A Bandwidth Sharing Theory for a Large Number of HTTP-like Connections

Cheng-Shang Chang, Zhen Liu

Abstract— **There has been tremendous progress in understanding how bandwidth is shared by TCP-like connections. By associating each TCP-like connection with a utility function, the bandwidth sharing problem of TCP-like connections can be modelled as a distributed optimization problem for utility functions. However, little is known on how bandwidth is shared by HTTP-like connections through their utility functions at the TCP level. One of the main objectives of this paper is to provide a theory for bandwidth sharing of a large number of HTTP-like connections. Based on certain technical assumptions, we show that there is a utility function at the HTTP level for an HTTP-like connection and such a utility function can be derived from the utility function at the TCP level. The bandwidth is then shared by HTTP-like connections through utility functions at the HTTP level. Moreover, there is a probabilistic interpretation for how the utility function at the HTTP level is related to the utility function at the TCP level. This is done by relating utility functions to large deviation rate functions.**

I. Introduction

Inspired by the seminal work by Kelly [21], many researchers (see e.g., [22], [25], [24], [23], [26], [7] and references therein) have addressed the bandwidth sharing problem of TCP-like connections as a distributed optimization problem for utility functions. In particular, Low [24] identified the corresponding utility functions used in TCP Reno [17], TCP Reno with RED [13], and TCP Vegas [8]. Massoulié and Roberts [25] also characterized various utility functions that achieve fairness, minimize delay, or maximize throughput. For these papers in the literature, it is generally assumed that there is a utility function associated with a TCP-like connection. If the capacity of the link is c, then the bandwidth received by the j^{th} TCP-like connection, denoted by x_j, is the solution of the following maximization problem:

$$\begin{aligned} &\max \sum_j U_j(x_j) \\ &\text{s.t.} \quad \sum_j x_j \le c, \end{aligned} \tag{1}$$

where U_j is the utility function of the j^{th} connection, which is in general increasing and concave. Though researchers have made tremendous progress in understanding how bandwidth is shared by TCP-like connections, little is known on how bandwidth is shared by HTTP-like connections through their utility functions at the TCP level.

One of the main objectives of this paper is to provide a theory for bandwidth sharing of HTTP-like connections. By an HTTP-like connection, we mean a sequence of ON-OFF periods, where the ON periods correspond to TCP-like transfers and the OFF periods correspond to think times.

Cheng-Shang Chang is with Department of Electrical Engineering and Institute of Communications Engineering, National Tsing Hua University, Hsinchu 300, Taiwan, R.O.C. Email: cschang@ee.nthu.edu.tw . This research is supported in part by IBM Faculty Partnership Award, 2001, and the National Science Council, Taiwan, R.O.C., under Contract NSC-90-2213-E007-034.

Zhen Liu is with IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598. Email: zhenl@us.ibm.com .

The study of such HTTP-like model is motivated by the fact that HTTP traffic is now the dominant one in the current Internet. Most of these HTTP sessions use TCP. Many studies have identified important Web traffic characteristics, see [28] for an overview. Realistic Web traffic models have been proposed such as SURGE [3], [4] and WAGON [18]. In all these models, the traffic is represented by sequence of TCP transfers separated by idle times which correspond to think times of physical users (as in WAGON) or of user equivalent (as in SURGE). In [2], such HTTP-like model is used for the performance evaluation of access networks.

We consider a large number of HTTP-like connections and assume that there are some laws of large numbers (Assumptions (A1-3) in Section II). Based on these assumptions, we obtain the following results:

(i) There is a utility function at the HTTP level for an HTTP-like connection and such a utility function can be derived from the utility function at the TCP level. The bandwidth is then shared by HTTP-like connections through utility functions at the HTTP level.

(ii) There is a probabilistic interpretation for how the utility function at the HTTP level is related to the utility function at the TCP level. This is done by relating utility functions to large deviation rate functions.

The paper is organized as follows: in Section II, we describe our model for HTTP-like connections. In Section III, we list our assumptions and report our main results for bandwidth sharing of HTTP-like connections. We provide a new model for bandwidth sharing of TCP-like connections based on large deviation rate functions in Section IV. This gives a probabilistic interpretation for how the utility function at the HTTP level is related to the utility function at the TCP level. In Section V, we consider a processor sharing system as a special case to verify our assumptions in Section III. In Section VI, we discuss possible extensions of the theory and future research topics along this line.

II. Model

As in [2], we consider N HTTP-like connections that shares a common link with capacity Nc (bits). All these connections alternate between *think times* and *transfer times* (see Figure 1). During the period of a think time, a connection does not require any bandwidth from the link. At the end of the period of a think time, a connection requests transmitting a random amount of data. This usually evokes a TCP-like connection. The transfer time is defined to be the time for the TCP-like connection to transmit the requested data. We assume there are J classes of connections. The number of class j connections is $N\alpha_j$ (with $\sum_{j=1}^J \alpha_j = 1$). The average period of a think time for class

j connections is $\bar{I}_j$ and the average amount of data for a class j request is $\bar{B}_j$, $j = 1, 2, \ldots, J$. Connections that are in the phase of transfer times, i.e., TCP-like connections, participate bandwidth sharing as in (1).

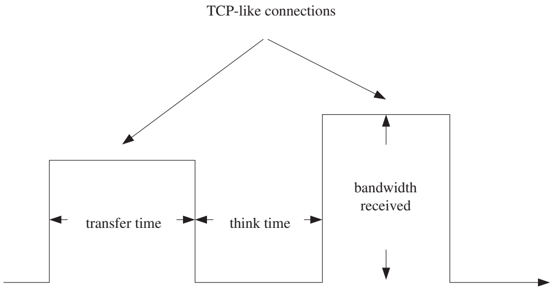

Fig. 1. A HTTP-like connection

To be precise, let $Q_j(t)$ be the number of class j connections that are in the phase of transfer times at time t, and $X_j(t)$ be the amount of bandwidth received by a class j connection that is in the phase of transfer times at time t. We assume the link is *fully utilized* and $X_j(t)$, $j = 1, \ldots, J$, are allocated according to the following optimization problem:

$$\begin{aligned} &\max \sum_{j=1}^{J} Q_j(t) U_j(X_j(t)) \\ &\text{s.t.} \quad \sum_{j=1}^{J} Q_j(t) X_j(t) = Nc. \end{aligned} \qquad (2)$$

The functions $U_j(\cdot)$'s are called the utility functions at the TCP level in this paper. Throughout the paper, we assume the utility functions $U_j(x)$, $j = 1, 2, \ldots, J$, are *strictly concave* and *twice differentiable* for all $x > 0$. Such an assumption ensures there is a unique positive optimum for the maximization problem. Note that our assumption is slightly stronger than the usual *continuously differentiable* condition. The reason for this is that it is easier for us to verify concavity for the utility functions at the HTTP level by differentiation.

