

DESIRABLE PROPERTIES OF TRAFFIC DESCRIPTORS FOR ATM CONNECTIONS IN A BROADBAND ISDN

Arthur W. Berger

AT&T Bell Laboratories, Room 1K-211, 101 Crawfords Corner Road, Holmdel, NJ 07733,
U.S.A.

For future B-ISDNs using Asynchronous Transfer Mode, traffic descriptors are investigated that provide the network operator with relevant information for connection admission control and resource management. For a traffic descriptor defined in terms of an algorithm that is a generalization of the well-known leaky-bucket algorithm, we show conditions on the algorithm that yield desirable properties, such as the existence of simple adjustments to the parameter values of the algorithm to account for cell delay variation in intermediate networks, and the conforming throughput being dependent on the burstiness of the source traffic.

Key words: traffic characterization, connection traffic descriptors, leaky-bucket algorithm, B-ISDN, ATM, connection admission control

1. INTRODUCTION

The description or characterization of traffic in future Broadband Integrated Services Digital Networks (B-ISDNs) using Asynchronous Transfer Mode (ATM) is an area of active research. One can roughly partition the work into two categories: (1) stochastic models, and (2) deterministic algorithms (and the interest of the present paper falls into the second category). These two categories of models address different needs, though there is much potential for cross fertilization of ideas. A stochastic model is typically tailored to a class of traffic (or a service type), such as variable rate video, or inter-LAN data traffic, e.g. see [1] or [2]. These models have many uses including guidance for the design of equipment and for the operation of networks. However, stochastic models are awkward for standardized traffic descriptors, where one wishes to specify a partial characterization that will pertain for, and during, a given realization of the traffic on a connection. In particular, requirements on a standardized traffic descriptor are: (1) that it be specified prior to the initiation of the cell flow, (2) that at all times during the connection, both the user and network operator can determine whether the submitted cell flow is in conformance with the traffic descriptor, and (3) that the traffic descriptor provide the network operator with relevant information for connection admission control and resource management, enabling the network to operate at greater useful utilization. One direction for research on standardized traffic descriptors is to begin with stochastic models that address requirement (3) and then to derive deterministic algorithms that approximate the relevant aspects of the stochastic model. For example, this direction for research could be pursued by the current work on alternative definitions of

"equivalent bandwidth" and the conditions for which they are effective or non-effective (e.g. see [3] [4] [5] [6] [7]). In contrast, herein we start with requirement (2) and note that the satisfaction of this requirement is greatly aided if the traffic descriptor is defined by an algorithm. Thus, we begin with algorithmically defined traffic descriptors, which we call Connection Traffic Descriptors (CTDs), and are interested in algorithms that are useful in satisfying requirement (3). Initial work on this topic is reported in [8]. A long-term goal for the work is its unification with the work that begins with requirement (3).

In this paper we examine a class of algorithms that guarantee that the CTD has properties that are useful in satisfying requirement (3) above. The class of algorithms we examine is a generalization of the leaky-bucket algorithm (a.k.a. the virtual-scheduling algorithm), which has been used by the Telecommunication Standardization Sector (formerly CCITT) to specify the peak cell rate and the associated tolerance for cell delay variation at a reference interface point, [9]. The leaky-bucket algorithm has also been used by the ATM Forum, an industry forum, to specify a "sustainable cell rate," which is an upper bound on the average rate of a connection, [10]. (The term "connection traffic descriptor" used herein was first used by the ATM Forum in [10].) The present work on more general CTDs is at a preliminary stage and does not propose any particular algorithms. However, the results herein are suggested for consideration in future work on algorithmic traffic descriptors.

1.1 Background

Broadband ISDNs using ATM are envisioned to support a wide variety of services, including services that provide transport for voice, image, video and data. Depending on the purpose of a given ATM connection, a user (person or terminal equipment) will be able to inform the public network operator to varying degrees of detail about the character of cell flow to be emitted on the connection. Roughly speaking, the more detail that users can specify about the traffic, the more efficiently the network can be operated. In addition, although tariff structures are unknown, users that can specify their traffic in greater detail may incur a lower cost for the connection, other factors being equal, such as requested quality of service for the connection.

