In Search of the Worst Case Arrivals of Independent Leaky Bucket Constrained Sources revised version*
UCL/CS Research Note 97/69

Philippe Oechslin
Department of Computer Science, University College of London
p.oechslin@cs.ucl.ac.uk

Keywords: leaky bucket, worst case, loss rate, statistical multiplexing

Abstract

In a previous paper [10] we have shown numerical evidence that there are traffic patterns wich generate more loss in a multiplexer than the on-off source. In this paper we derive mathematical upper and lower bound for two loss measures (P_{loss}, the cell loss rate and P_{sat}, the fraction of time where the multiplexer is saturated) in a buffered multiplexer fed by homogeneous leaky bucket constrained sources. Using these bounds we prove that the on-off source is not the worst case traffic pattern. We demonstrate that the so-called sym pattern generates both higher P_{loss} and P_{sat}, at least in the domain of validity of the bounds. We end the paper by a discussion on the implications of this new finding, especially in respect to CAC algorithms which are based on the assumption that on_off patterns are the worst case.

1 Introduction

Much work has been carried out lately on Connection Acceptance Control algorithms for networks with Quality of Service (QoS) guarantees (see [11]). Be it for ATM networks or for the Integrated Services Internet, traffic flows are described by leaky bucket parameters which include mean rate m, peak rate p and maximum burst size b. One practice found in many CAC algorithms is the estimation of the maximum amount of resources (bandwidth or buffer space) a connection with given leaky bucket parameters could utilise. To estimate this utilisation it is necessary to know what the worst case behaviour of a source with given leaky bucket parameters is. Here we have to differentiate between correlated and independent sources.

- If sources are correlated, we have to consider the case where they all send a burst at the same time with the same importance as the case where the bursts are equally distributed. The correlated case, also referred to as deterministic case, has been extensively studied, resulting in the theory of network calculus [2],[3],[6] with known worst case patterns for different study cases [4], [7].
- For the case of independent sources, also called the statistical case, less results have been published. It is a more rewarding case however, since a large multiplexing gain can be achieved if the sources are known to be independent. A CAC algorithm for the independent case can be found in [5]. The same basic idea has been extended in [1] and [12]. In this paper we only consider the case of independent sources.

*The original version of this paper appeared in the 97 UK Workshop on Performance Engineering Of Computer And Telecommunication Systems (UKPEW '97). In this version some formulas have been corrected and numerical results updated accordingly.

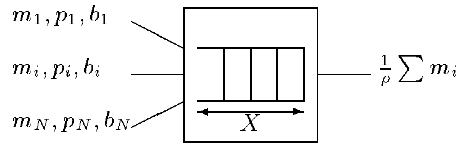

Figure 1: the multiplexer

A common method to analyse the independent case is to use the large deviation theory [8]. It allows to find bounds on probability that a sum of random variables exceeds a given limit. One of these bounds is the Chernoff bound. It is maximised if the random variables only take extreme values. It is thus tempting to believe that on_off sources are the worst case traffic pattern of independent leaky bucket constrained sources, since they maximise this bound. As we will see, the fact that on_off patterns maximise the probability that the instantaneous rate of a sum of leaky bucket constrained sources exceeds a limit does not imply that on_off patterns also maximise the loss rate in a buffered multiplexer.

Our goal is to prove that the on_off pattern is not the worst case pattern in regard both to cell loss rate and saturation probability.

The rest of this paper is organised as follows. Section 2 describes the experiment we simulated and for which we want to find the worst traffic pattern. Section 3 describes the traffic patterns investigated in this paper. Sections 4 and 5 derive bounds for the saturation probability and the cell loss probability respectively and prove that sym pattern are worse than on_off patterns. Finally Section 6 discusses the implications of this findings and concludes the paper.

2 The Experiment

We consider a multiplexer with N incoming connections, all described by their leaky bucket parameters m for mean rate, p for peak rate and b for burst size. The multiplexer has a FIFO buffer of size X and outputs its traffic on a link with capacity $\frac{1}{\rho}\sum m_i$. For reasons of stability ρ must be smaller than 1.

