Traffic Shaping for Congestion Control in High Speed ATM Networks

A Thesis

Submitted to the College of Graduate Studies and Research

in Partial Fulfillment of the Requirements

for the Degree of

Master of Science

in the

Department of Computational Science

University of Saskatchewan

Saskatoon, Saskatchewan

by

Heng Liu

August, 1992

Distribution Notice

Traffic Shaping for Congestion Control in High Speed ATM Networks

Candidate: Heng Liu

Supervisors: Richard Bunt and Carey Williamson

Master of Science Thesis

Presented to the College of Graduate Studies and Research

University of Saskatchewan

August, 1992

ABSTRACT

Congestion control will be an essential part of B-ISDN in order to guarantee its performance level. The high transmission rate of such networks and the support of new services with diverse traffic characteristics make the congestion control problem more complicated and more challenging than in conventional networks.

Voice, data and video traffic have radically different performance requirements and are of a very bursty nature. The bursty nature of the traffic is the primary cause of congestion. Congestion results in undesirable performance degradation such as long packet delay, high packet loss and large buffer requirements. Congestion control is a network function that attempts to avoid or control congestion and limit its effects.

Traffic shaping has been proposed as a congestion control mechanism to control the input rate and smooth the burstiness of traffic. Shaping the traffic when it enters the network makes the traffic smoother and more predictable, thus reducing the probability of congestion.

This thesis studies the effects of traffic shaping on the performance of the network using simulation. An ATM-based B-ISDN network model and a traffic model are developed, as well as a traffic shaper model. Simulation results show that traffic shaping can provide good network-level performance, i.e., low queueing delay and low cell loss inside the network. User-level performance, however, may be adversely affected. Traffic shaping is basically a conservative approach to congestion control. With traffic shaping, cell loss can be reduced, but at the cost of extra delay at the shaper. Traffic shaping can reduce buffer requirements at switches with the provision of an input buffer at the shaper. It can provide protection to the network from misbehaving sources when all traffic sources are properly shaped.

ACKNOWLEDGEMENTS

A special thanks to my supervisors, Dr. Rick Bunt and Dr. Carey Williamson, for their guidance, encouragement, and financial support throughout the whole duration of this thesis. Dr. Derek Eager and Dr. Winfread Grassmann are also to be thanked for their help from the very beginning of their involvement as committee members. Thanks to Dr. Bob Gander for carefully reading my thesis and providing valuable comments.

Many thanks go to the members of our wonderful DISCUS group: Sung, Darryl, Karen, Elina, Ker Hong. I would also like to thank my fellow graduate students and friends in Saskatoon for making my life here an especially enjoyable one. I will always miss this beautiful city, and the beautiful people living and working here.

A special thank you is reserved for my parents, my sister and my brother. Their constant love and care are always very important in my life.

Finally, I owe a great deal to my husband, Yaming Ye, who has been so understanding and supportive all the way. Without his encouragement and help, I would not have come to Canada to further my studies. I am happy to be able to be together with him again.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Congestion control in packet-switched networks has been an active area of research over the past few decades. A variety of congestion control strategies have been proposed, and some of them have been applied successfully to conventional packet-switched networks.

Packet-switched networks have changed considerably in recent years. The advent of high speed packet-switching and high reliability fiber-optic transmission makes it possible to develop a Broadband Integrated Services Digital Network (B-ISDN), an all-purpose digital network that will support a wide variety of applications in a flexible, integrated and cost-effective manner. An essential part of the success of such networks is congestion control. The high transmission rate of B-ISDN, and the different service requirements of the traffic it carries, have generated new challenges for congestion control.

Advances in optical transmission have dramatically increased the capacity of the communication links. Fiber optic media have pushed the transmission speed of communication links to over a gigabit per second, representing a significant increase over typical links in today's packet switched networks. In networks with such high transmission rates, the propagation delay may dominate other delays in the network, thereby rendering ineffective many of the traditional congestion control schemes that rely on feedback signals from the network.

The services supported in such a fully evolved multiservice network will have a wide range of performance requirements and can produce a wide range of traffic flow characteristics. High-bit-rate traffic, which includes fixed-bit-rate or variable-bit-rate video and real-time services, requires a low delay and a guaranteed low cell-loss probability. Low-bit-rate traffic, including various types of data, does not require such stringent delay characteristics but is more loss-sensitive. With the traditional window-based congestion control schemes, it is not easy to meet such diverse performance requirements.

According to CCITT[1] Draft Recommendation I.121, the Asynchronous Transfer Mode (ATM) will form the basis of B-ISDN for transmitting voice, data, and video through a common set of backbone nodes and links. ATM is a high-bandwidth, low-delay, packet switching and multiplexing technique. It provides a unified means by which all traffic, connectionless as well as connection-oriented, can be transported on a common network [31, 32].

A basic characteristic of packet-based transport is the sharing of critical network resources (processing resources, transmission bandwidth, and buffers) among a large number of information flows that are highly bursty and unpredictable. Since bursty traffic sources are allowed simultaneous access to the network, user demands may occasionally exceed the available network resources. Short overloads can be accommodated by the provision of buffers, which can absorb excess packets but at the expense of delay. If the condition persists, congestion will occur. The presence of congestion leads to undesirable consequences such as long packet delay and packet

[1]CCITT: The International Telegraph and Telephone Consultative Committee.

loss. Congestion control schemes try to protect the shared resources from saturation by dynamically adjusting the traffic of the network. A congestion control scheme should result in fair and efficient allocation of network resources and limit the effects of congestion. With poor, or no, congestion control, increased load in the network may result in severely degraded throughput, accompanied by unacceptable delay and packet loss rates due to buffer overflow.

Congestion control schemes can be classified into two categories: reactive control and preventive control. With reactive control, sources adjust their traffic flows based on feedback information received from the network about the presence of congestion. With preventive control, sources have to reserve network resources before they can access the network, and they are required to remain within their allocated resources.

Most congestion control mechanisms developed for conventional packet-switched networks are reactive approaches [16, 24, 25, 33]. Senders adjust their traffic flows based on packet loss, delay, queue length, or other performance criteria. Unfortunately, reactive approaches may encounter difficulties when applied to the high speed B-ISDN. In high speed wide area networks, the propagation delay may dominate other delays in the network, rendering reactive control ineffective. For example, by the time a source receives a feedback signal from the network to regulate its traffic, it may already have transmitted a large number of packets; or by the time it does react, the congestion may already have disappeared. Preventive congestion control approaches are thus suggested for high speed networks to guarantee a certain performance level and a reasonable network utilization.

Preventive congestion control usually involves two control functions: access admission control and bandwidth enforcement. By controlling the number of connections admitted to the network, admission control ensures that the load presented to the network will not exceed its capacity. If a new connection's requirements can be met by the network, and the new connection will not affect the performance of the existing connections, it will be accepted, and allocated resources. When a connection is accepted, bandwidth enforcement will function to ensure that the traffic generated by the source conforms to its allocated bandwidth so it will not adversely affect other connections. By performing admission control and bandwidth enforcement, preventive approaches try to avoid the occurrence of congestion inside the network.

The primary cause of congestion in networks is the burstiness of the traffic. Burstiness is the dynamic behaviour of traffic sources that may generate packets at a near-peak rate for a short period of time (active period) and then enter a period of no transmission (idle period). Burstiness has a negative effect on network performance. It increases the probability of congestion and results in large buffer requirements.

Traffic shaping has been proposed to regulate the input rate and smooth the burstiness of traffic [3, 4, 7, 21]. It is the action of buffering the input traffic at the end systems and then releasing it into the network at a controlled rate. Traffic shaping also acts as a bandwidth enforcement function. Shaping the traffic when it enters the network can make the traffic smoother, more predictable and controllable, and thus reduces the probability of congestion inside the network. Traffic shaping can also force the traffic to remain within its allocated resources. The performance gains of traffic shaping are lower delay inside the network and lower probability of congestion. These gains, however, come at the risk of extra delay introduced by traffic shaping before the traffic enters the network.

The goals of this thesis are:

1. to investigate the effects of traffic shaping on performance in an ATM network environment, and

2. to obtain a better understanding of the characteristics of bursty traffic: in particular, how these characteristics affect the performance of the network and how they can be altered by a traffic shaper.

The thesis will focus on the evaluation of tradeoffs associated with traffic shaping: the tradeoff between user level performance and network level performance; the tradeoff between trying to improve cell loss performance and the extra queueing delay at the shaper; and the tradeoff between reducing buffer requirements at switches and buffer requirements at the shaper.

1.2 Thesis Outline

The remainder of this thesis is organized as follows:

Chapter 2 presents background material for the work addressed in this thesis. It discusses issues that are important to the congestion control task in an ATM-based network environment. Previous work on congestion control and traffic shaping is reviewed as well.

Chapter 3 describes the network model and the traffic model that are used in the simulation study. It then describes the experiments conducted with these models. The influences of traffic characteristics on the performance of the network without traffic shaping are presented in this chapter.

Chapter 4 describes the traffic shaper model used and investigates the effects of traffic shaping on the performance of the network. The effects of various shaper parameters are studied as well. The merits and limitations of traffic shaping are presented.

The last chapter provides a summary of the results presented in the thesis. It also gives conclusions of the work and looks into possible future research directions.

Chapter 2

Background and Related Work

2.1 B-ISDN with ATM

The purpose of the Broadband Integrated Services Digital Network (B-ISDN) is to provide an integrated services network that supports a wide range of digital traffic, from low-bit-rate traffic to high-bit-rate traffic. It is intended to support both connection-oriented services such as telephone conversations and connectionless services such as electronic mail. In recent years, there has been a growing demand for high-speed services, such as full-bandwidth High Definition Television/Advanced Television (HDTV/ATV), voice or video conferencing, video phone, and high-speed facsimile. The combination of fast packet-switching technologies and highly reliable fiber-optic transmission makes it possible to build such a high speed network.

A major advantage of integrated broadband transport for network access is its flexibility. B-ISDN is designed to support a wide spectrum of present and future services with varying (and possibly unknown) transmission rates and requirements. Another advantage is the efficiency gained by statistically multiplexing bursty traffic generated from several sources and service types.

According to Study Group XVIII of CCITT [5], the broadband user-network interface (UNI) will have a gross bit rate of approximately 150 Mb/s. They also indicate the possibility for standardizing two interface rates, one being in the 150 Mb/s range and the other at a higher rate of 600Mb/s. In order to handle such high bit rates, the basic data transmission must be very simple. Resource management and congestion control schemes also have to be simple and have to be easily implementable.

Asynchronous Transfer Mode (ATM) is a multiplexing technique proposed by the CCITT and T1[1] standards committees as the preferred transport technique for B-ISDN. CCITT Recommendation I.121, a guideline for future B-ISDN standardization, designates ATM as "the target transfer mode solution for B-ISDN" [5]. It recognizes that ATM will influence the standardization of digital hierarchies and multiplexing structures, switching and interfaces for broadband signals.

The flexibility offered by ATM makes it well suited as a B-ISDN connection-oriented transport scheme. Before a user can start transmitting over the network, a connection has to be established to provide a path between the sender and the receiver. The path will involve a number of ATM switches and transmission lines. At each ATM switch node along the path, resources (bandwidth, buffer space, etc) have to be reserved and allocated to the new connection. These resources will be released when the connection is terminated. Figure 2.1 depicts a general ATM network environment. Voice, data and video traffic flows are multiplexed together. The integrated traffic flow will be forwarded by the local ATM switch and then transmitted through the backbone ATM network to its respective destinations.

In an ATM network, bandwidth is divided into small and fixed-size packets, called *cells*, which are dynamically allocated to connections on demand. A cell contains a 5-octet (byte) header followed by 48 octets of information. Routing of cells is based on a virtual connection identifier (VCI) contained in the header, which is translated into the next physical port identifier at each switching node. Switching speed can

[1] The T1 Committee is sponsored by the Exchange Carriers Standards Association and accredited by the American National Standards Institute (ANSI).

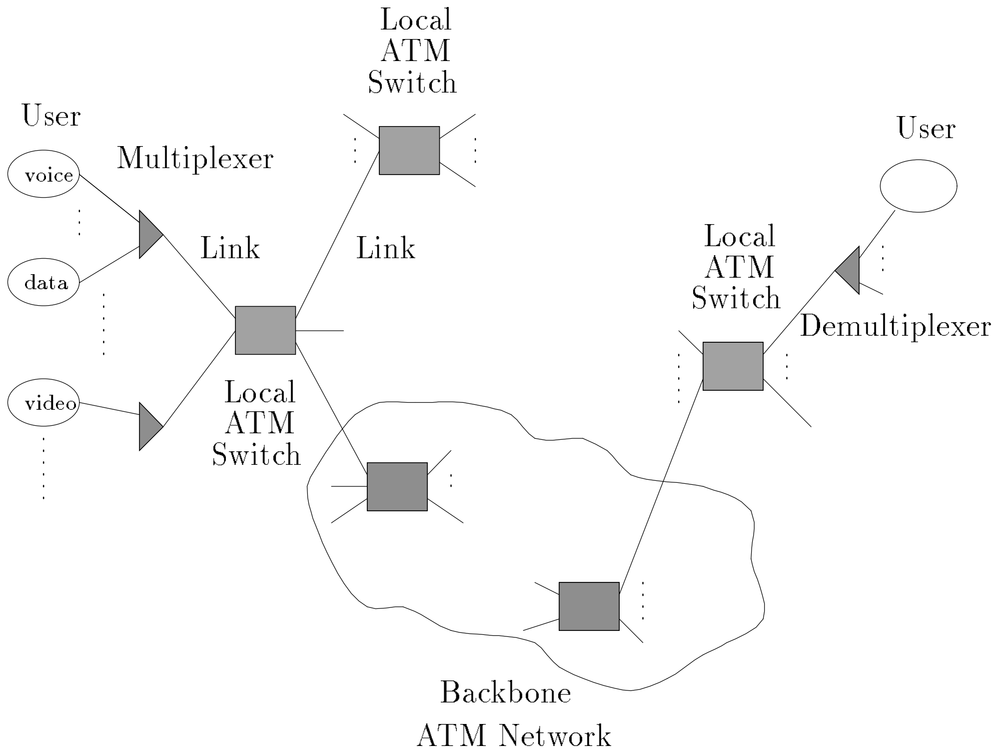

Figure 2.1: An ATM network environment

be maximized by minimizing the per-node processing, achieved through synchronous cell operation and by performing no link-by-link error recovery. Since cells belonging to a connection are identified by their header fields and there is no dedicated physical channel, the connection is only a virtual connection. In this way, bandwidth can be allocated to services on demand.

ATM fits into the OSI reference model[2] between the physical and data link layers [39], actually overlapping the bottom portion of the data link layer, as shown in Figure 2.2. The ATM layer can be further divided into an upper ATM Adaptation Layer (AAL) and a lower ATM Transport Layer (ATL). The ATM Adaptation Layer

[2]Open Systems Interconnection (OSI) Reference Model is a layered network model proposed by the International Standards Organization.

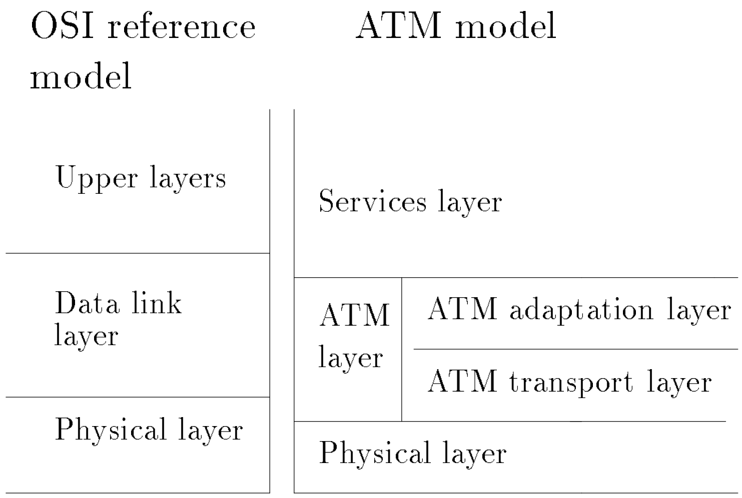

Figure 2.2: A comparison of the OSI and ATM network models

performs the interface between the data link layer (which is non-ATM) and the lower ATM layer. This interface includes such functions as segmentation and reassembly. Traffic flows originating at the services layer from user terminals or end systems are split up into ATM cells by the sender's AAL. At the receiving end, these cells are reassembled by the adaptation layer before being passed to the higher layer. The ATM Transport Layer, which is supported by the physical layer, is in charge of multiplexing and demultiplexing several streams of traffic onto the network.

The advantage of ATM is its flexibility to integrate the transport of services with varying bandwidth requirements by assigning cells to virtual connections on demand. Since the traffic of all services looks the same at the ATM level, switches can be very simple. ATM also provides the potential to improve bandwidth efficiency through statistically multiplexing variable bit rate (*i.e.*, bursty) traffic streams.

2.2 Congestion Control

2.2.1 The Problem

An ATM-based network tries to make efficient use of resources by allowing many bursty traffic sources to have access to the network simultaneously. This simultaneous access, however, may result in a state where the aggregate network traffic occasionally exceeds the capacity of the network, thereby causing a degradation in network performance. This state is referred to as *congestion*. The presence of congestion brings about undesirable consequences such as long queueing delay, buffer overflow, and high cell loss. *Congestion control* is a network layer function that attempts to avoid or control congestion and limit its undesirable effects. Approaches to congestion control include reserving resources in advance of their use [7, 10, 34], discarding excessive cells when overload occurs [9, 30], and shaping the traffic before it enters the network [3, 29, 38].

Most conventional congestion control schemes rely on end-to-end exchange of control messages in order to regulate traffic flow. The control messages (sometimes with additional congestion information added by the intermediate nodes) are used as feedback by the source node, which then adjusts its traffic level. This is called a *reactive* scheme. There are several potential problems with a reactive control approach in a broadband ATM network:

1. With the high speed of B-ISDN (potentially in the gigabit per second range) and the small sizes of cells, the propagation delay (the time for the first bit to travel the network) is considerably higher than the cell transmit time (the time between the first and the last bit of the packet) and switch processing time. The number of cells traveling in the high speed links (also called "in the pipe") is also considerably higher than in conventional networks [26]. Thus, any feedback from the network is usually outdated. By the time a source gets the feedback

signal and reacts to it, several thousand cells may have been lost owing to buffer overflow. The feedback control loop dynamics may also be sensitive to transient traffic behaviour, making it difficult to avoid oscillations or over-reactions while at the same time ensuring the guaranteed performance requirements.

2. The congestion control mechanism should operate at the speed of the communication link. For this reason, computationally intensive control schemes are less desirable than simple schemes that can be implemented easily in high-speed hardware.

3. The nature of traffic also affects the design of the congestion control mechanism. For example, real-time traffic in broadband networks usually requires some level of bandwidth guarantee. The ability to slow down such sources is usually very limited. It would also be difficult to determine which terminals or end systems can or should be throttled back when each may have such diverse bandwidth requirements.

