ABSTRACT

This thesis reports on the performance of a Packet Discard Strategy used as basis for congestion management of data traffic associated with Unspecified Bit Rate transfer service in ATM data networks. With this Packet Discard Strategy, when the cell occupancy of the ATM switch buffers is such that overflow is threatened, congestion is relieved by accepting into the congested buffers only cells from data frames already in transit through the switch and by rejecting all cells from new data frames. This regime continues until the congestion is relieved. In this thesis, we provide a conservative queueing system, where group arrivals is assumed, for modelling the behaviour of and ATM switch operating with and without the Packet Discard Strategy. Using a simulation model we compare the numerical results with those from the queueing model and investigate the factors affecting the relationship existed between the two results. For comparison, we also use a variant of the queueing system to model a Cell Discard Strategy where cells are discarded whenever the buffer is full. With the results from the queueing models, we also investigate the parameters affecting the performance of the discard strategies in terms of goodput and packet loss probability.

Simulations show that our simplification of using Group arrivals in the queueing analysis is reasonable and the results from the queueing model represents a conservative estimation of the performance of the system. We also observe that the switch operating with the Packet Discard Strategy remains 100% goodput even under sustained overload while the goodput of the switch operating with the Cell Discard Strategy degrades once the switch is overloaded.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

The evolution of high speed workstations and servers, and the emerging requirement for multi-media applications, is producing a new demand for high speed communication networks. Multi-media applications themselves bring the need for integrated network operation to handle both real-time, delay-sensitive traffic as well as data traffic that is loss-sensitive but delay-tolerant. Asynchronous Transfer Mode (ATM) is a cell based switching and multiplexing technique which can provide such integrated operation. It has been adopted by the International Telecommunications Union (ITU) for the implementation of the Broadband Intergrated Services Digital Networks (B-ISDN). It has also been adopted by the ATM Forum for private workplace, backbone and campus-wide networks.

With Asynchronous Transfer Mode, traffic is carried by ATM cells that are 53 octets long containing a 5 octet header and a 48 octet payload. These cells flow along entities called Virtual Channels (VCs) that are identified by their respective virtual channel identifiers (VCIs). These VCs are multiplexed within Virtual Paths (VPs) that are identified by their respective Virtual Path identifiers (VPIs). These identifiers are multiplexing labels used to identify different cell streams and are contained in ATM cell headers. In terms of protocol architecture, the ATM layer consists of the lower VP sublayer and the higher VC sublayer. The ATM layer sits above the physical layer and below the ATM Adaptation Layer (AAL). The various AALs provide different services through segmentation and reassembly of higher-layer data units, error detection and re-transmission. AAL1 and AAL2 are intended for delay-sensitive traffic while AAL3/4 and AAL5 are intended for loss-sensitive data traffic[1][2].

Communication in ATM networks is connection-oriented requiring a virtual connection, associated with a particular VPI/VCI, to be established at call set-up time. In accepting connections over public ATM networks, the network provider needs to confirm that sufficient network resources are available to sustain the connection end-to-end at its required Quality of Service and without affecting the service guarantees of the existing connections. At the ATM layer, Quality of Service (QoS) is expressed in terms of cell delay, cell delay variation, and cell loss ratio. To assist in the process of Connection Admission Control the network user is required to characterise its traffic and specify the required Quality of Service. The traffic may be characterised by its average bit rate, burstiness, peak duration and other such measures. However for purposes of ATM connection admission, traffic on a requested Deterministic Bit Rate (DBR) connection should be characterised solely by the peak cell rate (PCR), while traffic on a requested Statistical Bit Rate (SBR) connection should be characterised by peak cell rate, sustainable cell rate (SCR) and intrinsic burst tolerance (IBT) [3]. If adequate network resources are available to sustain a new connection at its required Quality of Service at the indicated peak cell rate, or peak and sustainable cell rate and intrinsic burst tolerance, the connection would be accepted by the network which would then monitor the traffic flow to ensure that the agreed rates are not violated. The provision applies to connections equally at the VP and the VC sublayers. Connections being set up this way are called resourced connections.

