Optimizing
Integrated Broadband Network Bandwidth Utilization
Through Connection Admission Control
Using Weighted Round Robin Queue Server Measurements

by

Stephen A. Oliva
B.S. University of Rhode Island, 1967
M.S. Rutgers University, 1974

Submitted to the Department of Electrical Engineering and Computer Science and the Faculty of the Graduate School of the University of Kansas in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

<Victor S. Frost>________________________

Chairman

<Joseph R. Evans>______________________

<James D. Church>_____________________

<Swapan Chakrabati>___________________

<David W. Petr>_________________________

Committee Members

March 12, 1999______________________

Date dissertation defended

Abstract

This research addresses improvements to network resource utilization through a specific means of measurement-based control of traffic admission in broadband integrated services digital networks. A method and system for Connection Admission Control (CAC) in a communications network, such as an ATM formatted B-ISDN network, is defined, analyzed, and tested. Using end-to-end virtual path (VP) structures and class-of-service separation, various network connections, or virtual channels (VC), are administered using a weighted round robin connection server. Network users aware of this network structure and the means by which queue lengths are determined may easily calculate a Sustainable Cell Rate (SCR) for the traffic they wish to introduce into the network path for transmission to a given destination. The user declared SCR, in addition to other user declared traffic parameters, determines the queue lengths allocated in the network switches, such that a required level of Quality of Service (QoS) is maintained. A measurement of certain types of transmitted cells in a VP, such as idle, unassigned, or low-priority cells, is made at the source of the VP. This measurement of cells that may be considered to be "empty", i.e., available for use by high-priority cells, is used as the basis for a determination of additional allowable connections that can be admitted to a VP. Relationships and bounds are determined through analysis of the mean number of "empty" cells per cycle of the connection server and the cell rate bandwidth that may be allowed to enter a VP that is already "full", based on the sum of the SCR of its existing component connections. These relationships are used as the basis of CAC, allowing connections to be admitted to a "full" virtual path if the cell rate of the requested connection is less than the allowable cell rate statistically determined from the mean number of "empty" cell timeslots in the path. Thus, a network VP can support bandwidth greater than the sum of the SCR of its included connections while maintaining required QoS.

Acknowledgements

This dissertation could not have been completed without the support and inspiration of many people. I would like to acknowledge the contributions of, and thank, the following:

My wife, Marianne, for her encouragement and support of my efforts throughout my studies and especially through the actual creation of this document. I can never thank her enough for her patience and understanding on those days during this undertaking where my approach to life in general was less than optimum.

Dr. Victor Frost, for his guidance, support and encouragement, both as my dissertation advisor, and as my academic advisor.

My dissertation committee members, Dr. Swapan Chakrabarti, Dr. James Church, Dr. Joseph Evans, and Dr. David Petr, for their overall encouragement, valuable suggestions, and careful reading of the document.

Dennis Schnack, for his encouragement and support of my participation in the Sprint Corporation Advanced Technical Degree Sabbatical Program, as well as Karl Kramer, Marty Kaplan, and George Fuciu, whose confidence in my abilities and approval of my participation in the program made this work possible.

Al White, for his steadfast encouragement of my research efforts, and understanding when things went a little slower than originally anticipated.

Chuck Weber, Steve Drumm, Scott Neal, Antwan Piggott, Ed Wortmann, and other members of the Sprint Technology Integration and Operations Center team, as well as Roel Jonkman, of the University of Kansas Information and Telecommunication Technology Center staff, whose expertise in keeping computer networks working properly insured the timely and successful completion of simulations described in this document.

Sprint Corporation, through the Advanced Technical Degree Sabbatical Program, provided support for this research.

I dedicate this work to my mother, Elizabeth Oliva, and to the memory of my father, Michael J. Oliva, to whom I owe everything I have accomplished in life.

Table of Contents

Chapter 3

Chapter 4

Chapter 5

List of Figures

List of Tables

Chapter 1

Introduction

CHAPTER CONTENTS

LIST OF FIGURES

List of Tables

Chapter

1

Introduction

This research addresses improvements to network resource utilization through a specific means of measurement-based control of traffic admission in broadband integrated services digital networks. To date, these networks have principally been implemented using asynchronous transfer mode network technology.

Broadband integrated services digital networks (B-ISDN) using asynchronous transfer mode (ATM) have been primarily designed to insure that a specified quality of service is provided to users within the parameters set for the various classes of traffic carried by the network. Of secondary importance during the initial deployment of these relatively new networks has been maximizing the overall efficiency of the allocation of network resources. The need today is both for quality, in terms of minimizing congestion and guaranteeing quality of service (QoS), and for maximizing the efficient use of the network by making effective use of available bandwidth and buffer resources. Accomplishing both simultaneously is often a problem.

Connection admission control (CAC), being the first level of congestion control for ATM networks, is critical in insuring required quality of service in a network, for all types of service, including voice, picture, data, and video. An effective CAC scheme is also necessary as the first step in insuring effective utilization of network resources by allowing the maximum use of available network link bandwidth, as well as queuing and computational resources in network nodes. This project proposes a CAC scheme based on the assignment of virtual path connections between each source node and destination node in a network. It makes use of class of service separation and measurements made at *weighted round robin* queue servers in the first network node traversed by the user traffic (i.e. the source node) to determine the amount of available bandwidth on a path. This approach provides a single measurement point per virtual path that may be used to facilitate admission of connections in such a way that optimization of bandwidth, buffer and computation resources occurs.

This research is based on the assumptions that (1) ATM-based B-ISDN networks will become a significantly large proportion of worldwide network capacity, (2) Given the current levels of increasing service requirements, effective use of bandwidth and switching capacity will be critical to providing adequate levels of service in the future, (3) Maintaining a guaranteed quality of service for defined traffic classes will continue to be a primary design objective in such networks, and (4) It is desirable to base a measurement-based scheme on taking measurements at as few points as possible.

1.1 Problem Statement

The initial design and deployment of today's new telecommunications capabilities, in the form of B-ISDN was primarily focused on insuring high quality of service (QoS) levels for all types of traffic. Although also important, the initial development of methods for insuring high levels of network resource utilization appeared to take a secondary role.

The efficient allocation of network resources through various management frameworks has now become the focus of much research under current B-ISDN architectures. In particular, since ATM was chosen as the technology for implementing B-ISDN, the means of initiating both bandwidth and buffer allocation in a network has become the role of CAC. Today's ATM based systems provide for virtual connections (VC) and virtual paths (VP) consisting of data placed in 53 byte packets, called cells. In these networks, CAC consists of actions taken by the network at the time establishment of a VC or VP is requested in order to determine whether a connection can be progressed or should be rejected (See [TMS4] for details, and [CHE94] for an overview). Although the functions and specifications of parameters useful for implementing CAC have been standardized as part of the overall ATM architecture, the specific CAC schemes in today's networks are not defined by specifications, but are dependent on a given network or given equipment design.

The CAC function in a given network may use traffic parameters, such as peak and sustainable cell rates, as traffic descriptors. Descriptors are evaluated and assumptions are made about the independence and distribution of the cell arrival process, with the distribution being evaluated to determine a usable connection admission parameter, for example: the equivalent virtual bandwidth. If evaluation of this quantity exceeds a threshold, such as available bandwidth remaining on a link or path, the request to establish a virtual channel is rejected. The request may alternatively be rejected if admitting it will cause QoS parameters to be exceeded. If the parameters of the requested connection are within the required quantity, the connection is accepted and network resources such as bandwidth and buffer space are allocated. However, the process may not be optimal in that the evaluation of the source traffic descriptor may or may not accurately specify the cell arrival process, since arrival processes having the characteristics of the source traffic descriptor are not necessarily unique.

In many ATM based network bandwidth allocation schemes proposed to date, CAC systems using a static resource reservation or allocation scheme have been defined. The bandwidth allocation process uses a source traffic descriptor to determine a reference model, which is then used by the CAC. Once allocated, the resources put in place to support the connection are left undisturbed for the life of the connection.

It is also possible to have a dynamic (or adaptive) allocation strategy. This type of CAC may monitor the number of cells arriving at a switch, the length of the switches cell buffer, and/or the number of cells lost, or measure some other parameter, to make decisions on whether to admit the connection and allocate bandwidth to it. For example, dynamic CAC strategies have been proposed that observe cell streams by counting the number of cells arriving at the output buffer of a virtual connection during a certain time interval, deriving a probability distribution of the number of cells arriving during each such period. An overall probability distribution is calculated and used to estimate a Cell Loss Ratio (CLR) assuming that the VC is admitted. If the CLR is acceptable (better than the requested CLR of the new VC), the VC is admitted and the estimate of the distribution is modified and the process repeated for additional VC's.

Additional CAC and bandwidth allocation strategies have been proposed, based on particular architectures, that throttle the traffic entering a VP at the "edge" of a network, so that loss and delay in the "core" of the network will be at a minimum, allowing the CAC decision to be made at the edge point, without requiring information from other nodes. Overall network architecture frameworks of this nature both guarantee quality of service and simplify CAC for ATM-based B-ISDN networks.

These methods, and many others, rely on estimated bandwidth calculations or estimated cell loss ratios, or specific measurements followed by calculations based on probabilities, to admit VC’s to a given switch or given path through many switches. Specific methods of resource allocation have been proposed by many researchers, using a large variety of methods, for various types of dynamic resource allocation and CAC.

Possible improvements in the efficiency of network resource allocation or in reducing the complexity of implementation methods compared to some or all of the above schemes can result from a procedure that provides a dynamic measurement of actual bandwidth remaining in a given end-to-end VP.

With this in mind, this research will focus on using ATM virtual channel and virtual path concepts as integral components in formulating a measurement-based scheme for connection admission and end-to-end bandwidth management. Analysis and computer simulation will be used in developing and evaluating a specific method of measurement that facilitates effective use of bandwidth in ATM networks.

1.2 Background

The purpose of this section is to describe the basic functions of ATM as used in B-ISDN today, and present possible solutions for some of the control problems faced by these networks. These relatively new networks are currently in the process of being implemented, and many aspects of their features, characteristics, and control are still being defined. The need for CAC to avoid congestion and improve the efficiency of allocating network resources will be discussed, and existing methods of

CAC will be described. The advantages of being able to dynamically determine allocated bandwidth of a connection between network nodes will also be discussed. This will be done in terms of determining how much bandwidth remains in a given path for the acceptance of additional service, which comprises the basic function of CAC. Alternatives in network implementation methods that, in conjunction with CAC, effect the overall efficiency of using network resources will be reviewed.

1.2.1 Asynchronous Transfer Mode

Today's B-ISDN systems, which are based on an ATM architecture, allow the combination of various types of services on one link, and also allow statistical multiplexing of services in which communications occur at a variable bit rate. ATM networks provide certain guarantees of the QoS allocated to various types of traffic served by the network. In short, ATM networks can serve voice, picture, data, and video traffic over the same physical and logical links in the network more effectively than previously deployed communications networks.

ATM networks use a cell (i.e. fixed length packet) based system having a five-byte header and a 48-byte payload in each cell. Network efficiency in ATM networks can be improved over previous networks through a statistical multiplexing process. This is accomplished by allowing traffic sources having a variable bit rate onto a link where the combination of their peak rates may exceed the maximum link capacity. This is possible since it may be done in a manner that makes it statistically unlikely that sufficient sources would transmit at the same time so as to cause congestion and loss of cells.

ATM networks also gain efficiency when compared to previous networks by their ability to combine various types of services into one link, rather than having to carry several "overlay" networks, each having their own inefficiencies. Procedures and parameters are defined in ATM networks that allow the ability to provide required QoS for the various network applications or types of service.

1.2.1.1 ATM Service Categories and Quality of Service.

Services provided in ATM based networks, as defined by the ATM Forum[1], presently consist of the following five service categories as defined in [TMS4]:

- Constant Bit Rate (CBR)
- Real-time Variable Bit Rate (rt-VBR)
- Non-real-time Variable Bit Rate (nrt-VBR)
- Unspecified Bit Rate (UBR)
- Available Bit Rate (ABR)

The above service categories relate traffic characteristics and QoS requirements to network behavior. Their characteristics are summarized here.

---

[1] The ATM Forum was formed in 1991 to accelerate the deployment of ATM products and services through rapid convergence of interoperability specifications and promotion of industry cooperation. The forum consists of communications industry, communications users, government, and university/research members.

CBR: Used by connections requiring a fixed amount of bandwidth, as defined by a peak cell rate (PCR). Once a connection is established, the required QoS is assured to all cells when they do not exceed the PCR. Supports constant rate real time services (ex: voice, video, dedicated circuit emulation) which are constrained by cell transfer delay.

rt-VBR: Intended for applications that are tightly constrained in regards to cell delay and cell delay variation (CDV), for example, voice and video applications which vary with time. These are described as "bursty" and characterized in terms of PCR, Sustainable Cell Rate (SCR), and Maximum Burst Size (MBS).

nrt-VBR: Intended for applications that do not require real-time transmission (i.e. delay and delay variation are acceptable), and have "bursty" characteristics which are characterized in terms of PCR, SCR, and MBS. No delay bounds are applied, but low cell loss ratio (CLR) is expected.

UBR: Intended for non-real-time applications not requiring tightly constrained delay and delay variation (ex: computer file transfer and e-mail). No numerical commitments are made with respect to CLR or delay.

> ABR: Intended for use where a specific cell rate is not required, but the cell stream can adapt to what is available using feedback from the network. Has no bound on delay or delay variation, but ABR service is expected to have low CLR if it conforms to the flow control feedback messages from the network. A minimum cell rate (MCR), which may be zero, is defined.

ATM networks establish virtual paths (VP) and virtual channels (VC) within the VPs by placing appropriate channel and path identification bytes within the cell headers. Virtual paths and channels may be either permanent (i.e. "provisioned" in a traditional network sense) or switched. The service categories above are defined by using the following QoS parameters, which are defined fully in [TMS4]:

- Peak-to-Peak Cell Delay Variation
- Maximum Cell Transfer Delay (maxCTD)
- Cell Loss Ratio (CLR)

Cells may be marked as high or low priority using a Cell Loss Priority (CLP) indicator, which, through appropriate network management controls, helps insure appropriate QoS to the various service classes.

1.2.1.2 ATM Traffic Descriptors

The concept of a user-network *traffic contract* has been introduced by the ATM Forum to insure that QoS requirements of the various classes of service are met. This is done by having the parameters under which the connection will be accepted into the network defined for use in establishing the resources needed to serve the connection and also for use in "policing" the connection at the point of acceptance. Policing insures that cell streams that violate the traffic contract are accommodated in such a way that the QoS of traffic streams that are conforming to their traffic contract is not degraded. The *connection traffic descriptor* consists of a *source traffic descriptor*, the cell delay variation tolerance (CDVT), and a conformance definition. The source traffic descriptor parameters are PCR, SCR, MBS, and MCR, as described in Section 1.2.1.1. The CVDT is specified at the user-network interface (UNI) and is defined in relation to the PCR provided by the user. The burst tolerance (BT) is also specified at the UNI and is defined in relation to the SCR provided by the user. A conformance definition based on a generic cell rate algorithm (GCRA), as defined by the ATM Forum, is used to determine whether cells in a given connection conform to the traffic contract. The descriptor parameters are either provisioned (for permanent connections) or are signaled (for switched connections) through a UNI when a connection is established.

Thus, traffic descriptors, provided through or derived from the traffic contract, can be used by the network in determining whether or not to accept a given connection while maintaining the QoS parameters of the existing traffic.

Traffic parameters and associated QoS parameters, as well as requirements for feedback of information from the network, for each ATM service category, are shown in Table 1-1.

Attribute	**ATM Service Category**				
	CBR	**rt-VBR**	**nrt-VBR**	**UBR**	**ABR**
Traffic Parameters					
PCR and CDVT [4,5]	specified			specified [2]	specified [3]
SCR,MBS,CDVT [4,5]	N/A	specified		N/A	
MCR [4]	N/A				specified
QoS Parameters					
peak-to-peak CDV	specified		unspecified		
max CTD	specified		unspecified		
CLR [4]	specified			unspecified	see Note [1]
Other Attributes					
Feedback	unspecified				specified

Note 1: CLR is low for sources that adjust cell flow in response to control information. Whether a quantitative value is specified is network specific.

Note 2: May not be subject to admission control and policing procedures

Note 3: Actual cell rate may not be peak cell rate, subject to feedback control information.

Note 4: May be either explicitly or implicitly specified for PVCs or SVCs.

Note 5: CDVT is not signaled and does not have a unique value for a given connection. There may be different values of CVDT at each interface along a path.

Table 1-1: ATM Source Category Attributes

1.2.2 Congestion Control in ATM Networks

Networks consist of physical devices and connections, each having their own characteristics, such as bandwidth capability and buffer size. Should the required connection characteristics at any point in a network approach or exceed the physical capabilities, congestion results. ATM networks are no exception to this, and congestion control in ATM networks has been extensively addressed. Mechanisms for accomplishing congestion control are, to a great extent, built into the infrastructure of these networks.

The primary role of an ATM network congestion control procedure is to preserve a desired Quality of Service (QoS) for the users of the network, in order to achieve the required performance objectives for the various classes of service.

Congestion control procedures may be either *preventive* or *reactive* in nature. Preventive congestion control requires strategies that prevent or reduce the occurrence of congestion. In reactive congestion control, information that is fed back through the network controls various parameters or devices, thus controlling the amount of congestion. In ATM networks, effective congestion control requires the use of both techniques. Effective control of CBR and VBR traffic in an ATM network makes use of preventive schemes, since no method is presently defined for CBR and VBR traffic to react to network conditions through feedback. Control of ABR and UBR traffic makes use of reactive schemes, such as rate control through feedback in the case of ABR traffic, and low priority cell discarding when congestion conditions occur in the case of UBR traffic.

Preventive congestion control may make use of the techniques of CAC, source traffic bandwidth enforcement, or reallocation of bandwidth among network paths. Effective network link utilization, and therefore effective overall resource utilization, is dependent on the efficiency of a CAC scheme. Since the ATM Forum already defines effective procedures for enforcement of source traffic bandwidth, and re-allocation of bandwidth among network paths would be based on evaluation of the same parameters and techniques used by an effective CAC, this project will focus only on CAC.

1.2.3 Connection Admission Control (CAC)

Prior to a network user transmitting over an ATM network, a connection between sending and receiving locations must be created, using a connection (or call) setup procedure. The path established between the sending and receiving locations will involve at least one, and possibly more, ATM switches. Some of the available resources of these switches will be assigned to each new connection as it is accepted. Should sufficient resources not be available, the request for connection will be rejected. The functions of CAC have been defined [TMS4], and are recognized as being network specific in their implementation. It is important to note that the decision by the CAC function as to whether to accept or reject the new connection is based on whether or not the network can provide the QoS required by the new connection without affecting the QoS of the connections already being carried in the network.

1.2.3.1 Types of Connection Admission Control

Many methods of implementing CAC are possible. Within these methods, the CAC decision may be accomplished on a node-by-node basis, as each switch determines if its currently available resources can support the requested new connection, and, if accepted, passes the request to the next node. Alternatively, if the network has sufficient knowledge of the state of the resources along the required path CAC may be accomplished on an end-to-end-basis.

CAC schemes have been classified [PER96] as *non-statistical* (or peak bandwidth) allocation and *statistical* allocation. Non-statistical allocation assigns a bandwidth equal to the peak bandwidth of the connection no matter what the statistical characteristics of the signal may be. In statistical allocation, bandwidth for a connection is assigned at less than the peak bandwidth of the connection, depending on the statistical distribution of the arriving cells in the connection. This distribution may be difficult to predict accurately. Statistical allocation results in statistical multiplexing gain, when dealing with sources that arrive in “bursts”, since the sum of all peak rates may be greater than the capacity of the output link. Statistical allocation is difficult to accomplish effectively because of the uncertainty in distribution of arriving cells.

CAC schemes have also been classified [SAI94] as *non-parametric* and *parametric*. The non-parametric approach does not assume a cell arrival process model, and is therefore the equivalent of the non-statistical allocation process described above. Since the cell arrival process is not modeled, performance measures such as Cell Loss Ratio (CLR) cannot be specified, but a boundary on the measure, such as maximum CLR, can be determined. Parametric allocation approaches are similar to the statistical allocation approaches described above.

In addition, CAC proposals may be classified [ZUK97] in three categories: (1*)* *Peak Rate Allocation* which considers neither sharing of bandwidth nor of buffer capacity; (2) *Rate Envelope Multiplexing* which is based on a zero buffer approximation and assumes sharing of bandwidth but not of buffering capacity; and (3) *Rate Sharing*, which considers sharing of both bandwidth and buffering.

Any of the above categories may also be implemented as *static* or *dynamic* CAC schemes. In static schemes, resources which are allocated to a connection upon its creation in the network remain in use unchanged for the duration of the connection. Dynamic schemes may change resource allocations from time to time during the connection time, based on the actual connection characteristics measured while it is in use, which in turn affects the resources available for additional connections through the network. Dynamic CAC schemes are also referred to as *adaptive.*

See Figure 1-1 for relationships between these various classification schemes.

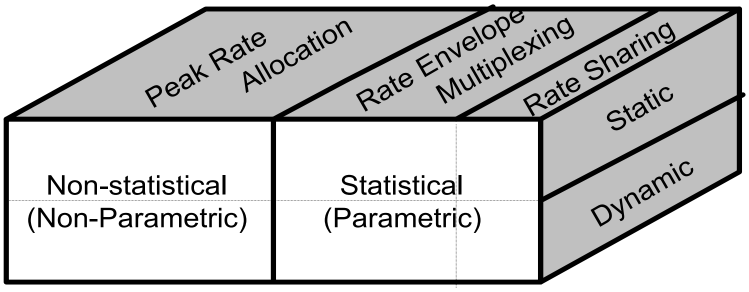

Figure 1-1: Classification of CAC Schemes

In ATM networks, the resources required for specific connections are not always clear, since the statistical distribution of arriving cells in the connection is not always known. Also, the solution to statistical allocation of bandwidth may require the application of large quantities of computer processing power, which may not be possible in the short times involved. A wide variety of call admission schemes have therefore been proposed, and several will be summarized in the following sections.

1.2.3.2 Specific Approaches for CAC

1.2.3.2.1 Peak Rate Allocation

Peak rate allocation is a method of CAC that is straightforward and simple to implement. In general, for a given link bandwidth, each admitted connection is allocated a bandwidth corresponding to its peak rate, either specified by the customer in the PCR when the contract is negotiated, or, in some dynamic allocation schemes, measured by the network when the connection has been in place for a sufficient time. On admission, the requested peak rate is compared to the remaining bandwidth in the link after all existing peak bandwidth have been subtracted, and if the requested peak bandwidth is larger, the connection is rejected.

This approach is effective for CBR traffic where peak rate is close to or matches actual rate of traffic transmission, however, poor efficiency due to low utilization of assigned bandwidth is the result of implementing this method for VBR traffic, since no statistical multiplexing gain can be realized. This method is therefore very conservative when applied to VBR or mixed VBR and CBR paths.

1.2.3.2.2 Equivalent Capacity

The equivalent capacity of a two-state source (or group of sources) that has Markov characteristics, which is feeding a queue of finite capacity, is the service rate of the queue that provides a specified cell loss probability (i.e. buffer overflow probability) of ε. Fluid flow approximations and other simplifications have been used to develop equations for equivalent capacity of a single source and also multiple sources [GUE91]. For a single source the equivalent capacity $\hat{c}$ is determined from:

$$\hat{c} = \frac{\alpha B(1-\rho)R_{peak} - L + \sqrt{[\alpha B(1-\rho)R_{peak} - L]^2 + 4L\alpha B\rho(1-\rho)R_{peak}}}{2\alpha B(1-\rho)} \tag{1.1}$$

where

$\alpha = \ln\left(\frac{1}{\varepsilon}\right)$

$\varepsilon =$ The desired cell loss probability for which $\hat{C}$ is being calculated.

$B =$ Mean burst period of the Markov source.

$\rho =$ Utilization (offered load) of the Markov source flowing into the queue.

$R_{peak} =$ Peak rate of the Markov source.

$L =$ Queue length.

The complete approach is to choose equivalent capacity $\hat{C}$ as the minimum of two values, one derived under a fluid flow model for multiplexed sources, and the other derived when the effect of statistical multiplexing is the dominant factor, that is:

$$\hat{C} = \min\left\{ m + \sigma\sqrt{-2\ln(\varepsilon) - \ln(2\pi)}, \sum_{i=1}^{N} \hat{c}_i \right\} \tag{1.2}$$

where

$N =$ The number of multiplexed connections.

$m =$ The mean aggregate bit rate of the combined sources.

$\sigma =$ The standard deviation of the aggregate bit rate.

Others, such as Zhang and Acampora [ZHA94] and de Veciana, Kesidis, and Walrand [VEC95] have also studied equivalent bandwidth methods, presenting various approximations for computing equivalent bandwidths for different types of source traffic systems sharing various types of buffers. Mang and Gelenbe [MAN96], and Gelenbe, Mang and Onvural [GEL97] proposed an improved CAC and bandwidth allocation procedure based on statistical bandwidth, taking into consideration the factors of cell loss requirement, traffic characteristics, and buffer size at statistical multiplexers, computed using diffusion approximations, which are continuous approximations of the arrival and service processes in discrete queuing models. Perros and Elsayed [PER96] have summarized the contributions of additional researchers to the equivalent capacity method, and similar methods, such as the Heavy Traffic Approximation method. These methods are based on the asymptotic behavior of the queue length distribution, which is also the basis for the equivalent capacity method. They also summarize studies that indicate the equivalent capacity method may be inaccurate in some situations. These inaccuracies have also been demonstrated by Braun, Petr, and Kelley [BRA94] showing that in general, the equivalent capacity method proposed by Guerin is more generous in estimating the capacity required for a source as the burstiness of the source increases. Choudhury, Lucatoni, and Whitt [CHO96] have also shown that the effective bandwidth method is generally conservative in estimating bandwidth if the sources are Poisson sources or more bursty than Poisson sources. However, if the sources are less bursty than Poisson sources, then the effective bandwidth approximations are no longer conservative and the method may then be a poor estimator of bandwidth requirements. They propose an approximation that is almost as simple to implement as the original effective bandwidth approach, but is applicable in a wider region. The equivalent capacity expression, Equation (1.1), has therefore been shown to be an upper bound on the actual bandwidth requirements of most traffic sources, and is generally a conservative approach to allocating bandwidth, but under some circumstances may be inappropriate.

1.2.3.2.3 Upper Bound on Cell Loss Probability

Saito, in [SAI91] [SAI92], proposed an admission control scheme using an upper bound for the cell loss probability. This method is based on customer specified parameters such as the average number of cells arriving during a fixed interval and the maximum number of cells arriving during the same fixed interval. The fixed interval is taken to be half of the maximum allowable delay in a buffer. In this method, T (in milliseconds) is the maximum allowable delay in a buffer; the buffer size K is dimensioned such that the maximum delay in a buffer is less than T under a FIFO queue discipline, i.e., $K = 1000TC / L$, where C is the transmission link speed in bits/sec and L is the cell length in bits. Making buffers this size always guarantees that cell delay requirements will be met, allowing the CAC procedure to be focused on meeting a cell loss requirement. If cells of N calls are sent by a transmission link, and the probability that j cells of the i^{th} call arrive during $T/2$ ms is $p_i(j)$, the Cell Loss Probability (CLP) is upper bounded by B which is a function of $p_i(j)$ and $K/2$ such that:

$$CLP \le B(p_1,\Lambda\ ,p_n(k);K/2) \overset{\Delta}{=} \frac{\sum_{k=0}^{\infty}\left[k-K/2\right]^{+} p_1 \otimes \Lambda \otimes p_N(k)}{\sum_{k=0}^{\infty} k p_1 \otimes \Lambda \otimes p_N(k)} \tag{1.3}$$

where $\otimes$ denotes convolution. Proof of the above is in [SAI92].

Letting $\theta_i(j)$ be the following function, where a_i is the average rate in given period and R_j is the maximum rate in given period, quantities which can be derived from customer supplied parameters of SCR and PCR in the call:

$$\theta_i(j) = \left\{ \begin{array}{ccc} a_i/R_i & if & j = R_i \\ 1 - a_i/R_i & if & j = 0 \\ 0 & & otherwise \end{array} \right\}$$

Then

$$CLP \le B(p_1, \Lambda\ , p_n(k); K/2) \le B(\theta_1, \Lambda\ , \theta_n(k); K/2) \overset{\Delta}{=} \frac{\sum_{k=0}^{\infty} \left[k - K/2\right]^{+} \theta_1 \otimes \Lambda \otimes \theta_N(k)}{\sum_{k=0}^{\infty} k \theta_1 \otimes \Lambda \otimes \theta_N(k)} \tag{1.4}$$

Using equation (3.3), a new connection is admitted if the resulting $B(\theta_1, \Lambda\ , \theta_n(k); K/2)$ is less than the admissible cell loss probability. Saito proposes a method of calculating $\theta_1 \otimes \Lambda \otimes \theta_N$ efficiently, based on peak and average cell rates, as well as a method of calculating the upper bound based on the average cell rate and variance of the cell rate.

1.2.3.2.4 Fast Resource Allocation

Other methods of controlling congestion through the control of cell admission, rather that connection admission, have also been included in discussions of CAC. Such methods include fast buffer/bandwidth allocation, where “bursty” sources are admitted by establishing a virtual circuit, but only allocating resources of bandwidth and/or buffers to it when requested at the time a burst is actually ready to be transmitted.

Turner [TUR92] proposed methods for fast buffer reservation where the reservation is made as the data is sent eliminating the need for explicit control messages. This was done by allocating two states to each virtual circuit passing through a given link buffer, where the two states are “idle” and “active”.

When a given circuit is active, it is allocated a pre-specified number of buffer slots in the link buffer until it becomes inactive through a transition to the idle state. User cells that are specifically marked as either start-of-burst or end-of-burst signal transition between idle and active state. Forced transitions to idle state are made if no cells are received after a given time period.

When determining whether a connection that uses this approach may be accepted, it must still be decided whether a new virtual circuit may be multiplexed with existing ones. Turner considers two methods for accomplishing this; called the excess demand probability and contention probability methods.

The analysis using these criteria does not depend on the connection’s cell burst lengths, and Turner shows that systems based on the above methods could be effectively implemented with relatively modest commitment of computational resources. CAC decisions are based on the network’s knowledge of the user’s peak and average data rate, making the effective implementation of this method a measurement-based scheme.

Additional methods using some form of fast resource allocation are outlined in [PER96].

1.2.3.2.5 Time Frames

Connection admission methods have also been proposed based on the concept that a transmitter will only be allowed to send up to a maximum number of bits (or packets or cells) within a fixed period of time.

Golestani [GOL91] proposed a method where the number of cells in each connection on any link is bounded by a time frame for each connection. The frame was proposed as not being adjustable and as being the same for every link in the network. This means that the total number of cells transmitted by all connections on a link is upper-bounded, thus maintaining a smooth traffic flow throughout the network. A cell arriving at an input in a given time frame will be transmitted from the switch at the end of the adjacent frame, thus putting a bound on cell delay. The proposed method requires buffering and has also been called stop-and-go queuing.

This strategy is best suited for periodic bit streams, moderately bursty services, and services with stringent packet loss and delay requirements. The approach was proposed to be shared with other approaches for services that are highly bursty, since it is very conservative in allocating resources. With its higher bandwidth utilization it allows fewer connections, but guarantees quality better than approaches more oriented to statistical multiplexing.

Additional schemes involving time frames (or time windows) are summarized in [PER96].

1.2.3.3 Dynamic/Measurement Based CAC Methods

CAC schemes maintain a record of allocated bandwidth on connections active in the network, and based on methods such as those discussed in the previous sections, allow new connections into the network, providing QoS standards can be met. A dynamic or adaptive CAC is capable of adjusting the record of resources in use, and, if necessary depending on the method, actually adjusting the bandwidth or buffer space allocated to the connection if it is not being fully utilized.

Resources freed through this procedure may then be used in the admission of new connections, or the re-negotiation of resources for connections requesting greater bandwidth or buffer space. Dynamic CAC requires that the resources allocated to a connection be adjusted, if needed, during successive fixed time periods. The method of re-allocation of the estimated or actual resources used must necessarily be based on the measurement of one or more network parameters, such as arriving cells, buffer occupancy, cells lost, point-to-point delay, etc. Several of these measurement-based approaches are summarized in the following sections.

1.2.3.3.1 Distribution of Arriving Cells

Saito has investigated dynamic CAC [SAI91] and dynamic resource allocation [SAI97] based on the distribution of the number of cells arriving during a fixed interval. The distribution is estimated from the measured number of cells arriving at an output buffer, and the traffic parameters specified by users. Call acceptance is decided on the basis of on-line evaluation of the upper bound of cell loss probability, a technique described in section 3.3.2. The process is independent of the classification of the calls and arrival process modeling. If the $p_i(j)$ from equations (3.1) can actually be observed, the upper bound of cell loss probability can be evaluated, allowing the call to be accepted if the result is less than the required cell loss probability.

1.2.3.3.2 Aggregate Cell Arrival Statistics

Zukerman and Tse [ZUK97] have proposed a method where the aggregate statistics of all calls in progress are measured and used in a real time estimation of CLR. New calls are considered for admission at their peak rate, and considered to be CBR at their peak rate during a “warming up” period. A new call is admitted if admitting it at its peak rate will not cause a threshold CLR to be exceeded. This method includes an adaptive mechanism which updates the warming up period based on traffic conditions, and a specific mechanism to predict cell loss. This scheme was simulated and compared to Peak Rate Allocation CAC and Equivalent Capacity CAC, showing increased utilization of at least 30% over the other schemes at large (1000 cell) buffer sizes. This CAC, requiring only a peak rate to be declared, using traffic sources based on variable bit rate video encoding, performs more efficiently than schemes that require a full set of declared traffic statistics.

1.2.3.3.3 Measured Buffer Occupancy

Courcoubetis, et. al. [COU95] have also proposed a method for call routing and acceptance that guarantees bounds on the fraction of cells lost by the connections because of buffer overflows. By monitoring its buffer occupancy, each switch constantly estimates its space capacity to accept new calls. To estimate its space capacity, the switch must evaluate the value that the cell loss probability would have if more calls were accepted using that switch. In developing an algorithm for this purpose, using Markov fluid sources, F(N,B,C) is the fraction of cells lost due to buffer overflows where N is the number of virtual circuits (>0), B is a given buffer size and C is the link capacity in cells/second. The fraction of cells lost when a fraction ε more connections are added is F($N(1+\varepsilon),B,C$). Using a general source traffic model, the authors show the probability of buffer overflow, Φ, has the following property for large B:

$$\Phi(N(1+\varepsilon),B,C) \approx \Phi(N,B,{}^{C}\!/_{1+\varepsilon})$$

By expressing the fraction of cells lost, F, in terms of the probability of cell loss Φ, the authors conclude that F has the above property as well. The problem is thus to estimate F($N,B,C/1+\varepsilon$) by traffic monitoring. Because the cell loss probability is designed to be very small, an estimator based on F has a large variance. To improve the estimator, a device can be added to a switch that tracks the buffer occupancy when the service rate is $C/1+\varepsilon$. Virtual buffers, which are smaller than the actual buffers, are used to increase the frequency of buffer overflows and speed up the collection of "important" samples, thus reducing the variance. The virtual buffer estimate $\mathrm{F}(N, B/k, C/1+\varepsilon)$ is related to the actual buffer $\mathrm{F}(N, B, C/1+\varepsilon)$ such that

$$\mathrm{F}(N,B,C/1+\varepsilon)=A(B)^{-\xi}e^{-BI(N,C/1+\varepsilon)}$$

then

$$\mathrm{F}\left(N,{B}/{k},{C}/{1+\varepsilon}\right)=A\left({B}/{k}\right)^{-\xi}e^{-\frac{B}{k}I\left(N,{C}/{1+\varepsilon}\right)}$$

where there are three unknowns; A, ξ, and I. If three virtual buffer sizes are selected by specifying three values of k: $k_0>k_1>k_2>1$, the resulting three equations can be solved for A, ξ, and $I(N, C/1+\varepsilon)$ which can then be used to determine the quantity

$$\mathrm{F}\left(N,B,{C}/{1+\varepsilon}\right)\approx \mathrm{F}\left(N(1+\varepsilon),B,C\right)$$

Thus, by monitoring the three virtual buffer overflows, estimates of $\mathrm{F}(N, B/k_i, C/1+\varepsilon)$ are constantly available to the call acceptance algorithm, which will base its decision on whether F increases unacceptably with the addition of ε more connections.

This method could be used to handle multiple types of calls sharing a buffer. However, many idealizations were made in its formulation, requiring further testing in real networks.

1.2.3.3.4 Large Deviation Rate-Function

Crosby, et. al. [CRO97] have also proposed a CAC method based on measurements of traffic activity. Using Large Deviation Theory [WEI95], the traffic measurements allow the estimation of a large deviation rate-function, or *entropy*, of bursty ATM traffic (See [DUF95] for further discussion of entropy). The entropy can be used to determine the bandwidth required by the traffic to achieve a given QoS, thus allowing CAC to be accomplished.

The bandwidth requirement of a connection is defined to be the minimum line-rate s at which a target cell-loss ratio c is not exceeded in a buffer of capacity b.

$$BWR(b,c) = \min\{s\colon CLR(b,s) \le c\}$$

An arrival process $\{A_t\}$ includes connections having bandwidth requirements, where A_t denotes the total cells that have arrived up to time t. Using large deviation theory, it can be shown that if the arrival process has a rate-function $I(a)$ on a time scale t such that:

$$\Pr\left(\frac{A_t}{t} > a\right) \propto e^{-tI(a)}$$

then the cell loss ratio in a buffer of size b with fixed line-rate s decays exponentially in b when b is large:

$$CLR(b,s) \propto e^{-b\delta}$$

and the *decay-rate* δ is determined by the rate-function I(a):

$$\delta = \min\left\{\frac{I(a+s)}{a} : a > 0\right\}$$

What is called the *entropy* of the system is the rate-function *I(a)*. If a two-state Markov model is used for modeling ATM traffic, the large deviation estimate can be improved by the use of a factor $e^{-\mu}$ leading to:

$$CLR(b,s) \leq e^{-\mu(s)-b\delta(s)}$$

which in turn leads to:

$$\log CLR(b,s) \leq -\mu(s) - b\delta(s)$$

for buffer sizes above a small buffer threshold size. If estimates of $\mu(s)$ and $\delta(s)$ can be obtained through measurement, the *estimated bandwidth requirement* is defined as:

$$B\hat{W}R(b,c) = \min\left\{s{:}-\hat{\mu}(s) - b\,\hat{\delta}(s) \leq \log c\right\}$$

The CAC algorithm for this method combines estimation of the bandwidth requirement of the existing load using on-line measurement to determine the parameters μ and δ, with a prediction of the bandwidth requirement of the offered traffic based on its declared parameters, after setting μ to zero. If the sum of these bandwidths exceeds the line rate, the offered traffic is rejected: otherwise it is accepted.

1.2.3.3.5 Virtual Buffers

Liu, Petr, and Braun [LIU97] have proposed CAC based on the measurement of the occupancy of a virtual buffer, and renegotiating UPC parameters, or "policing", based on the measurements. The method is based on: partitioning the ATM network into a VC-based edge network and a VP based core network; ensuring at least one CBR/VBR VP and one ABR/VBR VP between each edge-node-pair. The VP bandwidth used in this scheme is semi-permanently based on long-term bandwidth management needs and the use of weighted round robin (WRR) cell scheduling (see Section 1.2.4.2for a discussion of WRR). The proposed method achieves both bandwidth allocation and dynamic sharing of bandwidth, and also throttles the VC traffic entering a VP at the edge, so that loss and delay in the core will be minimal.

The virtual buffer described by this proposal simulates a FIFO queue with fixed serving rate equal to the allocated bandwidth, providing worst case delay and loss estimates. It also simulates a GCRA based cell policer, allowing UPC adjustments. The actual CAC strategy is based on estimation of actual usage of bandwidth, using a sliding window of pre-defined measurement window width, T_w, and for each VC letting T_s be the time elapsed since admission of the connection. Two sets are defined where S_1 is the set of VC's where $T_s > T_w$ and S_2 is the set of VC's where $T_s < T_w$; S_2

therefore contains VC’s for which measurements are not available, since insufficient time has elapsed since admission. The network constantly monitors the actual bandwidth usage of all VC’s belonging to S_1. The actual VP bandwidth usage is then estimated as the SCR value that satisfies the most stringent delay and CLR requirement of all the VC’s in the VP under consideration.

For those VC’s belonging to S_2, the bandwidth is always estimated as the claimed SCR. The bandwidth of the new connection request is also estimated as its SCR. The admission criterion is then: If the sum of the estimated bandwidth of both VC’s already admitted and the new request is greater than the VP bandwidth, reject the new call, otherwise the call can be accepted. In general, a larger T_w means a more conservative admission strategy, and a smaller T_w means a more aggressive strategy. The authors proposed finding a good rule through further experiments and analysis.

1.2.3.3.6 Measurement of End-to-End Delay

Frost and Mullen [FRO92] have proposed a methodology that uses end-to-end delay estimates to infer the available unused bandwidth along a VP connecting the source to the destination. Bursts of traffic exceeding the negotiated rate from the source are violation tagged, and allowed to enter the network if unused bandwidth is available. Negotiated bandwidth request and release procedures are based on estimates of the traffic flow. Many dynamic bandwidth allocation schemes depend on feedback information; this scheme is proposed such that it is robust relative to the overall feedback conditions in high speed networks. Capacity is released by monitoring the input traffic rate. When the input rate drops below the current negotiated peak rate by a certain percentage, bandwidth is released and a new negotiated peak rate is established. An estimate of the available free bandwidth is used in lieu of the explicit transmission of this information from each switch along the source-to-destination path. Probe cells, which are used to calculate the delay in a given path, are sent through the network and returned to the source. They are then used to estimate for the average end-to-end delay. The available free bandwidth is then estimated using:

$$\text{Available free BW} = \left(r_{\max} - r_e\right) * C$$

$$r_e = \frac{2(r-1)}{2r-1}$$

$$r = \frac{\left[D - T_p - \left(T_t * N_h\right)\right]}{T_t}$$

where

r_{max} = maximum utilization
r_e = estimated utilization
C = capacity of the link (bits/sec)
r = normalized estimated queuing delay
D = end-to-end delay
T_p = link propagation delay
T_t = transmission time
N_h = number of hops between source and destination access points

Simulation results indicated the good performance of the algorithm. Further work was recommended.

1.2.3.4 Summary of CAC

Of the many methods of CAC examined since ATM has become a significant method in B-ISDN networking, none has yet been judged fully and completely effective in handling all types of ATM traffic, or, more specifically, effective at handling VBR traffic in all of its possible manifestations. Peak rate allocation is adequate for systems having primarily CBR traffic, but is ineffective in providing high efficiency in VBR paths, since no statistical multiplexing is possible using peak allocation. The effective bandwidth method, perhaps because of its early development and relatively easy implementation, has gained a foothold in the management of VBR traffic in many actual and proposed systems.

Other methods of determining resource requirements in order to formulate a CAC scheme, such as boundary on cell loss, and dynamic bandwidth allocation methods, have been shown to be theoretically improved over peak allocation and early equivalent bandwidth methods. Some improvements have been proposed [DOU97] which are based on the techniques discussed in this section but use neural networks and fuzzy logic as a means of implementation. The implementation of many of these techniques imposes a cost of increased computational complexity, in comparison to relatively simple computations such as that required by the equivalent capacity approach.

In addition to all considerations discussed so far, the possible choices of a method of CAC are often determined by overall network architectures, specific designs of network switching elements, methods of overall network management control, and other considerations. Some alternatives for these choices will be examined next.

1.2.4 Some Network Implementation Alternatives Affecting CAC

ATM has evolved from a proposal for integrating a wide variety of services carried on Broadband Integrated Service Digital Networks to a reasonably well defined technology supporting four major service classes. ATM has proven itself to be effective in providing high-speed communications services in real commercial networks over the past several years. This structure of many services, where connections having traffic with very different statistical characteristics must satisfy specific QoS requirements, has led to the development and application of control processes targeted at insuring desired QoS levels while maximizing the utilization of network resources to some extent.

VBR traffic offers the greatest challenge in determining the proper bandwidth needed to insure required QoS levels, since this traffic is by nature difficult to pre-define in a statistical sense. The efficiency gains of statistically multiplexing these “bursty” VBR traffic connections have generated many schemes for determining bandwidths of given connections, as can be seen in the prior sections of this chapter. Methods based on user definitions of traffic have been implemented to control congestion by using both admission control and for “policing” connections to insure that parameters negotiated at traffic acceptance are not exceeded when the network becomes congested. Weighted round robin (WRR) and other fair queuing schemes have been proposed to schedule cells in ATM network switches to insure that connections receive sufficient bandwidth.

Many proposed methods of connection bandwidth allocation, although theoretically improved over early peak allocation or equivalent bandwidth methods, are reliant on estimates of the statistical characteristics of the traffic, or on models based on statistical characteristics. Such methods have been improved upon by methods less dependent on estimates or models, using actual measurements of available resources, such as bandwidth or buffer space. These methods can be effective, both in assuring QoS and in obtaining efficient link utilization, providing the collection and processing of needed data does not require a complex or extensive computation method.

The methods of bandwidth allocation schemes for implementing CAC may vary based on the types of traffic offered. CBR services may be controlled effectively by peak rate allocation methods, whereas VBR services, being statistically multiplexed to gain efficiency, may not be effectively served by peak rate allocation, as has been discussed in previous sections. VBR services gain efficiency by *over-allocating* link capacity, i.e., by admitting a number of VBR sources the sum of whose peak bit rates is greater than the total link capacity. Bandwidth allocation of user sources may be based on measurement of some current traffic parameter that may indicate actual cell arrival processes, or may be based on user declared traffic descriptors. If over-allocation is too aggressive, the result is congestion and cell loss when too many sources simultaneously transmit. If too conservative, the result is under utilization of allocated links. ABR and UBR services are “best effort services”. ABR traffic may be admitted to a connection by simply considering the sum of the Minimum Cell Rates in relation to the available bandwidth and buffer resources. UBR traffic, since it is always discarded when there is network congestion, may be admitted based on limiting total numbers of UBR connections so as not to exceed processing power of network elements, given overall traffic loads. Alternatively, UBR traffic admission may be simply limited to a network provider’s estimate of the lowest performance users will be willing to accept.

VBR services, being by nature “bursty”, in many cases provide users with a challenge in providing accurate traffic descriptors to the network. Various policing methods have been proposed for enforcing the user declared source mean rate and peak rate, as well as maximum burst sizes and intervals between bursts, and the ATM Forum has defined a Generic Cell Rate Algorithm (GCRA) for this purpose. GRCA policing either discards cells or “tags” them with a lower CLP if they do not conform to the user specified parameters. This cell tagging, if the GRCA policer determines it to be necessary, may lead to cell discard at other points in the network should the network become congested. Enforcement of cell rates by the GCRA will insure minimum congestion in the network, but may adversely effect the QoS provided to sources that do not actually conform to user predicted values by transmitting at a higher than predicted cell rate. If user predicted values are set with parameters guaranteed to be sufficiently high to avoid this, inefficiency in the network may result as network resources go unused if the source does not actually transmit at this high rate.

With this in mind, some alternatives for placing the various classes of traffic into service in the network are discussed below.

1.2.4.1 Path Sharing/Separation

One possibility for placing traffic into the network is that all classes of service share in a given virtual path's allocated resources. This may lead to inefficiency in network utilization if required QoS performance objectives must be achieved in common for all classes of service, requiring the use of more bandwidth than may actually be necessary if individual QoS objectives could be considered. On the other hand, if link efficiency is maximized through over-optimistically allocating bandwidth in the shared path to accomplish statistical multiplexing, the allocation may lead to QoS degradation for all the various classes of service. This alternative requires that the virtual path resources be fixed over a relatively long term, since short term dynamic re-allocation of path bandwidth is made more difficult by the presence of ABR traffic. An alternative is the separation of various classes of service into different virtual paths.

Some approaches to the problems noted in previous sections require certain structuring in allocating resources. This typically requires the division of services into classes that have common performance requirements or statistical characteristics. Bandwidth is allocated within the various classes, allowing statistical multiplexing only within groups with common characteristics. Veciana, Kesidis and Walrand [VEC95], in their evaluation of call admission schemes using effective bandwidth, or equivalent capacity, (see Section1.2.3.2.2), discuss separation of service by priority, allowing differing QoS requirements to be met efficiently. Elawid and Mitra [ELA95] have also used service separation by priority and refinements to effective bandwidth techniques to guarantee QoS. Bolla, Davoli, and Marchese [BOL97] describe a service separation philosophy where control of resources required by different types of traffic is decomposed from a complex problem having different time scales and requirements at both cell and connection levels into smaller and independent problems. They describe a methodology based on equivalent bandwidth that takes into account constraints on all loss probability and delay, and yields the maximum number of connections that each class can support, based on the definition of a “feasibility region” for each class. Four different schemes, all based on service separation, are described, and are evaluated by the authors as performing well.

1.2.4.2 Switch Architectures

Since packets may arrive at a switch in an unscheduled manner, two or more packets may arrive on different inputs destined for the same output. The switch architecture must allow one of the packets access to the output, but the others must be held in a queue, if possible, for later output.

The characteristics of various switching fabrics connecting inputs to outputs are widely known [PRY96], and many approaches have been proposed for providing the queuing needed to facilitate the flow when packets destined for the same output arrive at a switch. Approaches are: *input queuing*, where a separate buffer is provided at each input; *input smoothing*, where a frame of b packets is stored at each of the N inputs to a switch, and simultaneously launched into a switching fabric of size $Nb \times Nb$; *output queuing* where packets are held in first-in first-out (FIFO) buffers at each output of the switch; and *completely shared buffering* where all queuing is done at the outputs and all buffers are completely shared among all output lines. The throughput-delay performance characteristics and tradeoffs made in choosing one of these basic architectures over another have been examined and analyzed (see for example [HLU88]).

In addition, variations on the basic architectures, such as the two-stage output buffering and utilization of *weighted round robin* (WRR), weighted fair queuing, or other server schemes, have been proposed in an effort to optimize network ability to fairly meet overall QoS objectives while optimizing overall network efficiency characteristics. WRR multiplexing consists of sequentially serving traffic queues in proportion to their weight, which is calculated based on the ratio of the required rate of the particular traffic scheme to the overall rate of the path in which it is included. See [KES96] for a discussion of WRR. Recently, WRR schemes have become popular in proposed switch architectures, due to their ability to insure bandwidth allocation, fairness in serving queues, and minimal delay variation (see [MEZ94] for a discussion of the advantages of WRR versus weighted fair queuing). Various scheduling algorithms have been proposed for WRR servers to insure smoothness in serving the associated queues (see [WAN94] for an example).