4. Congestion control schemes in almost all existing computer networks utilize window mechanisms to adjust traffic flow [16, 24, 25, 33, 42]. *Window-based* control schemes use the window to limit the maximum number of outstanding packets (*i.e.*, sent but not yet acknowledged). For a network speed of 150 Mb/s, window-based control schemes would require an inordinately large window size and an inordinate amount of buffer space in intermediate nodes for long distance connections.

All the above arguments lead to the suggestion of simple open-loop approaches that do not require feedback and are preventive in nature. Preventive controls try to operate the network more efficiently by controlling the traffic flow admitted into the network at the access nodes. A new connection would be accepted only if its required performance could be met. Each connection's traffic flow is enforced to remain within its allocated bandwidth.

For preventive congestion control schemes for high speed networks, many researchers suggest that rate-based approaches are preferable to window-based approaches [3, 29, 38, 46, 47, 48]. Since most network resources, such as transmission links and switch processors, are rate limited (*i.e.*, they cannot sustain packets arriving at a rate faster than their capacity), transmission rate can provide a proper match between user demands and the resources available. Another argument for rate-based controls is that much of the high speed traffic will be stream oriented (such as voice or video), requiring a guarantee based on rate rather than on number of packets. Rate-based approaches seem to be a natural choice for congestion control in high speed networks, and so far, most proposed approaches are rate-based approaches [3, 7, 10, 12, 18, 20].

Any appropriate congestion control scheme should be able to ensure that during the information transfer phase, the performance requirements of all services are met, and the network resources are allocated in an efficient and fair manner among diverse services. A successful congestion control policy for ATM networks at broadband speed should also meet the following design principles [13]:

- *Simplicity*: The control should be kept as simple as possible since it is expected to function in a high speed network environment. The challenge is to achieve high resource efficiency while maintaining relatively simple control structures.

- *Flexibility*: The control scheme should be flexible enough to accommodate a range of broadband services that have different performance requirements. The congestion control scheme should make minimal assumptions about the service mixture since it is still unknown. It is important that the scheme has the flexibility to adapt as needed to new situations that arise.

2.2.2 Call Level and Cell Level Congestion Control

ATM networks are designed on a connection-oriented transport principle, which implies a setup of a virtual connection with end points. The request for the establishment of a connection is termed a *call*. Congestion control can be exercised at the call level and at the cell level [18].

At the call level, access admission control decides whether or not to accept an incoming call into the network. The purpose of admission control is to avoid long term network congestion. During the call setup phase, network resources are checked and reserved along the connection route before the call starts. The call is accepted if network resources are available to satisfy the required class of service. Otherwise, the call is rejected.

At the cell level, bandwidth enforcement functions located in network access nodes monitor the traffic status. The bandwidth enforcement function regulates the input traffic rate to ensure that the traffic stays within reasonable limits of its allocated capacity so it will not affect other connections in progress. This input traffic rate regulation is also called traffic shaping.

ATM connections can be allocated resources based on either deterministic multiplexing or statistical multiplexing. With deterministic multiplexing, each connection is allocated sufficient resources to accommodate its peak cell transfer rate. In this case, congestion at the cell level is totally eliminated. This mode of operation is appropriate only for smooth, constant rate traffic sources. The disadvantage of deterministic multiplexing is that good bandwidth utilization cannot be achieved when a large proportion of traffic is bursty. Since a bursty traffic source is not always in the active state sending cells at its peak rate (for example, when the source is in the idle state) the reserved bandwidth will be wasted.

Alternatively, connections may be multiplexed statistically to take advantage of the variable bit rates generated by different applications. Statistical multiplexing allows many connections to share resources on the assumption that each connection requires these resources only for a small fraction of the time. Whenever statistical multiplexing is used there is a possibility that the offered load to a multiplexing node will exceed the node's capacity. Congestion will occur due to the simultaneous arrival of cell bursts from independent sources. In order to make good use of statistical multiplexing, there should be a good understanding of traffic characteristics and carefully designed congestion control schemes.

Call-level Congestion Control: Admission Control

It is natural that the network may wish to accept any and all connection requests in order to increase network utilization. However, the network also has to ensure that the admission of a new connection does not adversely affect the performance of other connections sharing these network resources and will not cause long term network congestion. Admission control, including bandwidth allocation, is essential to maintain the balance between network performance and network utilization.

The decision to accept or reject a connection request should be made based on a description of the new connection's anticipated traffic characteristics, its performance requirements, and the current network status. Many admission control schemes have been proposed [18, 21, 37]. They differ mainly in how they characterize traffic and how they compute the appropriate bandwidth quota.

For admission control, any connection should be depicted by its traffic descriptors, a set of parameters characterizing a connection's traffic flow pattern. Traffic descriptors should correspond to measurements that are easily enforced. CCITT Recommendation I.311 [6] suggests using average cell rate, peak cell rate and mean peak (burst) duration as traffic descriptors for a connection. These three parameters, however, may not be able to describe a connection's characteristics adequately [17].

The determination of an appropriate set of traffic descriptors and an accurate estimation of required bandwidth are challenging problems. Exact prediction of traffic characteristics is difficult, so the admission control should be robust. The basic principle of admission control is: if all the estimated bandwidths sharing the same communication link sum to less than a predefined maximum value, the new call can be accepted [10, 18, 43]. This estimated bandwidth of the new connection then becomes the amount of bandwidth that is allocated to this connection, and the user is expected to keep within the bounds of this "contract" or suffer the consequences in terms of degraded services.

The most common admission control scheme was proposed by Habib [21], based on the above principle. The scheme allocates an amount of bandwidth that is between the peak and the average rates. This amount is called the effective bit rate, which is a fraction multiplied by the peak rate. This fraction reflects the variance of the traffic characteristics. The sum of all allocated bandwidth, *i.e.*, the most achievable load level, is limited to be 80% of the total link capacity.

The principle of admission control is intuitively simple. The actual implementation, however, involves a tradeoff between the network utilization and the effectiveness of congestion control.

Cell-level Congestion Control: Bandwidth Enforcement

Once a connection has been admitted to the network, the traffic flow must be monitored and enforced at the network access point to ensure that it conforms to its allocated bandwidth, or equivalently, to its traffic descriptor values.

Most traffic is inherently bursty and unpredictable, and is described only by its statistical measures such as average number of cells over a time interval, mean burst length, etc. It is possible to monitor statistical moments, given a sufficiently long time interval, but it is very difficult to enforce them since it must be decided whether a momentary burst of traffic is due to a violation, or is due to natural statistical variation. Upper-bounding parameters such as the peak rate and maximum burst length are more likely to be put under explicit control.

Regardless of the specific descriptors used, if the discrepancy between natural and required traffic flow is significant, there will be a need to sufficiently shape a statistically varying traffic flow to match its descriptors. Traffic shaping or traffic smoothing helps to avoid congestion by controlling the rate at which cells are allowed to enter the network from a given source. Many traffic shaping mechanisms have been proposed and studied, such as the sliding window, the jumping window, the exponentially weighted moving average (EWMA), and the leaky bucket scheme [40]. All these schemes shape the traffic for users, enforcing the contract for the network.

The *leaky bucket* mechanism is the most commonly proposed mechanism for traffic shaping [12, 18, 38, 40]. In the leaky bucket approach, cells are dropped from traffic sources violating their bit rate contract. The basic leaky bucket algorithm enforces a mean bit rate and a maximum burst length by means of a token pool (see Figure 2.3). Tokens are generated at a fixed rate. If the token pool is full, generated tokens are discarded. An arriving cell takes a token and then goes through to the network. If no tokens are available, the cell is called a violation cell and is dropped immediately. The size of the token pool (*i.e.*, the maximum number of available tokens) limits the maximum possible length of cell bursts that are allowed to enter the network. The token generation rate enforces the mean cell rate. This leaky bucket scheme can be modified [38] to provide input buffers for violation cells so as to better control the tradeoff between cell loss and queueing delay.

The jumping window mechanism [11] limits the maximum number of cells accepted from a source within a fixed time interval (window) to a maximum N. The first N cells arriving in a given time interval are allowed to enter the network. The following cells will be dropped if they arrive in this time interval. The new interval starts immediately at the end of the preceding interval (*i.e.*, jumping window) and the

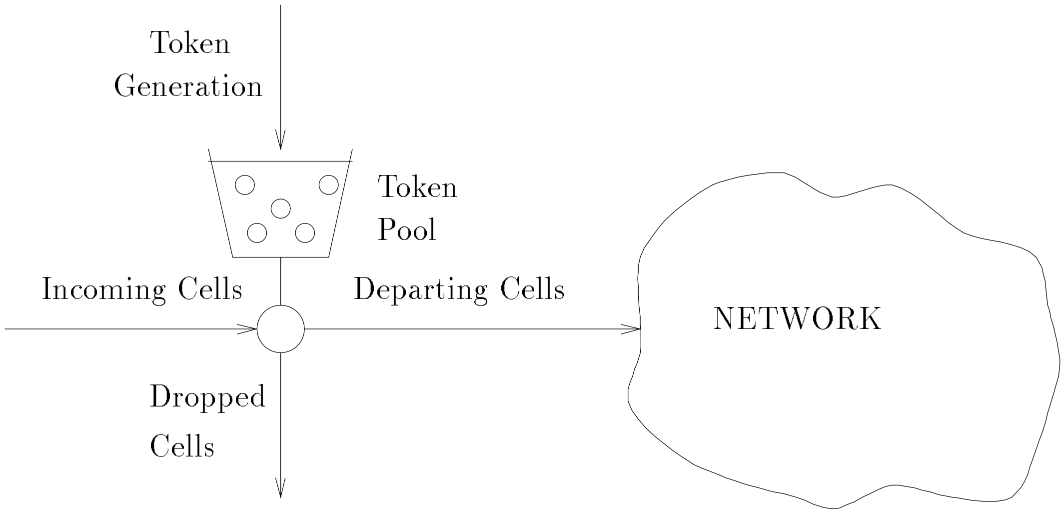

Figure 2.3: The original Leaky Bucket mechanism

associated counter is restarted with an initial value of zero.

The EWMA mechanism [35] uses fixed consecutive-time windows like the jumping window scheme. The difference is that the maximum number of accepted cells in the ith window (N_i) is a function of the allowed number of cells N per interval and an exponentially weighted sum S_i of the number of accepted cells X_i in the preceding intervals, *i.e.*,

$$N_i = \frac{N - \gamma S_{i-1}}{1-\gamma}, 0 \leq \gamma < 1$$

where $S_{i-1} = (1-\gamma)X_{i-1} + \gamma S_{i-2}$.

The factor γ controls the flexibility of the algorithm with respect to the burstiness of the traffic. If $\gamma = 0$, N_i is constant and the algorithm is identical to a jumping window scheme. A larger value of γ allows more variable source behaviour.

In a comparison of shaping mechanisms, Rathgeb [35] claims that EWMA and leaky bucket are the most promising with respect to mean cell rate policing. The jumping window scheme is not flexible enough to cope with the short-term fluctuations of the traffic. The implementation complexity of EWMA is higher than other mechanisms. Leaky bucket is claimed to be the most preferred policing scheme.

2.3 Traffic Characteristics

2.3.1 Service Categories

Successful integration of diverse services into a unified network requires a good understanding of traffic characteristics and service needs. The wide range of bit rates, the statistical profile of the information flow, and the variety of connection configurations (point-to-point, multi-point, symmetrical connections, etc.) make this characterization more complex than that required for conventional networks.

It is not easy to characterize broadband traffic since most of it still remains unknown. Nevertheless, it is essential that B-ISDN should at least meet the requirements of service categories that can be identified now, such as real time video and voice, and various types of data traffic. While data traffic is loss-sensitive (*i.e.*, cannot tolerate many cell losses), real time video and voice tend to be more delay-sensitive (*i.e.*, have a stringent bound on delay). Video traffic can also be both loss-sensitive and delay-sensitive. Typically, an acceptable cell loss probability is 10^{-4} for voice cells and 10^{-9} for video cells [22]. More diversified and sophisticated services are expected to emerge as broadband capabilities become more wide-spread. To simplify the discussion, the variety of services is grouped into three basic categories: data, video and voice.

There are two types of data traffic: interactive data traffic and bulk data traffic, corresponding to remote terminal access (*e.g.*, telnet) and file transfer (*e.g.*, FTP) services. Interactive data traffic between terminals and hosts, and transactions between diskless workstations and remote servers in client-server computer systems, require a high peak bit rate to achieve low response time. The average bit rate, however, may be low. Traffic is extremely sporadic with long silent intervals separating rapid high speed bursts. Bulk data transfer, such as large host-to-host file transfer between diskless workstations and remote file servers, generates bursty data traffic with long burst lengths. Usually such applications are loss sensitive but can tolerate a certain level of delay.

Video traffic can also be divided into two groups: image retrieval and real-time video. Access to image banks for a variety of applications, including medical imaging, computer-aided design, and an enhanced photograph videotext service, could generate a significant amount of traffic. Each image is coded in a number of bits ranging from some hundreds of kilobits to tens of megabits, depending on the resolution and the use of image coding techniques. Image retrieval applications typically require the transmission of a succession of images at irregular intervals. The traffic is of the on/off type where traffic varies between activity periods at peak rate and silence intervals, depending on user/database interactions. Display time compatible with user comfort requires a bit rate in the range of 1 to 10 Mb/s [36]. The traffic has moderate constraints on the delay it can incur and is highly loss-sensitive.

Real-time video, such as interpersonal video communication in an enhanced visual telephone service or in a video-conferencing application, requires bit rates of up to several megabits per second. The bit rate of such traffic varies from frame to frame in a continuous range. Distribution video for entertainment generates bit rates ranging from tens of megabits per second to more than 100 Mb/s for High Definition Television (HDTV). Both fall under the category of real-time traffic and are extremely sensitive to delays and variations in delay.

Voice traffic is also a source of real-time traffic but requires a low bit rate. It is characterized by a low volume of traffic and short burst lengths. Voice traffic is sensitive to delay, but can tolerate a certain level of cell loss. The delay limit for voice may be longer than the delay video traffic can tolerate.

Figure 2.4 illustrates the ranges of the end-to-end traffic performance requirements that might be expected for the services described. Performance requirements are specified in terms of maximum cell delay variation and probability of cell loss. Interactive services will have stringent delay requirements. Image retrieval and compressed video applications may have stringent loss requirements. The actual values of the requirements depend on the type of coding scheme, subjective testing scale assumed, etc., and are a subject for further research.

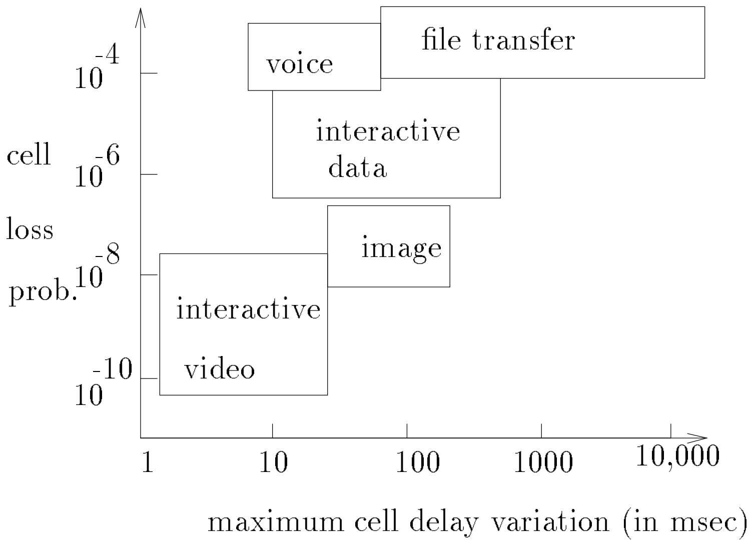

Figure 2.4: ATM traffic performance requirements (from [44])

The performance requirements stated in Figure 2.4, however, may be over engineered. Real application requirements may not be so stringent. Many researchers have looked at networks operating at a high load level and with cell loss rates higher than 10^{-4} [18, 32, 38]. A network under high load levels is more likely to suffer from congestion, which is the topic of interest in this study. Experiments presented in this thesis focused on regions of heavy network load, hence resulting in high cell loss. It is believed that the qualitative aspects of the results should also be applicable under conditions of lighter loads.

2.3.2 Traffic Models

Any study of B-ISDN networks requires a workload model to test the performance of designs. Traffic models that approximate the assumed traffic characteristics of potential services are used in these studies.

The cell stream arriving at a B-ISDN network is the superposition of the cells of individual calls. Usually there are different call types to consider. The randomness of the call mix and the fluctuation of cell generation in individual calls make traffic modeling a difficult problem. A number of models have been proposed as approximations of individual and/or aggregate traffic sources in a high speed network.

The simplest and most commonly used traffic model is the simple Poisson model that assumes that the arrival of cells is a Poisson process. It has the advantage of being analytically tractable. The actual traffic of B-ISDN, however, is believed to be more bursty than in a Poisson process due to the simultaneous presence of many calls of possibly different types and also due to the inherent burstiness of some call types [14, 23, 45].

Discrete-time Interrupted Poisson Process (IPP) is also a commonly used traffic model that can model bursty traffic. The IPP model was originally introduced by Kuczura in a continuous-time domain to model overflow traffic in telephone trunks [28]. It is also reported that IPP can nicely describe bursty traffic in ATM networks [32]. IPP is a process with two states, on (active) and off (idle), which alternates continuously, as shown in Figure 2.5. Arrivals during the active state are Poisson, and holding times in the active and idle states are exponential. Gallassi has used the IPP model to model various classes of traffic that have varying degrees of burstiness [18]. The IPP model, however, cannot model the correlation in the overall cell arrival process. This correlation appears, in general, to be non-negligible [17].

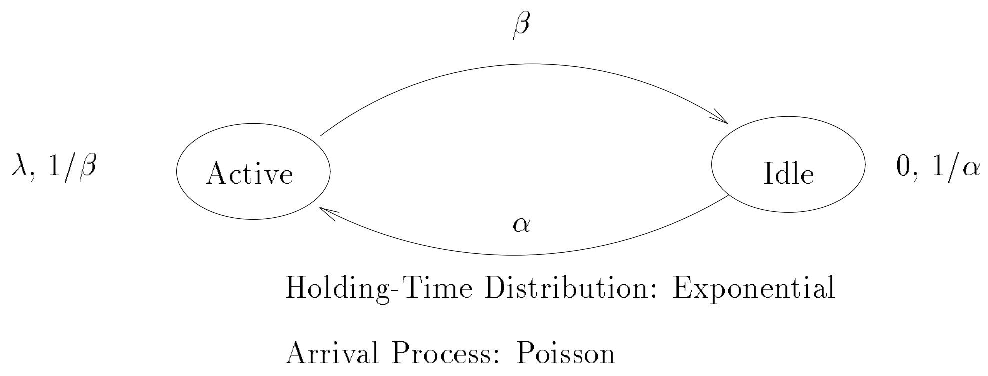

Figure 2.5: IPP traffic model

The most powerful model, and one used in many ATM research papers, is the Markov Modulated Poisson Process (MMPP) proposed by Heffes and Lucantoni [23]. An MMPP is essentially an m-state continuous-time Markov chain (see Figure 2.6). Associated with each state i is a bit rate λ_i and a mean holding time $1/r_i$. The arrival process in each state is Poisson, and the holding times are exponentially distributed. The simplest form of MMPP is a two-state Markov chain where the mean holding times in state 1 and state 2 are $1/r_1$ and $1/r_2$, respectively. The arrival processes of cells in state 1 and state 2 are Poisson with rates λ_1 and λ_2, respectively. A two-state MMPP has been used to approximate the superposition of packetized voice sources together with data traffic arrival at a multiplexer [23]. The IPP is actually a special case of the two-state MMPP, where the bit rate in state 1 is λ and the bit-rate in state 2 is 0. Both IPP and MMPP can model short-term correlated behaviour reasonably well, but less able to model the long-term correlations.

