Alcom-FT Technical Report Series
ALCOMFT-TR-01-16

MODELING MAX-MIN FAIRNESS FOR ELASTIC FLOWS IN TELECOMMUNICATION NETWORKS

ALAIN DUPUIS, FABRICE GUILLEMIN, AND PHILIPPE ROBERT

Abstract. An analytical model is developed in this paper to study max-min fairness in telecommunication networks. Max-min fairness is a popular bandwidth sharing principle, which is used to allocate bandwidth to elastic flows sharing the capacity of of the transmission links of a network. The model considered in the present paper is a linear single network link, on which originating points of data flows appear as a Poisson process; the length of a data is exponential with unit mean and the link transmission capacity is taken as unity. While the definition of the max-min fair share is algorithmically clear, any analytical computation rapidly proves intractable. To overcome this technical difficulty, we develop simple heuristic approximations. It turns out that the obtained formulae yield a satisfactory approximation of the max-min fair share. These approximations prove less accurate when considering the amount of unused bandwidth due to max-min fairness.

Keywords: Max-min fairness, $M/M/\infty$ queue, elastic flows, Charlier polynomials.

1. Introduction

Max-min fairness is a very popular bandwidth sharing principle, which has been advocated for a long time in the domain of data networks (see for instance the definition given by Bertsekas and Gallager [2]). The objective of max-min fairness is to maximize, under link transmission capacity constraints, the minimum bandwidth that a data flow can obtain when traversing a network. Specifically, in the domain of feasible bandwidth allocations, the rate λ_r allocated to the data flow r cannot be increased without decreasing the rate $\lambda_{r'}$ allocated to some other data flow r' such that $\lambda_{r'} < \lambda_r$.

Max-min fairness can be formulated in terms of an optimisation problem as follows: for every data flow r, there is at least one link l along the route of this data flow such that

$$\sum_{r' \in \mathcal{A}_l} \lambda_{r'} = C_l \quad \text{and} \quad \lambda_r = \max\{\lambda_{r'}, r' \in \mathcal{A}_l\}, \tag{1}$$

where C_l is the capacity of the link l and $\mathcal{A}_l$ is the set of data flows crossing the link l.

The problem of bandwidth sharing has recently been the subject of great interest, especially in connection with the problem of bandwidth allocation to elastic flows in the Internet (see for instance [6]). Those flows represent today the prevalent part of Internet traffic carried over telecommunications networks and understanding how bandwidth is shared between different flows crossing a link of the network is of crucial interest, for both users and network operators. Users appreciate the bandwidth sharing principles used within the network through the throughput of their applications, while network operators are interested in optimising the use of the bandwidth of their transmission links. Fairness is thus a design criterion for data networks. Fairness can also be used to introduce, for instance, tariff options.

Max-min fairness is not the only fairness principle, which has been introduced in the literature. Several fairness policies are discussed, for instance in [7, 6]. Some of these new fairness criteria are intended to remedy some weak points of max-min fairness, for instance by introducing the concept of weighted fairness and minimum bandwidth per flow. However, max-min fairness is the fundamental principle, which has been used in the design of popular data transmission protocols

Partially supported by the contract 00-1B320 with France Telecom and the IST Programme of the EU under contract number IST-1999-14186 (ALCOM-FT).

(for instance TCP). More recently, max-min fairness has been advocated in the design of the available bit rate service by the ATM Forum (see [1] for a survey of bandwidth sharing algorithms introduced in the context of the available bit rate service). Max-min fairness thus still appears as a fundamental design criterion, even though recent studies (for instance [4]) has questioned the appropriateness of max-min fairness for sharing bandwidth among elastic flows.

In spite of the fact that max-min fairness has existed for a long time in the literature on data networks, mathematical modeling of the behavior of max-min fairness in a network is not very frequently investigated. In general, max-min fairness is studied via extensive simulation in complex network configurations. In this paper, we study max-min fairness from an analytical point of view. For this purpose, we consider a simplistic network configuration, namely a single link network. Data flows appear on this link as a Poisson process with intensity u and have exponentially distributed lengths.

While the algorithmic definition of the max-min fair share is clear, any analytical computation rapidly turns out to be intractable. This is why we develop some heuristic approximations, which are quite fair when considering the max-min fair share, but which are less accurate for estimating the amount of unused bandwidth. This latter quantity corresponds to the proportion of the link bandwidth, which is not used because of the bandwidth sharing principle.