Our assumption on full utilization of the link makes this link a congested one. As such, there is no need to assume that the utility functions are increasing, which is a general assumption to encourage connections to consume bandwidth as much as possible. Also, for the link to be congested, there should be enough traffic generated by the connections. The maximum amount of traffic that can be generated is the case that requests are transmitted in zero time. In that case, the maximum rate of the traffic generated by a class j connection is

$$\rho_j = \bar{B}_j / \bar{I}_j. \qquad (3)$$

Thus, in order for the link to be congested, we need

$$\sum_{j=1}^{J} \alpha_j \rho_j > c, \qquad (4)$$

and this is assumed to be true throughout the paper.

III. Bandwidth sharing of HTTP-like connections

To analyze the bandwidth sharing problem, we make the following assumptions:

(A1) Ensemble ergodicity: for some finite constant t_0,

$$\lim_{N\to\infty} \sup_{0 \le t \le t_0} \left| \frac{Q_j(t)}{N\alpha_j} - \pi_j \right| = 0, \quad a.s., j = 1, 2, \ldots, J, \qquad (5)$$

for some $\pi_j > 0$.

(A2) Time ergodicity: let T_j and $\bar{T}_j$ be the transfer time and the average transfer time for a class j connection. Assume that

$$\lim_{N\to\infty} \bar{T}_j = \mathrm{E} \lim_{N\to\infty} T_j, \qquad (6)$$

and that

$$\lim_{N\to\infty} \frac{\bar{T}_j}{\bar{I}_j + \bar{T}_j} = \pi_j, \quad j = 1, \ldots, J. \qquad (7)$$

(A3) Independence: the random amounts of data requested by connections are independent of the system and they are bounded by t_0.

The physical interpretation of π_j is the probability that a class j connection is in the phase of transfer times at a random time (when $N \to \infty$). Assumption (A1) can then be interpreted as follows: if one takes the ensemble average of the number of class j connections that are in the phase of transfer times, it should be equal to the probability that a class j connection is in the phase of transfer times at a random time. Equation (7) is known as the cycle formula and it holds for most Markovian systems. Assumption (A2) says that the probability that a class j connection is in the phase of transfer times at a random time should be equal to the ratio of the average transfer time to the sum of the average think time and the average transfer time. The condition in (6) is a technical one which ensures that the limit and the expectation can be interchanged. In view of these, the first two assumptions are natural ones as we are assuming that the statistics gathered by observing a single connection for a long period of time are the same as those from observing many connections at a random instant of time. The third assumption is also a technical one that is needed in the proof of our main results in Theorem 1. In this paper, we do not provide a general condition on the traffic model for the assumptions to hold, but we will provide partial verification by a special case that uses the processor sharing policy in Section V. We also note that the assumptions may not hold in certain circumstances. For example, it is reported in [14] that there might be impatient users who interrupt TCP connections before they are completed.

Let

$$Y_j(t) = \frac{Q_j(t)}{N\alpha_j} X_j(t) \qquad (8)$$

be the (ensemble) average bandwidth received by class j connections at time t.

Theorem 1 *Assume that (A1-3) hold. For $j = 1, 2, \ldots, J$,*

$$\lim_{N\to\infty} Y_j(t) = y_j^*, \quad a.s. \qquad (9)$$

and y_j^*'s are the unique solution of the following concave programming problem:

$$\max \sum_{j=1}^{J} \alpha_j \hat{U}_j(y_j)$$
$$\text{s.t.} \quad \sum_{j=1}^{J} \alpha_j y_j = c, \tag{10}$$

where

$$\hat{U}_j(y_j) = \frac{\rho_j - y_j}{\rho_j} U_j\left(\frac{y_j \rho_j}{\rho_j - y_j}\right) \tag{11}$$

is strictly concave and twice differentiable for $0 < y_j < \rho_j$. Furthermore,

$$\lim_{N\to\infty} X_j(t) = x_j^*, \quad a.s. \tag{12}$$

and

$$x_j^* = \frac{\rho_j y_j^*}{\rho_j - y_j^*}, \quad j = 1, 2, \ldots, J. \tag{13}$$

The proof of Theorem 1 is deferred to the end of this section. The functions $\hat{U}_j(y_j)$'s can be viewed as the utility functions at the HTTP level. The concave programming problem in (10) shows that the bandwidth sharing for HTTP-like connections is in a similar form to the bandwidth sharing for TCP-like connections in (2). The difference is to replace the utility functions. Equation (11) shows how the utility functions at the HTTP level can be derived from the utility functions at the TCP level.

Also, one can view x_j^* as the TCP-like throughput and y_j^* as the HTTP-like throughput of a class j connection. Equation (13) shows how the TCP-like throughput x_j^* and the HTTP-like throughput y_j^* are related. As the maximum rate for the traffic generated by a class j connection is ρ_j, clearly the HTTP-like throughput y_j^* is bounded above ρ_j. To get the HTTP-like throughput y_j^* close to ρ_j, one needs a very large TCP-like throughput x_j^*.

To further illustrate how bandwidth is shared by HTTP-like connections, we consider the following example with two classes, i.e., $J = 2$. The utility function of the class j connections at the TCP level is

$$U_j(x_j) = d_j \log x_j, \quad j = 1, 2. \tag{14}$$

Such utility functions are used for TCP Vegas/DropTail in [24] and d_j is the round trip propagation delay of the class j connections. In this case, the utility function of the class j connections at the HTTP level is

$$\hat{U}_j(y_j) = d_j \frac{\rho_j - y_j}{\rho_j} \log\left(\frac{y_j \rho_j}{\rho_j - y_j}\right), \quad j = 1, 2. \tag{15}$$

In Figure 2, we plot the bandwidth received by class 1 HTTP-like connections as a function of $\log \rho_1$. In this figure, we choose $\alpha_1 = \alpha_2 = 0.5$, $c = 1.0$, $d_1 = 1.0$ and $d_2 = 2.0$. There are three curves in this figure, which correspond to the three cases (from the top to the bottom) (i) $\rho_2 = 0.5\rho_1$, (ii) $\rho_2 = \rho_1$ and (iii) $\rho_2 = 2\rho_1$. It is interesting to see that as the maximum rate ρ_1 increases, the bandwidth received by class 1 HTTP-like connections decreases from $c\frac{\rho_1}{\rho_1+\rho_2}$ to the limit $c\frac{d_1}{d_1+d_2} = 2/3$, which corresponds to the case with infinite duration TCP-like connections. These curves show that the round trip propagation delay d_j's only play an important role when the system is extremely overloaded ($\rho_1 \to \infty$). When the system is only slightly overloaded, i.e., $\alpha_1\rho_1 + \alpha_2\rho_2$ is slightly larger than c, the bandwidth is shared almost proportional to the load of each class.