There are some traffic models proposed that are more complex than IPP and MMPP models [15]. It is argued, however, that complex traffic models are useful only when their parameters can be estimated accurately [44]. An MMPP model will be used in this thesis because of its flexibility in simulating different traffic sources.

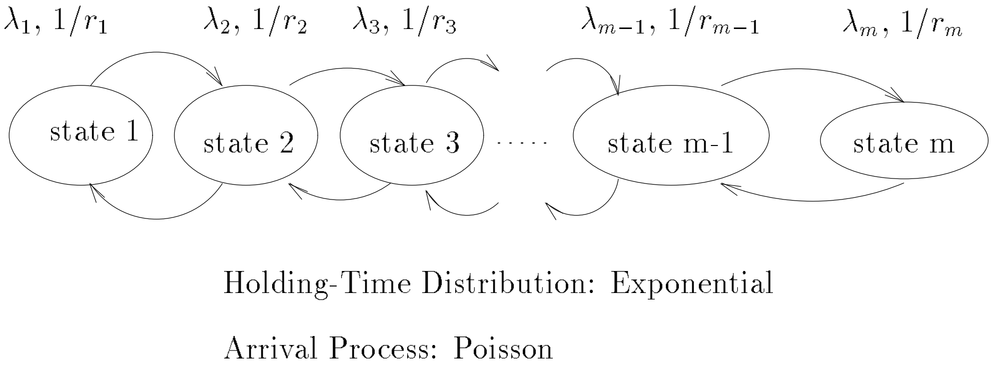

Figure 2.6: MMPP traffic model

2.3.3 Burstiness and Smoothness

Burstiness is an important traffic characteristic, but defining it is not easy. A simple way to define burstiness is as the ratio of the source's peak bit rate to its average bit rate [10, 18, 21]. There are other ways to characterize the burstiness, however. For example, variable bit rate video sources can be represented by time domain characteristics such as the coefficient of variation, the autocorrelation function, and the distribution of the autocorrelation function. Burstiness can also be represented by the squared coefficient of variation of the interarrival times, which is the ratio of the variance of the cell interarrival times to the square of the average value of the cell interarrival times. Akhtar [2] defines burstiness in the IPP model as $\lambda/(\alpha+\beta)$, where λ is the bit rate in the active state and $1/\beta$ and $1/\alpha$ are the mean holding times in the active and idle states, respectively.

Similar to the case of the definition of burstiness, there is no commonly accepted definition for smoothness either. Golestani [20] defines a traffic stream to be *(r, T)-smooth* if during each time period T, the arriving packets collectively have no more than $r * T$ bytes. In the case of ATM, where cells have a fixed size of 53 bytes, a stream of cells is (r, T)-smooth when the number of cells received during each time period T is bounded by $r * T/53$. Evidently, r represents an upper bound to the average rate of traffic (in bytes/sec) and T represents the time period during which this average is obtained. In fact, (r, T)-smoothness can also be viewed as a two-parameter characterization of the traffic, where r is an indication of the average rate and T is an indication of the burstiness.

Sidi *et al.* [38] characterize the smoothness of traffic by the variance of the inter-departure times of the cells leaving a traffic shaper for the network. The smaller the variance, the smoother the traffic can be viewed when it enters the network. They use this definition to demonstrate the effect of traffic shaping on the characteristics of the traffic.

It may not be possible to depict burstiness with a single parameter. In this thesis, the definition of burstiness level as the ratio of the peak rate over the mean rate is adopted, because of its simplicity.

2.4 Related Work on Traffic Shaping

Although it is widely acknowledged that traffic shaping is essential for congestion control [18, 21, 32, 36, 38, 43], few researchers have actually investigated the effect of traffic shaping on the performance of the network. This section provides a brief review of some previous work that is relevant to this thesis.

Sidi *et al.* [38] studied the smoothing effects of input rate regulation mechanisms (traffic shapers) on the characteristics of the traffic after shaping. The smoothing effect was characterized by the variance of the inter-departure time of the packet departure process from the traffic shaper. An analytic model was used to model a traffic shaper that employed a scheme similar to the leaky bucket. Under the assumption of Poisson arrivals, they studied the tradeoff between the smoothness of the shaped traffic and the packet waiting time in the shaper. Results were given for both finite and infinite buffer size cases.

From numerical results they concluded that the token generation rate should be close to the source arrival rate in order to have a significant effect on the system performance. They also found that the system performance was very sensitive to changes in the token generation rate. Decreasing the generation rate would result in an increase in the expected waiting time of cells in the shaper while the coefficient of variation of this waiting time would be decreased. Last, they observed that increasing the buffer size at the shaper decreased the cell loss at the shaper. Increasing the token pool size had a similar effect.

A major deficiency of this study is that they assumed Poisson cell arrivals, which, as already mentioned, is not sufficient in depicting the variable degree of burstiness of ATM traffic. The effect of traffic shaping on performance at the network-level was not examined.

Murata *et al.* [32] modeled an ATM switch as a discrete-time single-server queueing network (a $GI + M^{[x]}/D/1$ model) where a new call joined existing calls. They investigated how the network performance depends on the statistical characteristics of a new call and the existing calls, and demonstrated the effectiveness of admission control and traffic smoothing. The IPP traffic model was used to approximate the cell arrival of the new call, and a Bernoulli process with batches was used to approximate the aggregated arrivals of cells from the existing calls. They also considered both finite and infinite buffer cases.

The results they gave were the waiting time distribution and cell loss probability for a new call and for the existing calls in an ATM switch. They observed that the performance diminished proportionally to the on and off durations of the new connection. The longer the on and off durations, the worse the performance became, *i.e.*, higher cell loss and longer waiting time in the switch. They also showed that smoother traffic (traffic with the same load but longer on and off durations, so cell arrivals were spread over a longer time period, giving smoother traffic) obtained better performance, *i.e.*, lower waiting time and cell loss. They concluded that smoothing the input traffic could reduce the chance of network congestion. They didn't discuss, however, how to perform traffic shaping and the tradeoff between extra delay and improved performance inside the network.

Leland [29] used a cell-level simulation of an ATM network to study the performance of three proposed window-based congestion control schemes. They were a rate-based window scheme (similar to the leaky bucket), a traditional end-to-end acknowledgment window scheme, and a congestion feedback window scheme. He modeled the network as a single logical switching node. Eight identical independent traffic sources transmitted through the switch, each to its own destination. One access control window was assigned to each of the eight sources. The traffic sources generated cell trains of length 10. The performance in terms of throughput, delay and cell loss was studied. Leland found that, when assessed at the network level, the leaky bucket scheme performed well while neither acknowledgment-based mechanisms performed adequately. When assessed from the user perspective, the leaky bucket also failed to provide satisfactory service. He suggested that further study of intranetwork queueing and routing mechanisms was required to control congestion within the network and provide assured quality of service to the end users.

Butto *et al.* [4] analyzed the original leaky bucket scheme as a $G/D/1/N$ queue with finite waiting room. Since the complexity of the arrival patterns made it impossible to derive a formula for system performance, they used a fluid flow approximation, by which the bit flow of traffic was considered as a continuous variable. The traffic model they used was an on/off type two-state model. The burst duration and the inactive state duration were exponentially distributed. During the burst, the cells were assumed to arrive periodically. Following the fluid flow approximation, they derived an explicit formula for the cell loss probability.

Their conclusion was that the leaky bucket parameter settings should be based on the source traffic characteristics. They claimed that the peak rate could be controlled easily by setting the token generation rate near the source peak rate, but enforcing the mean bit rate required a token generation rate near the source mean rate. They also concluded that the burst duration was not easily controlled by the leaky bucket as long as the source increased in the same proportion as the burst and silence durations.

In the work of Butto *et al.*, the leaky bucket scheme did not have an input buffer to accommodate cells arriving in a burst, which, as will be discussed later in Chapter 4, is very important in providing a tradeoff between cell loss and queueing delay for traffic sources. Their work was focused on the policing effect of the leaky bucket scheme. This is also different from the viewpoint of the work presented in this thesis, which puts the emphasis on the shaping effect of a modified leaky bucket scheme.

The research presented in this thesis focuses on the concept of traffic shaping. The traffic model used is powerful enough to simulate traffic of different burstiness levels. The study investigates extensively the effects of traffic shaping on performance, both at the user level and at the network level.

Chapter 3

Network Performance Without Traffic Shaping

3.1 Introduction

In this chapter, simulation experiments are described that were performed to investigate the performance of an ATM-based B-ISDN network. The choice of simulation was due to the unavailability of a real ATM-supported B-ISDN testbed and limitations of analytical techniques. An attractive feature of simulation is its flexibility to explore a full range of parameters and models.

The objective of the simulation experiments in this chapter was to investigate how the input traffic characteristics affect the performance of the network when there is no traffic shaping. A network model and a traffic model were developed for this purpose.

Two basic criteria are used in the evaluation of overall performance of the network: average end-to-end queueing delay and cell loss.

3.2 Network Configuration Model

First, a simple network configuration model was developed for an ATM-based B-ISDN network. The model has four switches and connects eleven source hosts (src) to eleven destination hosts (dest), as shown in Figure 3.1. Each source can represent a single call or an aggregation of a number of calls. Calls are routed deterministically to their destination hosts. Each link is a 150 Mb/s communication link, and all links are assumed to be error free. Eleven virtual connections (src i to dest i, for $i = 1, 2, 3, \ldots, 11$) were provided, with path lengths varying from two to four switch hops (*i.e.*, crossing one to three switch-to-switch links). The shorter connection (connection 11) can be viewed as providing cross traffic to the longer ones.

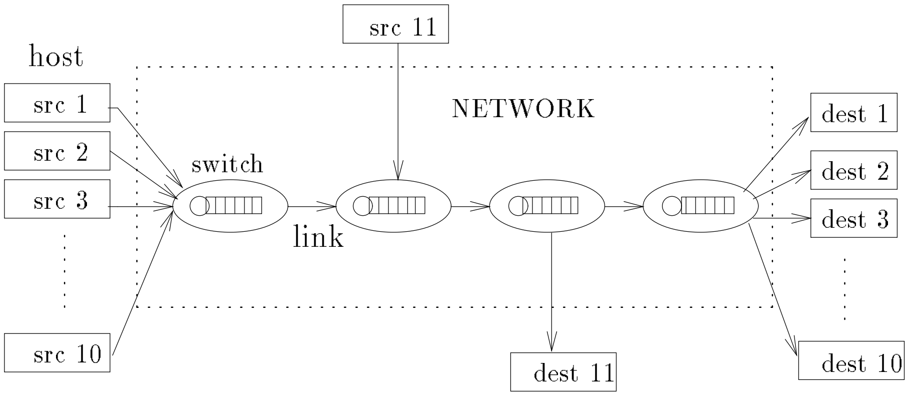

Figure 3.1: Network configuration for simulation

The model assumes internally non-blocking, output-buffered fast packet switches. Such switches can perform much better than input-buffered switches and achieve the optimal throughput-delay performance [1, 8]. For simplicity, a switching delay of zero is assumed. The service discipline in switches is FCFS.

The transmission time for a 53-byte ATM cell is approximately 2.8 μs on a 150 Mbps link. This transmission time is referred to as a *slot* and is the time unit used in the simulation. Bandwidth of links is allocated in units of slots.

3.3 Traffic Model

A two-state MMPP traffic model was adopted for the simulation experiments. The choice of MMPP is due to its popularity and its flexibility to model different traffic.

As described in Subsection 2.3.2, a two-state MMPP model is a Markov chain of two states (see Figure 3.2). Cells are generated according to a Poisson process with mean arrival rates λ_1 and λ_2, respectively. Holding times in each state are exponentially distributed with means $1/\beta$ and $1/\alpha$, respectively.

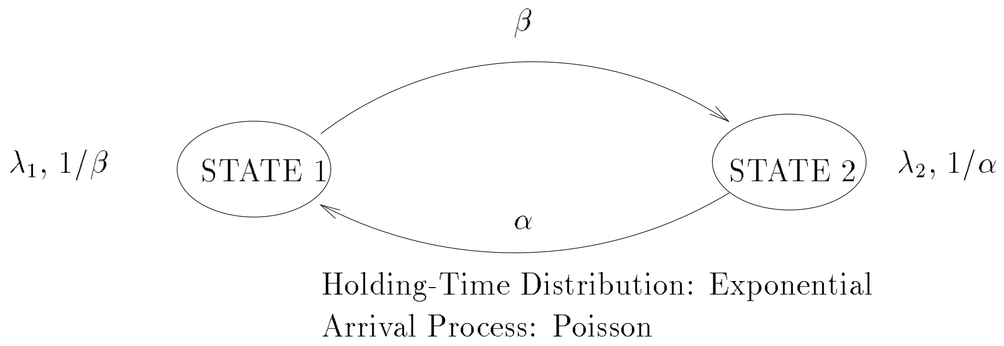

Figure 3.2: Two-state MMPP traffic model

In this thesis, the peak bit rate is referred to as B_p, defined as the rate of arrivals during the active state, and the mean bit rate as B_m. The burst duration is T. If state 1 is specified as the active state, then the following three parameters completely characterize the traffic:

- mean bit rate: $B_m = (\alpha\lambda_1 + \beta\lambda_2)/(\alpha + \beta)$,
- peak bit rate: $B_p = \lambda_1$, and
- mean burst duration: $T = 1/\beta$.

Different levels of burstiness can be simulated by varying the parameters of the model. Different levels of workload can be modeled as well.

All traffic sources used in the simulation are modeled using the MMPP model. In most scenarios, all sources have the same MMPP parameters, *i.e.*, the network is fed by homogeneous sources. In other scenarios, different sources may have different parameters.

3.4 Statistical Aspects of the Simulations

Since simulation is used in this study, various errors, such as statistical variation, may be present in the results. To obtain statistically meaningful results, sample sizes for random variables and the simulation run length must be chosen carefully. According to estimation theory, if a sample mean $\overline{X}$ is used as an estimate of a random variable μ, in order to be $(1-\alpha)100\%$ confident that the error will not exceed a specified amount e, the required sample size is calculated by:

$$n = \left(\frac{z_{\alpha/2}\sigma}{e}\right)^2$$

where σ is the standard deviation [41].

In this thesis, the desired run length of each simulation run is calculated according to estimation theory. In each simulation run, statistics on successfully transmitted cells are calculated on batches, with each batch consisting of 10,000 time slots. To reduce the effect of initial bias on the results, the first 2 batches of each run are ignored. In order to be 95% confident that the average queueing delay estimated by the simulation is within 1% of the mean queueing delay, the number of batches, n, can be calculated according to the formula above. For a 95% confidence interval, $z_{\alpha/2} = 1.96$. The standard deviation σ is estimated by taking a preliminary simulation run with the number of batches $n \geq 30$. Then, the sample standard deviation s is used as an approximation for σ.

The following equations are used to calculate the required number of batches:

$$N_{new} = \left(\frac{z_{\alpha/2} \quad s}{e}\right)^2$$

$$s = \sqrt{\frac{\sum_{i=1}^{N_{old}}(X_i - m)^2}{N_{old} - 1}}$$

$$m = \frac{\sum_{i=1}^{N_{old}} X_i}{N_{old}}$$

$$e = 0.01 * m$$

X_i is the average queueing delay within batch i, m is the average queueing delay over all batches, N_{old} is the previously determined number of batches required, and N_{new} is the newly calculated number of batches required.

In each simulation run, N_{old} was assigned to be 30 initially. After N_{old} batches of time slots were passed (in addition to the first 2 batches whose statistics were ignored to reduce the initial bias), N_{new} was calculated. If $N_{new} > N_{old}$, N_{old} was assigned the value of N_{new} and the simulation proceeded until the number of batches again reached N_{old}; otherwise, the simulation run was terminated. This iterative process continued until finally $N_{new} \leq N_{old}$. Note that it is assumed that the results of the other performance metrics obtained are as accurate as the average queueing delays obtained, as the average queueing delay was the metric used for the calculation of the simulation run length. The typical batch size in the simulation experiments presented in this thesis is 30.

In all the experiments conducted in this thesis, the confidence interval is 1% of the mean value of the performance measure. If a measured value of queueing delay is 100, for example, this means the 95% confidence interval of the delay is 100±1.

3.5 Simulation Results and Analysis

In this section, experiments to determine the influence of three traffic characteristics on the performance of the network without traffic shaping are described. These characteristics are:

- the burstiness level b, defined as the ratio of the peak rate to the mean rate ($b = B_p/B_m$),
- the load level ρ, defined as the sum of the mean rates of all sources, and
- the mean burst duration T, defined as the average time spent in the active state.

Switch buffer requirements for traffic sources with different burstiness levels and burst durations were examined in order to obtain a better understanding of the dimensioning of buffers at the switches.

The performance of the network operating at heavy load (*i.e.*, ρ in the range from 0.8 to 1) was the focus of this study. Since networks under congestion are the subject of interest, sufficient load is needed to result in congestion. Burstiness levels ranged from 2 to 10. When the load level is high, a burstiness level of 10 can give a very high peak cell rate such as 0.95 cells/slot for a single source. The mean burst duration T was varied from 10 time slots to 100 time slots. Traffic with a very long burst duration (such as 1000 cells) cannot obtain reasonable performance with a buffer size on the order of 100 cells [32]. Buffer size ranged from 10 cells to 600 cells. Buffer sizes around 100 cells are the most frequently used in ATM research [17]. Table 1 gives the range of parameter values determined.

In all scenarios considered in this chapter, homogeneous traffic sources were used, and their aggregate performance was determineded.

Parameters	Values
burstiness level b	2, 5, 10
aggregate load ρ	0.8 .. 1
individual load	0.08 .. 0.1
peak rate B_p	0.8 .. 1
switch buffer size B	10 .. 600 (cells)
burst duration T	10 .. 100 (slots)

Table 3.1: Simulation Parameter Values

3.5.1 Influence of Bursty Traffic on Performance

The bursty nature of traffic is the primary cause of congestion and cell loss. In this subsection, the influence of the burstiness level and the burst duration of the traffic on performance is examined.

In the first experiment, the effect of the burstiness level was examined. Ten homogeneous sources were used in this experiment. Cells from each source had to travel through the four switch nodes to reach their separate destinations (see the network configuration in Figure 3.1). Each source had a burst duration of 10 time slots. The burstiness level was varied from 2 to 10 while the aggregate load of the network increased progressively, from 0.8 to 1. No cells were generated in the idle state. The switch buffer size was fixed at 60.

Figures 3.3, 3.4 and 3.5 show the effect of the burstiness level of the traffic on its delay, coefficient of variation of delay, and cell loss performance, respectively. The performance measures shown in these figures are the mean values over all sources. The horizontal axis of the three figures is the aggregate load level of the network. The vertical axis of Figure 3.3 is the average end-to-end queueing delay of all sources (expressed in time slots). The vertical axis of Figure 3.4 is the coefficient of variation of the delay averaged over all sources. The vertical axis of Figure 3.5 is the cell loss, expressed as the ratio of the lost cells of all sources to the total number of cells generated by all sources.