The organization of this paper is as follows: In Section 2, the linear model network model is presented and some basic results for the $M/M/\infty$ system are recalled. In Section 3, different approximations of the max-min fair share are developed. These approximations are then used to estimate the amount of lost bandwidth. Some concluding remarks are presented in Section 4.

2. Network model and preliminary results

We consider a linear network (for instance a bus with infinite length), which is represented by the real half line $[0, \infty)$. The transmission capacity of this single link network is taken as unity. We assume that data flows randomly appear on the network. A data flow is then characterized by an originating point and a length. We assume that the originating points of data flows are distributed according to a Poisson process with intensity u and that the length of data flows is exponentially distributed with unit mean.

In this model, several data flows may share the bandwidth of the transmission link, giving rise to clumps (see Figure 1). By definition, a clump is a set of data flows, which share a common portion of the transmission link.

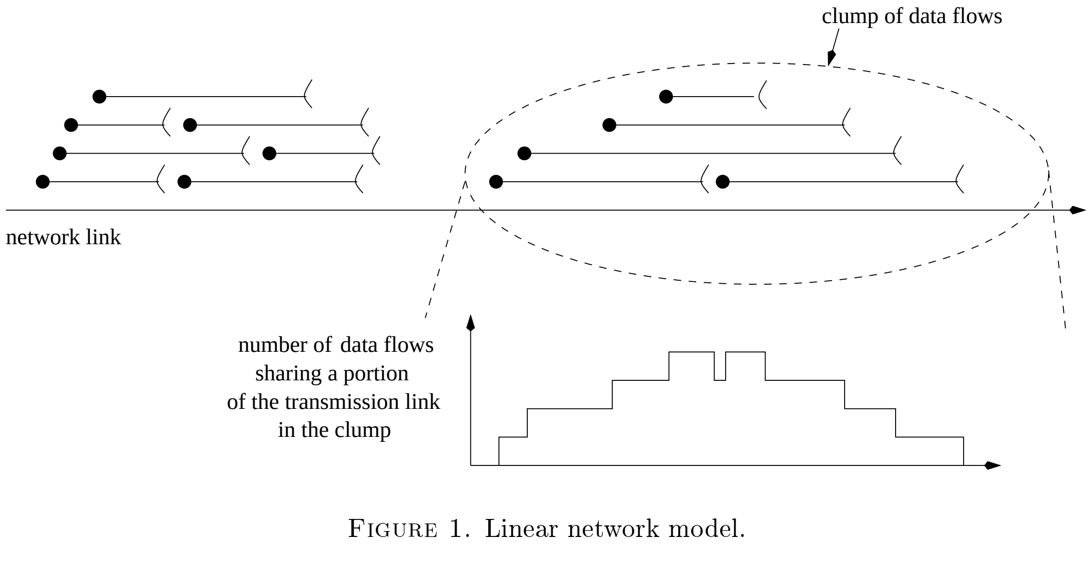

Figure 1. Linear network model.

We assume that in a clump, the bandwidth of the transmission link is shared among the different data flows according to max-min fairness. Moreover, we suppose that the rate adaptation of the data flows is immediate. There are thus no time aspects in the model considered. In other words, we consider a configuration at a given time and we study the allocation of bandwidth to each data flow.

The difficulty in max-min fairness consists of computing the bandwidth allocated to a data flow. While the algorithmic definition of max-min fairness is simple and quite intuitive, the analytical derivation of the exact max-min fair share rapidly proves intractable, even in extremely simplistic configurations as in the above linear network. The goal of this paper is to develop approximations of the max-min fair share as well as the amount of unused bandwidth in a clump. This last quantity may be interpreted as a performance measure for the network because it corresponds to the amount of the network bandwidth, which is not used because of the bandwidth allocation principle adopted for sharing the bandwidth among the different active connections.

For this purpose, we note that clumps are distributed as busy periods in an $M/M/\infty$ queueing system with input rate u and unit service rate. Before proceeding with some analytical results on the above linear network model, let us recall some basic elements about the $M/M/\infty$ queue.

