

Simple Models of Network Access, with Applications to the Design of Joint Rate and Admission Control

Michel Mandjes ^{*,†,**} Debasis Mitra ^{*} Werner Scheinhardt ^{†,**}
michel@cwi.nl mitra@lucent.com werner@math.utwente.nl

^{*} Bell Laboratories/Lucent Technologies, P.O. Box 636, Murray Hill, NJ 07974, United States

[†] Faculty of Mathematical Sciences, University of Twente, P.O. Box 217, 7500 AE Enschede, the Netherlands

^{**} CWI, Kruislaan 413, P.O. Box 94079, 1090 GB Amsterdam, the Netherlands

Abstract—At the access to networks, in contrast to the core, distances and feedback delays, as well as link capacities are small, which has network engineering implications that are investigated in this paper. We consider a single point in the access network which multiplexes several bursty users. The users adapt their sending rates based on feedback from the access multiplexer. Important parameters are the user's peak transmission rate p , which is the access line speed, the user's guaranteed minimum rate r , and the bound ϵ on the fraction of lost data. Two feedback schemes are proposed. In both schemes the users are allowed to send at rate p if the system is relatively lightly loaded, at rate r during periods of congestion, and at a rate between r and p , in an intermediate region. For both feedback schemes we present an exact analysis, under the assumption that the users' job sizes and think times have exponential distributions. We use our techniques to design the schemes jointly with admission control, i.e., the selection of the number of admissible users, to maximize throughput for given p , r , and ϵ . Next we consider the case in which the number of users is large. Under a specific scaling, we derive explicit large deviations asymptotics for both models. We discuss the extension to general distributions of user data and think times.

I. INTRODUCTION

In today's communication networks, design and control of the network core and access are different, primarily because of the differences in scale in bandwidth and distance. Quite often the bottleneck is the access, rather than the core. This may happen because the access network is characterized by relatively low line speeds and the limited ability of users to buffer and shape traffic (think of the extreme case of a user with a wireless handset). Access control, supported by the use of feedback, is an important mechanism to address this problem. Since distances between users/clients and network access points are relatively short, feedback delay due to propagation is negligible, which contributes to the efficacy of feedback control. In this paper we investigate the problems of access control by introducing simple models and techniques for their evaluation, design and performance optimization.

We make three main contributions. First, we present two simple models of network access. The models provide a framework for the joint design of feedback-based schemes for the adaptation of source rates and admission control. Second, we show how to compute the stationary behavior of the aforementioned feedback queues. We illustrate the use of these

techniques to solve the design problem. Finally we show how to use the theory of large deviations to obtain explicit results when the system and the number of sources is large.

In our model each user alternates between 'on' periods of transmission, and 'off' periods or 'think times'. The user model here differs from the familiar on-off source models, e.g. [2], in that *file sizes* (where a file size is the amount of data the source transmits during an active period) are independent, identically distributed (iid) random variables, but the *on periods* are not specified *a priori*. The on-periods are determined by the combination of the file sizes and the transmission rates allocated by the access control scheme, to be described below, which depend on the interaction of the multiplexer with the collective behavior of users. In contrast, the think times are iid random variables. The lengths of on periods and the throughputs of the individual users are key performance quantities to be obtained from an analysis of the model.

The access line speed (p for each of the users) is typically small compared to the output rate of the access multiplexer, and therefore constitutes an important constraint.

Another model feature is the minimum throughput rate (r) that is guaranteed to the users. In a number of applications clients derive zero utility if the throughput is below a threshold. This point has been made by Massoulié and Roberts [20] for the case of TCP traffic, in which performance collapse may ensue. As soon as the notion of a minimum guaranteed rate is introduced, admission control needs to be considered, together with the calculation of the capacity of the network.

We present two schemes for feedback-based adaptation of the users' rates. In both schemes users are allowed to transmit at rate p if the system is relatively lightly loaded, at rate r during periods of congestion, and at a rate between r and p , which is determined by the processor sharing discipline, in an intermediate region. In our model the feedback signal from the buffer to the sources is assumed to arrive with negligible delay, which is reasonable when the round trip distances are small. In the first scheme the feedback is based on the number of active users and whether the queue is empty or not. The second scheme utilizes a threshold B^* on the buffer content. Depending on whether the buffer content is less

than, greater than, or equal to B^* , the user rate is p , r , or determined by the processor sharing discipline, respectively. Importantly, the second scheme does not require knowledge of the (activity) state of the users, and is simpler to implement on that count. However, the analysis of the first queue is simpler and our results for it are in closed form. As discussed below, the combined use of the threshold and the number of admissible users in the second model provides a greater facility for regulating important trade-offs.

We note a feature of the behavior of the models, which is counter-intuitive. Our analysis shows that the effect of feedback is to *increase* the mean time that the buffer is empty, while simultaneously *increasing* the throughputs of the users. The explanation of this apparently paradoxical behavior is that the rates allocated to the users are *higher* during the periods that the buffer is empty. Hence feedback has the effect of reducing the on periods and consequently the cycle time of each user, and thereby it increases individual users' throughput and the system throughput, which is defined as the sum of the individual users' throughputs.

Next consider the role of the feedback parameter B^* , the threshold level in the buffer content in the second model. By reasoning as above, we infer that increasing B^* has the effect of increasing the throughput of the users. However, this is at the cost of increasing the probability of buffer overflow. Similarly, with all other parameters held fixed, increasing the number of users N has the effect of increasing both the total system throughput as well as the probability of buffer overflow. Hence, this model allows the study of interesting trade-offs between several important quantities, including the individual users' throughput, the system throughput, the loss probability and the number of users. By proper design the parameters B^* and N can regulate the trade-off.

There are several possible frameworks in which the trade-offs may be studied and quantified. We require that the fraction of source data that is lost does not exceed a given QoS parameter ϵ . Then we may seek a joint design of the feedback control scheme and admission control, i.e., selection of B^* and N such that the system throughput is maximized. The numerical procedure that is developed in this paper allows such design questions to be addressed. Indeed one of the highlights of the numerical results that we present later in the paper is the computed solution for an instance of the above design problem.

The two proposed feedback schemes both fall into the category of feedback fluid queues, which were introduced in [26] as generalizations of the well-known Markov-modulated fluid models in e.g. [2], [17]. In the latter, a fluid buffer receives or depletes fluid at rate r_i (positive or negative) at times when a background continuous-time Markov chain is in state i . Typically, stochastic fluid models are characterized by the generator (Q) of the background process and a diagonal rate matrix (R) which contains all fluid rates r_i . In feedback fluid queues most of the above remains true, except that the behavior of the background process (i.e. the matrix Q), and possibly the matrix R as well, depends on the current buffer content. As a result the background process is no longer

an autonomous Markov process. In this paper we confine ourselves to feedback fluid queues in which Q and R are piecewise constant functions of the buffer content level, see also [27]. Notice the crucial difference with [11], in which there are also thresholds, like here, but these only affect R , and not also Q . The analysis of these feedback fluid queues, based on spectral expansions [2], [10], [11], [17], is one of the main contributions of this paper.