We consider the homogeneous case where all connections have the same leaky bucket parameter. We apply a deterministic periodical pattern to all inputs of the multiplexer and measure its average cell loss rate and saturation probability. The cell loss rate is defined as the amount of data lost due to buffer overflow, divided by the amount of data offered to the multiplexer. The saturation probability is the fraction of time the occupancy of a system with an infinite queue would be larger than the size of the actual system. Since the same pattern is applied to all inputs, the only difference between the traffic in the connections is the phase of the patterns. The connections are assumed to be independent and thus the phases have a uniform random distribution.

The two traffic patterns we want to compare are the on_off and the sym pattern. An on-off pattern consists of a burst at peak with size b, followed by a silence of duration b/m after which a new burst can be sent. The sym pattern looks like an on_off pattern except for two symmetrical plateaus at mean rate before and after the burst. The sym pattern is parametrised by a form factor f which expresses the length of the pattern in units of the on_off pattern. The duration of the burst and silence periods being given, the form factor only influences the length of the plateaus.

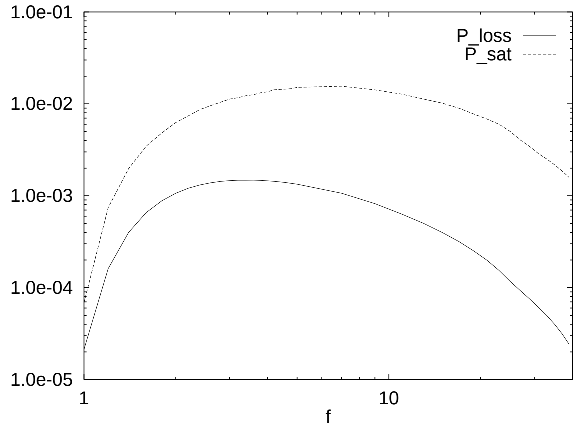

Figure 2: Cell loss rate and saturation probability as a function of f. For $f = 1$ the traffic pattern is an on_off pattern.

on_off

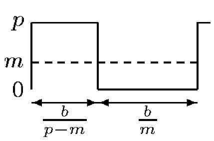

The maximum allowed burst followed by the shortest period of silence allowing a new burst.

sym(f)

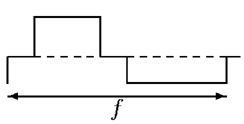

A symmetrical pattern which corresponds to an on-off pattern with two plateaus of the same duration at mean rate inserted before and after the burst. f denotes the length of the pattern in units of the the length of the on_off pattern. Note that sym(1) = on_off

2.1 Simulation Results

As an illustration we present simulation results for a system with $N = 20$, $p = 5$, $m = 1$, $b = 20$, $X = 180$ and $\rho = 0.99$. Figure 2 shows both cell loss rate and saturation probability as a function of f. Remember that for $f = 1$ the pattern is an on_off pattern. We see that as f increases, both measures for loss probability increase and have a maximum for f somewhere between 3 and 10. This shows that sym patterns can be worse than on_off patterns. The next Sections will try to formally prove this numerical result, at least for some limited range of the parameters of the experiment.

3 Saturation Probability P_{sat}

In this section we prove that sym patterns can generate a P_{sat} higher than on_off patterns by deriving a lower bound of P_{sat} of sym and an upper bound on P_{sat} of on_off.

