The multi-source model for dimensioning data networks

Thomas Bonald, Céline Comte

► To cite this version:

Thomas Bonald, Céline Comte. The multi-source model for dimensioning data networks. Computer Networks, 2016, 10.1016/j.comnet.2016.03.019 . hal-01314992

HAL Id: hal-01314992

<https://hal.science/hal-01314992>

Submitted on 12 May 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.

The Multi-Source Model for Dimensioning Data Networks

T. Bonald¹, C. Comte

Telecom ParisTech, Université Paris-Saclay, France

Abstract

Traffic modeling is key to the dimensioning of data networks. Usual models rely on the implicit assumption that each user generates data flows in series, one after the other, the ongoing flows sharing equitably the considered network link. We relax this assumption and consider the more realistic case where users may generate several data flows in parallel, these flows having to share the user's access line as well. We qualify this model as *multi-source* since each user now behaves as an independent traffic source. Usual performance metrics like mean throughput and congestion rate must now be defined at user level rather than at flow level. We derive explicit expressions for these performance metrics under the assumption that flows share bandwidth according to balanced fairness. These results are compared with those obtained by simulation when max-min fairness is imposed, either at flow level or at user level.

Keywords: Flow-level model, mean throughput, congestion rate, balanced fairness, max-min fairness.

1. Introduction

Internet service providers need to predict the impact of traffic load on the quality of service perceived by their customers. This is increasingly important with the advent of high-speed internet access that tends to move congestion from the access to the backhaul, where resources are shared by several users.

Internet traffic is most often modeled at flow level², assuming some ideal bandwidth sharing between ongoing flows [15, 3, 2, 1, 19, 27, 17]. Modeling traffic at packet level proves too complex and is hardly effective, given that users typically perceive quality of service at flow level [12]. In fact, the flow-level models of data networks can be considered as the analogues for the Erlang model

Email addresses: thomas.bonald@telecom-paristech.fr (T. Bonald),
celine.comte@telecom-paristech.fr (C. Comte)

¹Thomas Bonald is a member of LINCS. See www.lincs.fr.

²A flow is here defined as the set of packets having the same 5-tuple: IP source and destination addresses, IP source and destination ports, protocol.

of telephone networks and its extensions to multi-rate circuit-switched networks [5]. They have proved essential for both dimensioning [3, 6, 29, 25, 24, 16] and traffic engineering [15, 23, 21, 28].

These models rely on the implicit assumption that each user generates data flows in series, one after the other, so that bandwidth sharing occurs on the considered backhaul link only, and not on the user's access line. In this paper, we relax this assumption and consider the more realistic case where users may generate several data flows in parallel, these flows having to share both the backhaul link and the user's access line. It is not obvious how bandwidth is shared by end-to-end congestion control in this context. Proportional fairness is often considered as an adequate model [18, 26, 22]. It turns out to coincide with max-min fairness in the networks considered in the present paper. Since we look for closed-form expressions, we consider a slightly different allocation known as balanced fairness [7]. While the resulting expressions have already been presented in [4], we here resort to simulations to compare our results to those obtained under max-min fairness. We also use simulations to assess the impact of packet schedulers and buffer management schemes that impose max-min fairness at user level on the backhaul link.

Existing models assume that flows are generated either according to a Poisson process (the so-called *infinite-source* model) or by n users, each alternating between the active state and the idle state (the so-called *finite-source* model) [5]; in both cases, each user has at most one flow in progress at any given time. Our model consists of n users, each generating data flows according to a Poisson process; in particular, there is no limit on the number of flows in progress coming from the same user. We refer to this model as the *multi-source* model, since each user can now be viewed as an independent source of flows, as opposed to previous models where there is a unique source of flows, able to generate either an infinite number or a finite number of flows in parallel.

Since the model allows each user to generate multiple flows in parallel, it is not sufficient to focus on the flow level to evaluate user-level performance. In particular, the throughput of each user is the *total* throughput of her flows in progress. The corresponding performance results can differ significantly from those obtained under the infinite-source model and the finite-source model. They coincide only in the limit of an infinite number of access lines.

We consider a single backhaul link of fixed capacity (in bit/s) shared by a population of n users. The total throughput of each user is also constrained by the rate of her access line. The models apply equally to the uplink (from the users to the Internet) and to the downlink (from the Internet to the users). All traffic is elastic, meaning that each flow generated by a user corresponds to some data transfer and remains active as long as the corresponding data have not been fully transferred. Due to the insensitivity property, we do not specify the flow size distribution beyond the mean, nor the distribution of the idle times in the finite-source model [5]. We provide formulas that can be used in planning tools or directly by network engineers to get insights into the impact of traffic on user-level performance [16].

In the rest of the paper, we first review the results obtained with the infinite-

source model and the finite-source model. We then present the multi-source model and compare numerically the results obtained with the three models in some typical traffic scenarios. Finally, we present the simulation results obtained under max-min fairness and conclude the paper.