The Charlier polynomials form the fundamental orthogonal polynomial system associated with the $M/M/\infty$ with input rate u and unit service rate. Charlier polynomials $\{Q_m(x;u)\}$, $m \geq 0$, are defined by the recursion: $Q_{-1}(x,u) \equiv 0$, $Q_0(x,u) \equiv 1$, and for $m \geq 0$,

$$uQ_{m+1}(x,u) + (x-m-u)Q_m(x,u) + mQ_{m-1}(x,u) = 0. \tag{2}$$

The generating function of Charlier polynomials is given by

$$\sum_{j=0}^{\infty} \frac{u^j}{j!} Q_j(x,u) z^j = e^{uz}(1-z)^x. \tag{3}$$

An easy consequence of the above relation is that Charlier polynomials satisfy

$$\sum_{j=0}^{k} \frac{u^j}{j!} Q_j(x,u) = \frac{u^k}{k!} Q_k(x-1,u). \tag{4}$$

For numerical purposes, we note that $Q_0(x,u) \equiv 1$ and that the above relation implies that

$$Q_k(-1,u) = \frac{k!}{u^k} \sum_{k=0}^{k} \frac{u^j}{j!}. \tag{5}$$

From the general relation

$$Q'_n(0,u) = -\sum_{k=0}^{n-1} \frac{k!}{u^{k+1}} \sum_{j=0}^{k} \frac{u^j}{j!}$$

we deduce that

$$Q'_n(0,u) \stackrel{\text{def}}{=} \frac{\partial}{\partial x} Q_n(x,u)\Big|_{x=0} = -\frac{1}{u} \sum_{k=0}^{n-1} Q_k(-1,u),$$

and hence, we have the recursion

$$Q'_n(0,u) = Q'_{n-1}(0,u) - \frac{1}{u} Q_{n-1}(-1,u). \tag{6}$$

From relation (4), we have

$$Q'_k(-1,u) - \frac{k}{u} Q'_{k-1}(-1,u) = Q'_k(0,u) \tag{7}$$

and

$$Q'_k(-2,u) - \frac{k}{u} Q'_{k-1}(-2,u) = Q'_k(-1,u). \tag{8}$$

The above relations allow us to recursively compute the quantities $Q'_k(0,u)$, $Q'_k(-1,u)$ and $Q'_k(-2,u)$, which will be of important use in the following.

For a given initial state $a \in \mathbb{N}$, let T_b denote the first time the occupation process $\{\Lambda_t\}$ of the $M/M/\infty$ queue hits the level $b \in \mathbb{N}$:

$$T_b = \inf\{t > 0 : \Lambda_t = b \mid \Lambda_0 = a\}. \tag{9}$$

When $a \leq b$, the Laplace transform of the random variable T_b is given by Takàcs formula [5] (see also [3]):

$$\widetilde{T}_b(s) \stackrel{\text{def}}{=} \mathbb{E}_a\left[e^{-sT_b}\right] = \frac{Q_a(-s,u)}{Q_b(-s,u)}, \tag{10}$$

where $Q_n(x,u)$ is the nth Charlier polynomial.

3. Analysis of max-min fairness for a linear network

3.1. Approximation of the max-min fair share. The determination of the max-min fair share of one data flow requires the computation of the number of flows simultaneously sharing a portion of the link with the date flow considered. Such a computation requires the transient analysis of the number of active data flows at a given point of the link. Once a clump is clearly identified, a complete recursive procedure is required to determine the max-min fair share of a data flow. This nested procedure makes any mathematical analysis extremely difficult. Nevertheless, a simple lower bound for the max-min fair share can easily be obtained by computing the number ν of data flows, which have a portion of the network in common with the data flow under consideration. Once ν is computed, $1/(\nu+1)$ is a lower bound for the max-min fair share. The random variable ν is characterized as follows.