1.2.4.3 Flow Control/Connection Policing

The network alternatives here involve the level of policing accomplished by Usage Parameter Control (UPC). Cells that are not conforming to the traffic contract as determined by the Generic Cell Rate Algorithm (GCRA) may both be "tagged" with a low priority CLP indicator and admitted to the network, or the cells may be dropped. If the cells are admitted, they may later be dropped as they traverse the network if congestion is encountered at a particular node.

Schemes which base admission on bandwidth allocation determined through user supplied parameters, such as PCR and/or SCR, use the GCRA to insure compliance with the parameters, through UPC. This can allow users to achieve higher than negotiated resource commitment in certain types of switch during periods of low usage, if the non-conforming cells are merely “tagged” by the UPC. Another possibility is that under-utilization of committed resources can result in cases where a VBR traffic stream uses fewer resources than those committed to it based on its user supplied PCR, SCR, MBS parameters. This can occur if the user, lacking detailed knowledge of the true statistics of the traffic, typically defines parameters in a conservative manner sufficient to always insure a desired level of quality (i.e., cells are seldom or never “tagged” with a low priority CLP indicator).

1.2.4.4 Network Configuration

Networks may be defined as those where all switches can serve both as virtual channel switches and as virtual path switches. Alternatively, networks may be defined such that some switches serve both virtual channels and virtual paths, while other switches serve only virtual paths.

The virtual channel *source* and *destination* nodes under this concept may be any network node which accepts or delivers user traffic, in smaller networks, or may be a boundary node between an edge network and a core network, in larger network architectures. Such switches necessarily serve both virtual channels and virtual paths.

Figure 1-2 shows a network model where all switches serve as user source/destination nodes, thus possibly being both VC and VP switches.

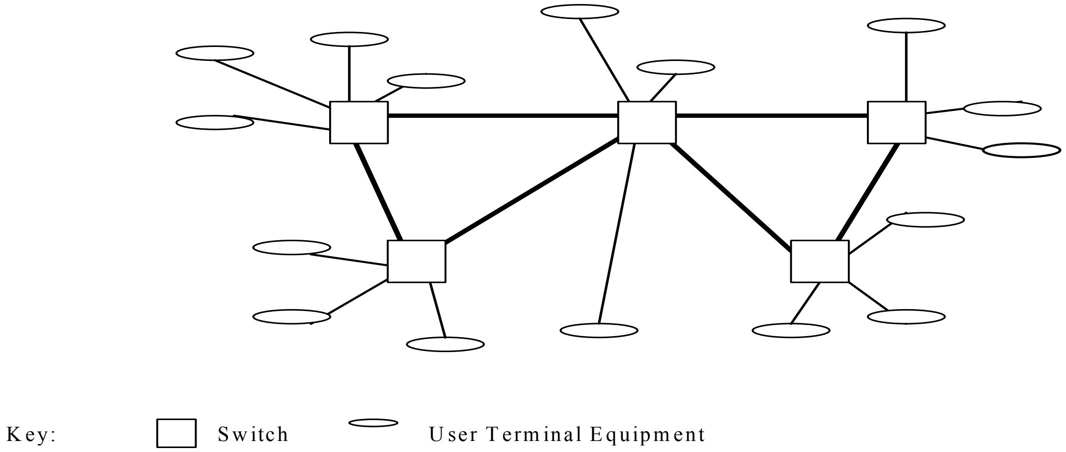

Figure 1-2: Network Model, No Core Switches

Figure 1-3 shows a network model where some switches do not serve users, and thus switch only VPs. This is the “core” and “edge “ switch model, where edge switch nodes serve users, or are gateways to user serving networks, and core switch nodes do not serve users directly, being only connected to edge nodes. Core switches therefore switch only virtual paths.

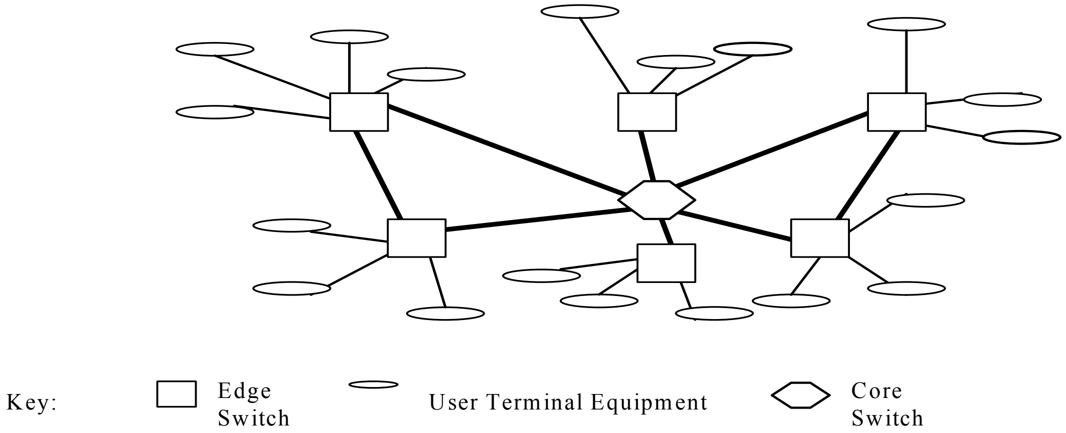

Figure 1-3: Network Model, Using Core Switches

1.2.4.5 Network Control Frameworks

A network-wide method and infrastructure for controlling the admission of requested connections and timely re-allocation of bandwidth between nodes is needed by any network that is to be both effective in congestion control and truly efficient in resource utilization. In some network frameworks, a connection request that can be allowed at a source switch in a network, based on the resource commitment required there, may not be able to be supported at a switch further along the source-to-destination path, based on its existing resource commitments. A path bandwidth re-allocation request must be able to be supported at all switches along the path before it can be allowed. There are several network control framework alternatives that may be employed in implementing CAC.

At least three major alternatives for network-wide control of CAC are possible. These include:

(1) CAC oriented toward using criteria for admission to a specific single switch output link, given a certain set of conditions in the individual switch, (buffer occupancy, utilization, etc.), with connection admission proceeding switch by switch, one hop at a time, from source to destination.

(2) CAC where a set of conditions for each physical link in the network are maintained and used to determine available end-to-end bandwidth for the admission decision.

(3) Those CAC schemes oriented toward end-to-end paths where resources to serve the entire virtual path at a specific bandwidth are pre-allocated and virtual connections are not admitted into the path unless sufficient bandwidth resources remain, as determined at the source.

In case (1), call admission proceeds in a step-by-step manner with each switch along the source-destination route accepting or rejecting the connection based on the switch's existing resource allocation at the time, as shown in Figure 1-4. This method leads to inefficiencies as calls may be accepted by an admitting switch and some intermediate switches, tying up resources, but the calls may be eventually rejected by a switch further along the route between the source and destination.

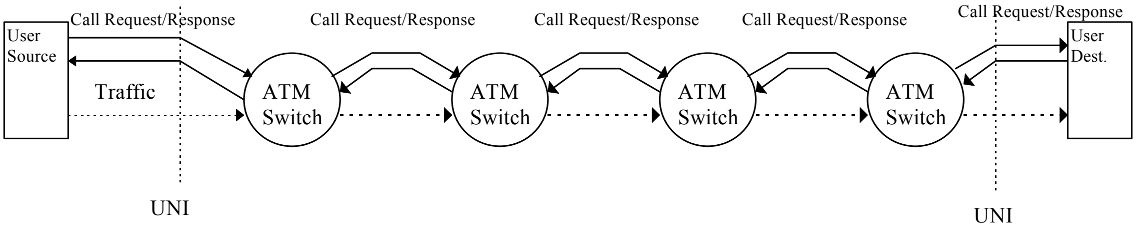

Figure 1-4: Step By Step CAC

In case (2), a central or distributed database of the parameters of each link must be maintained and used by the CAC, as shown in Figure 1-5.

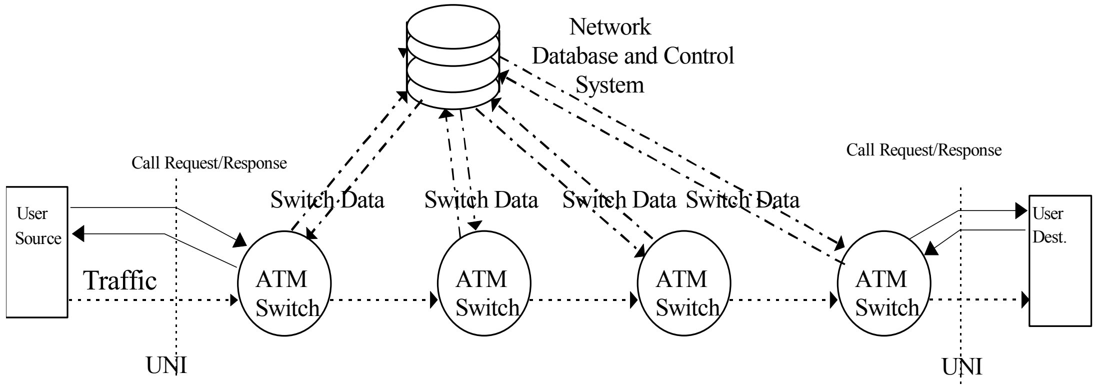

Figure 1-5: Parameter Database and Control System CAC

Although methods such as this have worked well in certain single-use, single-rate applications (e.g. the national voice telephone system), their use in multi-rate, multiple class-of-service applications on a large scale may be difficult because of the system complexity required. Thus, methods based on complex network resource monitoring schemes requiring large parameter databases and complex calculations to determine if admission is allowed have generally not been favorably considered when such schemes are proposed for actual implementation.

In case (3), resources to support a path of a given bandwidth have been committed from end-to-end, and connections are admitted to the path, based on some criteria such as peak rate, or sustained rate. Such a system is shown in Figure 1-6.

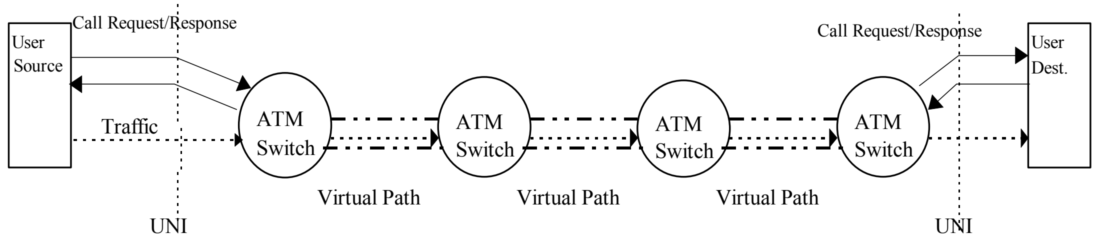

Figure 1-6: End-to-End Path Based CAC

1.2.4.6 Measurement Time Framework

In implementing CAC, the recognition that a certain amount of bandwidth is available may be based on measurements taken over time periods of various duration. The measurement may be instantaneous, and, if VBR traffic is involved, therefore subject to considerable error as the traffic rates change. The measurement may be short term, with predictions of the actual available bandwidth, or long term, where a safe (i.e. conservative) available bandwidth estimate is possible, but short term opportunities to pass higher rates of traffic are missed. Thus, measurement times may not be easily defined in general, and may possibly be dependent on the types of traffic involved.

Consideration of measurement time frameworks is sometimes necessary within certain CAC schemes as well. Some schemes admit a connection at a given rate, and later allocate the connection a different rate, based on measurements after some specified or adjustable time-frame has occurred. An example is the method discussed in Section1.2.3.3.2, which includes an adjustable “warming up” period implemented as each connection is admitted.

1.2.4.7 Summary of Alternatives

Starting from network architectures where all traffic classes might share in a given virtual path, in order to efficiently use bandwidth through statistical multiplexing, ATM concepts have evolved to where separation of traffic service by virtual paths, with all traffic in the path sharing similar characteristics, has also become an accepted and efficient method. Various switch architectures are possible, with many choices of switching fabrics and queuing schemes. Alternatives exist in the methods of providing flow control to best control possible network congestion. Concepts have also been proposed and are being implemented in real networks which involve the partitioning of the network elements, or nodes, into those which directly handle connections to the user versus those that handle only connections of network paths between the user-connected nodes. These concepts allow improvements in overall efficiencies while still maintaining required QoS for the respective traffic classes. However, since continued improvement in the efficiency of allocating network resources is desirable, a new model for CAC could possibly be even more effective in providing the best match of QoS for each traffic type, at the best utilization of network resources, for the least cost in application of computational resources. Choosing particular combinations of alternatives, coupled with novel methods of evaluating the amount of bandwidth remaining on a given path, could yield benefits in further improving efficiency in the use of network resources. The evaluation of such an overall combination, the model for which is presented in chapter 2, is the purpose of this research. An overview is presented in the next section.

1.3 Dissertation Overview

1.3.1 Organization

This dissertation is organized into five chapters, including this chapter of introduction. Chapters 2 through 4 each include a summary of the results obtained in the chapter. All summaries have been gathered together in the next section, in order to provide an overview of the work, and its significance.

1.3.2 Chapter Summaries

1.3.2.1 Chapter 1: Introduction

This chapter outlines the various aspects of ATM, methods of congestion control, and specific means of connection admission control (CAC). Network implementation alternatives affecting the methods used for CAC are discussed as background to the research.

1.3.2.2 Chapter 2: A New Method for Determining Available Bandwidth

In Chapter 2, a new method for determining available bandwidth in a virtual path that is carrying existing traffic and is considered to be “full” by existing CAC methods is broadly outlined. The result includes:

1) The development of a new general model for the analysis of traffic serving mechanisms designed to allow the introduction of "additional" traffic into a virtual path that is considered to be "full" by CAC algorithms currently in use. The model is based on use of weighted round robin servers having queues sized in accordance with a described network framework, associated with a particular method of determining required queue sizes, on dedicated end-to-end virtual paths. The proposed method is a means of guaranteeing required quality of service to admitted CBR and VBR traffic.

2) The use of the model and framework described above, where the probabilities of having an empty queue when serving "bursty" variable bit rate traffic is shown by simple illustration to lead to an ability to insert "additional" traffic into a virtual path server schedule without affecting the existing traffic.
3) The need to obtain estimates of the number of empty timeslots per scheduled server cycle through an adequate measuring technique, in order to make use of the new method.
4) The choices made in picking implementation alternatives, assumptions, and switch architectures needed to implement the proposed new method.
5) An outline of a possible CAC technique based on the use of the proposed new method.

6) A broad methodology for evaluating the capability of the proposed new method.

The chosen method for determining available bandwidth, and an associated CAC technique, is further developed and analyzed in detail in Chapter 3

1.3.2.3 Chapter 3: Analysis of WRR Queue Service Measurements

This chapter is focused on specific methods for adding traffic to broadband network paths. It contains an analysis of WRR operation in a source switch that is part of an overall network framework that requires end-to-end virtual paths, and separates the paths serving VBR and CBR traffic from those serving ABR and UBR traffic.

The analysis shows how a network user, or customer, may calculate SCR of requested traffic connections, in spite of lack of specific knowledge regarding the size of the queue serving the traffic. This may be accomplished through solving two simultaneous equations involving queue length and SCR, if the user has knowledge that the network provider is using the method suggested in this chapter. The analysis shows how the mean number of empty timeslots per cycle in a source switch WRR will be a function of the SCR declared by the user, and suggests two methods for determining available capacity in a virtual path considered by normal CAC methods to be full. In both methods, the available or “additional” capacity is a function both of the mean number of empty timeslots per WRR cycle, and the desired CLR. In an actual system, where the true mean number of empty timeslots per cycle is unknown, analysis shows that the sampled mean number of empty timeslots per cycle, and the associated estimate of the variance, may be used in creating statistics needed for hypotheses used to test the critical admission criteria. The methods used are conservative in allocating allowable “additional” bandwidth, being based on using the boundaries of the required relationships. Finally, this chapter describes a method for admission of “additional” connections to a “full” virtual path, without exceeding required QoS for any connection. The next chapter will address WRR server and traffic simulations used to support the analyses presented here.

1.3.2.4 Chapter 4: Evaluation and Implementation

In this chapter, the analysis developed in Chapter 3 was tested through simulation of traffic and of the WRR mechanism serving the end-to-end virtual path of the overall model framework. The simulations, using models of on-off traffic sources and WRR server mechanisms of up to twenty queues and twenty cell timeslots, were conducted over a wide range of traffic scenarios. The scenarios span a range of traffic conditions for both existing and “added” traffic connections designed to identify ranges of “added” traffic which could cause queue overflows, and the conditions under which the “added” traffic connection queue sizes are required to be augmented above the size needed in the “normal” traffic connections for similar traffic.

Simulations were also designed to validate the use of an allowable mean number of empty timeslots per WRR cycle as the basis of an allowable “additional” traffic sustainable cell rate (SCR). The allowable “additional” SCR is based on measurement of samples of the mean number of empty timeslots per WRR cycle and an estimate of the variance of the mean number of timeslots per WRR cycle, which is shown to be dependent on the correlation of the existing traffic stream in the virtual path.

The results of 15,374 simulations, spanning a significant portion of the design space, conducted over a duration of at least 100,000 virtual path timeslots each, support the analysis presented in Chapter 3. The next chapter will present conclusions, including a detailed discussion of the CAC algorithm, and possible implementation methods.

1.3.2.5 Chapter 5: Conclusions and Future Directions

In chapter 5, the implications of the results obtained in earlier chapters are discussed. Issues remaining open from this research are identified, with discussion of possible areas of future research.

Chapter 2

A New Method for Determining Available Bandwidth

CHAPTER CONTENTS

LIST OF FIGURES

Chapter

2

A New Method for Determining Available Bandwidth

This chapter proposes a new method for determining available path bandwidth for use in connection admission control (CAC) of integrated broadband networks. The method allows the introduction of additional traffic channels into a virtual path that other methods of CAC would consider as being full, thus improving bandwidth utilization.

In the previous chapter an argument was made that existing CAC schemes are based on methods which, due to lack of exact knowledge of the specific statistical nature of the traffic, may be overly conservative in allocation of bandwidth to connections; or may require a complex data collection effort as certain network link parameters are maintained for use by the CAC at each node; or may be based on certain compromises to obtain both required service quality guarantees and reasonable network efficiency; or may have some other hindrance to effective implementation, such as excessive hardware complexity, which would lead in turn to excessive equipment cost. The model proposed here is specifically an attempt to address the above issues.

In the model proposed here, mechanisms are developed for the use of measurements made within weighted round robin (WRR) queue serving structures as a means of facilitating CAC. This will be done within existing ATM based network architectures, standards, and techniques, such as those promulgated by the ATM Forum and ITU-T, or in such a manner as to facilitate the addition of new standards without precluding network operation under existing standards. The development of suitable techniques will require the refinement of the model presented in this chapter, through analysis of system models to bound system performance, and simulation under a wide variety of possible network traffic types and operating conditions, to confirm the analysis. This chapter will describe a general approach to the goal of improving bandwidth efficiency through estimation of actual available bandwidth.

2.1 A Model for Determining Available Bandwidth

In reviewing the approaches used in various CAC schemes, outlined in Chapter 1, it is apparent that, broadly speaking, they are based on either bandwidth allocation or buffer evaluation schemes, i.e., how much of a given bandwidth or buffer is in use. Bandwidth allocation schemes evaluate an effective bandwidth, or equivalent capacity, from the parameters describing the traffic and base connection admission on the sum of the equivalent capacities versus the total capacity of the path, or link, depending on queuing schemes. Buffer evaluation schemes determine the point at which a given quality of service (QoS) objective will be adversely affected if the connection is admitted to a buffer of a given size. Since the equivalent capacity approaches are based on a calculation that includes the size of the queue receiving the traffic, it may be said that these approaches are also based on the amount of buffer fill. Measurement based schemes are also primarily based on measuring buffer fill or overflow levels, i.e., how much of a buffer is in use. Many schemes for determining whether a requested connection is admissible depend on extensive or complex calculations that could preclude their use in an actual system, which is subject to computational resource limitations. Schemes, such as Measurement of End-to-End Delay, described in Chapter 1, that attempt to directly determine available free bandwidth on an end-to-end path without direct knowledge of channel buffer size or fill, are dependent on extensive knowledge of overall network physical parameters. These schemes therefore require an extensive means of providing and maintaining a record of the parameters for all network nodes, to be available at each possible source node, as well as requiring somewhat complex computation, leading to management system complexity across an entire network structure.

The relatively recent introduction of weighted round robin (WRR) queuing schemes, described in Chapter 1, to specific ATM network architectures for the purpose of insuring fairness in serving certain traffic types suggests a *new method of evaluating available bandwidth* on a path for use in the implementation of CAC. The method is based on evaluating bandwidth available to users that remains on a path by evaluating, over time, the number of *empty* virtual channel queues, rather than the *amount of fill* of virtual channel queues. In instances where no additional connections would be admitted because the allocated bandwidth of the individual channels, or connections, equals the available path bandwidth, the timeslots corresponding to empty queues may be used to accommodate additional traffic, rather than only being available to existing active traffic, essentially serving low priority cells. An illustration of this may be based on the operation of a WRR cell scheduler having the configuration illustrated by Figure 2-1.

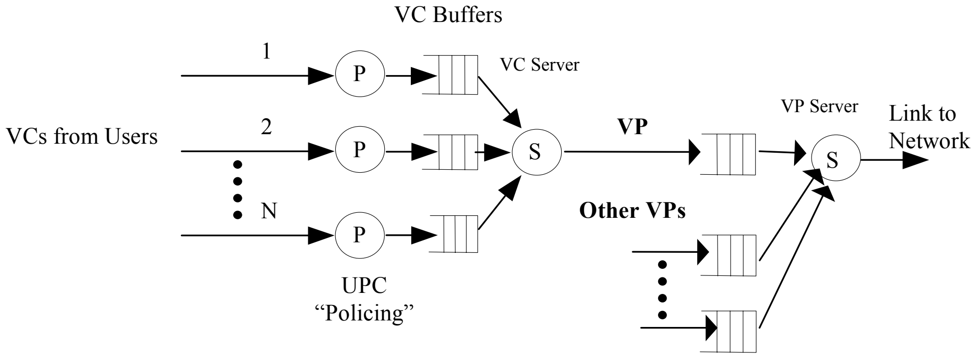

Figure 2-1. Basic WRR Path Queuing Model

In this model each VC has an associated usage parameter control (UPC) policing function, using a generic cell rate algorithm (GCRA), described in Chapter 1 and [UNI3.1], at the interface to the user to insure that the incoming traffic conforms to the traffic contract. In general, conforming cells are marked with a high cell loss priority (CLP) indicator and non-conforming cells are marked with a low CLP indicator

Normally the VC queues are push-out queues, so a low-CLP cell in a queue is discarded if a high-CLP cell arrives at a full queue. Also, bandwidth sharing normally takes place when a connection does not have enough cells in its queue during a serving cycle, allowing the empty slots to be used in serving queues from other active connections.

Buffers that are sized to insure required QoS and served as indicated by this model will allow cell-loss and cell-delay parameters conforming to the traffic contract, providing that GCRA-based UPC policing, with parameters derived from the user's traffic descriptor, is in place. Should the SCR be an underestimate of the rate of arriving cells, the UPC will police the excess cells, causing them to be tagged as low priority, or dropped, thereby insuring the quality of other connections in the VP. Should the SCR be an overestimate of the rate of arriving cells, i.e., a conservative estimate, buffers set to the sizes above will be underutilized, as indicated by timeslots scheduled for the queue being "skipped" with a higher than expected probability as the WRR cycles through its schedule. Should this occur, it implies that additional capacity is available in the VP, and therefore that additional connection(s) may be admitted to the VP. Thus, the amount of available bandwidth in the VP may be estimated from *measurement of the number of scheduled timeslots that are "skipped*" when the corresponding queue is found to be empty as the WRR cycles.

It is also possible to determine the amount of over- or under- estimation of the resources allocated to a VC by monitoring the amount of queue fill for *each* VC in the VP, resulting in *multiple monitoring points* per VP. However, the method proposed here allows the monitoring of a *single point*, the output of the VP WRR server in each CBR/VBR VP.

Under the scenario proposed for evaluation by this research, each VC queue has a certain probability of being empty (or of having a low-CLP cell at the head of the queue, which means it could be skipped if necessary). The probability that the queue is empty is based on its traffic characteristics and the rate at which it is being served. *Additional queues could then be added to the WRR and served at a rate corresponding to the rate that existing queues were measured as being empty or having low-CLP cells at the head of the queue*, without affecting the overall flow of the high-CLP cells. Such a measurement would be taken directly from the WRR in terms of bandwidth in cells per second, since each queue skipped corresponds to a cell which could be served. For use in analysis of this model, the probability that a given input channel queue is empty (or contains only low CLP cells) may be calculated based on the average cell arrival rate.

Assuming for the moment that low-CLP cells are dropped prior to entering the queue and are never served, measurements may be taken in the WRR to determine the rate at which queues are found empty when served. The WRR algorithm will skip a queue that is empty, continuing on to the next queue in the schedule. At some point in time an idle cell will be generated for every queue skipped, since the overall service rate will be effectively higher than the total of arrival rates, and other queues will be served faster for every queue skipped. It will be possible to admit another connection having a SCR equal to the rate of skipped queues, using a queue properly sized to handle the MBS of the connection. Illustrating the process, if a path is being served at cell rate $R_{VP} = 6$ cells/second, and three traffic streams having $\mathrm{SCR}_1 = 1$

cell/second, SCR_2 = 2 cells/second, and SCR_3 = 3 cells second are entering the WRR, normal CAC procedures would not allow another new connection to be admitted until one of the existing connections was terminated. The schedule under these conditions might be as shown in Figure 2-2.

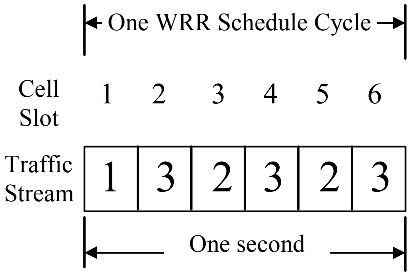

Figure 2-2: WRR Scheduling of Three Traffic Streams (Example)

Now let us suppose that measurement of empty slots taken during the WRR schedule cycle indicates that every three cycles one connection buffer (queue) is empty. In a schedule such as described above, occurrence of an empty buffer slot in the cycle would result in the transmission of an idle cell, if a rigid schedule is in place, as illustrated in Figure 2-3.

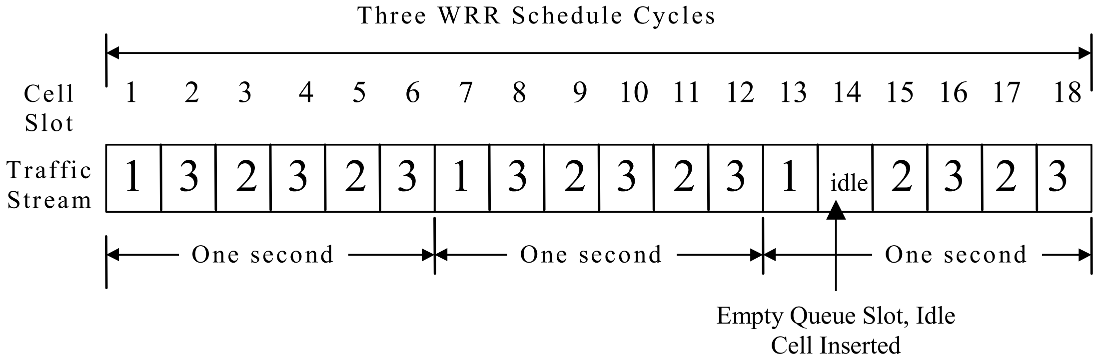

Figure 2-3: Rigid Cycle Scheduling with Idle Cell in Empty Slot (Example)

Another possibility is that, using true WRR scheduling, the empty cell slot could be skipped, with the WRR going on to serve the next scheduled slot as shown in Figure 2-4.

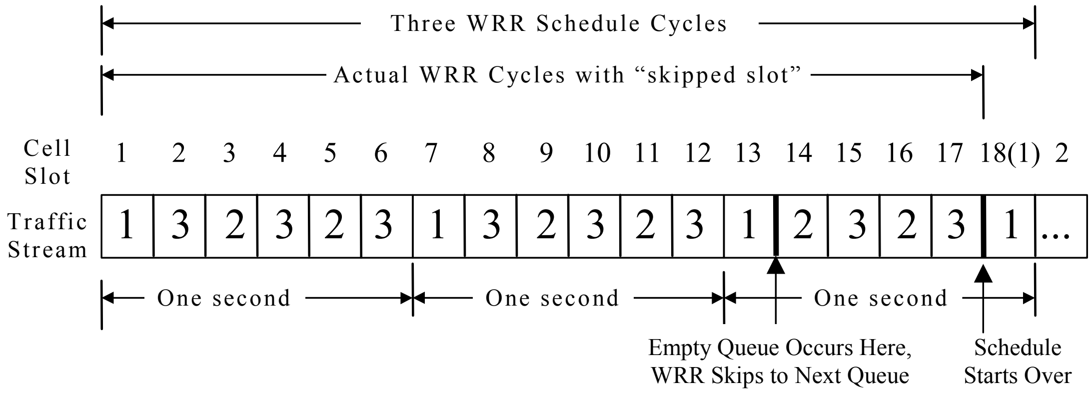

Figure 2-4. WRR Scheduling with Empty Slot (Example)

This would allow a new traffic stream (connection 4) having $SCR_4 = 1/3$ to be admitted without adversely affecting the existing traffic. The WRR schedule could be modified by the addition of a timeslot to accommodate the new connection, as shown in Figure 2-5.

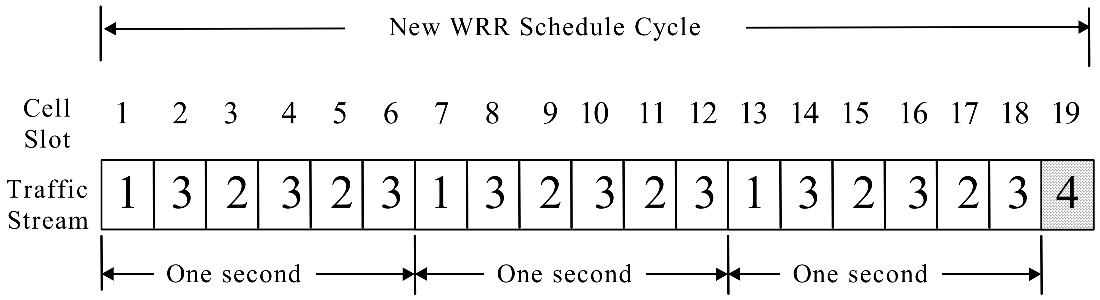

Figure 2-5: WWR Schedule with Added Traffic Stream (Example)

However, in a given cycle, where, based on the previous empty cell timeslot measurement, there is a high probability that one cell timeslot would be skipped during a cycle due to a queue being empty, the schedule might appear as follows:

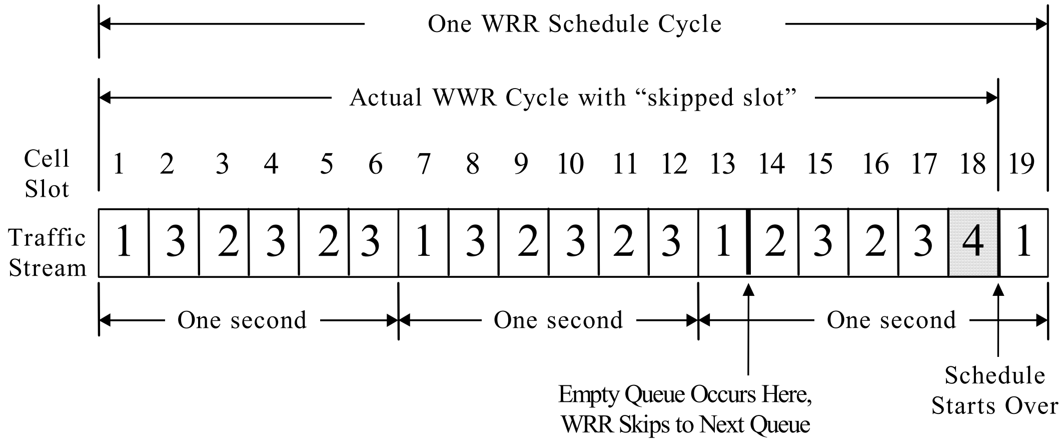

Figure 2-6: WRR Schedule with Queue Skipped (Example A)

Since there is a finite probability of no “normal” queue being empty, and therefore skipped in a given cycle, a queue would be required for each “additional” connection being served by the WRR. The desired CLR of the additional connection would determine the required queue size.

The previous method places an additional timeslot in the cycle schedule, with the probabilities being that an existing queue will be skipped often enough that the actual new cycle time will not exceed the previous cycle time, thereby assuring the QoS of both the “normal” and “additional” connections. An alternative WRR serving scenario would be to hold the additional queue(s) separate, serving it (them) at the time one of the “normal” WRR queues was skipped. This would require buffer space for the "additional" traffic queues based on the rate at which none of the "normal" queues were skipped, and the PCR, SCR, and MBS of the "additional" traffic. Under this scheme the original schedule of the "full" WRR would be maintained, with the additional traffic stream cells transmitted in the appropriate skipped queue slot, as shown in Figure 2-7.

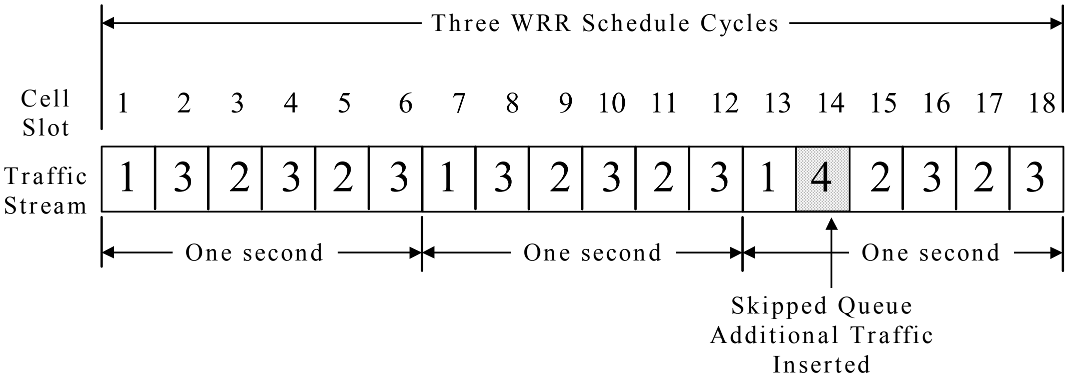

Figure 2-7: WRR Schedule with Queue Skipped (Example B)

Under this alternative, the available bandwidth, once obtained through measurement of the skipped timeslot rate, may then be used as the decision criteria for CAC, with no effect on cell loss or cell delay parameters of the already existing traffic.

Now considering the case where low-CLP cells may be admitted to the queues and reach the head-of-queue, an algorithm for handling the low-CLP cells would have to hold them until a schedule cycle was completed, and only send them if other queues, including the additional traffic added to the path using this model, were empty. If no slots were available in the cycle, low-CLP cells at the head of queues would be discarded. Since there would still be some probability of empty cells occurring within a given schedule cycle, when the “additional” connection queue was empty, some low-CLP cells would still be sent on a path using this CAC model.

To implement an actual network wide system for optimizing resources based on this scheme for using WRR queue serving, choices must also be made from among the other alternatives discussed in Chapter 1.

2.2 Choice of Additional Implementation Alternatives

Having chosen WRR as the queuing/serving scheme for CBR/VBR traffic, choices must also be made regarding other possible alternatives for implementing ATM networks. The concept of some level of service separation, coupled with a specific network model utilizing WRR queue servers, has been evaluated by Liu, et al., [LIU97-2]. In order to reach a reasonable compromise between the two conflicting goals of guaranteeing performance levels for each service class while still allowing sufficient statistical multiplexing such that the network is effectively used overall, [LIU97-2] propose a specific framework for the network which addresses the majority of implementation alternatives discussed Chapter 1.

The proposed network framework developed in [LIU97-2] includes routing of the CBR/VBR service types and ABR/UBR service types into separate virtual paths, and the partitioning of the physical network into “core” and “edge” sub-networks. The main purpose of the edge network is to provide network access to the user of broadband services, whereas the core network carries the traffic concentrated in the edge networks and passed to the core network through gateways. The core network under this concept is based on virtual paths (VP), where a virtual path starts at an edge gateway and terminates at another edge gateway. In the core network, the available bandwidth and buffer space are managed efficiently on a per-VP basis. At least two VP's are assigned between each edge-node-pair, one for VBR and CBR service, and the other for ABR and UBR service. Separating ABR traffic from the VBR/CBR VP insures that rate changes in large quantities of ABR circuits in a given VP do not affect the QoS performance of CBR and VBR service classes.

This framework proposed in [LIU97-2] requires that VP bandwidth be allocated on a relatively long-term basis, with the size of the VP based on physical link traffic forecasts and estimates, as bounded by physical link capacity. This provides the advantage of simplifying the VC-level CAC. The CAC is simplified because the decision to accept a CBR or VBR connection can be made at the corresponding source edge gateway switch by comparing the bandwidth requirement of the new call and the amount of allocated bandwidth on the VP which is to carry the new call. CBR connections are admitted if the PCR can be accommodated by the VP, and VBR connections are admitted if the SCR can be admitted by the VP. UPC policing insures that SCR will not be exceeded under congestion conditions.

The maximum queuing delay and maximum queue length for CBR/VBR services are evaluated in [LIU97-2], using source traffic descriptors, based on WRR servicing at the ingress and egress CBR/VBR VC’s. Numerical analysis by the authors shows that the proposed bandwidth management scheme provides reasonably small buffer sizes and can obtain satisfactory queuing delay performance on a typical large (coast to coast, USA) ATM network.

This possible network framework therefore potentially allows all ATM service classes to be served within QoS guarantees, CAC procedures can be easily implemented, and rate-based ABR congestion control easily implemented. Thus, the overall framework proposed in [LIU97-2] addresses most of the implementation alternatives proposed in Chapter 1, providing an efficient base for additional improvements in the efficiency of overall network bandwidth allocation.

In summary, the overall framework specifically used with the new model proposed here will be based on choice of the following alternatives to obtain the advantages outlined above:

- Service separation into two types - CBR/VBR paths, and UBR/ABR paths.
- Non-blocking switch architecture, with separate WRR queues for channels and paths.
- An edge/core network configuration.
- Flow control is required, and will be based on UPC using the ATM Forum GCRA, or an equivalent algorithm.

- Admission will be based on CAC at the source switch only (i.e. type 3 CAC from Chapter 1).

Choice of the alternatives discussed above will result in an overall network model well suited for the implementation of the proposed scheme for the utilization of WRR timeslot vacancy measurements to determine available bandwidth on CBR/VBR paths.

2.3 Implementing the New Model in a Network

2.3.1 Additional Assumptions

In order for this method based on measurement of empty queue statistics in WRR servers to function as an effective basis for CAC, the following specifics to the overall framework are assumed:

2.3.1.1 Virtual Paths

Virtual paths (VP) start at a user source node and end at a user destination node, traversing any intermediate switching nodes as a whole. The resources committed to the VP at each node are sufficient to guarantee common network-wide maximum cell delay, cell delay variation, and cell loss ratio parameters to VBR and CBR traffic, and common maximum cell loss ratio parameter to ABR traffic, under non-congested conditions.

2.3.1.2 Current VP Bandwidth

The decision to admit a connection to a given CBR/VBR VP will be based on the current VP bandwidth of the CBR/VBR VP in place between the source and destination nodes, regardless of the state of other unused bandwidth in the network at that time.

2.3.1.3 Bandwidth Allocation of Connections

New connection bandwidth and buffer sizes sufficient to maintain the required QoS are allocated based on the network user's source traffic parameters, provided by signaling the appropriate descriptors across the user-network interface (UNI). Knowledge of the bandwidth of the traffic, its class, and the requirements of the CGRA, based on the parameters provided by the traffic descriptors, will allow the establishment of channel buffer queue lengths sufficient to support the required QoS, using relationships developed by Liu, et al., as described in Section 2.3.2. The use of these descriptors will allow resource allocation and connection admission as follows:

- The value allotted to the bandwidth of a new VBR channel connection request will be the SCR, which will be based on an equivalent bandwidth calculation, or other method where the size of the queue in the traffic stream is needed in the solution of the allowable quantity of traffic at a given CLR.

- The value allotted to the bandwidth of a CBR connection will be the PCR.

- The value allotted to the bandwidth of an ABR connection will be the MCR.
- UBR connections will be admitted based on criteria other than bandwidth, as determined by network providers.

2.3.1.4 Network Connectivity

To allow for CAC at the source node only, each source node will be connected to every possible destination node with at least one VBR/CBR VP and one ABR/UBR VP, i.e., the ATM network is “flat” rather than hierarchical, and uses service separation. This assumption will require the creation of large numbers of VPs in a network, but may be supported even in large-scale networks as follows:

- The ATM network-to-network interface (NNI) allows for 12-bit virtual path identifiers in the cell headers, thereby allowing a particular switch a maximum of 4096 virtual paths in any given physical link through an associated physical port. Provisioning source-to-destination VPs from switch to switch on a long-term basis (see Section 2.3.1.3, above) would provide end-to-end paths for user traffic.
- Using an architecture that requires both service separation, as suggested here, and physically redundant paths for use in case of network failure, would require at least four VPs be directed from a given edge node to all other edge nodes; two carrying CBR/VBR traffic, and two carrying UBR/VBR traffic,

one of each type over separate physical ports. Thus, up to 2048 redundant VP connections to other edge switches could be supported on two links through two physical ports in the switch. Additional ports would increase the number of edge nodes that can be supported, or increase the total bandwidth available to a smaller number of edge nodes. This allows support of an extensive “flat” VP based network, given the port capabilities of current ATM switching equipment. Should physical limitations on the number of ports in a given switch become a constraint on the number of nodes that could be accommodated in a flat network, as a network became very large, the network could be divided in two, with each half being a flat network using the architecture described here, with each half being connected to the other through gateway edge switches. This would allow each sub-network to perform with the efficiencies described for this architecture.

2.3.1.5 Traffic Parameter Negotiation

No re-negotiation of source traffic parameters, *once the connection is admitted*, will be accomplished between the network and the user. The method described in this proposal is intended to be an alternative to methods that use renegotiating the traffic descriptor parameters during the time the connection is in existence as a means of increasing the number of connections on a given path through measurements of the existing traffic streams.

2.3.2 Switch System Architecture

A VC source switch model incorporating the WRR functions proposed in Section 2.1 for VBR/CBR traffic is shown in Figure 2-8.

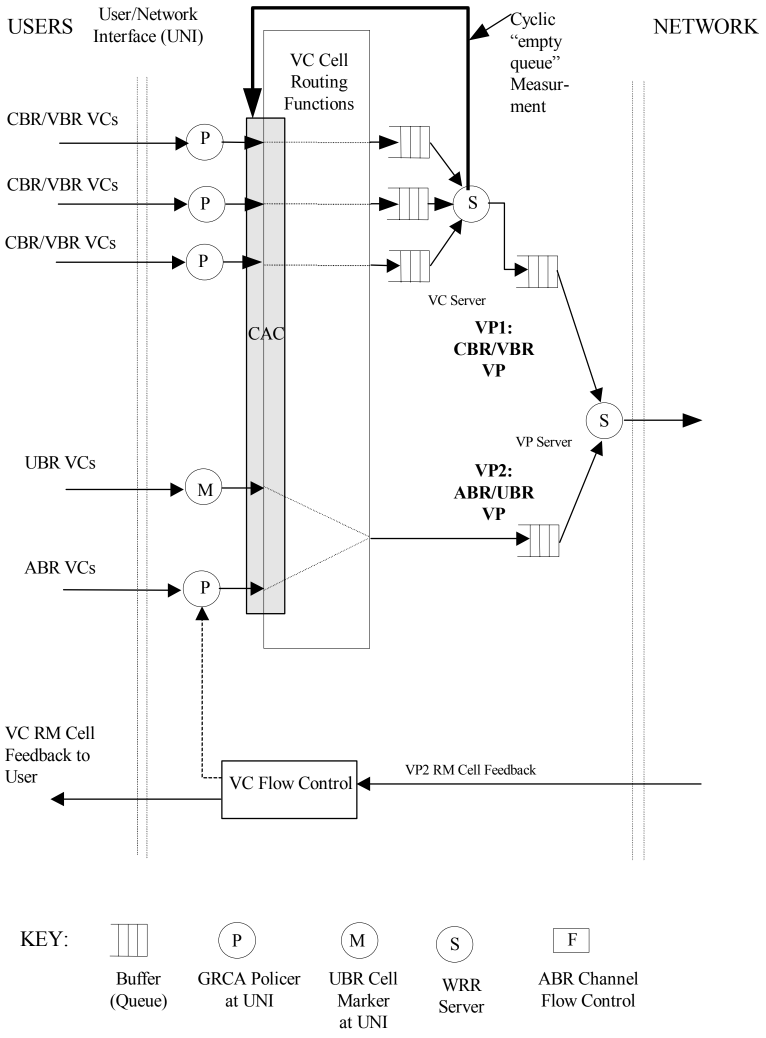

Figure 2-8: Source Switch Functionality

The architecture of Figure 2-8 is similar to that discussed by Liu [LIU97-1] and [LIU97-2] with the exception that no source re-negotiation is needed to effectively utilize path bandwidth, and only one measurement point per VP is needed, rather than one for each VC in a VP. As can be seen from Figure 2-8, VPs dedicated to user ABR/VBR traffic also exist in the switch, but will not be discussed in detail, as the admission criteria for this type of traffic is straightforward, as noted in Chapter 1.

Queue lengths in this architecture are dependent on the maximum burst size (MBS), of the traffic, in cells, as declared by the user of the network. The conforming MBS for traffic streams having Burst Tolerance (BT) of τ_s with Peak Cell Rates (PCR) of $1/T$ and Sustained Cell Rates (SCR) of $1/T_s$ (from ATM Forum definitions outlined in [TMS4]) is:

$$MBS = \left\lfloor 1 + \frac{\tau_s}{T_s - T} \right\rfloor \tag{2.1}$$

This equates to a MBS of 1 for CBR source traffic and a MBS that is dependent on the SCR and BT of VBR source traffic.

In the VC server, each CBR or VBR VC has a separate queue, m, and is allocated a weight w_m corresponding to its respective PCR or SCR bandwidth. The allocated weight is $w_m = 1/(T \cdot N \cdot CS)_m$ in the case of CBR traffic, and $w_m = 1/(T_s \cdot N \cdot CS)_m$ in the case of VBR traffic, where N is the number of cell slots served in the WRR schedule, and CS is one cell slot, i.e. the time to serve one cell. The output rate of the VC WRR is set to the currently allocated VP bandwidth, which is designated by R_{vp}.

The maximum queue length (MQL) needed for zero cell loss for each VC served by this WRR server at a source node, based on its user provided parameters, PCR, SCR, and MBS, has been shown in [LIU97-2] to be:

$$MQL = 1 + \frac{\tau}{T} \tag{2.2}$$

This relates to MQL of 1 for CBR traffic and a MQL of $1+\tau_s/T_s$ for VBR traffic, where τ_s may be determined from the MBS of the traffic through solution of Equation(2.1). The MQL will be the size of the actual queue serving the connection, as long as queues of size MQL will not cause user required delay parameters to be exceeded.

Another outcome of using this proposed model is that the user may calculate the SCR for VBR traffic with knowledge of what the source switch channel buffer size allocation will be, using equivalent bandwidth techniques. The user must generate a value for the SCR of the traffic to be sent, signaling it to the source switch, but equivalent bandwidth methods require knowledge of the buffer size to accomplish the calculation. The buffer is located in the source switch, parameters of which are not necessarily known to the user. However, since the user will know the PCR of the traffic to be transmitted, know (or estimate) the MBS and average cell rate for the traffic, and know the desired CLR, then the size of the buffer that will be created in the source switch based on the PCR, SCR and MBS values passed across the UNI can be determined by the user based on Equations (2.1)and (2.2).

Possible destination switch functions, also similar to those in [LIU97-2], corresponding to those shown at the source switch, are shown in Figure 2-9.

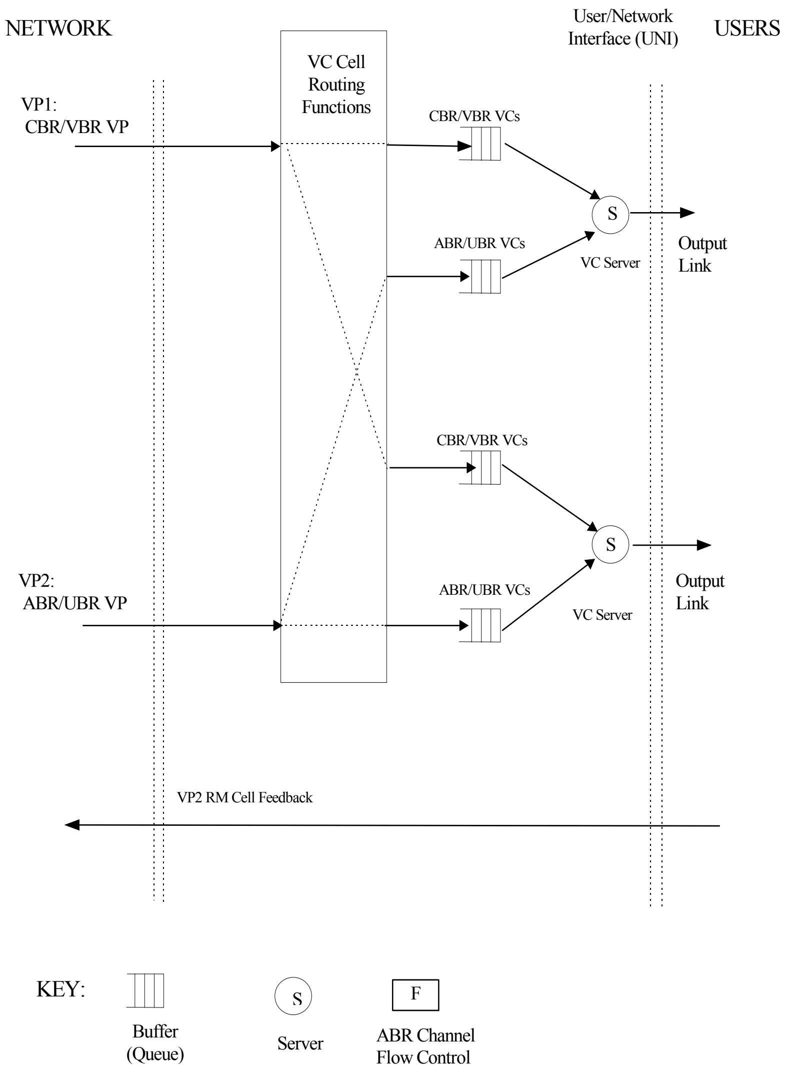

Figure 2-9: Destination Switch Functionality

MQL of the destination switch VC queues are based on maximum queuing delay calculations, as was the case for the source switch. In the case of the destination switch for CBR channels, MQL=2 cells, and for VBR channels, MQL=$2+\tau_s/T_s$ cells as shown in [LIU97-2]. Intermediate switches and the path server in the source switch will perform path-switching functions on end-to-end paths between source and destination switches and will use multiplexers consisting of WRR queue servers. Necessary weights and queue lengths at intermediate switch nodes can also be determined as described in [LIU97-2], i.e., WRR serving rates (weights) would be based on the current VP bandwidth, and the queue serving the path in the n^{th} intermediate node would have a queue length of n, where n can be between 1 and N, the total number of VP multiplexers in the end-to-end path. Combining all switching functions on an end-to-end virtual path provides an overall model, as described in Figure 2-10. CAC for a given VP is accomplished solely in the source switch through measurement of the number of empty timeslots at the output of the VC WRR.