For some applications, users will not know, or will prefer not to have to determine, anything about the traffic to be emitted, except the access-line speed. Other users may be able to commit to a peak rate that is below the access-line speed. Still other users may be able to specify more detail than the peak rate (particularly if the end-system or intervening customer premise equipment is capable of buffering and shaping the cell flow). The present paper is concerned with this last category of users.

1.2 Outline of the Paper

Section 2 contains the formulation of the problem. Section 3 then discusses basic properties that the CTD should have for constant bit rate sources. Section 4 examines properties such that the parameter values of the CTD may be easily adjusted to account for delay variation in intermediate networks. Section 5 considers the dependence of the conforming throughput on the burstiness of the source. The paper concludes with a brief summary in Section 6.

2. PROBLEM FORMULATION

The CTD is a conformance testing rule. Each cell that passes a reference interface point between customer-premise equipment and a network or between two networks is either conforming or not conforming to the CTD. We call the passage of the first bit of a cell across the reference interface point the "arrival time" of the cell. Let A_k denote the arrival time of the k^{th} cell of a given connections.

The algorithm or rule that determines conformance may be quite general and conceptually may depend on all previous cell arrival times of the given connection. Herein, we consider a class of algorithms where the conformance of the k^{th} cell depends on the past history of the cell arrival times via a "state" variable, denoted x_k , and on the interarrival time of the k^{th} and $k-1^{\text{st}}$ cells. We consider algorithms that have the form where at the arrival of the k^{th} cell, x_k is updated and compared against a threshold value. If x_k is below, or equal to, the threshold, then the cell is conforming. Otherwise, the cell is nonconforming and x_k is possibly given an alternative value. In particular, we consider algorithms of the following form. At the arrival of the k^{th} cell, first update x_k according to:

$$x_k = F(x_{k-1}, A_k - A_{k-1}), \quad x_0 \text{ and } A_0 \text{ given.} \quad (2.1)$$

Then, test for conformance: If $x_k \leq M$, then the k^{th} cell is conforming. Otherwise, the k^{th} cell is nonconforming, and x_k is set equal to $G(x_{k-1}, A_k - A_{k-1})$, which may or may not be the same function as $F(,)$.

Given the above structure, the choice of the functions $F(,)$ and $G(,)$, and the choice of the threshold M , along with the initial values x_0 and A_0 , determine the algorithm. In this paper, we focus on the function $F(,)$ and investigate conditions on $F(,)$ that yield desirable properties on the class of sources that are conforming. Herein, for convenience, we say a source is conforming if and only if all of the cells of the source are conforming. Thus, we do not address the important, practical issue that a network operator may choose to consider a user's connection to be conforming even though "a few" of the cells are nonconforming.

A special case of the above algorithm is:

$$F(x_{k-1}, A_k - A_{k-1}) = [ax_{k-1} - (A_k - A_{k-1})]^+ + f(A_k - A_{k-1}), \quad \text{and} \quad (2.2)$$

$$G(x_{k-1}, A_k - A_{k-1}) = [ax_{k-1} - (A_k - A_{k-1})]^+, \quad x_0 \text{ and } A_0 \text{ given,}$$

where $[z]^+ = \max(0, z)$ and $f()$ is some function. When $f()$ is simply a constant and $a = 1$, then (2.2) becomes the well-known leaky-bucket algorithm. For arbitrary $f()$ and for $a \in (0, 1]$, we call (2.2) the "generalized leaky-bucket algorithm." Note that a sequence $\{x_k\}$ given by $F(,)$ in (2.2) where $a = 1$ can be thought of as the content of the leaky bucket (or the work in system) just after the k^{th} arrival, where the k^{th} arrival adds $f(A_k - A_{k-1})$ units of work and where work is processed at the rate of one unit per time unit, given that there is work in the system. Moreover, for the case where $a < 1$, $\{x_k\}$ can still be thought of as the work in system at the k^{th} arrival after $f()$ units of work have been added, where the work is processed in a nonstandard manner. In particular, just after the k^{th} arrival, the work in system jumps from x_k to ax_k , and then, until the $k+1^{\text{st}}$ arrival, the work is processed at the rate of one

unit per time unit, given that there is work in the system.