On customer premises or private ATM networks, not all connections would be resourced. Certainly, connections set up to carry real-time traffic, such as voice and video, will require bandwidth guarantees so as to achieve constrained delay [3]. However connections set up for data traffic can achieve lower average latency if service is provided on a non-resourced, statistically-multiplexed basis using virtual channels where the bit rates are not specified. Accordingly, the Unspecified Bit Rate (UBR) transfer service is recommended in [3] for such services. With UBR connections, only peak cell rate is required to be specified with no Quality of Service guarantees (class U QoS)[3]. The UBR transfer service could be an important option for data traffic carried by protocols such as TCP and for data traffic emanating from legacy LANs. The use of UBR connections to reduce latency may also apply to ATM Wide-area Networks where non-resourced VC connections can be multiplexed within resourced VP connections [4]. Without a congestion management strategy such a data service may suffer cell-loss and cell-loss multiplication wherein even a single lost cell renders unusable all the other cells carrying in their payloads data segments from the same frame as the lost cell [3].

To address this problem, Frame Discard is proposed in [3] and [4] to be an additional traffic control and congestion control function for UBR transfer service provided at the ATM layer. The Frame Discard strategy operates such that if one cell of a frame (packet) needs to be discarded in a network element, the network may discard as many cells of the frame as possible. Different Frame Discard strategies have been suggested in the literature [4]-[9]. In this thesis, we report on the operation and performance of the Packet Discard Strategy suggested in [4] and our paper [5] to be used as basis for congestion management of UBR traffic in ATM data networks. In our analysis, we measure performance of a discard strategy in terms of the goodput characteristic of an output port and in terms of the packet loss probability of the respective output buffer. Here goodput refers to the output traffic intensity counting only cells which are part of undestroyed packets. With this Packet Discard Strategy, when the cell occupancy of the ATM switch buffers is such that overflow is threatened, congestion is relieved by accepting cells only from data frames already in transit through the switch and by rejecting all cells from new data frames as well as cells from previously rejected data frames. This regime continues until the congestion is relieved.

The Packet Discard Strategy is simple and maintains 100% goodput even in the presence of gross overload in both small and large networks and works well where the ATM data service is provided for terminals using adaptive, sliding window protocols such as TCP in their upper layers. A feature of the Packet Discard Strategy is that it operates at the ATM protocol layer. However, the use of the Packet Discard Strategy requires that data frame boundaries be identifiable and consequently must necessarily operate in conjunction with AAL5 where the end of frame is indicated in the ATM header [3].

In this thesis, our analysis of the Packet Discard Strategy assumes a wide range of data-traffic intensities, including gross overload, and with various threshold values. Different packet-size distributions are considered, including those with both fixed and variable sized packets, and both analytical and simulation results are presented. In our analytical model, we make the simplified assumption that cells belonging to the same packet arrive at the same instant. We refer this arrival pattern as group arrival. In our simulation model, cells carrying segmented packets arrive at the switch inputs contiguously. We refer this arrival pattern as cell arrival. In this thesis, we examine the relationship between cell arrivals and group arrivals in terms of buffer occupancy and investigate the various factors affecting this relationship. These factors include the number of links connected to the switch, the traffic intensities and the output link rate. For comparison, we also use a variant of the analytical model to consider a Cell Discard Strategy where cells are discarded whenever the buffer is full. In all our analysis, the Poisson arrival process is employed to help us understand the behaviour of an ATM switch and the two discard strategies.

In chapter two of the thesis, we describe the cause of congestion in ATM data networks where data traffic is non-resourced employing Unspecified Bit Rate transfer service. Then, we introduce the Packet Discard Strategy, and discuss the selection of parameters to achieve best performance. In this chapter, we also discuss our modelling of the Packet Discard Strategy and the Cell Discard Strategy.