P_{sat} is defined as the fraction of time during which the occupancy of a virtual system with an infinite queue is higher than the queue length of the real system. It is an estimation of the fraction of time during which loss will occur in the real system. Define P^+ as an upper bound on the probability that on_off patterns with a set of random phases generate loss, given buffer size X, number of connections N and the parameters of the connections m, p, b. Define t^+ as the maximum fraction of time during which loss occurs, given that loss does occur. An upper bound for P_{sat} is thus :

$$P_{sat} < P^+ t^+$$

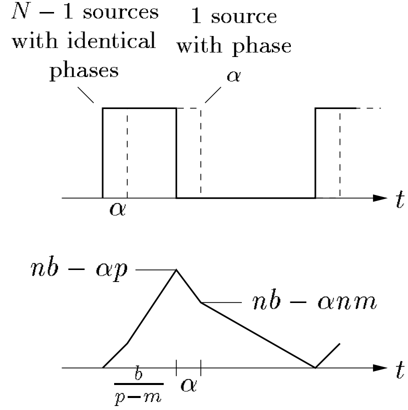

Figure 3: Buffer occupancy in function of time for N sources with phases within α. At the end of the busy period $\frac{b}{p-m}+\alpha$ the buffer occupancy is $Nb-\alpha Nm$. During the busy period a maximum occupancy of $Nb-\alpha p$ can be reached if $N-1$ sources are in phase. Note that $Nb-\alpha p$ is only larger than $Nb-\alpha Nm$ if $p<Nm$

Similarly, define P^- as the a lower bound of the probability that sym patterns with a set of random phases will generate loss, and t^- the minimum time fraction during which loss occurs with probability P^-. P_{sat} is thus lower bounded by:

$$P_{sat} > P^- t^-$$

Our goal is to show that there are cases where

$$P^+t^+ < P^-t^- \quad (1)$$

thus proving that the on_off pattern is not always the worst case in buffered systems. To do this we consider the case of a *full rate multiplexer*, meaning a multiplexer with an output link capacity equal to the sum of input mean rates. Furthermore we analyse P_{sat} for buffer sizes close to the total burst size of all sources. We thus have $\rho=1$ and $X=Nb-\epsilon$ with $\epsilon<b$

3.1 An upper bound for on_off sources

For $\epsilon=0$ on_off patterns can only fill a buffer of size Nb if all sources have exactly the same phase. For a small ϵ loss only occurs if all phases are within small interval α. P^+ is thus $N(\frac{\alpha}{\omega})^{(N-1)}$ where ω is the length of the on_off pattern, equal to $\frac{bp}{m(p-m)}$.

We now try to find the maximum α for which a buffer of $Nb-\epsilon$ can still be filled. As shown on Figure 3, when phases are within an interval α the buffer will fill up to $Nb-Nm\alpha$ during the active period. We thus have $\epsilon=Nm\alpha$ or $\alpha=\frac{\epsilon}{Nm}$. However, buffer occupancy can go above that level during the busy phase. The maximum occupancy happens if all but one sources are in phase. In that case buffer occupancy reaches $Nb-\alpha p$ which is greater than $Nb-Nm\alpha$ when $p<Nm$. ϵ is then equal to αp or $\alpha=\frac{\epsilon}{p}$. Since P^+ is an upper bound we take the maximum of both expressions which is $\alpha=\frac{\epsilon}{\min(p,Nm)}$ and $P^+=N(\frac{\epsilon(p-m)m}{bp\min(p,Nm)})^{(N-1)}$.

$\frac{1}{\omega}(\alpha+\frac{\epsilon}{N(p-m)})$ is a simple upper bound of t^+ since during the active period $b/(p-m)+\alpha$ it takes at least $b/(p-m)-\frac{\epsilon}{N(p-m)}$ time to fill the buffer up to $X-\epsilon$ and loss can only occur once the buffer is full. We thus have:

$$P^+t^+ = N\left(\frac{\alpha}{\omega}\right)^{(N-1)}\frac{1}{\omega}(\alpha+\frac{\epsilon}{N(p-m)})$$

$$P^+t^+ = \left(\frac{\epsilon(p-m)m}{bp\min(p,Nm)}\right)^N(N+\frac{\min(p,Nm)}{p-m}) \quad (2)$$