As an aside, note that the state variable x_k in (2.1) and (2.2) is updated at every arrival. In contrast, the definition of the leaky-bucket algorithm in [9] and [10] is phrased such that the state is updated only for cells that are determined to be conforming. However, the definition of the leaky-bucket algorithm herein is equivalent to the definition in [9] and [10] in the sense that for any sample path of arrival times $\{A_k\}$, both definitions determine the same cells to be conforming and hence the same cells to be nonconforming.

3. CONSTANT-BIT-RATE SOURCES

To begin the study of desirable properties of the CTD, we consider the special case of sources with equal spacing: $A_k - A_{k-1}$ equals a constant for all $k \geq 1$, and $A_0 = 0$. We call such sources constant bit rate (CBR) sources. Note that for any choice of the equal spacing, $A_k - A_{k-1} = T$, and for $\{x_k, k \geq 1\}$ given by $x_k = F(x_{k-1}, T)$ and $x_0 \leq M$, then either $x_k \leq M$ for all $k \geq 1$, or for some k , $x_k > M$. Thus, for each value of T , the CBR source is either conforming or not. A natural, desirable property for the CTD is that if a source with equal spacing T is conforming, then all sources with equal spacing greater than T are also conforming. Likewise, if a source with spacing T is nonconforming, then a desirable property for the CTD is that all sources with spacing less than T are also nonconforming. We show that the CTD has these properties if $F(x, t)$ is nondecreasing in x and nonincreasing in t . We begin with the following lemma.

Lemma 3.1: Consider two CBR sources, one with spacing T^1 and one with spacing T^2 where $T^1 > T^2$. If $F(x, t)$ is nondecreasing in x and nonincreasing in t , then:

$$x_k^1 \leq x_k^2 \quad k \geq 0,$$

where $\{x_k^i\}$ is given by $x_{k+1}^i = F(x_k^i, T^i)$, $i = 1, 2$, $k \geq 0$, and $x_0^1 = x_0^2$.

The proof is by induction. Noting that $x_0^1 \leq x_0^2$, suppose $x_j^1 \leq x_j^2$, and since $F(x, t)$ is nondecreasing in x and nonincreasing in t , then

$$x_{j+1}^1 = F(x_j^1, T^1) \leq F(x_j^2, T^1) \leq F(x_j^2, T^2) = x_{j+1}^2.$$

Lemma 3.1 and the definition of a conforming source imply:

Lemma 3.2: If $F(x, t)$ is nondecreasing in x and nonincreasing in t , then: (1) if a CBR source with spacing T is conforming, then all CBR sources with spacing greater than T are also conforming, and (2) if a CBR source with spacing T is nonconforming then all CBR sources with spacing less than T are also nonconforming.

A more complete statement is given in the following theorem.

Theorem 3.1: If $F(x, t)$ is nondecreasing in x and nonincreasing in t , then for a given M and $x_0 \leq M$ either:

1. all CBR sources are conforming, or
2. all CBR sources are nonconforming, or
3. there exists a positive constant, \hat{T} , such that CBR sources with spacing greater than \hat{T} are conforming and CBR sources with spacing less than \hat{T} are nonconforming.

The proof of Theorem 3.1 and of subsequent theorems has been omitted due to space limitations. However, for the interested reader, the proofs are available in [11].

Remark: Regarding the generalized leaky bucket, (2.2), if $F(x,t)$ has the form $[ax - t]^+ + f(t)$, $a > 0$, and if $f(t)$ is nonincreasing in t , then $F(x,t)$ is nondecreasing in x and nonincreasing in t and Theorem 3.1 pertains.