2.4 CAC Functions of the Model

In the model described, the key parameter is the number of empty queues measured as the WRR cycles through its assigned schedule. The mean number of empty queues (or queues with low priority traffic) encountered during WRR cycles represents some amount of unused bandwidth which could be used to support additional high priority traffic, and is denoted as BW_{EQ}. The number of empty queues associated with a given path WRR server will be continuously changing as conditions change because of the statistically changing nature of the bursty VBR traffic.

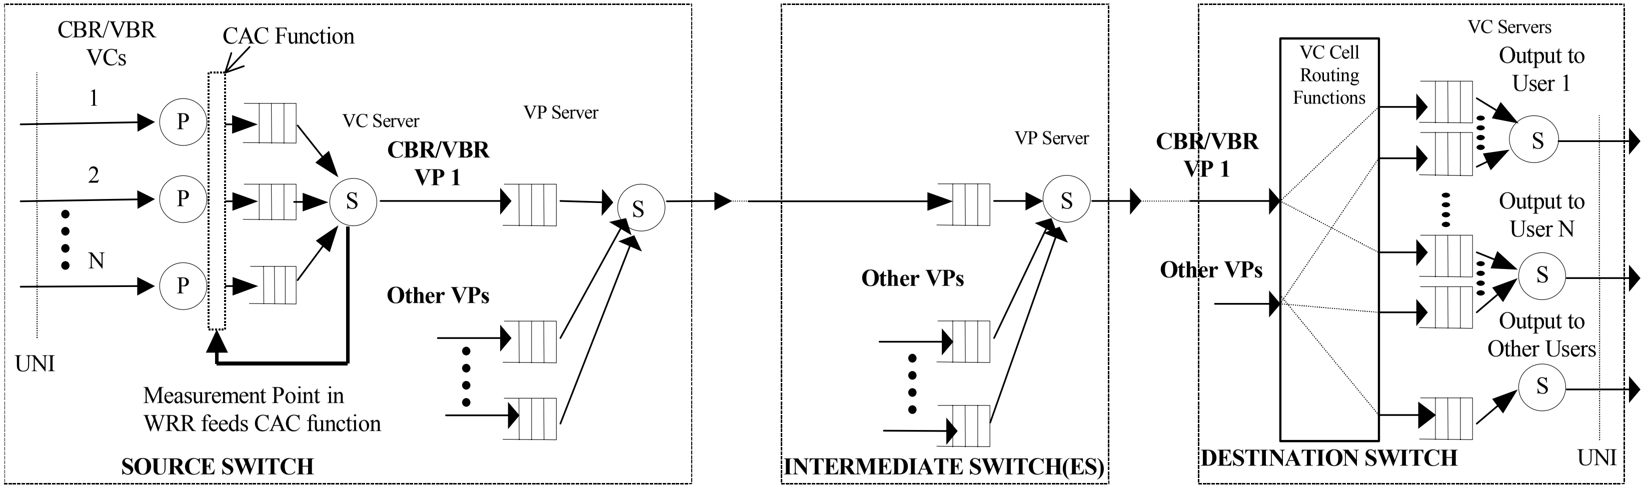

Figure 2-10. Virtual Path with N Virtual Channels to Given Destination

Thus, an effective method of determining a suitable value for the available bandwidth based on a simple statistical measure, such as average, minimum, or maximum quantity of empty queues, or based on some other more complex predictive measure of the usable available bandwidth given the variations of the empty queues over a given time period, must be determined. This research is focused on the development and validation of methods for predicting available bandwidth based on the measurement of the occurrence of empty queues, and their resultant empty timeslots in WRR based systems.

Using this model, the bandwidth-related CAC functions to be performed at the source switch are as follows:

(1) For VBR user traffic

a) Determine the bandwidth required for a newly requested connection in terms of cell bandwidth in cells per second, BW_{new}, using the SCR source parameter provided by the user.

b) Compare the BW_{new} resulting from (1) to the current bandwidth available in the path, in terms of $BW_{avail} = BW_{path} - \sum_{N} SCR - \sum_{M} PCR$, where N is the number of presently connected VBR channels, and M is the number of presently admitted CBR channels.

c) If $BW_{new} < BW_{avail}$ the connection is accepted, with appropriate resources being allocated along the path in accordance with the queue size and

WRR weighting parameters described in section 2.1. If $BW_{new}>BW_{avail}$, the current “empty queue (EQ)” measurement is evaluated at the WRR serving the path, providing a measure of the available, or “additional” bandwidth. If $BW_{new} < BW_{EQ}$ the connection is accepted, modifying the WRR schedule and allocating a channel queue as described in section 2.1, otherwise the connection is rejected (or re-negotiated, *prior to acceptance*, if the system is capable of such re-negotiation).

(2) For CBR user traffic

a) Determine the resources required for a newly requested connection in terms of cell bandwidth in cells per second, BW_{new}, using the PCR source parameter provided by the user.

b) Proceed as in steps (1b) and (1c) above.

(3) For ABR and UBR user traffic

ABR and UBR user traffic would be admitted into a separate source-to-destination path, not requiring feedback of any queue/server parameter for the implementation of CAC. Admission of ABR user traffic would be based on comparing the MCR signaled by the user through the UNI to the difference between the established ABR/UBR path capacity and the sum of the MCRs of the ABR user traffic occupying the path at the time of the admission request. UBR traffic would be admitted based on

> criteria established by individual network providers, since the characteristics of UBR traffic (UBR cells are always admitted at low priority, and may be discarded as necessary) make it unnecessary that it conform to CAC procedures as outlined above.

Having described the bandwidth-related CAC functions that must be performed at the source switch using the proposed model, the next section will describe an overall methodology for studying the proposed method of available bandwidth estimation and CAC.

2.5 Methodology for Studying WRR Measurement Based CAC

The issues being addressed here, and thus the associated goals of the research, are twofold. First, the QoS provided to all channels by the CAC method must be met, i.e., the cell loss ratio (CLR), maximum cell transfer delay (CTD), and cell delay variation (CDV) must at least meet the requirements of the traffic contract, with existing channels not being overly degraded by the admission of additional traffic to the WRR. Second, the traffic load carried in a established path must be maximized, or at least increased, over that which would be carried without using the proposed admission control method, thereby contributing to the overall efficiency of an integrated broadband network incorporating this method. Also, the technique for optimizing bandwidth utilization by the insertion of additional virtual channels to a path should be equal to or greater than the ability of other bandwidth optimization techniques.

To approach the first issue and meet the associated goal, the network models created by using the specified WRR output buffer queuing architectures to implement the network switches shown in Figure 2-10 will be used. Since all resources after the VC server are set to support the required VP, and are fixed in this role for long periods (see Sections 2.2 and 2.3.1), CAC analysis need only concern itself with the VC server and associated VP in the source switch. Analysis of the VC server model of Figure 2-11(a) will be accomplished, taking into account various combinations of CBR/VBR input traffic, having various peak and average rate parameters. For each combination of source traffic types, the probability of having an empty channel queue will be evaluated, and the overall distribution of empty timeslots under the associated WRR schedule can be determined. From these probabilities, the appropriate bandwidth for allowing additional traffic to be admitted to the path may be determined based on desired cell loss ratios. From the probabilities determined through analysis of the queues and WRR schedule, the CLR of cells in an end-to-end path using this model may be evaluated, both before and after the insertion of additional traffic channels. The CTD and CDV of all channels are readily evaluated based on numbers of traffic streams being served and traffic queue sizes, determined as described in [LIU97-2], and will not be addressed by this work. A detailed analysis focused on meeting required CLR is accomplished in Chapter 3.

Following analysis, a simulation model of the CBR/VBR VP in a source switch, as shown in Figure 2-11(a) will be created and, through appropriate simulations, executed using VBR traffic models. The servers and paths will be observed in simulated operation using various traffic parameters, while measuring CLR.

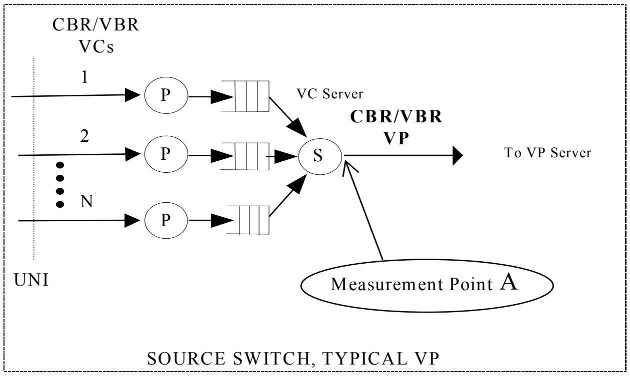

(a) Virtual Path with N Virtual Channels, Served by WRR

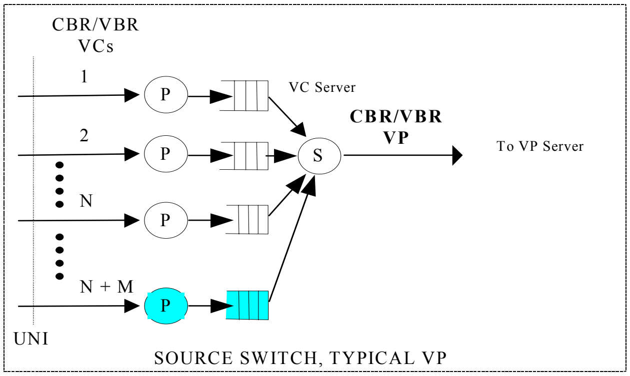

(b) Virtual Path with N + M Virtual Channels, served by WRR

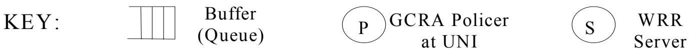

Figure 2-11: Network Model for Analysis and Evaluation Summary

To address the second issue and the associated goal, analysis of base system performance, using a CBR/VBR VP server in a source switch as shown in the path model, modified as shown in Figure 2-11(b), will be accomplished. The additional input at the source represents the additional traffic that could be admitted, based on a measurement of available unused bandwidth on the path using skipped queues as the criteria, as described in Section 2.1. Simulations incorporating the additional traffic will be created and run, to confirm the analysis.

As part of this research, the results of the analysis and simulation outlined above, as compared to results reported for other bandwidth evaluation and CAC schemes, will be reported. Chapters 3 and 4 will accomplish the analysis and describe the result of simulations required to support the research goals.

2.6 Summary

In Chapter 2, a new method for determining available bandwidth in a virtual path that is carrying existing traffic and is considered to be “full” by existing CAC methods is proposed and broadly outlined. Included are descriptions of:

1) The development of a new general model for the analysis of traffic serving mechanisms designed to allow the introduction of “additional” traffic into a virtual path that is considered to be “full” by CAC algorithms currently in use. The model is based on use of weighted round robin servers having queues sized in accordance with a described network framework, associated with a particular method of determining required queue sizes,

on dedicated end-to-end virtual paths. The proposed method is a means of guaranteeing required quality of service to admitted CBR and VBR traffic.

2) The use of the model and framework described above, where the probabilities of having an empty queue when serving “bursty” variable bit rate traffic is shown by simple graphic illustration to lead to an ability to insert “additional” traffic into a virtual path server schedule without affecting the existing traffic.
3) The need to obtain estimates of the number of empty timeslots per scheduled server cycle through an adequate measuring technique, in order to make use of the new method.
4) The choices made in picking implementation alternatives, assumptions, and switch architectures needed to implement the proposed new method.
5) An outline of a possible CAC technique based on the use of the proposed new method.
6) A broad methodology for evaluating the capability of the proposed new method.

The chosen method for determining available bandwidth, and an associated CAC technique, is further developed and analyzed in detail in Chapter 3, with results discussed in chapter 4.

Chapter 3

Analysis of WRR Queue Service Measurements

CHAPTER CONTENTS

LIST OF FIGURES

Chapter

3

Analysis of WRR Queue Service Measurements

This chapter provides an analysis of a WRR based system meeting the criteria outlined in Chapter 2. The system analyzed has an end-to-end framework as described in Section 2.4, and shown in Figure 2-9. Connection admission control is performed first by determining whether required delay parameters may be met, and then by determining whether required bandwidth is available. Delay parameters are a function of the established queue sizes, and are not analyzed here. Available bandwidth is initially determined based on the sum of the SCR of VBR traffic and sum of the PCR of CBR traffic in the virtual path. Once the virtual path is “full” based on this procedure, measurements of the number of empty timeslots in WRR cycles may then be used to determine the amount of “additional” bandwidth that is available, using the procedures derived from the analysis in the following sections. The results of the analysis show a relationship between the allowable cell rate of “additional” connections to the WRR server, and the mean number of available timeslots in the WRR cycle, as determined through measurement and statistical evaluation of idle, unassigned or low-priority cells in a virtual path cell stream.

3.1 Summary of the Connection Admission Control (CAC) Scenario

As described in the previous chapters, CAC (for VBR and CBR services) will be performed on VC's for which capacity is requested on a pre-established end-to-end VP, served by WRR queues. The VP has resources dedicated at the source node, each intermediate node, and the destination node, guaranteeing service to the properly admitted connections on the path. Required buffer sizes at source, destination and intermediate nodes may be calculated based on PCR, SCR (for VBR traffic), and MBS, provided by the user, as shown by Liu, et. al. [LIU97-2].

Focusing on the source node, where CAC takes place, the maximum queue length (MQL) for a buffer queue on a VC needed to guarantee zero cell loss (of conforming cells, after policing) as described in Section 2.3.2 (Equation 2.2) is:

$$MQL = 1 + \frac{\tau}{T} \tag{3.1}$$

This relates to MQL of 1 for CBR traffic and a MQL of $1+\tau_s/T_s$ for VBR traffic. The burst tolerance τ_s may be determined, as in Section 2.3.2 (Equation 2.1) from:

$$MBS = \left\lfloor 1 + \frac{\tau_s}{T_s - T} \right\rfloor \tag{3.2}$$

The source node WRR servers and served queues are arranged as in Figure 3-1.

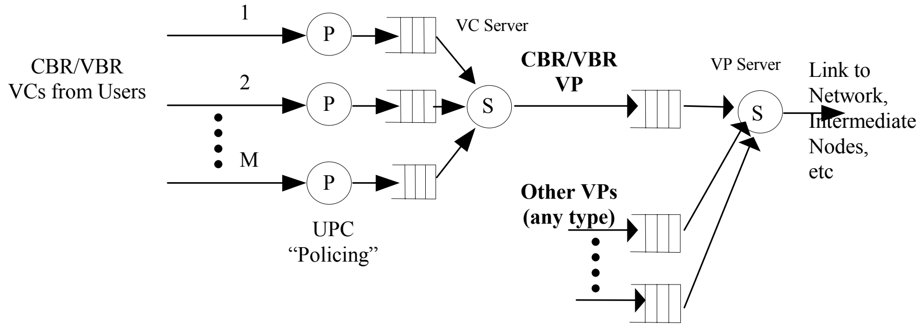

Figure 3-1: WRR Path Queuing Model for CBR/VBR Traffic

The required queue length and WRR schedule weighting for VBR VCs are dependent on the user provided parameters, PCR and MBS, which are relatively straightforward to determine, and SCR, which is not. SCR, for a given cell loss ratio (CLR), is dependent on queue size, which the user does not normally know directly, since it is in the domain of the service provider. However, if the user has knowledge of how the service provider will calculate required queue length based on PCR, MBS, and SCR, such as described by Equations (3.1) and (3.2), then the SCR may be calculated by the user. Connection admission methods depend on queue size and traffic characteristics such as mean burst size, peak rate, and utilization to determine the allowable cell rate of a traffic stream that will not cause queue overflows to exceed a required CLR. This calculated cell rate could be equated to the SCR of the traffic stream. Knowledge that the service provider network equipment will calculate queue size based on SCR and the equations for MQL, Equations (3.1) and (3.2), allows simultaneous solution of the allowable connection cell rate and MQL equations. The solution yields both the required SCR and queue size, if the allowable cell rate is set equal to the SCR. For purposes of this research, SCR will be equated to the specific equivalent capacity obtained by the method outlined in Section 1.2.3 and summarized in the following section.

3.2 An Example: Determination of SCR using Equivalent Capacity

The equivalent capacity of a two-state (on/off) source having Markov characteristics, which is flowing into a queue of finite capacity, is the queue service rate that provides a specified cell loss probability (i.e., buffer overflow probability), or CLR, of ε. Fluid flow approximations and other simplifications have been used by Guerin, *et. al.,* [GUE91] to develop equations for equivalent capacity of a single source and also multiple sources. For a single source having an average cell arrival rate of λ, each cell being served during period *T*, the equivalent capacity $\hat{c}$ is determined from:

$$\hat{c} = \frac{\alpha B(1-\rho)R_{peak} - L + \sqrt{\left[\alpha B(1-\rho)R_{peak} - L\right]^2 + 4L\alpha B\rho(1-\rho)R_{peak}}}{2\alpha B(1-\rho)} \tag{3.3}$$

where $L =$ Queue length

$\alpha = \ln\left(\frac{1}{\varepsilon}\right)$

$\varepsilon =$ The desired cell loss probability for which $\hat{c}$ is being calculated.

$B =$ Mean burst period of the source.

$\rho = \lambda T =$ The utilization of the source flowing into the queue.

R_{peak} = Peak rate of the source (PCR).

Since the WRR approach has one channel source per queue, the above equivalent capacity equation is set equal to SCR and solved for queue length, *L*, resulting in:

$$L = -SCR \cdot \log\left(\frac{1}{\varepsilon}\right) \cdot B \cdot \left(\frac{-PCR \cdot \rho + SCR \cdot \rho + PCR - SCR}{PCR \cdot \rho - SCR}\right) \tag{3.4}$$

The MQL Equations (3.1) and (3.2) may also be solved in terms of *PCR*, *SCR*, and *MBS*, yielding an equation for queue length, *L*, resulting in:

$$L = MQL = 1 + \left(\frac{(MBS-1) * (PCR - SCR)}{PCR}\right) \tag{3.5}$$

Equations (3.4) and (3.5) may be solved simultaneously for *L* and SCR, given ρ, ε, *B*, MBS, and PCR, thus giving a network user a means of calculating SCR without having prior knowledge of the source switch queue size, *L*.

The VC server queues, $m = 1,2,...,M$, of Figure 3-1 will be served by the WRR at a rate corresponding to the *PCR,* for CBR traffic, and *SCR*, for VBR traffic, to insure zero cell loss for conforming cells. The service period, *T*, of a given CBR queue will therefore be 1/*PCR*, and for a VBR queue will be *1/SCR*. The average cell arrival rate, λ_m, of traffic to a given queue, *m*, will be $\rho_m \cdot PCR_m$, where ρ_m is the utilization of the connection associated with the queue. The probability that queue *m* is empty, P_{Q0m}, may be then be obtained from $P_{Q0m} = 1 - \rho_m = 1 - \lambda_m T_m$, or:

$$P_{Q0m} = 1 - \lambda_m / SCR_m \tag{3.6}$$

Equation (3.6) has been shown to be valid [ROB94] for G/G/1 type queues, using standard queuing theory notation. For CBR traffic, $P_{Q0} = 0$, since λ, SCR and PCR are the same.

If SCR rates of connections entering the queues of the WRR server are such that a path is "full" based on the sum of the SCR rates, the probabilities that no queue is empty may be determined from the product of the $(1 - P_{Q0m})$ terms of the queues. The probability that at least one queue is empty may then be determined from (1-Probability that no queue is empty). Similarly, in a WRR schedule cycle, since timeslots are allocated to serving the queues in proportion to their traffic rates, the probability of no WRR timeslot being empty and at least one WRR timeslot being empty may be determined from the P_{Q0m} of each queue and the weight of its associated traffic in the cycle. Distributions of the probability of empty queues may be used to determine the number of WRR cycles needed to allow the service of an "additional" timeslot in the cycle, with "success" defined as finding at least one empty timeslot, corresponding to an empty queue, in a WRR schedule cycle.

3.3 Determination of Mean Empty Timeslots per WRR Cycle

The mean number of empty timeslots per cycle in a WRR of N timeslots may be calculated by taking the expected value of a random variable X, the number of empty timeslots per cycle,

$$E(X) = \sum_{x=1}^{N} x \cdot \Pr(X = x) \quad \text{empty slots per cycle} \tag{3.7}$$

The calculation of Pr(X=x) is computationally intensive, involving many permutations of the different probabilities of a certain number of empty and non-empty timeslots. However, recognizing the nature of the system, i.e., each timeslot in a WRR cycle is either empty or not empty, the mean number of empty timeslots per cycle may also be determined directly from P_{T0n}, the probability that the n^{th} timeslot is empty, as follows.

Let Y_n be a discrete random variable defined such that Y_n=1 when the n^{th} timeslot in a cycle is empty and Y_n=0 when the n^{th} timeslot is occupied by a cell requiring transmission. Since the objective is to use empty timeslots to carry additional traffic, "success" occurs whenever the server finds a queue empty, which results in an empty timeslot, and each instance of the WRR server checking a queue to determine whether a timeslot will be empty or occupied is a *Bernoulli trial* [HOG95]. The probability distribution function (also referred to as a probability mass function), of Y_n is thus

$$f_n(y) = \Pr(Y_n = y) = P_{T0n}^{y}(1 - P_{T0n})^{1-y}\,,\; y=0,1$$

where each Y_n has a Bernoulli distribution. If there is a one-to-one relationship, or mapping, between WRR queues m and the resulting virtual path timeslots, n, in a WRR cycle, with the traffic arriving into each queue assumed to be statistically independent, then the random variables Y_n will be statistically independent. Should the mapping be such that in a given cycle a queue, m, is visited more than once, then cells from one queue will occupy many timeslots in a WRR cycle. As a result, the random variables, Y_n representing the state of the timeslots, will not be independent, leading to correlation in the total number of empty timeslots observed at the output of the WRR from cycle to cycle. The effects of correlation in the number of empty timeslots per WRR cycle will be discussed in Section 3.5.2, and, for purpose of evaluation here, the Y_n will be treated as being statistically independent.

When Y_n is defined in this manner, a random variable X, the number of empty timeslots per cycle of N timeslots, is the sum of N random variables Y_n in a cycle,

$$X = \sum_{n=1}^{N} Y_n \tag{3.8}$$

For each Y_n the moment generating function (mgf), where M(t)=E[e^{ty}], is

$$M_{Y_n}(t) = (1 - P_{T0n}) + P_{T0n} e^t$$

If X is the sum of statistically independent Y_n, the mgf of X is the product of the individual terms, in this case, $M_X(t) = \prod_{n=1}^{N} [(1 - P_{T0n}) + P_{T0n} e^t]$, as shown in [HOG95].

This mgf may be rearranged to $M_X(t) = \prod_{n=1}^{N} [1 - P_{T0n}(1 - e^t)]$ and expanded to a series by multiplying the product terms, leading to

$$M_X(t) = 1 - \sum_{n=1}^{N} P_{T0n}(1 - e^t) + P_{T01} \sum_{n=2}^{N} P_{T0n}(1 - e^t)^2 + P_{T02} \sum_{n=3}^{N} P_{T0n}(1 - e^t)^2 + \mathrm{K} + \text{higher order terms}$$

The mean and variance of X may then be determined from the derivatives of the moments [HOG95], where $\mathrm{E}[X] = M'(0)$ and $\mathrm{VAR}(X) = M''(0) - [M'(0)]^2$. Thus,

$$M'_X(t) = \sum_{n=1}^{N} P_{T0n} e^t - 2P_{T01} \sum_{n=2}^{N} P_{T0n} (1 - e^t)(e^t) - 2P_{T02} \sum_{n=3}^{N} P_{T0n} (1 - e^t)(e^t) + \ldots + \text{higher order terms}$$

and

$$M''_X(t) = \sum_{n=1}^{N} P_{T0n} e^t - 2P_{T01} \sum_{n=2}^{N} P_{T0n} (e^t - 2e^{2t}) - 2P_{T02} \sum_{n=3}^{N} P_{T0n} (e^t - 2e^{2t}) + \ldots + \text{higher order terms}$$

Using the above moments,

$$\mu_X = E(X) = M'(0) = \sum_{n=1}^{N} P_{T0n} \quad \text{cells} \tag{3.9}$$

$$Var(X) = \sigma_X^2 = M''(0) - [M'(0)]^2 = \sum_{n=1}^{N} P_{T0n} (1 - P_{T0n}) \tag{3.10}$$

The distribution of X depends on Bernoulli distributions associated with the individual timeslots. If the probabilities of being empty are identical for each timeslot, $P_{T0n} = P$ for all n, $n = 1,2,\ldots,N$, then the mean and variance reduce to $\mu_X = NP$ and $Var(X) = NP(1-P)$, which are characteristics of a binomial distribution [HOG95]. In the scenario being examined here, P_{T0n} may not be the same for each n, so the distribution of X will be referred to as "binomial-like".

Other characteristics of this "binomial-like" distribution include a variance that, for a given mean, is at a maximum when all probabilities are equal, i.e., when the "binomial-like" distribution is actually a binomial distribution. This may be shown as follows.

If all probabilities of empty timeslots in the cycle are equal, $P_{T0n} = P$, then from Equations (3.9) and (3.10) the mean and variance of empty timeslots per WRR cycle are from the binomial distribution, where

$$\mu_{X_B} = \sum_{n=1}^{N} P_{T0n} = P_{T01} + P_{T02} + \Lambda + P_{T0N} = NP$$

and

$$\sigma_{X_B}^2 = \sum_{n=1}^{N} P_{T0n}(1 - P_{T0n}) = P_{T01}(1 - P_{T01}) + P_{T02}(1 - P_{T02}) + \Lambda + P_{T0N}(1 - P_{T0N}) = NP(1 - P)$$

The smallest change in probabilities that results in a "binomial-like" distribution having the same mean as a binomial distribution with mean μ_{X_B} occurs when one of the N equal timeslot probabilities is reduced by some infinitesimally small increment δ, while another is increased by the same amount, resulting in

$$\mu_{X_{BL}} = \sum_{n=1}^{N} P_{T0n} = (P_{T01} + \delta) + (P_{T02} - \delta) + \Lambda + P_{T0N} = NP + \delta - \delta = NP = \mu_{X_B}$$

The same incrementally different probabilities are used to calculate the variance of this "binomial-like" distribution, resulting in

$$\sigma^2_{X_{BL}} = (P_{T01}+\delta)(1-P_{T01}-\delta)+(P_{T02}-\delta)(1-P_{T02}+\delta)+\Lambda\ +P_{T0N}(1-P_{T0N}) = NP(1-P)-2\delta^2$$
$$= \sigma^2_{X_B} - 2\delta^2$$

From the above $\sigma^2_{X_{BL}} \le \sigma^2_{X_B}$

Thus, the variance of a binomial distribution having a given mean is the upper bound on the variance of a "binomial-like" distribution having the same mean.

The number of empty timeslots per cycle in WRR servers having the characteristics described in this work will have "binomial-like" distributions as described above. These distributions will be used in determining the available capacity in a virtual path considered as being "full" based on the SCR's of the traffic streams already being served, as will be seen in the next section.

3.4 Determination of Available Capacity in a "Full" Virtual Path

Assume that a virtual path is "full" based on the sum of the SCR of the incoming traffic streams. The WRR overall is being served at the virtual path rate , R_{VP}, with each of M queues being served at its SCR, as determined by the WRR schedule cycle.

Thus, $\sum_{m=1}^{M} SCR_m = R_{VP}$ cells/second.

In a WRR consisting of M traffic streams, the queues will be served in proportion to their SCR, based on the overall path service rate. The number of timeslots the traffic stream occupies in a WRR cycle of N timeslots will depend on its SCR in relation to the SCR other queues. In general, as illustrated in Section 2.1, there are *N* WRR virtual path timeslots serving *M* queues in accordance with a scheduling algorithm based on the rates of the incoming traffic to the queues. The probability that a given timeslot is empty is the probability that the queue being served by the scheduled WRR timeslot is empty. The probability P_{T0n} that a given timeslot *n* is empty is therefore the probability that the queue *m*, which is being served in accordance with the WRR schedule, is empty, where the probability that the queue is empty, P_{Q0m}, is determined to be $1-\lambda_m/SCR_m$, from Equation (3.6). Some proportion of the *N* timeslots of the WRR server cycle will serve each of the *M* queues. The proportion of the WRR cycle devoted to serving queue *m* is called the *weight* of the queue, w_m. The weight will be proportional to the ratio of the SCR of the queue to the overall cell transmission rate of the virtual path, R_{VP}, and represents the number of times a queue is served in a single cycle. The weight of a queue is thus

$$w_m = N\frac{SCR_m}{R_{VP}}, \quad m = 1,2,\ldots,M \tag{3.11}$$

The number of timeslots, *N*, in a cycle must be an integer, and is chosen such that the weight of each queue, w_m is also an integer.

Thus, for example, there could be five queues being served at individual queue service rates of 100, 100, 200, 200, and 400 cells per second, with R_{VP} being 1000 cells/second. This would require a minimum cycle of 10 timeslots, and a WRR service rate of 100 cycles per second, and the respective queue weights would be 1, 1, 2, 2, and 4. This schedule requires the first queue to be served once per cycle, the second queue once per cycle, the third queue twice per cycle, etc.

There are two possible methods for calculation the SCR of "additional" traffic that may be served based on the probabilities that queues allocated to existing connections are empty. One method takes into account the probability that at least one timeslot will be empty every c cycles of the WRR, and bases the allowable SCR of the traffic entering an "additional" queue on this probability. The second method takes into account the probabilities that one or more timeslots will be empty every c_1 cycles of the WRR, two or more timeslots will be empty every c_2 cycles, three or more timeslots will be empty every c_3 cycles, etc., where $c_1<c_2<c_3<\ldots<c_N$.. The single empty timeslot method is computationally less intensive than the multiple empty timeslot method, and will be more conservative in determining allowable "additional" traffic cell rates. However, the multiple empty timeslot method allows the admission of "additional" traffic with a larger SCR than the single empty timeslot method.

3.4.1 The Single Empty Timeslot Method

The first of the two methods considered for calculation of the cell rate of allowable "additional" traffic will be illustrated by example, followed by detailed analysis.

3.4.1.1 Example

As an example of the single empty timeslot method, consider Table 3-1, which illustrates a scenario where the path capacity is 1000 cells/ sec, and the VBR sources transmit at the peak path rate when on, and have the SCR values shown.

m	Type	PCR	SCR	ρ	B	MBS	L	P_{Q0m}	w_m
1	CBR	100	100	N/A	N/A	N/A	N/A	0	1
2	VBR-a	1000	100	.05	---	---	---	.5	1
3	VBR-b	1000	200	.05	---	---	---	.75	2
4	VBR-c	1000	200	.05	---	---	---	.75	2
5	VBR-d	1000	400	.05	---	---	---	.88	4

Table 3-1. Traffic Scenario Example

For a given connection, the mean burst size, B, max burst size, MBS, and queue size, L, may take many different values that can result in the given SCR at a given cell loss ratio and utilization, ρ, of the connection. The probability that the queue associated with connection m is empty, P_{Q0m}, is dependent on the SCR and utilization of the connection, not on B, L, or MBS of the connection. A WRR schedule cycle for this scenario requires N=10 timeslots, with the queue service weights, w_m, determined from Equation (3.11) being the number of times a given queue is served in a WRR cycle. Taking into account these SCR-based weights of the traffic streams, a schedule of visits to timeslots of 1-3-5-4-5-2-5-3-4-5 could be chosen. In this scenario, if X is the number of empty timeslots in a cycle, the probability that no timeslot is empty during a cycle is:

$$\Pr(X=0)=$$
$$(1-P_{Q01})(1-P_{Q03})(1-P_{Q05})(1-P_{Q04})(1-P_{Q05})(1-P_{Q02})(1-P_{Q05})(1-P_{Q03})(1-P_{Q04})(1-P_{Q05})$$

or,

$$\Pr(X=0)=(1-P_{Q01})^1(1-P_{Q02})^1(1-P_{Q03})^2(1-P_{Q04})^2(1-P_{Q05})^4 = 4.05x10^{-7}$$

The probability that some slot (at least one) is empty is 1-Pr($X=0$) =.999999595

Note that the exponent integers in the above expression are the queue weights, as defined by Equation (3.11), and that the probability that X=0 can be expressed as

$$\Pr(X=0)=(1-P_{Q01})^{w1}(1-P_{Q02})^{w2}(1-P_{Q03})^{w3}(1-P_{Q04})^{w4}(1-P_{Q05})^{w5}$$

or, since $P_{Q0m}=1-\lambda_m/SCR_m$,

$$\Pr(X=0)=\left(\frac{\lambda_1}{SCR_1}\right)^{w_1}\bullet\left(\frac{\lambda_2}{SCR_2}\right)^{w_2}\bullet\Lambda\bullet\left(\frac{\lambda_5}{SCR_5}\right)^{w_5}$$

In a general case, combining the above with Equation (3.11) results in

$$\Pr(X=0)=\prod_{m=1}^{M}\left[\left(\frac{\lambda_m}{SCR_m}\right)^{w_m}\right]=\prod_{m=1}^{M}\left[\left(\frac{\lambda_m}{SCR_m}\right)^{N\frac{SCR_m}{R_{VP}}}\right] \qquad (3.12)$$

Now if Z is a random variable corresponding to the number of successes (i.e., at least one empty timeslot) in c statistically independent trials (WRR schedule cycles), then

$$\Pr(Z \geq 1) = 1 - \Pr((Z = 0) = 1 - [\Pr(X = 0)]^c = 1 - (4.07x10^{-7})^c$$

The smallest number of c that yields $\Pr(Z \geq 1) > 1 - \varepsilon$ for $\varepsilon = 10^{-10}$ may be found from

$$1 - (4.07x10^{-7})^c > 1 - 10^{-10}$$
$$c = 1.6 \text{ cycles} = 16 \text{ timeslots}$$

For this example, a cycle is 10 timeslots, allowing an additional timeslot (serving an additional traffic queue) to be introduced to the schedule every 16 original timeslots. This results in the ability to add a traffic connection with a cell rate of 1/17 of 1000 cells/sec or 58 cells/sec. This represents, in this case, a 5.8% increase in traffic carried by the given 1000 cell/second virtual path, over that which could be carried using the sum of SCR as the CAC criteria.

The probability that no timeslot is empty to allow the service of the additional queue every 1.6 cycles is $(4.07x10^{-7})^{1.6} = 5.915\text{x}10^{-11}$. This is a worst case probability of cell loss, since the cells destined for the 17th timeslot are actually buffered in the queue of the "additional" connection. There would be cell loss only when the "additional" connection buffer was full and a new cell arrives. Thus, the calculation of the cell rate of an "additional" traffic stream, served whenever there is an empty slot in the original WRR schedule cycle, is a conservative estimate, since it provides for the required CLR without regard to the size of the queue associated with the “added” connection. The parameter ε, used in determining the number of cycles needed in assuring a certain probability that a timeslot will be available, and thus the corresponding probability that a timeslot will not be available, is therefore a bound on the probability of cell loss in the “additional” traffic stream, resulting in a $CLR \le \varepsilon$.

From Equation (3.9), the mean number of empty timeslots per cycle for the timeslots involved in the five-queue, ten-timeslot example scenario described in Table 3-1, based on the WRR queue service schedule of 1-3-5-4-5-2-5-3-4-5, is the sum of the timeslot probabilities, or,

$$\mu_X = 0 + .75 + .88 + .75 + .88 + .5 + .88 + .75 + .75 + .88 = 7.02 \text{ empty timeslots/cycle}$$

Consolidating the above in terms of the queue probabilities from which the timeslot probabilities are determined,

$$\mu_X = 1(0) + 1(.5) + 2(.75) + 2(.75) + 4(.88) = 7.02 \text{ empty timeslots/cycle}$$

where the integers preceding the probabilities above are the queue weights, w_m. The mean number of empty timeslots may thus be expressed in terms of the empty queue probabilities as

$$\mu_X = w_1 P_{Q01} + w_2 P_{Q02} + w_3 P_{Q03} + w_4 P_{Q04} + w_5 P_{Q05}$$

Note that in general, during a cycle of N timeslots, each queue m of M queues will be visited w_m times, where w_m is determined by Equation (3.11). Whenever a queue is visited, the probability that it is empty is P_{Q0m}. This allows the mean number of empty timeslots per cycle to be calculated from the probabilities of empty queues, using Equation (3.9), as shown in the above calculations for μ_X, and Equation (3.11),

$$\mu_X = \sum_{n=1}^{N} P_{T0n} = \sum_{m=1}^{M} w_m \cdot P_{Q0m} = \sum_{m=1}^{M} N \cdot \frac{SCR_m}{R_{VP}} \cdot P_{Q0m}$$

Since $P_{Q0m} = 1 - \lambda_m / SCR_m$, this results in

$$\mu_X = \frac{N}{R_{VP}} \cdot \sum_{m=1}^{M} \left(SCR_m - \lambda_m\right) \tag{3.13}$$

The variance of X for this scenario may also be determined, from Equations (3.10),

$$\sigma_X^2 = 0(1) + .75(.25) + .88(.12) + .75(25) + .88(.12) + .5(.5) + .88(.12) + .75(.25) + .75(.25) + .88(.12)$$

$$\sigma_X^2 = 1(0)(1) + 1(.5)(.5) + 2(.75)(.25) + 2(.75)(.25) + 4(.88)(.12)$$

where the integers preceding the probabilities above are the queue weights, w_m. The variance of empty timeslots per cycle may thus be expressed in terms of the empty queue probabilities as

$$\sigma_X^2 = w_1 P_{Q01}(1 - P_{Q01}) + w_2 P_{Q02}(1 - P_{Q02}) + w_3 P_{Q03}(1 - P_{Q03}) + w_4 P_{Q04}(1 - P_{Q04}) + w_5 P_{Q05}(1 - P_{Q05})$$

which leads to

$$\sigma_X^2 = \sum_{n=1}^{N} P_{T0n}(1 - P_{T0n}) = \sum_{m=1}^{M} w_m P_{Q0m}(1 - P_{Q0m}) = \frac{N}{R_{VP}} \cdot \sum_{m=1}^{M} \left(SCR_m - \lambda_m\right)\left(\frac{\lambda_m}{SCR_m}\right) \tag{3.14}$$

Thus, an accurate estimate of the mean number of empty timeslots per cycle (7.02 timeslots in this case), made at the output of the WRR, would result in the knowledge that an additional connection of up to 58 cells/second could be admitted, without affecting the quality of the existing connections, and assuring the QoS of the new connection (CLR$\leq 10^{-10}$ in this case). In the next section a detailed analysis will be accomplished, using the concepts laid out in this example.

3.4.1.2 Detailed Analysis of Single Empty Timeslot Method

The probability that a queue is not empty is $1\text{-}P_{Q0m}$, or λ_m/SCR_m. The probability that a given timeslot is not empty depends on the probability that the queue it serves is not empty, as discussed in Section 3.3. The probability that no timeslot is empty, Pr($X=0$), in a WRR cycle of N timeslots that is serving M queues, as determined by Equation (3.12) is

$$\Pr(X=0)=\prod_{m=1}^{M}\left[\left(\frac{\lambda_m}{SCR_m}\right)^{N\frac{SCR_m}{R_{VP}}}\right]$$

The probability that a slot is empty, Pr($X\geq 1$) in a WRR cycle of N timeslots that is serving M queues is then 1-Pr$(X=0)$, or

$$\Pr(X\geq 1)=1-\prod_{m=1}^{M}\left[\left(\frac{\lambda_m}{SCR_m}\right)^{N\frac{SCR_m}{VPSR}}\right]$$

Now, if Z is a random variable representing the number of “successes”, where success is defined as having at least one empty time-slot in c statistically independent trials (a trial corresponds to a WRR cycle c, consisting of N timeslots), then

$$\Pr(Z \geq 1) = 1 - \Pr(Z = 0) = 1 - [\Pr(X = 0)]^c$$

The smallest *c* that yields $\Pr(Z \geq 1) > 1 - \varepsilon$, where the desired $CLR \leq \varepsilon$, may be found by

$$1 - [\Pr(X = 0)]^c \geq 1 - \varepsilon$$

$$[\Pr(X = 0)]^c \leq \varepsilon$$

$$\left\{ \prod_{m=1}^{M} \left[\left(\frac{\lambda_m}{SCR_m} \right)^{N \frac{SCR_m}{R_{VP}}} \right] \right\}^c \leq \varepsilon$$

where *M* is the number of queues, *N* is the number of timeslots in the WRR cycle, *c* is the number of cycles needed to achieve the desired bound on the CLR, λ_m is the average arrival rate and SCR_m is the sustainable cell rate, in cells/second, of the connection entering queue *m*, and R_{VP} is the virtual path service rate in cells/second.

From the above, $$c \cdot \ln \left\{ \prod_{m=1}^{M} \left[\left(\frac{\lambda_m}{SCR_m} \right)^{N \cdot \frac{SCR_m}{R_{VP}}} \right] \right\} \geq \ln(\varepsilon)$$

$$c \left\{ \sum_{m=1}^{M} N \cdot \frac{SCR_m}{R_{VP}} \cdot \ln \left(\frac{\lambda_m}{SCR_m} \right) \right\} \geq \ln(\varepsilon)$$

$$c \cdot \frac{N}{R_{VP}} \left\{ \sum_{m=1}^{M} SCR_m \left[\ln \left(\frac{\lambda_m}{SCR_m} \right) \right] \right\} \geq \ln(\varepsilon)$$

$$c \geq \frac{R_{VP}}{N} \cdot \ln(\varepsilon) \cdot \frac{1}{\sum_{m=1}^{M} SCR_m \left[\ln\left(\frac{\lambda_m}{SCR_m} \right) \right]}$$

which results in a bound on *c* of

$$\tilde{c} = \left\lceil \frac{R_{VP}}{N} \cdot \ln(\varepsilon) \cdot \frac{1}{\sum_{m=1}^{M} SCR_m \left[\ln\left(\frac{\lambda_m}{SCR_m} \right) \right]} \right\rceil \tag{3.15}$$

Based on this result, one "additional" cell may be sent for each $\tilde{c}$ cycles (i.e., every $\tilde{c} \cdot N$ timeslots), with the resulting additional traffic having a CLR of no more than ε. Now, the SCR of an "additional" connection served by the WRR when a regularly scheduled timeslot in a WRR cycle is empty is

$$R_{ASCR} = \frac{1}{(\tilde{c} \cdot N) + 1} \cdot R_{VP} \quad \text{cells/second,}$$

Substituting $\tilde{c}$ from Equation (3.15) yields

$$R_{ASCR} = \frac{R_{VP}}{\left\lceil R_{VP} \cdot \ln(\varepsilon) \cdot \frac{1}{\sum_{m=1}^{M} SCR_m \left[\ln\left(\frac{\lambda_m}{SCR_m} \right) \right]} \right\rceil + 1} \quad \text{cells/sec.} \tag{3.16}$$

The "additional" connection cell rate, R_{ASCR}, is thus a function of the virtual path cell rate, R_{vp}, the bound on the desired CLR, ε, and both the mean cell arrival rate, λ, and SCR of the existing connections. Since R_{vp}, ε, and, for a given connection *m,* SCR_m are known at the source switch, but λ_m is not, a means of evaluating the "additional" connection cell rate based on measurable parameters must be determined. This is addressed in the next section.

3.4.1.3 Determining Allowable "Added" Connection Rate

In order to determine an allowable cell rate for a connection that could be added to a "full" WRR-served virtual path, a means of evaluating Equation (3.16) must be found, in spite of the lack of knowledge regarding the mean cell arrival rates, λ_m, of the existing connections. This section accomplishes an analysis that determines a bound on the allowable cell rate of an "additional" connection.

From Equation (3.13), the mean number of empty slots per cycle is

$$\mu_X = \frac{N}{R_{VP}} \cdot \sum_{m=1}^{M} \left(SCR_m - \lambda_m \right).$$

Therefore

$$\sum_{m=1}^{M} \left(SCR_m - \lambda_m \right) = \frac{R_{VP}}{N} \cdot \mu_X \qquad (3.17)$$

Now recalling the previous expression for the “additional” connection SCR, Equation (3.16),

$$R_{ASCR} = \frac{R_{VP}}{\left\lceil R_{VP} \cdot \ln(\varepsilon) \cdot \frac{1}{\sum_{m=1}^{M} SCR_m \left[\ln\left(\frac{\lambda_m}{SCR_m} \right) \right]} \right\rceil + 1} \text{ cells/second}$$

The summation terms from Equations (3.16) and (3.17) respectively are

$$f_1 = \sum_{m=1}^{M} \left(SCR_m - \lambda_m \right) \qquad \text{and} \qquad f_2 = -\sum_{m=1}^{M} SCR_m \left[\ln\left(\frac{\lambda_m}{SCR_m} \right) \right]$$

Re-arranging f_1 and letting $\rho_m = \frac{\lambda_m}{SCR_m}$, $0 \le \rho_m \le 1$, the summation terms become

$$f_1 = \sum_{m=1}^{M} SCR_m (1 - \rho_m) \qquad \text{and} \qquad f_2 = -\sum_{m=1}^{M} SCR_m [\ln(\rho_m)]$$

Expanding the natural logarithm in f_2,

$$\ln(\rho_m) = 2\left[\left(\frac{\rho_m - 1}{\rho_m + 1} \right) + \frac{1}{3}\left(\frac{\rho_m - 1}{\rho_m + 1} \right)^3 + \frac{1}{5}\left(\frac{\rho_m - 1}{\rho_m + 1} \right)^5 + \Lambda \right]$$

Since $(\rho_m - 1) \le 0$

$$\ln(\rho_m) \le 2\left(\frac{\rho_m - 1}{\rho_m + 1} \right)$$

$$\ln(\rho_m) \le (\rho_m - 1)\left(\frac{2}{\rho_m + 1}\right), \text{ but } \frac{2}{1+\rho_m} \ge 1$$

$$\ln(\rho_m) \le (\rho_m - 1)$$

$$-\ln(\rho_\mathrm{m}) \ge (1 - \rho_m)$$

Substituting this result into the expression for f_2,

$$-\sum_{m=1}^{M} SCR_m[\ln(\rho_m)] \ge \sum_{m=1}^{M} SCR_m[1 - \rho_m]$$

$$-\sum_{m=1}^{M} SCR_m\left[\ln\left(\frac{\lambda_m}{SCR_m}\right)\right] \ge \sum_{m=1}^{M} SCR_m\left[1 - \frac{\lambda_m}{SCR_m}\right]$$

Therefore

$$-\sum_{m=1}^{M} SCR_m\left[\ln\left(\frac{\lambda_m}{SCR_m}\right)\right] \ge \sum_{m=1}^{M}(SCR_m - \lambda_m)$$

Substituting from Equation (3.17) and multiplying by –1,

$$\sum_{m=1}^{M} SCR_m\left[\ln\left(\frac{\lambda_m}{SCR_m}\right)\right] \le -\frac{R_{VP}}{N} \cdot \mu_X$$

Substituting the above into Equation (3.16),

$$R_{ASCR} \geq \frac{R_{VP}}{R_{VP} \cdot \ln(\varepsilon) \cdot \frac{1}{-\frac{R_{VP}}{N} \cdot \mu_X} + 1}$$

resulting in $$R_{ASCR} \geq \frac{R_{VP} \cdot \mu_X}{-N \cdot \ln(\varepsilon) + \mu_X}$$

Thus, an estimate of the allowable connection rate in terms of the SCR of the "additional" connection using this method has a bound of

$$\hat{R}_{ASCR} = \frac{R_{VP} \cdot \mu_X}{N \cdot \ln\left(\frac{1}{\varepsilon}\right) + \mu_X} \tag{3.18}$$

This value is the SCR of an "added" connection that can safely be admitted (at a CLR $\leq \varepsilon$) to a "full" VP served by a WRR having N timeslots per cycle, where the mean number of empty timeslots per cycle, μ_x is an unknown, but R_{vp}, N, and ε are known from system or user provided parameters. This result may also be expressed in terms of a normalized mean number of empty slots per N-timeslot cycle, $\mu_{XN} = \mu_X / N$, resulting in

$$\hat{R}_{ASCR} = \frac{R_{VP} \cdot \mu_{XN}}{\ln\left(\frac{1}{\varepsilon}\right) + \mu_{XN}} \text{ cells/sec}, \ 0 \leq \mu_{\mathrm{XN}} \leq 1 \tag{3.19}$$

The mean number of empty timeslots per cycle, μ_X (or μ_{XN}), in an actual system may be estimated by measuring the number of empty slots per cycle over multiple

WRR cycles, thus providing the means of determining an allowable rate for admitting additional traffic.

The result shown by Equation (3.19) is independent of any particular means of calculating SCR, or particular type of traffic, as long as the traffic may be characterized by the parameters PCR, SCR, and MBS, and is policed by UPC using the GCRA, as described in Section 1.3.4, to insure that the user declared parameters are not exceeded. However, the type of traffic (characterized by utilization, mean burst size, max burst size, etc.) will effect the accuracy of the estimate of the mean number of empty slots per cycle when it is measured, by affecting the number of cycles over which measurement must take place to insure the confidence interval of the estimate.