Proposition 1. *The number ν of data flows, which share a portion of the transmission link with a data flow with an arbitrary originating point on the half line $[0,\infty)$, has the generating function $\nu(z)$ given by*

$$|z| \leq 1, \quad \nu(z) = \frac{e^{-u(1-z)}}{1+u(1-z)}. \tag{11}$$

and mean value

$$\mathbb{E}[\nu] = 2u. \tag{12}$$

Proof. Consider a data flow with an arbitrary originating point. The number ν_1 of data flows with originating points before the point considered and sharing a portion of the link with the data flow considered is distributed as a Poisson random variable with parameter u. Thus,

$$\mathbb{P}\{\nu_1 = k\} = \frac{u^k}{k!}e^{-u}. \tag{13}$$

Knowing the length d of the data flow considered, say, $d=\sigma$ for $\sigma>0$, the number ν_2 of data flows, which share a portion of the link with the data flow is distributed as Poisson random variable with parameter $u\sigma$

$$\mathbb{P}\{\nu_2 = k \mid d=\sigma\} = \frac{(u\sigma)^k}{k!}e^{-u\sigma}. \tag{14}$$

and hence,

$$\mathbb{P}\{\nu_2 = k\} = \int_0^\infty \frac{(u\sigma)^k}{k!}e^{-u\sigma}e^{-\sigma}\,d\sigma = \frac{u^k}{(u+1)^{k+1}}. \tag{15}$$

Since the random variables ν_1 and ν_2 are independent of each other, the result follows. This completes the proof. □

Corollary 1. *The mean value of the rate $\overline{\lambda}$ obtained by an arbitrary data flow under max-min fairness is such that*

$$\overline{\lambda} \geq \mathbb{E}[b_{\min}] = \mathbb{E}\left[\frac{1}{\nu+1}\right] = \frac{e}{u}\left[E_1(1) - E_1(1+u)\right], \tag{16}$$

where $E_1(z)$ is the exponential integral defined by

$$E_1(z) = \int_z^\infty \frac{e^{-t}}{t}\,dt, \quad |\arg z| < \pi. \tag{17}$$

Proof. We have

$$\mathbb{E}\left[\frac{1}{\nu+1}\right] = \int_0^1 \mathbb{E}[z^\nu]\, dz = \int_0^1 \frac{e^{-ut}}{1+ut}\, dt = \frac{e}{u}\left[E_1(1) - E_1(1+u)\right].$$

This completes the proof. □

As it will be shown by simulation results presented in the following, the lower bound given by $1/(\nu+1)$ may be too drastic. Another lower bound for the max-min fair share of a data flow is obtained when computing the maximum number κ of data flows sharing a portion of the transmission link with the data flow considered. The lower bound $1/(\kappa+1)$ is exact, when the data flow belongs to the group of data flows realizing the maximum number of simultaneous data flows in a clump. This lower bound is not exact, when in the group of data flows simultaneously active with the data flow considered, there is a data flow sharing a portion with another group of data flows such that the maximum number of simultaneous data flows is greater than the maximum number of simultaneous data flows in the former group. This phenomenon is illustrated in Figure 2.

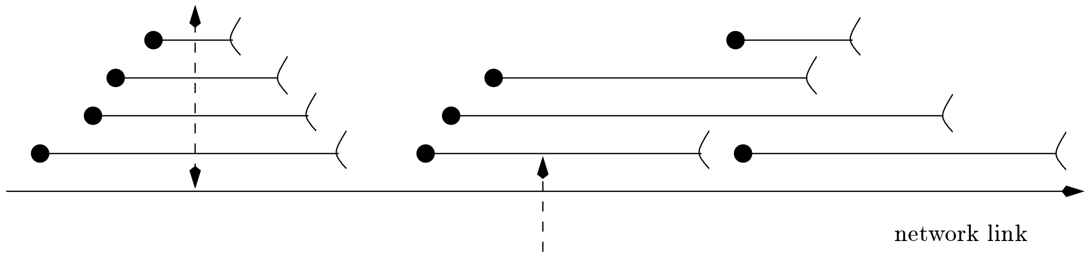

FIGURE 2. Comparison between the max-min fair share and $1/(\kappa+1)$.

Proposition 2. *The maximum number of data flows simultaneously sharing a portion of the transmission link with an arbitrary data flow is κ characterized by be generating function*

$$|z| \le 1, \quad \kappa(z) = e^{-u(1-z)} - (1-z)\mathcal{Q}(z,u), \tag{18}$$

where $\mathcal{Q}(z,u)$ is defined by

$$\mathcal{Q}(z,u) = e^{-u}\sum_{h=0}^{\infty} \frac{z^h u^h}{h!}\frac{Q_h(-2,u)}{Q_{h+1}(-1,u)}. \tag{19}$$