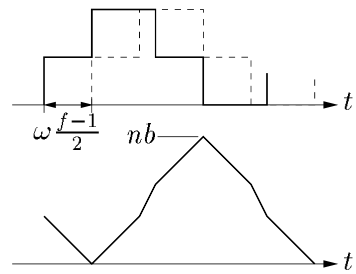

Figure 4: Buffer occupancy resulting of two sym with different phases. As long as the phase difference is smaller than $\omega\frac{f-1}{2}$ the buffer will fill up to Nb

3.2 A lower bound for sym sources

As can be seen from Figure 4, sym patterns can fill a buffer of size Nb only if all phases are within an interval of β which is equal to $\omega\frac{f-1}{2}$, f being the form factor of the pattern, or its length in units of the on_off pattern. For any ϵ, $N(\frac{\beta}{\omega f})^{(N-1)}$ is thus a lower bound for the probability of loss occurring. If all phases are in an interval of β, the buffer will eventually reach Nb. The shortest time to go from $Nb-\epsilon$ to Nb is $\frac{\epsilon}{N(p-m)}$ (all sources at peak rate). The shortest time to empty the buffer from Nb to $Nb-\epsilon$ is $\frac{\epsilon}{Nm}$ (all sources at zero). A lower bound of fraction of time where the buffer occupancy is above $Nb-\epsilon$ is thus $\frac{1}{f\omega}\frac{\epsilon p}{Nm(p-m)}$. We thus have:

$$P^-t^- = N\left(\frac{\beta}{f\omega}\right)^{(N-1)}\frac{\epsilon p}{f\omega Nm(p-m)}$$

$$P^-t^- = \left(\frac{f-1}{2f}\right)^{(N-1)}\frac{\epsilon}{fb} \qquad (3)$$

3.3 Comparing the bounds

To prove that the sym pattern can be worse than the on_off pattern we need to show that there exist cases where $P^-t^- > P^+t^+$ or in other words:

$$\left(\frac{f-1}{2f}\right)^{(N-1)}\frac{\epsilon}{fb} > \left(\frac{\epsilon(p-m)m}{bp\min(p,Nm)}\right)^N(N+\frac{\min(p,Nm)}{(p-m)}) \qquad (4)$$

We can rewrite 4 depending on the result of $\min(p, Nm)$

$$p > Nm: \left(\frac{f-1}{2f}\right)^{(N-1)}\frac{1}{f} \quad > \quad \left(\frac{\epsilon(p-m)}{bpN}\right)^{(N-1)} \qquad (5)$$

$$p \le Nm: \left(\frac{f-1}{2f}\right)^{(N-1)}\frac{1}{f} \quad > \quad \left(\frac{\epsilon(p-m)m}{bp^2}\right)^{(N-1)}\frac{N(p-m)m+pm}{p^2} \qquad (6)$$

We first see that there is still a free parameter f in the left hand side of the equations. We can arbitrarily chose f to be 2, which makes the left hand side $(\frac{1}{4})^{(N-1)}\frac{1}{2}$, a small positive number. Second and more important, we see that the right hand side of the equation can be made arbitrarily small by choosing ϵ small enough (or p large enough or $(p-m)$ small enough). This proves that there are cases in which the lower bound of P_{sat} of a multiplexer fed by sym patterns is higher than an upper bound of P_{sat} of the same system fed by on_off patterns, *QED*.

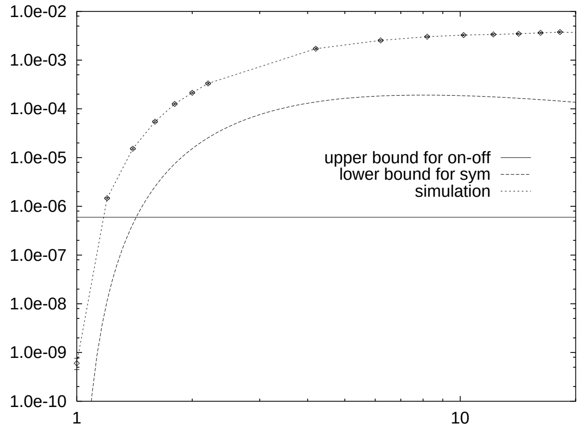

Figure 5: bounds and simulated P_{sat} as a function of f

3.4 Numerical Results

For a numerical validation of the above result we first examine the example simulated in [10] for a full rate multiplexer. The parameters where the following: $N = 20$, $p = 5$, $m = 1$, $b = 20$, and $\rho = 1$. We chose $f = 2$ for the reason cited above and a small ϵ equal to 10 ($X = 390$). We are in the case where $p \leq Nm$ thus we refer to equation 6 and find