Chapter three introduces our fundamental queueing system for modelling an ATM switch, namely the $M^{[x]}/D/1/\infty$ queue. Here, we present the analysis of the queue which forms the basis for the analysis of its variants, namely the $M^{[x]}/D/1/B$ queue and the $M^{[x]}/D/1$ queue with threshold. An algorithm for effective numerical calculation of queueing statistics is also presented.

Chapter four presents the modelling of the Packet Discard Strategy using the $M^{[x]}/D/1$ queue with threshold. Through the queueing analysis, we explore the relationship of the equilibrium queue occupancy distributions between this queue and the $M^{[x]}/D/1/\infty$ queue. The goodput and the loss probabilities of the queue are also derived.

Chapter five presents the modelling of the Cell Discard Strategy using the $M^{[x]}/D/1/B$ queue. Through the analysis, we explore the relationship between the equilibrium queue occupancy distribution of this queue and that of the $M^{[x]}/D/1/\infty$ queue. Moreover, we derived the goodput and the loss probabilities of the queue.

Chapter six presents the numerical results from both queueing and simulation models. The implications of the results are also discussed. In this chapter, we compare the buffer occupancy distribution of the queueing and simulations models to investigate their relationship and the factors affecting this relationship. For the case with the Packet Discard Strategy, we compare the packet loss probability of the queueing and simulations models and a similar relationship is observed. Then, with the results from the queueing model we discuss the selection of the threshold value for the Markovian traffic pattern. In addition, we investigate the parameters affecting the performance of the Packet Discard Strategy and the Cell Discard Strategy. Finally, with the queueing models we compare the goodput of the two discard strategies.

Chapter seven concludes this thesis and discusses possible further work. With our analysis, we observe that the performance of an ATM switch operating with the Packet Discard Strategy maintains 100% goodput even under severe overload. On the other hand, the performance of that operating with Cell Discard Strategy degrades once the port is overload. We also observed that the our simplified queueing models are good approximations and are conservative. These results provide a basis for selecting the threshold for Packet Discard Strategy.

CHAPTER 2

CONGESTION IN ATM DATA NETWORKS

In this chapter, we first discuss the congestion problem associated with Unspecified Bit Rate (UBR) transfer service of ATM data networks. We then consider in detail the Packet Discard Strategy (PDS), which we are investigating for use in conjunction with the UBR service. The parameters to be selected for the PDS to achieve the best performance are also discussed.

2.1 Data Traffic in ATM Networks

In ATM networks, data in packets are segmented at the AAL layer into ATM cells so that the packets can be transmitted through the ATM layer. At the ATM layer, cells carrying real-time, delay sensitive traffic are transmitted using the CBR or SBR transfer service with resourced connections. Such connections are managed by Connection Admission Control and Usage Parameter Control to ensure the agreed QoS is met and the agreed transmission rate of cells is not violated [3]. On the other hand, cells carrying loss-sensitive but delay-tolerant data traffic could be transmitted using the UBR transfer service with unresourced connections [3]. Such connections are not managed in the same manner. The transmission rate of cells through these connections is not controlled, however, the cells on the connections must be transmitted at a lower priority than that used for the CBR and SBR services. Our analysis and simulations only concern UBR data traffic.

To illustrate the congestion problem with data traffic, we first introduce the most general topology of an ATM network as shown in Figure 2.1. Such a topology comprises ATM switches and Terminal Equipment (TEs). The switches with TEs connected to them are local switches while those with no TEs connected to them are transit switches. Data traffic is generated in TEs and transmitted over the network to other TEs.