The mean value of the random variable κ is

$$\mathbb{E}[\kappa] = u - e^{-u}\mathcal{Q}(1,u) = u - e^{-u}\sum_{h=0}^{\infty}\frac{u^h}{h!}\frac{Q_h(-2,u)}{Q_{h+1}(-1,u)}. \tag{20}$$

Proof. Consider a given data flow and let n_0 be the number of data flows sharing the same portion of the transmission link with the data flow considered at its originating point. Then, by denoting by σ the length of the data flow considered,

$$\mathbb{P}\{\kappa < h \mid n_0 = k\} = \mathbb{P}_k\{T_h > \sigma\} = \int_0^\infty \mathbb{P}_k\{T_h > s\}e^{-s}\, ds = 1 - \mathbb{E}_k\left[e^{-T_h}\right] = 1 - \frac{Q_k(-1,u)}{Q_h(-1,u)}.$$

Since n_0 has a Poisson distribution with parameter u, we have

$$\mathbb{P}\{\kappa < h\} = e^{-u}\sum_{k=0}^{h-1}\frac{u^k}{k!}\left(1 - \frac{Q_k(-1,u)}{Q_h(-1,u)}\right),$$

and by using relation (4),

$$\mathbb{P}\{\kappa < h\} = e^{-u}\left(\sum_{k=0}^{h-1}\frac{u^k}{k!} - \frac{u^{h-1}}{(h-1)!}\frac{Q_{h-1}(-2,u)}{Q_h(-1,u)}\right)$$

and hence,

$$\mathbb{P}\{\kappa = h\} = e^{-u}\left(\frac{u^h}{h!} - \frac{u^h}{h!}\frac{Q_h(-2,u)}{Q_{h+1}(-1,u)} + \frac{u^{h-1}}{(h-1)!}\frac{Q_{h-1}(-2,u)}{Q_h(-1,u)}\right)$$

and equation (18) follows. The mean value is obtained by taking the derivative at point $z = 1$. This completes the proof. □

Corollary 2. *The mean value of the rate $\overline{\lambda}$ obtained by an arbitrary data flow under max-min fairness is such that*

$$\overline{\lambda} \geq \mathbb{E}\left[\frac{1}{\kappa+1}\right] = \frac{1-e^{-u}}{u} - e^{-u}\sum_{h=0}^{\infty}\frac{u^h}{(h+2)!}\frac{Q_h(-2,u)}{Q_{h+1}(-1,u)}, \tag{21}$$

where $Q_h(x,u)$ is the hth Charlier polynomial.

Remark. When the load u of the system is very small ($u \ll 1$), then

$$\mathbb{E}\left[\frac{1}{\nu+1}\right] \sim \mathbb{E}\left[\frac{1}{\kappa+1}\right] \sim 1-u$$

and thus, the approximations (16) and (21) lead to the same asymptotic behaviour.

The two approximations (16) and (21) are illustrated in Figure 3 for different values of the load u. The exact value of max-min fair share has been computed by simulation. It turns out that for small values of the load u (namely $u \leq 1$), these two approximations are quite accurate. The accuracy rapidly degrades when the load u increases but they still give a good idea for the value of the max-min fair share. Moreover, approximation (21) is a little bit better than approximation (16). For $u \sim 5$, the relative error is about 93% for approximation (16) and 22% for approximation (21). It thus appears that the best approximation is the latter one, whose computation is however much more complicated than that of approximation (16).