$$9.09\ 10^{-13} > 2.45\ 10^{-21}$$

which confirms that at least for small values of ϵ the sym pattern produce a higher saturation probability than on_off patterns. Unfortunately, the values are too small to be verified by simulation. However we can chose example for which it will be possible. Take $N = 8$, $p = 2$, $m = 1$, $b = 20$, and $\rho = 1$. We plot P_{sat} in function of f for $\epsilon = 10$ ($X = 150$).

Taking equation 4 and replacing the given parameters we find

$$(\frac{f-1}{2f})^7 \frac{1}{2f} > 5.96\ 10^{-7}$$

The right hand side is the upper bound for P_{sat} for on_off patterns while the left hand side is a lower bound of P_{sat} for sym patterns in function of the form factor f. Figure 3.4 shows both bounds as a function of f as well as simulation results of P_{sat} for this example. We see that for f between 1.5 and 20, the lower bound for sym patterns is higher than the upper bound of on_off patterns, proving that sym pattern can be worse than on_off patterns. The result is confirmed by the simulation which shows P_{sat} to be 6 10^{-10} for the on_off pattern and 3 10^{-3} for sym pattern with form factor of 10.

4 Cell Loss Rate P_{loss}

Having proven that on_off patterns are not the worst case in a buffered multiplexer with regard to the saturation probability, we now try to prove the same for the cell loss rate P_{loss}. The cell loss rate is defined as the amount of data lost due to buffer overflow divided by the amount of data offered.

Again, we identify an upper bound for P_{loss} of the on_off pattern and a lower bound for P_{loss} of the sym pattern. We decompose P_{loss} into P^+l^+ and P^-l^- accordingly, where P is a bound on the probability that loss occurs and l is a bound on the loss rate, given that loss occurs.

4.1 An upper bound for on_off sources

Using the same reasoning as for P_{sat} we find the same probability P^+ that loss can occur, namely $P^+ = N(\frac{\epsilon(p-m)m}{bp\min(p,Nm)})^{(N-1)}$.

An upper bound of the loss rate that can occur is $\frac{\epsilon}{B}$, where $B = \frac{Nbp}{(p-m)}$ is the total amount of data transmitted during one period. We thus have:

$$P^+l^+ = \left(\frac{\epsilon(p-m)m}{bp\min(p,Nm)}\right)^{(N-1)} \frac{\epsilon(p-m)}{bp} \tag{7}$$

4.2 A lower bound for sym sources

Again, using the same reasoning as with P_{sat} we know that $N(\frac{f-1}{2f})^{(N-1)}$ is a lower bound of the probability that loss occurs. In fact it is the probability that a loss of exactly ϵ occurs, thus $\frac{\epsilon}{B}$ is a lower bound for the amount of loss. Note that for sym patterns, the amount of data transmitted per period is $B = \frac{fNbp}{(p-m)}$ We have:

$$P^-l^- = \left(\frac{f-1}{2f}\right)^{(N-1)} \frac{\epsilon(p-m)}{fbp} \tag{8}$$

4.3 Comparing the bounds

For P^-l^- to be greater than P^+l^+ the following condition must hold:

$$\left(\frac{f-1}{2f}\right)^{(N-1)} \frac{\epsilon(p-m)}{fbp} > \left(\frac{\epsilon(p-m)m}{bp\min(p,Nm)}\right)^{(N-1)} \frac{\epsilon(p-m)}{bp} \tag{9}$$