2. Infinite-source model

Like the Erlang model for telephone networks, which relies on the assumption of Poisson call arrivals [14], it is common practice to assume Poisson flow arrivals in data networks. This is referred to as the infinite-source model since it corresponds to the finite-source model (presented in the next section) in the limiting case where the number of users n grows to infinity.

2.1. No access rates

We start with the simplest case where there is no rate limit at the access: each user has full access to the backhaul link, which is assumed to be equitably shared by ongoing flows. Flows arrive according to a Poisson process of intensity λ and have i.i.d. sizes of mean σ bits, corresponding to a traffic intensity of $A = \lambda\sigma$ bit/s. Denoting by C the capacity of the backhaul link in bit/s, the link load is $\rho = A/C$. Under the assumption of perfect fair sharing, the traffic model corresponds to an $M/G/1$ processor-sharing queue of load ρ . It is stable if and only if $\rho < 1$, in which case the stationary distribution of the number of flows in progress X is given by:

$$\pi(x) = (1 - \rho)\rho^x.$$

Recall that this distribution is insensitive to the flow size distribution beyond the mean [20, 2, 5].

It turns out that the stationary distribution seen by a user having a flow in progress is different. Since there are x flows in progress in state x , which are assumed to be generated by different users, each active user sees the *size-biased* probability distribution of the random variable X ,

$$\pi'(x) \propto x\pi(x).$$

Observe that $\pi'(0) = 0$. We will denote by P' and E' the corresponding probability measure and expectation, respectively. We derive two key performance metrics on this basis.

Mean throughput. The first performance metric is the mean throughput experienced by users. Assume there are x ongoing flows, with $x > 0$. The throughput of each flow is then C/x . Thus the mean throughput experienced by users, normalized by the maximum throughput C , is given by

$$\gamma = E' \left(\frac{1}{X} \right).$$

Replacing E' by its expression, we obtain

$$\gamma = \frac{\sum_{x>0} \pi(x)}{E(X)} = \frac{\rho}{E(X)},$$

that is

$$\gamma = 1 - \rho. \quad (1)$$

Observe that the mean throughput decreases linearly with the link load.

Congestion rate. The second performance metric is the congestion rate, defined as the probability that an active user gets a throughput less than the maximum throughput C . Since there is no rate limit at the access, the congestion rate is the probability seen by an active user that there are other active users:

$$\eta = P'(X > 1).$$

Replacing E' by its expression, we obtain

$$\eta = \frac{\sum_{x>1} x\pi(x)}{E(X)} = 1 - \frac{P(X = 1)}{E(X)},$$

that is

$$\eta = \rho(2 - \rho). \quad (2)$$

As expected, the congestion rate grows from 0 to 1 as the link load grows from 0 to 1.

2.2. Same access rates

We now consider the practically interesting case where each flow has a rate limit $r < C$ corresponding to the capacity of the user's access line in bit/s. For convenience, we assume that the capacity of the backhaul link is some multiple of this access rate, that is $C = mr$ for some integer $m \geq 1$. We denote by $\alpha = A/r$ the traffic intensity expressed in units of the access rate. This would correspond to the mean number of flows if the backhaul link were of infinite capacity. The model corresponds to an $M/G/m$ processor-sharing queue of load $\rho = A/C$. Under the stability condition $\rho < 1$, the stationary distribution of the number of flows X is given by [5]:

$$\pi(x) = \frac{1}{G} \begin{cases} \frac{\alpha^x}{x!} & \text{for } x \leq m, \\ \frac{\alpha^m}{m!} \rho^{x-m} & \text{for } x > m, \end{cases}$$

where G denotes the normalization constant:

$$G = \sum_{x=0}^m \frac{\alpha^x}{x!} + \frac{\alpha^m}{m!} \frac{\rho}{1 - \rho}.$$

Both performance metrics extend to this case.

2.3. Different access rates

Finally, we consider the general case of K different access rates $r_1, \dots, r_K$. We denote by $A_1, \dots, A_K$ the respective traffic intensities in bit/s generated by each class of users, and by $\alpha_1 = A_1/r_1, \dots, \alpha_K = A_K/r_K$ the traffic intensities expressed in multiples of the access rates; these would correspond to the mean number of flows of each class if the backhaul link were of infinite capacity. The corresponding loads on the backhaul link are $\rho_1 = A_1/C, \dots, \rho_K = A_K/C$, and the total load is $\rho = \rho_1 + \dots + \rho_K$.

