, *Telecommunication Systems* 15 (2000) 185–201.
- [21] L. Massoulié and J.W. Roberts, Bandwidth sharing: Objectives and algorithms, *IEEE/ACM Transactions on Networking* 10-3 (2002) 320–328.
- [22] M. Miyazawa, R. Schassberger and V. Schmidt, On the structure of an insensitive generalized semi-Markov process with reallocation and with point-process input, *Adv. Appl. Prob.* 27 (1995) 203–225.
- [23] J. Mo and J. Walrand, Fair end-to-end window-based congestion control, *IEEE/ACM Transactions on Networking* 8-5 (2000) 556–567.
- [24] V. Paxson and S. Floyd, Difficulties in Simulating the Internet, *IEEE/ACM Transactions on Networking* 9-4 (2001) 392–403.
- [25] R.F. Serfozo, *Introduction to Stochastic Networks*, Springer, 1999.