An intrinsic drawback of the analytical approach described so far is that it is computationally intensive. This provides motivation for simpler, asymptotic approaches to deal with the access models. In both access models, we study large-deviations asymptotics for the scaling introduced by Weiss [31], i.e. the regime in which the number of users grows large and resources (buffer and bandwidth) scale proportionally. We derive 'exponential approximations', comparable to those in [3], [4], [9] for the ordinary FIFO discipline. Exponential approximations of the first feedback model, i.e., without the threshold, were obtained earlier by Ramanan and Weiss [23] for exponential file sizes and think times. Our major contribution is that these results are explicit and the computations are simple. Some of them have nice insensitivity properties, i.e., depend on the distributions of the think times and job sizes only through their means.

The role of feedback in packet networks has a long history, see [28, Ch. 7] for a review. Recently there has been a resurgence of interest, driven in part by work on explicit congestion notification (ECN) marking schemes by Gibbens and Kelly [12], [13], [16], and others. In [22] a feedback model with feedback delays is considered based on the marking scheme of [10]. However, the feedback considered there regulates the marking of packets to be dropped later on, whereas in our model the feedback regulates the actual source rate. The processor sharing discipline has been highlighted in recent work on QoS delivery by Roberts *et al.* [24], [25], albeit in the bufferless framework.

The organization of the paper is as follows. Section II deals with the first access model. Section III considers the more advanced feedback scheme that utilizes a threshold. The large deviations asymptotics are described in Section IV. Finally Section V reports on several numerical examples.

For reasons of space we have had to omit several general results, specifically a treatment of general feedback fluid queues, including finite buffers and multiple thresholds, and the large deviations asymptotics for the feedback model with a threshold. These results are in [19].

II. FEEDBACK MODEL WITHOUT THRESHOLD

A. Model

We model a single aggregation or multiplexing point in the network. The output trunk speed is C , and there are N access links. The peak (or line) rate of each access link is p . The minimum rate given to each access user is r . We make the simplifying modeling assumption that traffic can be considered to be continuous fluid. For the benefit of the reader, we note that in fluid models there may be considerable transfer of fluid,

TABLE I

Allowed user transmission rates (and sign of the buffer drift), as functions of number of active sources and buffer content.

Number of active sources	$W(t) = 0$	$W(t) > 0$
$0 \leq Y(t) \leq N'$	p (0)	r (-)
$N' < Y(t) \leq N''$	$C/Y(t)$ (0)	r (-)
$N'' < Y(t) \leq N$	NA	r (+)

i.e., high throughput, even during periods when the buffer is empty. This is because in fluid models, once the buffer is empty, it remains empty for as long as the total input rate does not exceed the output capacity.

In this section we describe the scheme without threshold. When the multiplexer buffer is empty, the access rate will be determined by dividing equally the trunk rate between the (small) number of active users, truncated to p . At the other extreme, when the number of active users exceeds a critical number, N'' in this work, then the fair share of the trunk rate for an active user drops below the guaranteed rate of r , and the buffer is no longer empty and its content grows. As long as the buffer content is positive, each source is assigned the rate of just r . Clearly, N'' is the largest integer not exceeding C/r .

While the buffer is empty and the number of active users is small, say below a critical number N' , then each active user transmits at peak access rate p . It is easily seen that N' is the largest integer not exceeding C/p . When the number of active users is between N' and N'' , which is the middle range, and the buffer is empty, then the trunk speed is shared equally, i.e., the processor sharing regime holds. Hence, when the buffer is empty there are two regimes characterized respectively by the access line speed and the processor sharing rate. In contrast, when the buffer is not empty, there is a unique transmission rate, namely r , the guaranteed minimum.

Table I summarizes the feedback protocol. Let $Y(t)$ denote the number of active sources at time t , and $W(t)$ the buffer content at time t . In the table the allocated rate, as well as the sign of the 'drift' of the buffer content are given, as functions of $Y(t)$ and $W(t)$.

An important aspect of our model is the behavior of the homogeneous sources. Each source alternates between activity ('on') and inactivity ('off'). The inactivity periods are independent, exponentially distributed random variables with mean λ^{-1} . Each source transmits a file during its activity period, whose size (in bits) is independent of everything else and exponentially distributed with mean μ^{-1} . The length of the induced activity period is not given *a priori*, since this depends on the rate(s) at which the file is transmitted. An important parameter for the Quality of Service is ϵ , which the buffer overflow probability is required not to exceed.

We follow a conventional approach in inferring finite-buffer performance from an infinite-buffer model. Since the long run fraction of time a source is on, given that the buffer level is large, is $\lambda/(\lambda + \mu r)$ (note that $(\mu r)^{-1}$ is the maximum time to transmit a file of size μ^{-1}), the stability condition of this

infinite-buffer model is given by

$$N \frac{\lambda r}{\lambda + \mu r} < C. \quad (1)$$

In practice, there is a need to prevent N from being too large. One of the objectives of the analysis is to calculate the capacity of our system, which is the largest value of N such that the overflow probability of the queue does not exceed ϵ . In particular, the number $N_r(\epsilon)$ of connections admissible as a function of ϵ and the guaranteed rate r , is an important design parameter. We return to the calculation of this quantity in Section V-B.

B. Preliminary results

To analyze the model described in Section II-A, we first review some results from the model *without* feedback. Anick, Mitra, and Sondhi [2] consider the model in which the sources transmit at a constant rate, say r , while in the on-state, i.e., the allowable transmission rate does not depend on current occupancy of the system. Detailed results on this model are available [2], [10], [17].

Buffer content distribution. For this case without feedback we denote the state of the background process (i.e. the number of sources that are transmitting) and the buffer content process at time t as $X(t)$ and $V(t)$ respectively. In [2] the stationary distribution of the content process is given. It is computed as follows. Clearly $X(\cdot)$ constitutes a continuous-time Markov chain on the state space $\{0, \dots, N\}$. The (i, j) th element ($i \neq j$) of its generator Q is given by

$$Q(i, j) := \begin{cases} (N - i)\lambda & \text{if } j = i + 1, \\ ir\mu & \text{if } j = i - 1, \\ 0 & \text{otherwise.} \end{cases}$$

The diagonal elements ($i = j$) are such that the rowsums are zero. Element (i, j) represents the probability flux of the continuous-time Markov chain from state i to state j . Define by R the diagonal matrix $\text{diag}\{r_0, \dots, r_N\}$ with r_i the net input rate if there are i sources in the on state, i.e., $r_i = ir - C$. To find the stationary buffer content distribution, we first define $F_i(x) := \mathbb{P}(X = i, V \leq x)$. It is not hard to show that the vector $\mathbf{F}(\cdot) := (F_0(\cdot), \dots, F_N(\cdot))$ satisfies $\mathbf{F}'(x)R = \mathbf{F}(x)Q$. The spectral expansion of the solution is given by

$$\mathbf{F}(x) = \sum_{j=0}^N a_j \mathbf{v}_j \exp[z_j x], \quad (2)$$

where the a_j are coefficients determined later, and (z_j, \mathbf{v}_j) is an eigenvalue-eigenvector pair, i.e., obtained from $z_j \mathbf{v}_j R = \mathbf{v}_j Q$.

The coefficients a_j are calculated as follows. Define D , the set of states with downward drifts, by all states i such that $ir < C$, and U , the set of states with upward drifts, by all other states. Assume that there is no state i such that $ir = C$. Let $Q_{DD}, Q_{DU}, Q_{UD}, Q_{UU}$ be the submatrices that are obtained by partitioning Q . The vector $\mathbf{F}_D(x)$ consists of the $F_i(x)$ with $i \in D$; $\mathbf{F}_U(x)$ is defined analogously. It is easily seen that $a_j = 0$ if $\text{Re}(z_j) > 0$, as the distribution should range

between 0 and 1. The remaining a_j follow from $\mathbf{F}_U(0) = \mathbf{0}$. It turns out that there are just as many unknowns as equations. Clearly, $\mathbb{P}(V \leq x) = \sum_i F_i(x)$.