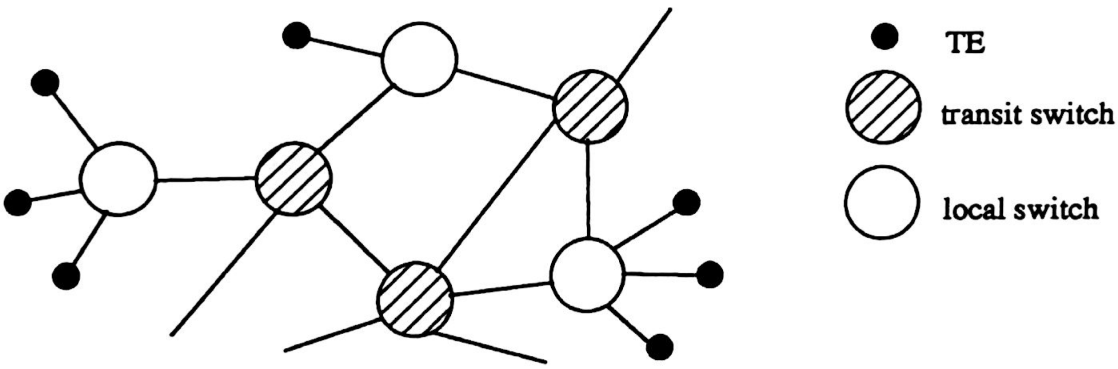

Figure 2.1: The most general topology of an ATM Network.

Next, we consider a generic ATM switch architecture, for either the local or transit switch, as shown in Figure 2.2. At the ATM layer, the switch can be divided into three component functional blocks namely the cross connect, the output buffers and the input/output links. This switch has N input links feeding into it and N output links going out from it with each input and output link having the same bandwidth. A cell coming from an input link is transmitted to the appropriate output link according to the virtual connection identifier, i.e. the VPI/VCI label, specified in the header of the cell. Because of the switching action, cells from different packets become interleaved. The output buffers are for temporary storage of cells that cannot be transmitted immediately to an output link. We consider in our architecture that the cross-connecting function is non-blocking. That is, cells presented to the input ports are delivered immediately to their destination output buffers without loss. However, cells from multiple input links can converge into one output link causing the associated output buffer to fill. In the most extreme case, the aggregate rate of cells coming from multiple input links to a particular output link can be N times the bandwidth of that link. If the output buffers do overflow, cells will be lost with attendant cell loss multiplication. Cell loss multiplication refers to the process whereby a whole packet is destroyed due to the loss of one of its constituent cells. Therefore, a congestion management strategy based on packet discard is recommended for connections used by UBR transfer service in ATM networks [3] [4].