4. ACCOUNTING FOR DELAY VARIATION IN INTERMEDIATE NETWORKS

For ATM connections that span multiple networks, a desirable quality for a CTD is that if a traffic stream at the input to a network is conforming to a CTD with given parameter values and if the traffic stream is subject to variable delay within the network and if some bounds or characterization of the possible delay is known to the network operator, then the network operator can easily determine a CTD that the traffic stream is conforming to when the traffic exits the network (and either enters another public or private network or enters the destination terminal). Ideally, the CTD for the exiting traffic has the same form as the CTD for the input traffic and the only change is in the value of a few parameters of the CTD. Furthermore, it is desirable that the change in the values is easy to determine. In this section, we show conditions on $F(,)$ and $f(,)$ of the CTDs in equations (2.1) and (2.2) such that only the threshold parameter, M , needs be updated and such that the updated value is easy for the network to determine, given bounded delays within the network.

Note that Cruz in [12] and [13] develops a calculus for network delays where the arrival processes are characterized in a manner that can be interpreted as a particular CTD - one that is isomorphic to the standard leaky-bucket algorithm. The present section is in the same spirit as Cruz's work, except herein the CTD is more general and the network is not modeled in any detail.

Consider a traffic flow on an ATM connection that traverses multiple networks. Let $\{A_k, k \geq 1\}$ be the times that cells of this connection pass an interface reference point at the input to one of these networks, and suppose these emission times are conforming to the CTD (2.1). Thus,

$$x_k \leq M, \quad k \geq 1, \quad (4.1)$$

where,

$$x_{k+1} = F(x_k, A_{k+1} - A_k) \quad k \geq 0, \quad x_0 = A_0 = 0. \quad (4.2)$$

The emitted traffic then passes through the given network. Let d_k denote the delay through the network of the k^{th} cell. The only assumption made about the network is that the delay of any cell on the connection is bounded, and in particular is within the interval:

$$d_k \in [d_{\min}, d_{\max}], \quad k \geq 1. \quad (4.3)$$

Now consider the process of cells exiting the network. Let $\{B_k, k \geq 1\}$ denote the times the cells of the given connection exit the network and cross an interface reference point. Then:

$$B_k = A_k + d_k, \quad k \geq 1. \quad (4.4)$$

Lastly, consider the process $\{B_k, k \geq 1\}$ as input to the function $F(\cdot, \cdot)$ with state variable z_k :

$$z_{k+1} = F(z_k, B_{k+1} - B_k) \quad k \geq 0, \quad z_0 = B_0 = 0. \quad (4.5)$$

We are interested in a bound on $\{z_k, k \geq 1\}$, and in particular the properties of $F(\cdot, \cdot)$ that guarantee such a bound. As a first step, we have the following result.

Theorem 4.1: If $F(\cdot, \cdot)$ is nondecreasing in its first argument and nonincreasing in its second argument and if $F(\cdot, \cdot)$ satisfies a Lipschitz condition for each of its arguments, with Lipschitz constant a for the first argument and constant b for the second argument, then for delays $\{d_k\}$ satisfying (4.3)

$$z_k \leq x_k + b(d_{\max} - d_{\min}) \cdot \sum_{i=0}^{k-1} a^i, \quad k \geq 1. \quad (4.6)$$

Theorem 4.1 implies the following corollary.

Corollary 4.1: If in addition to the conditions of Theorem 4.1, a is less than 1, then

$$z_k \leq x_k + \frac{b}{1-a}(d_{\max} - d_{\min}), \quad k \geq 1.$$

Note that the additional condition in corollary 4.1 makes $F(\cdot, \cdot)$ a contraction mapping with respect to its first argument.

Since the arrival process $\{A_k, k \geq 1\}$ is given to be conforming, then we know from (4.1) that x_k is less than or equal to M for $k \geq 1$. Thus, from Corollary 4.1 we obtain a bound on z_k that does not depend on k , and we label the result a theorem because of its applied interest.