The results demonstrate that queueing delay and cell loss increase as the load increases. Increasing the load level causes a sharp increase in cell loss (see Figure 3.5), while queueing delay is limited by the switch buffer size (see Figure 3.3). Since all sources are of an on/off type, a queue builds up at the switch when there are simultaneous arrivals of bursts from different sources. Cells are dropped when the combination of bursts exceeds the capacity of the switch. Queues may gradually clear in the following silent period. For sources with a burstiness level of 10, when the load level approaches 1, the average queueing delay is about half the buffer size at the switch, and cell loss is 9%.

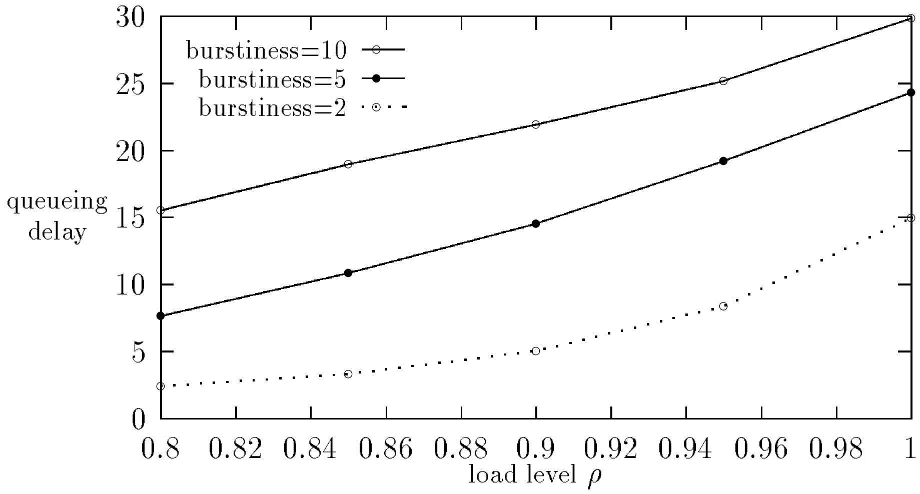

Figure 3.3: The effect of burstiness on end-to-end queueing delay

The coefficient of variation of delay decreases when the load level increases. The coefficient of variation of delay decreases because of the finite buffer size and the high load level. When the load level is high, most cells experience high delay, and the variation of delay is low. In all the experiments presented in this thesis, the coefficient of variation of delay follows the same pattern: the higher the queueing delay, the lower the coefficient of variation of the delay. For this reason, the results in the remaining sections will present only the mean queueing delay and cell loss performance metrics; coefficient of variation of delay will be ignored.

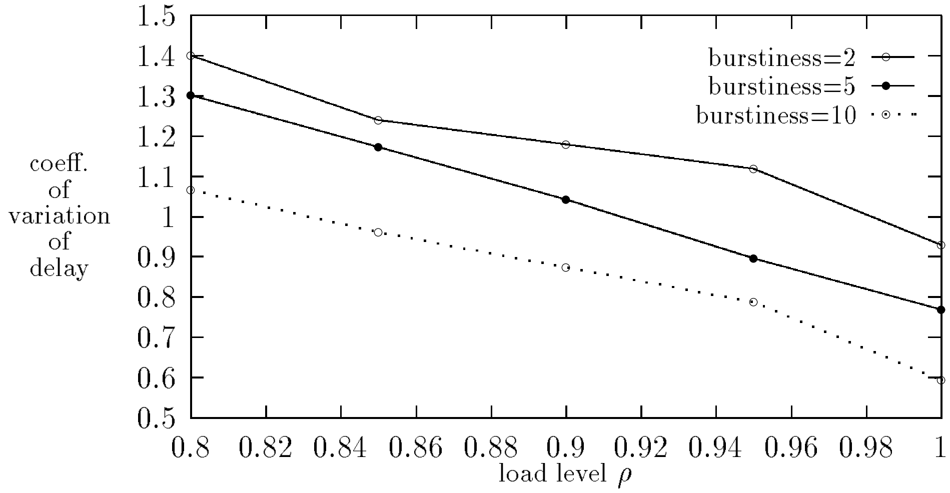

Figure 3.4: The effect of burstiness on coefficient of variation of queueing delay

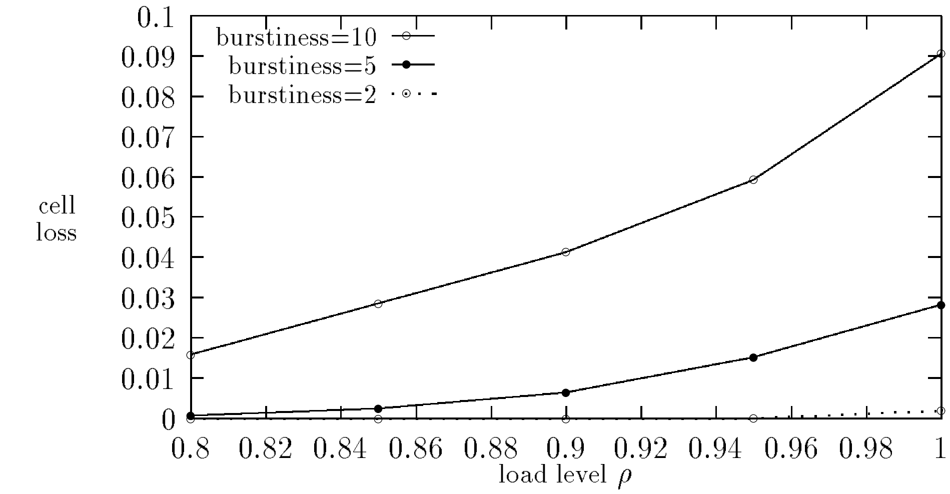

Figure 3.5: The effect of burstiness on cell loss

The burstiness level of the traffic has a large influence on performance. The higher the burstiness level, the larger is the end-to-end queueing delay and the higher is the cell loss. When the load level is low, the network can tolerate some burstiness in traffic. When the load level increases, cell loss for traffic with a high burstiness level increases faster than for traffic with a low burstiness level.

In this simulation setting, the switch buffer size should be able to accommodate a burst of cells since the mean burst duration of sources is 10. That is, the mean burst length is about $10 * B_p$, which is substantially smaller than 60. But due to the randomness and burstiness of the traffic sources, they still suffer a high percentage of cell loss, especially when the load level is high and/or the burstiness level is high.

The burst duration T also has a significant influence on the performance of the network. If a burst has more cells than the switch buffer can handle, the excess cells will be lost. In the second experiment, the performance of bursty traffic with different burst durations is examined. The same traffic sources as in the previous simulation experiment were used, but with different burst durations. Similarly, the switch buffer size was fixed at $B = 60$. The burst duration was varied from 10 to 100. The burstiness level was fixed at $b = 10$.

Figures 3.6 and 3.7 show the effect of burst duration on performance. It is clear that performance strongly depends on the burst duration. That is, the higher the load level, or the longer the burst duration, the poorer is the delay and loss performance. Traffic with long burst durations, even when the network load level is not very high, suffers high cell loss. For example, traffic with a burst duration of $T = 50$ loses 11% of its cells when the load level is $\rho = 0.8$.

It is clear that the switch buffer size of $B = 60$ cannot provide acceptable loss performance for sources with burst durations longer than 10. With increasing load level, the situation becomes even worse. The source with $T = 100$ loses more than 26% of its cells when the load level ρ is equal to 1.

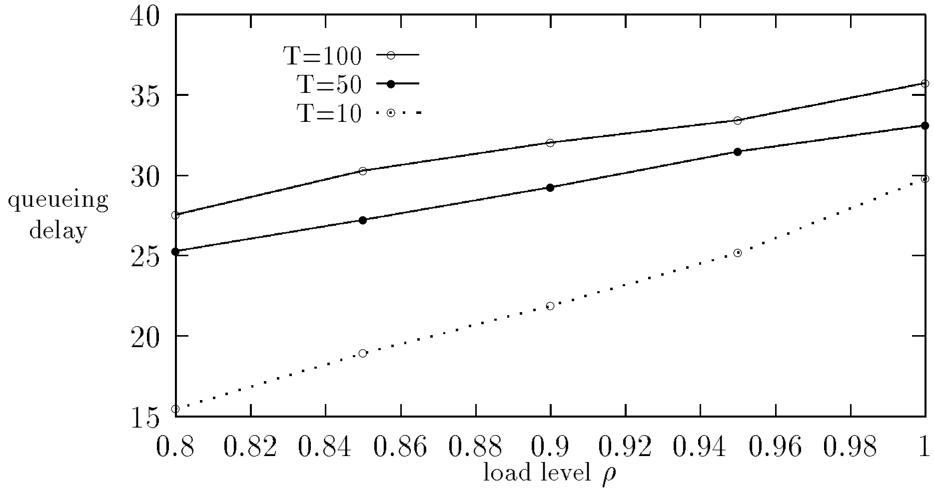

Figure 3.6: The effect of burst duration on queueing delay

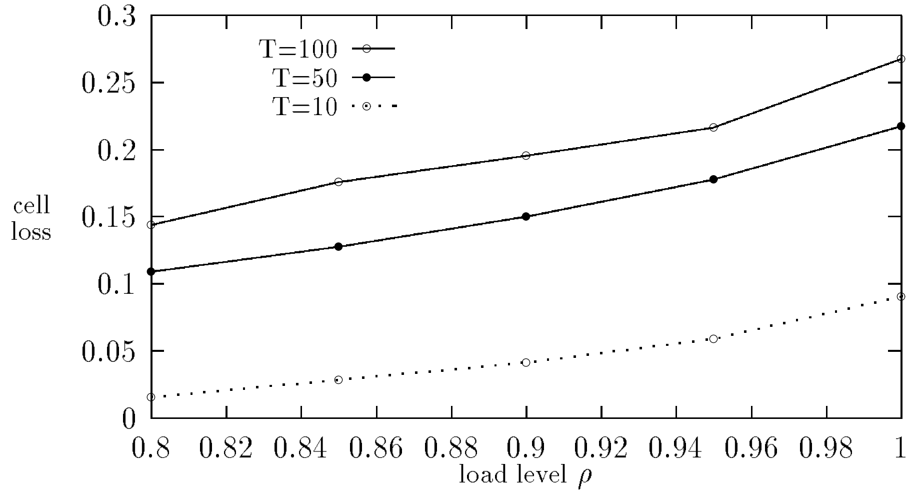

Figure 3.7: The effect of burst duration on cell loss

Comparing Figures 3.3 and 3.5 with Figures 3.6 and 3.7 (noting the difference in scales), it is clear that the burst duration has a greater impact on both delay performance and loss performance than does the burstiness level. Traffic with long burst durations is more difficult to handle than traffic with short burst durations, even though the latter may have a higher burstiness level.

3.5.2 Buffer Requirements at Switches

Buffer space at switches provides a means to absorb a temporary overload of the network. It reduces cell loss inside the network, but at the cost of increased queueing delay. In this subsection, switch buffer requirements for traffic with different burstiness levels and different burst durations are studied. The same traffic sources as in the previous subsection were used. The load was kept constant at $\rho = 0.95$. The choice of $\rho = 0.95$ is based on the following reasoning: a low load level (*e.g.*, below 0.8) is not interesting because the probability of congestion is too low, and a very high load level (*e.g.*, 0.99 or 1) is an unrealistic operating point for a network.

In the first experiment, the influence of traffic burstiness level on switch buffer requirements was studied. The burstiness level was varied from 2 to 10. The burst duration was fixed at $T = 10$, and the switch buffer size ranged from 10 to 300 cells.

Figures 3.8 and 3.9 show the effect of buffer size B on queueing delay and cell loss performance, respectively. Clearly, switch buffers can reduce cell loss to a great degree. Cell loss drops sharply at first as the buffer size increases, without incurring too much queueing delay. Since the burst duration of traffic sources used here was short ($T = 10$), increasing the buffer size was effective at reducing the cell loss.

The effect of increasing switch buffer size diminishes with the increase. For example, for traffic with a burstiness level $b = 10$, increasing the buffer size from 20 to 120 reduces cell loss from 14.3% to 2.5%, while increasing it from 200 to 300 reduces cell loss from 1.02% to 0.28%.

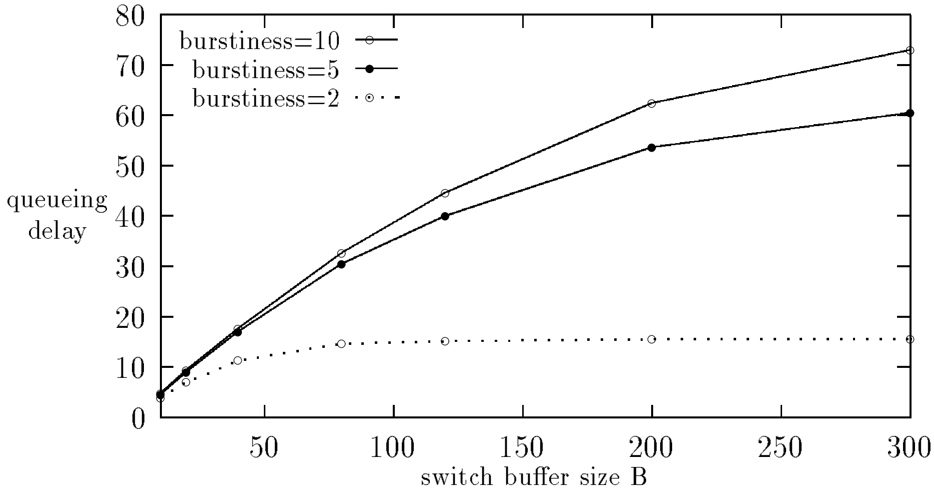

Figure 3.8: The effect of burstiness level and buffer size on delay

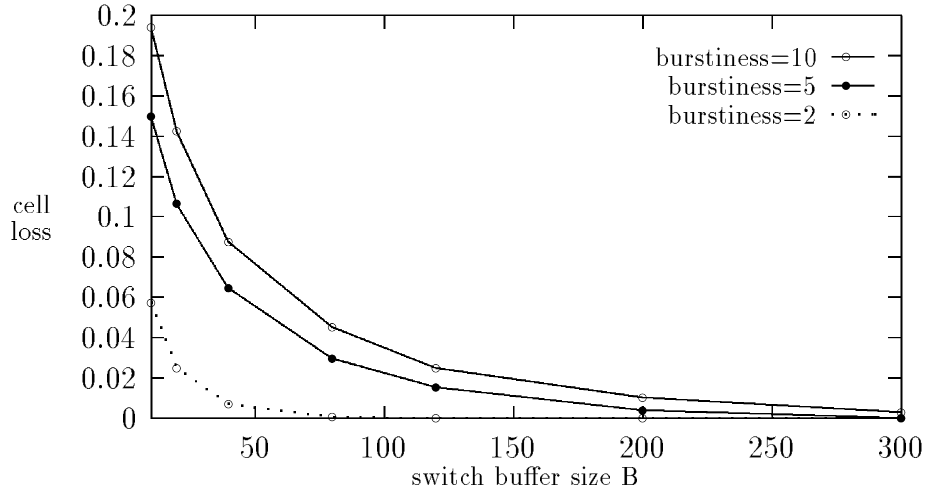

Figure 3.9: The effect of burstiness level and buffer size on cell loss

The burstiness level b has a major impact on buffer requirements at the switches. For a source with a low burstiness level and short burst durations, *e.g.*, $b = 2$ and $T = 10$, a switch buffer size of 40 can reduce cell loss to 0.7%. For a source with a high burstiness level, more switch buffer space is needed to handle cells arriving in a burst. For example, for a source with a burstiness level $b = 10$ and mean burst duration $T = 10$, a switch buffer size of 240 is required to limit the cell loss to 0.7%.

In the second experiment, the influence of burst duration on performance was examined. The same traffic sources as in the first experiment were used, but with different burst durations.

Figures 3.10 and 3.11 show the delay and loss performance of sources with different burst durations under different switch buffer sizes. For each of these curves, the buffer size at the switch is varied from 10 to 600. The burstiness level is fixed at $b = 10$. The mean burst duration is increased from 10 slots to 100 slots.

From the figures it is clear that the longer the burst durations in the traffic, the poorer is the network performance. At the same load level and with the same switch buffer size, traffic with a long burst duration has higher queueing delay and much higher cell loss than traffic with short burst durations. For example, when the switch buffer size is 100, traffic with a burst duration $T = 100$ loses about 20% of its cells, while the traffic with a burst duration $T = 10$ loses about 4% of its cells.

In order to achieve the same cell loss performance, traffic with long burst durations requires far more buffer space than traffic with short burst durations. For example, to reduce cell loss to 5%, traffic with a burst duration $T = 10$ requires a switch buffer of about 80 cells, while traffic with burst duration $T = 100$ requires a switch buffer of more than 600 cells.

Buffer space at the switch can accommodate some bursts of traffic. The switch buffer size represents a clear tradeoff between queueing delay and cell loss. Increasing the buffer sizes reduces cell loss, but at the expense of increased queueing delay.

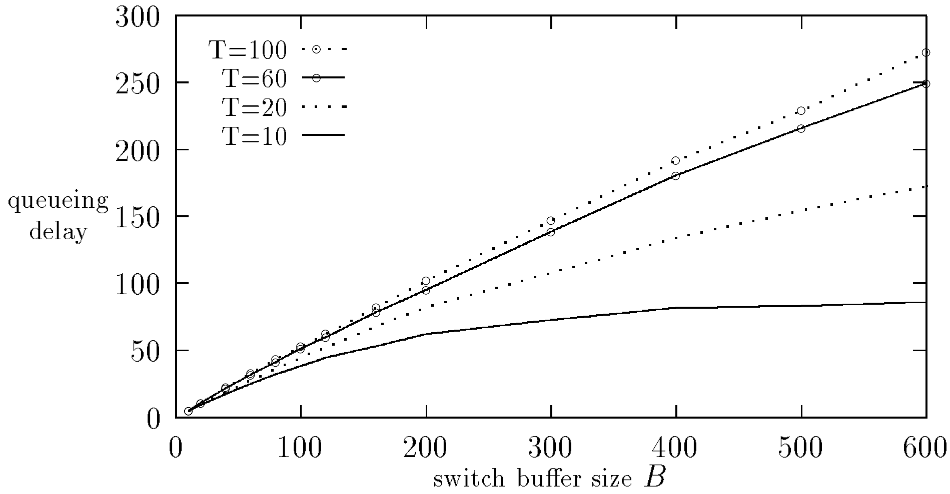

Figure 3.10: The effect of burst duration and buffer size on delay

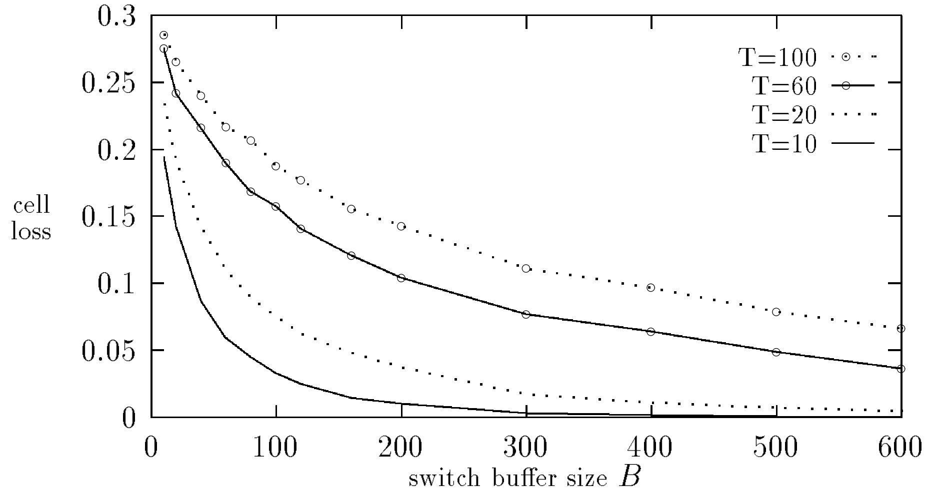

Figure 3.11: The effect of burst duration and buffer size on cell loss

When the buffer size increases, the cell loss drops sharply at first, then the curve becomes fairly flat (see Figure 3.11). The switch buffer size increase has very little impact on cell loss performance when the increase passes a certain point. This means that just increasing the buffer size at switches cannot solve the problem of cell loss satisfactorily when traffic is highly bursty and burst durations are long.