Idle and busy periods. Elwalid and Mitra [10] give explicit expressions for a number of quantities that are related to the busy and idle periods of the queue. A busy period is defined as a period in which the buffer content is positive, whereas an idle period is a period in which the buffer is empty. It is easily seen that at the beginning of a busy period the number of sources in the on-state is equal to $N'' + 1$; at the end of the busy period the number of sources in the on state is in D . The lengths of consecutive busy and idle periods are independent. Denote by \mathbf{P} the distribution at the end of the busy period. Then it is not hard to prove that

$$\mathbf{P} = \frac{1}{\langle \mathbf{F}_D(0)Q_{DD}, \mathbf{1} \rangle} \mathbf{F}_D(0)Q_{DD}, \quad (3)$$

see Equation 5.9 of [10]; $\langle \cdot, \cdot \rangle$ denotes the inner product of two vectors. The mean idle period $\mathbb{E}I$ is given by

$$\mathbb{E}I = -\frac{\sum_{i \in D} F_i(0)}{\langle \mathbf{F}_D(0)Q_{DD}, \mathbf{1} \rangle}.$$

Finally, the mean busy period $\mathbb{E}B$ can be calculated using $\sum_{i \in D} F_i(0) = \mathbb{E}I / (\mathbb{E}I + \mathbb{E}B)$:

$$\mathbb{E}B = \mathbb{E}I \cdot \frac{1 - \sum_{i \in D} F_i(0)}{\sum_{i \in D} F_i(0)}. \quad (4)$$

C. Analysis

The model analyzed in this section has been described in Section II-A. Recall that $Y(t) \in \{0, \dots, N\}$ denotes the number of sources that are transmitting at time t in the feedback model of Section II-A. Notice that this does *not* constitute a Markov chain, unlike $X(t)$ in Section II-B. This is because the sojourn times and transition probabilities depend on the amount of fluid stored in the buffer. However, as long as the buffer is empty, $Y(t)$ behaves as a continuous-time Markov chain.

Denote the stationary buffer content distribution in the feedback model by $\mathbb{P}(W < x)$. A busy period in this model is distributed as the random variable B' , and an idle period as I' . The sequence of busy periods is i.i.d., and so is the sequence of idle periods, as can be seen easily. The distribution of $Y(t)$ at the end of the busy period is denoted by (the vector) P' . The next lemma links B' and P' to the corresponding quantities in the model without feedback.

Lemma II.1: Busy periods B and B' have the same distribution. Also, the distributions P and P' are identical.

Proof. Both in the models with (Section II-A) and without (Section II-B) feedback, during busy periods on-periods terminate at a rate $ir\mu$ when there are i sources in their on-state. In both models the busy period starts when there are $N'' + 1 = \lceil C/r \rceil$ sources transmitting. Hence, the buffer dynamics in both models have the same probabilistic properties during a busy period. This immediately implies both assertions. \square

Corollary II.2: With the same argument as in the proof of the previous lemma, we find

$$\mathbb{P}(W \leq x | W > 0) = \mathbb{P}(V \leq x | V > 0), \quad x > 0.$$

This immediately implies that $\mathbb{P}(W \leq x)$ equals

$$\mathbb{P}(V \leq x | V > 0)\mathbb{P}(W > 0) + \mathbb{P}(W = 0) = \frac{\sum_i F_i(x) - \sum_i F_i(0)}{1 - \sum_i F_i(0)}\mathbb{P}(W > 0) + \mathbb{P}(W = 0).$$

As $\mathbb{P}(W = 0) = 1 - \mathbb{P}(W > 0)$, the only quantity that is left to compute is the probability of an empty buffer in our feedback model. This is given by

$$\mathbb{P}(W = 0) = \frac{\mathbb{E}I'}{\mathbb{E}I' + \mathbb{E}B'} = \frac{\mathbb{E}I'}{\mathbb{E}I' + \mathbb{E}B},$$

applying Lemma II.1. As we know $\mathbb{E}B$ from Section II-B, we only have to find $\mathbb{E}I'$. This will be done in the next lemma, but first we introduce some required notation.

Q'_{DD} is a square matrix of dimension $N'' + 1$. For $i \neq j$:

$$Q'_{DD}(i, j) := \begin{cases} (N - i)\lambda & \text{if } j = i + 1, \\ ip\mu & \text{if } j = i - 1, \\ 0 & \text{otherwise,} \end{cases}$$

if $i \leq N'$, and

$$Q'_{DD}(i, j) := \begin{cases} (N - i)\lambda & \text{if } j = i + 1, \\ C\mu & \text{if } j = i - 1 \\ 0 & \text{otherwise,} \end{cases}$$

if i is between $N' + 1$ and N'' . The diagonal elements are such that the rowsums are zero, except for $Q'_{DD}(N'', N'')$, which equals $-C\mu - (N - N'')\lambda$. Notice that, as long as the buffer is empty, $Y(t)$ is a Markov chain which obeys the transition rates of Q'_{DD} .

Lemma II.3: With \mathbf{P} defined in (3), the mean idle time in the feedback-based model is given by

$$\mathbb{E}I' = \langle -\mathbf{P}(Q'_{DD})^{-1}, \mathbf{1} \rangle. \quad (5)$$

Proof. This follows directly from standard results of mean passage times [15]. It says that the mean time spent by $Y(t)$ in j before the set D is left, given that the process starts in i , is given by the (i, j) entry of $-(Q'_{DD})^{-1}$.

Then the reasoning is analogous to Equations (5.11) and (5.14) of [10], as follows. The vector $-\mathbf{P}(Q'_{DD})^{-1}$ gives the mean time spent in all states in D during an idle period of the buffer, applying Lemma II.1. The sum of its entries is the mean length of the idle period. \square

We arrive at the following result for the buffer content distribution in the feedback-based model. A similar proportionality result was found, although not explicitly mentioned, in Adan *et al.* [1], where another feedback fluid queue is analyzed that has two types of behavior, depending on whether the buffer is empty or not.

Theorem II.4: In the feedback-based model, the stationary buffer content distribution $\mathbb{P}(W \leq x)$ is given by

$$\frac{\sum_i F_i(x) - \sum_i F_i(0)}{1 - \sum_i F_i(0)} \frac{\mathbb{E}B}{\mathbb{E}I' + \mathbb{E}B} + \frac{\mathbb{E}I'}{\mathbb{E}I' + \mathbb{E}B},$$

TABLE II

Allowed user transmission rates (and sign of the buffer drift), as functions of number of active sources and buffer content.