Theorem 4.2: If the input process $\{A_k\}$ is conforming to the CTD with given $F(\cdot, \cdot)$ and M and if $F(\cdot, \cdot)$ is nondecreasing in its first argument and nonincreasing in its second argument and if $F(\cdot, \cdot)$ satisfies a Lipschitz condition for each of its arguments, with Lipschitz constant $a < 1$ for the first argument and with constant b for the second argument, then for delays $\{d_k\}$ satisfying (4.3)

$$z_k \leq M + \frac{b}{1-a}(d_{\max} - d_{\min}), \quad (4.7)$$

and thus the output process $\{B_k\}$ is also conforming to the CTD where the threshold M is replaced with $M + \frac{b}{1-a}(d_{\max} - d_{\min})$.

Theorem 4.2 has the following applied interest. Suppose a source requests to set up a connection and signals given values for the parameters of $F(,)$ and a given value for M , and the connection is to traverse multiple networks. When the first network signals the requested parameter values to the second network, the first network knows that if it increases the requested value of M by $\frac{b}{1-a}(d_{\max} - d_{\min})$ then the second network has better information on how "stressful" the incoming, conforming traffic could be and also the second network will not judge originally conforming traffic to be nonconforming. (In practice, when a signaling message is passed from one network to another, it might contain both the original parameter values requested by the user and the modified values that apply at the given interface.)

Although Theorem 4.2 provides a bound on $\{z_k\}$, an open issue is how tight is the bound for "typical" cases.

4.1 Generalized Leaky Bucket

Suppose $F(x, t)$ has the form:

$$F(x, t) = [ax - t]^+ + f(t), \text{ where } a \in (0, 1]. \quad (4.8)$$

Note that with respect to x , $F(x, t)$ satisfies a Lipschitz condition where the minimal value for the Lipschitz constant is a . Thus:

Corollary 4.2: If $f(t)$ in (4.8) is nonincreasing and satisfies a Lipschitz condition with constant c , then $F(x, t)$ in (4.8) is nondecreasing in x and nonincreasing in t and satisfies a Lipschitz condition for each of its argument with Lipschitz constants a and $c + 1$ respectively, and thus Theorem 4.1 pertains. Moreover, if $a < 1$, then Theorem 4.2 pertains as well.

The case where $a = 1$ is also of interest, and from Theorem 4.1 we have a hint that there may not be a finite bound on $\{z_k\}$ that is independent of k , unless we require "extra" conditions on $f(t)$. Indeed, if one makes the simple and natural assumption that $f(t)$ is nonincreasing and satisfies a Lipschitz condition with constant c , then one can still construct examples where $\{z_k\}$ grows without bound. For example, consider a constant bit rate (CBR) connection where $A_{k+1} - A_k$ equals some constant, T , for $k \geq 0$. Suppose $f(t)$ is chosen to have a fixed point at T and further

$$f(t) > T \quad \text{for } t < T, \text{ and} \quad (4.9)$$

$$f(t) = T \quad \text{for } t \geq T.$$

(One can easily choose an $f(t)$ that satisfies (4.9) and satisfies a Lipschitz condition.) Suppose further that the traffic passes through a network that introduces non-constant delay and for simplicity of construction of this example suppose the delays alternate between two values. In particular, suppose $d_{2k-1} = d_1$ and $d_{2k} = d_2$ for $k \geq 1$, where $d_1 \geq d_2$. Then one can show that after a cycle of two inter-exit times from the network, z_k increases by a constant amount. In particular, $z_1 = T$ and

$$z_{2k+1} - z_{2k-1} = f(T - (d_1 - d_2)) - T, \quad k \geq 1,$$

where from (4.9) $f(T - (d_1 - d_2)) - T > 0$. Thus, $\{z_k\}$ grows without bound.