Let X be the K -dimensional vector of the number of flows of each class in progress. Denote by $\phi_k(x)$ the total throughput of class- k users in state x . The capacity constraints are

$$\forall k = 1, \dots, K, \quad \phi_k(x) \leq x_k r_k$$

and

$$\sum_{k=1}^K \phi_k(x) \leq C.$$

Now let r be the K -dimensional vector of access rates. Under balanced fair sharing [7, 10], all users get their maximum throughput, in the sense that $\phi_k(x) = x_k r_k$ for all $k = 1, \dots, K$, if and only if $x \cdot r \leq C$ (the access lines are limiting); otherwise, no user gets her or his maximum throughput and the total throughput is $\sum_{k=1}^K \phi_k(x) = C$ (the backhaul link is limiting). The stability condition is $\rho < 1$ and the vector X has the stationary distribution:

$$\pi(x) = \begin{cases} \frac{1}{G} \frac{\alpha_1^{x_1}}{x_1!} \dots \frac{\alpha_K^{x_K}}{x_K!} & \text{for } x \cdot r \leq C, \\ \sum_{k=1}^K \rho_k \pi(x - e_k) & \text{for } x \cdot r > C, \end{cases}$$

where e_k is the unit vector on component k and G denotes the normalization constant. Here and in the rest of the paper, we adopt the convention that $\pi(x) = 0$ for any $x \notin \mathbb{N}^K$. Performance now depends on the user's class. Both metrics can be computed through a recursive formula [10], which is the analogue of the Kaufmann-Roberts formula for circuit-switched networks.

3. Finite-source model

When the user population is relatively small, flow arrivals cannot be considered as Poisson. Each user is still assumed to generate flows in series, with a random idle time between the end of a flow and the beginning of the next flow. This is the analogue of the Engset model used for telephone networks [13]. We only give the stationary distribution of the number of active users; the corresponding performance metrics can be derived as for the infinite-source model.

3.1. No access rates

Consider n users having full access to the backhaul link. Any idle user tends to become active at rate $\nu > 0$, while any active user tends to become idle at rate $\mu = C/\sigma$ when no other users are active. We deduce that any user alone in the system is active a fraction of time $\beta/(1 + \beta)$, with $\beta = \nu/\mu$.

Now assume active users share the backhaul link in a fair way. The stationary distribution of the number of active users X is then given by [3, 9]:

$$\pi(x) = \frac{1}{G} \frac{n!}{(n-x)!} \beta^x, \quad x \leq n,$$

where G is the normalization constant:

$$G = \sum_{x=0}^n \frac{n!}{(n-x)!} \beta^x.$$

The infinite-source model corresponds to the case $n \rightarrow \infty$ and $\beta \rightarrow 0$, with $n\beta \rightarrow \rho$. A key difference with the infinite-source model is that traffic intensity is no longer an exogenous parameter but given by $A = CP(X > 0)$. We deduce the link load:

$$\rho = P(X > 0) = \frac{G-1}{G}.$$

3.2. Same access rates

Now assume all users have the same access rate r . The link capacity is $C = mr$ for some integer $m \geq 1$, with $n > m$. Any active user tends to become idle at rate $\mu = r/\sigma$ when no other users are active. The stationary distribution of the number of flows X becomes [6]:

$$\pi(x) = \frac{1}{G} \begin{cases} \binom{n}{x} \beta^x & \text{for } x \leq m, \\ \frac{n!}{(n-x)!m!} \frac{\beta^x}{m^{x-m}} & \text{for } m < x \leq n, \end{cases}$$

where $\beta = \nu/\mu$ and G is the normalization constant. Traffic intensity is $A = E(\min(X, m))r$, corresponding to load

$$\rho = \frac{E(\min(X, m))}{m}.$$

3.3. Different access rates

Finally, consider the general case of K different access rates $r_1, \dots, r_K$. There are n_k users with access rate r_k , mean flow size σ_k and mean idle time $1/\nu_k$ between two flows. Under balanced fair sharing, the stationary distribution of the system state X is given by:

$$\pi(x) = \begin{cases} \frac{1}{G} \prod_{k=1}^K \binom{n_k}{x_k} \beta_k^{x_k} & \text{for } x.r \leq C, \\ \sum_{k=1}^K \frac{\beta_k r_k}{C} (n_k - x_k + 1) \pi(x - e_k) & \text{for } x.r > C, \end{cases}$$

where $\beta_k = \nu_k/\mu_k$, $\mu_k = r_k/\sigma_k$ and G the normalization constant. Traffic intensity is $A = E(\min(X.r, C))$, corresponding to load:

$$\rho = \frac{E(\min(X.r, C))}{C}.$$

4. Multi-source model

We now introduce the multi-source model where data flows must share *both* the backhaul link and the user's access line. We consider n users, with user i generating flows according to an independent Poisson process of intensity λ_i , corresponding to the traffic intensity $a_i = \lambda_i \sigma$ in bit/s. We are interested in the total throughput obtained by each user.