Number of active sources	$W(t) < B^*$	$W(t) = B^*$	$W(t) > B^*$
$0 \leq Y(t) \leq N'$	p (-)	NA	r (-)
$N' < Y(t) \leq N''$	p (+)	$C/Y(t)$ (0)	r (-)
$N'' < Y(t) \leq N$	p (+)	NA	r (+)

where the vector $\mathbf{F}(x)$ is given by (2), $\mathbb{E}B$ by (4), and $\mathbb{E}I'$ by (5). Equivalently,

$$\mathbb{P}(W > x) = \mathbb{P}(V > x) \frac{\mathbb{E}I + \mathbb{E}B}{\mathbb{E}I' + \mathbb{E}B}.$$

An important interpretation of the above theorem is the following: the gain with respect to the model without feedback (where the sources always send at rate r when active) is expressed by the term

$$\frac{\mathbb{E}I + \mathbb{E}B}{\mathbb{E}I' + \mathbb{E}B} \leq 1. \quad (6)$$

The fact that this ratio is less than 1 is due to the fact that $\mathbb{E}I' \geq \mathbb{E}I$, which can be understood as follows. Clearly, I can be interpreted as a first entrance time in the birth-death process with generator Q'_{DD} , namely as the first entrance time to state $N'' + 1$, starting from a state $i \leq N''$ that is drawn from the distribution P' . Similarly I is the corresponding entrance time in the birth-death process with generator Q , due to the fact that $P = P'$. Since the death rates in Q'_{DD} are larger than those in Q , while the birth rates are equal, it follows that $\mathbb{E}I' \geq \mathbb{E}I$.

In many situations, particularly when the number of sources is large, $\mathbb{E}B$ will be much smaller than $\mathbb{E}I$ and $\mathbb{E}I'$. In that case (6) is well approximated by $\mathbb{E}I \cdot (\mathbb{E}I')^{-1}$. That is, the ratio of mean idle times of the buffer in the models without and with feedback effectively quantifies the performance gain from feedback.

III. FEEDBACK MODEL WITH THRESHOLD

A. Model

In this section we consider a generalization of the feedback model presented in Section II-A. As before, $Y(t)$ is the number of active users at time t (with state space $\{0, \dots, N\}$) and $W(t)$ is the buffer content at time t . We now introduce a threshold level $B^* > 0$ such that the sources are allowed to send at peak rate p as long as $W(t) < B^*$. When $W(t) > B^*$ the sources are allowed to transmit data only at the guaranteed rate r . When $W(t) = B^*$ the processor sharing policy applies. In that case the buffer content process will 'stick' at level B^* for as long as the number of active users lies in the set $N' + 1, \dots, N''$. The algorithm is summarized in Table II. Notice that the model described in Section II-A is a limiting case of this model, as it is obtained by letting $B^* \downarrow 0$. We will describe the exact solution of the stationary buffer content distribution.

Notice that the stability condition for this model is the same as (1). Define the matrix $Q^{(r)}$ to be Q (as defined in Section II-B); generator $Q^{(p)}$ is also defined as Q , but with rate r

replaced by rate p . Finally, $Q^{(*)}$ is similarly defined, but now $Q(i, i-1) = C\mu$. The idea is that $Y(t)$ behaves like a Markov chain with generator $Q^{(p)}, Q^{(*)}, Q^{(r)}$, whenever the buffer content process $W(t)$ is respectively below, at, or above level B^* . Furthermore $R^{(r)}$ is defined as in Section II-B, i.e. as a diagonal matrix of dimension $N + 1$, its i th diagonal element being given by $ir - C$. $R^{(p)}$ is similarly defined, except that its i th diagonal element is given by $ip - C$. The entries in these matrices are the net fluid rates into the buffer for the guaranteed rate ($W(t) > B^*$) and peak rate ($W(t) < B^*$) regimes respectively.

B. Analysis

Our purpose is to find the joint distribution $G_i(x)$ defined as $G_i(x) = \lim_{t \rightarrow \infty} \mathbb{P}(Y(t) = i, W(t) \leq x)$. To do this we determine the Kolmogorov forward equations for

$$\begin{aligned} G^{(p)}_i(t, x) &:= \mathbb{P}(Y(t) = i, W(t) \leq x), \quad 0 \leq x < B^*; \\ G^{(r)}_i(t, x) &:= \mathbb{P}(Y(t) = i, W(t) \leq x), \quad x \geq B^*. \end{aligned}$$

For $x < B^*$, $G^{(p)}_i(t + h, x)$ equals

$$\begin{aligned} &(1 - h(q^{(p)}_{i,i-1} + q^{(p)}_{i,i+1})) G^{(p)}_i(t, x - hr^{(p)}_i) \\ &+ h q^{(p)}_{i-1,i} G^{(p)}_{i-1}(t, x) + h q^{(p)}_{i+1,i} G^{(p)}_{i+1}(t, x) + o(h), \end{aligned}$$

cf. [2]. By taking $h \downarrow 0$ we find, in matrix form,

$$\frac{\partial}{\partial t} \mathbf{G}^{(p)}(t, x) + \frac{\partial}{\partial x} \mathbf{G}^{(p)}(t, x) R^{(p)} = \mathbf{G}^{(p)}(t, x) Q^{(p)},$$

where $\mathbf{G}^{(p)}(t, x)$ is an N -dimensional row vector.

However, for $x > B^*$ the Kolmogorov equations take a less simple form. With $G^{(p)}_i(t, B^* -) := \lim_{x \uparrow B^*} G^{(p)}_i(t, x)$, we find that $G^{(r)}_i(t + h, x)$ equals

$$\begin{aligned} &G^{(r)}_i(t, x - hr^{(r)}_i) \\ &- h(q^{(r)}_{i,i-1} + q^{(r)}_{i,i+1}) (G^{(r)}_i(t, x - hr^{(r)}_i) - G^{(r)}_i(t, B^*)) \\ &- h(q_{i,i-1}^{(*)} + q_{i,i+1}^{(*)}) (G^{(r)}_i(t, B^*) - G^{(p)}_i(t, B^* -)) \\ &- h(q^{(p)}_{i,i-1} + q^{(p)}_{i,i+1}) G^{(p)}_i(t, B^* -) \\ &+ h q^{(r)}_{i-1,i} (G^{(r)}_{i-1}(t, x) - G^{(r)}_{i-1}(t, B^*)) \\ &+ h q^{(r)}_{i+1,i} (G^{(r)}_{i+1}(t, x) - G^{(r)}_{i+1}(t, B^*)) \\ &+ h q_{i-1,i}^{(*)} (G^{(r)}_{i-1}(t, B^*) - G_{i-1}^{(p)}(t, B^* -)) \\ &+ h q_{i+1,i}^{(*)} (G^{(r)}_{i+1}(t, B^*) - G_{i+1}^{(p)}(t, B^* -)) \\ &+ h q^{(p)}_{i-1,i} G^{(p)}_{i-1}(t, B^* -) \\ &+ h q^{(p)}_{i+1,i} G^{(p)}_{i+1}(t, B^* -) + o(h). \end{aligned}$$

After letting $h \rightarrow 0$, this leads to

$$\begin{aligned} \frac{\partial}{\partial t} \mathbf{G}^{(r)}(t, x) + \frac{\partial}{\partial x} \mathbf{G}^{(r)}(t, x) R^{(r)} = & \\ &(\mathbf{G}^{(r)}(t, x) - \mathbf{G}^{(r)}(t, B^*)) Q^{(r)} + \\ &(\mathbf{G}^{(r)}(t, B^*) - \mathbf{G}^{(p)}(t, B^* -)) Q^{(*)} + \mathbf{G}^{(p)}(t, B^* -) Q^{(p)}. \end{aligned}$$