3.4.2 The Multiple Empty Timeslot Method

The single empty timeslot method described in the previous section is computationally efficient since it requires knowledge only of the $\Pr(Z \geq 1)$, where Z was a random variable representing the number of successes(empty timeslots) in c trials (WRR cycles), which was easily calculated from $\Pr(X=0)$.

In dealing with multiple empty timeslots, the probabilities $\Pr(Z \geq 1)$, $\Pr(Z \geq 2)$,…, $\Pr(Z \geq N)$ must be considered to determine the number of cycles c_1, c_2,…, c_N respectively needed to insure the availability of one, two,…, N timeslots in a WRR cycle with probability sufficient to guarantee the CLR. As was the case in the single empty timeslot method, these probabilities are determined from

$$\Pr\left(Z \geq 1\right) = 1 - \Pr(Z = 0) = 1 - \left[\Pr(X = 0)\right]^{c_1}$$

$$\Pr\left(Z \geq 2\right) = 1 - \Pr(Z \leq 1) = 1 - \left[\Pr(X \leq 1)\right]^{c_2}$$

$$\vdots$$

$$\Pr\left(Z \geq i\right) = 1 - \Pr(Z \leq i-1) = 1 - \left[\Pr(X \leq i-1)\right]^{c_i}, \ 1 \leq i \leq N \tag{3.20}$$

The probability distribution of *Z* depends in turn on the probability distribution of *X*, which may be determined from the binomial-like distribution of *X* described in Section 3.3, having a mean and variance described by Equations (3.9) and (3.10). The distribution of *X* may be approximated, if it is binomial, and symmetrical (i.e.. the probability of success is .5), for values of *N* as small as ten, using a normal pdf [HOG95], as suggested by Figure 3-2. For skewed distributions, where the probability of success is closer to zero or one, larger values of *N* are required in order for the normal distribution to be a close approximation to the binomial distribution.

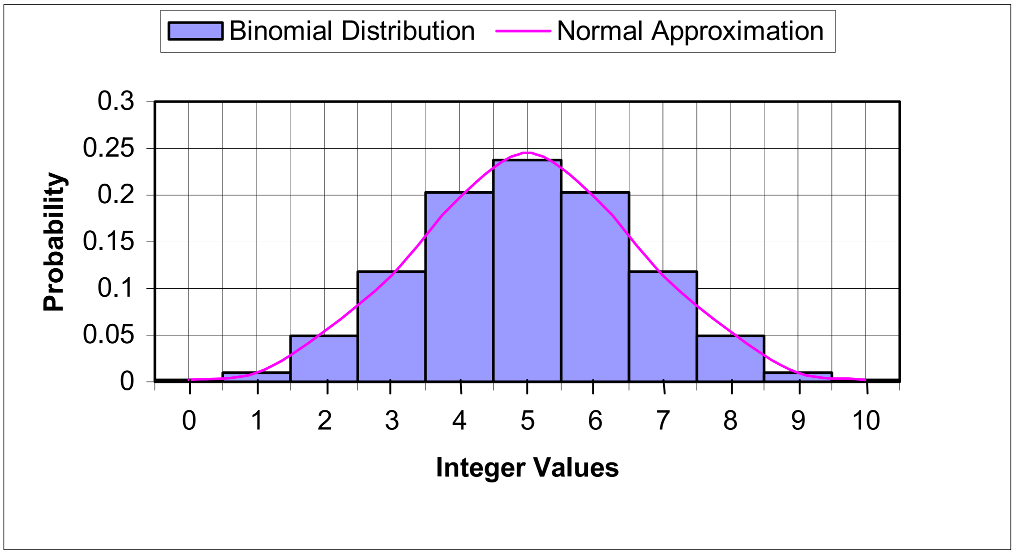

Figure 3-2: Normal Approximation to a Binomial Distribution, N=10, p=.5

Note that in Figure 3-2 the heights of the rectangles represent the probabilities of the integers i=0, 1, 2, …, 10, and that the area of the rectangle with a base between $(i-.5)$ and $(i+.5)$ is approximately the same as the area under the normal pdf between $(i-.5)$ and $(i+.5)$. Also, the area of the set of rectangles occupying the region $(0-.5)$ to $(i+.5)$ is approximately the same as the area under the normal pdf in the region from $(-\infty)$ to $(i+.5)$. Thus, the probability that $X \le i$ in a binomial distribution of any size N may be approximated using the probability that $X \le (i+.5)$ in a normal distribution having the same mean and variance as the binomial distribution.

The mean and variance of X, the number of empty timeslots in a WRR cycle, where X has a binomial-like distribution as shown in Section 3.3, will therefore be used in determining the probabilities of X using a normal approximation. The mean and variance of the number of empty timeslots per cycle, X, were determined from Equations (3.9) and (3.10) to be

$$E(X) = \sum_{n=1}^{N} P_{0n} = \mu_X \qquad \text{and} \qquad Var(X) = \sum_{n=1}^{N} [P_{0n}(1-P_{0n})] = \sigma_X^2$$

The probabilities that X is greater than x, x=0,1,…,N may be estimated by taking the normal probabilities at a point half of the distance between the x being considered and the next x, as indicated in the discussion of Figure 3-2. For x=0 this results in

$$\Pr(X \geq 0) = 1 - \Pr(X \leq 0) \approx 1 - \Pr\left(\Psi \leq \frac{0 + .5 - \mu_X}{\sigma_X}\right),$$

where Ψ has standard normal distribution, i.e., Ψ is N(0,1).

In general,

$$\Pr(X \geq x) \approx 1 - \Pr\left(\Psi \leq \frac{x + .5 - \mu_X}{\sigma_X}\right), \; x = 0,1,2,\ldots,N$$

Letting $\psi_x = \dfrac{x + .5 - \mu_X}{\sigma_X}, x = 0,1,2,\ldots,N$, then $\Pr(\Psi \leq \psi_x)$ is determined by the normal distribution, and

$$\Pr(\Psi \leq \psi_x) = \Phi(\psi_x) = \int_{-\infty}^{\psi_x} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx$$

$$\Pr(X \geq x) \approx 1 - \Phi(\psi_x)$$

Note that for the normal distribution, $\Phi(-\psi_x) = 1 - \Phi(\psi_x)$.

Accomplishing a similar calculation to that described in Section 3.4.1 for the c values associated with each Pr($X \geq x$), i.e., $c_1, c_2, \ldots, c_{x+1}$, using Equation (3.20), the number of cycles needed to insure the CLR for one empty slot, two empty slots, etc.,

$$[1 - \Pr(X \geq x)]^{c_{x+1}} \leq \varepsilon, \quad x = 0,1,2,\ldots,N$$

$$[\Phi(\psi_x)]^{c_{x+1}} \leq \varepsilon$$

$$c_{x+1} \ln[\Phi(\psi_x)] \geq \ln \varepsilon$$

$$c_{x+1} \geq \frac{\ln \varepsilon}{\ln[\Phi(\psi_x)]}, x = 0,1,\ldots,N, \psi_x = \frac{x + .5 - \mu_X}{\sigma_X}$$

To simplify notation, let $i = x + 1$, then

$$c_i \geq \frac{\ln \varepsilon}{\ln[\Phi(\psi_i)]}, i = 1,2,\ldots,N+1, \text{and}\, \psi_i = \frac{i - .5 - \mu_X}{\sigma_X},$$

The bound on c_i is therefore

$$\widetilde{c}_i = \frac{\ln \varepsilon}{\ln[\Phi(\psi_i)]}, i = 1,2,\ldots,N+1, \psi_i = \frac{i - .5 - \mu_x}{\sigma_x} \tag{3.21}$$

where $\widetilde{c}_i$ is the number of cycles that must take place in order to have a probability of 1-ε that there are *i* timeslots available to serve the "additional" connection.

Thus, in a manner similar to that used by the single empty timeslot method, every $\widetilde{c}_i$ cycles *i* timeslots will be empty and available to serve "additional" traffic with probability 1-ε, and the incremental portion of the overall SCR of the "additional" connection will be

$$R_{ASCR_i} = \frac{1}{(\widetilde{c}_i \cdot N) + 1} \cdot R_{VP}, i = 1,2, N+1$$

for each $\widetilde{c}_i$, where *N* is the number of timeslots in the WRR cycle. The total supportable rate of the "additional" connection will be the sum of the incremental rates, where

$$R_{ASCR} = R_{VP} \cdot \sum_{i=1}^{N} \frac{1}{(\widetilde{c}_i \cdot N) + 1} \quad \text{cells/second}$$

Now, substituting the value for $\tilde{c}_i$ from Equation (3.21),

$$R_{ASCR} = R_{VP} \cdot \sum_{i=1}^{N} \left\{ \frac{1}{\left\lceil \frac{\ln \varepsilon}{\ln[\Phi(\psi_i)]} \cdot N \right\rceil + 1} \right\} \text{ cells/second} \tag{3.22}$$

where $\psi_i = \dfrac{i - .5 - \mu_X}{\sigma_X}$

In equation (3.22), ψ_i may be expressed in terms of normalized components, such that

$$\psi_i = \frac{\frac{i}{N} - \frac{.5}{N} - \mu_{XN}}{\sigma_{XN}} \tag{3.23}$$

where μ_{XN} and σ_{XN} are the mean and standard deviation of *X* normalized to *N* timeslots, i.e., divided by *N*.

For a random variable *X* having a binomial-like distribution, the variance of *X* will be less than the variance of a random variable with binomial distribution (i.e., where all of the N probabilities of success are equal) as shown in Section 3.3. In other words, the variance of a binomial-like distribution will be bounded by the variance of a binomial distribution having the same mean as the binomial-like distribution.

Thus, if a random variable *X,* having binomial-like distribution, has a mean, μ_X, such that the mean is the same as a binomial distribution having mean NP and variance $NP(1-P)$, P will then be

$$P = \frac{\mu_X}{N}$$

and the variance of the binomial-like distribution will be such that

$$\sigma_X^2 \le NP(1-P)$$

The normalized mean and variance of a binomial-like distribution are then such that

$$\mu_{XN} = \frac{\mu_X}{N} = P \quad \text{and} \quad \sigma_{\mathrm{XN}}^2 = \frac{\sigma_X^2}{N} \le P(1-P)$$

Therefore,

$$\sigma_{XN}^2 \le \frac{\mu_X}{N}\left(1-\frac{\mu_X}{N}\right)$$

The normalized standard deviation of a binomial-like distribution is then such that

$$\sigma_{XN} \le \sqrt{\frac{\mu_X}{N}\left(1-\frac{\mu_X}{N}\right)}$$

which, when substituted into Equation (3.23), leads to

$$\psi_i \ge \frac{\frac{i}{N}-\frac{.5}{N}-\frac{\mu_X}{N}}{\sqrt{\frac{\mu_X}{N}\left(1-\frac{\mu_X}{N}\right)}}$$

This in turn leads to a bound on ψ_i,

$$\hat{\psi}_i = \frac{\frac{i}{N}-\frac{.5}{N}-\frac{\mu_X}{N}}{\sqrt{\frac{\mu_X}{N}\left(1-\frac{\mu_X}{N}\right)}}$$

This then leads to a bound on R_{ASCR}, from Equation (3.22), of

$$\hat{R}_{ASCR} = R_{VP} \cdot \sum_{i=1}^{N} \left\{ \frac{1}{\left\lceil \frac{\ln \varepsilon}{\ln[\Phi(\hat{\psi}_i)]} \cdot N \right\rceil + 1} \right\} \text{ cells/second} \tag{3.24}$$

where $\hat{\psi}_i = \dfrac{\frac{i}{N} - \frac{.5}{N} - \frac{\mu_X}{N}}{\sqrt{\frac{\mu_X}{N}\left(1 - \frac{\mu_X}{N}\right)}}$

The result shown in Equation (3.24) was derived from a normal approximation to a binomial distribution. This approximation is only completely accurate for infinitely large values of *N*, since the binomial distribution approaches the normal distribution as *N* increases without limit [RIC68]. For smaller values of *N*, the probability in the negative "tail" of the normal distribution will make the estimates of the probability of the smaller values of *X* slightly larger than the actual value. This over-estimate of the probabilities in turn causes a larger estimate of the value of c_I than would be true for larger values of *N*. This in turn causes a smaller estimate of $\hat{R}_{ASCR}$ for smaller values of *N*. Thus, a conservative estimate of $\hat{R}_{ASCR}$ is created at the smallest acceptable value of *N*. This relationship can be seen in Table 3-2, which shows the results of calculating $\hat{R}_{ASCR}$ as a percentage of total virtual path rate, using Equation (3.24), for *N* between 10 and 3000, and $\varepsilon = 10^{-3}$, where the row corresponding to N=10 contains the most conservative values of $\hat{R}_{ASCR}$.

	Mean Empty Timeslots per WRR Cycle								
N	.1	.2	.3	.4	.5	.6	.7	.8	.9
3000	3.29	5.38	7.33	9.33	11.56	14.28	18.04	24.36	40.19
2000	3.29	5.38	7.33	9.33	11.56	14.28	18.04	24.36	40.19
1000	3.29	5.38	7.33	9.32	11.55	14.27	18.04	24.35	40.17
200	3.28	5.38	7.33	9.32	11.54	14.26	18.02	24.30	40.04
100	3.28	5.37	7.32	9.31	11.53	14.24	17.97	24.24	39.89
20	3.26	5.34	7.27	9.24	11.41	14.09	17.72	23.81	38.71
10	**3.23**	**5.28**	**7.20**	**9.11**	**11.32**	**13.86**	**17.38**	**23.33**	**38.10**

Table 3-2. Additional SCR Percent, $\hat{R}_{ASCR} / R_{VP}$, N=10 to 3000

As was discussed earlier in this section, N=10 will be used as the smallest acceptable number of timeslots to be used in developing this method.

The estimate for $\hat{R}_{ASCR}$ resulting from Equation (3.24) will be inaccurate when the normalized mean number of empty slots per cycle is close to one, due to the nature of the normalized variance comprising the denominator of $\hat{\psi}_i$. When $\mu_X / N \approx 1$ a very small variance results, which in turn causes very small probabilities to be associated with the small-valued *i* terms in Equation(3.24). The small-valued *i* terms contribute the most to the total $\hat{R}_{ASCR}$, therefore small inaccuracies in the very small probabilities resulting from the normal approximation to the binomial distribution may lead to large inaccuracies in $\hat{R}_{ASCR}$. For this reason, $\hat{R}_{ASCR}$ for μ_{XN} close to 1 were evaluated, using Equation(3.24). A qualitative estimate in this region led to the conclusion that the $\hat{R}_{ASCR}$ for all $\mu_{XN} > .95$ will be considered to be equal to the $\hat{R}_{ASCR}$ resulting from $\mu_{XN} = .95$, as a conservative estimate.

Equation (3.24) describes a multiple timeslot relationship, subject to the limits discussed above, based on normal probabilities, between μ_X, the mean number of empty slots per cycle, and $\hat{R}_{ASCR}$, the approximate SCR of an "additional" connection that can be admitted to the network and carried at the required CLR or better. As was the case in the single empty timeslot method described in Section 3.4.1, $\hat{R}_{ASCR}$ is considered to be an "additional" connection since it is admitted even after the virtual path is considered "full" based on the sum of the SCRs of the other traffic in the path. As was also the case in the single empty timeslot method, μ_X is the only unknown, and it may be estimated through measurements of empty timeslots as the WRR cycles.

3.4.3 Determination of Added Queue Size

The required maximum queue length, *MQL*, of the queues serving "normal" traffic streams accepted prior to the virtual path being "full" is determined as discussed in Sections 3.1 and 3.2 "Additional" traffic stream cell rates are determined through the procedures described in Section 3.4. The "additional" traffic is served whenever timeslots are available in the "normal" queue schedule, rather than being served on a regular basis, as is the case for the "normal" traffic streams. Thus, a queue serving an "additional" traffic stream will have a buffer size greater than the MQL determined for the "normal" queues serving traffic streams having the same cell rate.

Over “long” periods, the empty timeslot probabilities resulting from the characteristics of the “normal’ traffic streams is sufficient to ensure the overall availability of virtual path timeslots needed to support the “additional” traffic cell rate, provided that cells exceeding the required CLR have not been lost due to the “additional” traffic queue size being too small. Over “short” periods, however, there is some probability that the “additional” traffic may be served for a time at a rate less than that required to maintain the “additional” cell rate. Queues of “additional” traffic streams must be at least *MQL* in length even when being served regularly on a schedule corresponding to the “additional” traffic rate. Since “additional” traffic is not served on a regular basis, the length of queues in “additional” traffic streams admitted using methods described in Section 3.4 must be greater than the *MQL* of traffic streams of the same cell rate served by a regularly scheduled WRR cycle.

An “added” queue is served only when there are empty timeslots in the WRR cycle, and thus there is some probability of the “added” queue being served later than required by its SCR. This probability contributes to the probability of queue overflow if the queue is sized only to handle the MBS of the “added” traffic stream, using Equation 3.5. The availability of empty timeslots during worst case periods (periods where no cells from the “additional” traffic stream are being served by the WRR for a significant time) is used in determining the queue size needed to support the “additional” connection at the required QoS.

Traffic that is policed using the GCRA algorithm, as discussed in Chapter 1 and [TMS4], will comply with GCRA(T,0) and GCRA(T_S, τ_S) for peak cell rate, *PCR*, and sustainable cell rate, *SCR*, respectively. A worst case scenario occurs when a source exclusively transmits maximum size bursts of MBS at the maximum periodic rate allowed by the GCRA policing algorithm. When this occurs, WRR cell slots associated with that traffic stream are unavailable for the longest possible time, before becoming available for the remainder of the overall period. A periodic cell stream with period $MBS \cdot T_S$ which transmits a burst of *MBS* cells at the peak rate, with inter burst spacing of $T_1 = MBS \cdot (T_S - T) + T$ has *PCR*=1/*T*, sustainable cell rate *SCR*=1/T_S, and is compliant with the GCRA for both *PCR* and *SCR* [TMS4]. The relationship for this transmission pattern is shown in Figure 3-3.

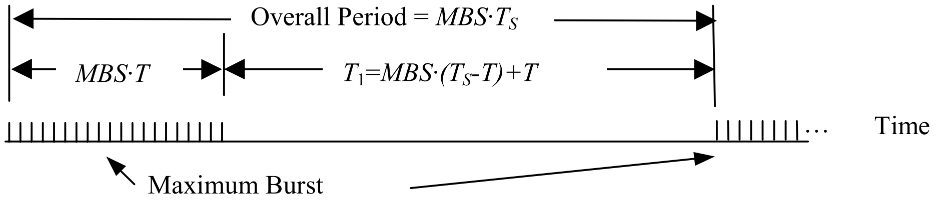

Figure 3-3: Worst Case Maximum Burst Relationships for GCRA Policed Traffic

In a WRR having *M* queues, the worst case for the overall WRR corresponds to all *M* traffic sources transmitting at the worst case cell rate described above, with the bursts exactly overlapping where all sources are transmitting at the same rate with identical MBS, as shown in Figure 3-4.

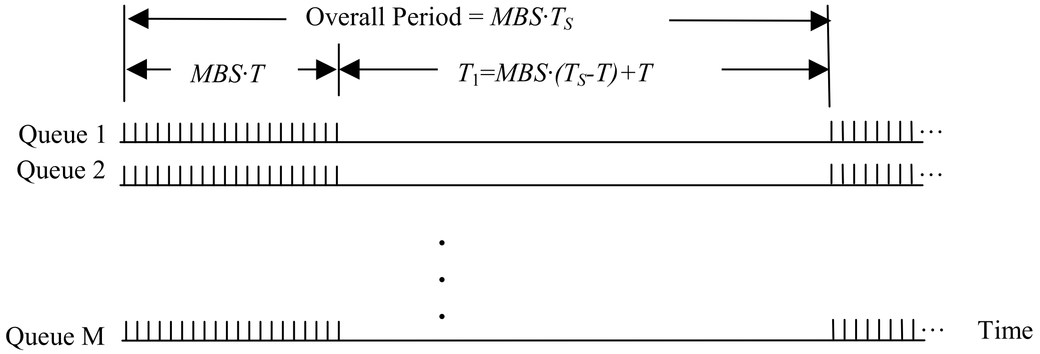

Figure 3-4: Worst Case Scenario for M Queues of GCRA Policed Traffic

The above worst case scenario results in each of N timeslots in a WRR cycle being occupied by "normal" traffic for MBS cycles, leaving no timeslots available to serve the "additional" traffic stream. Considering an N timeslot WRR to consist of N queue equivalents each having SCR of

$$SCR_E = \frac{1}{N}\sum_{i=1}^{M} SCR_i = \frac{R_{VP}}{N}$$

and each queue equivalent having a *MBS* of

$$MBS_E = \frac{1}{N}\sum_{j=1}^{M} MBS_j$$

will result in a worst case scenario for the given traffic pattern.

The "additional" traffic stream will be served at a rate of N cells per WRR cycle during the inter-burst period, T_1, and will not be served at all during the burst period,

$MBS_E \cdot T$. For an "additional" traffic stream having a cell rate of SCR_A, the queue length needed in addition to the MQL of Equation 3.5 may be found as follows.

The number of "additional" traffic cells arriving during the overall period $MBS \cdot T_S = MBS / SCR$ of the existing traffic using the worst case equivalents discussed above will be

$$N_A = \frac{MBS_E}{SCR_E} \cdot SCR_A \quad \text{cells}$$

The number of "additional" traffic stream cells not served during the burst period $MBS \cdot T = MBS / PCR$ of the existing traffic is

$$N_{NS} = \frac{MBS_E}{PCR_E} \cdot SCR_A \quad \text{cells}$$

Based on the desired CLR, the number of "additional" traffic stream cells which can be lost, while still maintaining the required QoS, is

$$N_L = \varepsilon \cdot N_A = \varepsilon \cdot \frac{MBS_E}{SCR_E} \cdot SCR_A \quad \text{cells}$$

where ε is the bound on CLR, as discussed in section 3.4. The "additional" traffic stream buffer length, L', must, in addition to the MQL, be sufficient to insure that no more than N_L cells are lost, or

$$L' = \frac{MBS_E}{PCR_E} \cdot SCR_A - \varepsilon \cdot \frac{MBS_E}{SCR_E} \cdot SCR_A$$

which may be rearranged as

$$L' = MBS_E \cdot SCR_A \cdot \left(\frac{1}{PCR_E} - \frac{\varepsilon}{SCR_E} \right) \tag{3.25}$$

The required total maximum queue length required for the "additional" traffic stream is therefore $MQL' = L' + MQL$, with L' being the buffer component associated with the burstiness of the serving mechanism, and MQL being the buffer component associated with the burstiness of the "additional" traffic stream. The use of "additional" traffic stream queue lengths of MQL' ensures that the required CLR, as bounded by ε, will not be exceeded.

3.4.4 Summary

Both the single and multiple empty timeslot methods described in the previous sections result in an estimate of an allowable SCR rate of traffic entering a queue served by a WRR, after the served path is considered "full" based on normal CAC methods. The only unknown parameter required to determine $\hat{R}_{ASCR}$ in each case, based on Equations (3.19) and (3.24) is the mean number of empty slots per cycle, μ_X. This mean may be estimated in real time by measurement of the number of empty timeslots being sent to the virtual path as the WRR cycles. The relationships between μ_X and $\hat{R}_{ASCR}$ are shown in Figure 3-5 for the single empty timeslot method, and in Figure 3-6 for the multiple empty timeslot method.

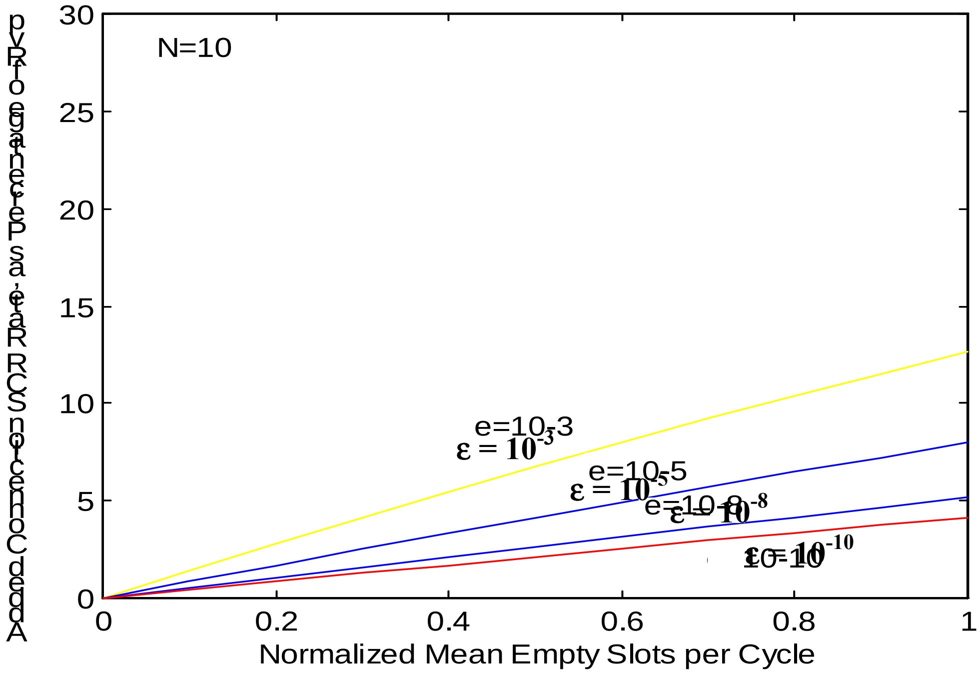

Figure 3-5: Added Connection Cell Rate, Single Empty Timeslot Method

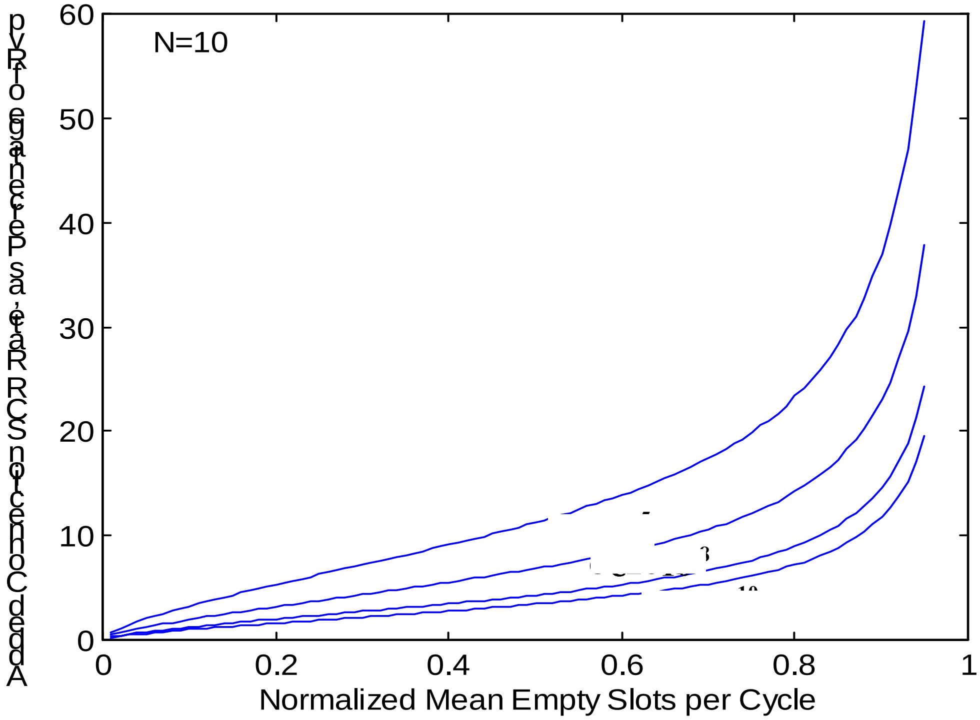

Figure 3-6: Added Connection Cell Rate, Multiple Empty Timeslot Method

Both methods will be conservative in allowing the admission of "additional" traffic to a "full" path. The single empty timeslot method is the more conservative of the two approaches, allowing the admission of a smaller traffic stream for a given mean number of empty timeslots, but requires fewer computing resources than the multiple empty timeslot method. Admission of significant amounts of "additional" traffic are possible with both methods, with "additional" traffic rates averaging from 3.5 to 11.5 percent of the virtual path rate for CLR's of 10^{-10} to 10^{-3} for the multiple empty timeslot method. Average ranges from 2 to 6.75 percent of the virtual path rate are possible using the single empty timeslot method. When "additional" VBR connections are included in the virtual path, the server will still allow some amount of non-conforming or low-priority cells to be sent in virtual path cell slots. This is possible since at a given point in time there is some probability that the "additional" traffic queue(s) will be empty, thus not requiring service, when only non-conforming or low priority cells are at the head of some "normal" queues. Because of the methods used in determining R_{ASCR}, the multiple empty timeslot method will fill a virtual path more fully with conforming or high-priority cell traffic than the single empty timeslot method, in return for the additional computing complexity. Both methods require the use of larger buffers than required for "normal" traffic streams having the same cell rate. Thus, either method may be used depending on specific equipment design and network configuration tradeoffs.

The next step in developing a workable CAC will be to determine a method for estimating the actual mean number of empty timeslots based on measurements of observed empty timeslots during cycles of the WRR.

3.5 Admission of "Additional" Traffic

The CAC procedure broadly described in Section 2.4 requires knowledge of the available, or allowable, "additional" bandwidth in the virtual path when a connection request is made and the virtual path serving the connection is "full" based on the sum of the cell rates of the existing traffic. In the previous sections of this chapter a relationship between the mean number of empty slots per cycle in a WRR, μ_X, and an allowable cell rate for "additional" traffic, R_{ASCR}, was determined. The CAC procedure will therefore actually determine the mean value corresponding to the desired cell rate of the "additional" connection, using the relationships of the previous sections, and then determine whether this desired mean value, μ_D, is less than or equal to the actual mean number of empty timeslots per cycle, μ_X. Unfortunately, the actual number of empty timeslots per cycle is unknown, and estimates of μ_X, using a measured mean number of empty timeslots per cycle, or sample mean, must be used to form hypotheses and make decisions regarding admission of an "additional" connection. The admission procedure must determine whether a given sample mean value, $\overline{X}$, justifies a conclusion that the true mean number of empty timeslots per cycle is greater than the desired mean number of empty timeslots per cycle. The CAC process thus requires that a test be made between two possibilities, which are

$$\mu_X \leq \mu_D \quad \text{(in which case the connection request will be rejected)}$$
$$\mu_X > \mu_D \quad \text{(in which case the connection request will be accepted)}$$

The formulation of the test will determine a procedure for connection admission.

3.5.1 Statistical Procedure for Admission

Hypothesis testing provides a statistical procedure for making decisions between contrasting possibilities, such as those presented above. Using standard statistical analysis procedures, hypotheses may be formed from the possibilities, resulting in

$$H_0 : \mu_X = \mu_D, \text{ the hypothesis}$$
$$H_A : \mu_X > \mu_D, \text{ the alternate hypothesis}$$

where the hypothesis is simplified by using the bound $\mu_X = \mu_D$ rather than $\mu_X \leq \mu_D$, with μ_D, which separates the hypotheses, being the *null value*, and H_0 being the *null hypothesis*. With the hypotheses defined in this manner, if H_0 is accepted, the connection request will be rejected, and if H_A is accepted (H_0 is rejected) the connection request will be accepted. The connection request will be rejected unless data strongly suggests that H_A is true.

The sample mean $\overline{X}$ is an unbiased estimator of the true mean μ_X [HOG95], and thus may be used as a basis for a testing procedure. The sample mean is evaluated at every cycle of the WRR, and thus is a function of the number of observed cycles, *n*, that have elapsed since measurement started, and the values of X at each cycle, or

$$\overline{X}(n) = \frac{\sum_{i=1}^{n} X(i)}{n} \tag{3.26}$$

Data will favor H_A only if the observed value of the sample mean is larger than μ_D, so H_0 will be rejected, and H_A accepted only if $\overline{X}(n) - \mu_D$ is a large positive number.

Determining whether $\overline{X}(n) - \mu_D$ is large enough to justify the rejection of H_0 depends on the variance of the sample mean values, and therefore a test statistic of

$$A = \frac{\overline{X}(n) - \mu_D}{\sqrt{\hat{Var}[\overline{X}(n)]}} \tag{3.27}$$

where, *if the X's are independent, and therefore uncorrelated*, an unbiased estimator [LAW91] of the variance of the sample mean would be

$$\hat{Var}[\overline{X}(n)] = \frac{S^2(n)}{n} = \frac{\sum_{i=1}^{n} [X(i) - \overline{X}(n)]^2}{n(n-1)} \tag{3.28}$$

Using the above statistic *A*, the rejection region is $A \geq \phi$, where ϕ is the critical value to be used in making the decision concerning the hypothesis. Critical values are determined by controlling the probability of making a Type I error. A Type I error results if H_0 is rejected when it is actually true, resulting in a connection being admitted when it should not have been. The probability of Type I error is denoted by $\alpha = \Pr(A \geq \phi \text{ when } \mu_X = \mu_D)$. The critical values corresponding to various values of α are derived from normal probabilities.

As an example, an α of .01 corresponds to 99% confidence that the true value of the mean is greater than the value of the critical point, which is the mean needed to support the desired SCR on the "additional" connection, for which admission to the network is being requested. For $\alpha = .01$, the corresponding critical value ϕ is 2.326, from normal probability distributions. In general this may be expressed as $\phi = \Phi^{-1}(1-\alpha)$ where $\Phi^{-1}(1-\alpha)$ is the value of the standard normal distribution associated with cumulative probability of less than $1-\alpha$. Any value of α between 0 and 1 may be chosen, however, the particular α chosen will have an effect on cell loss of the "additional" connection, since α, through its association with the critical value, has a role in the acceptance or rejection of the hypothesis. It is also possible to have a Type II error, where H_0 is not rejected (i.e., the connection request is rejected) when H_0 is actually false (i.e., the true mean is greater than the desired mean). This type of error, although having significant probability when the desired mean is close to the true mean, actually contributes to the conservative nature of this method of CAC, since an error of this type will result in the rejection of a connection request which could have been accepted.

The procedure for admission using these concepts will be as follows. Starting from when the last connection was admitted, $\overline{X}(n)$ and $\hat{Var}[\overline{X}(n)]$ will be updated with measured data at each WRR cycle, *n*, until a new connection is admitted, thus changing the value of the statistic *A* (having a μ_D corresponding to the desired connection cell rate) at each cycle. Thus, for a given α

$$\frac{\overline{X}(n)-\mu_D}{\sqrt{\hat{Var}[\overline{X}(n)]}} \geq \Phi^{-1}(1-\alpha)$$

is the rejection region, where rejecting the hypothesis because the statistic is not large enough corresponds to accepting the connection. Thus, solving the above for the desired mean, a requested connection having a desired cell rate corresponding to a desired mean number of empty cells per WRR cycle, μ_D, *will be accepted if*

$$\overline{X}(n)-[\Phi^{-1}(1-\alpha)]\cdot\sqrt{\hat{Var}[\overline{X}(n)]} \geq \mu_D$$

and the requested connection *will be rejected* if

$$\overline{X}(n)-[\Phi^{-1}(1-\alpha)]\cdot\sqrt{\hat{Var}[\overline{X}(n)]} < \mu_D$$

This method of admission requires a calculation of the statistic and a decision as to whether or not to accept the hypothesis at the time admission to a “full” VP is requested. In order for the statistic to be evaluated, knowledge of both the sample mean and the variance of the sample mean is required. The sample mean is obtained from Equation (3.26). However, variable bit rate traffic, characterized by PCR, SCR, and MBS parameters, arrives in bursts and is queued, resulting in samples that are not independent, but are correlated to an extent determined by the burst sizes and cell rates of the individual traffic streams. Since Equation (3.28) requires independent samples, *X*(*i*), to determine the variance of the sample mean, there is a need for a modified estimate of the variance of the sample mean, as discussed in the next section.

3.5.2 Estimating Variance of the Sample Mean of Correlated Samples

If the actual mean and variance of the process producing the samples of the number of empty timeslots per cycle is stationary over time, but the samples are correlated, rather than using Equation (3.28) the variance of the mean number of empty slots per cycle may be estimated [LAW91] from

$$\hat{Var}\left[\bar{X}(n)\right] = \frac{S^2(n)}{n}\left[1 + 2\sum_{j=1}^{n-1}\left(1 - \frac{j}{n}\right)\cdot \rho_j\right] \tag{3.29}$$

where ρ_j is the correlation between the first sample and the j^{th} sample in a sequence of n samples. The ρ_j terms associated with a sequence of n samples may be estimated from the following relationships [LAW91]:

$$\hat{\rho}_j = \frac{\hat{C}_j}{S^2(n)}, \text{ where } \hat{C}_j = \frac{\sum_{i=1}^{n-j}\left[X(i) - \bar{X}(n)\right]\cdot\left[X(i+j) - \bar{X}(n)\right]}{n-j}, \quad j = 1,2,\ldots n-1 .$$

These estimates of ρ_j are biased and have a large variance themselves, requiring very large *n,* and small *j* relative to *n* to get good estimates. To determine the characteristic behavior of the correlation coefficients of *X*(*i*) in a WRR, simulations of a "full" VP served by a WRR having a variable number of VBR connections were run, using BONeS® DESIGNER™ [BON96]. The simulated connections had on-off traffic characteristics with diverse SCR and MBS parameters, taken from Appendix A. Details of the traffic and virtual path models used for the simulations are discussed in Section 4.1.3. The overall MBS was considered to be the average of the MBS of each connection served. A typical simulation result is shown in Figure 3-7.

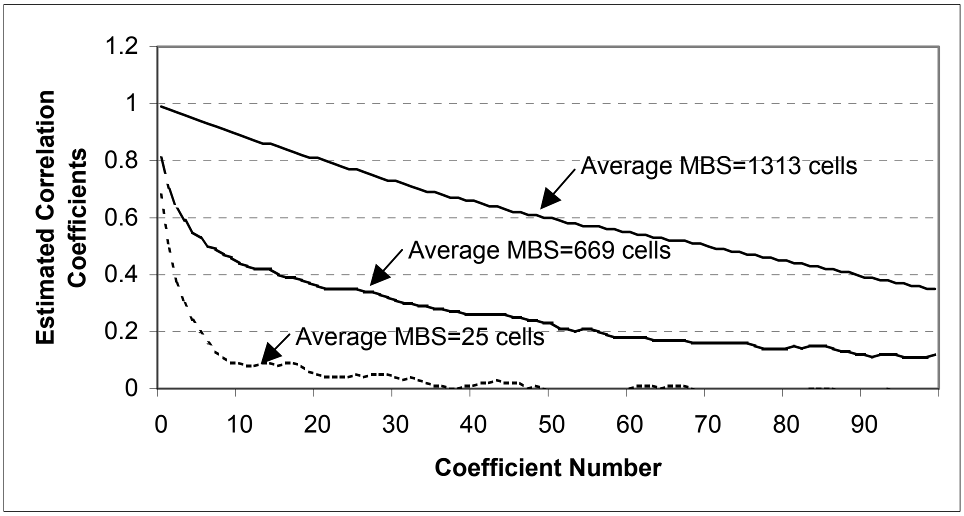

Figure 3-7: Typical Correlation Coefficients of Samples of *X*

Based on the simulations, a bound on the correlation coefficients was estimated to be that shown in Figure 3-8.

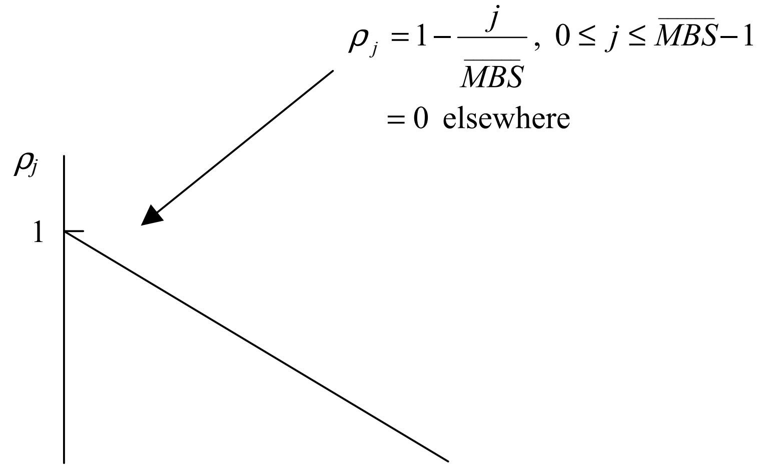

$$\rho_j = 1 - \frac{j}{\overline{MBS}},\ 0 \le j \le \overline{MBS} - 1$$
$$= 0 \text{ elsewhere}$$

Figure 3-8: Estimated Correlation Coefficients of X

Using the function described in the above figure as an estimate of ρ_j in Equation (3.29), the correlation coefficients will be zero when $j \geq \overline{MBS}$. Therefore, in order to insure that the correlation is accounted for and is no longer affecting the independence of the samples, it will be required that at least $\overline{MBS}$ cycles elapse prior to the samples being used in an evaluation of $\hat{Var}\left[\overline{X}(n)\right]$. If this is the case, then from Equation (3.29)

$$\hat{Var}\left[\overline{X}(n)\right] = \frac{S^2(n)}{n}\left[1 + 2\sum_{j=1}^{\overline{MBS}-1}\left(1 - \frac{j}{n}\right)\cdot\left(1 - \frac{j}{\overline{MBS}}\right)\right], n \geq \overline{MBS} \tag{3.30}$$

Considering the summation term in Equation (3.30),

$$
\begin{aligned}
&\sum_{j=1}^{\overline{MBS}-1}\left(1-\frac{j}{n}\right)\cdot\left(1-\frac{j}{\overline{MBS}}\right)\\
&=\left(1-\frac{1}{n}\right)\left(1-\frac{1}{\overline{MBS}}\right)+\left(1-\frac{2}{n}\right)\left(1-\frac{2}{\overline{MBS}}\right)+\mathrm{K}+\left(1-\frac{\overline{MBS}-1}{n}\right)\left(1-\frac{\overline{MBS}-1}{\overline{MBS}}\right)\\
&=1\text{-}\left(\frac{1}{n}+\frac{1}{\overline{MBS}}\right)+\frac{1}{n\cdot\overline{MBS}}+1\text{-}2\left(\frac{1}{n}+\frac{1}{\overline{MBS}}\right)+\frac{2^2}{n\cdot\overline{MBS}}+1\text{-}3\left(\frac{1}{n}+\frac{1}{\overline{MBS}}\right)+\frac{3^2}{n\cdot\overline{MBS}}+\mathrm{K}\\
&=\left(\overline{MBS}-1\right)(1)+\left(1+2+\Lambda+(\overline{MBS}-1)\right)\left(\frac{1}{n}+\frac{1}{\overline{MBS}}\right)+\left(1^2+2^2+\Lambda+(\overline{MBS}-1)^2\right)\left(\frac{1}{n\cdot\overline{MBS}}\right),\\
&\hfill n\geq\overline{MBS}
\end{aligned}
$$

Substituting the results of the series $1+2+...+(\overline{MBS}-1)=(\overline{MBS}-1)(\overline{MBS})/2$,

and $1^2+2^2+...+(\overline{MBS}-1)^2=(\overline{MBS}-1)(\overline{MBS})(2\,\overline{MBS}-1)/6$ [HAN94],

$$
\begin{aligned}
&\sum_{j=1}^{\overline{MBS}-1}\left(1-\frac{j}{n}\right)\cdot\left(1-\frac{j}{\overline{MBS}}\right)\\
&=\left(\overline{MBS}-1\right)+\frac{1}{2}\left(\overline{MBS}-1\right)\left(\overline{MBS}\right)\left(\frac{1}{n}+\frac{1}{\overline{MBS}}\right)+\frac{1}{6}\left(\overline{MBS}-1\right)\left(\overline{MBS}\right)\left(2\,\overline{MBS}-1\right)\left(\frac{1}{n\cdot\overline{MBS}}\right)\\
&=\left(\overline{MBS}-1\right)\left(1-\frac{\overline{MBS}+n}{2n}+\frac{2\,\overline{MBS}-1}{6n}\right),\ n\geq\overline{MBS}
\end{aligned}
$$

Substituting this back into Equation (3.31),

$$
\hat{Var}\left[\overline{X}(n)\right]=\frac{S^2(n)}{n}\left[1+2\left(\overline{MBS}-1\right)\left(1-\frac{\overline{MBS}+n}{2n}+\frac{2\,\overline{MBS}-1}{6n}\right)\right],n\geq\overline{MBS} \qquad (3.31)
$$

Having determined an expression for $Var[\overline{X}(n)]$ for use in testing the hypothesis that $\overline{X}(n) \leq \mu_X$, knowledge of how many cycles *n* are needed to insure the accuracy of the estimate is required. In accordance with the Central Limit Theorem, if a sufficient number of samples is taken, the distribution of sample means can be assumed to be approximately normal, N($\overline{X}(n), \hat{Var}\left[\overline{X}(n)\right]$), and it can be converted to a standard normal function, N(0,1), using the sample mean, and the estimate of the variance of the sample mean.

The use of the standard normal function results in an accurate estimate only if *n*, the number of WRR cycles over which the estimate is being formed, is "large". The validity of the estimate of the variance $Var[\overline{X}(n)]$ for a particular value of *n* depends on the actual characteristics of the probability distribution of *X*, the number of empty timeslots per cycle. If the underlying distribution is close to a normal distribution and the samples are independent or have small correlation, then the approximation will be good for small *n*. If the underlying distribution is skewed compared to a normal distribution, or the samples are not independent, having high correlation, then the approximation will be good for large values of *n*. A minimum value of *n* has been estimated at $n \leq 30$ [DEV82]. However, to insure minimal correlation, the value of *n* must be determined based on the MBS characteristics of the connections being served by the WRR, which can be inferred from the correlation estimate shown in Figure 3-8. Thus, a conservative approach is to sample for at least the average MBS of all the connections served by the WRR, providing this quantity is greater than 30. This requirement places further restrictions on the estimate of the variance of the mean, and Equation (3.31) becomes

$$\hat{Var}\left[\overline{X}(n)\right]=\frac{S^2(n)}{n}\left[1+2\left(\overline{MBS}-1\right)\left(1-\frac{\overline{MBS}+n}{2n}+\frac{2\overline{MBS}-1}{6n}\right)\right], n\geq \max\left(30,\overline{MBS}\right). \quad (3.32)$$

The sample mean, $\overline{X}(n)$, and variance of the sample mean, $\hat{Var}\left[\overline{X}(n)\right]$, will be evaluated at each WRR cycle, i.e., as n increases, until a new connection is admitted. These sampled parameter estimates can then be used in making decisions for the admission of additional connections in a virtual path, at any time after the required 30 cycles has elapsed. Procedures for accomplishing this are discussed in the next section.

3.5.3 The Admission Process

To accomplish admission of "additional" traffic streams(s) when the virtual path is considered to be "full", measurements of the number of empty slots per cycle will be taken at each WRR cycle, and used to calculate $\overline{X}(n)$, the sample mean, at each cycle, as described by Equation (3.26). For the purpose of these measurements, timeslots counted as empty will include timeslots containing low priority cells, since these cells may be replaced with the high priority cells of the "additional" connection. The sample mean may then be used in conjunction with the estimate of the variance of the sample mean, $\hat{Var}\left[\overline{X}(n)\right]$ described in Section 3.5.2, to test the hypothesis used in admitting new connections. Measurements of $\overline{X}(n)$ can be used as if they were from a normal distribution after $n = 30$ cycles of the WRR have elapsed, however, conditions described in the previous section require that no evaluation of the true mean be accomplished until $\overline{MBS}$ cycles have elapsed, where $\overline{MBS}$ is the average maximum burst size of the traffic streams entering the queues that make up the WRR, as required by Equation (3.32).

Knowledge of $\overline{X}(n)$ and $\hat{Var}\left[\overline{X}(n)\right]$ gained from Equations (3.26) and (3.32) leads to the ability to test the given hypothesis, with a specified level of probability of error, α, corresponding to a desired confidence level, $1-\alpha$. This overall process defines the measurement-based CAC that is the subject of this analysis.

3.6 Considerations in Admitting "Additional" Traffic

Any CAC method must insure that the required QoS is maintained for both the existing connections and the newly admitted connections. The network framework used in this analysis, consisting of end-to-end-virtual paths, class of service separation, WRR queue service, and queues for normal connections sized to insure no cell loss for conforming traffic, has been described in Chapter 2. Since the "additional" queues are served only whenever a timeslot is available while the WRR is cycling through the "normal" queues, no affect on "normal" traffic can result using this method, other than a possible reduction in the number of non-conforming or low-priority cells transmitted into the network from the "normal" traffic connections.

Thus, the result of admitting “additional” traffic is necessarily evaluated by determining whether or not the “additional” traffic will meet its QoS objective, which in terms of this analysis is measured in terms of cell loss.

Cell loss in an “additional” connection will occur whenever the connection is not served often enough to insure that no more than an allowed number of cells are lost. In determining the allowable rate of an additional connection, the characteristics of the existing connections will determine the rate at which the additional connection can be served. The evaluation of the characteristics of the existing connections presented here is conservative at every step to assure the QoS of the additional connection.

The allowable SCR of an “additional” connection was taken to be the lower bound on the possible range of allowable SCR in both the single empty timeslot method, resulting in Equation (3.19), and the multiple empty timeslot method, resulting in Equation (3.24). The number of cycles needed to insure the service of an additional timeslot, on which Equations (3.19) and (3.24) are based, was also taken as a bound in both methods. The bound on the number of cycles needed to admit an “additional” timeslot was calculated based on the assumption that the queue length was one rather than a much longer queue, calculated based on Equations (3.1) and (3.2) and discussed in Section 3.4.1.1. Thus, the relationship between the connection rate of additional traffic, R_{ASCR}, and the mean number of empty timeslots per cycle, μ_X, is conservative and will allow the QoS to be met as long as μ_X may be accurately estimated. In order to estimate μ_X, the mean of samples of X is used, along with an estimate of its variance.

The methods used in estimating the variance and testing a statistical hypothesis regarding whether the sample mean is less than the actual mean, described in Section 3.5, are also conservative in that they insure all correlation in the input traffic streams is accounted for, before allowing a decision on the hypothesis.