These results demonstrate that for sources with long burst durations, switch buffer sizes about an order of magnitude longer than T are required to achieve about 5% cell loss probability. As will be seen later in Chapter 4, traffic shaping provides a way to reduce the buffer requirements at the switches.

3.6 Summary of Results

The experiments presented in this chapter illustrate the performance impact of bursty traffic and the tradeoff between queueing delay and cell loss introduced by the switch buffer. Several observations can be made based on these results.

First, the aggregate load level of the network is a significant factor. When the load level is fairly low, the network can tolerate a certain degree of burstiness by buffering cell bursts at the switches. There is a low chance of congestion. When the load level is high, however, bursty traffic becomes the primary cause of congestion and cell loss.

Second, the burstiness level has a major impact on the network performance. The higher the burstiness level, the larger is the end-to-end queueing delay and the higher is the cell loss.

Third, the burst duration also has a significant impact on the performance of the network. The longer the burst duration, the poorer is the delay and loss performance. Traffic with long burst durations, even when the network load level is not very high, suffers high cell loss.

Fourth, providing buffer space at switches can offset a certain degree of bursty traffic. It cannot, however, eliminate cell loss totally. Too much buffer space may result in queueing delays that are unacceptably long.

Fifth, traffic with long burst durations requires more buffer space than traffic with short burst durations to achieve the same cell loss performance. Generally speaking, reasonable cell loss performance requires a switch buffer space that is at least an order of magnitude longer than the traffic's burst duration. Increasing switch buffer size is desirable only for short burst duration traffic, since the queueing delay incurred by the switch buffer limits the buffer dimension.

Since the bursty nature of traffic is the primary cause of congestion and cell loss, traffic shaping may be desirable to smooth this burstiness. In the next chapter, the effect of traffic shaping on performance will be described.

Chapter 4

Effects of Traffic Shaping on Performance

4.1 Introduction

This chapter describes the experiments conducted to evaluate the effect of traffic shaping on the performance of the network. The purpose of traffic shaping is to regulate the input traffic, to smooth its burstiness and to ensure that it adheres to the declared traffic parameters. Traffic shaping is performed in each individual source prior to entering the network.

The network configuration model with traffic shaping is shown in Figure 4.1. The network model is the same one as in the previous chapter, except that there is a traffic shaper for every traffic source. A modified leaky bucket scheme is used as the traffic shaper model. Traffic shaping is applied to the traffic sources studied in the previous chapter.

The purpose of these experiments is to study the effect of traffic shaping on network performance, the impact of different shaper parameters, and the nature of the traffic shaper.

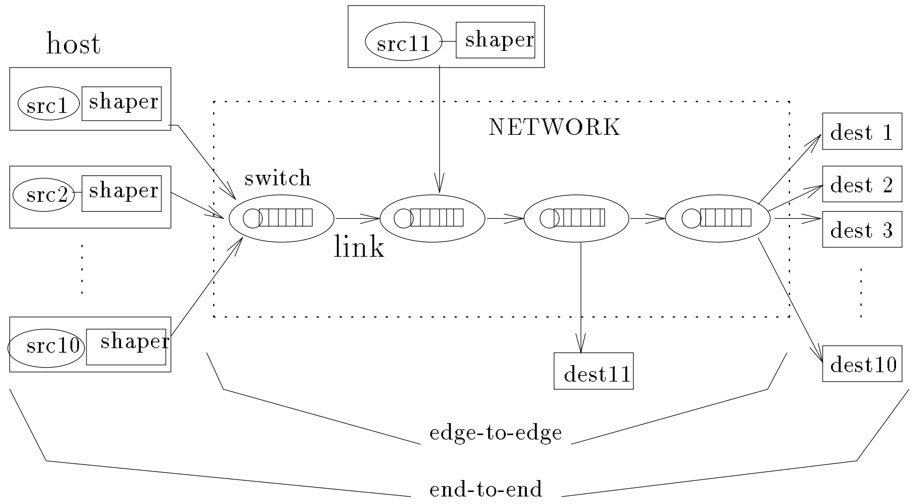

Figure 4.1: Network configuration with traffic shaping for simulation

4.2 Traffic Shaper Model

A traffic shaper shapes the input traffic to match its claimed traffic descriptors: mean bit rate, peak bit rate and burst length. A traffic shaper acts like a filter. It allows cells that arrive smoothly to proceed into the network, while it delays or drops cells that violate the bandwidth "contract".

In the traffic shaper employed (see Figure 4.2), a pool of tokens is available that contains at most M tokens. The pool capacity M limits the maximum possible length of cell bursts. The larger M is, the more cells that can be allowed to enter the network back-to-back at a time. The generation process of tokens is deterministic. If the token generation rate is r tokens/slot, one new token arrives every $1/r$ slots. The token generation rate determines the mean rate of the cells entering the network. A newly generated token will be stored in the pool if the pool contains fewer than M tokens. Otherwise, the token will be discarded. An arriving cell will enter a finite input buffer first. If a cell in the input buffer finds the token pool non-empty, it will

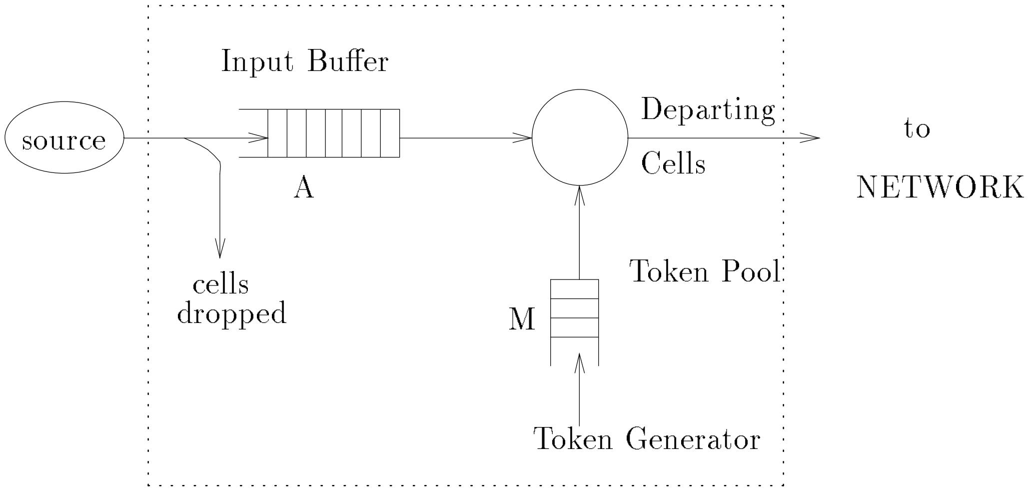

Figure 4.2: Traffic shaper model

take a token and enter the network immediately. If the input buffer is full, the cell will be dropped. The size of the input buffer (denoted as A) represents a tradeoff between cell delay and cell loss in the shaper. The larger the input buffer size, the higher is the maximum queueing delay, but the lower is the cell loss probability.

4.3 Performance Measures

Performance can be assessed at the network-level (edge-to-edge performance) or at the user-level (end-to-end performance), as shown in Figure 4.1. Network-level performance is measured from network entry access point to exit access point, *i.e.*, queueing delay and cell loss that connections experience inside the network. User-level performance is measured from the time of cell generation at the source (src) until its delivery to the destination host (dest). The user-level performance measures are queueing delay and cell loss that connections incur before they enter the network (*i.e.*, at the shaper) plus after they enter the network (*i.e.*, including network-level performance).

These two levels of performance may differ since the introduction of the traffic shaper may cause extra delay and cell loss. Both levels of the performance measure are studied in this chapter.

4.4 Description of Experiments

The results presented in this chapter were obtained with the network model and the traffic shaper model described in the previous sections, and the traffic model described in the previous chapter.

The general effect of traffic shaping is discussed first, followed by the impact of traffic shaping with different shaper parameters. Shaper parameters are: the input buffer size A, the token pool size M, and the token generation rate r. These parameters determine the characteristics of the traffic after shaping, and hence determine the performance, both network-level and user-level.

The buffer requirements at switches were studied with traffic shaping. Various features of traffic shaping, such as its protective nature and conservative nature, are also discussed. In most scenarios, except where specified to the contrary, homogeneous traffic sources were used, and their aggregate performance was studied.

Every simulation experiment ran long enough so that the statistical variation of a performance measure was within 1% of the mean value of this performance measure with 95% confidence, as explained in the previous chapter.

4.5 General Effects of Traffic Shaping

In this section, the general effect of traffic shaping on network performance is considered. Ten homogeneous sources were used in this simulation experiment. Cells from each source must travel through the four switch nodes to reach their destinations (see the network configuration in Figure 4.1). Each source had a burst duration of $T = 10$

and a burstiness level of $b = 10$. The individual load level ranged from 0.08 to 0.1, so the aggregate load level increased from 0.8 to 1.0. The switch buffer size was fixed at $B = 60$.

In this experiment, all traffic sources were subjected to traffic shaping. The traffic shaper used an input buffer size $A = 60$, and a token pool size $M = 20$. Each shaper generated tokens at a rate $r = 0.11$ tokens/slot.

Figure 4.3 shows the effect of traffic shaping on queueing delay. For comparison purposes, the delay performance of traffic with the same characteristics in the network environment with the same setting but subjected to no traffic shaping is also included. This case is indicated by the label "no-shaping".

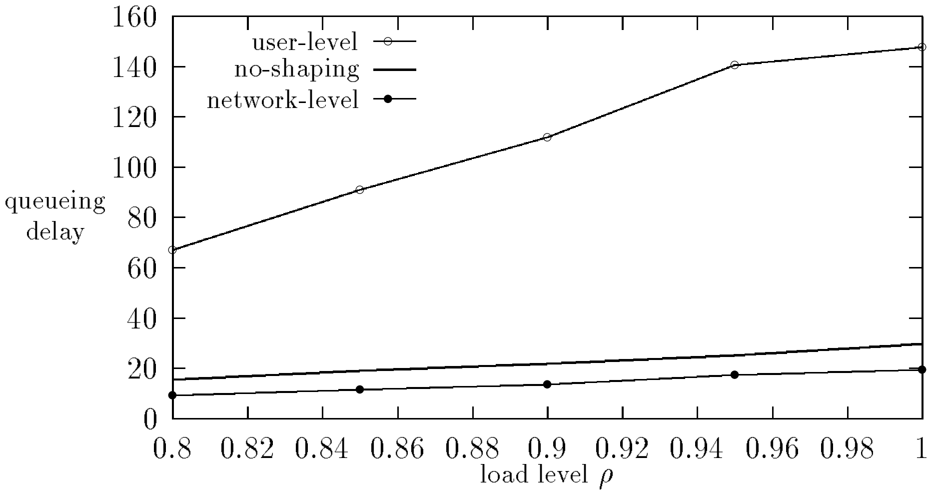

Figure 4.3: The effect of traffic shaping on queueing delay

Traffic shaping incurs a high queueing delay for traffic at the shaper. Since the traffic was very bursty, many cells arriving in a burst were forced to wait at the shaper before they were eligible to enter the network. Delays will be accumulated for cells generated towards the end of each burst since the token pool is usually empty then. The shaper buffer gradually clears between bursts. From the user-level perspective, traffic experiences a higher delay than without traffic shaping. Much of this delay happens in the shaper.

From the network-level perspective, since traffic shaping provides a network environment where all traffic has been smoothed, cells experience little delay inside the network. This is shown clearly in Figure 4.3.

Figure 4.4 shows the effect of traffic shaping on cell loss. Cell loss inside the network is reduced to a great extent. With the help of an input buffer, cell loss in the shaper is also limited. Results concerning the impact of the input buffer on performance are presented in Section 4.6.1 (Figures 4.5 and 4.6). From the network-level perspective, traffic shaping, with or without an input buffer, provides better loss performance than when there is no traffic shaping.

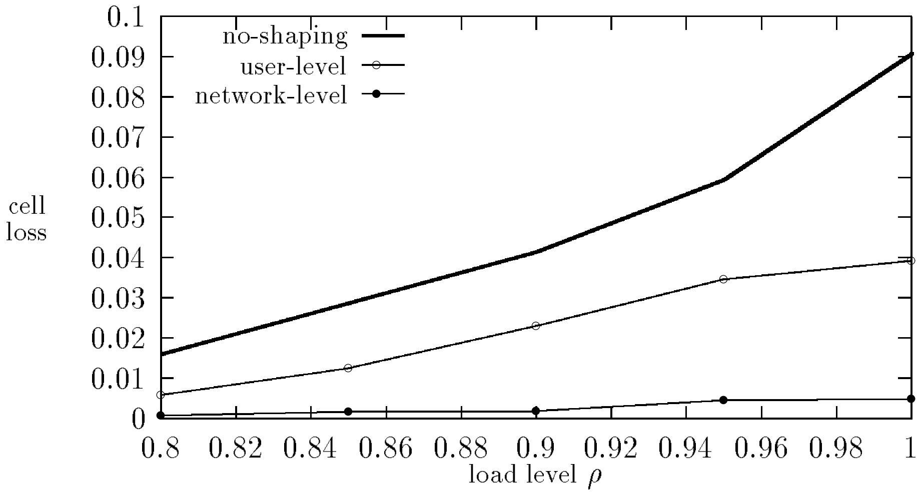

Figure 4.4: The effect of traffic shaping on cell loss

From the network-level perspective, the increase in the aggregate network load level has little effect on performance. Since all traffic is smooth, the network can accommodate a high load level without incurring too much delay or cell loss inside the network.

These results demonstrate that traffic shaping is very desirable for loss-sensitive traffic, such as data transfer, since it can greatly reduce cell loss. The extra delay, however, makes it less desirable for delay-sensitive traffic, such as real-time video, since this traffic may not be able to tolerate the extra delay introduced by traffic shaping.

4.6 Effects of Traffic Shaper Parameters

Shaper parameters determine the characteristics of traffic after shaping, which in turn determines the performance of the network, both user-level and network-level. This section presents the results of experiments run to determine the effects of different shaper parameters, namely, the input buffer size A, the token pool size M, and the token generation rate r.

In this section, the network load level is fixed at $\rho = 0.95$. Each traffic source has a burstiness level of $b = 10$ and a burst duration of $T = 10$. The switch buffer size is fixed at $B = 60$.

4.6.1 Input Buffer Size

The input buffer size A controls the maximum number of cells that are waiting for access to the network. It provides a tradeoff between the waiting time at the shaper and cell loss.

For the first experiment, the token pool size at the shaper was fixed at $M = 10$ and the token generation rate was fixed at $r = 0.11$. The input buffer size A was varied from 20 to 120 in steps of 20 so as to study its effect on network performance.

Figure 4.5 shows how the input buffer size A affects queueing delay. When A increases, the user-level queueing delay increases rapidly while the network-level queueing delay stays fairly constant. That means that most of the increase in queueing delay occurs at the shaper.

When a long burst of cells arrives, tokens saved in the token pool will soon be used up, and all the remaining cells in the input buffer have to enter the network at the token generation rate. For a given cell, its waiting time is the number of cells waiting before it in the buffer multiplied by the interval between token generations. If the token generation rate is 0.1, that means tokens are generated at an interval of 10 time slots. A cell waiting at the end of an input buffer which has a size of $A = 100$ will have to wait for 1000 time slots to enter the network! This fact argues that a small input buffer size might be more suitable.

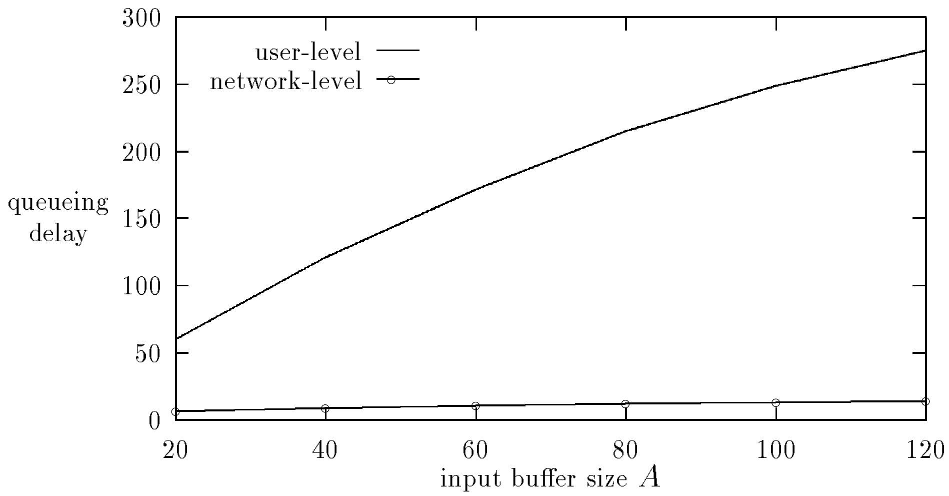

Figure 4.5: The effect of the input buffer size on queueing delay

Since cells are slowed down at the shaper and enter the network smoothly, they experience little delay inside the network. The network-level delay is limited by the switch buffer size. If the switch buffer size is small, the delay inside the network cannot be too high. The input buffer size has little effect on network-level delay. This is shown clearly in Figure 4.5.

Figure 4.6 shows the effect of input buffer size A on cell loss. As A increases, user-level cell loss drops quickly at first, then begins to flatten out. That is, increasing the input buffer size has a strong effect at first, but increasing it past a certain point has little effect in improving loss performance. For example, increasing A from 20 to 40 reduces cell loss from 14.7% to 7.6%, an improvement of 48%, while increasing A from 100 to 120 reduces cell loss from 1.7% to 1.1%, an improvement of 35%.