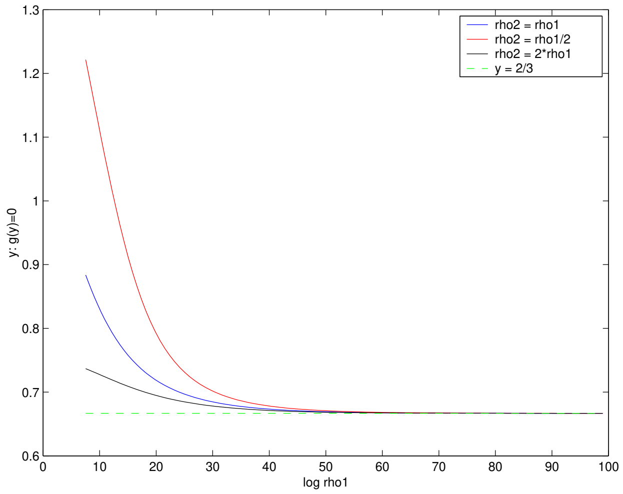

Fig. 2. Bandwidth received by class 1 HTTP-like connections

The rest of this section is devoted to the proof of Theorem 1. The proof of Theorem 1 requires results in Lemma 2 and Lemma 3 below. To see the intuition of these two lemmas, note that $X_j(t)$, $j = 1, 2, \ldots J$, are obtained from the following maximization problem:

$$\max \sum_{j=1}^{J} \frac{Q_j(t)}{N} U_j(X_j(t))$$
$$\text{s.t.} \quad \sum_{j=1}^{J} \frac{Q_j(t)}{N} X_j(t) = c. \tag{16}$$

In view of the limits in (5), it is expected that $X_j(t)$ converge uniformly over the interval $[0, t_0]$ to some constants x_j^*, $j = 1, 2, \ldots, J$, as $N \to \infty$. If the utility functions $U_j(\cdot)$ are smooth enough, one might expect that x_j^* are the optimal solution of the following maximization problem:

$$\max \sum_{j=1}^{J} \alpha_j \pi_j U_j(x_j)$$
$$\text{s.t.} \quad \sum_{j=1}^{J} \alpha_j \pi_j x_j = c. \tag{17}$$

This is stated formally in the following lemma and its proof is given in Appendix A.

Lemma 2 *Under (A1), we have for all $j = 1, \ldots, J$,*

$$\lim_{N\to\infty} \sup_{0\le t\le t_0} |X_j(t) - x_j^*| = 0, \quad a.s., \tag{18}$$

where x_j^*, $j = 1, 2 \ldots, J$, are the unique solution of the maximization problem in (17).

Since class j connections that are in the phase of transfer times receive a constant amount of bandwidth x_j^* (when N is large), the average transfer time of a class j connection $\bar{T}_j$ might be expressed as $\bar{B}_j/x_j^*$. This is shown in Lemma 3 and its proof is given in Appendix B.

Lemma 3 *Under (A2) and (A3), as $N \to \infty$, the average transfer time of a class j connection $\bar{T}_j$ converges to $\bar{B}_j/x_j^*$. As a direct consequence of (7),*

$$\pi_j = \frac{\bar{B}_j/x_j^*}{\bar{I}_j + \bar{B}_j/x_j^*} = \frac{\rho_j/x_j^*}{1+\rho_j/x_j^*}, \quad j = 1, \dots, J. \tag{19}$$

Note that π_j's in (19) are also functions of x_j^*'s. In conjunction with (17), we have a fixed point limit. Such a fixed point limit was previously reported in [2] for the single class case. However, in general, there might not exist a unique solution for such a fixed point limit. The key to the proof of Theorem 1 is to convert the fixed point limit problem into the concave programming problem in (10) so that the existence of a unique solution can be shown.

Now we prove Theorem 1.

Proof. (Theorem 1) We first show that there exists a unique positive solution of the concave programming problem in (10). From the Kuhn-Tucker theorem, an optimal solution to the optimization problem in (10) satisfies

$$\hat{U}_j'(y_j) = \lambda, \quad j = 1, 2, \dots, J, \tag{20}$$

and

$$\sum_{j=1}^{J} \alpha_j y_j = c. \tag{21}$$

Note from (11) that

$$\hat{U}_j''(y_j) = (\frac{\rho_j}{\rho_j - y_j})^3 U_j''(\frac{\rho_j y_j}{\rho_j - y_j}). \tag{22}$$

Since $U_j(x_j)$'s are strictly concave and twice differentiable for all $x_j > 0$, we have from (22) that $\hat{U}_j(y_j)$'s are also strictly concave and twice differentiable for $0 < y_j < \rho_j$. As such, $(\hat{U}_j')^{-1}(\lambda)$ is strictly decreasing and continuous, and it ranges between 0 and ρ_j. As we assume that $\sum_{j=1}^{J} \alpha_j \rho_j > c$ in (4), there is a unique λ_0 that satisfies

$$\sum_{j=1}^{J} \alpha_j (\hat{U}_j')^{-1}(\lambda) - c = 0. \tag{23}$$

Thus, the unique optimal solution for (10) is

$$y_j = (\hat{U}_j')^{-1}(\lambda_0), \quad j = 1, 2, \dots, J. \tag{24}$$

Now we convert the fixed point limit to the concave programming problem in (10). From (A1) and Lemma 2, it follows that for all $j = 1, 2, \dots, J$,

$$\begin{aligned} y_j^* &= \lim_{N\to\infty} Y_j(t) \\ &= \lim_{N\to\infty} \frac{Q_j(t)}{N\alpha_j} \lim_{N\to\infty} X_j(t) \\ &= \pi_j x_j^*, \quad a.s., \end{aligned} \tag{25}$$

where x_j^*, $j = 1, 2 \dots, J$, are the unique solution of the maximization problem in (17). Using (19), we have that

$$y_j^* = \frac{\rho_j x_j^*}{\rho_j + x_j^*}, \quad j = 1, 2, \dots, J. \tag{26}$$

Solving (26) yields (13). Replacing (13) in (19) yields

$$\pi_j = 1 - \frac{y_j^*}{\rho_j}, \quad j = 1, 2, \dots, J. \tag{27}$$

Using (13) and (27) in (17), we show that y_j^*'s are the unique solution of the following maximization problem:

$$\begin{aligned} &\max \sum_{j=1}^{J} \alpha_j \frac{\rho_j - y_j}{\rho_j} U_j(\frac{y_j \rho_j}{\rho_j - y_j}) \\ &\text{s.t.} \quad \sum_{j=1}^{J} \alpha_j y_j = c. \end{aligned} \tag{28}$$

Replacing (11) in (28) yields the maximization problem in (10). ■

IV. Modelling for bandwidth sharing of TCP-like connections with large deviation rate functions

In this section, we provide a new model for bandwidth sharing of TCP-like connections based on large deviation rate functions. The large deviation rate functions in the new model not only have a natural mapping to the utility functions, but also provide a probabilistic interpretation how the two utility function $U_j(\cdot)$ and $\hat{U}_j(\cdot)$ derived in the previous section are related.