Assuming stationarity, we now set $G^{(r)}_i(t, x) \equiv G^{(r)}_i(x)$ and $G^{(p)}_i(t, x) \equiv G^{(p)}_i(x)$ for $i = 0, \dots, N$, and take all

derivatives with respect to t equal to 0. In matrix form,

$$\frac{d}{dx} \mathbf{G}^{(p)}(x) R^{(p)} = \mathbf{G}^{(p)}(x) Q^{(p)}, \text{ and} \quad (7)$$

$$\begin{aligned} \frac{d}{dx} \mathbf{G}^{(r)}(x) R^{(r)} &= (\mathbf{G}^{(r)}(x) - \mathbf{G}^{(r)}(B^*)) Q^{(r)} \\ &+ (\mathbf{G}^{(r)}(B^*) - \mathbf{G}^{(p)}(B^*-)) Q^{(*)} + \mathbf{G}^{(p)}(B^*-) Q^{(p)}. \end{aligned} \quad (8)$$

Solving (7) is immediate and leads to

$$\mathbf{G}^{(p)}(x) = \sum_{j=0}^N a^{(p)}_j \mathbf{v}^{(p)}_j \exp \left[z^{(p)}_j x \right], \quad (9)$$

where $(z^{(p)}_j, \mathbf{v}^{(p)}_j)$ is an eigenvalue-eigenvector pair of $z^{(p)}_j \mathbf{v}^{(p)}_j R^{(p)} = \mathbf{v}^{(p)}_j Q^{(p)}$, and the $a^{(p)}_j$ are coefficients. The solution of (8) can be found as follows. To deal with the inhomogeneous terms we first differentiate with respect to x , so that we find homogeneous equations for

$$\mathbf{g}^{(r)}(x) \equiv \frac{d}{dx} \mathbf{G}^{(r)}(x).$$

We write down the solution for the resulting system of equations using the spectral method.

In the case that all eigenvalues are different, we find the solution to the differentiated system to be of the form

$$\mathbf{g}^{(r)}(x) = \sum_{j=1}^N \tilde{a}^{(r)}_j \mathbf{v}^{(r)}_j \exp \left[z^{(r)}_j x \right],$$

where $(z^{(r)}_j, \mathbf{v}^{(r)}_j)$ is an eigenvalue-eigenvector pair of $z^{(r)}_j \mathbf{v}^{(r)}_j R^{(r)} = \mathbf{v}^{(r)}_j Q^{(r)}$, and the $\tilde{a}^{(r)}_j$ are coefficients. As $Q^{(r)}$ is a generator, it has an eigenvalue 0, and hence one of the eigenvalues $z^{(r)}_j$ is zero, say $z^{(r)}_{j_*} = 0$, cf. [21]. With this in mind integration immediately yields that $\mathbf{G}^{(r)}(x)$ equals

$$\sum_{j \neq j_*} a^{(r)}_j \mathbf{v}^{(r)}_j \exp \left[z^{(r)}_j x \right] + a^{(r)}_{j_*} \mathbf{v}^{(r)}_{j_*} x + \mathbf{w}, \quad (10)$$

where $a_j^{(r)} = \tilde{a}^{(r)}_j / z^{(r)}_j$ for $j \neq j_*$, $a^{(r)}_{j_*} = \tilde{a}^{(r)}_{j_*}$, and the components w_i of \mathbf{w} are integration constants.

Now the vectors $\mathbf{a}^{(p)}$, $\mathbf{a}^{(r)}$, and \mathbf{w} can be found by considering the following boundary conditions:

(i) $G^{(p)}_i(0) = 0$ for all $i \geq N' + 1$ (i.e., the buffer cannot be empty when it fills), leading to $N - N'$ equations.

(ii) Similarly, $G^{(p)}_i(B^*-) = G^{(r)}_i(B^*)$ for all $i \in \{N'' + 1, \dots, N'\}$. This gives $N + N' - N'' + 1$ equations.

(iii) For all $z^{(r)}_j$ with a non-negative real part, the corresponding $a^{(r)}_j$ is zero, since the $G^{(r)}_i(x)$ should remain bounded for $x \rightarrow \infty$. There are $N'' + 1$ such eigenvalues [21]. Notice that this also entails that the equilibrium distribution of $Y(t)$ is given by \mathbf{w} .

(iv) By letting $x \rightarrow \infty$ in (8), setting the left hand side equal to zero, we find the $N + 1$ global balance equations for $\mathbf{w} = \lim_{x \rightarrow \infty} \mathbf{G}^{(r)}(x)$. What we do here in fact, is to substitute the integrated solution (10) back into the inhomogeneous (undifferentiated) equations (8) to find the integration

constants. From the global balance equations we can derive the N local balance equations for $i \in \{1, \dots, N\}$,

$$\begin{aligned} (N - i + 1)\lambda w_{i-1} &= ir\mu \left(w_i - G^{(r)}_i(B^*) \right) + \\ &C\mu \left(G^{(r)}_i(B^*) - G^{(p)}_i(B^*-) \right) + ip\mu G^{(p)}_i(B^*-). \end{aligned}$$

(v) Finally we normalize: $\sum_{i=0}^N w_i = 1$.

Noticing that there are just as many boundary conditions as coefficients (namely $3N + 3$), we conclude that the system is solvable. We have proved:

Theorem III.1: The above procedure gives the exact solution to the buffer content distribution.

The above solution enables the computation of several key quantities. Denoting the throughput per user by τ , it is straightforward to obtain that $N\tau$ equals

$$\begin{aligned} \sum_{i=0}^N ip G^{(p)}_i(B^*-) &+ \sum_{i=N'+1}^{N''} C(G^{(r)}_i(B^*) - G^{(p)}_i(B^*-)) \\ &+ \sum_{i=0}^N ir(w_i - G^{(r)}_i(B^*)). \end{aligned}$$

The mean file transfer delay $\mathbb{E}T$ is found from

$$\tau = \frac{1/\mu}{\mathbb{E}T + 1/\lambda}. \quad (11)$$

IV. MANY SOURCES

The intrinsic drawback of the technique of the previous sections is its computational complexity. When the size of the system (i.e., the number of sources) grows, a large eigensystem needs to be solved. This explains the interest in simpler asymptotic approaches. In this section we will focus on the so-called ‘many-sources scaling’, which was introduced by Weiss [31]. In this regime, we derive explicit results on the overflow probability.

In the many-sources scaling, buffer and bandwidth resources are scaled by the number of users N . In other words, if we scale $C \equiv Nc$, the exponential decay rate of $\mathbb{P}(W \geq Nx)$ can be determined *explicitly* in terms of x and model parameters r , p , λ , μ , and c . Because W is now implicitly parametrized by N , we write W_N . The random variable V_N is defined as the buffer content in the corresponding model without feedback (as in Section II-B).

A. Feedback model without threshold

Asymptotics. First we estimate the average behavior of the number of active sources in the asymptotic limit (large N). It is not hard to show that there are two regimes.

(A) In the first regime $\lambda p / (\lambda + \mu p) < c$, cf. (1). In this case, in the asymptotic limit, on average the sources are allowed to transmit at peak rate; the buffer will be empty nearly always. The number of active sources on average will be Nm , with $m := \lambda / (\lambda + \mu p)$.

(B) In the second regime, c is between $\lambda r / (\lambda + \mu r)$ and $\lambda p / (\lambda + \mu p)$. In this case, the network will in general be in the

processor sharing regime. The average number of active users simultaneously in the system is Nm with $m := 1 - \mu c \cdot \lambda^{-1}$. The sources are allowed to transmit at a rate m' between r and p , where $m' := c/m = \lambda c/(\lambda - \mu c)$. The following lemma gives the decay rate of the probability of a non-empty buffer in each of the regimes. A similar result can be found in Ramanan and Weiss [23].