4.1. No access rates

As above, we first consider the case where each user has full access to the backhaul link. Under fair sharing between flows in progress, the model reduces to an $M/G/1$ multi-class processor-sharing queue. Denoting by $\rho_i = a_i/C$ the load due to user i and by $\rho = \rho_1 + \dots + \rho_n$ the total load, the stationary distribution of the number of flows of each user X is given by

$$\pi(x) = (1 - \rho) \binom{x_1 + \dots + x_n}{x_1, \dots, x_n} \rho_1^{x_1} \dots \rho_n^{x_n},$$

under the stability condition $\rho < 1$.

Now user i sees the stationary distribution $\pi_i(x) \propto \pi(x) 1_{x_i > 0}$ when active. We denote by P_i and E_i the corresponding probability measure and expectation.

Mean throughput. The total throughput of user i is proportional to the number of ongoing flows of this user, that is $(x_i / \sum_j x_j) \times C$ in any state x such that $x_i > 0$. We deduce the mean throughput of user i , normalized by the maximum throughput C ,

$$\gamma_i = E_i \left(\frac{X_i}{\sum_j X_j} \right).$$

By work conservation,

$$\rho_i = E \left(\frac{X_i}{\sum_j X_j} 1_{X_i > 0} \right),$$

so that

$$\gamma_i = \frac{\rho_i}{P(X_i > 0)}.$$

Since

$$P(X_i > 0) = \frac{\rho_i}{1 - \rho + \rho_i},$$

we obtain $\gamma_i = 1 - \rho + \rho_i$. Note that the mean throughput is larger than that obtained with the infinite-source model, given by (1), with equality when $\rho_i \rightarrow 0$ (in which case user i generates flows in series, as in the infinite-source model). Observe also that the mean (normalized) throughput of user i is larger than the load ρ_i generated by this user, with equality when $\rho \rightarrow 1$ (in which case the system is saturated and the throughput of each user corresponds to her bandwidth share). For homogeneous traffic distribution, all users get the same throughput,

$$\gamma = 1 - \rho + \frac{\rho}{n} \geq \frac{1}{n}.$$

Congestion rate. The congestion rate seen by user i is the probability that the total throughput of this user is less than C , that is the probability that there are other active users:

$$\eta_i = P_i(\sum_j X_j > X_i).$$

We get

$$\eta_i = \frac{P(X_i > 0, \sum_{j \neq i} X_j > 0)}{P(X_i > 0)}.$$

Since

$$P(X_i > 0, \sum_{j \neq i} X_j = 0) = \frac{\rho_i(1 - \rho)}{1 - \rho_i},$$

we obtain

$$\eta_i = (2 - \rho) \frac{\rho - \rho_i}{1 - \rho_i}.$$

This congestion rate is smaller than that obtained with the infinite-source model, given by (2), with equality when $\rho_i \rightarrow 0$.

4.2. Same access rates

Now assume all users have the same access rate r , with $C = mr$ for some integer m such that $1 \leq m < n$. We denote by $\varrho_i = a_i/r$ the load of user- i access line. The load of user i on the backhaul link is $\rho_i = a_i/C = \varrho_i/m$.

Let $\phi_i(x)$ be the total throughput of user i in state x . The capacity constraints are

$$\forall i = 1, \dots, n, \quad \phi_i(x) \leq r$$

and

$$\sum_{i=1}^n \phi_i(x) \leq C.$$

Let $n(x) = \sum_{i=1}^n 1_{x_i > 0}$ be the number of active users in state x . Under balanced fair sharing, all active users get their maximum throughput, that is $\phi_i(x) = r$ for all $i = 1, \dots, n$ such that $x_i > 0$, if and only if $n(x)r \leq C$ (the access lines are limiting); otherwise, no user gets the maximum throughput and the total throughput is $\sum_{i=1}^n \phi_i(x) = C$ (the backhaul link is limiting). Under the stability condition $\rho < 1$ and $\varrho_i < 1$ for all $i = 1, \dots, n$, the stationary distribution of the network state X is

$$\pi(x) = \begin{cases} \frac{1}{G} \prod_{i=1}^n \varrho_i^{x_i} & \text{for } n(x) \leq m, \\ \sum_{i=1}^n \rho_i \pi(x - e_i) & \text{otherwise,} \end{cases}$$

where G is the normalization constant. Since the network has a tree topology, we deduce from [8] that

$$G = \sum_{I \subset \{1, \dots, n\}, |I| \leq m} \prod_{i \in I} \frac{\varrho_i}{1 - \varrho_i} + \sum_{I \subset \{1, \dots, n\}, |I| = m} \prod_{i \in I} \frac{\varrho_i}{1 - \varrho_i} \frac{\sum_{i \notin I} \varrho_i}{m - \sum_{i \in I} \varrho_i}.$$