Suppose that *when there is no constraint on the link capacity*, the amounts of bandwidth received by class j connections in the phase of transfer times are independent and identically (i.i.d.) random variables (r.v.'s) with a common moment generating function $\phi_j(\theta)$. This is the *only* randomness in this model (readers might think other randomness, such as the random amount of data requested, has been incorporated into the randomness of these i.i.d. r.v.'s). From (A1), there are (roughly) $N\alpha_j\pi_j$ class j connections that are in the phase of transfer times (TCP-like connections). Let X_{j,n_j}, $n_j = 1, \dots, N\alpha_j\pi_j$, $j = 1, 2, \dots, J$, be the amount of bandwidth received by the n_j^{th} class j connection when it is in the phase of transfer times. Let $S_j(m) = \sum_{n_j=1}^{m} X_{j,n_j}$ be the partial sum. As X_{j,n_j}'s are i.i.d., it follows from Crämer's theorem (see e.g., [12], [31]) that $\{(S_j(m)/m, m \geq 1\}$ satisfies the large deviation principle with the rate function $W_j(x_j)$, where

$$W_j(x_j) = \sup_{\theta}[\theta x_j - \log \phi_j(\theta)] \tag{29}$$

is the Legendre transform of $\log \phi_i(\theta)$. Roughly speaking, when m is large, we have that

$$\mathrm{P}(\frac{S_j(m)}{m} \approx x_j) \approx e^{-mW_j(x_j)}. \tag{30}$$

Since X_{j,n_j}'s are all independent,

$$\begin{aligned}
&\mathrm{P}(\frac{S_j(N\alpha_j\pi_j)}{N\alpha_j\pi_j} \approx x_j, j = 1, \dots J) \\
&= \prod_{j=1}^{N} \mathrm{P}(\frac{S_j(N\alpha_j\pi_j)}{N\alpha_j\pi_j} \approx x_j) \\
&\approx \exp\Big(-N \sum_{j=1}^{J} \alpha_j\pi_j W_j(x_j)\Big). \qquad (31)
\end{aligned}$$

Now there is a constraint on the total capacity of the link. Conditioning on the link is fully utilized, i.e., $\sum_{j=1}^{N} S_j(N\alpha_j\pi_j) = Nc$, we can find the most likely state by solving the following minimization problem:

$$\begin{aligned}
&\min \sum_{j=1}^{J} \alpha_j\pi_j W_j(x_j) \\
&\text{s.t.} \quad \sum_{j=1}^{J} \alpha_j\pi_j x_j = c. \qquad (32)
\end{aligned}$$

In view of the minimization problem in (32), there is a natural mapping between the large deviation rate function $W_j(x_j)$ and the utility function $U_j(x_j)$. One can simply let

$$W_j(x_j) = -U_j(x_j) \qquad (33)$$

to obtain the same maximization problem in (17). However, a large deviation rate function requires more than convexity. It should also be nonnegative. This implies that some constant should be added to make the large deviation rate function nonnegative.

To derive the optimization problem corresponding to (10), let $A_j(\tau) = \sup\{m : S_j(m) < \tau\}$ be the inverse mapping of S_j. In the setting of point processes, usually $A_j(\tau)$ is referred to as the cumulative number of arrivals by time τ and $S_j(m)$ is referred to as the arrival time of the m^{th} arrival. It is well known (see e.g., [15]) that if $\{S_j(m)/m, m \geq 1\}$ satisfies the large deviation principle with the rate function $W_j(\cdot)$, then $\{A_j(\tau)/\tau, \tau > 0\}$ also satisfies the large deviation principle with the rate function

$$W_j^*(z_j) = z_j W_j(\frac{1}{z_j}). \qquad (34)$$

An intuitive argument for this goes as follows: observe that the event $\{A_j(\tau) \leq m\}$ is equivalent to the event $\{S_j(m) > \tau\}$. When τ is large,

$$\begin{aligned}
&\mathrm{P}(\frac{A_j(\tau)}{\tau} \approx z_j) \approx \mathrm{P}(S_j(\tau z_j) \approx \tau) \\
&= \mathrm{P}(\frac{S_j(\tau z_j)}{\tau z_j} \approx \frac{1}{z_j}) \\
&\approx e^{-\tau z_j W_j(\frac{1}{z_j})} = e^{-\tau W_j^*(z_j)}.
\end{aligned}$$

Via the inverse mapping, $A_j(\tau)$ can also be interpreted as the total transfer time for a class j connection to complete transmission of τ units of data. Let

$$\hat{A}_j(\tau) = \tau/\rho_j + A_j(\tau), \quad j = 1, 2, \dots, J. \qquad (35)$$

We claim that $\hat{A}_j(\tau)$ is the time to transmit τ units of data for a class j connection. To see this, consider a particular class j connection for a long period of time such that τ units of class j data has been transmitted. During that period of time, the number of class j requests transmitted is (roughly) $\tau/\bar{B}_j$, and the time that the connection is in the phase of think times is (roughly) $\bar{I}_j\tau/\bar{B}_j = \tau/\rho_j$. Thus, $\hat{A}_j(\tau)$ is simply the sum of the total think time and the total transfer time of transmitting τ data units.

As the difference of $\hat{A}_j(\tau)/\tau$ and $A_j(\tau)/\tau$ is a constant, we have from (34) that $\{\hat{A}_j(\tau)/\tau, \tau > 0\}$ satisfies the large deviation principle with the rate function

$$W_j^{**}(z_j) = W_j^*(z_j - \frac{1}{\rho_j}) = (z_j - \frac{1}{\rho_j}) W_j(\frac{1}{z_j - \frac{1}{\rho_j}}). \qquad (36)$$

Let $\hat{S}_j(m) = \sup\{\tau : \hat{A}_j(\tau) < m\}$ be the inverse mapping of $\hat{A}(\cdot)$. Then $\hat{S}_j(m)$ can be interpreted as the amount of bandwidth consumed by m class j connections. From the inverse mapping formula and (36), we know that $\{(\hat{S}_j(m)/m, m \geq 1\}$ satisfies the large deviation principle with the rate function

$$\hat{W}_j(y_j) = y_j W_j^{**}(\frac{1}{y_j}) = \frac{\rho_j - y_j}{\rho_j} W_j(\frac{y_j\rho_j}{\rho_j - y_j}). \qquad (37)$$

Since there are $N\alpha_j$ class j connections (HTTP-like connections), $j = 1, \dots, J$, we have from the large deviation principle that

$$\begin{aligned}
&\mathrm{P}(\frac{\hat{S}_j(N\alpha_j)}{N\alpha_j} \approx y_j, j = 1, \dots, J) \\
&\approx \exp\Big(-N \sum_{j=1}^{J} \alpha_j \hat{W}_j(y_j)\Big). \qquad (38)
\end{aligned}$$

Conditioning on the link is fully utilized, i.e.,

$$\sum_{j=1}^{J} \hat{S}_j(N\alpha_j) = Nc,$$

we can find the most likely state by solving the following minimization problem:

$$\begin{aligned}
&\min \sum_{j=1}^{J} \alpha_j \hat{W}_j(y_j) \\
&\text{s.t.} \quad \sum_{j=1}^{J} \alpha_j y_j = c. \qquad (39)
\end{aligned}$$

The minimization problem in (39) and (37) provide a probabilistic interpretation why and how the utility function $\hat{U}_j(\cdot)$ at the HTTP level should be related to the utility function $U_j(\cdot)$ at the TCP level.