Lemma IV.1: The decay rate of the probability of a non-empty buffer is given by

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(W_N > 0) = I(0).$$

When $\lambda p/(\lambda + \mu p) < c$ (regime (A)), $I(0)$ is given by

$$\left(1 - \frac{c}{r}\right) \log \left(\frac{1 - c/r}{c\mu/\lambda}\right) + \log \left(\frac{c(\lambda + p\mu)}{p\lambda}\right) + \frac{c}{r} - \frac{c}{p}$$

and when $\lambda r/(\lambda + \mu r) < c \leq \lambda p/(\lambda + \mu p)$ (regime (B)), $I(0)$ is given by

$$\left(1 - \frac{c}{r}\right) \log \left(\frac{1 - c/r}{c\mu/\lambda}\right) + \frac{c\mu}{\lambda} - \left(1 - \frac{c}{r}\right).$$

Proof. Directly from Theorem 11.15 of [29], the decay rate of the probability of a non-empty buffer equals

$$\int_m^{c/r} \log \left(\frac{\mu_x x}{\lambda(1-x)}\right) dx.$$

Here μ_x is the (downward) probability flux per source, when the number of sources in the on state is Nx :

$$\mu_x := \begin{cases} p\mu & \text{if } xp < c, \\ c\mu \cdot x^{-1} & \text{otherwise.} \end{cases}$$

Direct calculation yields the stated expression. \square

Define $A(t)$ as the amount of fluid generated in the interval $[0, t)$ by one source with off-periods that are i.i.d $\text{Exp}(\lambda)$ random variables, and on-periods that are i.i.d $\text{Exp}(r\mu)$, and constant generation rate r . Let \mathbb{E}_0 (\mathbb{E}_1 , respectively) denote expectation given that the source start in the off (on) state at time 0.

Proposition IV.2: The decay rate for positive buffer content values is given by

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(W_N > Nx) =: I(x) = J(x) + I(0),$$

with $J(x)$ equal to

$$\inf_{t>0} \sup_{\theta} \left(\theta(x + ct) - \frac{c}{r} \log \mathbb{E}_1 e^{\theta A(t)} - \left(1 - \frac{c}{r}\right) \log \mathbb{E}_0 e^{\theta A(t)} \right). \quad (12)$$

Proof. Evidently, $\mathbb{P}(W_N > Nx)$ equals

$$\mathbb{P}(W_N > Nx | W_N > 0) \mathbb{P}(W_N > 0).$$

As shown in Section II-C,

$$\mathbb{P}(W_N > Nx | W_N > 0) = \mathbb{P}(V_N > Nx | V_N > 0).$$

Immediately from Theorem 1 of Duffield [8], we have

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(V_N > Nx | V_N > 0) = J(x).$$

Together with Lemma IV.1 this proves the stated. \square

The variational problem in (12) cannot be solved analytically; numerical methods have to be applied. Fortunately, for large x asymptotics are available.

Simple approximations for large buffers. Let θ^* the equation $\lim_{t \rightarrow \infty} t^{-1} \log \mathbb{E} e^{\theta A(t)} = c\theta$. Define, for $i = 0, 1$,

$$a_i := \lim_{t \rightarrow \infty} \log \mathbb{E}_i e^{\theta^* A(t)} - c\theta^* t.$$

In Duffield [8] it is proven that, for $x \rightarrow \infty$,

$$J(x) = \theta^* x - \frac{c}{r} a_1 - \left(1 - \frac{c}{r}\right) a_0 + o(x).$$

Following the *Chernoff Dominant Eigenvalue method* of [9], we propose an even simpler approximation:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(W_N > Nx) \approx \theta^* x + I(0). \quad (13)$$

Here $\theta^* = r\mu/(r - c) - \lambda/c$, and $I(0)$ is given in Lemma IV.1. In [9] it is shown that this approximation is conservative for all x (in fact it is the best possible linear estimate that is conservative for all x). Notice that the analysis of [9] requires the sources to be time-reversible, which condition is trivially met for exponentially distributed file sizes and user think times. In [9] it is concluded that the approximation is usually not overly conservative.

General think-time and file-size distributions for large buffers The nature of approximation (13) is, for large x ,

$$\mathbb{P}(V_N > Nx) \approx \mathbb{P}(V_N > 0) e^{-\theta^* Nx}. \quad (14)$$

We may follow the same approach in case of general (rather than exponentially distributed) think-time and file-size distributions. The questions then are: how to compute the counterparts of the probability of a non-empty buffer $\mathbb{P}(V_N > 0)$ and the exponential decay rate θ^* ?

The probability of a non-empty buffer is computed as follows. As long as the buffer is empty, the process behaves like an infinite-server queue if the number of jobs is below N' , and like a processor sharing queue if the number of jobs is between N' and $N'' + 1$. Let D_N be the number of jobs in the system. It is easily verified that in the case of exponential think-times and file-sizes, the blocking probability (i.e., the probability of $D_N = N'' + 1$) equals

$$\frac{1}{\text{Norm}} \left(\frac{\mu}{\lambda}\right)^{N''+1} \left(r^{N''-N'+1} p^{N'}\right) \frac{N'! \cdot (N - N'' - 1)!}{N!}, \quad (15)$$

where the normalizing constant Norm is given by

$$\begin{aligned} \text{Norm} &= \sum_{j=0}^{N'} \left(\frac{\mu}{\lambda}\right)^j (p^j) \frac{(N-j)!}{N!} \\ &+ \sum_{j=N'+1}^{N''+1} \left(\frac{\mu}{\lambda}\right)^j \left(r^{j-N'} p^{N'}\right) \frac{N'!(N-j)!}{N!}. \end{aligned}$$

It is tedious but straightforward to verify that the decay rate of (15) is indeed $I(0)$, as was defined in Lemma IV.1.

Importantly, formula (15) also holds in the case of *general* think-time and file-size distributions, with respective means λ^{-1} and μ^{-1} . This is due to *insensitivity* results for networks of generalized processor sharing queues, as was shown by Cohen [5].

We now focus on the exponential decay rate θ^* . Let T be the distribution of the think-time and F the distribution of the file-size. During busy periods, the buffer is fed by a superposition of N on-off sources with off-times distributed like T , and on-times distributed like F/r (with peak rate r). Let $\bar{A}(t)$ be the amount of traffic generated by such a source, in steady state, in an interval of length t . Then (under mild technical conditions) the exponential decay rate in this model is the solution of

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{E} e^{\theta \bar{A}(t)} = c\theta, \quad (16)$$

see Glynn and Whitt [14]. Equation (16) does not necessarily have a solution. In case of heavy-tailed on-times there is no solution – approximation (14) does not apply.

General think-time and file-size distributions for small buffers. Above we concentrated on loss behavior under general think-time (with mean $1/\lambda$) and file-size distributions (with mean $1/\mu$), and large buffers. For small buffers a result from Mandjes and Kim [18] is applicable: for $x \downarrow 0$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(W_N > Nx) = \frac{2\sigma}{r} \sqrt{x} + I(0) + O(x),$$

with constant σ given by

$$\sqrt{(cr\mu + (r-c)\lambda) \log\left(\frac{cr\mu}{(r-c)\lambda}\right) - 2(cr\mu - (r-c)\lambda)}.$$

Strikingly, σ depends on the distributions of the think times and file sizes only through their means, making this an *insensitivity* result.