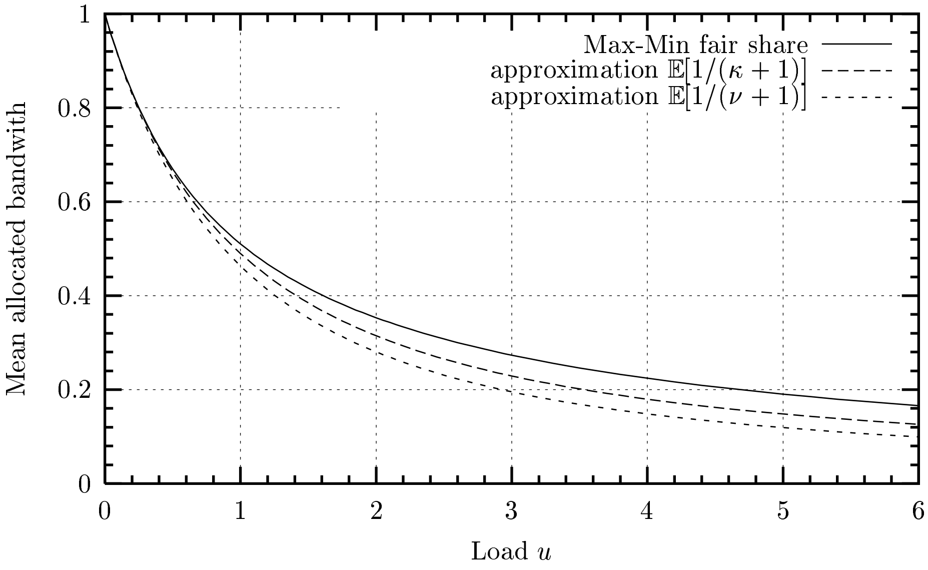

FIGURE 3. Mean value of the max-min fair share for the linear network model.

3.2. **Amount of lost bandwidth.** We define the lost bandwidth as the long range ratio of bandwidth of the link, which is not used in a clump due to the max-min bandwidth sharing principle. Let o_n and f_n denote the originating and the terminating points of the nth data flow. Then, the proportion of bandwidth, which is not used up to point t because of max-min fairness, is mathematically defined by

$$B_t = \frac{\int_0^t \mathbb{1}_{\{n_s>0\}}\, ds - \int_0^t \sum_{n=1}^{\infty} \lambda_n \mathbb{1}_{\{o_n \leq s \leq f_n\}}\, ds}{\int_0^t \mathbb{1}_{\{n_s>0\}}\, ds}, \tag{22}$$

where n_s is the number of active data flows at point s and λ_n is the rate allocated to the n-th data flow. Of course, the random variables $d_n = f_n - o_n$ and λ_n are correlated in an intricate way and any analysis of their interdependence seems to be out reach, when λ_n is computed by using the exact man-min fairness allocation.

When considering the long range average, it is easily checked that

$$\overline{B} \stackrel{\text{def}}{=} \lim_{t\to\infty} B_t = 1 - \frac{u}{1-e^{-u}}\mathbb{E}[d\lambda] \tag{23}$$

where d and λ are the length and the rate allocated to an arbitrary data flow, respectively.

$\overline{B}$ is clearly upper bounded by the value obtained when taking $\lambda = 1/(\nu+1)$ in definition (23), where ν is defined by Proposition 1. We specifically have the following result.

Proposition 3. *The proportion of lost bandwidth $\overline{B}$ is upper bounded as*

$$\overline{B} \leq \overline{B}_1 = -\frac{u}{(1+u)(e^u-1)} + \frac{e}{1-e^{-u}}\left[E_1(1) - E_1(1+u)\right], \tag{24}$$

where $E_1(z)$ is the exponential integral defined by equation (17).

Proof. With the notation of the proof of the Proposition 1, we have

$$\mathbb{E}[d\lambda] = \mathbb{E}\left[\mathbb{E}\left[\frac{d}{\nu+1} \mid d\right]\right]$$

It is easily checked that

$$\mathbb{E}[z^\nu \mid d=\sigma] = e^{-u(1+\sigma)(1-z)},$$

and then,

$$\mathbb{E}\left[\frac{1}{\nu+1} \;\middle|\; d=\sigma\right] = \int_0^1 e^{-u(1+\sigma)(1-z)}\, dz.$$

It follows that

$$\mathbb{E}[d\lambda] = \int_0^1 \frac{e^{-u(1-z)}}{[1+u(1-z)]^2}\, dz = \frac{1}{u}\left[1 - \frac{e^{-u}}{1+u}\right] - \frac{e}{u}\left[E_1(u) - E_1(1+u)\right].$$

Equation (24) is obtained by using relation (23). This completes the proof. □

Now, we consider the second approximation of the max-min fair share, which is given by equation (21).