Mean throughput. The mean throughput of user i , normalized by the maximum throughput r , is given by

$$\gamma_i = E_i \left(\frac{\phi_i(X)}{r} \right).$$

By work conservation, $E(\phi_i(X)) = a_i$ so that

$$\gamma_i = \frac{\varrho_i}{P(X_i > 0)}.$$

Now

$$P(X_i = 0) = \frac{G_i}{G},$$

where G_i denote the normalization constant in the absence of user i ,

$$G_i = \sum_{I \subset \{1, \dots, n\} \setminus \{i\}, |I| \leq m} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} + \sum_{I \subset \{1, \dots, n\} \setminus \{i\}, |I| = m} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} \frac{\sum_{j \notin I, j \neq i} \varrho_j}{m - \sum_{j \neq i} \varrho_j}.$$

We deduce

$$\gamma_i = \frac{G \varrho_i}{G - G_i}. \quad (3)$$

Congestion rate. The congestion rate seen by user i is

$$\eta_i = P_i(\phi_i(X) < r).$$

We get

$$\eta_i = \frac{P(X_i > 0, \sum_j 1_{X_j > 0} > m)}{P(X_i > 0)},$$

that is

$$\eta_i = \frac{F - F_i}{G - G_i}, \quad (4)$$

with

$$F = \sum_{I \subset \{1, \dots, n\}, |I| = m} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} \frac{\sum_{j \notin I} \varrho_j}{m - \sum_{j=1}^n \varrho_j}$$

and

$$F_i = \sum_{I \subset \{1, \dots, n\} \setminus \{i\}, |I| = m} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} \frac{\sum_{j \notin I, j \neq i} \varrho_j}{m - \sum_{j \neq i} \varrho_j}.$$

4.3. Different access rates

We now consider the general case where user i has access rate r_i . The load of user- i access line becomes $\varrho_i = a_i/r_i$. The capacity constraints are

$$\forall i = 1, \dots, n, \quad \phi_i(x) \leq r_i$$

and

$$\sum_{i=1}^n \phi_i(x) \leq C.$$

Under balanced fair sharing, all active users get their maximum throughput if and only if

$$\sum_{i=1}^n r_i 1_{x_i > 0} \leq C.$$

Under the stability condition $\rho < 1$ and $\varrho_i < 1$ for all $i = 1, \dots, n$, the stationary distribution of the network state X is

$$\pi(x) = \begin{cases} \frac{1}{G} \prod_{i=1}^n \varrho_i^{x_i} & \text{for } \sum_{i=1}^n r_i 1_{x_i > 0} \leq C, \\ \sum_{i=1}^n \rho_i \pi(x - e_i) & \text{otherwise,} \end{cases}$$

where G denotes the normalization constant [8]

$$G = \sum_{I \subset \{1, \dots, n\}, e_I \cdot r \leq C} \prod_{i \in I} \frac{\varrho_i}{1 - \varrho_i} + \sum_{I \subset \{1, \dots, n\}, e_I \cdot r \leq C} \prod_{i \in I} \frac{\varrho_i}{1 - \varrho_i} \frac{\sum_{i \notin I, r_i + e_I \cdot r > C} \varrho_i r_i}{C - \sum_{i=1}^n \varrho_i r_i},$$

with e_I the vector of ones on all components $i \in I$ and zeros elsewhere.

Mean throughput. The mean throughput of user i normalized by the maximum throughput r_i of this user is

$$\gamma_i = E_i \left(\frac{\phi_i(X)}{r_i} \right).$$

We obtain the same expression (3), with

$$G_i = \sum_{I \subset \{1, \dots, n\}, i \notin I, e_I \cdot r \leq C} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} + \sum_{I \subset \{1, \dots, n\}, i \notin I, e_I \cdot r \leq C} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} \frac{\sum_{j \notin I, j \neq i, r_j + e_I \cdot r > C} \varrho_j r_j}{C - \sum_{j \neq i} \varrho_j r_j}.$$

Congestion rate. The congestion rate seen by user i is

$$\eta = P_i(\phi_i(X) < r_i),$$

that is (4) with

$$F = \sum_{I \subset \{1, \dots, n\}, e_I \cdot r \leq C} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} \frac{\sum_{j \notin I, r_j + e_I \cdot r > C} \varrho_j r_j}{C - \sum_j \varrho_j r_j}$$

and

$$F_i = \sum_{I \subset \{1, \dots, n\}, i \notin I, e_I \cdot r \leq C} \prod_{j \in I} \frac{\varrho_j}{1 - \varrho_j} \frac{\sum_{j \notin I, j \neq i, r_j + e_I \cdot r > C} \varrho_j r_j}{C - \sum_{j \neq i} \varrho_j r_j}.$$