Thus, being conservative in all aspects of its application, the methods analyzed here should prove adequate to allow required quality of service to "additional" connections in a "full" virtual path.

3.7 Summary

This chapter is focused on specific methods for adding traffic to broadband network paths. It contains an analysis of WRR operation in a source switch that is part of an overall network framework that requires end-to-end virtual paths, and separates the paths serving VBR and CBR traffic from those serving ABR and UBR traffic.

The analysis shows how a network user, or customer, may calculate SCR of requested traffic connections, in spite of lack of specific knowledge regarding the size of the queue serving the traffic. This may be accomplished through solving two simultaneous equations involving queue length and SCR, if the user has knowledge that the network provider is using the method suggested in this chapter. The analysis shows how the mean number of empty timeslots per cycle in a source switch WRR will be a function of the SCR declared by the user, and suggests two methods for determining available capacity in a virtual path considered by normal CAC methods to be full. In both methods, the available or “additional” capacity is a function both of the mean number of empty timeslots per WRR cycle, and the desired CLR. In an actual system, where the true mean number of empty timeslots per cycle is unknown, analysis shows that the sampled mean number of empty timeslots per cycle, and the associated estimate of the variance, may be used in creating statistics needed for hypotheses used to test the critical admission criteria. The methods used are conservative in allocating allowable “additional” bandwidth, being based on using the boundaries of the required relationships. Finally, this chapter describes a method for admission of “additional” connections to a “full” virtual path, without exceeding required QoS for any connection. The next chapter will address WRR server and traffic simulations used to support the analyses presented here.

Chapter 4

Evaluation and Discussion of Results

CHAPTER CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter

4

Evaluation and Discussion of Results

In this chapter, connection admission using the methods described in the previous chapters will be evaluated. The network framework described in Chapter 2 and analyzed in Chapter 3 allows connection admission to an end-to-end virtual path be accomplished at the source only, and guarantees the QoS of any connection that is properly admitted to the path, through WRR-based queue service.

4.1 Evaluation: Traffic and Path Model Simulation

Simulation of both traffic and a virtual path served by a WRR mechanism will be used to evaluate the analytical results of Chapter 3. Several tasks must be accomplished for this evaluation.

First, models must be created and tested to insure they properly simulate the traffic and WRR-served path parameters needed to support the analysis. Once a traffic and path simulation model is created, several steps must be accomplished to show the validity of the model. Measurements of simulated traffic flowing through a simulated WRR server must result in a mean number of empty timeslots per WRR cycle, X, on the virtual path that reflects the predicted mean of X based on the applied traffic parameters. Also, WRR server queues sized in accordance with the proposed network framework must have no overflows when the virtual path is “full” based on the sum of the SCR of the served traffic streams. In this regard, Section 4.1.1 will describe the traffic parameters used in this study, and Section 4.1.2 will describe the path and server models used, as well as the technique used to evaluate the performance of the model in regard to queue overflows and measurements of X.

Next, the validity of the connection acceptance criteria determined theoretically in Chapter 3, which is based on the sampled mean of X and the estimate of the variance of the sampled mean of X, must be determined. To accomplish this, the estimate of allowable “additional” traffic, developed in Section 3.6, must be validated. The simulation model of a “full” virtual path serving traffic with known characteristics (i.e., PCR, SCR, mean burst size, maximum burst size and utilization) must be shown to have an estimated mean of X that approaches, but never exceeds, the actual mean of X predicted by the empty queue probabilities, as analyzed in Chapter 3. Since this estimate depends on an evaluation of the correlation characteristics of X, as was discussed in Section 3.6.2, an estimate of correlation is required. The procedures used in accomplishing the above will be described in Sections 4.1.3 and 4.1.4.

Finally, a WRR-served virtual path that is already “full” must be shown to be capable of properly serving an “additional” connection, where the “additional” connection is introduced using the available capacity and queue sizing criteria developed in Chapter 3. This will be accomplished as described in Section 4.1.5.

In order to accomplish the tasks outlined above, models simulating multiple traffic channels entering a WRR-served virtual path were created using the Block Oriented Network Simulator, BONeS® DESIGNER™ [BON96]. Traffic parameters for the required simulations were chosen from a range of fixed utilization and fixed probability traffic types (defined in the next section), designed to fully explore the bounds of the design space. A total of 15, 374 simulations were run for durations ranging from 100,000 to 300,000 virtual path cell timeslots. The results of these simulations confirmed the operation of the modeled traffic sources and the path models, and provided the data used in evaluating correlation characteristics and estimates of the mean of *X*. The simulations also fulfilled a primary purpose of the research, by providing data on the operational characteristics of an "additional" connection introduced to a "full" virtual path in accordance with the proposed technique. The results of these simulations are summarized in the following sections.

4.1.1 Selection of Source Traffic Scenarios

On-off traffic sources, which send bursts of cells at a designated PCR for a random "on" period and send no cells for a random "off" period, were used in this analysis. On-off traffic source models have been shown to be capable of modeling voice, bursty sources including images, and video, among other sources [SCH96].

The characteristics of the on-off traffic sources are defined by setting their PCR, mean burst length, and offered load, or utilization. Appropriate queue lengths can be determined, using equivalent capacity techniques as the basis for SCR, for a given desired maximum cell loss ratio, CLR, and sustainable cell rate, SCR, as described in Section 3.2, Equations 3.4 and 3.5. On-off sources having a peak cell rate, PCR, equal to the virtual path rate, were used throughout this analysis. Two sets of simulation source traffic parameters were used, covering a wide range of utilization, burst size, and queue length alternatives. Since testing using small cell loss ratios would require extremely long simulation run times to support the validity of the model, all server, path, and traffic models were based on a CLR bound of $\varepsilon = 10^{-3}$. To obtain reasonable simulation cell granularity while making results easily scalable, a virtual path rate of 1000 cells/second was used in all simulations. On-off traffic source models having geometrically distributed on-time and exponentially distributed off-time characteristics were used.

The first set of source traffic parameters used *fixed utilization* values of $\rho = .05$, .1, .2, .4 and .8, as appropriate to the SCR, to generate various probabilities of empty queues, which would result in various mean number of empty timeslot per cycle values for different scenarios. Solutions providing traffic parameters for a SCR of 100 through 900 cells per second, having the utilization values described above, and mean burst lengths of 2, 5, 10, 50, and 100 cells, are shown at Appendix A1.

The second set of simulation source traffic parameters was generated with the traffic stream utilization values chosen such that the resulting probabilities of empty queues occur at fixed values of .1, .5, and .9. This in turn results in the corresponding normalized mean number of empty timeslot per cycle being fixed at values of .1, .5, and .9. These values were chosen as being the low end, middle and high end of the allowable rate vs. mean functions shown in Figures 3-3 and 3-4. Solutions having these *fixed probability* characteristics, providing traffic parameters over the same SCR and mean burst lengths used for the fixed utilization case discussed above, are shown at Appendix A2.

Simulated traffic having the characteristics noted above was used in creating all traffic scenarios used in this study. Actual traffic combinations used in specific scenarios are shown in Appendix B. This set of simulated traffic parameters consists of traffic with mean burst sizes of from 2 to 100 cells, maximum burst sizes of from 2 to almost 7,000 cells, and a range of utilization values resulting in sustainable cell rates (SCR) of 1% to 90% of the virtual path rate. These traffic parameters form an extensive design space, and result in a range of empty-queue probabilities from .1 to .9, (resulting in a range of possible *X* of from .1 to .9 of the timeslots in the WRR schedule). This design space is sufficient to extensively test the proposed method, being capable of determining where it is least conservative in its evaluation of available bandwidth through evaluation of the results.

4.1.2 The WRR Server with “Full” Virtual Path

The basic path model, consisting of multiple queues served by a WRR, as shown in Figure 2-11(a), is implemented using independent on-off traffic sources as described in the previous section, with the peak traffic rate equal to the virtual path rate. The on-off sources are policed, using a GCRA-based algorithm, as described in Chapter 1, by discarding cells that do not conform to designated PCR, SCR and MBS characteristics. A designated number of on-off source traffic streams enter model FIFO queues having designated queue length parameters that can be determined from Equations 3.1 and 3.2. The traffic streams are served by a WRR that examines each queue in accordance with a schedule based on the SCR of the associated traffic, skipping to the next queue in the schedule if the current queue is empty. For the purposes of this simulation, cells tagged as low-priority by the GCRA policer are discarded immediately, to reduce model complexity. In an actual system, the low priority cells could be allowed to enter the queues, and would be discarded only if a high priority cell in the "additional" traffic stream required the timeslot. Basic models capable of handling up to ten queues in a ten-timeslot schedule, and up to twenty queues in a twenty-timeslot schedule were constructed. Models having an "additional" queue, as shown in Figure 2-11(b), which is served whenever a "normal" queue was skipped, based on both ten and twenty normal timeslot models, were also constructed. In simulations where an "added" traffic queue is present, the "added" queue length is a model parameter that can be set to a size determined through procedures discussed in Section 3.4.3.

Simulations using the traffic and server models discussed above were run to confirm the characteristics of the analysis when traffic conforming to the scenarios described in Appendix B1 was introduced. The characteristics to be confirmed were the value of the mean number of empty timeslots per cycle resulting from the given traffic load, and the number of queue overflows. Simulations of ten-timeslot models were performed for each scenario where each queue was served on a *strictly scheduled* basis, and also where each queue was *skipped if empty*. A strictly scheduled server serves each queue in sequence, creating an empty cell timeslot each time an empty queue is encountered, whereas a skip-if-empty server will skip to the next queue if a given queue is empty.

Each of the simulations was run for 10,000 WRR cycles. Typical results of tabulating the empty timeslots per cycle for the strictly scheduled queue and skipped queue models are shown below for traffic streams having short burst and long burst characteristics, as illustrated by Figure 4-1 and Figure 4-2. The result of tabulating the mean number of empty timeslots per WRR cycle between the scheduled queue and skipped queue models using the fixed utilization scenarios of Appendix B1 is shown in Table 4-1. The results shown in Figures 4-1, and 4-2, and Table 4-1 show that the mean number of empty timeslots per WRR cycle for both *strictly scheduled* and *skipped if empty* models converge to the predicted mean value, thus validating the skipped queue model for use in simulations supporting this research.

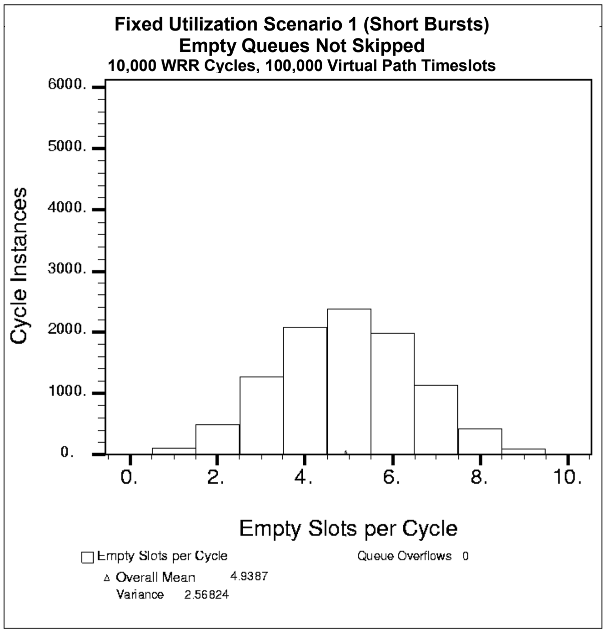

a) Queue Strictly Scheduled

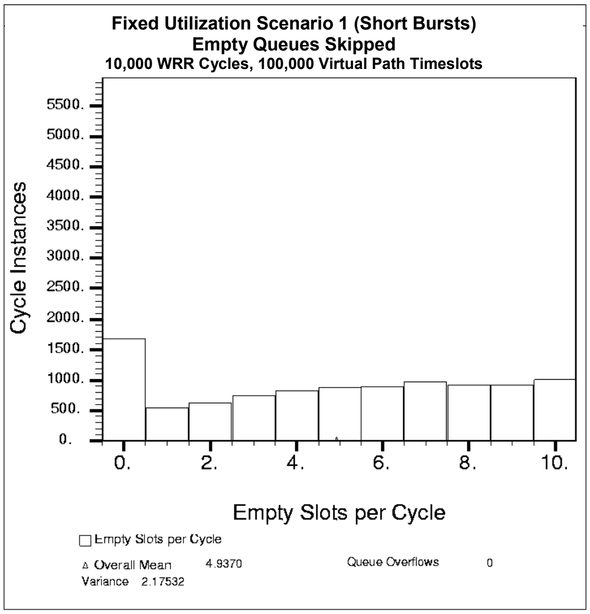

b) Queue Skipped if Empty

Figure 4-1. Empty Timeslots per Cycle, Short Burst Traffic Stream Scenario

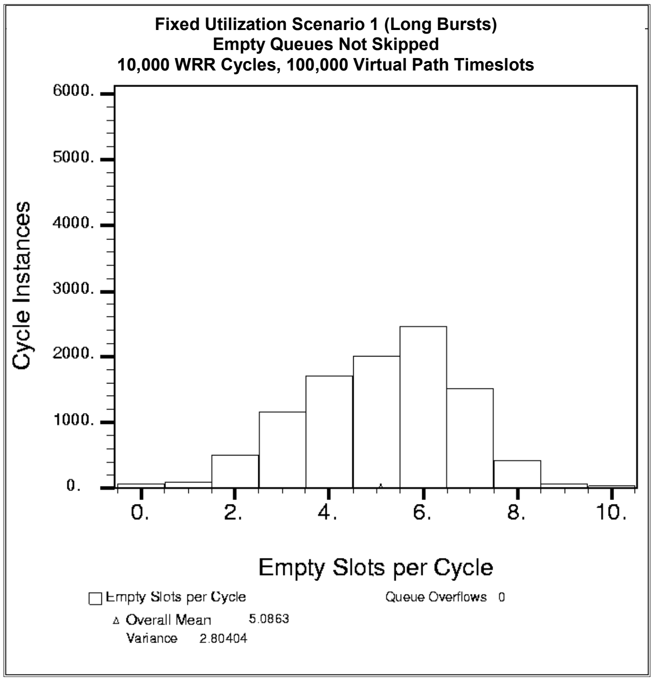

a) Queue Strictly Scheduled

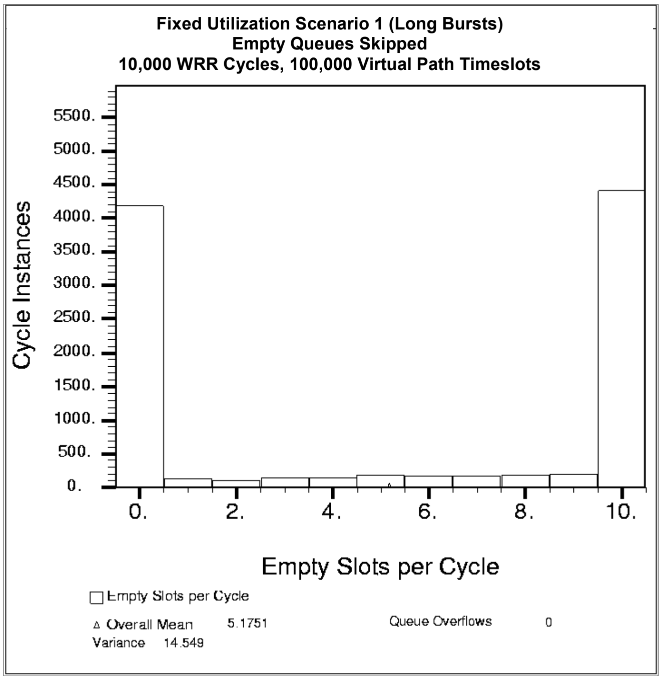

b) Queue Skipped if Empty

Figure 4-2. Empty Timeslots per Cycle, Long Burst Traffic Stream Scenario

Scenario (Fixed Utilization)	Predicted Mean	Burst Type	Scheduled Queue Mean	Skipped Queue Mean
1	5	short	4.95	4.94
		mixed	4.87	4.95
		long	5.00	5.17
2	5.5	short	5.44	5.43
		mixed	5.39	5.39
		long	5.40	5.61
3	5	short	4.91	4.91
		mixed	4.89	4.88
		long	5.15	5.11
4	7.5	short	7.44	7.47
		mixed	7.47	7.70
		long	7.29	7.69
5	5	short	4.81	5.63
		mixed	5.15	5.34
		long	5.18	4.89
6	7.98	short	7.95	7.94
		mixed	8.01	8.05
		long	7.84	8.12
7	4.98	short	4.20	4.20
		mixed	4.44	4.64
		long	4.58	4.76
8	3.5	short	2.87	2.96
		mixed	3.49	3.65
		long	3.52	3.91
9	9	short	8.98	8.97
		mixed	8.96	8.90
		long	8.93	8.93
10	2	short	.948	.966
		mixed	1.57	1.61
		long	2.28	1.83
11	8.96	short	9.08	9.07
		mixed	8.95	8.90
		long	8.94	8.93
12	1.49	short	.803	.794
		mixed	1.45	1.44
		long	1.45	1.47

Table 4-1. Mean Empty Timeslots per Cycle, Scheduled and Skipped Queue Models

Results of creating and testing simulation models for the WRR served virtual path showed that the models met their predicted characteristics in terms of the mean number of empty slots per cycle observed when traffic of known characteristics was applied to the model. The traffic applied was such that the virtual path was “full” based on the sum of the SCR of the traffic streams being served. This result was consistent over a wide range of input traffic parameters.

4.1.3 Traffic Correlation

Simulations designed to provide an estimate of the correlation characteristics of the number of empty timeslots in the virtual path traffic stream, using traffic parameters from Appendix A, in simulation scenarios from Appendix B, were carried out using ten- and twenty-timeslot models. Simulations for each fixed utilization scenario, for long, mixed, and short bursts, from Appendix B1, as well as fixed probability scenarios 1, 2, and 3 for long, mixed, and short bursts, from Appendix B2, were accomplished, where the number of empty timeslots per cycle over 10,000 cycles (100,000 timeslots) was recorded. The recorded values were then used in calculating the correlation coefficients of the number of empty timeslots per cycle. The need for determining the correlation between samples used in estimating the variance of the sample mean is discussed in Section 3.5.2.

The results of these simulations showed that the correlation coefficients of samples of X always reached a value close to zero within a number of samples corresponding to the average MBS of the traffic sources being served. These results were used in Chapter 3 in developing an estimate of a bound on the correlation coefficients X for use in determining an estimate of the variance of the sample mean of X. A typical example of the result of the simulation was shown as Figure 3-5, and will not be repeated here.

4.1.4 Estimate of Allowable "Additional" Traffic

Simulations of a "full" VP, as defined in Chapter 3, in a WRR with traffic characterized by PCR, SCR, and MBS, were run to test the validity of the estimate of the variance of the mean of X as a bound on the true variance of the mean of X. This is necessary in order to show that the estimate of the variance of the mean of X as defined by Equation (3.32) may be used in determining the acceptance and rejection criteria defined in Section 3.5. In order for the prediction of the variance of the mean of X , from Equation (3.32),to be valid for this analysis, it must always be greater than the actual variance of sample means taken from multiple simulations. To test this, simulations of scenarios 1 and 2 of Appendix B1, having the MBS characteristics designated by Appendix A1 for each SCR, were run 20 times per scenario in simulations with different seeds for the random traffic streams. The simulations were run for various numbers of cycles depending on the average MBS of the given scenario. In all cases where the simulation ran at least a number of WRR cycles equal to the average MBS of all the queues, the prediction of the variance of the mean number of empty slots was greater than the actual variance observed in the 20 iterations of each scenario. The results of evaluating typical simulation runs from Scenario 1 are shown at Table 4-2.

Table 4-2 shows the results of estimating the variance of the mean number of empty timeslots after the various numbers of WRR cycles tabulated have elapsed. The number of cycles over which the estimate is made is designated as *n* in the table.

Fixed Utilization Scenario 1
10 Queue, 10 Timeslots, Twenty Trials
[True Mean Number of Empty Timeslots/Cycle = 5.0]

Number of WRR Cycles, *n*	Measurement Type	Average MBS (cells)		
		$\overline{MBS} = 25$	$\overline{MBS} = 656$	$\overline{MBS} = 1313$
25	Mean of $\overline{X}$	5.670	8.880	9.410
	Var of $\overline{X}$	.606	.614	.484
	Min. Est. Var of $\overline{X}$	.830	.140	.000
50	Mean of $\overline{X}$	5.250	8.200	8.880
	Var of $\overline{X}$	.455	.581	.621
	Min. Est. Var of $\overline{X}$	.820	.550	.140
100	Mean of $\overline{X}$	5.090	7.390	8.200
	Var of $\overline{X}$	.159	.829	.585
	Min. Est. Var of $\overline{X}$	.720	.950	.550
500	Mean of $\overline{X}$	4.970	5.730	6.370
	Var of $\overline{X}$	.047	.727	.849
	Min. Est. Var of $\overline{X}$	.18	1.16	1.260
656	Mean of $\overline{X}$	X	5.520	X
	Var of $\overline{X}$		.531	
	Min. Est. Var of $\overline{X}$		.990	
1000	Mean of $\overline{X}$	4.940	5.270	5.730
	Var of $\overline{X}$	.022	.437	.725
	Min. Est. Var of $\overline{X}$	.100	.900	1.160
1313	Mean of $\overline{X}$	X	X	5.520
	Var of $\overline{X}$			.530
	Min. Est. Var of $\overline{X}$			.990
5000	Mean of $\overline{X}$	4.940	5.05	5.100
	Var of $\overline{X}$	.006	.062	.122
	Min. Est. Var of $\overline{X}$	.020	.480	.820

Key:

░ Point where number of WRR cycles, *n*, are equal to $\overline{MBS}$

▓ Points where the Minimum Est. Var $\overline{X}$ < Var $\overline{X}$

X Point not tested. Tests were conducted after 25, 50, 100, 500, 1000, and 5000 cycles had elapsed, and at the point where the number of WRR cycles, *n*, are equal to $\overline{MBS}$

Table 4-2. Typical Mean, Variance and Minimum Estimated Variance of Sample Means

For the particular scenario shown, the actual mean number of timeslots calculated from the empty queue probabilities of the traffic types being served, or the true mean of the scenario, was 5.0. As can be seen in the table, as the number of WRR cycles, *n*, increases, the mean of the sample means from the 20 trials converges to the true mean expected for this WRR-served traffic configuration. The results at a point where *n* is equal to the average MBS of the existing queues, $\overline{MBS}$, are shaded with light gray. The points where the minimum values of estimated variance of $\overline{X}$ from 20 trials are less than the actual variance of $\overline{X}$ over 20 trials are shown as dark gray. These dark gray areas occur, in all cases, at values of N that are less than the average MBS of the existing queues, thus supporting the predictions of the theory developed in Chapter 3.

Thus, the results described in Table 4-2 indicate that the estimate of the variance of the sample mean, from Equation 3-32,

$$\hat{Var}\left[\overline{X}(n)\right]=\frac{S^2(n)}{n}\left[1+2\left(\overline{MBS}-1\right)\left(1-\frac{\overline{MBS}+n}{2n}+\frac{2\overline{MBS}-1}{6n}\right)\right], n\geq \max\left(30,\overline{MBS}\right)$$

where the estimate is taken in the region of *n* that is greater than $\overline{MBS}$ cycles or 30 cycles, whichever is larger, may be used in forming confidence intervals for the mean number of empty timeslots per cycle, which may in turn be used with appropriate confidence in determining the SCR of the “additional” traffic that may be admitted to a “full” virtual path.

Having used simulation to confirm that the prediction of the variance of the sample mean is an acceptable bound on the actual variance of sample means of a given traffic type for the cases considered, the use of this estimate in determining an allowable "additional" traffic stream cell rate was examined. In Section 3.5.1 the relationship used in accepting an "additional" traffic flow to a "full" path was shown to be

$$\overline{X}(n) - [\Phi^{-1}(1-\alpha)] \cdot \sqrt{\hat{Var}[\overline{X}(n)]} \geq \mu_D$$

where μ_D is the mean value corresponding to a desired "additional" traffic stream cell rate, as determined by Equation 3-19 or Equation 3-24, and α is a desired confidence level. From the above, an estimate of the "allowable mean", $\hat{\mu}_A$, based on the above estimate of the variance of the mean number of empty timeslots per cycle may be defined to be a mean value such that

$$\hat{\mu}_A = \overline{X}(n) - [\Phi^{-1}(1-\alpha)] \cdot \sqrt{\hat{Var}[\overline{X}(n)]} \tag{4.1}$$

The time series of the number of empty timeslots per cycle obtained through simulation of the WRR from the same scenarios described in Section 4.1.3 were used to calculate a sample mean and estimate of the variance of the sample mean for each scenario, resulting in an estimate of $\hat{\mu}_A$. In every case $\hat{\mu}_A$ was less than the actual mean of the scenario for all α between $2\text{x}10^{-1}$ and $1\text{x}10^{-10}$. A typical example is shown in Figure 4-3.

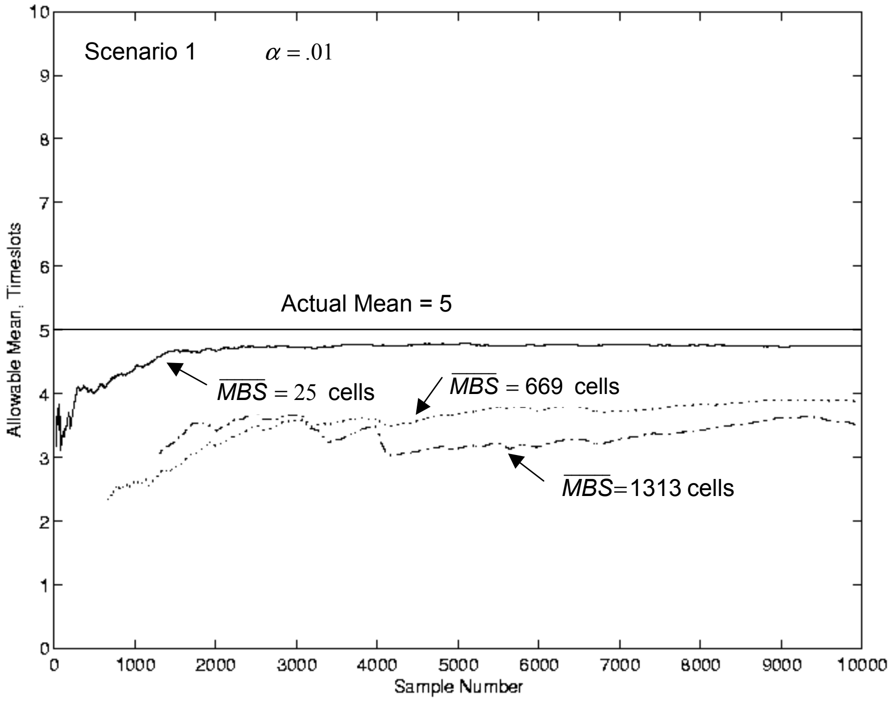

Figure 4-3. Typical Estimated Allowable Mean Over 10,000 WRR Schedule Cycles

Another aspect of the use of an estimate of the variance of the sample mean is its relationship to the number of traffic streams being served. As discussed in Chapter 3, the normalized variance of X will remain constant (in the binomial case) regardless of the number of traffic streams, assuming cell arrivals in each stream are independent. However, VBR traffic is by definition bursty, being characterized by a distribution having a mean burst size and maximum burst size. Assuming independence between traffic streams results in varying amounts of "overlap" in the on- and off-periods of the traffic streams, rather than the worst case scenario of full overlap, as discussed in Section 3.4.3. If the burst arrivals in each stream are independent, the amount of overlap with other traffic streams is random, and the variance in the mean of X will become smaller as the number of traffic queues increases. This will in turn result in a decrease in the sample variance of *X*, leading to a decrease in the estimated variance of the mean of *X*. Thus, based on Equation (4.1), an increase in the number of traffic streams will lead to a more rapid convergence of the allowable mean to the actual mean as n increases. To test this hypothesis, the results of estimating an allowable mean of ten- and twenty-queue scenarios were examined, with the SCR for all traffic streams entering the WRR queues being equal, i.e., equal to 1/10 and 1/20 of the virtual path rate, respectively. The results showed a more rapid convergence of the allowable mean to the true mean in the scenario with the larger number of queues, as shown by comparing Figure 4-5 and Figure 4-5.

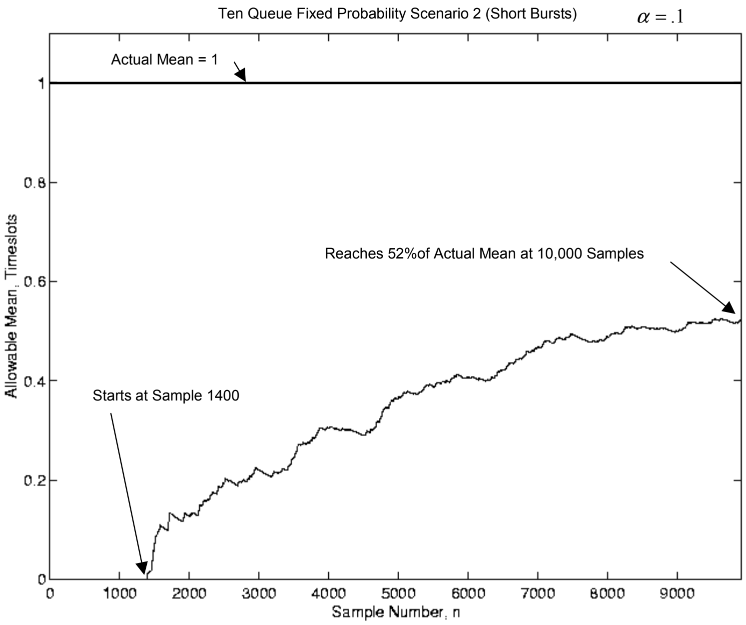

Figure 4-4. Allowable Mean from 10 Queue Scenario

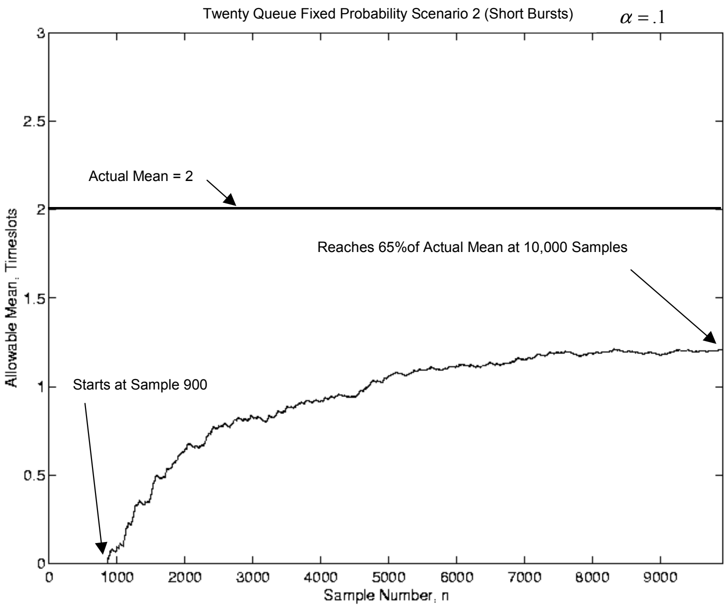

Figure 4-5. Allowable Mean from 20 Queue Scenario

The results shown in Figure 4-3 and the results of comparing Figure 4-4 and Figure 4-5 show that the allowable mean developed through the technique proposed here is an effective means of estimating the true mean number of empty timeslots per WRR cycle. The estimate rapidly converges to the true mean when existing traffic has short burst characteristics, taking longer to converge when the existing traffic has long burst characteristics. Thus, as the number of served queues increases, improvements in the time of convergence results.

4.1.5 Introduction of "Additional" Traffic

In order for the proposed method of connection admittance to be effective, traffic at rates determined from the multiple and single empty timeslot methods of Chapter 3 must be supported by the virtual path for each value of μ_X, the mean number of empty timeslots per cycle. To determine if this is the case, the path model of a “full” virtual path is modified by the addition of a traffic stream and queue having definable traffic parameters and queue size parameters. To insure validity across the range of possible mean values and the corresponding “added” connection sustainable cell rates, the fixed probability scenarios listed in Appendix B2 were used.

Since the queue sizes associated with the “added” connections, as discussed in Section 3.4.3, are based on a worst case analysis, the initial simulations were run first with “added” queue lengths of MQL, as determined by Equation (3.5), rather than the augmented queue lengths, MQL’, specified in Section 3.4.3. This allows for the development of a baseline where there are no augmented “added” queue sizes for use in future evaluations of the need for specific “added” queue sizes larger than MQL. It also allows the total number of required simulations to be reduced, since the results of initial simulations will indicate areas of mean and burst size parameters where additional simulation are actually needed.

Simulations were run such that within each “existing” traffic scenario, taken from Appendix B2, having various SCR and short, mixed, or long burst size characteristics, “added” traffic was introduced through a queue sized at either MQL or MQL’, as discussed above. Simulations were run 20 times each, using a different random seed, for each “added” traffic characteristic type within each SCR tested. The “added” traffic characteristic types used mean burst characteristics of 2, 10, and 100 for each set of parameters within the empty queue probability parameters of .1, .5, and .9

shown in Appendix B2. The simulations were run such that the “added” traffic SCR was greater than the predicted allowable “added” traffic SCR, and was increased on successive runs until the allowed CLR (10^{-3} in this case) was exceeded. In cases where the CLR was exceeded at the first value tested, the SCR was reduced on successive simulation runs until the CLR was below that required, or until the SCR was below the allowable “added” SCR for that “existing” traffic scenario. A CLR threshold was defined such that it existed at the boundary between two added traffic SCR values, the higher value being where the CLR was exceeded, and the lower one where the CLR was not exceeded. A typical graph resulting from this procedure is shown at Figure 4-6, with detailed results being shown in Appendices C, D, and E.

Results of initial simulations of 10 queue, 10 timeslot fixed probability scenarios 1, 2, and 3, for short, mixed, and long burst existing traffic types taken from Appendix B2 are shown in Appendix C, Figures C-1 through C-9. The results show the cell loss ratios were better than required for all “existing” traffic scenarios and burst types other than Scenario 2, where the normalized mean number of empty timeslots per cycle is .1, with existing traffic having long burst characteristics, as seen in Figure C-6. This characteristic, where CLR did not meet requirements for low mean, long burst existing traffic scenarios, was apparent in other scenario sequences, as well. This is illustrated by Scenario 5, Figures C-10 through C-12, where the long burst scenario causes excessive CLR at additional traffic rates that should be acceptable, as seen in Figure C-12. This was also true in the case of Scenario 8, long burst existing traffic, as illustrated by Figure C-13.

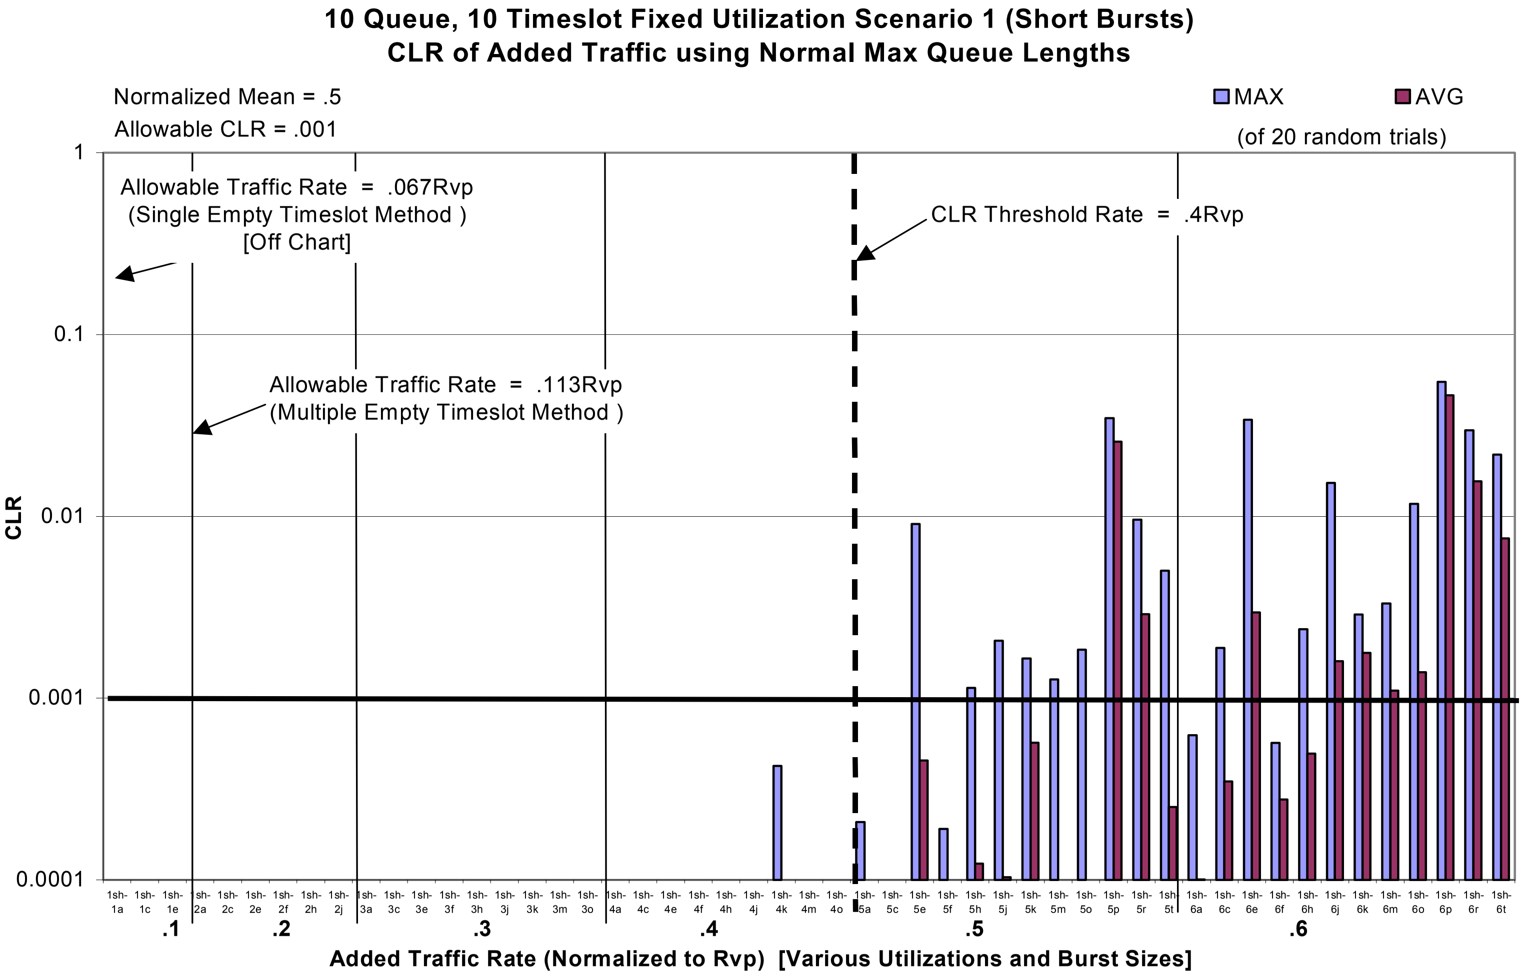

Figure 4-6. Typical “Added” Queue Simulation Results

The initial simulation runs described above achieved their objective of determining the traffic characteristics that would cause “additional” traffic connection queue overflows in queues sized by Equation 3.5, thus indicating which scenarios should be run with the augmented “added” queue lengths required by Equation (3.25). Scenarios 2, 5, and 8 were then run with the augmented “added” queue lengths required by section 3.4.3, with the results shown in Appendix D, Figures D-1 through D-3. As illustrated in these figures, all results of simulations using augmented “added” queues met the required CLR.

To further evaluate the model, 20-queue models were tested, using the scenarios shown in Appendix B3. Results from 20-queue, 20-timeslot Scenario 2, for short, mixed, and long existing traffic burst characteristics are shown at Appendix E, Figures E-1 through E-3. All 20-queue models tested met the required CLR without the need for augmented queues on “added” traffic connections. The implication of the results obtained above will be discussed in the next section.

4.2 Discussion of Results

The results obtained through the simulations described in Section 4.1 show the following:

- For the traffic type listed in Appendix B1, a queue-skipped-when-empty WRR server has the same mean (see Table 4-1) as a rigidly scheduled timeslot server, which was the basis of the analysis conducted in Chapter 3.

- The prediction of the variance of the sample mean number of empty timeslots per WRR cycle is a conservative estimate of the actual variance, as shown in Table 4-2. In addition, Figure 4-3 shows that estimates of an allowable mean, based on measuring the sample mean and estimating the sample variance, may be used as the parameter for admitting "additional" traffic, since such estimates are less than the actual mean of the existing traffic.

- "Additional" traffic admitted to a virtual path based on the actual mean of the existing traffic will not cause required CLR to be exceeded as long as queues are sized as required by the analysis given in Chapter 3.

- Traffic admitted based on the actual mean number of empty timeslots per cycle will not exceed the required CLR as long as the "added" traffic rate does not exceed the rate indicated by the single or multiple empty timeslot method described in Chapter 3.

4.3 Summary

In this chapter, the analysis developed in Chapter 3 was tested through simulation of traffic and of the WRR mechanism serving the end-to-end virtual path of the overall model framework. The simulations, using models of on-off traffic sources and WRR server mechanisms of up to twenty queues and twenty cell timeslots, were conducted over a wide range of traffic scenarios. The scenarios span a range of traffic conditions for both existing and "added" traffic connections designed to identify ranges of "added" traffic which could cause queue overflows, and the conditions under which the "added" traffic connection queue sizes are required to be augmented above the size needed in the "normal" traffic connections for similar traffic.

Simulations were also designed to validate the use of an allowable mean number of empty timeslots per WRR cycle as the basis of an allowable "additional" traffic sustainable cell rate (SCR). The allowable "additional" SCR is based on measurement of samples of the mean number of empty timeslots per WRR cycle and an estimate of the variance of the mean number of timeslots per WRR cycle, which is shown to be dependent on the correlation of the existing traffic stream in the virtual path.

The results of 15,374 simulations, spanning a significant portion of the design space, conducted over a duration of at least 100,000 virtual path timeslots each, support the analysis presented in Chapter 3. The next chapter will present conclusions, including a detailed discussion of the CAC algorithm, and possible implementation methods.

Chapter 5

Conclusions and Future Directions

CHAPTER CONTENTS

LIST OF FIGURES

Chapter

5

Conclusions and Future Directions

This final chapter is used to draw conclusions regarding the proposed methods of determining available bandwidth and associated connection admission control, which have the potential of increasing the amount of traffic carried in a virtual path by up to 20 to 60 percent, depending on existing traffic characteristics and desired CLR. Issues resulting from the use of this method, and possible directions for further research in this area, are discussed.

5.1 Conclusions

Conclusions regarding the validity of the proposed technique are presented, and results compared to those reported for other methods of connection admission control (CAC). Implementation methods are discussed, and specific algorithms resulting from this method are presented in detail.

5.1.1 Validity of the Proposed Technique

The proposed method for determining bandwidth available on end-to-end virtual paths having class of service separation has been shown by analysis to enable the addition of traffic to virtual paths previously considered to be full. The analysis applies to all possible normalized mean number of empty timeslots per WRR cycle, μ_{XN}, and is supported by the results of the simulations.

The simulations showed that queues sized in accordance with Equation (3.5) were able to support the traffic in scenarios where the mean number of empty timeslots per WRR cycle, μ_X, was not at the low end of its possible range. In scenarios where μ_X is at the low end of its possible range, "added" traffic queues augmented as described in Section 3.4.3 are sufficient to insure that the required QoS is met. Thus, queues supporting "added" traffic, sized in accordance with Equation (3.25), will support the allowed "added" traffic over the entire range of mean values, which supports the results of the analysis.

In addressing the use of the proposed method as the number of queues being served increases, simulations having 20 queues are compared to simulations having 10 queues. The differences seen by comparing Figures C-4 through C-6 with Figures E-1 through E-3 show that the CLR performance of the system improves as the number of queues increases, even without augmenting the queues serving the "added" traffic. This occurs because, as the number of queues supporting VBR traffic increases, the probability that no timeslots are available in a given cycle decreases, assuming that the existing traffic streams are independent. The amount of augmentation actually needed in an "added" traffic queue as the number of "normal" traffic streams increases is thus reduced, since the augmented queue size for an "added" queue was based on a worst case where the traffic streams were not independent and always overlapped (see Section 3.4.3). This result indicates that future research into the queue size augmentation needed for "added" traffic queues would be beneficial in reducing the amount of "added' queue augmentation needed to support the QoS requirements of "added" traffic.

Scenarios having the mean number of empty timeslots per cycle, μ_X, at the higher end of its possible range, as illustrated by Figures C-1 through C-3 and C-7 through C-9, illustrate that the proposed method becomes more conservative as μ_X increases. This conclusion is supported since required QoS in these scenarios is maintained even at "added" traffic rates much higher than predicted by the analysis. In addition, the method of estimating an allowable mean as the WRR operates is also conservative in that the estimate of the mean is less than, or, at an infinite number of WRR cycles, equal to the true mean, as discussed in Section 4.1.

Thus, the connection admission control method itself will be conservative, admitting "additional" connections at a rate that is less than could actually be supported by the virtual path. The simulation results support the analysis in that no cells in excess of those established by the desired CLR parameter are lost in the "added" queues, and no cells at all are lost in the "existing" queues of a "full" WRR-served virtual path. Thus, being both conservative in determining available bandwidth, and capable of supporting a required QoS, the proposed method of determining available bandwidth appears to be a useful method of CAC in broadband networks. The use of the method developed here can allow more efficient use of the available bandwidth in an end-to-end path, by at least the percentages by a factor of up to 20% to 60%, depending on desired CLR, as illustrated by Figures 3-5 and 3-6, thus improving overall network efficiency.

5.1.2 Comparison with Other Methods.

The method developed, analyzed, and evaluated here can meet quality of service (QoS) objectives, in terms of CLR, cell delay, and CDV, for CBR and VBR traffic connections. These include the traffic connections admitted to the virtual path up until the path is considered “full” based on the sum of traffic stream SCR parameters, as well as those “added” traffic connections admitted to the “full” path using the mean number of empty timeslots per cycle versus allowable “additional” cell rate criteria developed in Chapter 3. Although this research primarily addresses the CLR aspects of QoS, maximum cell delay and CDV can be accounted for in a CAC through consideration of the queue size required by the method, as determined by Equations (3.5) and (3.25).

In the bandwidth evaluation and CAC method developed here, all traffic conforming to the traffic contract requested by the network user would meet its CLR objectives. Other methods, based on multiple traffic streams entering a single buffer, have been extensively investigated, however, comparison with other methods is difficult, since results from many investigations are reported across a wide range of possible parameter formats. Some of these formats for reporting results include:

- CLR vs. buffer size for a given traffic load [CH096] [PET96][ZUK97] [ZHA94]
- Maximum connections of a given type and size vs. cell loss probability [PER96]
- Maximum number of connections versus buffer size for a given CLR [PER96]
- Admission regions in terms of quantities of sources from two source classes vs. loss probability [ELW95] [MAN96] [GEL97] [PER96]
- Percentage of blocked calls vs. traffic load for various allocation strategies [BOL97]
- Cell loss probability vs. call arrival rate for various peak bit rates [SAI91]

In comparing various existing methods of CAC that are based on evaluating available bandwidth or evaluating cell loss probabilities, the equivalent capacity method was shown by Perros and Elsayed [PER96] to be superior to heavy traffic approximations and upper bound on cell loss probability methods for CLR = 10^{-6} and two traffic classes having differing peak bit rate, mean bit rate, and mean burst length. The use of diffusion-based techniques has been shown by Golestani [GEL97] to be conservative with respect to cell loss, but more economical in bandwidth allocation, leading to larger admission regions for both homogenous and heterogeneous traffic. Techniques based on the large duration rate function (entropy) of bursty ATM traffic

[CRO97] through measurement of traffic activity have been initially evaluated as having 5% of connections experiencing CLR in excess of their QoS requirements, while admitted bandwidth was approximately 10% higher than the bandwidth allocated by an equivalent capacity approximation. The method proposed by this research allows the admission of bandwidth at varying percentages of up to 20 % higher (at a CLR of 10^{-10}) than that allowed by equivalent capacity techniques (see Figures 3.5 and 3.6), while assuring the QoS of all connections.

The use of other measurement-based strategies for CAC, such as that proposed by Liu, et al., [LIU97-1], have been shown to be potentially effective in efficiently allocating buffer resources, at a cost of requiring re-negotiation of source traffic parameters once the true characteristics of a traffic source is determined. However, in [LIU97-1] it was also recommended that since this technique is risky, no more than 90% of physical VP bandwidth should be made available for allocation by this procedure. The methods proposed by this research have the potential to utilize an actual 100% of each physical VP while maintaining QoS on all connections.

Several authors have noted that the equivalent capacity approach to bandwidth allocation has some limitations when the source does not conform to exponentially distributed packet lengths and silence times [BRA94], and is therefore not always a conservative estimate of available bandwidth [CHO96]. Although the method proposed by this dissertation has used equivalent capacity calculation to determine a value for SCR, any method which determines a SCR in which the desired CLR is not exceeded based on a given peak cell rate, mean burst period, queue length and utilization is usable. Should the method chosen be too conservative, i.e., cells to not arrive often enough for the buffers and service rate allocated, then measurements of the mean number of empty timeslots in the virtual path, as proposed by the methods outlined in this paper, will detect this and allow the introduction of additional traffic. Should the method chosen not be conservative, i.e., cells arrive too often for the buffers and service rate allocated, then the GCRA "policer" will mark cells for possible discard. Thus, the CAC method proposed by this research is useful over a wide range of possible methods for determining SCR, and in fact for any method which accurately creates an SCR that guarantees QoS to the satisfaction of the network user based on the peak cell rate, utilization, mean burst period, and desired CLR of his traffic.

5.1.3 Implementation Methods

The development of the two methods of equating a mean number of empty timeslots per WRR cycle to an allowable "additional" traffic connection SCR, as described in Chapter 3, leads to the need to choose between them when implementing a system. The choice will be based on the computing resources available for implementing the methods in a given hardware and software design. Note that the single empty timeslot method is less computationally intensive, while the multiple empty timeslot method allows the admittance of a larger "added" traffic stream for a given measured mean at the cost of more computational complexity.