V. Processor sharing as a verifying special case

Our theory relies heavily on the assumptions on ergodicity in (A1) and (A2). The purpose of this section is to consider a special case that can verify those assumptions and the results derived based on those assumptions.

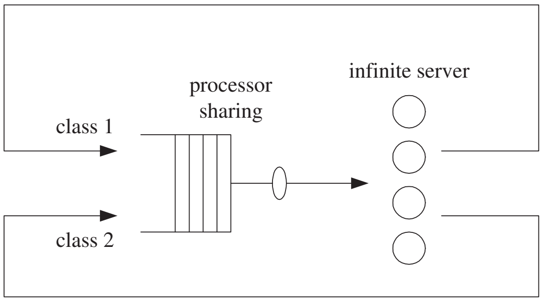

Fig. 3. The closed queueing networks with two classes of customers

For this, we consider a special case with two classes of connections, i.e., $J = 2$. We assume that think times and requests of class j connections are i.i.d. r.v.'s with means $\bar{I}_j$ and $\bar{B}_j$ respectively. Most importantly, we further assume that the two utility functions in (2) are the same, i.e., $U_1(x) = U_2(x)$ for all x. Then we have from the concave property of the utility functions that

$$\begin{aligned}
&\frac{Q_1(t)U_1(X_1(t)) + Q_2(t)U_2(X_2(t))}{Q_1(t) + Q_2(t)} \\
&= \frac{Q_1(t)}{Q_1(t)+Q_2(t)} U_1(X_1(t)) \\
&\quad + \frac{Q_2(t)}{Q_1(t)+Q_2(t)} U_2(X_2(t)) \\
&\le U_1\Big(\frac{Q_1(t)}{Q_1(t)+Q_2(t)} X_1(t) + \frac{Q_2(t)}{Q_1(t)+Q_2(t)} X_2(t)\Big) \\
&= U_1(\frac{Nc}{Q_1(t)+Q_2(t)}).
\end{aligned}$$

Thus, the optimization problem in (2) is achieved when

$$X_1(t) = X_2(t) = \frac{Nc}{Q_1(t)+Q_2(t)}. \tag{40}$$

As bandwidth is shared equally for all the TCP-like connections, the bandwidth sharing policy in (40) is known as the processor sharing policy.

Now we can view the bandwidth sharing system as a closed two-stage queueing networks (see Figure 3) with two classes of customers. The population of class j customers is $N\alpha_j$, $j = 1$ and 2. The first stage is a single server queue running with the processor sharing policy, and the second stage is an infinite server queue. The average service time of class j customers is $\bar{B}_j/Nc$ at the first stage and $\bar{I}_j$ at the second stage. For such a closed queueing network, we have from the well-known product form solution (see e.g., [19], [27]) that

$$\begin{aligned}
&\mathrm{P}(Q_1(t) = n_1, Q_2(t) = n_2) \\
&= G(N\alpha_1, N\alpha_2) \frac{(\bar{I}_1)^{N\alpha_1 - n_1}}{(N\alpha_1 - n_1)!} \frac{(\bar{I}_2)^{N\alpha_2 - n_2}}{(N\alpha_2 - n_2)!} \\
&\quad \frac{(n_1+n_2)!}{n_1! n_2!} (\frac{\bar{B}_1}{Nc})^{n_1} (\frac{\bar{B}_2}{Nc})^{n_2},
\end{aligned} \tag{41}$$

where

$$\begin{aligned}
&G(N\alpha_1, N\alpha_2) \\
&= \Big(\sum_{n_1=0}^{N\alpha_1} \sum_{n_2=0}^{N\alpha_2} \frac{(\bar{I}_1)^{N\alpha_1 - n_1}}{(N\alpha_1 - n_1)!} \frac{(\bar{I}_2)^{N\alpha_2 - n_2}}{(N\alpha_2 - n_2)!} \\
&\quad \frac{(n_1+n_2)!}{n_1! n_2!} (\frac{\bar{B}_1}{Nc})^{n_1} (\frac{\bar{B}_2}{Nc})^{n_2} \Big)^{-1}
\end{aligned} \tag{42}$$

is the normalization constant.

Theorem 4 *For the processor sharing system with two class of customers, if $\alpha_1\rho_1 + \alpha_2\rho_2 > c$, then the following strong laws of large numbers hold:*

$$\lim_{N\to\infty} \frac{Q_1(t)}{N\alpha_1} = \frac{\rho_1}{x^* + \rho_1}, \quad a.s. \tag{43}$$

$$\lim_{N\to\infty} \frac{Q_2(t)}{N\alpha_2} = \frac{\rho_2}{x^* + \rho_2}, \quad a.s. \tag{44}$$

where

$$\begin{aligned}
x^* = &\frac{1}{2(\alpha_1\rho_1 + \alpha_2\rho_2 - c)} \Big(c(\rho_1+\rho_2) - \rho_1\rho_2 + \\
&\sqrt{(c(\rho_1+\rho_2) - \rho_1\rho_2)^2 + 4(\alpha_1\rho_1 + \alpha_2\rho_2 - c)c\rho_1\rho_2} \Big).
\end{aligned} \tag{45}$$

The Theorem was proved by Pittel [29] based on asymptotic expansion of the product form in (41). It was originally stated in the form of convergence in probability. However, the extension to almost sure convergence in Theorem 4 is straightforward. More recent results along this line can be found in [5], [6]. For completeness, we outline the proof of Theorem 4 in Appendix C.

Now we show the results in Theorem 4 are consistent with those in (A1) and (A2). For this, we first solve the optimization problem in (17) with $x_1 = x_2 = x$ as the two utility functions U_1 and U_2 are identical. Thus, we have from (19) that

$$\pi_j = \frac{\rho_j/x}{1+\rho_j/x}, \quad j = 1, 2. \tag{46}$$

Using this in (17) yields the following quadratic equation

$$\alpha_1 \frac{\rho_1/x}{1+\rho_1/x} x + \alpha_2 \frac{\rho_2/x}{1+\rho_2/x} x = c. \tag{47}$$

Rewrite (47) as follows:

$$x^2(\alpha_1\rho_1 + \alpha_2\rho_2 - c) + x(\rho_1\rho_2 - c(\rho_1+\rho_2)) - c\rho_1\rho_2 = 0. \tag{48}$$

As we assume that $\alpha_1\rho_1 + \alpha_2\rho_2 - c > 0$, there are one positive root and one negative root of the equation. The positive root is x^* in (45), and (46) yields the same results as those in Theorem 4.