4.4. Different user classes

The previous formulas have exponential complexity in n . To keep the complexity polynomial in n , we need to group users in some finite number of classes K , as in the infinite-source model and the finite-source model. With some slight abuse of notation, we denote respectively by r_k and ϱ_k the rate and the load of the access line of each class- k user. There are n_k class- k users and we denote by n the vector $(n_1, \dots, n_K)$. The normalization constant is then given by

$$G = \sum_{\ell \leq n: \ell.r \leq C} \prod_{k=1}^K \binom{n_k}{\ell_k} \left(\frac{\varrho_k}{1 - \varrho_k} \right)^{\ell_k} + \sum_{\ell \leq n: \ell.r \leq C} \prod_{k=1}^K \binom{n_k}{\ell_k} \left(\frac{\varrho_k}{1 - \varrho_k} \right)^{\ell_k} \times \frac{\sum_{j: \ell_j < n_j, \ell.r + r_j > C} (n_j - \ell_j) \varrho_j r_j}{C - \sum_{k=1}^K n_k \varrho_k r_k}.$$

The mean throughput of each class- k user is

$$\gamma_k = \frac{G \varrho_k}{G - G_k},$$

where G_k is the constant G with n_k replaced by $n_k - 1$. Similarly, letting

$$F = \sum_{\ell \leq n: \ell.r \leq C} \prod_{k=1}^K \binom{n_k}{\ell_k} \left(\frac{\varrho_k}{1 - \varrho_k} \right)^{\ell_k} \times \frac{\sum_{j: \ell_j < n_j, \ell.r + r_j > C} (n_j - \ell_j) \varrho_j r_j}{C - \sum_{k=1}^K n_k \varrho_k r_k},$$

we obtain the congestion rate of class- k users

$$\eta_k = \frac{F - F_k}{G - G_k},$$

where F_k is the constant F with n_k replaced by $n_k - 1$.

In the limit where the number of users tends to infinity with traffic intensities $n_1 \varrho_1 r_1, \dots, n_K \varrho_K r_K$ tending to some fixed constants $A_1, \dots, A_K$ such that $A_1 + \dots + A_K < C$, the multi-source model reduces to the infinite-source model: there is an infinite population of users, each user generating flows according to a Poisson process of null intensity.

4.5. Numerical results

Figures 1 and 2 compare the performance metrics obtained with the three considered models when all users have the same access rate r , $C = mr$ with $m = 1, 10, 100$, $n = 2m$ and $n = 10m$, respectively. Traffic distribution is homogeneous.

We observe that the infinite-source model is overly pessimistic while the finite-source model is overly optimistic compared to the multi-source model, especially for the mean throughput. For the congestion rate, the infinite-source model is a good approximation of the multi-source model only for $n = 10m$ while the finite-source model is a very good approximation in both cases.

Figure 1: Performance metrics under the three models for $n = 2m$.

Figure 2: Performance metrics under the three models for $n = 10m$.

Figure 3 shows the results obtained for two classes of users, class-2 users generating 10 times more traffic than class-1 users. There is the same number of class-1 and class-2 users. All users have the same access rate r , $C = mr$ with $m = 1, 10, 100$, and there is a total of $n = 10m$ users. We observe that neither the infinite-source model nor the finite-source model is able to predict the performance of both user classes: the underlying assumption of flows generated in series by each user is not satisfactory.

The same conclusion can be drawn from Figure 4, showing the results for two classes of users with different access rates, $r_1 = 1$ and $r_2 = 4$. We take $n_1 = 4n_2$ and the same load for all access lines so that the total traffic intensity is the same for each class. Here $m = C/r_1$ takes the values 5, 50, 500 and the total number of users is $n = 2m$.

Figure 3: Mean throughput for $m = 1, 10, 100$ (from bottom to top) and $n_1 = n_2 = 5m$.

Figure 4: Mean normalized throughput for $m = 5, 50, 500$ (from bottom to top), $n = 2m$, $n_1 = 4n_2$.

5. Imposing fairness

In this section, we study by simulation the impact of the sharing policy on the performance results obtained with the multi-source model. Specifically, we consider max-min fairness applied either at flow level or at user level.

Since the insensitivity property is no longer valid with max-min fairness [7], we need to specify the flow size distribution. We assume that this distribution is exponential so that the evolution of the network state is described by a Markov process. The transitions rates are λ_i for user- i flow arrivals and $\phi_i(x)/\sigma$ for user- i flow departures in state x , where $\phi(x)$ denotes the vector of bandwidth shares in state x under max-min fairness at flow level or at user level. This Markov process is ergodic under the usual stability condition $\rho < 1$ and $\varrho_i < 1$ for all $i = 1, \dots, n$ [11].

5.1. Flow-level fairness

We first consider the case of max-min fairness at flow level, which is representative of the sharing achieved by TCP in the considered network. In practice,

the Internet service provider can enforce max-min fairness at flow level by identifying flows through the usual 5-tuple³ in the IP header of each packet and by applying some adequate packet scheduler and buffer management scheme.