B. Feedback model with threshold

As before, we scale the resources buffer and bandwidth: $C \equiv Nc$, and $B^* \equiv Nb^*$. Again, this regime allows explicit results, which we describe below.

As in Section IV-A, it turns out that there are two possible regimes: (A) the buffer is empty with an overwhelming probability and the active sources transmit at rate p almost all the time, and (B) the buffer occupancy is approximately nb^* on average, and the active sources send at a rate between r and p .

(A) $\pi^{(p)}p < c$, where $\pi^{(p)}$ is the on-probability in the ‘peak rate regime’: $\lambda(\lambda + p\mu)^{-1}$. In this case the ‘average state’ of the system is a (nearly) empty buffer, and the active sources transmit at peak rate. In order for the buffer to exceed level Nx , where x exceeds b^* , four events have to occur in order: (1) The buffer must become non-empty, i.e., the number of sources in the on-state must exceed Nc/p . (2) Given that the buffer content is at the point of becoming positive, an amount of Nb^* of fluid has to accumulate. At the epoch the buffer content reaches Nb^* , let the number of sources transmitting be $N\alpha$. (3) If α is smaller than c/r , then the number of

sources in the on-state has to grow to Nc/r , in order for the buffer to exceed level Nb^* . If α is already at least c/r then in this phase nothing has to happen. Let $N\alpha'$ be the number of sources in the on-state at the end of this phase. (4) An amount of $N(x - b^*)$ of fluid has to accumulate in the buffer, where at the beginning of this phase the number of sources transmitting is $N\alpha'$. This makes the decay rate be the sum of four decay rates:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(W_N > Nx) \approx \sum_{i=1}^4 I_i. \quad (17)$$

Notice that this construction (the decay rate of a steady-state probability being formulated as the solution of a transient problem, namely the decay rate of the path from equilibrium towards the rare event) is essentially of the Freidlin-Wentzell type. Details on this approach are found in [29, Chapter 6]. In [19] explicit expressions for the four decay rates are given, in terms of the model parameters r , p , c , λ , and μ .

(B) $\pi^{(r)}r < c \leq \pi^{(p)}p$. In this case the ‘average state’ of the system is a buffer occupancy of Nb^* , and all active sources sending at rate $m' = c\lambda/(\lambda - c\mu)$. In equilibrium, the system is operating in the processor sharing regime. Clearly, the first two decay rates of case A do not apply anymore, as the queue is already operating at level nb^* on average. Hence, only the other two decay rates have to be computed. Again [19] provides the explicit expressions.

V. NUMERICAL RESULTS

In this section we provide numerical results. We first consider an example where we maximize the system throughput, using the procedure of Section III. The second example relies on the many sources asymptotics of Section IV.

A. System throughput maximization

We consider a buffer that is fed by 10 identical and independent on-off sources. Let $Y(t)$ denote the number of active sources at time t ; we choose the process $Y(\cdot)$ as the regulating process. As before, $W(t)$ will denote the buffer content at time t . The output trunk speed is $C = 11$. The off-times of the sources are exponentially distributed with $\lambda = 3$. File sizes are exponentially distributed with $\mu = 2$. The peak transmission rate of the sources $p = 4$, which is the actual transmission rate when the buffer content is below $B^* = 2$, and the rate drops to its minimum value $r = 2$ when the content is above B^* . Notice that the buffer content may stick at B^* for as long as 3, 4, or 5 sources are active; then the total actual transmission rate is 11, i.e., the output trunk speed, equally shared among the active sources. As a consequence the transition intensity of the process $Y(\cdot)$ from state i to $i-1$, given that the buffer process remains at 2, is equal to $i \cdot 2C/i = 22$. These considerations allow us to write down the diagonal matrices $R^{(p)}$ and $R^{(r)}$, as well as the tri-diagonal matrices $Q^{(p)}$, $Q^{(r)}$ and $Q^{(*)}$. After the numerical determination of the eigensystems of the matrices $Q^{(p)}(R^{(p)})^{-1}$ and $Q^{(r)}(R^{(r)})^{-1}$, we apply the appropriate boundary conditions to find the 33 unknowns (note that $N = 10$). After solving this (linear) set of equations, the stationary distribution \mathbf{G} of the joint process $(Y(t), W(t))$

can be found numerically. A graphical representation of $G(y) \equiv \mathbb{P}(W \leq y) = \sum_{i=0}^{10} G_i(y)$ is given in Figure 1.




Fig. 1. $B = \infty$: $\log_{10}(1 - G(y))$ as a function of y

The throughput τ of the system can be found as

$$\tau = \sum_{j=0}^{10} jr(p_j - G_j^{(r)}(B^*)) + jpG_j^{(p)}(B^* -) + C(G_j^{(r)}(B^*) - G_j^{(p)}(B^* -)) = 10.2652. \quad (18)$$

Finite-buffer model. The corresponding finite-buffer system can be solved similarly [19]. We did the calculations for $B = 5$, leaving all other parameters the same. Clearly, now the size of the ‘jump’ that $G(y)$ has at $y = 5$ is exactly the (time average) probability of a full buffer:

$$\mathbb{P}(W = B) = 1 - \sum_{j=0}^{10} G_j^{(r)}(B) = 2.01 \cdot 10^{-4}.$$

It is also not difficult to find the average amount of fluid sent into the buffer per unit of time, or fluid input rate τ^* :

$$\tau^* = \sum_{j=0}^{10} jr(p_j - G_j^{(r)}(B^*)) + jpG_j^{(p)}(B^* -) + C(G_j^{(r)}(B^*) - G_j^{(p)}(B^* -)) = 10.2656, \quad (19)$$

and the average amount of fluid sent over the link per unit of time, or throughput, as

$$\tau = \sum_{j=0}^{10} C(p_j - G_j^{(p)}(0)) + jpG_j^{(p)}(0) = 10.2651. \quad (20)$$

Notice that the numerical outcomes for these quantities are close to each other and to (18), which can be explained by the fact that we chose the (finite) size of the buffer quite large. Indeed by subtracting (20) from (19) we immediately find that the average amount of lost fluid per unit of time, or fluid loss rate,

$$\text{Fluid loss rate} = \tau^* - \tau = 5.83 \cdot 10^{-4}, \quad (21)$$

is small. Another way to find this result is to use

$$\text{Fluid loss rate} = \sum_{j=0}^{10} (jr - C)(p_j - G_j^{(r)}(B)). \quad (22)$$

The fraction of fluid that is lost can be found as the ratio of the fluid loss rate (21) and the fluid input rate (19):

$$\text{Fraction of fluid lost} = \frac{\text{Fluid loss rate}}{\text{Fluid input rate}} = 5.67 \cdot 10^{-5}. \quad (23)$$

Maximization of system throughput. As a final application we show how the threshold level B^* and the number of sources N may be jointly chosen such that the system throughput is maximized, with the fraction of lost fluid not to exceed $\epsilon = 10^{-6}$. The other parameters are as before. In Figures 2 and 3 we plot the throughput and the loss fraction as functions of the threshold level for $N = 7, \dots, 12$. From Figure 3 it can be seen that only for $N \leq 9$ the loss fraction criterion can be satisfied for some value of B^* . From Figure 2 we compare the corresponding throughputs for these pairs of (N, B^*) . It turns out that the system throughput is maximized by choosing $N = 9$ and $B^* = 2.05$, giving a throughput of 9.5725. This also allows us to compute the mean file transfer time, see (11): solve $\mathbb{E}T$ from

$$\frac{9.5725}{9} = \frac{0.5000}{\mathbb{E}T + 0.3333},$$

giving $\mathbb{E}T = 0.1368$. During the active period a source’s throughput is $0.5000/0.1368 = 3.6526$, which is between $r = 2$ and $p = 4$.