Another design choice is which of two possible methods to use in making the CAC decision. One method is to make the decision based on test of a hypothesis that $\mu_X = \mu_D$ at the time of admission, as discussed in Section 3.5.3. A second method is possible through tracking the lower bound of a confidence interval as it changes, thus leading to an ongoing estimate of μ_A that forms the basis for a decision, as discussed in Section 4.1.4. Both methods are based on the same statistic and require continuous evaluation of the mean of the number of empty timeslots per WRR cycle, X, evaluated over the period since the last new connection was admitted. The first method evaluates the statistic at the time the decision is needed, while the other method computes the statistic at each WRR cycle, keeping an ongoing estimate of μ_A immediately available to make the admission decision if required at that time. Thus, at the time of an admission request, the second method can make the CAC decision faster, at the cost of more intensive cycle-by-cycle computing load. The following section describes a detailed CAC algorithm, based on determination of probable bandwidth using an ongoing estimate of μ_A.

5.1.4 CAC Algorithm

Details of a CAC algorithm based on the ongoing evaluation of $\hat{\mu}_A$ are shown in Figure 5-1. The algorithm first determines whether there is sufficient bandwidth in the virtual path to allow the new connection to be admitted “normally”, i.e., where the sum of the SCR of all traffic is less than the virtual path rate, R_{VP}. If so, the

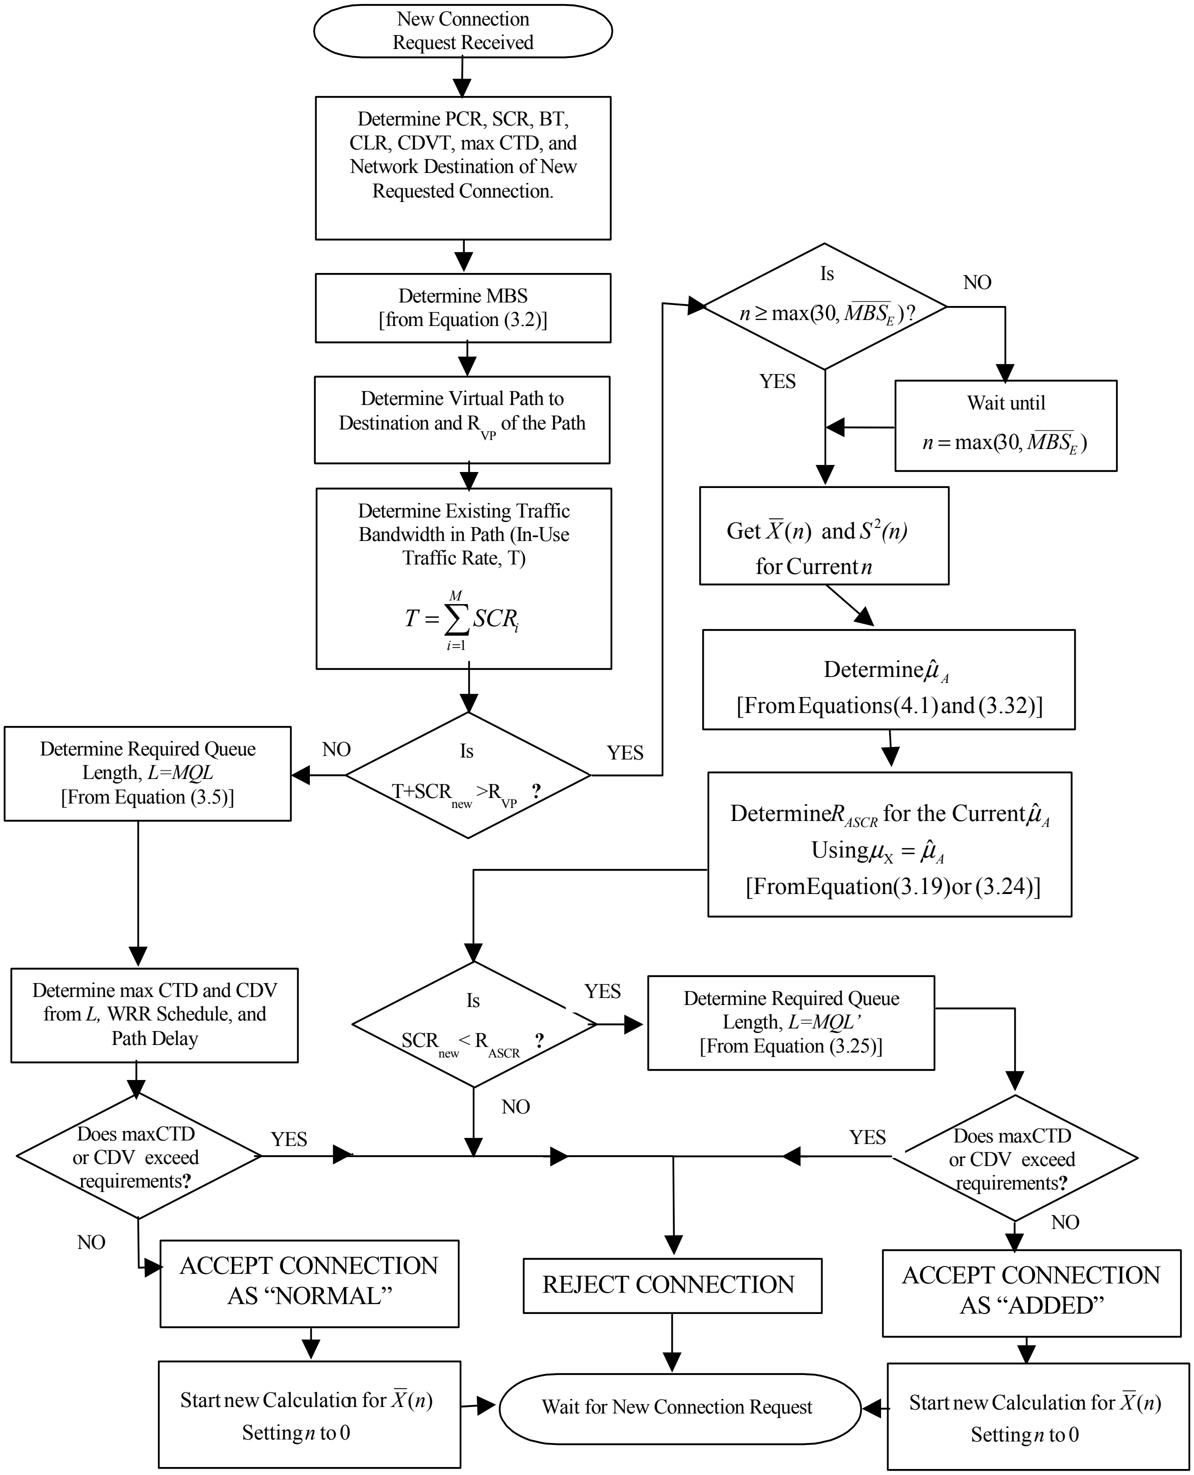

Figure 5-1. CAC Algorithm

algorithm goes on to determine whether or not the cell delay and delay variation QoS criteria are met. If not, the algorithm continues on to evaluate the available bandwidth based on the current value of $\hat{\mu}_A$. This evaluation depends on the current values of the mean and variance of the mean number of empty timeslots per cycle, as shown in Section 3.5.

In support of this function of the algorithm, the switch performing CAC must maintain an ongoing calculation of $\overline{X}(n)$, the sample mean number of empty timeslots per cycle at *n* WWR cycles, and $S^2(n)$, the sample variance of the sample mean (from Equation (3.28)), which are used in determining the estimated variance of the mean (from Equation (3.32)).

Once a current estimate of the mean is determined, it may be used to determine the allowable cell rate of “additional” connections, using the relationships determined for the single and multiple empty timeslot methods (from Equation (3.19) or (3.24)). If the required delay and delay variation criteria can be met in view of the queue sizes needed, then the connection is admitted.

5.1.5 General

The results of this research are directly applicable to integrated services broadband networks carrying large volumes of traffic, in particular the types of variable bit rate traffic characterized by real time video and data. Such traffic will be carried by the high speed ATM based networks presently being deployed worldwide. This research will allow higher utilization of network resources allowing these networks to carry additional traffic per physical link compared to techniques presently in use being proposed. This advantage, coupled with the relative simplicity of a CAC requiring knowledge of the characteristics of only one node, the entry node, on an end-to-end path, makes the method developed here a serious contender for implementation in future broadband networks.[1]

5.2 Future Directions

In this section, issues concerning the use of the method proposed by this research are discussed, as well as potential topics of further research.

5.2.1 Issues

The primary issue resulting from this investigation of increasing bandwidth utilization through an improved CAC results from consideration of how the proposed technique will be used. The proposed method of connection admission control should be targeted on networks that use switched virtual circuits (SVC) rather than permanent virtual circuits (PVC). The reason for this is based on the possible usage habits of network customers. Users of PVC may or may not be fully utilizing the channel at all times. Since this proposed method is dependent on measurements of empty cell timeslots, an erratic usage pattern of PVC connections could result in erroneous measurements. The actual effect of this type of user behavior depends on

---

[1] The methods described in this paper are the subject of a current patent application.

what percentage of the total virtual path is occupied by the effected traffic stream. Since the proposed method is very conservative, if there are a large number of PVC in a path, this type of behavior on the part of a small number of users will probably not have an unacceptable effect on other users. However, the potential for adverse effect on QoS exists because of this possible behavior by network users where PVC is involved.

On the other hand, it is unlikely that users of SVC capable networks would establish a connection and then not use it to the extent requested by the traffic contract, since the connection can easily be terminated and re-established when again required, or re-established under a different traffic contract should required connection characteristics change. The cost to network users of establishing a connection and then purposely not using it, or under-using it, should preclude this behavior [DAS98], or at least confine it to a small number of users. Assuming the worst case, that a significant number of users under-utilize their connections for some period, during which "added" connections are introduced to the network, and subsequently start using them, using this method the only connections affected will be the "added" connection. The connections that existed prior to the introduction of the "added" connection are all guaranteed to have the required QoS for conforming cells. Also, "added" connections will eventually be placed into the "normal" WRR schedule cycle, as existing "normal" switched traffic connections are concluded, thereby causing "added" connections with QoS affected by this particular circumstance (if any) to be returned to conformance after some period. Since QoS performance is required over a long period, and short term degradation of QoS is allowable, this method should be considered as conforming to applicable broadband network requirements.

Other issues that may be addressed by further research could involve possible modifications to this method of improving network bandwidth utilization, based on making different choices regarding worst case scenarios. Extensions of this method into other aspects of improving network bandwidth utilization are also possible.

5.2.2 Further Research

Further improvements to this model for determining available bandwidth and accomplishing CAC are possible, and could be the subject of further research. For example, the method of determining the size of the queue on an "added" connection was based on a worst case that all existing "normal" sources in a "full" virtual path are transmitting the maximum burst size allowed by their traffic contract, followed by the shortest allowed off period. The worst case further assumes that the existing sources were synchronized such that the on periods always overlapped. Since in real networks the traffic streams from various customers are effectively independent from each other, this worst case is extremely unlikely. Additional research in this area could result in improvements to the model where smaller "added" connection queue sizes are possible.

This research focuses on the admission of connections at the channel level in networks having a channel and path architecture. These results could be extended to the dynamic control of path bandwidth, i.e., adaptive control of the allocation of path bandwidth and buffer resources, based on feedback from a weighted round robin queue scheduler. This in turn would speed the improvement of overall network resource utilization in the future.

References

[BOL97] R. Bolla, F. Davoli, and M. Marchese, “Bandwidth Allocation and Admission Control in ATM Networks with Service Separation”, *IEEE Communications Magazine*, May 1997.

[BON96] ALTA Group^{TM} of Cadence Design Systems, Inc., *BONeS$^{®}$ DESIGNER™*, Product Number DES 8004, March 1996.

[BRA94] C. Braun, D. Petr, and T. Kelley, “Performance Evaluation of Equivalent Capacity for Admission Control”, *IEEE Wichita Conference on Communications, Networking and Signal Processing*, April 1994, pp. 155-158.

[CHE94] T. M. Chen and S. S. Liu. “Management and Control Functions in ATM Switching Systems,” *IEEE Network*, July/August 1994, pp. 27-40.

[CHO96] G. L. Choudhury, D. M. Lucantoni, and W. Whitt, “Squeezing the most out of ATM,” *IEEE Transactions on Communications*, Vol 44, Number 2, February 1996, pp.203-217.

[COU95] G. Courcoubetis, et. al, “Admission Control and Routing in ATM Networks using Inferences from Measured Buffer Occupancy,” *IEEE Transactions on Communications*, Volume 43, Numbers, 2, 3, 4, February, March, April 1995.

[CRO97] S. Crosby, I. Leslie, J. T. Lewis, R. Russell, F. Toomey, and B.McGurk, “Practical Connection Admission Control for ATM Networks Based on On-Line Measurements”, *IEEE Proceedings of ATM’97*, Lisbon, June 1997.

[DAS98] L. DaSilva, D. Petr, N. Akar, “Effects of Pricing in Multiple Service Networks with Resource Allocation”, *Proc of Conf. on Performance, and Control of Network Systems II, SPIE Int’l Symposium on Voice, Video, and Data,* November, 1998.

[DEV82] J.L.Devore, *Probability and Statistics for Engineering and the Sciences*, Brooks/Cole Publishing Company, Monterey, CA, 1982.

[DOU97] C. Douligeris and G. Develekos, “Neuro-Fuzzy Control in ATM Networks”, *IEEE Communications Magazine*, May 1997.

[DUF95] N. G. Duffield, J. T. Lewis, N. O’Connell, R.Russell, and F. Toomey, “Entropy of ATM traffic Streams: A Tool for Estimating QoS Parameters”, *IEEE Journal on Selected Areas in Communications*, Volume 13, Number 6, August, 1995.

[ELW95] A. Elwalid, D. Mitra, “Analysis, Approximations, and Admission Control of a Multi-Service Multiplexing Scheme with Priorities”, *Proceedings of the IEEE Infocom’95*, 1995, pp243-472.

[FRO92] V. Frost and M. Mullen, “Bandwidth Allocation for B-ISDN based on End-to-End Delay Estimates”, *Proceedings of IEEE ICC’92*, June 1992.

[GEL97] E. Gelenbe, X. Mang, and R. Onvural, “Bandwidth Allocation and Call Admission Control in High Speed Networks”, *IEEE Communications Magazine*, May 1997.

[GOL91] S.J.Golestani, “Congestion-Free Communications in High-Speed Packet Networks”, *IEEE Transactions on Communications*, Vol 39, December 1991.

[GUE91] R. Guerin, H. Ahmadi, and M. Naghshineh, “Equivalent Capacity and Its Application to Bandwidth Allocation in High-Speed Networks”, *IEEE Journal on Selected Areas in Communications*, Volume 9, Number 7, September 1991.

[HAN94] ______, *Handbook of Mathematical, Scientific, and Engineering Formulas, Tables, Functions, Graphs, Transforms*, Research and Education Association, Piscataway, New Jersey, 1994.

[HOG95] R.V.Hogg and A.T.Craig, *Introduction to Mathematical Statistics,* Prentice Hall, Englewood Cliffs, New Jersey, 1995.

[HLU88] M. G. Hluchyj and M. J. Karol, “Queueing in High-Performance Packet Switching”, *IEEE Journal on Selected Areas in Communications*, Volume 6, Number 9, December 1988.

[KES96] G. Kesidis, *ATM Network Performance*, Kluwer Academic Publishers, Boston, 1996.

[LAW91] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis*, McGraw Hill, Inc., New York, 1991.

[LIU97-1] K. Liu, D. Petr, and C. Braun “A Measurement-Based CAC Strategy for ATM Networks”, *Proceedings of IEEE ICC’97*, 1997.

[LIU97-2] K. Liu, D. Petr, V. Frost, et. al., “Design and Analysis of a Bandwidth Management Framework for ATM Based Broadband ISDN”, *IEEE Communications Magazine*, May 1997, pp. 138-145.

[MAN96] X. Mang and E.Gelenbe, “Call Admission Control in ATM Using the Diffusion Model”, *Proceedings of the IEEE Globecom’96*, 1996, pp 1700-1704.

[MEZ94] K. Mezger, D. W. Petr, and T. Kelley, “Weighted Fair Queuing vs. Weighted Round Robin: A Comparative Analysis”, *IEEE Wichita Conference on Communications, Networking and Signal Processing*, April 1994.

[PET96] D. W. Petr, V. F. Frost, T. Kelley, et. al., “Cell Loss Quality of Service in an Integrated Traffic ATM Network”, *Int’l Journal of Comm Systems*, Volume 9, Number 2, March-April 1996, pp 97-104.

[PER96] H. Perros and K. Elsayed, “Call Admission Control Schemes: A Review,” *IEEE Communications Magazine*, November 1996, pp. 82-91.

[PRY96] M. de Pryker, *Asynchronous Transfer Mode: Solution for Broadband ISDN, Third Edition*, Prentice Hall, 1996.

[RIC68] S. B. Richmond, *Operations Research for Management Decisions*, The Ronald Press Company, New York, 1968.

[ROB94] T.G.Robertazzi, *Computer Networks and Systems*, Springer-Verlag, 1994, p. 77

[SAI91] H. Saito and K. Shiomoto, “Dynamic Call Admission Control in ATM Networks,” *IEEE Journal on Selected Areas in Communications*, Volume 9, Number 7, September 1991.

[SAI92] H. Saito, “Call Admission Control in an ATM Network Using Upper Bound of Cell Loss Probability,” *IEEE Transaction on Communications*, Volume 40, Number 9, September 1992.

[SAI94] H. Saito, *Teletraffic Technologies in ATM Networks*, Artech House, 1994, pp. 77-98.

[SAI97] H. Saito, “Dynamic Resource Allocation in ATM Networks”, *IEEE Communications Magazine*, May 1997.

[SCH96] M. Schwartz, *Broadband Integrated Networks,* Prentice Hall, Inc., 1996.

[TMS4] *Traffic Management Specification version 4.0*, The ATM Forum, April 1996.

[TUR92] J. Turner, “Managing Bandwidth in ATM Networks with Bursty Traffic”, *IEEE Network*, September 1992.

[UNI3.1], *User-Network Interface (UNI) Specification version 3.1*, The ATM Forum, 1994

[VEC95] G. de Veciana, G. Kesidis, and J. Walrand, “Resource Management in Wide-Area ATM Networks Using Effective Bandwidths,” *IEEE Journal on Selected Areas of Communications*, Volume 13, Number 6, August 1995, pp1081-1090

[WAN94] Y.-T. Wang, T.P. Lin, K.C. Gan, “ An Improved Scheduling Algorithm for Weighted Round-Robin Cell Multiplexing in an ATM Switch”, *Proceedings of the IEEE ICC’94,* 1994, pp. 1032-1037

[WEI95] A. Weiss, “An Introduction to Large Deviations for Communications Networks”, *IEEE Journal on Selected Areas in Communications*, Volume 13, Number 6, August 1995.

References

[ZHA94] Z. Zhang and A.S. Acampora, “Equivalent Bandwidth for Heterogeneous Sources in ATM Networks,” *Proceedings of the IEEE ICC’94*, 1994, pp 1025-1031.

[ZUK97] M. Zukerman and P.W.Tse, “An Adaptive Connection Control Scheme for ATM Networks”, *Proceedings of the IEEE ICC’97*, June 1997.

Appendix A

VBR Traffic Types

This appendix contains two lists of VBR traffic types illustrating the results of simultaneously solving Equations 3.8 and 3.9 for, maximum burst size, MBS, and queue length, *L,* when the desired CLR, ε ,is given and the peak cell rate, PCR, sustained cell rate, SCR, mean burst size, *B*, and utilization (offered load), ρ , of the traffic are specified. In an actual network scenario, the equations would be solved simultaneously for SCR and *L* by the user, with PCR, *B*, MBS, desired CLR, and ρ being known to, or estimated by, the user. The PCR, SCR and MBS parameters provided by the user to the service provider would allow the establishment of a queue of the required size *L*, as described in Chapter 3.

The lists also show probabilities of empty queues, P_{Q0}, calculated from $P_{Q0} = 1 - \lambda_Q / SCR_Q$, as described in Chapter 3, where λ, the average cell arrival rate is known from $\lambda = \rho \cdot PCR$ for the on-off traffic types used.

The two lists differ in that in List A1 the utilization values are fixed at $\rho = [.05,.1,.2,.4,.8]$, leading to differing probabilities P_0, while in List A2, the probabilities are fixed at $P_0 = [.1,.5,.9]$, leading to the differing utilization values required to obtain the required probabilities of empty queue, P_0

The last column in the two lists is a shorthand notation used to describe the type of traffic in charts and graphs.

Appendix A1 Call Types (Fixed Utilization) CLR: 1e-3 .
PCR: 1000cells/sec .

SCR Ratio	SCR	Utilization	B	MBS	L	P0	Type ID
0.1	100	0.05	2	25	24	0.5	1 a
0.1	100	0.05	5	65	60	0.5	1 b
0.1	100	0.05	10	131	119	0.5	1 c
0.1	100	0.05	50	656	591	0.5	1 d
0.1	100	0.05	100	1313	1182	0.5	1 e
0.2	200	0.05	2	16	14	0.75	2 a
0.2	200	0.05	5	42	35	0.75	2 b
0.2	200	0.05	10	86	70	0.75	2 c
0.2	200	0.05	50	436	350	0.75	2 d
0.2	200	0.05	100	874	700	0.75	2 e
0.2	200	0.1	2	24	20	0.5	2 f
0.2	200	0.1	5	61	50	0.5	2 g
0.2	200	0.1	10	124	100	0.5	2 h
0.2	200	0.1	50	622	498	0.5	2 I
0.2	200	0.1	100	1243	995	0.5	2 j
0.3	300	0.05	2	15	12	0.83	3 a
0.3	300	0.05	5	38	28	0.83	3 b
0.3	300	0.05	10	78	56	0.83	3 c
0.3	300	0.05	50	392	276	0.83	3 d
0.3	300	0.05	100	787	552	0.83	3 e
0.3	300	0.1	2	18	14	0.67	3 f
0.3	300	0.1	5	45	33	0.67	3 g
0.3	300	0.1	10	92	66	0.67	3 h
0.3	300	0.1	50	466	327	0.67	3 i
0.3	300	0.1	100	932	653	0.67	3 j
0.3	300	0.2	2	32	24	0.33	3 k
0.3	300	0.2	5	82	59	0.33	3 l
0.3	300	0.2	10	165	117	0.33	3 m
0.3	300	0.2	50	828	581	0.33	3 n
0.3	300	0.2	100	1657	1161	0.33	3 o
0.4	400	0.05	2	13	9	0.88	4 a
0.4	400	0.05	5	36	23	0.88	4 b
0.4	400	0.05	10	73	45	0.88	4 c
0.4	400	0.05	50	373	225	0.88	4 d
0.4	400	0.05	100	748	450	0.88	4 e
0.4	400	0.1	2	14	10	0.75	4 f
0.4	400	0.1	5	39	25	0.75	4 g
0.4	400	0.1	10	81	50	0.75	4 h
0.4	400	0.1	50	413	249	0.75	4 i
0.4	400	0.1	100	828	498	0.75	4 j
0.4	400	0.2	2	21	14	0.5	4 k
0.4	400	0.2	5	54	34	0.5	4 l
0.4	400	0.2	10	109	67	0.5	4 m
0.4	400	0.2	50	552	332	0.5	4 n
0.4	400	0.2	100	1105	664	0.5	4 o
0.5	500	0.05	2	13	8	0.9	5 a
0.5	500	0.05	5	35	19	0.9	5 b
0.5	500	0.05	10	71	37	0.9	5 c
0.5	500	0.05	50	364	183	0.9	5 d

SCR Ratio	SCR	Utilization	B	MBS	L	P0	Type ID
0.5	500	0.05	100	728	365	0.9	5 e
0.5	500	0.1	2	13	8	0.8	5 f
0.5	500	0.1	5	37	20	0.8	5 g
0.5	500	0.1	10	75	39	0.8	5 h
0.5	500	0.1	50	388	195	0.8	5 i
0.5	500	0.1	100	776	389	0.8	5 j
0.5	500	0.2	2	17	10	0.6	5 k
0.5	500	0.2	5	45	24	0.6	5 l
0.5	500	0.2	10	91	47	0.6	5 m
0.5	500	0.2	50	460	231	0.6	5 n
0.5	500	0.2	100	920	461	0.6	5 o
0.5	500	0.4	2	39	21	0.2	5 p
0.5	500	0.4	5	101	52	0.2	5 q
0.5	500	0.4	10	205	104	0.2	5 r
0.5	500	0.4	50	1036	519	0.2	5 s
0.5	500	0.4	100	2072	1037	0.2	5 t
0.6	600	0.05	2	11	6	0.92	6 a
0.6	600	0.05	5	34	15	0.92	6 b
0.6	600	0.05	10	69	29	0.92	6 c
0.6	600	0.05	50	356	144	0.92	6 d
0.6	600	0.05	100	714	287	0.92	6 e
0.6	600	0.1	2	11	6	0.83	6 f
0.6	600	0.1	5	34	15	0.83	6 g
0.6	600	0.1	10	71	30	0.83	6 h
0.6	600	0.1	50	372	150	0.83	6 i
0.6	600	0.1	100	744	299	0.83	6 j
0.6	600	0.2	2	14	7	0.67	6 k
0.6	600	0.2	5	39	17	0.67	6 l
0.6	600	0.2	10	81	34	0.67	6 m
0.6	600	0.2	50	412	166	0.67	6 n
0.6	600	0.2	100	827	332	0.67	6 o
0.6	600	0.4	2	21	10	0.33	6 p
0.6	600	0.4	5	59	25	0.33	6 q
0.6	600	0.4	10	121	50	0.33	6 r
0.6	600	0.4	50	619	249	0.33	6 s
0.6	600	0.4	100	1242	498	0.33	6 t
0.7	700	0.05	2	11	5	0.93	7 a
0.7	700	0.05	5	31	11	0.93	7 b
0.7	700	0.05	10	68	22	0.93	7 c
0.7	700	0.05	50	352	107	0.93	7 d
0.7	700	0.05	100	705	213	0.93	7 e
0.7	700	0.1	2	11	5	0.86	7 f
0.7	700	0.1	5	31	11	0.86	7 g
0.7	700	0.1	10	68	22	0.86	7 h
0.7	700	0.1	50	358	109	0.86	7 i
0.7	700	0.1	100	722	218	0.86	7 j
0.7	700	0.2	2	11	5	0.71	7 k
0.7	700	0.2	5	34	12	0.71	7 l
0.7	700	0.2	10	74	24	0.71	7 m

SCR Ratio	SCR	Utilization	B	MBS	L	P0	Type ID
0.7	700	0.2	50	384	117	0.71	7 n
0.7	700	0.2	100	772	233	0.71	7 o
0.7	700	0.4	2	14	6	0.43	7 p
0.7	700	0.4	5	44	15	0.43	7 q
0.7	700	0.4	10	94	30	0.43	7 r
0.7	700	0.4	50	482	146	0.43	7 s
0.7	700	0.4	100	965	291	0.43	7 t
0.8	800	0.05	2	6	3	0.94	8 a
0.8	800	0.05	5	26	7	0.94	8 b
0.8	800	0.05	10	61	14	0.94	8 c
0.8	800	0.05	50	342	70	0.94	8 d
0.8	800	0.05	100	692	140	0.94	8 e
0.8	800	0.1	2	6	3	0.88	8 f
0.8	800	0.1	5	31	8	0.88	8 g
0.8	800	0.1	10	66	15	0.88	8 h
0.8	800	0.1	50	352	72	0.88	8 i
0.8	800	0.1	100	707	143	0.88	8 j
0.8	800	0.2	2	6	3	0.75	8 k
0.8	800	0.2	5	31	8	0.75	8 l
0.8	800	0.2	10	66	15	0.75	8 m
0.8	800	0.2	50	362	74	0.75	8 n
0.8	800	0.2	100	732	148	0.75	8 o
0.8	800	0.4	2	11	4	0.5	8 p
0.8	800	0.4	5	36	9	0.5	8 q
0.8	800	0.4	10	76	17	0.5	8 r
0.8	800	0.4	50	406	83	0.5	8 s
0.8	800	0.4	100	821	166	0.5	8 t
0.9	900	0.05	2	2	2	0.94	9 a
0.9	900	0.05	5	21	4	0.94	9 b
0.9	900	0.05	10	51	7	0.94	9 c
0.9	900	0.05	50	332	35	0.94	9 d
0.9	900	0.05	100	682	70	0.94	9 e
0.9	900	0.1	2	2	2	0.89	9 f
0.9	900	0.1	5	21	4	0.89	9 g
0.9	900	0.1	10	51	7	0.89	9 h
0.9	900	0.1	50	332	35	0.89	9 i
0.9	900	0.1	100	682	70	0.89	9 j
0.9	900	0.2	2	2	2	0.78	9 k
0.9	900	0.2	5	21	4	0.78	9 l
0.9	900	0.2	10	61	8	0.78	9 m
0.9	900	0.2	50	342	36	0.78	9 n
0.9	900	0.2	100	702	72	0.78	9 o
0.9	900	0.4	2	2	2	0.56	9 p
0.9	900	0.4	5	21	4	0.56	9 q
0.9	900	0.4	10	61	8	0.56	9 r
0.9	900	0.4	50	362	38	0.56	9 s
0.9	900	0.4	100	732	75	0.56	9 t
0.9	900	0.8	2	11	3	0.11	9 u
0.9	900	0.8	5	51	7	0.11	9 v

SCR Ratio	SCR	Utilization	B	MBS	L	P0	Type ID
0.9	900	0.8	10	111	13	0.11	9 w
0.9	900	0.8	50	612	63	0.11	9 x
0.9	900	0.8	100	1232	125	0.11	9 y

CLR bound: 1e-3 .
PCR: 1000 cells/sec .

SCR Ratio	SCR	Utilization	B	MBS	X	P0	Type ID
0.01	10	0.001	2	15	16	0.9	.01 a
0.01	10	0.001	5	37	38	0.9	.01 b
0.01	10	0.001	10	76	76	0.9	.01 c
0.01	10	0.001	50	383	380	0.9	.01 d
0.01	10	0.001	100	767	760	0.9	.01 e
0.01	10	0.005	2	27	28	0.5	.01 f
0.01	10	0.005	5	69	69	0.5	.01 g
0.01	10	0.005	10	137	137	0.5	.01 h
0.01	10	0.005	50	687	681	0.5	.01 I
0.01	10	0.005	100	1374	1361	0.5	.01 j
0.01	10	0.009	2	136	136	0.1	.01 k
0.01	10	0.009	5	341	339	0.1	.01 l
0.01	10	0.009	10	684	678	0.1	.01 m
0.01	10	0.009	50	3422	3389	0.1	.01 n
0.01	10	0.009	100	6846	6778	0.1	.01 o
0.02	20	0.002	2	15	16	0.9	.02 a
0.02	20	0.002	5	38	38	0.9	.02 b
0.02	20	0.002	10	77	76	0.9	.02 c
0.02	20	0.002	50	383	376	0.9	.02 d
0.02	20	0.002	100	766	751	0.9	.02 e
0.02	20	0.01	2	27	27	0.5	.02 f
0.02	20	0.01	5	68	68	0.5	.02 g
0.02	20	0.01	10	137	135	0.5	.02 h
0.02	20	0.01	50	684	671	0.5	.02 I
0.02	20	0.01	100	1368	1341	0.5	.02 j
0.02	20	0.018	2	135	133	0.1	.02 k
0.02	20	0.018	5	339	333	0.1	.02 l
0.02	20	0.018	10	678	665	0.1	.02 m
0.02	20	0.018	50	3391	3324	0.1	.02 n
0.02	20	0.018	100	6783	6648	0.1	.02 o
0.03	30	0.003	2	14	15	0.9	.03 a
0.03	30	0.003	5	38	38	0.9	.03 b
0.03	30	0.003	10	76	75	0.9	.03 c
0.03	30	0.003	50	382	372	0.9	.03 d
0.03	30	0.003	100	765	743	0.9	.03 e
0.03	30	0.015	2	27	27	0.5	.03 f
0.03	30	0.015	5	68	67	0.5	.03 g
0.03	30	0.015	10	136	133	0.5	.03 h
0.03	30	0.015	50	680	661	0.5	.03 I
0.03	30	0.015	100	1361	1321	0.5	.03 j
0.03	30	0.027	2	134	131	0.1	.03 k
0.03	30	0.027	5	335	326	0.1	.03 l
0.03	30	0.027	10	671	652	0.1	.03 m
0.03	30	0.027	50	3360	3260	0.1	.03 n
0.03	30	0.027	100	6721	6520	0.1	.03 o

SCR Ratio	SCR	Utilization	B	MBS	X	P0	Type ID
0.04	40	0.004	2	15	15	0.9	.04 a
0.04	40	0.004	5	38	37	0.9	.04 b
0.04	40	0.004	10	76	74	0.9	.04 c
0.04	40	0.004	50	382	367	0.9	.04 d
0.04	40	0.004	100	764	734	0.9	.04 e
0.04	40	0.02	2	26	26	0.5	.04 f
0.04	40	0.02	5	67	65	0.5	.04 g
0.04	40	0.02	10	134	130	0.5	.04 h
0.04	40	0.02	50	676	650	0.5	.04 I
0.04	40	0.02	100	1354	1300	0.5	.04 j
0.04	40	0.036	2	132	128	0.1	.04 k
0.04	40	0.036	5	332	320	0.1	.04 l
0.04	40	0.036	10	666	640	0.1	.04 m
0.04	40	0.036	50	3330	3197	0.1	.04 n
0.04	40	0.036	100	6659	6393	0.1	.04 o
0.05	50	0.005	2	15	15	0.9	.05 a
0.05	50	0.005	5	38	37	0.9	.05 b
0.05	50	0.005	10	76	73	0.9	.05 c
0.05	50	0.005	50	382	363	0.9	.05 d
0.05	50	0.005	100	764	726	0.9	.05 e
0.05	50	0.025	2	26	26	0.5	.05 f
0.05	50	0.025	5	66	64	0.5	.05 g
0.05	50	0.025	10	134	128	0.5	.05 h
0.05	50	0.025	50	673	640	0.5	.05 I
0.05	50	0.025	100	1347	1280	0.5	.05 j
0.05	50	0.045	2	132	126	0.1	.05 k
0.05	50	0.045	5	329	314	0.1	.05 l
0.05	50	0.045	10	659	627	0.1	.05 m
0.05	50	0.045	50	3298	3134	0.1	.05 n
0.05	50	0.045	100	6597	6268	0.1	.05 o
0.06	60	0.006	2	15	15	0.9	.06 a
0.06	60	0.006	5	37	36	0.9	.06 b
0.06	60	0.006	10	76	72	0.9	.06 c
0.06	60	0.006	50	381	359	0.9	.06 d
0.06	60	0.006	100	763	718	0.9	.06 e
0.06	60	0.03	2	27	26	0.5	.06 f
0.06	60	0.03	5	66	63	0.5	.06 g
0.06	60	0.03	10	133	126	0.5	.06 h
0.06	60	0.03	50	670	630	0.5	.06 I
0.06	60	0.03	100	1340	1260	0.5	.06 j
0.06	60	0.054	2	130	123	0.1	.06 k
0.06	60	0.054	5	327	308	0.1	.06 l
0.06	60	0.054	10	653	615	0.1	.06 m
0.06	60	0.054	50	3267	3072	0.1	.06 n
0.06	60	0.054	100	6534	6143	0.1	.06 o

SCR Ratio	SCR	Utilization	B	MBS	X	P0	Type ID
0.07	70	0.007	2	15	15	0.9	.07 a
0.07	70	0.007	5	38	36	0.9	.07 b
0.07	70	0.007	10	75	71	0.9	.07 c
0.07	70	0.007	50	381	355	0.9	.07 d
0.07	70	0.007	100	762	709	0.9	.07 e
0.07	70	0.035	2	26	25	0.5	.07 f
0.07	70	0.035	5	66	62	0.5	.07 g
0.07	70	0.035	10	132	124	0.5	.07 h
0.07	70	0.035	50	666	620	0.5	.07 I
0.07	70	0.035	100	1333	1240	0.5	.07 j
0.07	70	0.063	2	129	121	0.1	.07 k
0.07	70	0.063	5	323	301	0.1	.07 l
0.07	70	0.063	10	646	602	0.1	.07 m
0.07	70	0.063	50	3236	3010	0.1	.07 n
0.07	70	0.063	100	6472	6020	0.1	.07 o
0.08	80	0.008	2	15	15	0.9	.08 a
0.08	80	0.008	5	38	36	0.9	.08 b
0.08	80	0.008	10	76	71	0.9	.08 c
0.08	80	0.008	50	380	351	0.9	.08 d
0.08	80	0.008	100	761	701	0.9	.08 e
0.08	80	0.04	2	26	25	0.5	.08 f
0.08	80	0.04	5	66	62	0.5	.08 g
0.08	80	0.04	10	133	123	0.5	.08 h
0.08	80	0.04	50	663	611	0.5	.08 I
0.08	80	0.04	100	1327	1221	0.5	.08 j
0.08	80	0.072	2	127	118	0.1	.08 k
0.08	80	0.072	5	320	295	0.1	.08 l
0.08	80	0.072	10	640	590	0.1	.08 m
0.08	80	0.072	50	3204	2949	0.1	.08 n
0.08	80	0.072	100	6410	5898	0.1	.08 o
0.09	90	0.009	2	14	14	0.9	.09 a
0.09	90	0.009	5	37	35	0.9	.09 b
0.09	90	0.009	10	76	70	0.9	.09 c
0.09	90	0.009	50	380	347	0.9	.09 d
0.09	90	0.009	100	761	693	0.9	.09 e
0.09	90	0.045	2	26	25	0.5	.09 f
0.09	90	0.045	5	66	61	0.5	.09 g
0.09	90	0.045	10	132	121	0.5	.09 h
0.09	90	0.045	50	660	601	0.5	.09 I
0.09	90	0.045	100	1319	1201	0.5	.09 j
0.09	90	0.081	2	126	116	0.1	.09 k
0.09	90	0.081	5	316	289	0.1	.09 l
0.09	90	0.081	10	634	578	0.1	.09 m
0.09	90	0.081	50	3174	2889	0.1	.09 n
0.09	90	0.081	100	6348	5777	0.1	.09 o

SCR Ratio	SCR	Utilization	B	MBS	X	P0	Type ID
0.1	100	0.01	2	14	14	0.9	.1 a
0.1	100	0.01	5	38	35	0.9	.1 b
0.1	100	0.01	10	76	69	0.9	.1 c
0.1	100	0.01	50	379	342	0.9	.1 d
0.1	100	0.01	100	759	684	0.9	.1 e
0.1	100	0.05	2	25	24	0.5	.1 f
0.1	100	0.05	5	65	60	0.5	.1 g
0.1	100	0.05	10	131	119	0.5	.1 h
0.1	100	0.05	50	656	591	0.5	.1 I
0.1	100	0.05	100	1313	1182	0.5	.1 j
0.1	100	0.09	2	125	114	0.1	.1 k
0.1	100	0.09	5	313	283	0.1	.1 l
0.1	100	0.09	10	628	566	0.1	.1 m
0.1	100	0.09	50	3142	2829	0.1	.1 n
0.1	100	0.09	100	6286	5658	0.1	.1 o
0.2	200	0.02	2	15	13	0.9	.2 a
0.2	200	0.02	5	37	31	0.9	.2 b
0.2	200	0.02	10	75	61	0.9	.2 c
0.2	200	0.02	50	375	301	0.9	.2 d
0.2	200	0.02	100	752	602	0.9	.2 e
0.2	200	0.1	2	24	20	0.5	.2 f
0.2	200	0.1	5	61	50	0.5	.2 g
0.2	200	0.1	10	124	100	0.5	.2 h
0.2	200	0.1	50	622	498	0.5	.2 I
0.2	200	0.1	100	1243	995	0.5	.2 j
0.2	200	0.18	2	112	91	0.1	.2 k
0.2	200	0.18	5	282	227	0.1	.2 l
0.2	200	0.18	10	566	454	0.1	.2 m
0.2	200	0.18	50	2832	2266	0.1	.2 n
0.2	200	0.18	100	5664	4532	0.1	.2 o
0.3	300	0.03	2	14	11	0.9	.3 a
0.3	300	0.03	5	37	27	0.9	.3 b
0.3	300	0.03	10	74	53	0.9	.3 c
0.3	300	0.03	50	372	261	0.9	.3 d
0.3	300	0.03	100	744	522	0.9	.3 e
0.3	300	0.15	2	22	17	0.5	.3 f
0.3	300	0.15	5	58	42	0.5	.3 g
0.3	300	0.15	10	117	83	0.5	.3 h
0.3	300	0.15	50	587	412	0.5	.3 I
0.3	300	0.15	100	1174	823	0.5	.3 j
0.3	300	0.27	2	100	71	0.1	.3 k
0.3	300	0.27	5	251	177	0.1	.3 l
0.3	300	0.27	10	502	353	0.1	.3 m
0.3	300	0.27	50	2520	1765	0.1	.3 n
0.3	300	0.27	100	5042	3530	0.1	.3 o

SCR Ratio	SCR	Utilization	B	MBS	X	P0	Type ID
0.4	400	0.04	2	13	9	0.9	.4 a
0.4	400	0.04	5	36	23	0.9	.4 b
0.4	400	0.04	10	73	45	0.9	.4 c
0.4	400	0.04	50	368	222	0.9	.4 d
0.4	400	0.04	100	737	443	0.9	.4 e
0.4	400	0.2	2	21	14	0.5	.4 f
0.4	400	0.2	5	54	34	0.5	.4 g
0.4	400	0.2	10	109	67	0.5	.4 h
0.4	400	0.2	50	552	332	0.5	.4 I
0.4	400	0.2	100	1105	664	0.5	.4 j
0.4	400	0.36	2	88	54	0.1	.4 k
0.4	400	0.36	5	219	133	0.1	.4 l
0.4	400	0.36	10	441	266	0.1	.4 m
0.4	400	0.36	50	2210	1327	0.1	.4 n
0.4	400	0.36	100	4420	2653	0.1	.4 o
0.5	500	0.05	2	13	8	0.9	.5 a
0.5	500	0.05	5	35	19	0.9	.5 b
0.5	500	0.05	10	71	37	0.9	.5 c
0.5	500	0.05	50	364	183	0.9	.5 d
0.5	500	0.05	100	728	365	0.9	.5 e
0.5	500	0.25	2	19	11	0.5	.5 f
0.5	500	0.25	5	49	26	0.5	.5 g
0.5	500	0.25	10	101	52	0.5	.5 h
0.5	500	0.25	50	518	260	0.5	.5 I
0.5	500	0.25	100	1036	519	0.5	.5 j
0.5	500	0.45	2	73	38	0.1	.5 k
0.5	500	0.45	5	187	95	0.1	.5 l
0.5	500	0.45	10	377	190	0.1	.5 m
0.5	500	0.45	50	1898	950	0.1	.5 n
0.5	500	0.45	100	3798	1900	0.1	.5 o
0.6	600	0.06	2	11	6	0.9	.6 a
0.6	600	0.06	5	34	15	0.9	.6 b
0.6	600	0.06	10	69	29	0.9	.6 c
0.6	600	0.06	50	359	145	0.9	.6 d
0.6	600	0.06	100	719	289	0.9	.6 e
0.6	600	0.3	2	16	8	0.5	.6 f
0.6	600	0.3	5	46	20	0.5	.6 g
0.6	600	0.3	10	94	39	0.5	.6 h
0.6	600	0.3	50	482	194	0.5	.6 I
0.6	600	0.3	100	964	387	0.5	.6 j
0.6	600	0.54	2	61	26	0.1	.6 k
0.6	600	0.54	5	156	64	0.1	.6 l
0.6	600	0.54	10	316	128	0.1	.6 m
0.6	600	0.54	50	1586	636	0.1	.6 n
0.6	600	0.54	100	3177	1272	0.1	.6 o

CLR bound: 1e-3 .
PCR: 1000 cells/sec .

SCR Ratio	SCR	Utilization	B	MBS	X	P0	Type ID
0.7	700	0.07	2	11	5	0.9	.7 a
0.7	700	0.07	5	31	11	0.9	.7 b
0.7	700	0.07	10	68	22	0.9	.7 c
0.7	700	0.07	50	354	108	0.9	.7 d
0.7	700	0.07	100	712	215	0.9	.7 e
0.7	700	0.35	2	14	6	0.5	.7 f
0.7	700	0.35	5	41	14	0.5	.7 g
0.7	700	0.35	10	84	27	0.5	.7 h
0.7	700	0.35	50	444	135	0.5	.7 I
0.7	700	0.35	100	895	270	0.5	.7 j
0.7	700	0.63	2	48	16	0.1	.7 k
0.7	700	0.63	5	124	39	0.1	.7 l
0.7	700	0.63	10	251	77	0.1	.7 m
0.7	700	0.63	50	1274	384	0.1	.7 n
0.7	700	0.63	100	2552	767	0.1	.7 o
0.8	800	0.08	2	6	3	0.9	.8 a
0.8	800	0.08	5	31	8	0.9	.8 b
0.8	800	0.08	10	66	15	0.9	.8 c
0.8	800	0.08	50	346	71	0.9	.8 d
0.8	800	0.08	100	702	142	0.9	.8 e
0.8	800	0.4	2	11	4	0.5	.8 f
0.8	800	0.4	5	36	9	0.5	.8 g
0.8	800	0.4	10	76	17	0.5	.8 h
0.8	800	0.4	50	406	83	0.5	.8 I
0.8	800	0.4	100	821	166	0.5	.8 j
0.8	800	0.72	2	31	8	0.1	.8 k
0.8	800	0.72	5	91	20	0.1	.8 l
0.8	800	0.72	10	186	39	0.1	.8 m
0.8	800	0.72	50	962	194	0.1	.8 n
0.8	800	0.72	100	1926	387	0.1	.8 o
0.9	900	0.09	2	2	2	0.9	.9 a
0.9	900	0.09	5	21	4	0.9	.9 b
0.9	900	0.09	10	51	7	0.9	.9 c
0.9	900	0.09	50	332	35	0.9	.9 d
0.9	900	0.09	100	682	70	0.9	.9 e
0.9	900	0.45	2	2	2	0.5	.9 f
0.9	900	0.45	5	21	4	0.5	.9 g
0.9	900	0.45	10	61	8	0.5	.9 h
0.9	900	0.45	50	362	38	0.5	.9 I
0.9	900	0.45	100	742	76	0.5	.9 j
0.9	900	0.81	2	11	3	0.1	.9 k
0.9	900	0.81	5	51	7	0.1	.9 l
0.9	900	0.81	10	121	14	0.1	.9 m
0.9	900	0.81	50	642	66	0.1	.9 n
0.9	900	0.81	100	1302	132	0.1	.9 o

Appendix B

Simulation Scenarios

This Appendix contains three sets of scenarios used in simulations designed to test the analyses described in Chapter 3.

Appendix B1 contains a list of fixed utilization scenarios derived from the types of traffic described in Appendix A1. Scenarios were chosen to cover the range from two to ten queues in a ten-timeslot WRR server, having mean number of empty timeslots from 1.49 to 8.96 timeslots (.149 to .896 normalized timeslots).

Appendices B2 and B3 contain lists of fixed probability scenarios derived from the types of traffic described in Appendix A2. Ten queue scenarios requiring ten and twenty timeslots are listed in Appendix B2, while Appendix B3 lists twenty queue, twenty timeslot scenarios.

The columns titled Queue Type ID's in all three sections of this appendix represent a shorthand notation for describing each queue type used in that scenario. The queue types pertain to the queue type ID shown in the last column of the listed Appendix A sub-section.

Results of simulations using the scenarios described in this Appendix are described in Chapter 4.