Proposition 4. *The mean value of the amount of lost bandwidth $\overline{B}$ is upper bounded as*

$$\overline{B} \leq \overline{B}_2 = -\frac{u}{e^u-1}\frac{\partial}{\partial s}\left[\frac{1}{s}\sum_{h=0}^{\infty}\frac{u^h}{(h+2)!}\frac{Q_h(-s-1,u)}{Q_{h+1}(-s,u)}\right]\Bigg|_{s=1} = \frac{u}{e^u-1}\overline{b}_2, \tag{25}$$

where $\overline{b}_2$ can be written as

$$\overline{b}_2 = \sum_{h=0}^{\infty}\frac{u^h}{(h+2)!}\frac{Q_h(-2,u)}{Q_{h+1}(-1,u)} + \sum_{h=0}^{\infty}\frac{u^h}{(h+2)!}\frac{Q_h'(-2,u)}{Q_{h+1}(-1,u)} - \sum_{h=0}^{\infty}\frac{u^h}{(h+2)!}\frac{Q_h(-2,u)Q_{h+1}'(-1,u)}{Q_{h+1}^2(-1,u)}, \tag{26}$$

with $Q_h(s,u)$ being the hth Charlier polynomials and

$$Q_n'(x,u) = \frac{\partial}{\partial x} Q_n(x,u).$$

Proof. By using the same arguments as in the proof of Proposition 2, we have by conditioning on the length of the data flow considered,

$$\mathbb{E}\left[\frac{1}{\kappa+1} \;\middle|\; d=\sigma\right] = e^{-u}\left(\sum_{h=0}^{\infty}\frac{1}{h+1}\sum_{k=0}^{h}\frac{u^k}{k!}\mathbb{P}_k\{T_{h+1}>\sigma\} - \sum_{h=0}^{\infty}\frac{1}{h+1}\sum_{k=0}^{h-1}\frac{u^k}{k!}\mathbb{P}_k\{T_h>\sigma\}\right)$$

It follows that

$$\mathbb{E}[\lambda d] = e^{-u}\sum_{h=0}^{\infty}\frac{1}{h+1}\left(\mathcal{Q}_h - \mathcal{Q}_{h-1}\right),$$

where

$$\mathcal{Q}_h = \sum_{k=0}^{h}\frac{u^k}{k!}\int_0^\infty \sigma\mathbb{P}_k\{T_{h+1}>\sigma\}e^{-\sigma}\,d\sigma.$$

Noting that

$$\int_0^\infty \mathbb{P}_k\{T_h>\sigma\}e^{-s\sigma}\,d\sigma = \frac{1}{s}\left[1-\frac{Q_k(-s,u)}{Q_h(-s,u)}\right],$$

we have

$$\int_0^\infty \mathbb{P}_k\{T_h>\sigma\}\sigma e^{-\sigma}\,d\sigma = -\left.\frac{\partial}{\partial s}\frac{1}{s}\left[1-\frac{Q_k(-s,u)}{Q_h(-s,u)}\right]\right|_{s=1}$$

and then, by using relation (4),

$$\mathbb{E}[\lambda d] = \frac{1}{u}(1-e^{-u}) + e^{-u}\left.\frac{\partial}{\partial s}\left[\frac{1}{s}\sum_{h=0}^{\infty}\frac{u^h}{(h+2)!}\frac{Q_h(-s-1,u)}{Q_{h+1}(-s,u)}\right]\right|_{s=1}$$

Relation (25) is obtained by using definition (23). This completes the proof. □

Remark. When the load of the system is very small ($u \ll 1$), then

$$\overline{B}_1 \sim u$$

Moreover, when $u \ll 1$,

$$\frac{1}{s}\sum_{h=0}^{\infty}\frac{u^h}{(h+2)!}\frac{Q_h(-s-1,u)}{Q_{h+1}(-s,u)} \sim \frac{u}{2s(s+u)}$$

and then, $\overline{B}_2 \sim u$. Hence, the two approximations lead to the same asymptotic behavior when $u \ll 1$.

When $u \to \infty$, we see that $\overline{B}_1 \to e \times E_1(1) \sim 60\%$ while $\overline{B}_2$ decreases. When considering the exact max-min fairness allocation, we may reasonably conjecture that the amount of lost bandwidth should tend to 0 as $u \to \infty$. Indeed, when $u \to \infty$, the system is heavily loaded and the superposition of all data flows exhibits a smooth behavior (central limit theorem). Hence, the amount of lost bandwidth should be small.