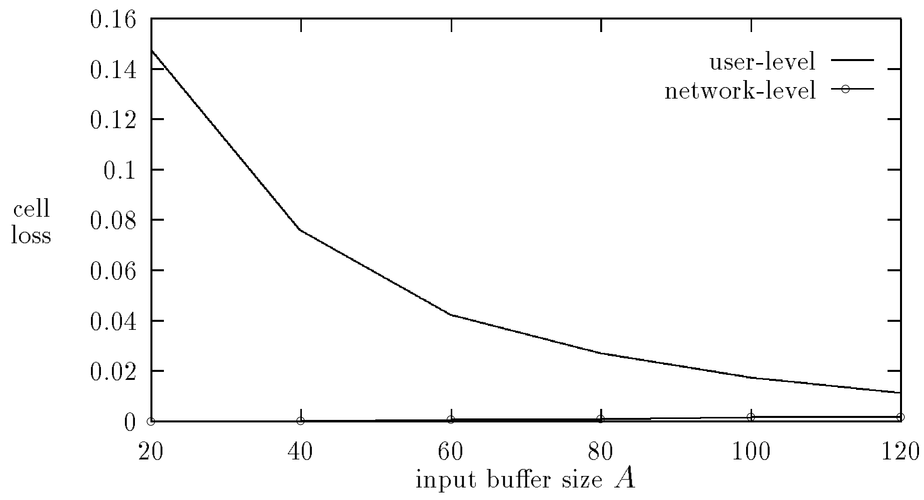

Figure 4.6: The effect of the input buffer size on cell loss

From the network-level perspective, since traffic is smoothed, it experiences little cell loss inside the network. As A increases, fewer cells are dropped at the shaper, meaning more cells are allowed to enter the network. So there is a slight increase in network-level cell loss as A increases.

These results demonstrate that an input buffer reduces cell loss at the cost of high queueing delay at the shaper. The choice of A determines the tradeoff between this cell loss and the queueing delay. If A is too big, it will result in an unacceptable queueing delay at the shaper. If A is too small, however, cell loss at the shaper will be too high. Choosing a suitable A depends on the characteristics of the traffic source such as its burstiness level and burst duration, and the other shaper parameters as well. For the traffic discussed here, around 60 to 80 is the suitable range for A.

4.6.2 Token Pool Size

The size of the token pool M determines how many tokens can be "saved up" for a given traffic source, and thus the maximum number of cells that can be injected into the network back-to-back at a time. The larger the token pool size, the more cells can enter the network at a burst, so the more bursty is the traffic. This subsection presents the results of an experiment conducted to study the effect of M on performance.

For this experiment, the input buffer size was fixed at 60 and the token generation rate was fixed at 0.11. The token pool size was varied from 10 to 80 in steps of 10.

Figure 4.7 shows how the token pool size M affects the queueing delay. When M increases, more tokens can be saved in the token pool, so cells arriving in bursts may spend less time waiting at the shaper, hence user-level queueing delay decreases. The network-level queueing delay is determined by the switch buffer size so M has little impact on it.

Figure 4.8 shows the effect of M on cell loss. Increasing M results in fewer cells lost in the shaper. Allowing more cells to enter the network in a burst, however, obviously leads to more cells lost inside the network. The larger M is, the bigger is the burst that is permitted to enter the network, and therefore the higher is the network-level cell loss.

It is interesting to note that in both figures, the user-level and network-level curves tend to be close together when M increases. It can be expected that when M is big enough, cells will experience little delay or loss in the shaper. Network-level performance becomes dominant. The situation is similar to the case of no-shaping.

The choice of a suitable M also depends on the characteristics of the traffic source, *i.e.*, its burstiness level and burst duration, and the other shaper parameters as well. For the traffic discussed here, the suitable range for M is from 10 to 30.

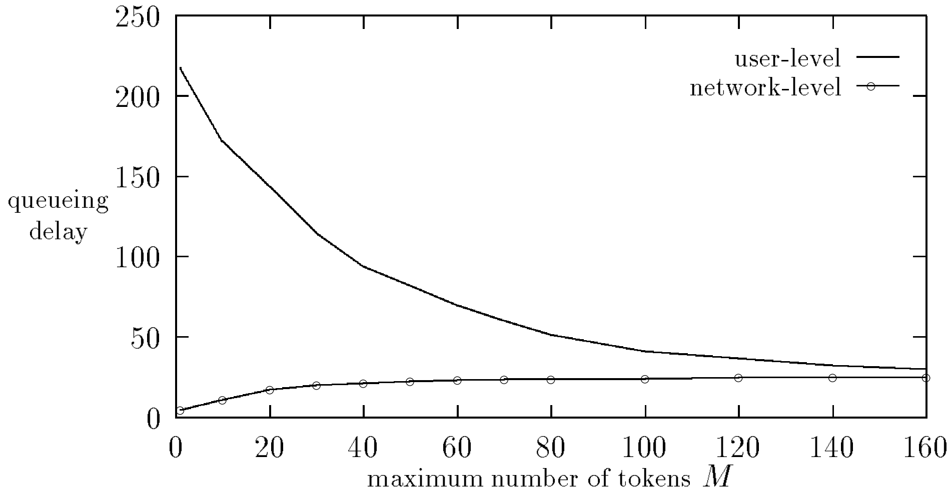

Figure 4.7: The effect of token pool size on queueing delay

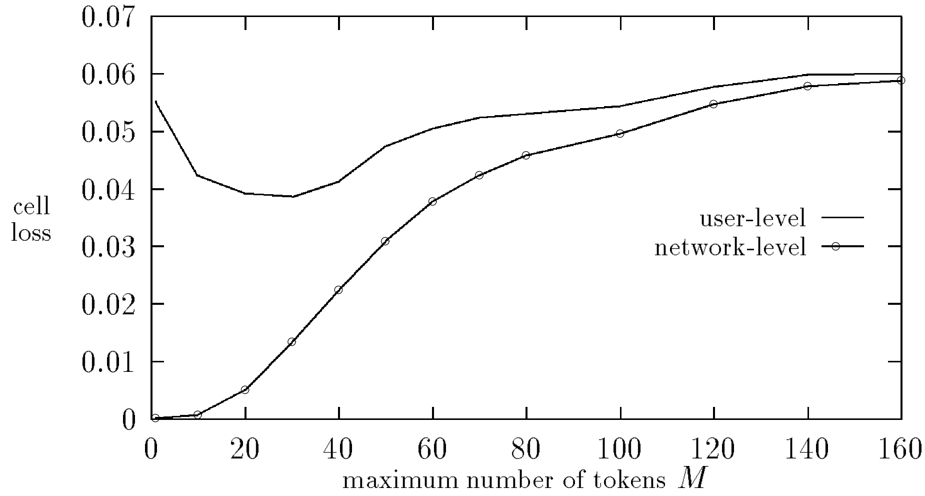

Figure 4.8: The effect of token pool size on cell loss

4.6.3 Combined Effect of Input Buffer Size and Token Pool Size

For the traffic shaper to function well and give optimal performance, the choice of shaper parameters is critical. An appropriate choice of shaper parameters depends on the characteristics of the traffic and the values of other shaper parameters. This subsection reports on an experiment conducted to study the combined effect of two shaper parameters, the input buffer size A and the token pool size M.

Figures 4.9 and 4.10 show the queueing delay and cell loss of the traffic with different shaper parameters. In these curves, M was varied from 10 to 40, while A was increased from 20 to 120. The token generation rate was fixed at $r = 0.11$.

For delay performance, the choice of M clearly represents a tradeoff between user-level performance and network-level performance. A larger M means that more tokens can be saved up, so more cells can enter the network at a burst. Larger M results in lower delay in the shaper but more delay inside the network. The total delay performance (*i.e.*, delay in the shaper plus delay in switches) is better for larger M than for smaller M.

User-level delay increases with the increase in the input buffer size A. A has little effect on the network-level delay performance, which is limited by the switch buffer size.

For loss performance, A and M have a combined effect. If A is small, a larger M results in lower cell loss in the shaper, although cell loss inside the network is higher. Since input buffer space is not enough to accommodate cell bursts, a large M will allow more of them to enter the network instead of being dropped at the shaper. These long cell bursts, however, may cause congestion and high cell loss inside the network.

When A is large, the shaper has enough space to buffer and smooth the cell bursts, and a relatively small M gives better performance. The smaller the value of M, the lower is the cell loss inside the network. This result demonstrates that A and M should not be chosen independently.

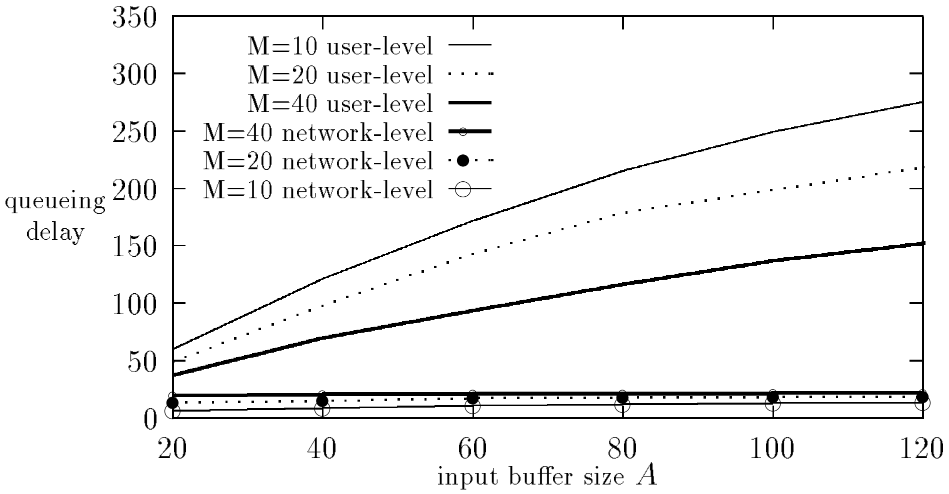

Figure 4.9: The joint effect of A and M on queueing delay

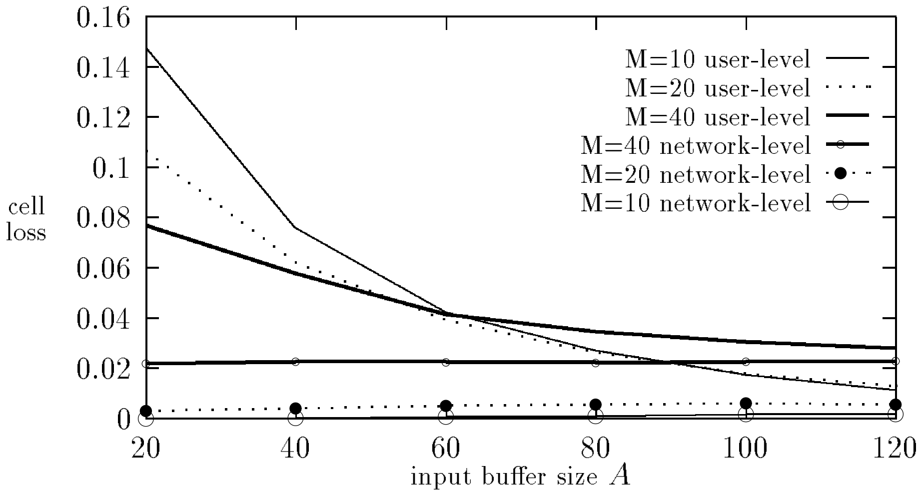

Figure 4.10: The joint effect of A and M on cell loss

4.6.4 Token Generation Rate

The token generation rate of a traffic shaper, r, determines the mean rate at which cells enter the network. If the token generation rate is lower than the mean input cell rate, obviously many cells will be dropped in the shaper. The token generation rate also represents the amount of the bandwidth allocated to the call, *i.e.*, the token generation rate is set corresponding to the bandwidth allocated to the source. For bursty traffic, if the call just reserves the bandwidth that is equal to its mean rate, it may lose many cells when they come in a burst. A bursty traffic source requires a token generation rate higher than its mean cell rate, *i.e.*, an amount of bandwidth higher than its mean rate. Since the traffic does not actually use all the bandwidth allocated to it, the network may allow a little over-reservation of the bandwidth in the sense that the sum of the token generation rates of all sources can be higher than 1. But too much over-reservation may cause the performance of the network to deteriorate. In this subsection, an experiment is described that studied the effect of the parameter r on overall performance.

In this experiment, the traffic sources were the same as in the previous section. That is, the load level was fixed at $\rho = 0.95$. The average cell rate of each source was $B_m = 0.095$. Each source had a burstiness level $b = 10$ and a burst duration $T = 10$.

For the traffic shaper, the input buffer size was fixed at $A = 80$ and the token pool size was fixed at $M = 10$. The token generation rate r was varied from 1.0 times the mean cell rate to 2.0 times the mean rate, that is, from 0.095 to 0.19. The sum of all r's thus ranged from 0.95 to 1.9. Bandwidth is over-reserved from the point $r = 0.1$, *i.e.*, when the ratio of the token rate over the mean rate is greater than 1.0.

Figures 4.11 and 4.12 show the delay and cell loss performance, respectively, as a function of r. For comparison, the performance of the case where there is no traffic shaping is included and indicated as "no-shaping".

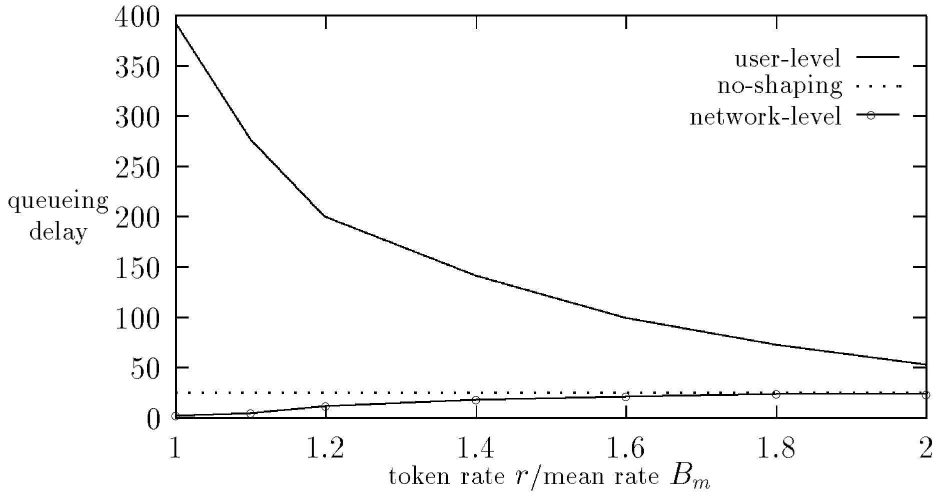

Figure 4.11: The effect of token generation rate on delay

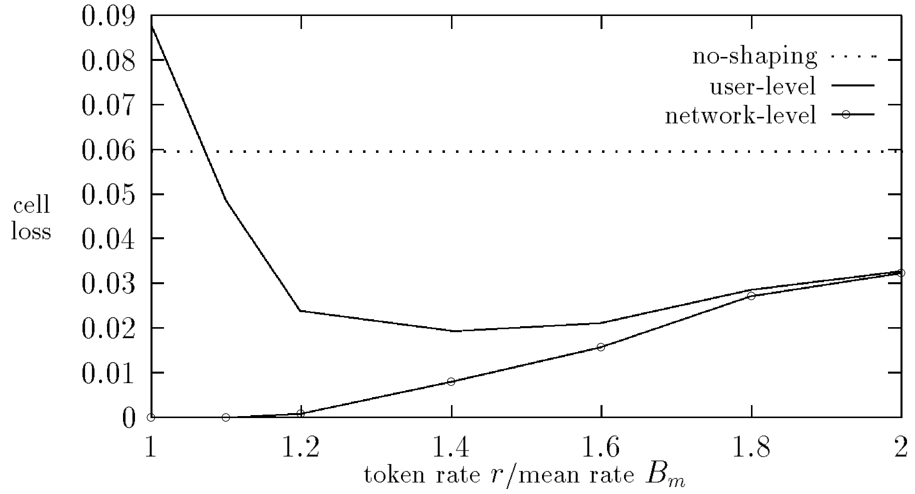

Figure 4.12: The effect of token generation rate on cell loss

When the token generation rate is equal to the mean rate of the input traffic, user-level queueing delay and cell loss are high (*e.g.*, a queueing delay of 375 slots and 8.9% cell loss). For bursty traffic sources, many cells arrive in bursts and at a rate much higher than the source's mean rate (*i.e.*, the peak cell rate). If the token generation rate is equal to the mean cell rate, cells in a burst will soon use up all the tokens in the token pool, and the remaining cells will have to wait at the input buffer of the shaper and enter the network at the token generation rate. The excess cells that the input buffer cannot handle will be dropped. This result demonstrates that making r equal to the mean rate is not a good choice.

Increasing r means more bandwidth is allocated to the call. The larger r is, the faster tokens are generated, and so cells spend less time waiting at the shaper and fewer cells are lost in the shaper. User-level queueing delay drops quickly as r increases. Network-level queueing delay increases a little when r increases, since more cells enter the network. When r increases, the two curves begin to be close to each other. Fewer cells are lost in the shaper, but more cells begin to be dropped inside the network. The total cell loss begins to increase when r passes a certain point since too many cells are lost inside the network. When r is very big, no cells are dropped at the shaper. All cells dropped are dropped inside the network. This does not mean, however, that traffic shaping has no effect on performance. Since the token generation rate is still less than the peak rate, the shaper is still effective in smoothing the bursts. The total cell loss is still lower than in the case where there is no traffic shaping.

These results demonstrate that bursty traffic needs to reserve bandwidth a little higher than its mean rate. It also shows that network bandwidth can be over-reserved to a certain extent. For the traffic source discussed here, a token rate that is about 1.2 times the mean rate is a suitable choice. In this case, the sum of the token generation rate of all sources is 1.14, which implies that the bandwidth is over-reserved by 14%.

4.7 Protective Nature of Traffic Shaping

Traffic shaping forces the traffic to observe its allocated bandwidth so that the user's traffic characteristics will not adversely affect the overall performance of the network. In this way, it provides protection to the network from bursty traffic sources or misbehaving users. The performance of one user will not be affected by the behaviour of others, since traffic shaping provides a good network environment and good network-level performance. Once the network-level performance is not dominant, the user-level performance of a traffic source will depend primarily on the delay and cell loss in the shaper, which is determined by the source's own traffic characteristics. With traffic shaping, all sources get the same network-level performance. From the user-level perspective, sources with smooth traffic can obtain better performance than sources with bursty traffic. This section describes an experiment to study the performance of heterogeneous traffic sources.

In this experiment, all sources had the same burst duration, $T = 10$. Nine sources had a burstiness level of $b = 10$, and one source had a burstiness level of $b = 1.1$ (much smoother). The aggregate network load level was increased from 0.8 to 1. The switch buffer size was fixed at $B = 60$. The traffic shapers of all sources had the same set of parameters: $A = 60$, $M = 10$, and $r = 0.11$.

Figures 4.13 and 4.14 show the delay and loss performance of these sources in two cases: one with traffic shaping, and one without.

When there is no traffic shaping, all sources suffer similar performance degradation in the presence of bursty traffic. With traffic shaping, smooth traffic sources experience little delay and cell loss, while bursty traffic sources suffer performance degradation caused by their own burstiness. Bursty traffic sources experience most of their queueing delay and cell loss at the shaper. In this way, traffic shaping protects smooth traffic sources from the effects of bursty ones.