Appendix B1 Fixed Utilization 10 Timeslot Scenarios

(ρ=.05, .1, .2, .4, .8 only)

Scenario	Number of Queues, N	Mean μ_X/N	Composition SCR ratio (ρ)[# of Queues]	*MBS* Burst Type	Queue Type ID's (Appendix A1)	$\overline{MBS}$
1 (1)	10	.5	.1(.05) [10]	Short	1a 1a 1a 1a 1a 1a 1a 1a 1a 1a	25
				Mixed	1a 1a 1a 1a 1a 1e 1e 1e 1e 1e	669
				Long	1e 1e 1e 1e 1e 1e 1e 1e 1e 1e	1313
2 (2)	9	.55	.1(.05) [8]	Short	1a 1a 1a 1a 1a 1a 1a 1a 2a	24
			.2(.05) [1]	Mixed	1a 1a 1a 1a 1a 1e 1e 1e 2e	549
				Long	1e 1e 1e 1e 1e 1e 1e 1e 2e	1264
3 (3)	9	.5	.1(.05) [8]	Short	1a 1a 1a 1a 1a 1a 1a 1a 2f	25
			.2(.1) [1]	Mixed	1a 1a 1a 1a 1a 1e 1e 1e 2j	590
				Long	1e 1e 1e 1e 1e 1e 1e 1e 2j	1305
4 (16)	5	.75	.2(.05) [5]	Short	2a 2a 2a 2a 2a	16
				Mixed	2a 2a 2c 2e 2e	374
				Long	2e 2e 2e 2e 2e	874
5 (21)	5	.5	.2(.1) [5]	Short	2f 2f 2f 2f 2f 2f	24
				Mixed	2f 2f 2h 2j 2j	532
				Long	2j 2j 2j 2j 2j	1243
6 (67)	4	.798	.2(.05) [2]	Short	2a 2a 3a 3a	16
			.3(.05) [2]	Mixed	2a 2e 3a 3e	423
				Long	2e 2e 3e 3e	831
7 (76)	4	.498	.2(.05) [1]	Short	2a 2f 3k 3k	26
			.2(.1) [1]	Mixed	2e 2f 3k 3o	647
			.3(.2) [2]	Long	2e 2j 3o 3o	1358
8 (179)	5	.35	.1(.05) [3]	Short	1a 1a 1a 2f 5p	28
			.2(.1) [1]	Mixed	1a 1a 1e 2f 5t	69
			.5(.4) [1]	Long	1e 1e 1e 2j 5t	1451
9 (240)	2	.9	.5(.05) [2]	Short	5a 5a	13
				Mixed	5a 5e	371
				Long	5e 5e	728
10 (246)	2	.2	.5(.4) [2]	Short	5p 5p	39
				Mixed	5p 5t	1056
				Long	5t 5t	2072
11 (331)	2	.896		Short	1a 9a	14
			.1(.05) [1]	Mixed	1a 9e	354
			.9(.05) [1]	Long	1e 9e	998
12 (335)	2	.149		Short	1a 9u	18
			.1(.05) [1]	Mixed	1a 9y	629
			.9(.8) [1]	Long	1e 9y	1272

Appendix B2 Fixed Probability 10 Queue Scenarios

(ρ = Utilization needed for empty timeslot probabilities of .1, .5, and .9)

Scenario	Number of Timeslots, N	Mean μ_X/N	Composition SCR ratio (ρ)[# of Queues]	*MBS* Burst Type	Queue Type ID's (Appendix A2)	$\overline{MBS}$
1	10	.5	.1(.05) [10]	Short	.1f x 10	25
				Mixed	.1f x 5, .1j x 5	669
				Long	.1j x 10	1313
2	10	.1	.1(.09) [10]	Short	.1k x 10	125
				Mixed	.1k x 5, .1o x 5	3205.5
				Long	.1o x 10	6286
3	10	.9	.1(.01) [10]	Short	.1a x 10	14
				Mixed	.1a x 5, .1e x 5	386.5
				Long	.1e x 10	759
4	20	.5	.2(.1)[1], .1(.05)[7], .05(.025)[2]	Short	.2f x 1, .1f x 7, .05f x 2	12.6
				Mixed	.2f x 1, .1f x 4, .1j x 3, .05h x 2	216.6
				Long	.2j x 1, .1j x 7, .05j x 2	656.4
5	20	.1	.2(.18)[1], .1(.09)[7], .05(.045)[2]	Short	.2k x 1, .1k x 7, .05k x 2	62.6
				Mixed	.2k x 1, .1k x 4, .1o x 3, .05m x 2	1039.4
				Long	.2o x 1, .1o x 7, .05o x 2	3143
6	20	.9	.2(.02)[1], .1(.01)[7], .05(.005)[2]	Short	.2a x 1, .1a x 7, .05a x 2	7.2
				Mixed	.2a x 1, .1a x 4, .1e x 3, .05c x 2	125
				Long	.2e x 1, .1e x 7, .05e x 2	379.7
7	20	.5	.5(.25)[1], .1(.05)[1], .05(.025)[8]	Short	.5f x 1, .1f x 1, .05f x 8	12.6
				Mixed	.5h x 1, .1h x 1, .05f x 4, .05j x 4	286.2
				Long	.5j x 1, .1j x 1, .05j x 8	656.3
8	20	.1	.5(.45)[1], .1(.09)[1], .05(.045)[8]	Short	.5k x 1, .1k x 1, .05k x 8	62.7
				Mixed	.5m x 1, .1m x 1, .05k x 4, .05o x 4	1396
				Long	.5o x 1, .1o x 1, .05o x 8	3143
9	20	.9	.5(.05)[1], .1(.01)[1], .05(.005)[8]	Short	.5a x 1, .1a x 1, .05a x 8	7.4
				Mixed	.5c x 1, .1c x 1, .05a x 4, .05e x 4	163.5
				Long	.5e x 1, .1e x 1, .05e x 8	380.0

Appendix B3 Fixed Probability 20 Queue Scenarios

(ρ = Utilization needed for empty timeslot probabilities of .1, .5, and .9)

Scenario	Number of Timeslots, N	Mean μ_X/N	Composition SCR ratio (ρ)[# of Queues]	*MBS* Burst Type	Queue Type ID's (Appendix A2)	$\overline{MBS}$
1	20	.5	.05(.025)[20]	Short	.05f x 20	26
				Mixed	.05f x 10, .05j x 10	686.5
				Long	.05j x 20	1347
2	20	.1	.05(.045)[20]	Short	.05k x 20	132
				Mixed	.05k x 10, .05o x 10	3364.5
				Long	.05o x 20	6597
3	20	.9	.05(.005)[20]	Short	.05a x 20	15
				Mixed	.05a x 10, .05e x 10	389.5
				Long	.05e x 20	764

Appendix C

Results
Simulated System CLR, 10 Queue Existing Traffic, Added Queue of Normal Maximum Queue Length

This Appendix contains the detailed results of simulations where an “additional” traffic queue was introduced to a WRR server where the existing “normal” traffic was sufficient to fill the virtual path, based on the sum of the SCR’s of the “normal” traffic. The existing traffic occupied 10 traffic streams, each sized such that the resulting WRR schedule was composed of ten or twenty timeslots, depending on the scenario represented. The “added” traffic consists of a single stream passing through a queue sized in accordance with a normal Maximum Queue Length (MQL) as determined from Equation 3.5. Details of the scenarios used to obtain these results are contained in Appendices A and B. The results are in the form of bar charts where each set of bars represents the maximum and average cell loss ratio resulting from twenty simulations, where the “normal traffic has the rate and burst characteristics indicated along the x-axis.

Each chart shows the allowable CLR threshold as a heavy horizontal line at the appropriate level on the CLR axis. Each chart also shows a CLR threshold rate that shows the boundary below which the CLR for all scenarios at the Added Traffic Rate shown is less than the allowable CLR.

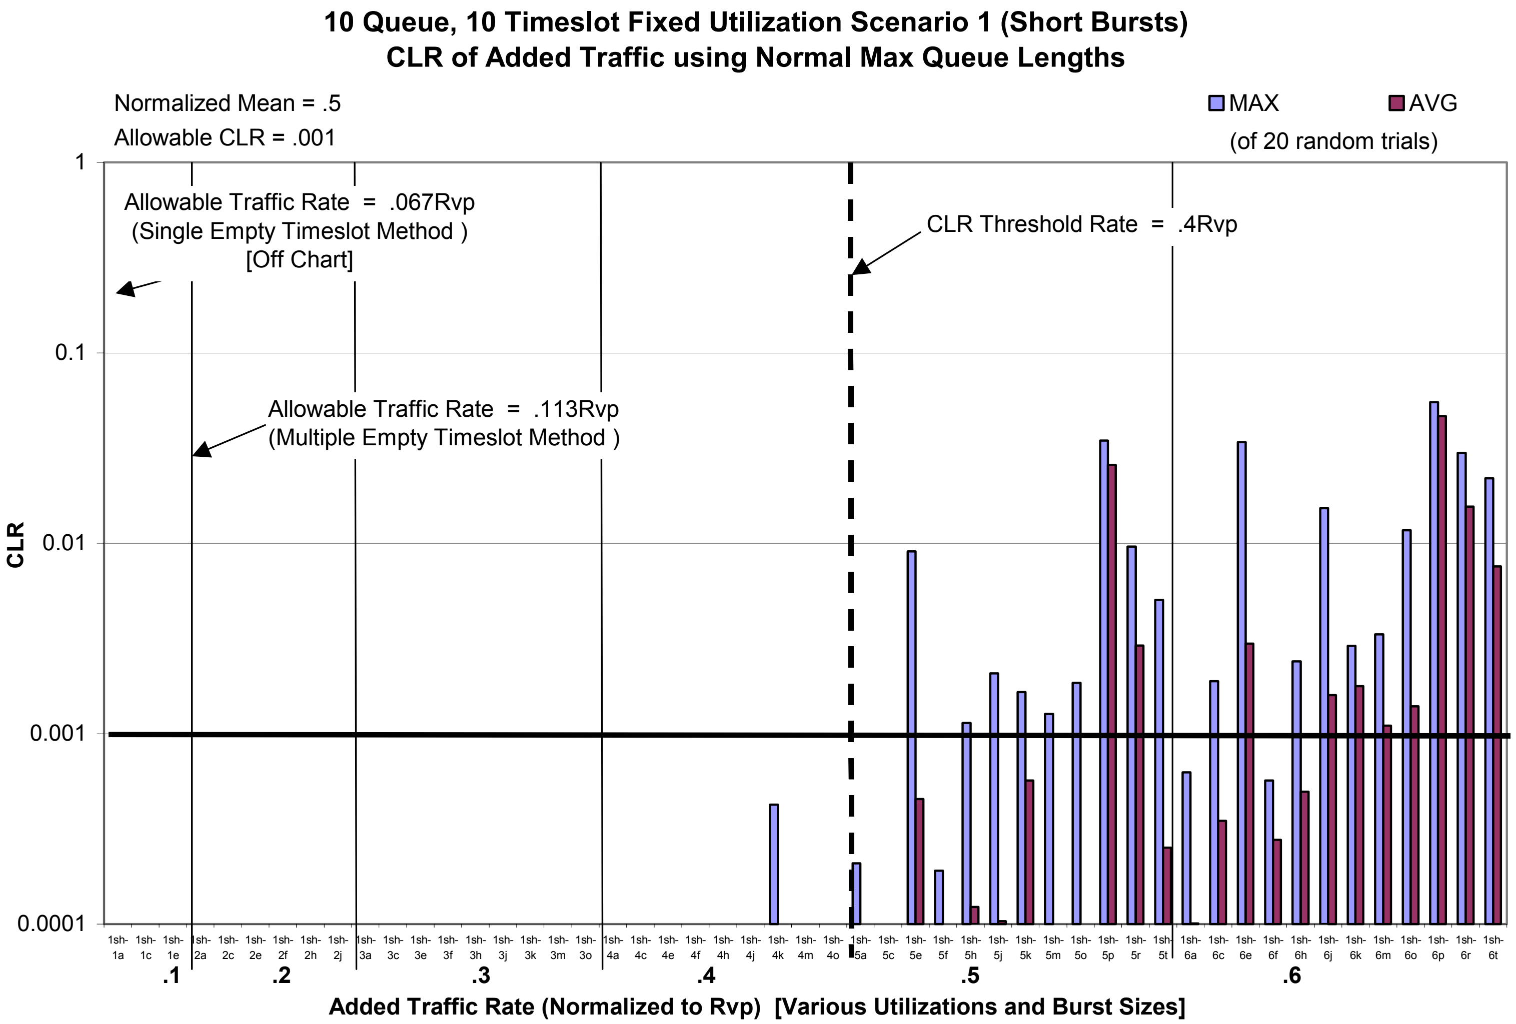

Figure C-1. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Utilization Scenario 1 (Short Bursts)

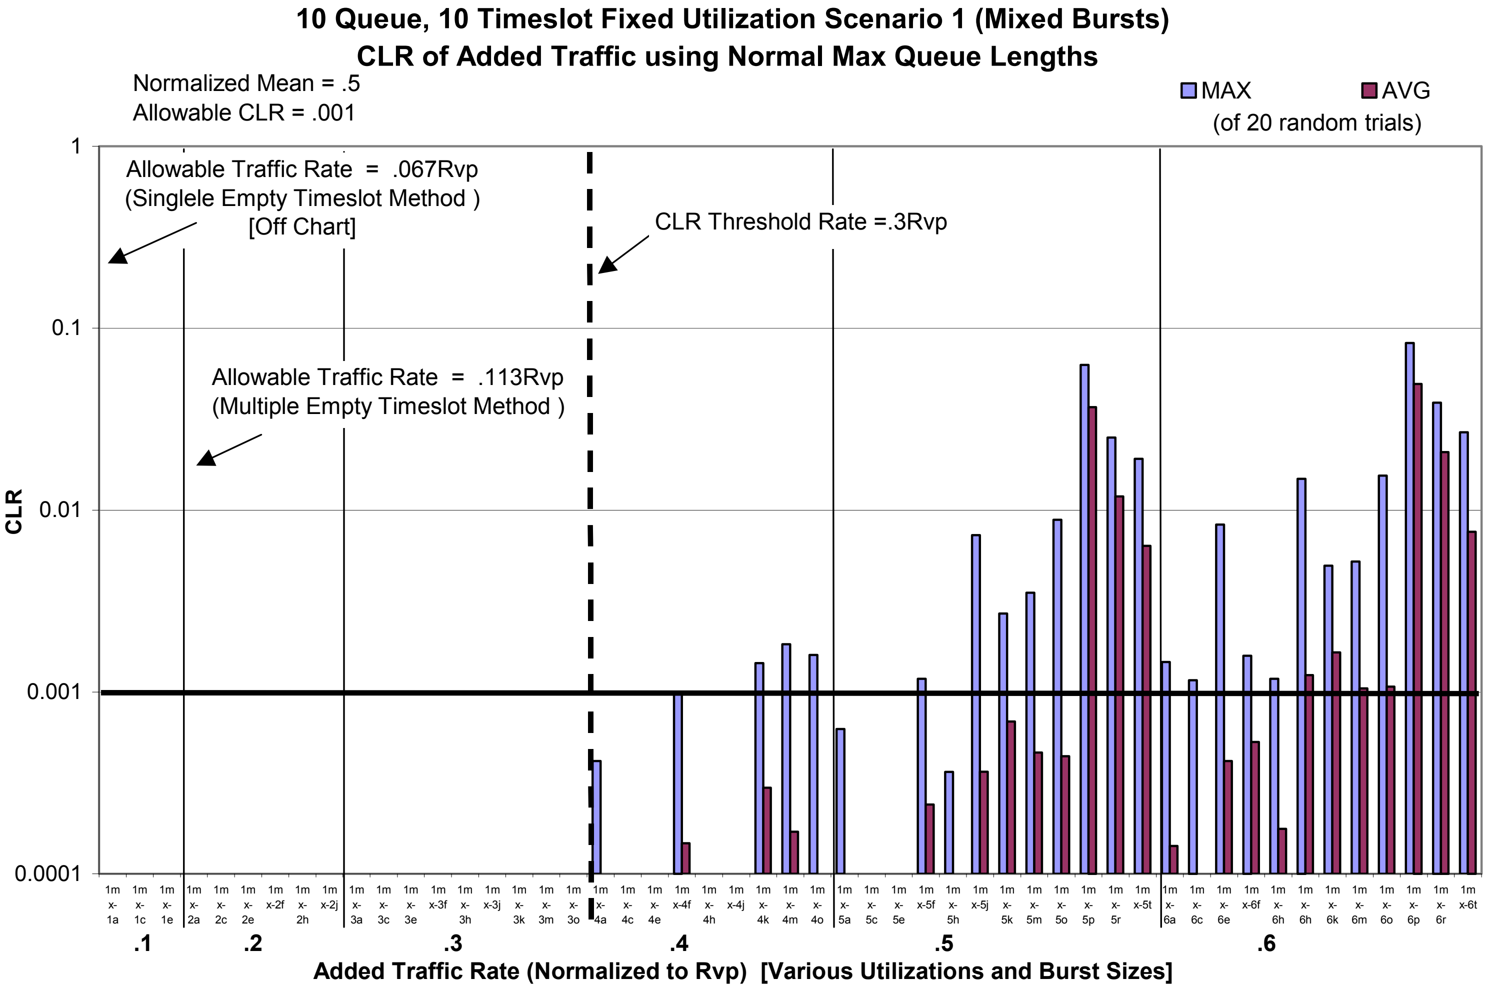

Figure C-2. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Utilization Scenario 1 (Mixed Bursts)

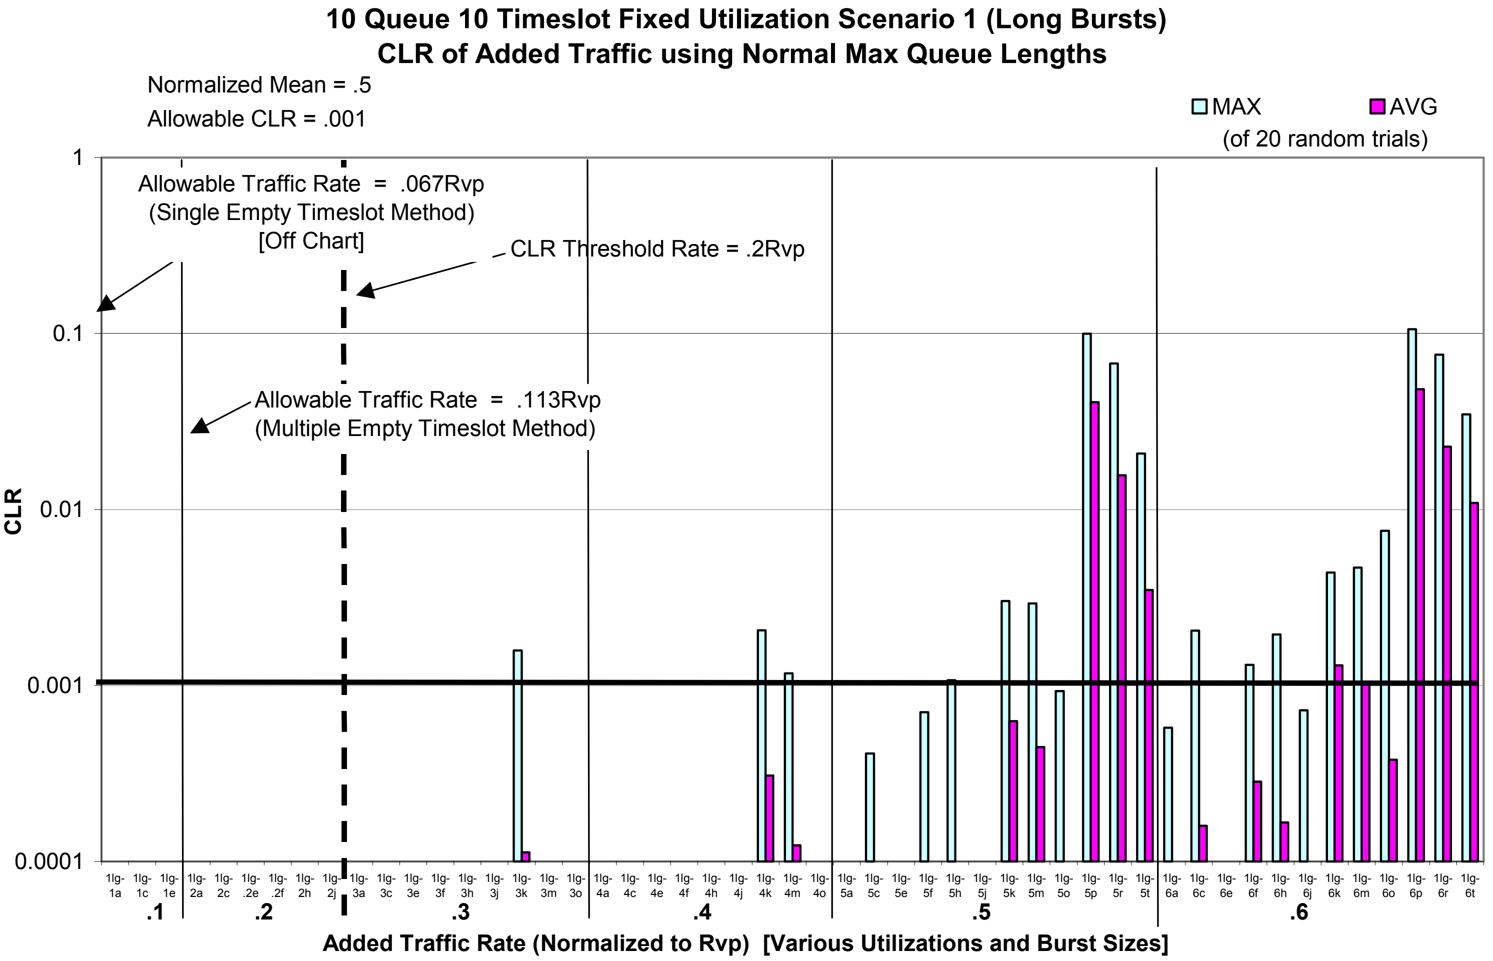

Figure C-3. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Utilization Scenario 1 (Long Bursts)

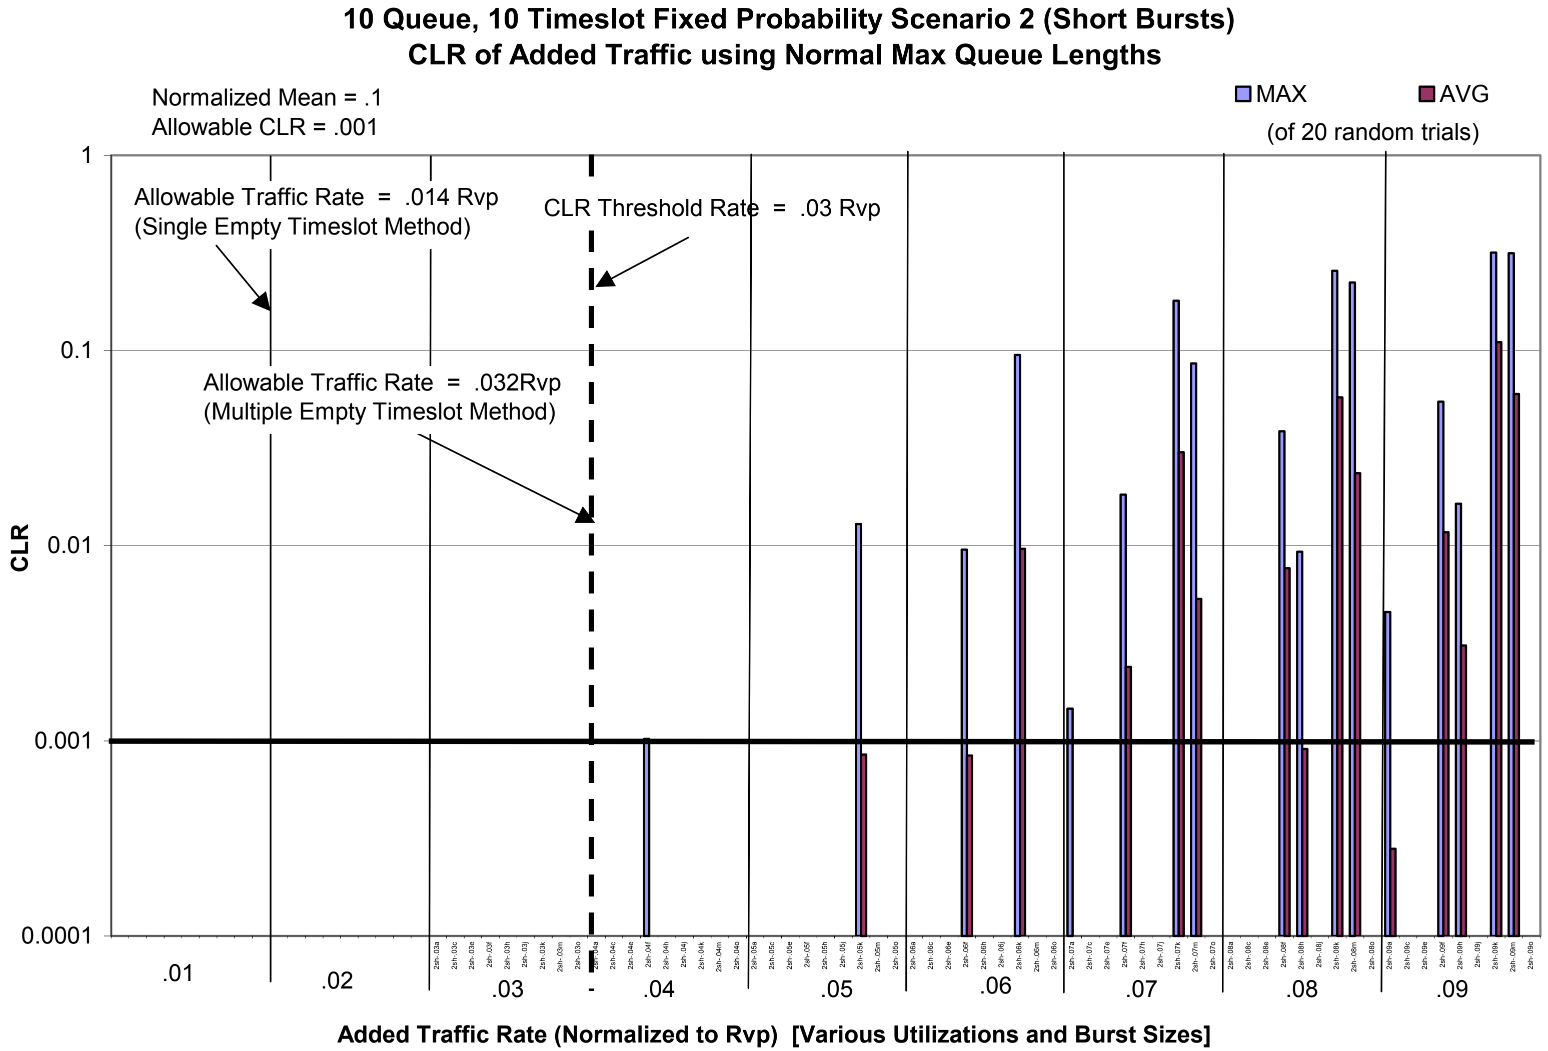

Figure C-4. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 2 (Short Bursts)

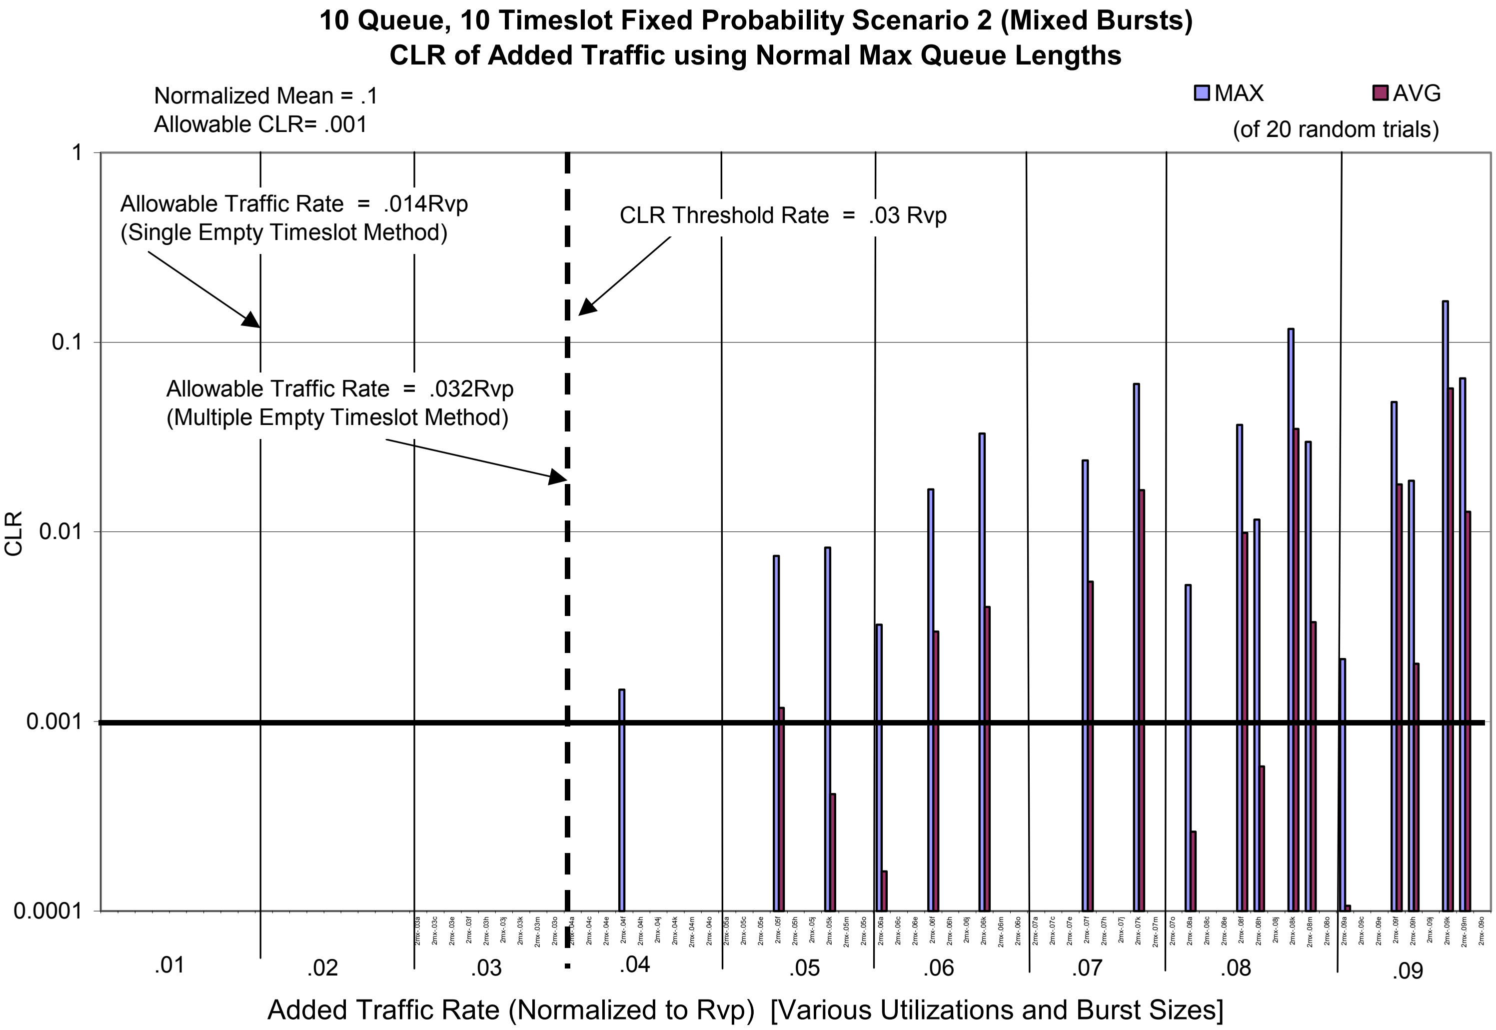

Figure C-5. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 2 (Mixed Bursts)

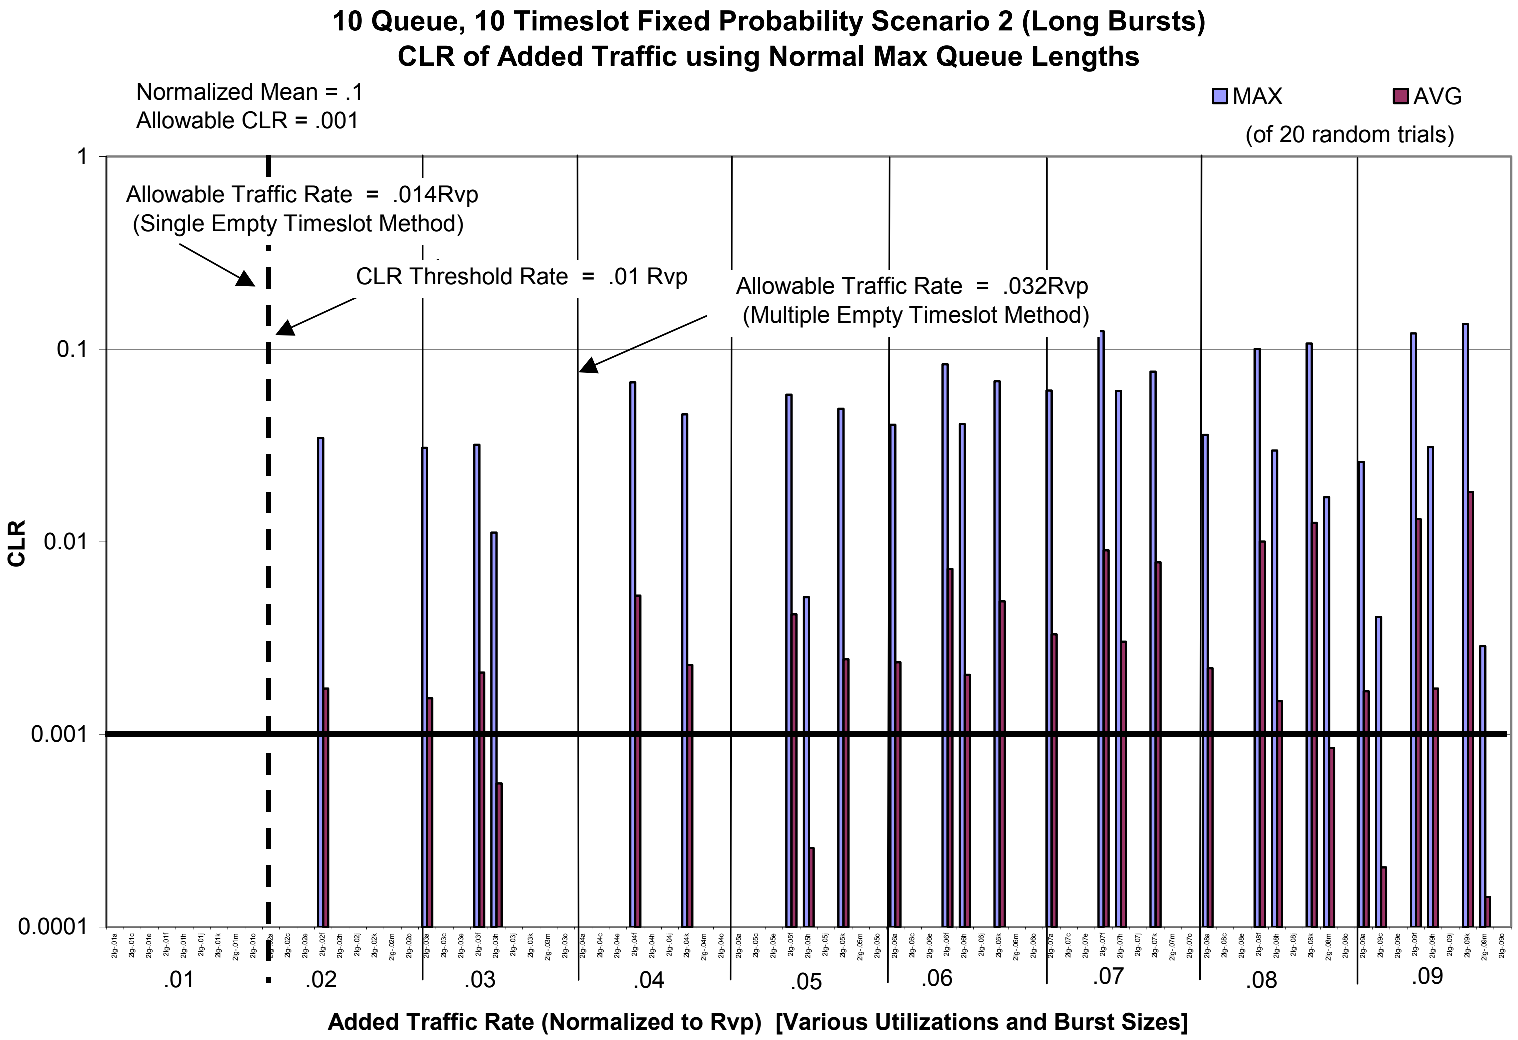

Figure C-6. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 2 (Long Bursts)

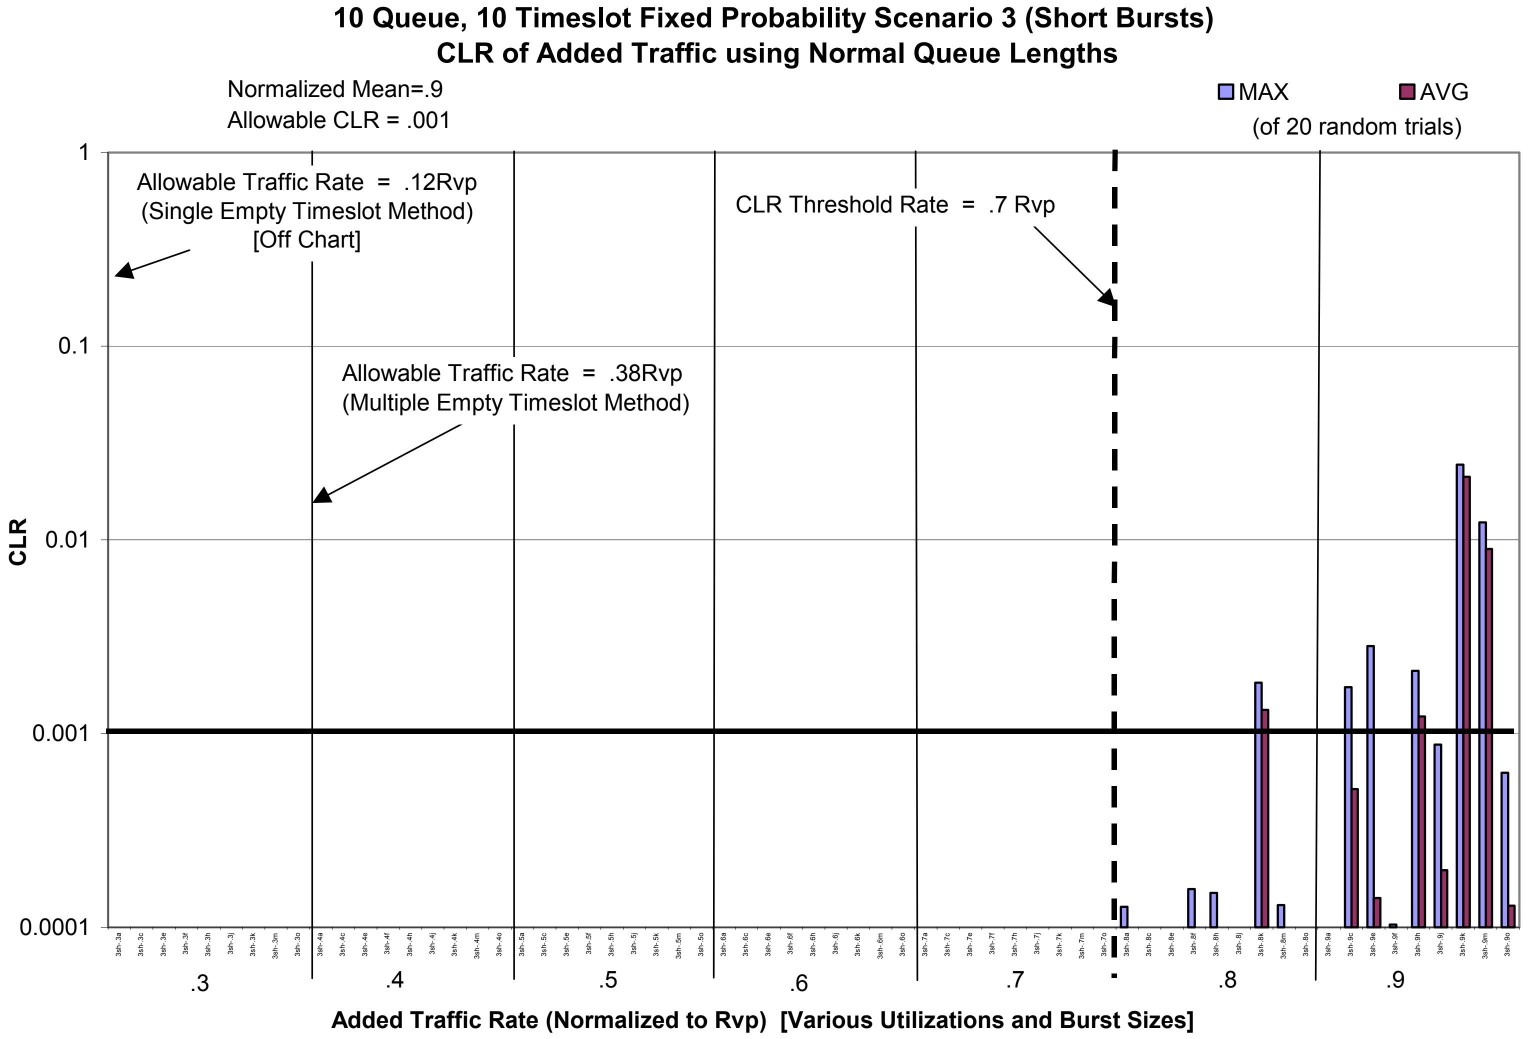

Figure C-7. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 3 (Short Bursts)

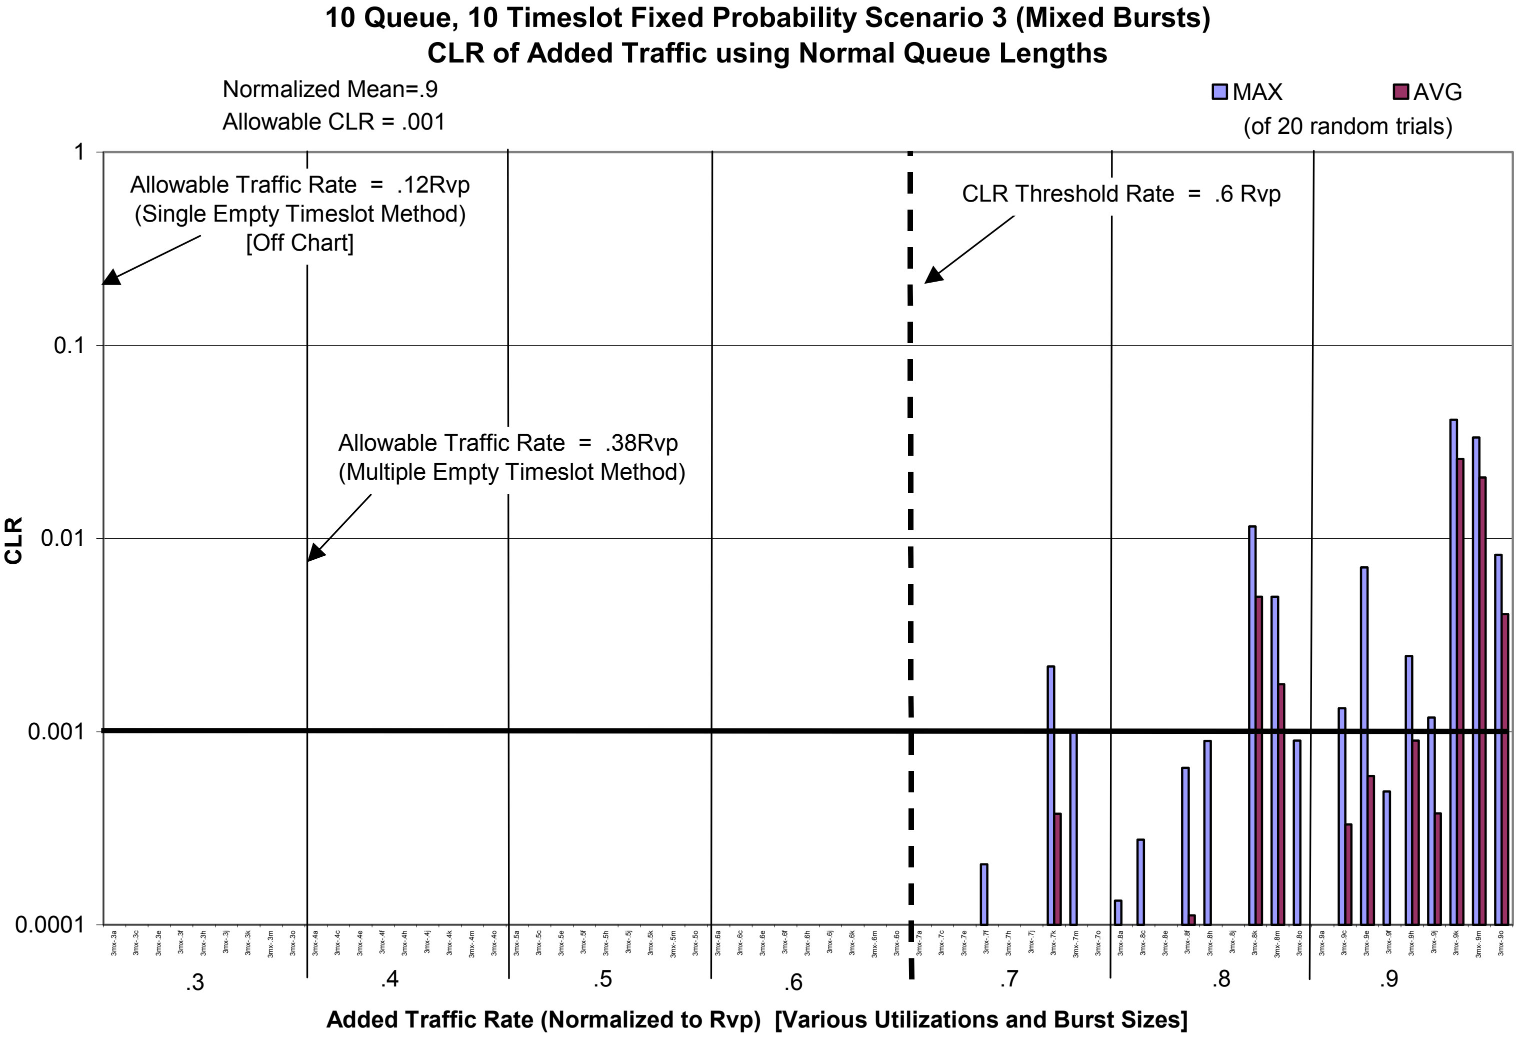

Figure C-8. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 3 (Mixed Bursts)

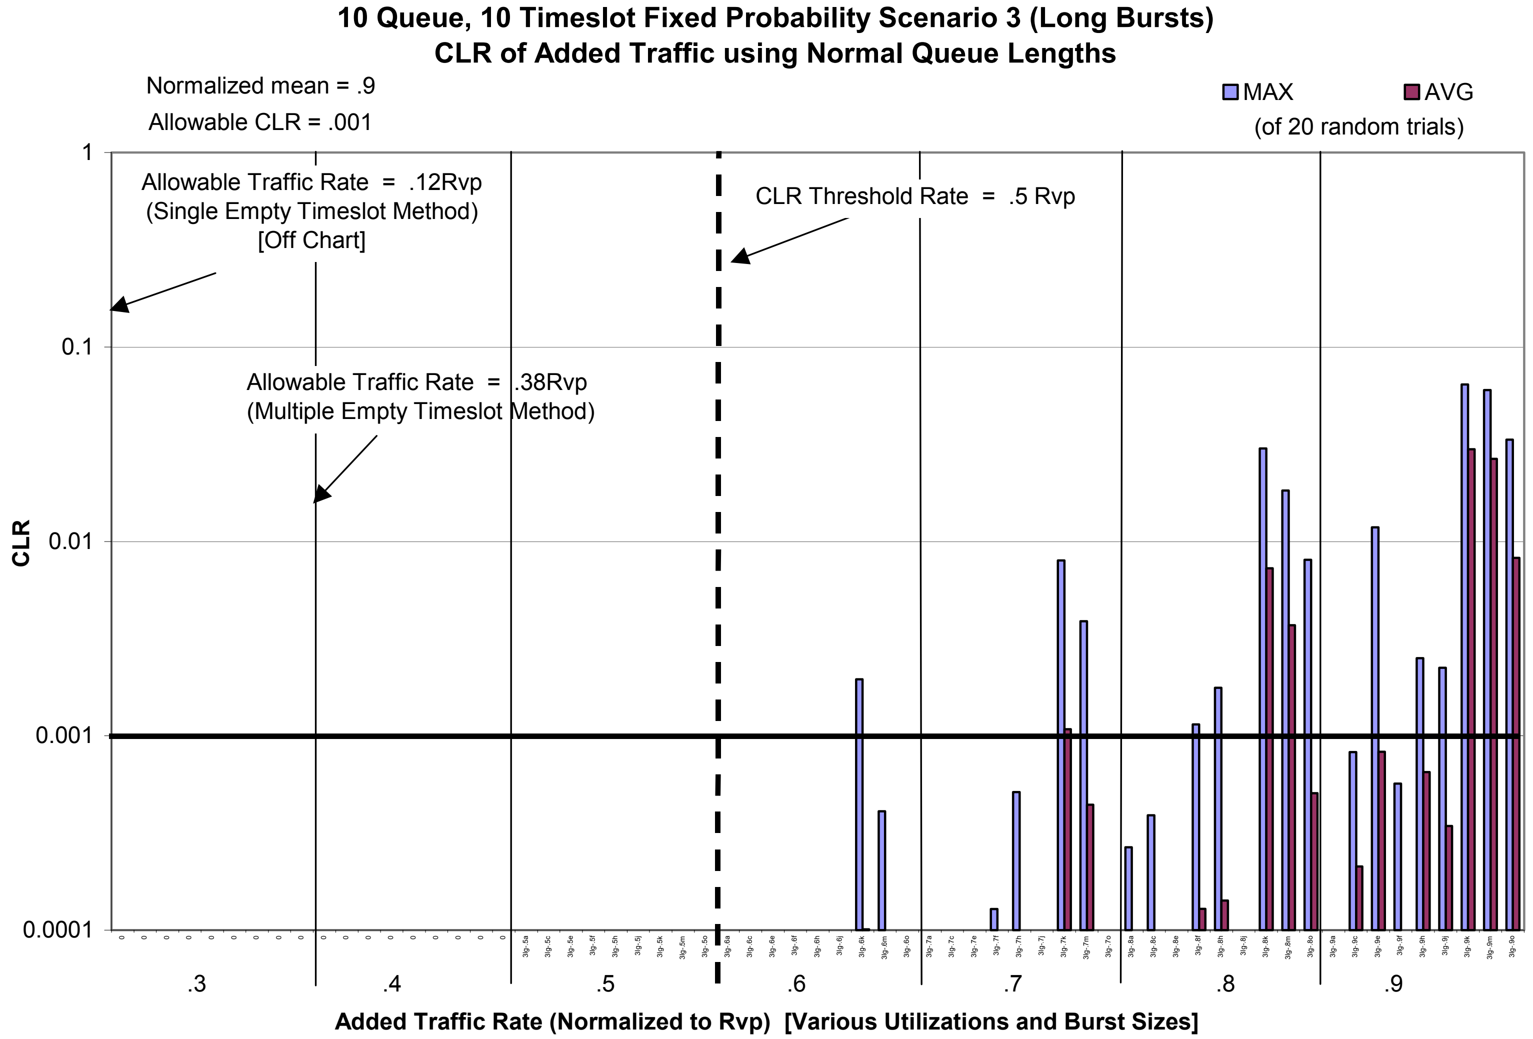

Figure C-9. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 3 (Long Bursts)