It follows that approximation $\overline{B}_1$ does not follow the general behavior of the exact amount of lost bandwidth under max-min fairness. Approximation $\overline{B}_2$ exhibits the correct behavior and yields a quite accurate estimate for the amount of lost bandwidth. For instance, when u is in the neighborhood of 6, the relative error is about 25% as illustrated in Figure 4. Approximation $\overline{B}_2$ is numerically computed by using the recursions (7) and (8) between the derivatives of Charlier polynomials.

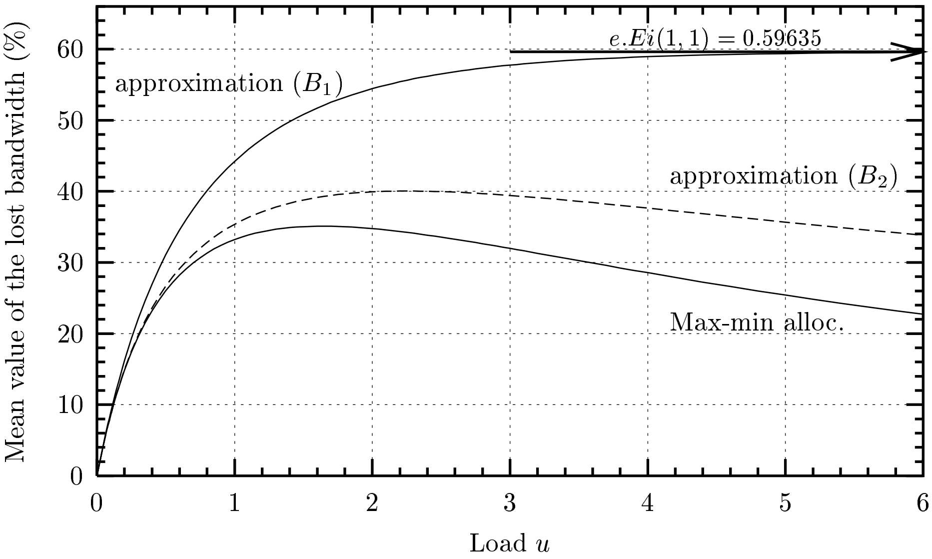

FIGURE 4. Mean value of the bandwidth lost in a clump due to max-min fairness.

4. CONCLUSION

Two approximations have been proposed in this paper to estimate the fair share as well as the amount of lost bandwidth when bandwidth is allocated to elastic flows according to max-min fairness in a single link linear network . It turns out that the approximation based upon the maximum number of flows sharing the transmission capacity of the link (equations (21) and (25)) yields the best results for both the max-min fair share and the amount of lost bandwidth. This opens the door to analytically modeling max-min fairness in more complex situations, for instance when the network is a graph. This will be investigated in further studies.

Finally, more accurate estimates for the max-min fair share and the amount of lost bandwidth could certainly be obtained when considering heavy traffic approximations (specifically when u becomes very large). The development of approximations under heavy traffic conditions is however made difficult by the complexity of max-min fairness, in particular by the algorithmic definition of the fair share. This point will be investigated in forthcoming papers.

REFERENCES

[1] A. Arulambalam and X.Q. Chen. Allocating fair rates for available bit rate service in ATM networks. *IEEE Communications Magazine.*, 92-100, 1996.
[2] D. Bertsekas and R. Gallager. *Data networks.* Prentice Hall, 1987.
[3] F. Guillemin and A. Simonian. Transient characteristics of an $M/M/\infty$ system. *Adv. Appl. Prob.*, 27:862–888, September 1995.
[4] F. 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.
[5] Ph. Robert. *Réseaux et files d'attente : méthodes probabilistes.* Mathématiques et Applications 35, Springer Verlag 2000.
[6] J. Roberts and L. Massoulié. Bandwidth sharing and admission control for elastic traffic. In *Proc. Infocom'99.*
[7] J. Ros, W.K. Tsai and Y. Kim. A theory of network optimization In *Proc. Globecom'99.*

A. Dupuis, France Telecom R&D, 2, Avenue Pierre Marzin, 22300 Lannion (France)
E-mail address: `Alain.Dupuis@rd.francetelecom.fr`

F. Guillemin, France Telecom R&D, 2, Avenue Pierre Marzin, 22300 Lannion (France)
E-mail address: `Fabrice.Guillemin@rd.francetelecom.fr`

P. Robert, Algorithms Project, INRIA, F-78150 Rocquencourt (France)
E-mail address: `Philippe.Robert@inria.fr`