As an aside, note that although the choice of alternating values for the delay is made to simplify the construction of this example, the choice is also a rough abstraction from plausible scenarios. For example, suppose the only significant variation in delay for the given connection occurs from the queueing delay in one of the buffers that the cells path through. Suppose further that multiple, CBR connections, with random initial phase shifts, pass through this buffer. Then the buffer can be modeled as a discrete time $\sum_i D_i/D/1$ queue, for which the queue length process is periodic. Moreover, if a particular CBR connection has an inter-cell spacing equal to one half the period of the queue length process, then the delays for cells of this connection can oscillate between two values.

We can obtain positive results if additional conditions are placed on $f(\cdot)$. In particular:

Theorem 4.3: If $F(x,t)$ has the form:

$$F(x,t) = [x - t]^+ + f(t) = [x - t]^+ - ct + e, \quad (4.10)$$

where $c \geq 0$, then for $\{d_k\}$ satisfying (4.3):

$$z_k \leq x_k + (c + 1)(d_{\max} - d_k). \quad (4.11)$$

Note that for the special case of $c=0$, (4.10) becomes the standard, leaky-bucket algorithm where e is the constant increment I in [10]. Thus,

Corollary 4.3: Under the conditions of Theorem 4.3 and for the leaky-bucket algorithm, where $F(x,t) = [x - t]^+ + I$,

$$z_k \leq x_k + d_{\max} - d_k, \quad k \geq 1.$$

Theorem 4.3 and Corollary 4.3 each imply a bound on z_k that is independent of k , which for the importance case of the standard leaky-bucket algorithm we explicitly state in the following theorem.

Theorem 4.4: If the input process $\{A_k\}$ is conforming to the CTD with the leaky-bucket algorithm, $F(x,t) = [x - t]^+ + I$, and given threshold M , then for delays $\{d_k\}$ satisfying (4.3)

$$z_k \leq M + d_{\max} - d_{\min}, \quad k \geq 1,$$

and thus the output process $\{B_k\}$ is also conforming to the CTD with the leaky-bucket algorithm where the threshold M is replaced with $M + d_{\max} - d_{\min}$.

5. CONFORMING THROUGHPUT TO DEPEND ON BURSTINESS

CTDs based on the popular leaky-bucket algorithm or the sliding-window algorithm have the property that the source can obtain the same conforming throughput regardless of whether the cells are sent with equal spacing (CBR) or sent with non-equal spacing, such as a burst followed by an idle period. A more useful descriptor for the network operator is one where the conforming throughput declines as the "burstiness" increases. Herein, we begin the examination of this issue by considering conditions on the generalized leaky-bucket CTD that

ensure that the CTD allows the greatest throughput when the traffic is equally spaced. (We do not address more general CTDs, nor the broader issue of the properties that ensure the CTD appropriately distinguishes between traffic flows that have different values of a chosen measure of burstiness.)

Suppose a source has n cells to send and suppose the source can choose to schedule the emission times, $\{A_k, k=1, \dots, n, A_0=0\}$. Suppose also that the source wishes to complete the emission of the cells as quickly as possible, subject to the constraint that the emitted cells are conforming to the CTD. Let $t_k = A_k - A_{k-1}$, $k=1, \dots, n$, and view the t_k 's as control variables of the source. Consider the following minimization problem:

$$\text{minimize}_{t_1, \dots, t_n} A_n = \sum_{k=1}^n t_k \quad \text{such that} \quad (5.1a)$$

$$x_k \leq M, \quad k=1, \dots, n, \quad \text{where} \quad (5.1b)$$

$$x_k = [ax_{k-1} - t_k]^+ + f(t_k), \quad x_0 = M, \quad a \in (0,1]. \quad (5.1c)$$

In (5.1c), x_0 , is set equal to the upper limit, M , in order to avoid the complication of some edge effects that can occur.

Note that (5.1) is a minimum time control problem, for which there is a large literature, see, e.g. [14]. As typically investigated in control theory, $f(\cdot)$ is a model of the plant, and the optimal control is expressed as a feedback function of the state. Herein, the viewpoint has shifted, where $f(\cdot)$ is a design entity, and we are interested in conditions on $f(\cdot)$ that would imply an optimal control of a particular form, namely a constant value. (Although the connection to control theory is not pursued further herein, it may be a promising direction to consider.)