Decomposing and simplifying for both outcomes of $\min(p, Nm)$ we find:

$$p > Nm: \left(\frac{f-1}{2f}\right)^{(N-1)} \frac{1}{f} > \left(\frac{\epsilon(p-m)}{bpN}\right)^{(N-1)} \tag{10}$$

$$p \leq Nm: \left(\frac{f-1}{2f}\right)^{(N-1)} \frac{1}{f} > \left(\frac{\epsilon(p-m)m}{bp^2}\right)^{(N-1)} \tag{11}$$

Again we can chose f to be 2, such that the left hand side becomes $(\frac{1}{4})^{(N-1)}\frac{1}{2}$, a small positive number. The right hand side can be made arbitrarily small by choosing ϵ small enough, thus proving that there are cases where on_off patterns can not generate a higher cell loss rate than sym pattern, *QED*.

4.4 Numerical Results

Again we first look at the example simulated in [10] for a full rate multiplexer with parameters $N = 20$, $p = 5$, $m = 1$, $b = 20$, $\rho = 1$, $f = 2$ and $\epsilon = 10$. Note that $p \leq Nm$, we thus we refer to equation 11 and find

$$7.62\ 10^{-6} > 5.76\ 10^{-22}$$

which confirms that at least for small values of ϵ the sym pattern produce a higher cell loss rate than on_off patterns. The values being too small to be verified by simulation we turn to the example of Section 3.4 with $N = 8$, $p = 2$, $m = 1$, $b = 20$, and $\rho = 1$. We plot the cell loss rate in function of f for $\epsilon = 10$

Taking equation 9 and replacing the given parameters we find

$$(\frac{f-1}{2f})^7 \frac{1}{4f} > 1.19\ 10^{-7}$$

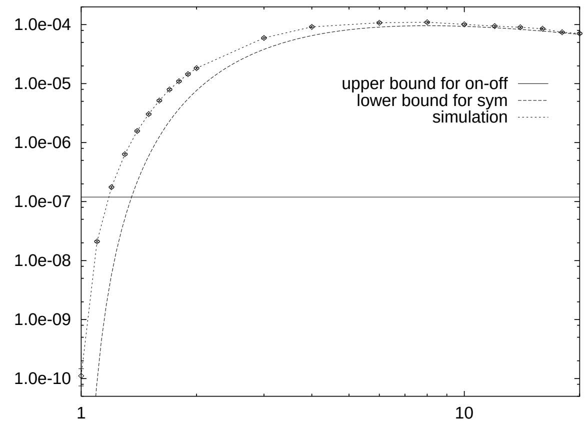

Figure 6: bounds and simulated cell loss rate as a function of f

The right hand side is the upper bound for P_{loss} for on_off patterns while the left hand side is a lower bound of P_{loss} for sym patterns in function of the form factor f. Figure 6 shows both bounds as a function of f as well as simulation results of P_{loss} for this example. We see than for f between 1.5 and 20, the lower bound for sym patterns is higher than the upper bound of on_off patterns, proving that sym pattern can be worse than on_off patterns. The result is confirmed by the simulation which shows P_{loss} to be $1\ 10^{-10}$ for the on_off pattern and $1.0\ 10^{-4}$ for sym pattern with form factor of 10.

5 Conclusions

We have given a formal proof that the on_off pattern is not the worst case pattern for independent leaky bucket constrained sources feeding a buffered multiplexer. We have proven that the sym pattern can generate more loss, be it measured as cell loss rate or probability of saturation. This result stands in opposition with the general belief and a formal proof ([9]) that on_off is the general worst case. Here is an explanation for the diverging results:

The proof in [9] proves that on_off is the worst case a bufferless multiplexer by means of the large deviation theory. This is a result of large deviation theory and is one of the reason for the general belief that on_off sources are a general worst case. The proof is extended to the bufferful multiplexer by using a transformation of a bufferful multiplexer into an equivalent bufferless one described in [5]. This transformation is very appealing since it reduces a two resource problem (buffer and bandwidth) into a single resource one. Without going into the details of that transformation we note that it hinges on a decomposition of the multiplexer into virtual trunk/buffer systems which are all allocated slices of the total bandwidth and buffer space. The buffer occupancy in each trunk/buffer system is then upper bounded by its maximal value during the periods where the buffer is not empty. This procedure is conservative in that the set of virtual trunk/buffer systems with bounded buffer usage can only generate more loss that the actual multiplexer. We conjecture that the on_off source may indeed maximise the loss in this conservative virtual systems but since the loss in the actual multiplexer can be smaller, another traffic pattern may still generate more loss than on_off in the real system. Because of this conservative upper bound one can not use the transformation of the buferfull multiplexer into a bufferless one to prove that bufferfull and bufferless multiplexers have the same worst case traffic pattern.

Being conservative, the decomposition into virtual trunk/buffer system is still a valid method for Connection Acceptance Control, even if it is not based on the actual worst case pattern. Building a CAC algorithm on sym patterns would be much harder since we would first have to devise a method of finding the f wihch maximises the loss probability.

Our proof is done by showing the existence of counter examples in a particular range of parameters (capacity of the multiplexer equal to the sum of mean input rates, large buffer). However, this does not imply that sym patterns are only worst than on_off patterns in these situations. It is rather a consequence of our upper and lower bounds having been derived and being valid in this range. Simulations show that sym patterns generate more loss in a much larger range of the parameters of the experiment. There are also parameters for which on_off actually is the worst case. There is no contradiction between this and our results since the on_off pattern is just a sym pattern where f is equal to 1. Interesting further work would be to identify the range of parameters for which on_off really is the worst case.

References

[1] J.-Y. Le Boudec and A. Ziedins. A CAC algorithmm for VBR connections over a VBR trunk. In *Proceedings for the 15th International Teletraffic Congress*, 1997. to appear.

[2] Rene L. Cruz. A calculus for network delay, part I: Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.

[3] Rene L. Cruz. A calculus for network delay, part II: Network analysis. *IEEE Transactions on Information Theory*, 37(1):132–141, January 1991.

[4] B. T. Doshi. Deterministic rule based traffic descriptors for broadband ISDN: Worst case behavior and connection acceptance control. In J. Labetoulle and J. W. Roberts, editors, *Proceedings of ITC 14*, pages 591–600. Elsevier Science B. V., 1994.

[5] A. Elwalid, D. Mitra, and R. Wentworth. A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6):1115–1127, August 1995.

[6] Jean-Yves Le Boudec. Network calculus made easy. Technical Report 96/218, EPFL/DI, Ecole Polythechnique Federale de Lausanne, December 1996.

[7] D. C. Lee. Effects of leaky bucket parameters on the average queueing delay: Worst case analysis. In *Proceedings of IEEE Infocom '94*, pages 482–489, 1994.

[8] J. Lewis and R. Russell. An introduction to large deviations. http://stp01.stp.dias.ie/ftpdir/LDtut96.ps, 1996.

[9] D. Mitra and J. Morrison. Multiple time scale regulation and worst case processes for ATM network control. In *Proceedings of the 34th Conference on Decision and Control CDC*, pages 353–358, 1995.

[10] Ph. Oechslin. Worst case arrival of leaky bucket constrained sources: The myth of the on-off source. In *Proceedings of the fifth IFIP International Workshop on Quality of Service*, pages 67–77, Columbia University, New York, USA, May 1997.

[11] H. G. Perros and K. M. Elsayed. Call admission control schemes: A review. *IEEE Communications Magazine*, 34(11):82–91, November 1996.

[12] F. Presti, Z. Zhang, J. Kurose, and D. Towsley. Source time scale and optimal buffer/bandwidth trade-off for regulated traffic in an ATM node. In *Proceedings of IEEE Infocom '97*, 1997.