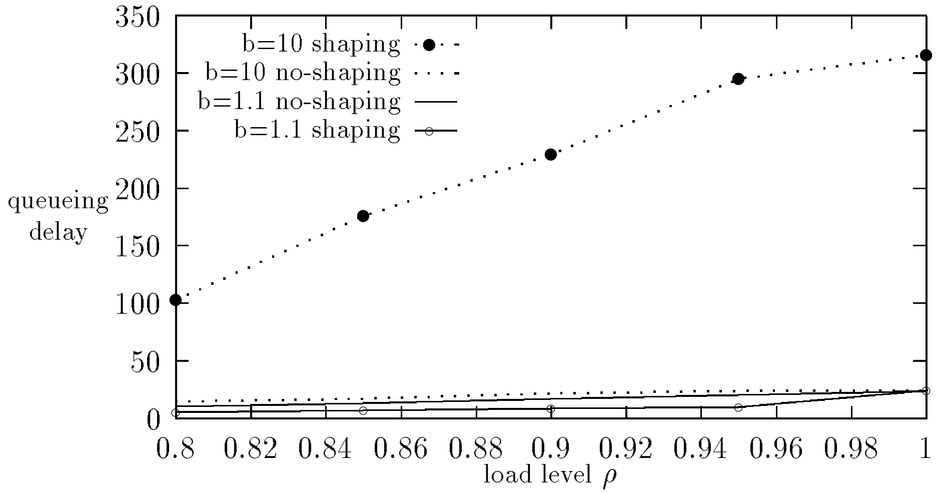

Figure 4.13: Queueing delay of different users

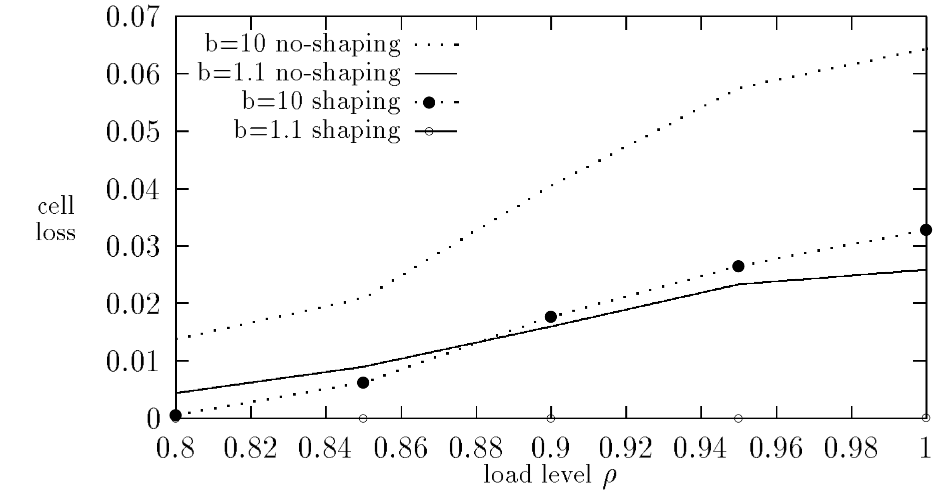

Figure 4.14: Cell loss of different users

The traffic shaper is transparent to smooth traffic while shaping bursty traffic to provide a network environment that has a low probability of congestion. It is desirable to be sure that once the traffic gets into the network, the traffic will not suffer performance degradation due to the misbehaviour of other traffic. The protective nature of traffic shaping makes traffic shaping very desirable in an environment that requires a certain degree of performance guarantee.

4.8 Buffer Requirements at Switches

Traffic shaping buffers excessive cells at the shaper instead of at the switches, *i.e.*, outside the network instead of inside the network. Since traffic shaping smoothes bursts of cells before they enter the network, traffic inside the network is very smooth, hence less buffer space is required at the switches. An experiment was performed to study buffer requirements at the switches with traffic shaping.

Figure 4.15 shows loss performance as a function of switch buffer size. With traffic shaping, a buffer size $B = 30$ can reduce the network-level cell loss to close to zero. By comparison, an 80 cell buffer is required to achieve the same user-level performance without traffic shaping. With traffic shaping, however, an input buffer is required to limit the cell loss probability at the shaper.

The results show that traffic shaping can greatly reduce the buffer requirements at the switches. For example, to obtain a cell loss probability of 0.01 without traffic shaping requires a 120 cell buffer space, while this same level of cell loss can be achieved with $B = 20$ with shaping.

Without traffic shaping, the buffer requirements at the switch depend on the burst duration of the traffic. Traffic with long burst durations needs large buffer space to accommodate bursts of cells (see Section 3.6.3). With traffic shaping, since cell flows inside the network are smoothed, buffer requirements at the switches are less dependent on the burst duration. The input buffer size at the shaper, however, is still dependent on the burst duration.

Figures 4.16 and 4.17 show the user-level delay and loss performance of traffic sources with different burst durations and under different input buffer sizes. The switch buffer size was fixed at $B = 60$. The input buffer size at the shaper was increased from 10 to 300. The burst duration was varied from 10 to 100.

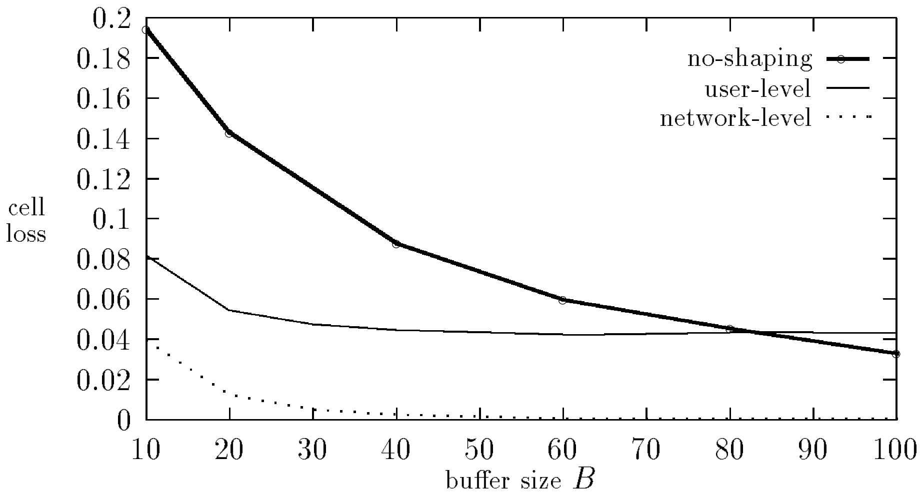

Figure 4.15: Cell loss performance with different buffer sizes

The burst duration has a significant impact on the performance achieved by an input buffer with different sizes. Sources with long burst durations cannot obtain an acceptable cell loss performance even with a large input buffer size.

Increasing the input buffer size reduces cell loss at the expense of an enormous delay in the shaper. For a traffic source with a given burst duration, increasing the input buffer size beyond a certain point has little effect on both delay and loss performance. For example, for the traffic source with a burst duration of $T = 10$, increasing the input buffer size from 10 to 100 reduces cell loss from 22% to 1.5%, while increasing A from 200 to 300 improves loss performance from 0.49% to 0.19%.

This result demonstrates that traffic shaping reduces buffer requirements at switches at the cost of buffer requirements at the shaper. It is also clear that the values of the shaper parameters should be based on the characteristics of the traffic source.

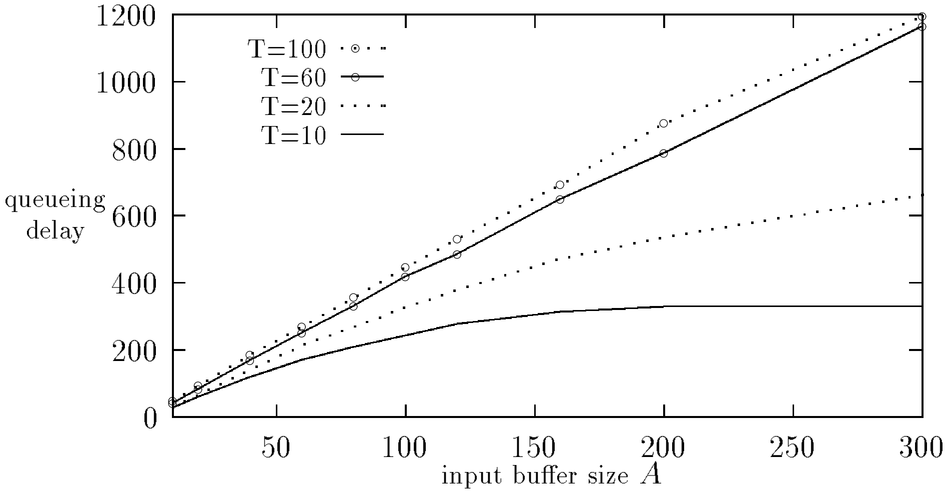

Figure 4.16: The effect of input buffer size and burst duration on delay

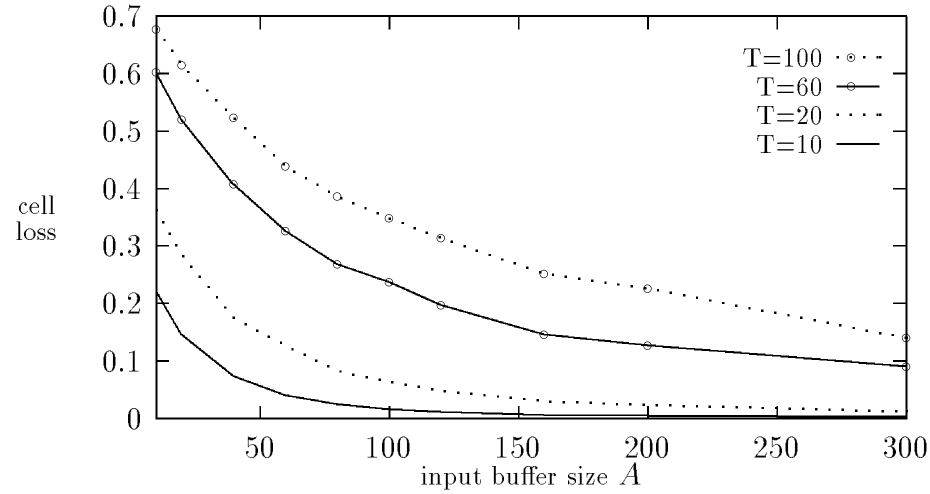

Figure 4.17: The effect of input buffer size and burst duration on cell loss

4.9 Conservative Nature of Traffic Shaping

In essence, traffic shaping slows down the source traffic to make it enter the network more gracefully. In some cases the network may be able to absorb long bursts of cells without being severely congested. For example, a long burst of cells from one source might happen to come to the network when all other sources are in their idle states. Without traffic shaping, such a burst could get through the network since there would be enough excess bandwidth to accommodate it. With traffic shaping, however, the burst will be delayed at the shaper and sent at the token generation rate or even be lost. In this sense, the traffic shaper may be too conservative in its approach.

In this experiment, all sources had the same burstiness level, $b = 10$. The burst duration was increased from 10 to 80. The traffic shapers of all sources had the same set of parameters: $A = 60$, $M = 10$, and $r = 0.11$.

Figure 4.18 shows the loss performance as a function of burst duration, both with shaping and without shaping. At first, when the burst duration is short, cell loss performance with traffic shaping, both from user perspective and network perspective, is better than the case without shaping. The user-level cell loss with shaping increases faster than the cell loss without shaping. The two curves cross at around $T = 40$. From that point on, traffic shaping results in poorer performance than in the case of no traffic shaping.

The longer the burst duration, the higher the cell loss. Almost all losses occur at the shaper. From the user-level perspective, when the burst duration is long, the performance with shaping is even worse than the performance without shaping.

The above results demonstrate that while traffic shaping limits the rate at which cells are allowed to enter the network, it also limits the chance of taking advantage of the burstiness of traffic.

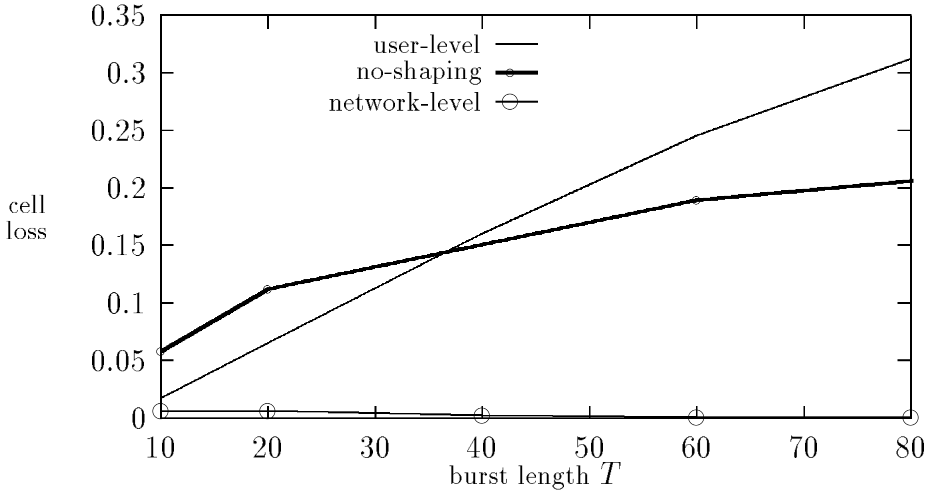

Figure 4.18: The effect of burst duration on cell loss

4.10 The Effect of Cross Traffic

The previous discussions in this chapter all concern ten connections travelling through four switch nodes. This section presents an experiment that considered the possibility of additional traffic arriving at later switches. This experiment constitutes a preliminary investigation of the impact of cross traffic on the network performance.

The cross traffic for this experiment was represented by the connection from src11 to dest11 (see Figure 4.1). Cells from src11 must travel through the second and third switches in the network to reach destination dest11. The characteristics of this connection, and whether it is shaped or not, will affect the performance of other connections.

In this experiment, ten long connections had the same burstiness level $b = 10$ and burst duration $T = 10$. Their individual load level was 0.08, so the ten sources contributed a load level of 0.8. The cross traffic source also had a burstiness level $b = 10$ and average burst duration $T = 10$. The load level of the cross traffic source was increased from 0.0 to 0.2 in steps of 0.05. The switch buffer size was fixed 60.

The ten long traffic sources all were shaped with the same shaper parameters: $A = 60$, $M = 10$ and $r = 0.1$. Two cases were considered for the cross traffic: shaping and no-shaping. In the case of shaping, the shaper parameters were set as $A = 60$, $M = 10$, and $r = 1.2 * B_m$. In the case of no-shaping, cells from the cross traffic source entered the network immediately after they were generated. Figures 4.19 and 4.20 show the effect of the cross traffic on the queueing delay and cell loss of the long connections. When the load level of the cross traffic source is 0, that means there is no cross traffic. Increasing the load level of the cross traffic causes the queueing delay and cell loss of the long connections to increase, whether the cross traffic is shaped or not.

If the cross traffic is not shaped, it causes greater performance degradation, both user-level and network-level, to the long connections than when the cross traffic is shaped, as shown in Figure 4.20. When the cross traffic is not shaped, both user-level and network-level cell loss performance degrade quickly with the increase in the load level. Actually, all performance degradation occurs inside the network. The queueing delay and cell loss that the long connections experience in the shaper are not affected by the cross traffic. Unshaped cross traffic causes very undesirable performance deterioration.

When the cross traffic is shaped (*i.e.*, all traffic accessing the network is smoothed), good network-level performance is maintained for all traffic sources, although there is a slight increase in cell loss due to the increase in the aggregate load level of the network. Notice that the differences between delays, whether user-level or network-level, are not significant. Since the aggregate load level considered here is very high, queueing delay is determined primarily by the switch buffer size and the input buffer size.

These results demonstrate that unshaped cross traffic may disturb the good network environment provided by traffic shaping, thus causing great performance degradation to other sources. Shaped cross traffic does not have such a damaging effect on other users. This result suggests that in a network environment that adopts a traffic shaping policy, all traffic sources accessing the network should be shaped.

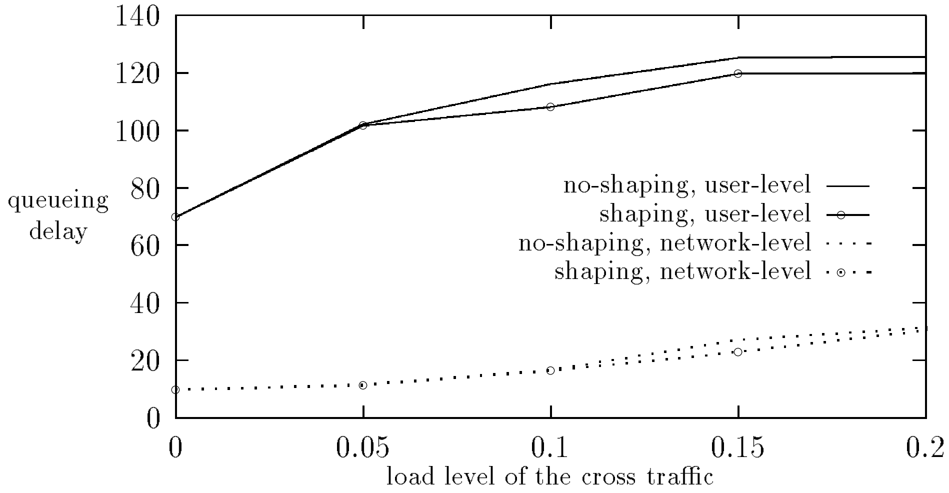

Figure 4.19: The effect of cross traffic on queueing delay

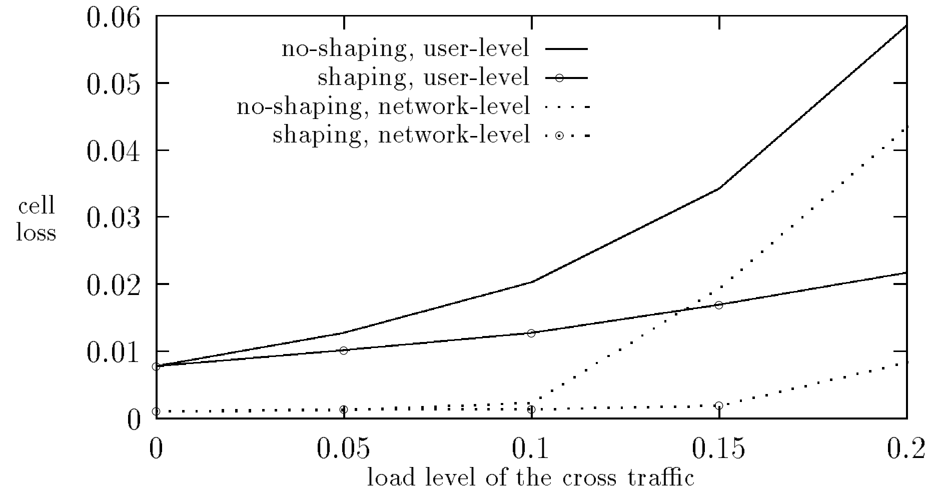

Figure 4.20: The effect of cross traffic on cell loss

4.11 Conclusions

The simulation results presented in this chapter reveal several consequences of traffic shaping in an ATM based network environment:

- Traffic shaping provides very good network-level performance. It can effectively reduce the probability of congestion inside the network which obviously leads to low queueing delay and low probability of cell loss. The process of traffic shaping itself is transparent to smooth traffic, while it shapes the bursty traffic.