VI. Conclusions and Future Research

In this paper, we considered the bandwidth sharing problem for a large number of HTTP-like connections. We showed that if bandwidth shared by TCP-like connections is based on optimizing the sum of their utility functions, then bandwidth shared by a large number of HTTP-like connections can be transformed as an optimization problem for certain utility functions that are related to the original utility functions for the TCP-like connections. Such results are based on assumptions on ergodicity and could be partially verified for the processor sharing case. We also provided a probabilistic interpretation for the relation between the utility functions at the TCP level and the utility functions at the HTTP level. This is done by a large deviation model for bandwidth sharing of TCP-like connections. There seems a natural mapping between the large deviation rate functions and the utility functions.

In the following, we discuss possible extensions of the theory.
(i) It is natural to ask whether one can extend the theory to the case with multiple classes of connections with multiple types of requests, where each type of requests is associated with a utility function. The motivation for this is that in practice there are short lived TCP connections (mice) and long lived TCP connections (elephants). Short lived connections could take advantage of the slow start phase and might receive more bandwidth than long lived connections. As such, short lived connections and long lived connections share the bandwidth quite differently and should be assigned different utility functions. We also note that the utility functions obtained in [24] are for long lived connections. The utility functions for short lived connections still require further study.
Based on similar assumptions on ergodicity to (A1-3), one can still expect the fix point limit to hold. However, it seems difficult to transform the fixed point limit in this case to a standard concave programming problem. Without this, it is not clear whether the fixed point limit has a unique solution. In fact, it would be more interesting to identify a case with multiple fix point solutions!
We note that if there is only *one* class of connections with multiple types of requests, then it is still possible to convert the fix point limit to a standard concave programming problem. Moreover, there is also a probabilistic interpretation based on large deviation rate functions. Due to space limitation, such results will be reported separately.
(ii) The relation between the utility functions and the large deviation rate functions developed in this paper might lead to more large deviation analysis of TCP-like connections. For instance, if one considers charging θ_j (dollars/sec) for one unit bandwidth used by the class j connections, then the utility function associated with the charging scheme, denoted by $U_j^c(\cdot)$, can be modified as follows:

$$U_j^c(x_j) = U_j(x_j) - \theta_j x_j. \tag{49}$$

If we use $W_j(x_j) = -U_j(x_j)$ as the large deviation rate function, then the large deviation rate function for the charging scheme, denoted by $W_j^c(\cdot)$, can be modified as follows:

$$W_j^c(x_j) = W_j(x_j) + \theta_j x_j, \tag{50}$$

and this corresponds to the well-known exponential change of measures (see e.g., [12], [31], [9]). There are techniques developed for fast simulation of large systems based on exponential change of measures. Such techniques might be applicable for simulating a large number of TCP-like (or HTTP-like) connections.
As expected, we also note that in this particular charging scheme, the utility function at the HTTP level, denoted by $\hat{U}_j^c(y_j)$, can be derived from (49) and (11) as follows:

$$\hat{U}_j^c(y_j) = \hat{U}_j(y_j) - \theta_j y_j. \tag{51}$$

(iii) The assumptions on ergodicity was partially verified for the case with processor sharing. It would be of interest to provide a general theoretical condition for these assumptions to hold. Also, bandwidth sharing for real TCP connections is based on a distributed control algorithm and it might take a while to reach the steady state. The assumptions on ergodicity also need further verification by simulation.
(iv) In this paper, we only consider *persistent* TCP connections. It is reported in [14] that there might be impatient users who interrupt TCP connections before they are completed. As in [14], one might assume that a TCP connection with request B will be interrupted when the transfer time exceeds a patient duration $\delta(B)$. The function $\delta(B)$ is an increasing and concave function of B. Based on this, it is possible to derive a fixed point limit. However, it seems difficult to transform the fixed point limit to a standard concave programming problem.

Appendix

I. Appendix A

In this section, we prove Lemma 2. From the Kuhn-Tucker theorem, an optimal solution to the optimization problem in (17) satisfies

$$U_j'(x_j) = \lambda, \quad j = 1, 2, \ldots, J, \tag{52}$$

and

$$\sum_{j=1}^J \alpha_j \pi_j x_j = c. \tag{53}$$

As the utility functions $U_j(x)$, $j = 1, 2, \ldots, J$, are strictly concave and continuously differentiable for all $x > 0$, $U_j'(x)$, $j = 1, 2 \ldots, J$, are strictly decreasing and continuous for all $x > 0$. As such, their inverse functions $U_j'^{-1}(\lambda)$, ranging between 0 and ∞, are also strictly decreasing and continuous (in their effective domain). As $U_j'^{-1}(\lambda)$ ranges between 0 and ∞, there is a unique λ^* that satisfies

$$f(\lambda) = \sum_{j=1}^J \alpha_j \pi_j U_j'^{-1}(\lambda) - c = 0. \tag{54}$$

Thus, the unique optimal solution for (17) is

$$x_j^* = U_j'^{-1}(\lambda^*), \quad j = 1, 2, \ldots, J. \tag{55}$$

Note that $f(\lambda)$ is strictly decreasing and continuous as both $U_j'^{-1}(\lambda)$'s are strictly decreasing and continuous. This implies that $f^{-1}(\cdot)$ is also strictly decreasing and continuous. So is the composite function $U_j'^{-1}(f^{-1}(\cdot))$. From the continuous property of $U_j'^{-1}(f^{-1}(\cdot))$, we have that for every $\epsilon > 0$, there exists a $\delta > 0$ such that for all $|\gamma| \le \delta$

$$\begin{aligned} &|U_j'^{-1}(f^{-1}(\gamma)) - U_j'^{-1}(f^{-1}(0))| \\ &= |U_j'^{-1}(f^{-1}(\gamma)) - U_j'^{-1}(\lambda^*)| \\ &= |U_j'^{-1}(f^{-1}(\gamma)) - x_j^*| < \epsilon. \end{aligned} \tag{56}$$

Following the same argument, we can show that the unique optimum for the maximization problem in (16) occurs at

$$X_j(t) = U_j'^{-1}(\lambda^{(N)}(t)), \quad j = 1, 2 \dots, J, \tag{57}$$

where $\lambda^{(N)}(t)$ is the unique solution of

$$\sum_{j=1}^{J} \frac{Q_j(t)}{N} U_j'^{-1}(\lambda) - c = 0. \tag{58}$$

From the strong law in (5), for every $\epsilon_1 > 0$, there is a (random variable) N_0 such that for all $N \geq N_0$ and $j = 1, \dots, J$,

$$\sup_{0 \leq t \leq t_0} |\frac{Q_j(t)}{N} - \alpha_j \pi_j| \leq \epsilon_1. \tag{59}$$

In view of (58),

$$U_j'^{-1}(\lambda^{(N)}(t)) \leq \frac{cN}{Q_j(t)} \leq \frac{c}{\alpha_j \pi_j - \epsilon_1}, \quad 0 \leq t \leq t_0. \tag{60}$$