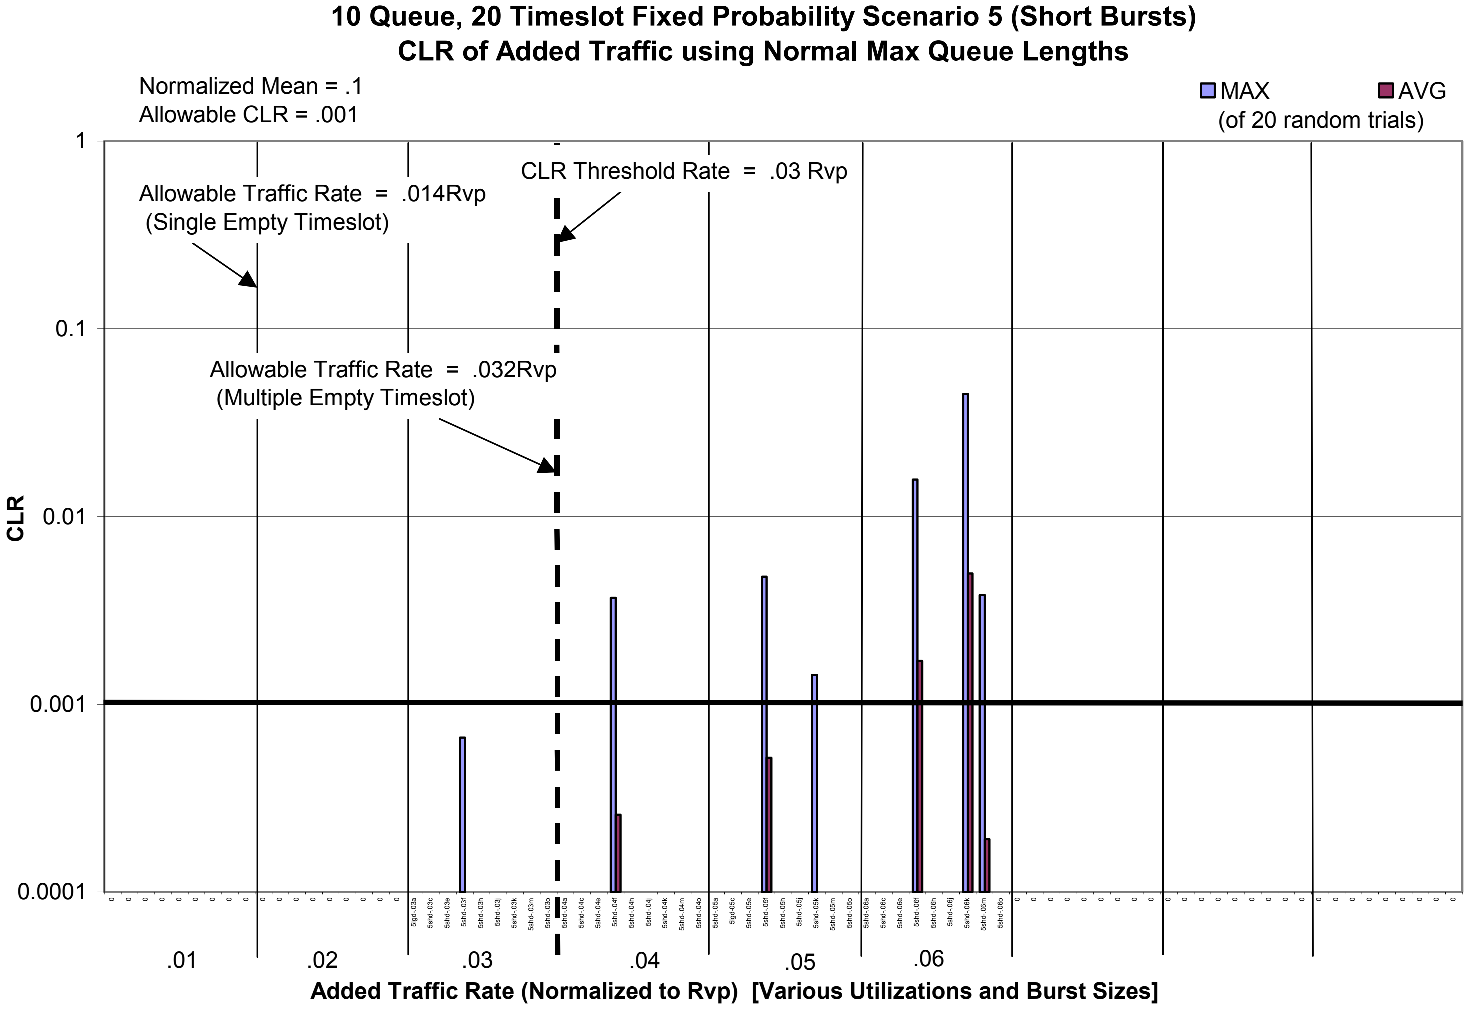

Figure C-10. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 5 (Short Bursts)

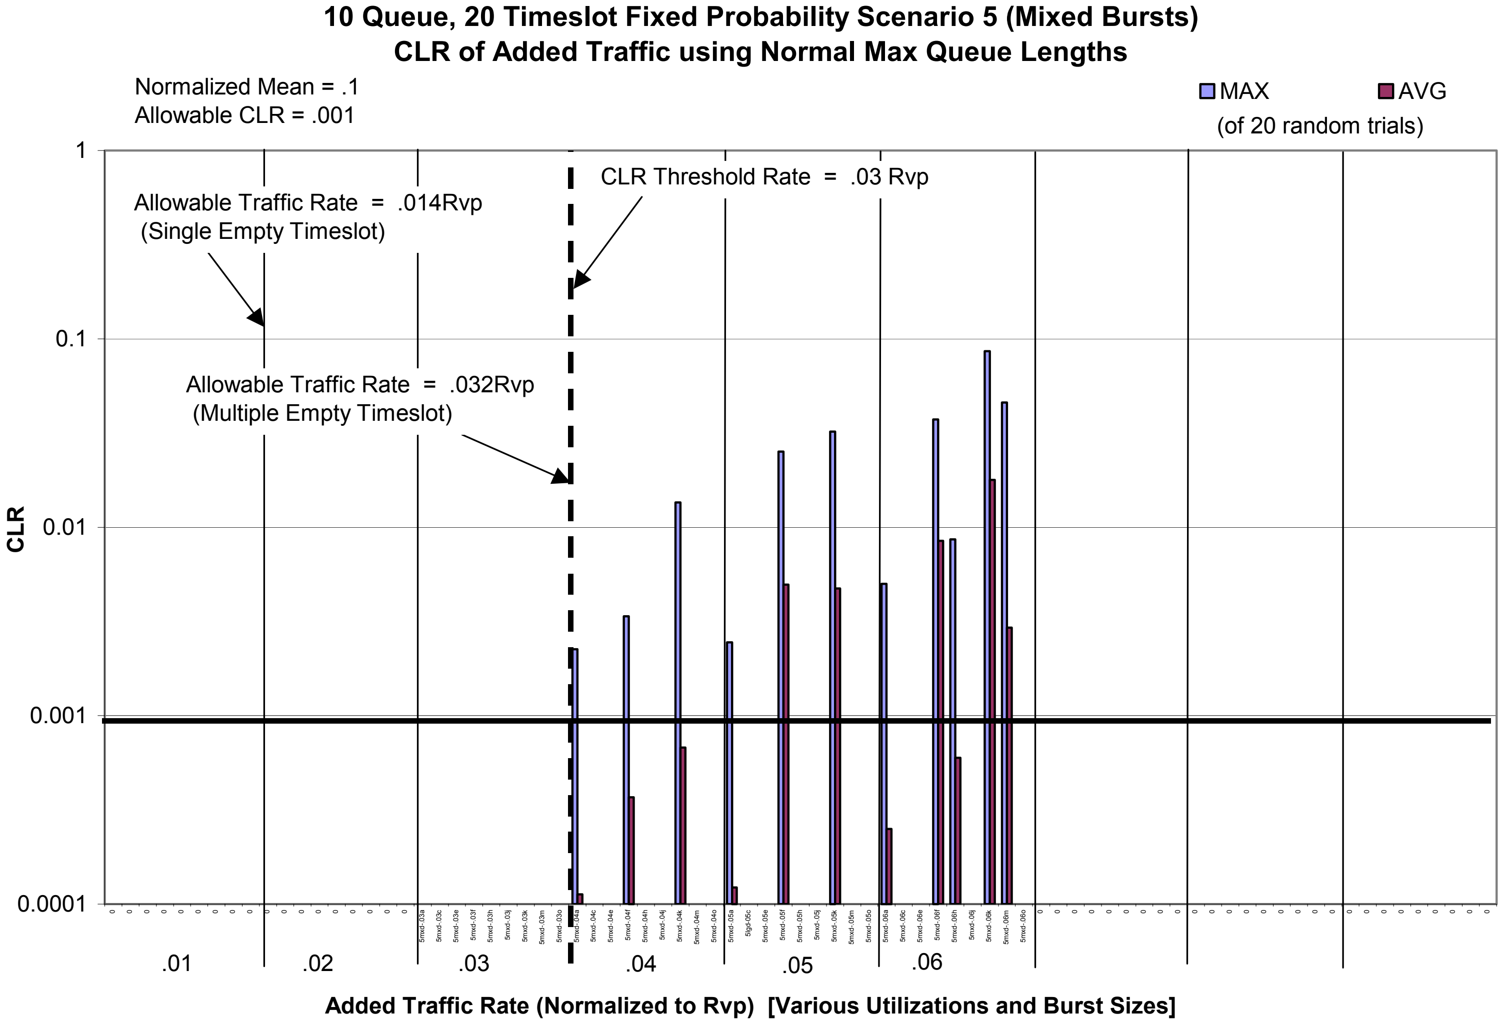

Figure C-11. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 5 (Mixed Bursts)

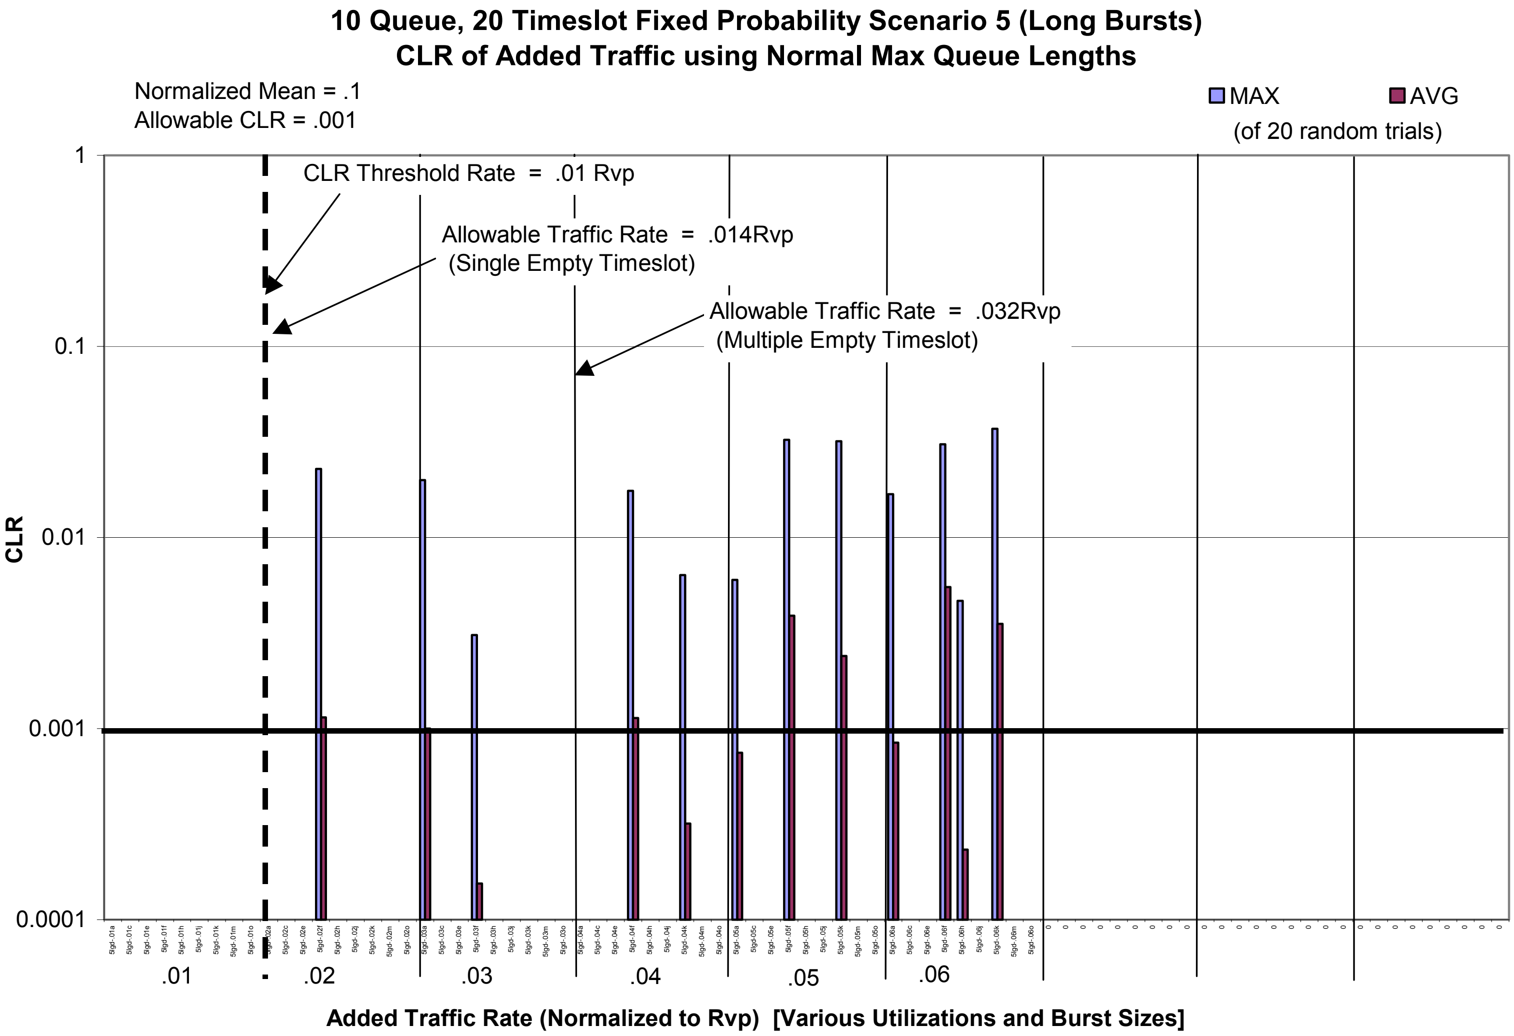

Figure C-12. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 5 (Long Bursts)

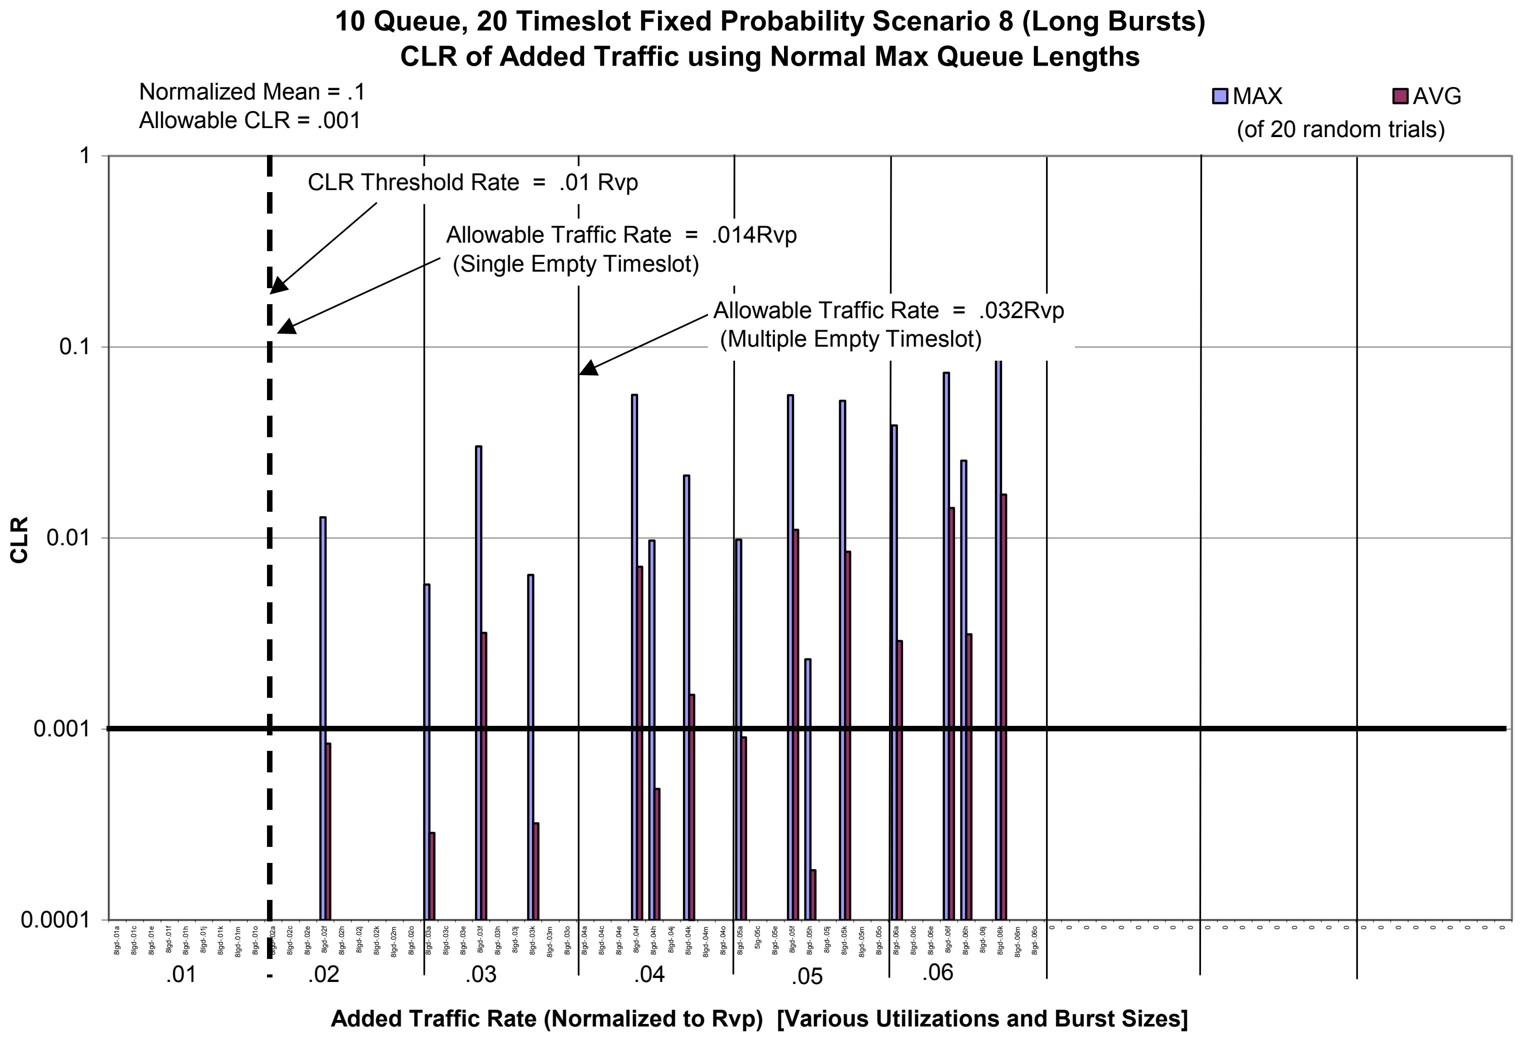

Figure C-13. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 8 (Long Bursts)

Appendix D

Results
Simulated System CLR, 10 Queue Existing Traffic, Added Queue of Augmented Maximum Queue Length

This Appendix contains results of simulations where an “additional” traffic queue was introduced to a WRR server where the existing “normal” traffic was sufficient to fill the virtual path, based on the sum of the SCR’s of the “normal” traffic. The existing traffic occupied 10 traffic streams, each sized such that the resulting WRR schedule was composed of ten or twenty timeslots, depending on the scenario represented. The “added” traffic consists of a single stream passing through a queue sized in accordance with an extended Maximum Queue Length (MQL) as determined from Equation 3.5. Details of the scenario used to obtain the results shown are contained in Appendices A and B. The results are in the form of bar charts where each set of bars represents the maximum and average cell loss ratio resulting from twenty simulations, where the “normal traffic has the rate and burst characteristics indicated along the x-axis.

Each chart shows the allowable CLR threshold as a heavy horizontal line at the appropriate level on the CLR axis. Each chart also shows a CLR threshold rate that shows the boundary below which the CLR for all scenarios at the Added Traffic Rate shown is less than the allowable CLR.

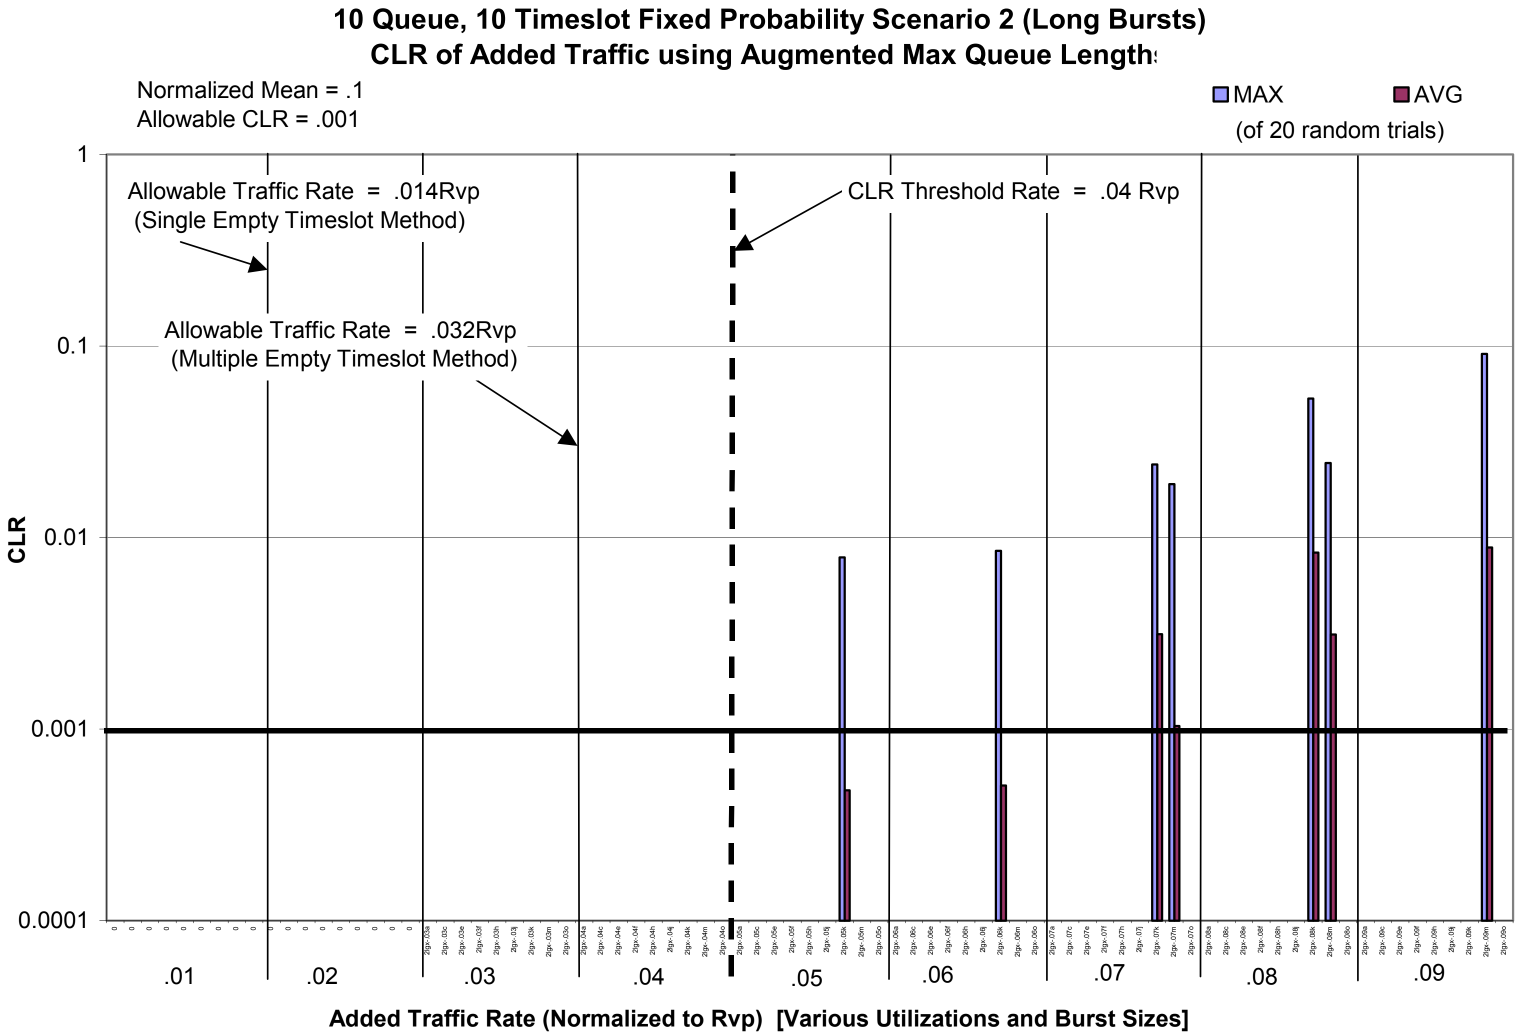

Figure D-1. Cell Loss Ratios, 10 Queue, 10 Timeslot Fixed Probability Scenario 2 (Long Bursts)

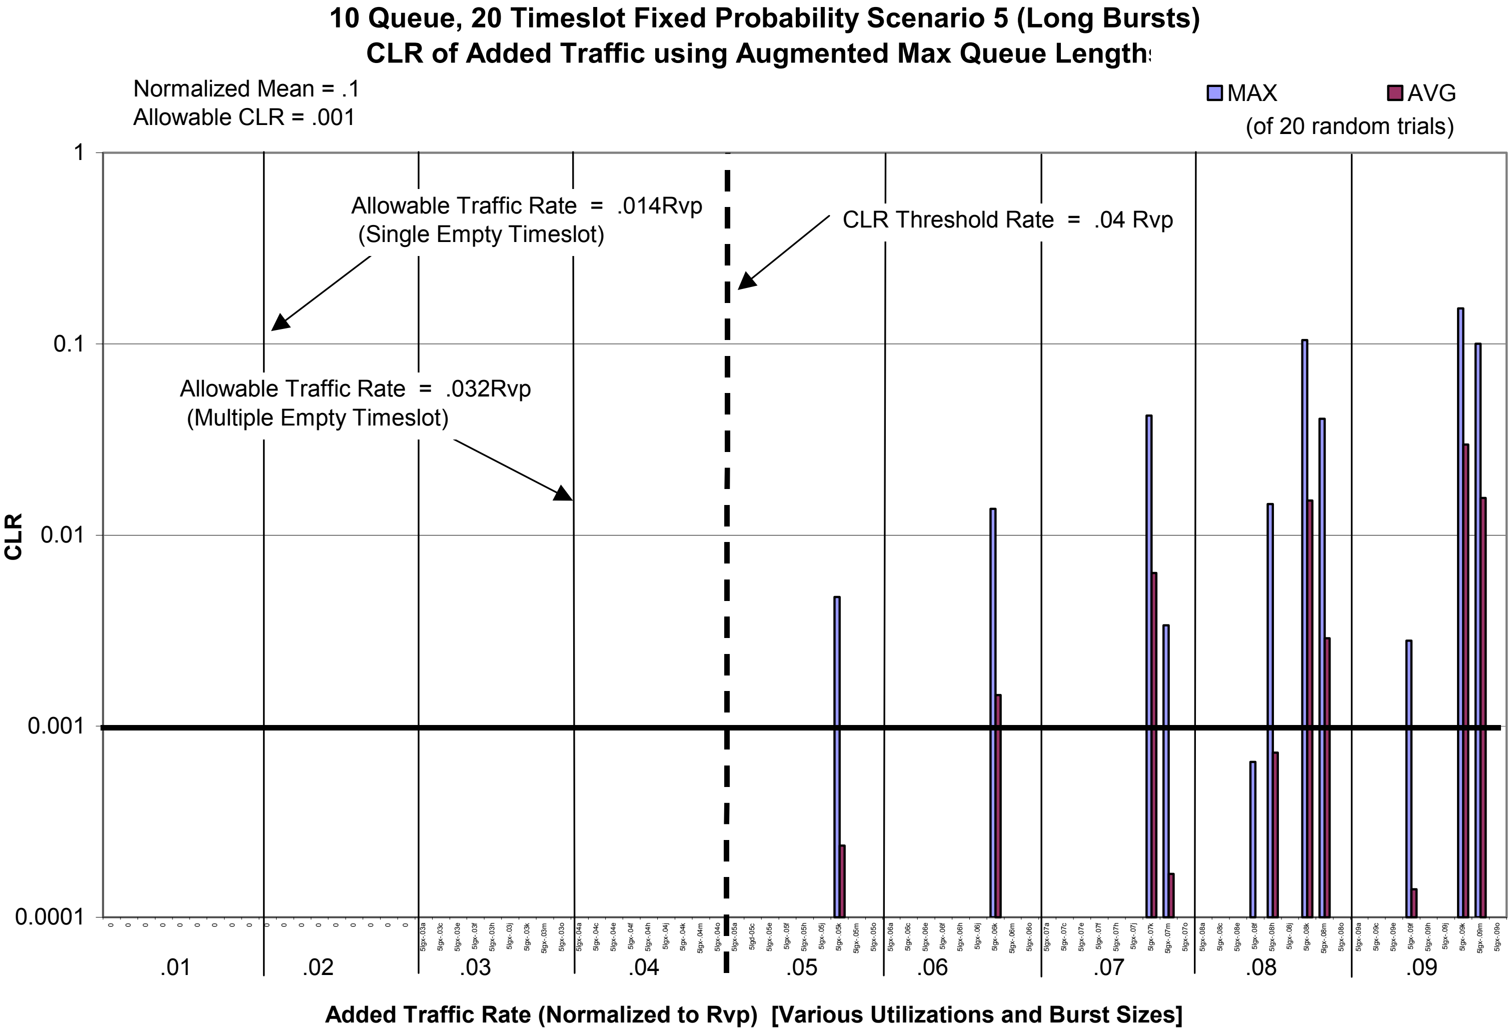

Figure D-2. Cell Loss Ratios, 10 Queue, 20 Timeslot Fixed Probability Scenario 5 (Long Bursts)

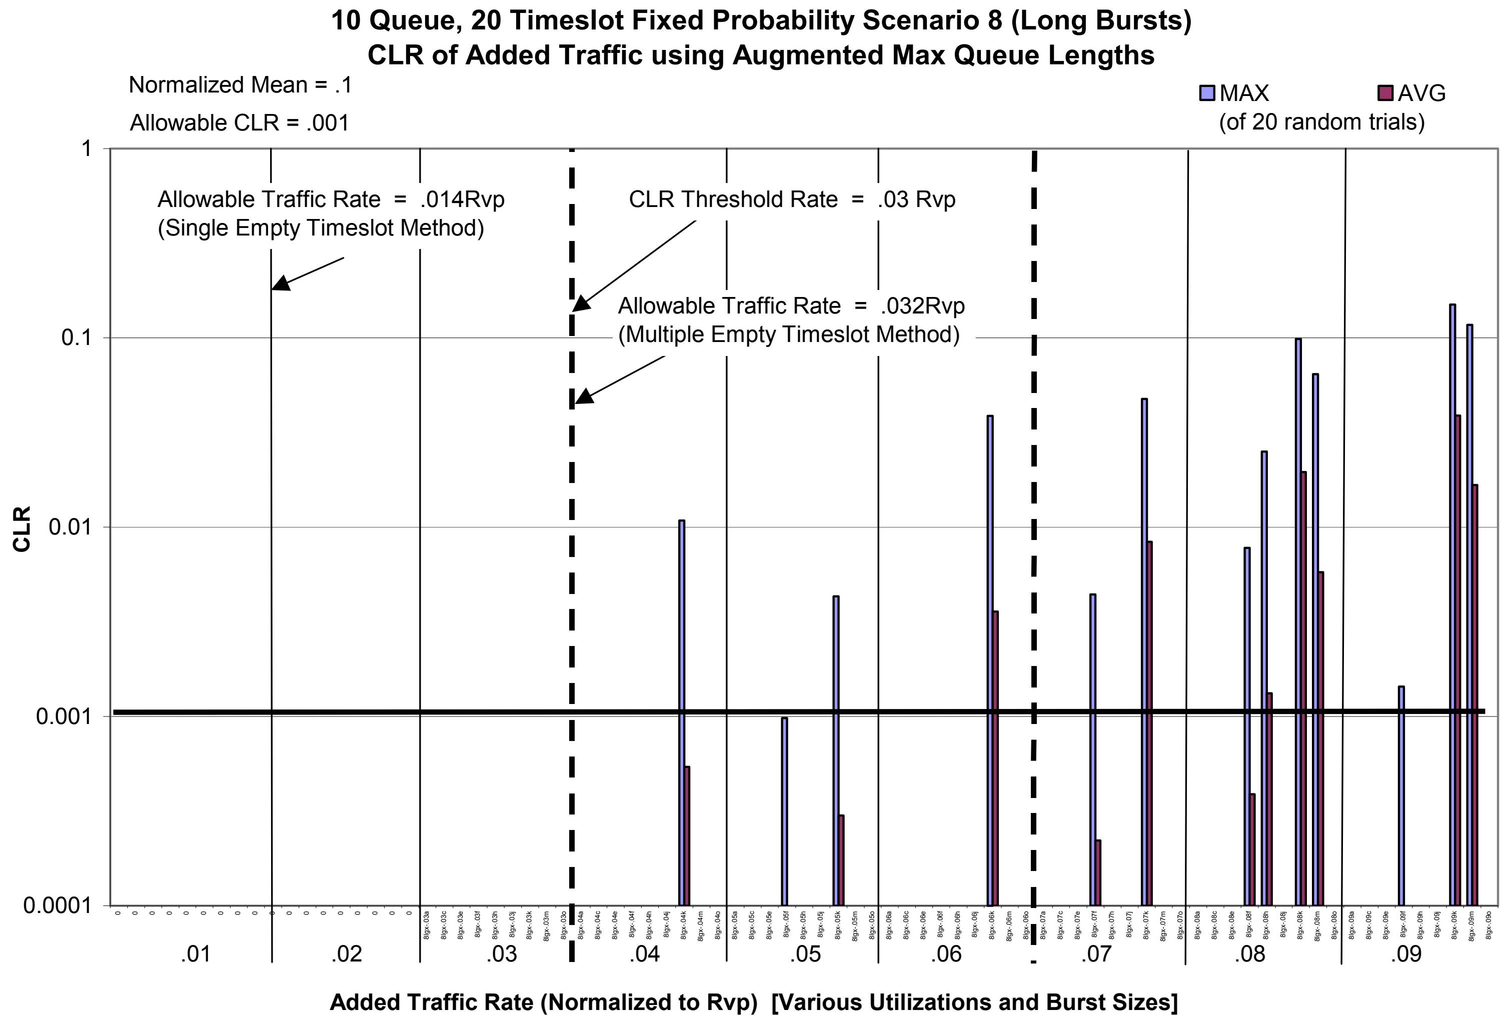

Figure D-3. Cell Loss Ratios, 10 Queue, 20 Timeslot Fixed Probability Scenario 8 (Long Bursts)

Appendix E

Results
Simulated System CLR, 20 Queue Existing Traffic, Added Queue of Normal Maximum Queue Length

This appendix contains results of simulations where an “additional” traffic queue was introduced to a WRR server where the existing “normal” traffic was sufficient to fill the virtual path, based on the sum of the SCR’s of the “normal” traffic. The existing traffic occupied 20 traffic streams, each sized such that the resulting WRR schedule was composed of twenty timeslots, based on the scenario represented. The “added” traffic consists of a single stream passing through a queue sized in accordance with a normal Maximum Queue Length (MQL) as determined from Equation 3.5. Details of the scenario used to obtain the results shown are contained in Appendices A and B. The results are in the form of bar charts where each set of bars represents the maximum and average cell loss ratio resulting from twenty simulations, where the “normal traffic has the rate and burst characteristics indicated along the x-axis.

Each chart shows the allowable CLR threshold as a heavy horizontal line at the appropriate level on the CLR axis. Each chart also shows a CLR threshold rate that shows the boundary below which the CLR for all scenarios at the Added Traffic Rate shown is less than the allowable CLR.

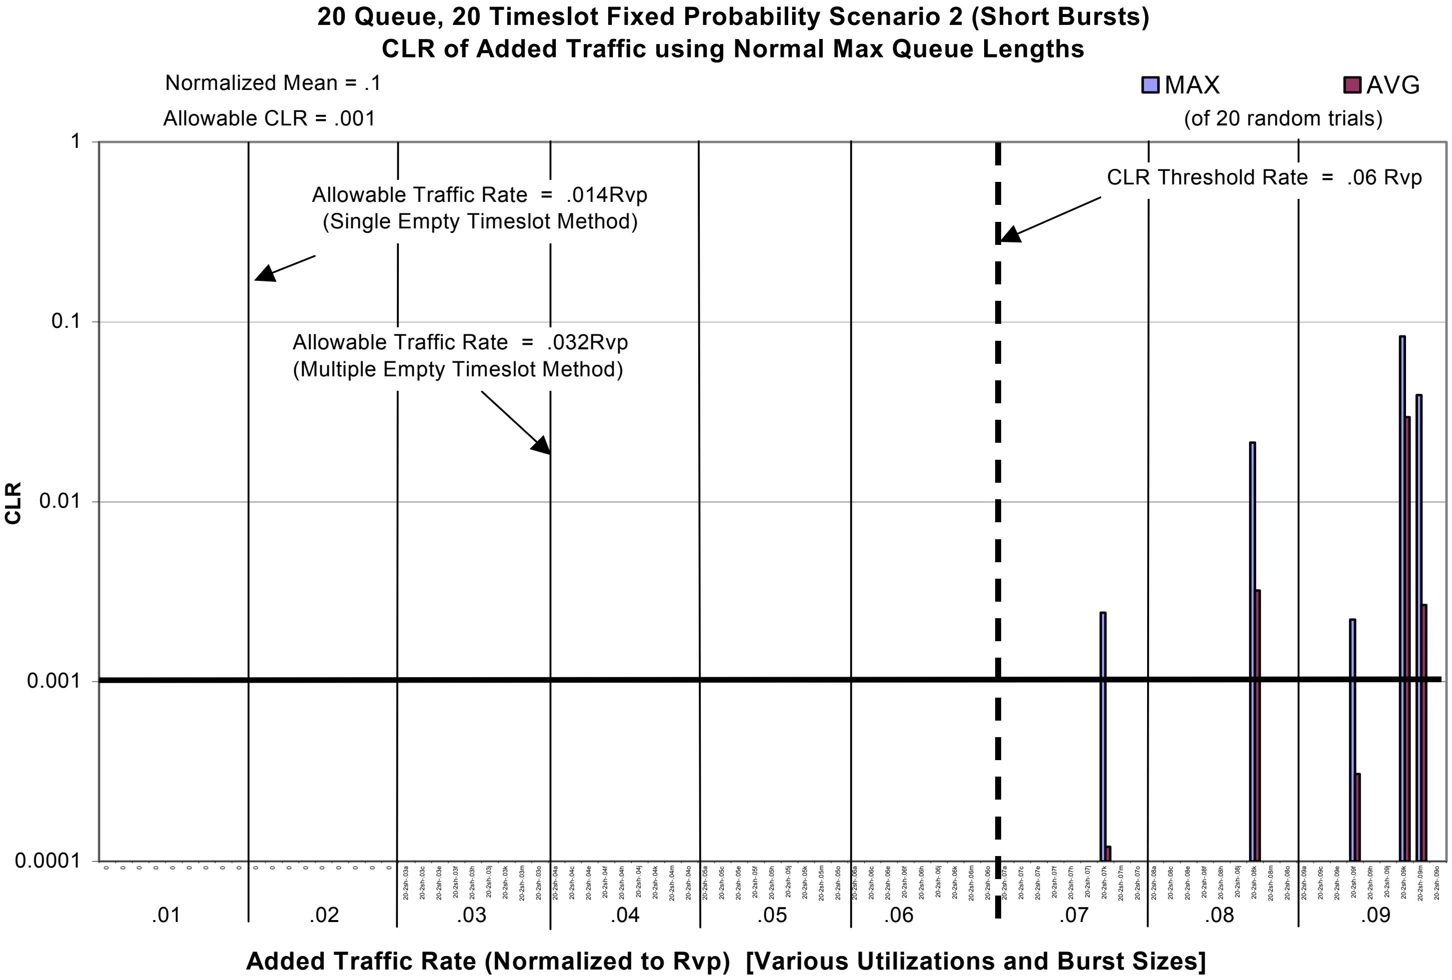

Figure E-1. Cell Loss Ratios, 20 Queue, 20 Timeslot Fixed Probability Scenario 2 (Short Bursts)

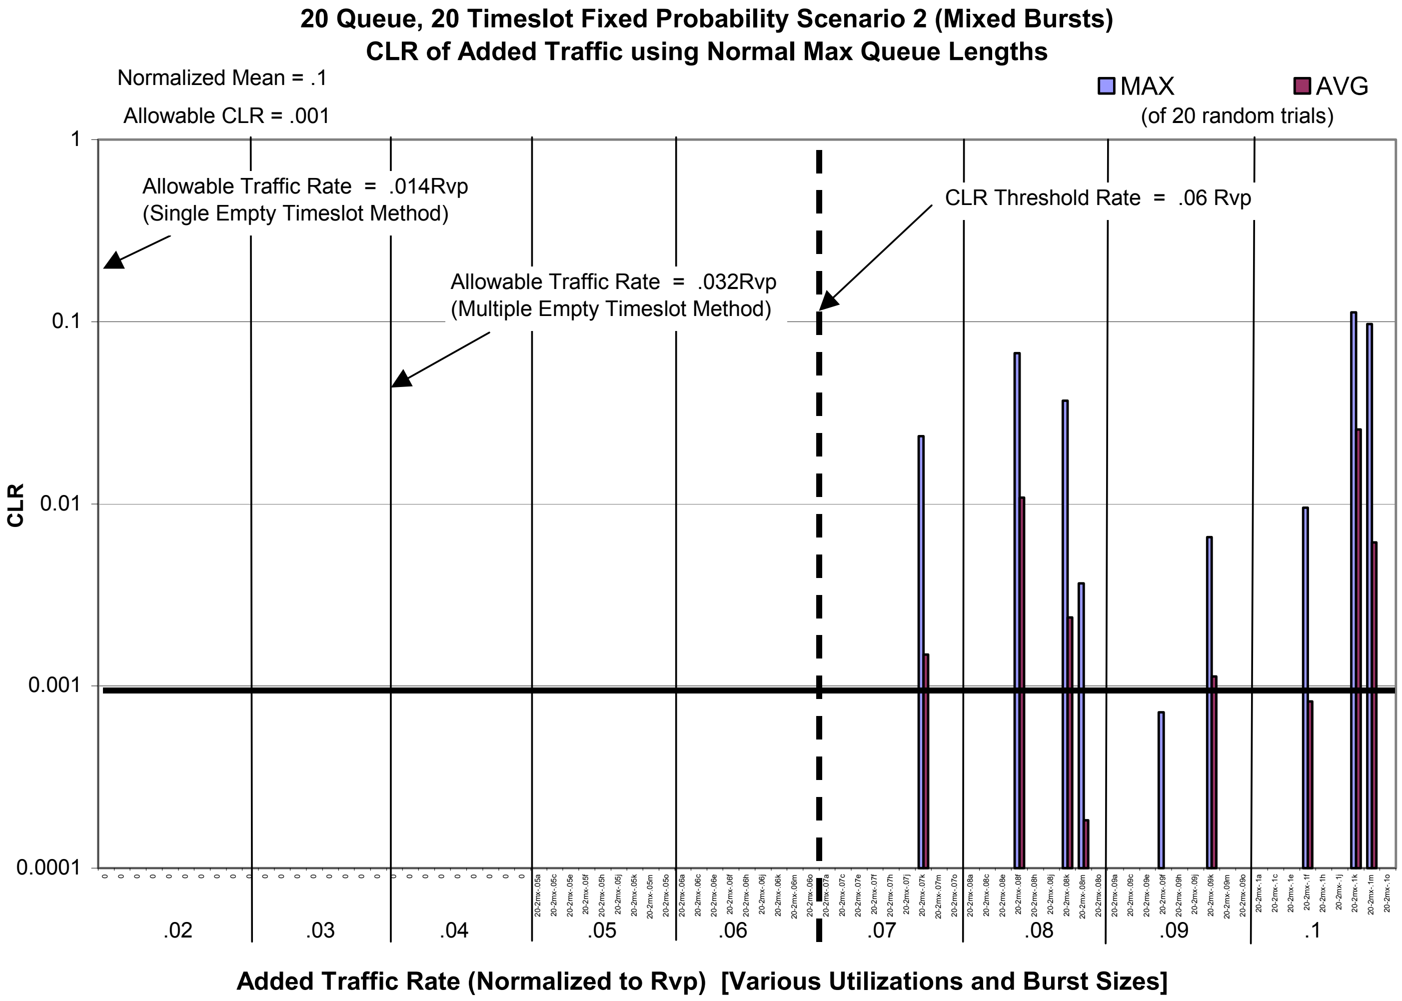

Figure E-2. Cell Loss Ratios, 20 Queue, 20 Timeslot Fixed Probability Scenario 2 (Mixed Bursts)

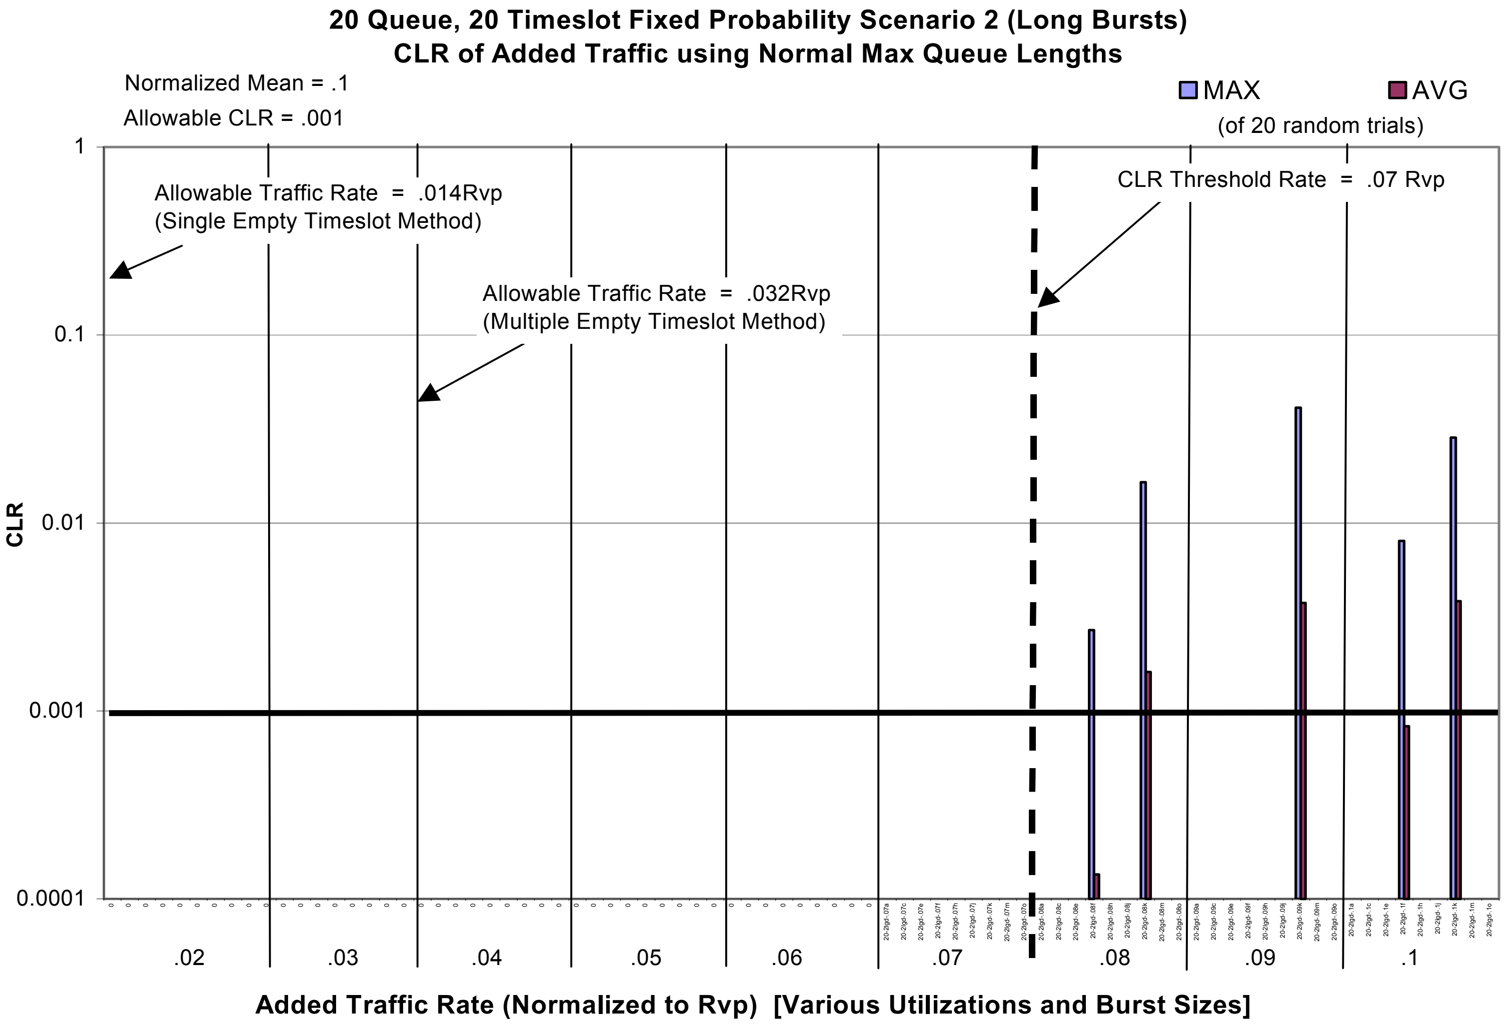
20 Queue, 20 Timeslot Fixed Probability Scenario 2 (Long Bursts)
CLR of Added Traffic using Normal Max Queue Lengths

Figure E-3. Cell Loss Ratios, 20 Queue, 20 Timeslot Fixed Probability Scenario 2 (Long Bursts)