Since there is no explicit expression for the stationary distribution of the network state, we use simulations to get the corresponding performance metrics.

5.2. User-level fairness

When fairness is imposed at flow level, users having a large number of flows in progress typically get a higher bandwidth share than other users. The Internet service provider may rather impose fairness at *user* level to avoid this bias. Since all packets generated by the same user generally share the same IP source address on the uplink and the same IP destination address on the downlink, it is sufficient to identify users through the corresponding field in the IP header of each packet and to apply some fair packet scheduler and buffer management scheme on this basis at the backhaul link.

Assuming for instance that the n users are active and indexed in increasing order of their access rates, the bandwidth share of user i in state x is given by her access rate r_i if $i \leq k$ and

$$\frac{C - \sum_{j=1}^k r_j}{n - k}$$

otherwise, where k is the highest index l such that $\sum_{i=1}^l r_i + (n - l)r_l \leq C$. Observe that the allocation is the same for all states x such that $x_i > 0$ for all $i = 1, \dots, n$. Again, the corresponding Markov process does not have a closed-form stationary distribution and we need simulations to estimate the performance metrics.

5.3. Numerical results

Each result obtained by simulation is derived from the average of the considered performance metric over 10 independent runs of the corresponding Markov process, each consisting of $5 \cdot 10^5$ jumps after a warm-up period of $5 \cdot 10^5$ jumps. This allows us to get for each result a 95% confidence interval included in the plotted value ± 0.02 .

Figure 5 shows the results obtained when all users have the same access rate r , $C = mr$, $m = 1, 10, 100$ and $n = 10m$. We observe that the simulation results obtained with max-min fairness, either at flow level or at user level, are very close to the analytical results derived under balanced fairness.

We now consider the heterogeneous scenario of Figure 3 with two classes of users, class-2 users generating 10 times more traffic than class-1 users. The results are presented in Figure 6. Balanced fairness still provides a very good approximation of throughput performance under *flow-level* max-min fairness,

³IP source address, IP destination address, source port, address port, protocol.

Figure 5: Impact of fairness on performance under the multi-source model for $n = 10m$.

but under-estimates the mean throughput of class-1 users under *user-level* max-min fairness. This is due to the fact that class-2 users, who typically have a larger number of flows in progress than class-1 users, are favored under both balanced fairness and flow-level max-min fairness. Imposing fairness at user level allows the Internet service provider to protect users generating less traffic. Balanced fairness does not capture this phenomenon but provides conservative estimates of performance and thus can be used for dimensioning purposes.

Figure 6: Impact of fairness on throughput performance under the multi-source model for $m = 1, 10, 100$ (from bottom to top) and $n_1 = n_2 = 5m$.

Finally, we give in Figure 7 the results corresponding to the scenario of Figure 4, with two access rates. The throughput performance as estimated by balanced fairness is slightly optimistic for class-2 users and very pessimistic for class-1 users, especially under user-level max-min fairness. Again, this can be explained by the fact that max-min fairness at user level tends to protect users with low access rates, a phenomenon that is not captured by balanced fairness.

Figure 7: Impact of fairness on throughput performance under the multi-source model for $m = 5, 50, 500$ (from bottom to top), $n = 2m$, $n_1 = 4n_2$.

6. Conclusion

We have proposed a new traffic model for evaluating user-level performance in data networks. The key characteristic of this model is to account for bandwidth sharing on the user's access line. The results turn out to be very different from those obtained with usual models in practically interesting cases, like $n = 100$ users having different traffic profiles or access rates. They coincide only for large values of n , say $n \geq 1000$. Simulations show that the results are approximately the same under *flow-level* max-min fairness. When max-min fairness is imposed at *user level*, the throughput performance of users with low traffic or low access rate tends to be better than that estimated by balanced fairness.

One of the key benefits of the proposed multi-source model is to account precisely for the number of access lines n without the complexity of the finite-source model. For instance, traffic intensity (and thus link load) is an exogenous parameter of the multi-source model but an endogenous parameter of the finite-source model. Moreover, the normalization constant is explicit in the multi-source model, which greatly simplifies the computation of the performance metrics.

A drawback of the multi-source model compared to the infinite-source model is the lack of a recursive formula for evaluating the normalization constant in the presence of a large number of different access rates. We let this for future work. Other interesting issues include the derivation of more accurate approximations in case fairness is imposed at user level and extensions of the model to non-elastic traffic (for instance, adaptive streaming traffic).

References

- [1] C. Barakat, P. Thiran, G. Iannaccone, C. Diot, and P. Owezarski. Modeling internet backbone traffic at the flow level. *Signal Processing, IEEE Transactions on*, 51(8):2111–2124, 2003.