Fig. 2. Throughput as a function of threshold B^* for $N = 7$ (bottom), \dots , 12 (top)

B. Impact of the choice of the guaranteed rate r

The purpose of this subsection is to provide an illustration of the effect of the feedback mechanisms proposed in a more practical situation. For reasons of convenience, we assume that the number of sources is sufficiently large to use the approximations of Section IV. We consider a 45 Mbit/s link,




Fig. 3. $\log_{10}(\text{loss fraction})$ as a function of threshold B^* for $N = 7$ (bottom), \dots , 12 (top)




Fig. 4. Number of admissible sources as a function of guaranteed rate r .

with a buffer of 10 Mbit. The sources have peak rates of 3 Mbit/s. The users alternately send files (exponentially distributed with mean 100 kbit) and ‘think’ (for an exponentially distributed time with mean 10 seconds). We again require that the loss probability is below $\epsilon = 10^{-6}$. We focus on the impact of r , the guaranteed minimum throughput advertised to customers, because this is an important parameter. Clearly, for a fixed value of B^* (possibly 0), the loss probability increases with r . On the other hand, the number of admissible sources N_r decreases in r , for fixed loss probability. In Figure 4, N_r is given as a function of r , for different values of B^* ($B^* = 0, 2, \dots, 8$ Mbit). We observe that for low values of r , N_r is quite sensitive to the threshold value B^* ; the opposite is true for r in the neighborhood of p .

REFERENCES

[1] I. Adan, E. van Doorn, J. Resing, and W. Scheinhardt. Analysis of a single-server queue interacting with a fluid reservoir. *Queueing Systems*, 29: 313 – 336, 1998.
 [2] D. Anick, D. Mitra, and M. Sondhi. Stochastic theory of a data-handling system with multiple sources. *The Bell System Technical Journal*, 61: 1871 – 1894, 1982.

[3] D. Botvich and N. Duffield. Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. *Queueing Systems*, 20: 293 – 320, 1995.
 [4] G. Choudhury, D. Lucantoni, and W. Whitt. Squeezing the most out of ATM. *IEEE Transactions on Communications*, 44: 203 – 217, 1996.
 [5] J. Cohen. The multiple phase service network with generalized processor sharing. *Acta Informatica*, 12: 245 – 284, 1979.
 [6] C. Courcoubetis and R. Weber. Buffer overflow asymptotics for a buffer handling many traffic sources. *Journal of Applied Probability*, 33: 886 – 903, 1996.
 [7] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Jones and Bartlett, Boston, 1993.
 [8] N. Duffield. Conditioned asymptotics for tail probabilities in large multiplexers. *Performance Evaluation*, 31: 281 – 300, 1998.
 [9] A. Elwalid, D. Heyman, T. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers with applications to videoconferencing. *IEEE Journal on Selected Areas in Communications*, 13: 1004 – 1016, 1995.
 [10] A. Elwalid and D. Mitra. Analysis and design of rate-based congestion control of high-speed networks, I: stochastic fluid models, access regulation. *Queueing Systems*, 9: 29 – 64, 1991.
 [11] A. Elwalid and D. Mitra. Statistical multiplexing with loss priorities in rate-based congestion control of high-speed networks. *IEEE Transactions on Communications*, 42: 2989 – 3002, 1994.
 [12] R. Gibbens and F.P. Kelly. Resource pricing and the evolution of congestion control. To appear in: *Automatica*.
 [13] R. Gibbens and F.P. Kelly. Distributed connection acceptance control for a connectionless network. In: *Proceedings ITC 16: Teletraffic engineering in a competitive world*, eds. P. Key and D. Smith, 941 – 952. Elsevier, Amsterdam, the Netherlands, 1999.
 [14] P. Glynn and W. Whitt. Logarithmic asymptotics for steady-state tail probabilities in a single-server queue. *Journal of Applied Probability*, 31A: 131 – 156, 1994.
 [15] R. Howard. *Dynamic probabilistic systems, Vol. 1*. Wiley, New York, NY, USA, 1971.
 [16] F.P. Kelly. Models for a self-managed Internet. Submitted to: *Royal Society*, 2000.
 [17] L. Kosten. Stochastic theory of data-handling systems with groups of multiple sources. In: H. Rudin and W. Bux (eds.), *Performance of Computer-Communication Systems*, 321 – 331, Elsevier, Amsterdam, the Netherlands, 1984.
 [18] M. Mandjes and J.H. Kim. Large deviations for small buffers: an insensitivity result. *Queueing Systems*, 37: 349 – 362, 2001.
 [19] M. Mandjes, D. Mitra, and W. Scheinhardt. Simple models of network access using feedback fluid queues. See: <http://cm.bell-labs.com/cm/ms/who/michel/pub.html>.
 [20] L. Massoulié and J. Roberts. Arguments in favour of admission control for TCP flows. In: *Proceedings ITC 16: Teletraffic engineering in a competitive world*, eds. P. Key and D. Smith, 33 – 44. Elsevier, Amsterdam, the Netherlands, 1999.
 [21] D. Mitra. Stochastic theory of a fluid model of producers and consumers coupled by a buffer. *Advances in Applied Probability*, 20: 646– 676, 1988.
 [22] R. Pazhyannur and R. Agrawal. Feedback-based flow control of B-ISDN/ATM networks. *IEEE Journal in Selected Areas in Communications*, 13: 1252 – 1266, 1995.
 [23] K. Ramanan and A. Weiss. Sharing bandwidth in ATM. In: *Proceedings Allerton Conference*, 732 – 740, 1997.
 [24] J. Roberts. Engineering for Quality of Service. *Preprint*.
 [25] J. Roberts and S. Oueslati-Boulahia. Quality of Service by flow aware networking. Submitted to: *Royal Society*, 2000.
 [26] W. Scheinhardt. *Markov-modulated and feedback fluid queues*. See: <http://www.ub.utwente.nl/webdocs/tw/1/t0000008.pdf>. Ph.D. Thesis, University of Twente, the Netherlands, 1998.
 [27] W. Scheinhardt. Analysis of feedback fluid queues. In: *Proceedings 14th ITC-specialists seminar on access networks and systems*, 215 – 220, Girona, Spain, April 25-27, 2001.
 [28] M. Schwartz. *Broadband Integrated Networks*. Prentice Hall, Upper Saddle River, NJ, USA, 1996.
 [29] A. Shwartz and A. Weiss. *Large deviations for performance analysis, queues, communication, and computing*. Chapman and Hall, New York, NY, USA, 1995.
 [30] A. Simonian and J. Guibert. Large deviations approximation for fluid queues fed by a large number of on/off sources *IEEE Journal of Selected Areas in Communications*, 13: 1017 – 1027, 1995.
 [31] A. Weiss. A new technique of analyzing large traffic systems. *Advances in Applied Probability*, 18: 506 – 532, 1986.