Theorem 5.1: If:

- (i) $f(\cdot)$ is convex, and
- (ii) if there exists a $T \in (0, aM]$ such that $f(T) = T + (1 - a)M$, and
- (iii) if $f(t) \geq f(T)$ for $t < T$,

then a solution to problem (5.1), for any n , is:

$$t_k = T, \quad \text{for } k=1, \dots, n. \quad (5.2)$$

If, in addition, $f(\cdot)$ is strictly convex, then solution (5.2) is unique.

Remark: Note that if $a=1$, then the second condition on $f(\cdot)$ becomes a fixed-point condition.

6. SUMMARY

We have investigated desirable properties of algorithmic based traffic descriptors that are generalizations of the well-known leaky-bucket algorithm. We have shown sufficient conditions on the algorithm that yield properties that are useful to the network operator for

connection admission control and resource management. In particular, for an ATM connection that spans multiple networks, if the arrival times of cells of the connection at a given network are conforming to the CTD with given $F(x,t)$ and M and if $F(x,t)$ has certain properties and if the delays within the given network are bounded, then the exit times of the cells from the given network are also conforming to the CTD where the threshold M is increased by a specified amount, as given in Theorem 4.2. Another result is for the generalized leaky-bucket, $F(x, t) = [ax - t]^+ + f(t)$, where given that $f(\cdot)$ satisfies certain conditions, see Theorem 5.1, then the conforming throughput is maximized if all of the cells have equal spacing, and moreover the equal spacing is the only choice for the emission times that maximizes the conforming throughput. Although the work herein is at a preliminary stage and does not propose any particular algorithms, the results are suggested for consideration in future work on algorithmic traffic descriptors.

REFERENCES

1. D. M. Lucantoni, M. F. Neuts, & A. R. Reibman, "Traffic Modeling for Broadband Services," *Symposium on Worldwide Advances in Communication Networks*, May, 1992.
2. W. E. Leland, M. S. Taqqu, W. Willinger, & D. V. Wilson, "On the self-similar nature of ethernet traffic," *SIGCOMM'93*, San Francisco, CA, Sept., 1993.
3. J. Roberts, editor, *Performance evaluation and design of multiservice networks*, Luxembourg: Commission of the European Communities, 1992.
4. F. P. Kelly, "Effective bandwidths at multi-class queues," *Queueing Systems*, 9 (1991) 5-16.
5. W. Whitt, "Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues," *Telecommunication Systems*, to appear.
6. G. L. Choudhury, D. M. Lucantoni, & W. Whitt, "On the effectiveness of effective bandwidths for admission control in ATM networks," *Proc. of the 14th Inter. Teletraffic Congress*, Antibes, France, June, 1994.
7. A. Elwalid & D. Mitra, "Effective bandwidths of general Markovian traffic sources and admission control of high speed networks," *IEEE/ACM Trans. Networking*, 1 (1993) 329-343.
8. A. Berger & A. E. Eckberg, "A B-ISDN/ATM traffic descriptor, and its use in traffic and congestion controls," *GLOBECOM*, Dec. 1991, 266-270.
9. "Traffic control and congestion control in B-ISDN," ITU-T Recommendation I.371 (formerly CCITT Recommendation I.371), Geneva, Switzerland, March, 1993.
10. ATM Forum, *ATM user-network interface specification*, Version 2.4, August, 1993.
11. A. W. Berger, "Algorithmic traffic descriptors for broadband ISDN," AT&T Bell Laboratories working paper, 1993.
12. R. L. Cruz, "A calculus for network delay, part I: network elements in isolation," *IEEE Trans. Information Theory*, 37 (1991) 114-131.
13. R. L. Cruz, "A calculus for network delay, part II: network analysis," *IEEE Trans. Information Theory*, 37 (1991) 131-141.
14. A. E. Bryson & Y. C. Ho, *Applied Optimal Control*, Washington D.C.: Hemisphere Publishing Corp., 1975.