- [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 *ACM SIGCOMM Computer Communication Review*, volume 31, pages 111–122. ACM, 2001.
- [3] A. W. Berger and Y. Kogan. Dimensioning bandwidth for elastic traffic in high-speed data networks. *IEEE/ACM Trans. Netw.*, 8(5):643–654, Oct. 2000.
- [4] T. Bonald. Traffic Models for User-Level Performance Evaluation in Data Networks. In *Proceedings of ITC*, Ghent, Belgium, 2015.
- [5] T. Bonald and M. Feuillet. *Network Performance Analysis*. Wiley, 2011.
- [6] T. Bonald, P. Olivier, and J. Roberts. Dimensioning high-speed IP access networks. In *Proceedings of ITC*, pages 241–251, 2003.
- [7] T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks. *Queueing systems*, 44(1):69–100, 2003.
- [8] T. Bonald, A. Proutière, J. Roberts, and J. Virtamo. Computational aspects of balanced fairness. *Teletraffic Science and Engineering*, 5:801–810, 2003.
- [9] T. Bonald and M.-A. Tran. Balancing elastic traffic sources. *IEEE communications letters*, 11(8):692–694, 2007.
- [10] T. Bonald and J. Virtamo. A recursive formula for multirate systems with elastic traffic. *IEEE Communications Letters*, 9(8):753–755, 2005.
- [11] G. De Veciana, T.-J. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting elastic services. *Networking, IEEE/ACM Transactions on*, 9(1):2–14, 2001.
- [12] N. Dukkupati and N. McKeown. Why flow-completion time is the right metric for congestion control. *ACM SIGCOMM Computer Communication Review*, 36(1):59–62, 2006.
- [13] T. Engset. On the calculation of switches in an automatic telephone system. In A. Myskja and O. Espvik, editors, *Tore Olaus Engset: The man behind the formula*, 1998.
- [14] A. Erlang. Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. In E. Brockmeyer, H. Halstrom, and A. Jensen, editors, *The life and works of A.K. Erlang*, 1948.
- [15] D. P. Heyman, T. Lakshman, and A. L. Neidhardt. A new method for analysing feedback-based protocols with applications to engineering web traffic over the internet. In *ACM Sigmetrics performance evaluation review*, volume 25, pages 24–38. ACM, 1997.

- [16] R. Indre, P. Olivier, B. Kauffmann, and N. Benameur. Evaluating user-perceived performance in high-speed backhaul networks. In *16th International Telecommunications Network Strategy and Planning Symposium (Networks 2014)*, 2014.
- [17] F. Kelly and R. Williams. Fluid model for a network operating under a fair bandwidth-sharing policy. *Annals of Applied Probability*, pages 1055–1083, 2004.
- [18] F. P. Kelly, A. K. Maulloo, and D. K. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society*, pages 237–252, 1998.
- [19] A. A. Kherani and A. Kumar. Stochastic models for throughput analysis of randomly arriving elastic flows in the internet. In *INFOCOM 2002*, volume 2, pages 1014–1023. IEEE, 2002.
- [20] L. Kleinrock. *Theory, Volume 2, Queueing Systems*. Wiley-Interscience, 1975.
- [21] A. Kortebi, L. Muscariello, S. Oueslati, and J. Roberts. On the scalability of fair queueing. In *ACM HotNets-III*, 2004.
- [22] S. Kunniyur and R. Srikant. End-to-end congestion control schemes: Utility functions, random losses and ecn marks. *IEEE/ACM Transactions on Networking*, 11(5):689–702, 2003.
- [23] A. Lakshmikantha, C. Beck, and R. Srikant. Impact of file arrivals and departures on buffer sizing in core routers. *IEEE/ACM Transactions on Networking (TON)*, 19(2):347–358, 2011.
- [24] P. Lassila, A. Penttinen, and J. Virtamo. Dimensioning methods for data networks with flow-level QoS requirements. In *Local Computer Networks, 2007. LCN 2007. 32nd IEEE Conference on*, pages 353–359, Oct 2007.
- [25] P. Lassila, A. Penttinen, and J. Virtamo. Dimensioning of data networks: a flow-level perspective. *European Transactions on Telecommunications*, 20(6):549–563, 2009.
- [26] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, Oct. 2000.
- [27] S. Rácz, B. P. Gerő, and G. Fodor. Flow level performance analysis of a multi-service system supporting elastic and adaptive services. *Performance Evaluation*, 49(1):451–469, 2002.
- [28] V. Shah and G. De Veciana. High-performance centralized content delivery infrastructure: models and asymptotics. *Networking, IEEE/ACM Transactions on*, 23(5):1674–1687, 2015.

- [29] H. Van den Berg, M. Mandjes, R. Van de Meent, A. Pras, F. Roijers, and P. Venemans. QoS-aware bandwidth provisioning for IP network links. *Computer Networks*, 50(5):631–647, 2006.