- Traffic shaping may adversely affect user-level performance. The good network-level performance comes at the expense of extra queueing delay and cell loss at the shaper.

- Traffic shaping can reduce cell loss to a great extent at the cost of a great degree of extra delay. It is more desirable for loss-sensitive traffic than for delay-sensitive traffic.

- Traffic shaping can reduce the buffer requirements at switches. Since traffic shaping makes all traffic travelling inside the network smooth, a small switch buffer (*e.g.*, one-fourth of the size of the switch buffer size without traffic shaping) is enough to provide the same level of cell loss performance as in the case where there is no shaping.

- Shaper parameters have a significant impact on the characteristics of the traffic flow, as well as the performance. The choice of an appropriate set of shaper parameters depends on the characteristics of the source traffic, and the values of other shaper parameters as well.

- Traffic shaping provides protection to the network from the misbehaviour of other users as long as all traffic flows are shaped. Unshaped cross traffic may cause great performance degradation to other users.

- Traffic shaping is a conservative approach to congestion control. Traffic shaping reduces the probability of congestion by limiting the rate and burst length by which cells are allowed to enter the network. This may result in waste of bandwidth.

Chapter 5

Summary and Conclusions

In this final chapter, a summary of the thesis is presented along with its conclusion. This is followed by suggestions for future research.

5.1 Thesis Summary

The research reported in this thesis was undertaken to address two main goals:

1. to investigate the effects of traffic shaping on performance of an ATM-based network environment

2. to get a better understanding of the characteristics of bursty traffic: how they affect the performance of the network, and how they can be altered by a traffic shaper.

In Chapter 3, an ATM network simulation model and a MMPP traffic model were described. These models formed the basis for a series of experiments. In these experiments, the impact of traffic characteristics on the performance of the network were examined. There are a number of observations that can be gathered from this investigation.

First, it is confirmed that the bursty nature of traffic is the primary cause of cell loss and queueing delay. The burstiness level and the burst duration determine the performance. Experiments show that increasing the burstiness level can greatly worsen the delay and loss performance. For instance, when the burstiness level is increased from 5 to 10, the cell loss doubles. Performance also depends strongly on the duration of bursts. The longer the burst duration, the poorer is the delay and loss performance. Traffic that has a high burstiness level and long burst durations suffers the most severe performance degradation. For example, traffic with an average burst duration of 100 and a burstiness level of 10 loses 22% of its cells when the network load level is 0.85 and the switch buffer size is 60 cells.

Second, increasing switch buffer size cannot solve the problem of cell loss satisfactorily. Buffer size increases have a great effect at first, but will have little effect on performance beyond a certain range. In one experiment presented in this thesis, for example, increasing the switch buffer size from 10 to 80 reduced cell loss from 19% to 3%, while increasing the buffer size from 500 to 600 only reduced cell loss from 0.03% to 0.01%.

The buffer requirements at switches depend on both the burstiness level and the burst durations of the traffic. Traffic with long burst durations requires far more buffer space than traffic with short burst durations to achieve the same cell loss performance. Generally speaking, reasonable cell loss performance requires buffer space that is at least an order of magnitude longer than the burst duration of the source traffic. The ability to increase switch buffer size is also limited, since large buffers will result in an unacceptable queueing delay.

The effects of traffic shaping on performance based on a traffic shaper model are considered in Chapter 4. The experimental results provide some insight into the merits and limitations of traffic shaping.

Traffic shaping provides very good network level performance. After shaping, traffic flows that enter the network are smooth, so they cause fewer chances for congestion and experience little delay or loss inside the network, even when the network load is high. By properly setting the shaper parameters, traffic shaping can effectively eliminate cell loss inside the network.

Good network level performance, however, may come at the expense of extra queueing delay and cell loss at the traffic shaper. Cells may have to wait a long time to be eligible for access to the network or may even be dropped in the shaper.

Shaper parameters have a significant impact on performance. The input buffer size provides a clear tradeoff between waiting time and cell loss at the shaper. Larger input buffers can greatly reduce cell loss at the shaper by buffering cells arriving in bursts. However, it also incurs additional queueing delay. This delay limits the desirability of a large buffer.

The token pool size controls the maximum number of tokens that can be saved in the token pool, hence the maximum number of cells that can enter the network back-to-back at one time (*i.e.*, the length of a burst). A big token pool can reduce queueing delay and cell loss in the shaper. A bigger token pool, however, will let more cells enter the network in a longer burst, so network level performance degrades with an increase in token pool size.

The choice of the input buffer size and token pool size should not be made independently. If the input buffer is large, the token pool can be made small to effectively smooth the bursts. When the input buffer size is small, however, the token pool should be a little larger so as to reduce cell loss in the shaper due to lack of space in the input buffer.

The token generation rate is determined by how much bandwidth the connection is allocated. The token generation rate of the shaper also determines the mean rate at which cells enter the network. Experiments demonstrated that bursty sources need a bandwidth that is higher than its mean rate (e.g. 1.2 times its mean rate). Bandwidth can be over-reserved since bursty traffic may not use all the bandwidth allocated.

Traffic shaping can reduce buffer requirements at switches since all traffic transmitted inside the network is smooth. For example, a 20 cell switch buffer with traffic shaping can obtain the same level of network-level cell loss performance as a 120 cell switch buffer without traffic shaping.

Traffic shaping provides protection to the network from the misbehaviour of some users. Traffic shaping is transparent to smooth traffic, while bursty traffic will suffer performance degradation due to its own burstiness. Without traffic shaping, all sources will suffer similar performance deterioration in the presence of bursty traffic from any source.

Traffic shaping is a conservative approach to congestion control. While limiting the rate and the burst length at which cells are allowed to enter the network, traffic shaping also limits the chances of taking advantage of the burstiness of traffic. With shaping, some cells may have to wait at the shaper even though the network may be able to accommodate them.

The good network-level performance provided by traffic shaping requires cooperation from all traffic sources. Unshaped bursty traffic can cause great performance deterioration to other users. Shaped traffic does not have such a damaging effect.

5.2 Conclusions

The research presented in this thesis leads to the following conclusions.

First, the bursty nature of traffic, represented by its burstiness level and burst duration, has a significant impact on the network performance. The higher the burstiness level and the longer the burst duration, the poorer is the network performance. Traffic with long burst durations requires more buffer space than traffic with short burst durations to achieve the same cell loss performance. Traffic that has a high burstiness level and with long burst durations is very difficult to handle.

Second, traffic shaping provides good network-level performance, *i.e.*, low queueing delay and low cell loss inside the network. Traffic shaping, however, may adversely affect user-level performance. The good network-level performance comes at the expense of extra queueing delay and cell loss at the shaper. Traffic shaping can reduce cell loss to a great extent, but it does so at the cost of a great deal of extra delay. It is more desirable for loss-sensitive traffic than for delay-sensitive traffic.

Third, traffic shaping can reduce the buffer requirements at switches with the provision of an input buffer at the shaper. It also provides protection to the network from the misbehaviour of other users. In essence, traffic shaping is a conservative approach to congestion control. It may result in wasted bandwidth in some situations.

Finally, allowing unshaped traffic to enter the network will cause performance degradation to all other sources. In a network environment that adopts the policy of traffic shaping, all sources should be shaped.

5.3 Future Work

The research presented in this thesis can be extended in many ways.

One way to gain more insight from the experiments would be by mixing different traffic models together to obtain a more realistic mix of traffic. The characteristics of the traffic have a significant impact on performance, so a more realistic traffic model would give a better understanding of the effects of traffic shaping. The traffic characteristics and service requirements that ATM networks are supposed to support are very diverse. Heterogeneous sources should be examined to determine how well traffic shaping works under various traffic conditions. With heterogeneous sources, fairness is also an important issue to address.

The traffic studied in this thesis is one-way traffic. Two-way traffic can introduce complicated second-order effects (*e.g.*, higher cell loss, ACK-compression) [42, 49] that can reduce the effectiveness of traffic shaping. It would be interesting to investigate the behaviour of the traffic shaper under the presence of two-way traffic.

For a preventive-based congestion control scheme such as traffic shaping, it might also be desirable to have the source adjust its traffic flow according to feedback signals from the network. One reason for the unsuitablity of feedback-based congestion control schemes is that it may take too long for a feedback signal from the network to reach the source. With traffic shaping, however, a feedback mechanism can be added to the shaper. When the shaper finds that the queue length of its input buffer reaches a certain threshold, the shaper can send a feedback signal to the source to ask the source to adjust the source's traffic flow. The source can then slow down to observe its contract with the network or ask the network for more bandwidth so that the source can adjust the token generation rate of its shaper. Since the shaper is very close to the source, the feedback would be more timely than in the case that feedback signals come from somewhere inside the network.

The preliminary study of the effect of cross traffic in this thesis shows that cross traffic has a significant impact on the performance of other sources. It might be desirable to study traffic shaping under a network configuration more complicated than the model presented in this thesis and with more cross traffic sources. For example, the impact of multiple switches on the smoothness of the traffic inside the network would be of interest. Traffic that enters the network smoothly may be clustered together by a slow switch and lose its smoothness. Mismatching link speeds may also interfere with the effect of traffic shaping. Studying the effect of traffic shaping in a more realistic network environment that has multiple switches and multiple cross traffic sources will give more insights into traffic shaping and congestion control.

References

[1] H. Ahmadi and W. Denzel, "A Survey of Modern High-Performance Switching Techniques", *IEEE Journal on Selected Areas in Communication*, Vol. SAC-7, No. 7, Sept. 1989, pp. 1091-1103.

[2] S. Akhtar, "Congestion Control in a Fast Packet Switching Network", M.S. Thesis, Washington University, St. Louis, Missouri, 1987.

[3] A. Berger, "Performance Analysis of a Rate Control Throttle Where Tokens and Jobs Queue", *Proceedings of INFOCOMM 1990*, June 1990, pp. 30-38.

[4] M. Butto, E. Cavallero and A. Tonietti, "Effectiveness of the 'Leaky Bucket' Policing Mechanism in ATM Networks", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No.3, April 1991, pp. 335-342.

[5] "Broadband Aspects of ISDN", *CCITT Recommendation I.121*, Geneva, Switzerland, 1989.

[6] *CCITT Proposed Recommendation I.311*, June 1991.

[7] I. Cidon, I. Gopal and R. Guerin, "Bandwidth Management and Congestion Control in plaNET", *IEEE Communications Magazine*, Oct. 1991, pp. 54-64.

[8] G. Daddis and H. Torng, "A Taxonomy of Broadband Integrated Switching Architectures", *IEEE Communications Magazine*, May 1989, pp. 32-42.

[9] J. Davin and A. Heybey, "A Simulation Study of Fair Queueing and Policy Enforcement", *Proceedings of ACM SIGCOMM '90*, Philadelphia, Pennsylvania, Sept. 1990.

[10] M. Decina, L. Faglia and T. Toniatti, "Bandwidth Allocation and Selective Discarding for Variable Bit Rate Video and Bursty Data Calls in ATM", *Proceedings of INFOCOMM 1991*, Bal Harbour, Florida, April 1991, pp. 12A.2.1-12A.2.8.

[11] L. Dittmann, S. Jacobsen and K. Moth, "Flow Enforcement Algorithms for ATM Networks", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No.3, April 1991, pp. 343-350.

[12] A. Eckberg, D. Luan and D. Lucantoni, "Bandwidth Management: A Congestion Control Strategy for Broadband Packet Networks – Characterizing the Throughput-Burstiness Filter", *Computer Networks and ISDN Systems*, Vol. 20, 1990, pp. 415-423.

[13] A. Eckberg, B. Doshe and R. Zoccolillo, "Controlling Congestion in B-ISDN/ATM: Issues and Strategies", *IEEE Communications Magazine*, Sept. 1991, pp. 64-70.

[14] T. Eliazov, V. Ramaswami, W. Willinger and G. Latouche, "Performance of an ATM Switch: Simulation Study", *Proceedings of INFOCOMM 1990*, June 1990, pp. 644-653.

[15] J. Filipiak, "Structure of Traffic Flow in Multiservice Networks", *Proceedings of INFOCOMM 1988*, pp. 13.4.1-13.4.5.

[16] G. Finn, "A Connectionless Congestion Control Algorithm", *ACM Computer Communication Review*, Vol. 19, No. 5, Oct. 1989, pp. 12-31.

[17] V. Friesen, *A Performance Study of Broadband Networks*, M.Math. Thesis, Dept. of Computer Science, University of Waterloo, 1991.

[18] G. Gallassi, G. Rigolio and L. Verri, "Resource Management and Dimensioning in ATM Networks", *IEEE Network*, Vol. 4, No. 3, May 1990, pp. 8-17.

[19] H. Gilbert, O. Aboul-Magd and V. Phung, "Developing a Cohesive Traffic Management Strategy for ATM Networks", *IEEE Communications Magazine*, Oct. 1991, pp. 36-45.

[20] S. Golestani, "Congestion-Free Transmission of Real-time Traffic in Packet Networks", *Proceedings of INFOCOMM 1990*, June 1990, pp. 527-536.

[21] I. Habib and T. Saadawi, "Controlling Flow and Avoiding Congestion in Broadband Networks", *IEEE Communications Magazine*, Oct. 1991, pp. 46-53.

[22] I. Habib and T. Saadawi, "Multimedia Traffic Characteristics in Broadband Networks", *IEEE Communications Magazine*, July, 1992, pp. 48-54.

[23] H. Heffes and D. Lucantoni, "A Markov Modulated Characterization of Packetized Voice and Data Traffic and Related Statistical Multiplexer Performance", *IEEE Journal on Selected Areas in Communication*, Vol. SAC-4, No. 6, Sept. 1986, pp. 856-868.

[24] V. Jacobson, "Congestion Avoidance and Control", *Proceedings of ACM SIGCOMM '88*, Stanford, California, Aug. 1988, pp. 314-329.

[25] R. Jain and K. Ramakrishnan, *Congestion Avoidance in Computer Networks with a Connectionless Network Layer Part IV: A Selective Binary Feedback Scheme for General Topologies*, Technical Report DEC-TR-510, Digital Equipment Corporation, Aug. 1987.

[26] R. Jain, "Congestion Control in Computer Networks: Issues and Trends", *IEEE Network Magazine*, Vol. 4, No. 3, May 1990.

[27] R. Jain, *Myths About Congestion Management in High-Speed Networks*, Technical Report DEC-TR-726, Digital Equipment Corporation, Oct. 1990.

[28] A. Kuczura, "Queues With Mixed Renewal and Poisson Inputs", *Bell System Journal*, Vol. 51, No. 6, July 1972, pp. 1305-1326.

[29] W. Leland, "Window-Based Congestion Management in Broadband ATM Networks: The Performance of Three Access Control Policies", *Proceedings of INFOCOMM 1989*, Ottawa, Canada, pp. 49.7.1-49.7.7.

[30] A. Mankin, "Random Drop Congestion Control", *Proceedings of ACM SIGCOMM '90*, Philadelphia, Pennsylvania, Sept. 1990, pp. 1-7.

[31] S. Minzer, "Broadband ISDN and Asynchronous Transfer Mode (ATM)", *IEEE Communications Magazine*, May 1989, pp. 17-24.

[32] M. Murata, Y. Oie, T. Suda and H. Miyahara, "Analysis of a Discrete-Time Single-Server Queue with Burst Inputs for Traffic Control in ATM Networks", *IEEE Journal on Selected Areas in Communication*, Vol. SAC-8, No. 3, April 1990, pp. 447-458.

[33] K. Ramakrishnan and R. Jain, "A Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer", *Proceedings of ACM SIGCOMM '88*, Stanford, California, Aug. 1988, pp. 303-313.

[34] G. Ramamurthy and R.S. Dighe, "Distributed Source Control: A Network Access Control for Integrated Broadband Packet Networks", *Proceedings of INFOCOMM 1990*, June 1990, pp. 896-907.

[35] E. Rathgeb, "Modeling and Performance Comparison of Policing Mechanisms for ATM Networks", *IEEE Journal on Selected Areas in Communication*, Vol. SAC-9, No. 3, April 1991, pp. 325-332.

[36] J. Roberts, "Variable-Bit-Rate Traffic Control in B-ISDN", *IEEE Communications Magazine*, Sept. 1991, pp. 50-56.

[37] H. Saito, M. Kawaraski and H. Yamada, "An Analysis of Statistical Multiplexing in an ATM Transport Network", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 3, April 1991, pp. 359-367.

[38] M. Sidi, W. Liu, I. Cidon and I. Gopal, "Congestion Control Through Input Rate Regulation", *Proceedings of INFOCOMM 1989*, Ottawa, Canada, pp. 49.2.1-49.2.5.

[39] A. S. Tanenbaum, *Computer Networks*, Second Edition. Prentice-Hall, 1989.

[40] J. Turner, "New Direction in Communications (or Which Way to the Information Age?)", *IEEE Communications Magazine*, Vol. 24, No. 10, 1986, pp. 8-15.

[41] R. E. Walpole and R. H. Myers, *Probability and Statistics for Engineers and Scientists*, Third Edition, Macmillan, New York, 1985.

[42] R. Wilder, K. Ramakrishnan and A. Mankin, "Dynamics of Congestion Control and Avoidance of Two-Way Traffic in an OSI Testbed", *ACM Computer Communication Review*, Vol. 21, No. 2, April 1991, pp. 43-49.

[43] G. Woodruff, R. Rogers and P. Richards, "A Congestion Control Framework for High-speed Integrated Packetized Transport", *Proceedings of INFOCOMM 1988*, pp. 7.1.1-7.1.5.

[44] G. Woodruff and R. Kositipaiboon, " Multimedia Traffic Management Principles for Guaranteed ATM Network Performance", *IEEE Journal on Selected Areas in Communication*, Vol. 8, No. 3, April 1990, pp. 437-446.

[45] H. Yamashi and S. Sumita, "A Traffic Measurement Method and its Application for Cell Loss Probability Estimation in ATM Networks", *IEEE Journal on Selected Areas in Communication*, Vol. SAC-9, No. 3, April 1991, pp. 315-323.

[46] L. Zhang, *A New Architecture for Packet Switching Network Protocols*, Ph.D Thesis, Dept. of Electrical Engineering and Computer Science, MIT, Technical Report LCS-TR-455, Aug. 1989.

[47] L. Zhang, "Virtual Clock: A New Traffic Control Algorithm for Packet Switching Networks", *Proceedings of ACM SIGCOMM '90*, Philadelphia, Pennsylvania, Sept. 1990, pp. 19-29.

[48] H. Zhang and S. Keshav, "Comparison of Rate-Based Service Disciplines", *Proceedings of ACM SIGCOMM '91*, Zürich, Switzerland, Sept. 1991, pp. 113-121.

[49] L. Zhang, S. Shenker and D. Clark, "Observation on the Dynamics of a Congestion Control Algorithm: The Effects of Two-Way Traffic", *Proceedings of ACM SIGCOMM '91*, Zürich, Switzerland, Sept. 1991, pp. 133-147.