Rewrite (58) as follows:

$$0 = \sum_{j=1}^{J} (\frac{Q_j(t)}{N} - \alpha_j \pi_j) U_1'^{-1}(\lambda^{(N)}(t)) + f(\lambda^{(N)}(t)). \tag{61}$$

Using (59) and (60) in (61), we have

$$|f(\lambda^{(N)}(t))| \leq J\epsilon_1 \max_{1 \leq j \leq J} [\frac{c}{\alpha_j \pi_j - \epsilon_1}]. \tag{62}$$

Choose ϵ_1 small enough such that

$$J\epsilon_1 \max_{1 \leq j \leq J} [\frac{c}{\alpha_j \pi_j - \epsilon_1}] \leq \delta.$$

Form (56), it then follows that

$$\begin{aligned} &\sup_{0 \leq t \leq t_0} |U_j'^{-1}\left(f^{-1}(f(\lambda^{(N)}(t)))\right) - x_j^*| \\ &= \sup_{0 \leq t \leq t_0} |X_j(t) - x_j^*| \leq \epsilon. \end{aligned} \tag{63}$$

This shows that

$$\lim_{N \to \infty} \sup_{0 \leq t \leq t_0} |X_j(t) - x_j^*| = 0, \quad a.s. \tag{64}$$

II. Appendix B

In this section, we prove Lemma 3. Consider a particular random request B_j from a class j connection. Denote by T_j the transfer time for transmitting the request. From Lemma 2, for every $\epsilon > 0$, there is a N_0 such that for all $N \geq N_0$,

$$\sup_{0 \leq t \leq t_0} |X_j(t) - x_j^*| < \epsilon. \tag{65}$$

From the independent assumption in (A3), we then have

$$\frac{B_j}{x_j^* + \epsilon} \leq T_j \leq \frac{B_j}{x_j^* - \epsilon}. \tag{66}$$

This then implies that

$$|T_j - \frac{B_j}{x_j^*}| \leq \frac{B_j}{x_j^*} \frac{\epsilon}{x_j^* - \epsilon} \tag{67}$$

$$\leq \frac{t_0}{x_j^*} \frac{\epsilon}{x_j^* - \epsilon}. \tag{68}$$

Hence,

$$\lim_{N \to \infty} T_j = \frac{B_j}{x_j^*}, \quad a.s. \tag{69}$$

It then follows from (6) in (A2) that

$$\lim_{N \to \infty} \bar{T}_j = \frac{\bar{B}_j}{x_j^*}. \tag{70}$$

III. Appendix C

In this section, we outline the proof for Theorem 4. Rewrite (41) as follows:

$$\begin{aligned} &\mathrm{P}(Q_1(t) = n_1, Q_2(t) = n_2) \\ &= \tilde{G}(N\alpha_1, N\alpha_2) \frac{1}{(N\alpha_1 - n_1)!} \frac{1}{(N\alpha_2 - n_2)!} \\ &\quad \frac{(n_1 + n_2)!}{n_1! n_2!} (\frac{\rho_1}{Nc})^{n_1} (\frac{\rho_2}{Nc})^{n_2}, \end{aligned} \tag{71}$$

where

$$\begin{aligned} &\tilde{G}(N\alpha_1, N\alpha_2) \\ &= \Big(\sum_{n_1=0}^{N\alpha_1} \sum_{n_2=0}^{N\alpha_2} \frac{1}{(N\alpha_1 - n_1)!} \frac{1}{(N\alpha_2 - n_2)!} \\ &\quad \frac{(n_1 + n_2)!}{n_1! n_2!} (\frac{\rho_1}{Nc})^{n_1} (\frac{\rho_2}{Nc})^{n_2} \Big)^{-1}. \end{aligned} \tag{72}$$

The idea of showing the strong laws in (43) and (44) is to show that for every $\epsilon > 0$,

$$\mathrm{P}(\max[|\frac{Q_1(t)}{N\alpha_1} - \frac{\rho_1}{x^* + \rho_1}|, |\frac{Q_2(t)}{N\alpha_2} - \frac{\rho_2}{x^* + \rho_2}|] > \epsilon)$$

converges to 0 exponentially fast so that the Borel-Cantelli lemma (see e.g., [11]) can be applied. Such an idea was previously used in loss networks by Kelly [20] to prove strong laws and by Chang and Wang [10] to prove large deviations principles.

Use Stirling's formula ($n! \approx n^n e^{-n}$) to show that

$$\begin{aligned} &\mathrm{P}(\frac{Q_1(t)}{N\alpha_1} = \pi_1, \frac{Q_2(t)}{N\alpha_2} = \pi_2) \\ &\approx \tilde{G}(N\alpha_1, N\alpha_2) e^{-N \log N} e^{N g(\pi_1, \pi_2)}, \end{aligned} \tag{73}$$

where

$$\begin{aligned} g(\pi_1, \pi_2) = &\alpha_1(1 - \pi_1) \\ &- \alpha_1(1 - \pi_1)\log(\alpha_1(1 - \pi_1)) \\ &+ \alpha_2(1 - \pi_2) - \alpha_2(1 - \pi_2)\log(\alpha_2(1 - \pi_2)) \\ &+ (\alpha_1\pi_1 + \alpha_2\pi_2)\log(\alpha_1\pi_1 + \alpha_2\pi_2) \\ &- \alpha_1\pi_1\log(\alpha_1\pi_1) - \alpha_2\pi_2\log(\alpha_2\pi_2) \\ &+ \alpha_1\pi_1\log(\frac{\rho_1}{c}) + \alpha_2\pi_2\log(\frac{\rho_2}{c}). \end{aligned}$$

(We note that the error term in Stirling's formula is much smaller than others and can be neglected for large N.) The term $\exp(-N \log N)$ is the common term and can be absorbed by the normalization constant $\tilde{G}(N\alpha_1, N\alpha_2)$. Thus, the most likely state can be found by maximizing $g(\pi_1, \pi_2)$.

Now we show that $g(\pi_1, \pi_2)$ is strictly concave. Let

$$H = \begin{bmatrix} \frac{\partial^2 g}{\partial \pi_1^2} & \frac{\partial^2 g}{\partial \pi_1 \partial \pi_2} \\ \frac{\partial^2 g}{\partial \pi_1 \partial \pi_2} & \frac{\partial^2 g}{\partial \pi_2^2} \end{bmatrix} \tag{74}$$

be the Hessian matrix of g. It is straightforward to verify that

$$\begin{aligned} &\begin{vmatrix} \frac{\partial^2 g}{\partial \pi_1^2} & \frac{\partial^2 g}{\partial \pi_1 \partial \pi_2} \\ \frac{\partial^2 g}{\partial \pi_1 \partial \pi_2} & \frac{\partial^2 g}{\partial \pi_2^2} \end{vmatrix} \\ &= \frac{\alpha_1 \alpha_2 (\alpha_1 \pi_1^2 + \alpha_2 \pi_2^2)}{\pi_1(1-\pi_1)\pi_2(1-\pi_2)(\alpha_1\pi_1 + \alpha_2\pi_2)} \\ &> 0, \end{aligned} \tag{75}$$

and that

$$\frac{\partial^2 g}{\partial \pi_1^2} = -\frac{\alpha_1^2 \pi_1^2 + \alpha_1 \alpha_2 \pi_2}{\pi_1(1-\pi_1)(\alpha_1\pi_1 + \alpha_2\pi_2)} < 0, \tag{76}$$

$$\frac{\partial^2 g}{\partial \pi_2^2} = -\frac{\alpha_2^2 \pi_2^2 + \alpha_1 \alpha_2 \pi_1}{\pi_2(1-\pi_2)(\alpha_1\pi_1 + \alpha_2\pi_2)} < 0. \tag{77}$$

Thus, the matrix H has two negative eigenvalues and it is negative definite. As such, the function $g(\pi_1, \pi_2)$ is strictly concave. Letting $\frac{\partial g}{\partial \pi_j} = 0$, $j = 1$ and 2 yields

$$\log\left(\frac{(1-\pi_j)(\alpha_1\pi_1 + \alpha_2\pi_2)}{\pi_j} \frac{\rho_j}{c}\right) = 0, \quad , j = 1, 2. \tag{78}$$

This implies that there is some $x^* > 0$ such that

$$x^* = \frac{1-\pi_1}{\pi_1}\rho_1 = \frac{1-\pi_2}{\pi_2}\rho_2 = \frac{c}{\alpha_1\pi_1 + \alpha_2\pi_2}. \tag{79}$$

From (79), it follows that x^* satisfies the equation in (47) and have the form in (45). This shows that $(\pi_1, \pi_2) = (\frac{\rho_1}{\rho_1 + x^*}, \frac{\rho_2}{\rho_2 + x^*})$ is the most likely state. Moreover, the states outside the neighborhood of the most likely state have exponentially smaller probabilities than that of the most likely state.

References

[1] M. Avriel, "Nonlinear Programming: Analysis and Methods", *Prentice-Hall*, 1976.

[2] F. Baccelli, D. Hong and Z. Liu, "Fixed point methods for the simulation of the sharing of a local loop by a large number of interacting TCP connections," *Proc. ITC Specialist Conference on Local Loop*, Barcelona, Spain, 2001. Also available as an *INRIA research report*, April, 2001.

[3] P. Barford and M. E. Crovella. Generating representative web workloads fot network and server performance evaluation. In *Proceedings of ACM Sigmetrics'98*, 1998.

[4] P. Barford and M. E. Crovella. A performance evaluation of hyper text transfer protocols. In *Proceedings of ACM Sigmetrics'99*, 1999.

[5] A. Berger, L. Bregman and Y. Kogan, "Bottleneck analysis in multiclass closed queueing networks and its application," *Queueing Systems*, Vol. 10, 1998.

[6] A. Berger and Y, Kogan, "Dimensioning Bandwidth For Elastic Traffic In High-Speed Data Networks," *IEEE/ACM Transactions on Networking*, Vol. 8, pp. 643-654, 2000

[7] T. Bonald and L . Massoulié, "Impact of fairness of Internet Performance," *Proc. of ACM Sigmetrics 2001*, pp. 82-91, Cambridge, Massachusetts, U.S.A.

[8] L.S. Brakmo and L.L. Peterson, "TCP Vegas: end-to-end congestion avoidance on a global Internet," *IEEE Jounral on Selected Area in Communications*, Vol. 13, October, 1995.

[9] C.S. Chang. *Performance Guarantees in Communication Networks*. London: Springer-Verlag, 2000.

[10] C.S. Chang and H.J. Wang, "Large deviations for large capacity loss networks with fixed routing and polyhedral admission sets," *Discrete Event Dynamic Systems*, Vol. 7, pp. 391-418, 1997.

[11] Y.S. Chow and H. Teicher, *Probability Theory: Independence, Interchangeability, Martingales*. New York: Springer-Verlag, 1988.

[12] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Boston: Jones and Barlett Publishers, 1992.

[13] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, Vol. 1, pp. 397-413, August, 1993.

[14] S.B. Fredj, T. Bonald, A. Proutiere, G. Régnié, J.W. Roberts, "Statistical bandwidth sharing: a study of congestion at flow level," *ACM SIGCOMM'01*, August 27-31, 2001, San Diego, California, U.S.A.

[15] P.W. Glynn and W. Whitt, "Large deviations behavior of counting processes and their inverses," *Queueing Systems*, Vol. 17, pp. 107-128, 1994.

[16] Hiriart-Urruty, J.-B., and Lemaaréchal, C. *Convex Analysis and Minimization Algorithms II*. New York: Springer-Verlag, 1991.

[17] V. Jacobson, "Congestion avoidance and control," *Proc. of ACM SIGCOMM'88*, August 1988.

[18] Z. Liu, N. Niclausse, C. Jalpa-Villanueva, "Traffic Model and Performance Evaluation of Web Servers", *Performance Evaluation*, Vol 46/2-3, pp 77-100, 2001.

[19] F.P. Kelly. *Reversibility and Stochastic Networks*. Wiley, 1979.

[20] F.P. Kelly, "Loss networks," *Ann. Appl. Probab.*, Vol. 1, pp. 319–378, 1991.

[21] F.P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, Vol. 8, pp. 33-37, 1997.

[22] F.P. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of Operations Research Society*, Vol 49, pp. 237-252, 1998.

[23] S. Kunniyur and R. Srikant, "End-to-end congestion control: utility functions, random losses and ECN marks," *Proc. of IEEE INFOCOM* 2000, Tel Aviv, Israel.

[24] S.H. Low, "A duality model of TCP flow control," *Proceedings of ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, Sept. 18-20, 2000. http://netlab.caltech.edu.

[25] L. Massoulié and J. Roberts, "Bandwidth sharing: objectives and algorithms," *Proc. of IEEE INFOCOM'99*, New York, 1999.

[26] J. Mo and J. Walrand, "Fair end-to-end window based congestion control," *IEEE/ACM Transactions on Networking*, Vol. 8, pp. 556-567, 2000.

[27] R. Nelson, *Probability, Stochastic Processes, and Queueing Theory: the Mathematics of Computer Performance Modeling*. Springer-Verlag: New York, 1995.

[28] J. E. Pitkow, "Summary of WWW Characterizations." *Proceedings of WWW7, Computer Networks* 30(1-7): 551-558 (1998).

[29] B. Pittel, "Closed exponential networks of queues with saturation: the Jackson-type stationary distribution and its asymptotic analysis," *Mathematics of Operations Research*, Vol. 4, pp. 357-378, 1979.

[30] R.T. Rockafeller. *Convex Analysis*. Princeton: Princeton University Press, 1970.

[31] A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis: Queues, Communication and Computing*. 1994.