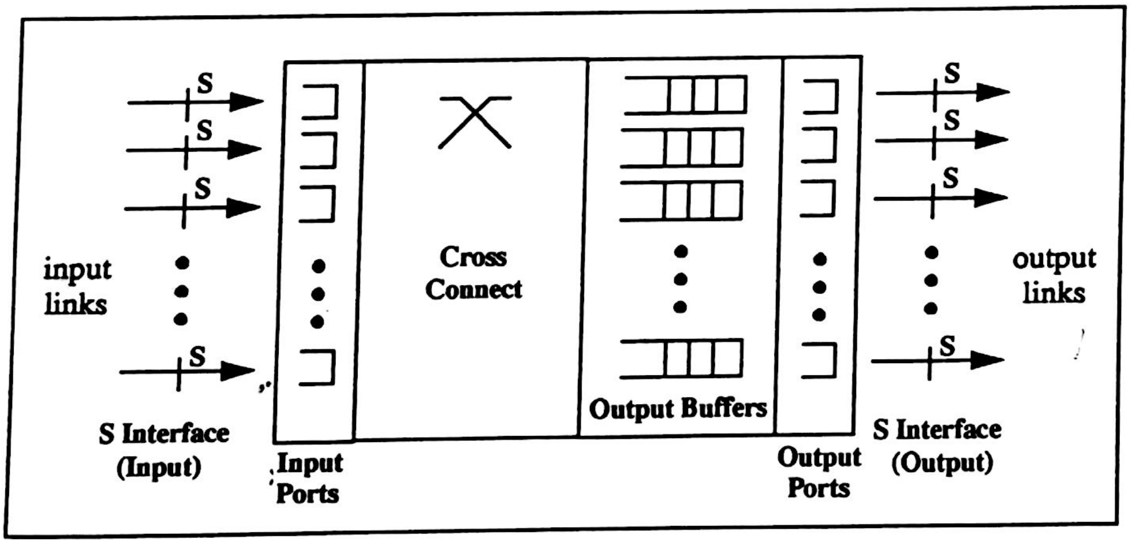

Figure 2.2: An ATM Switch

2.2 The Packet Discard Strategy

In order to handle the congestion problem, we propose a congestion management strategy, the Packet Discard Strategy (PDS) [4], to be used as the Frame Discard management scheme for UBR transfer service [3]. With the PDS, when the buffer fill is at or over a specified threshold, the output buffer rejects cells from new packets, i.e. packets which have had no component cell yet admitted. However, even though over threshold, it continues to accept cells from packets already in transit. When the buffer fill returns under the threshold, the output buffer will then only accept cells from packets from which no cell has previously been discarded. This regime continues until congestion is relieved. With such a strategy, the integrity of packets is conserved when they are transmitted through the switch. Each packet is either transmitted through the switch in full or discarded completely. In this thesis, we consider the performance of the PDS with fixed threshold values. In other work reported in [6], a fixed threshold was not employed. Instead a decision is made whether to accept or reject each new packet based on the likelihood of buffer overflow given the current buffer occupancy and the number of packets still in transit.

At ATM layer data is transmitted in ATM cells. The use of AAL5 in conjunction with the PDS obviates the need to reassemble the ATM cells into data packets since with AAL5 a start of packet indicator is carried with the ATM cell header. Of course packet discard can also be exercised when AAL3/4 is employed but in this case packet reassembling would be required.

The Packet Discard Strategy is relatively simple and maintains 100% goodput under sustain overload. The simplicity of the PDS also lies in the fact that it requires no control signals. Another feature of the PDS is its flexibility, as it can be applied with equal facility to both Customer Premises and Wide-area ATM Networks.

With the Packet Discard Strategy, the two parameters to be selected are the threshold value and the buffer size. In order to maximize the goodput, the proportion of the buffer below threshold is required to be large enough as to prevent buffer underflow when the buffer occupancy returns below threshold and the remnant of the packets which have had some cells discarded, continue to be discarded. In order to minimize buffer overflows, the proportion of the buffer above the threshold is required to be large enough to accommodate the packets in transit at the time the decision was taken to reject new packets. In other words, the threshold value and the buffer size should be engineered to avoid overflow and underflow. The occurrence of underflow and overflow is shown pictorially in Figure 2.3.

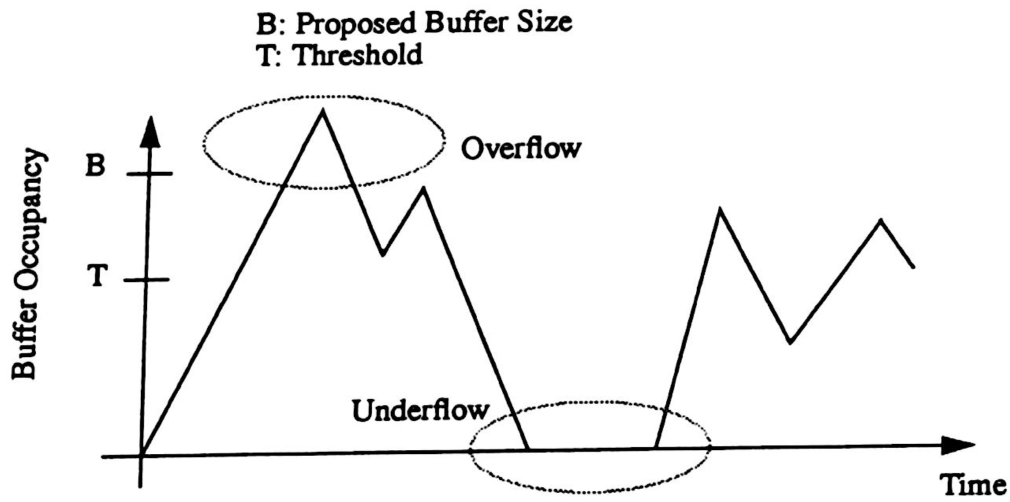

Figure 2.3: Buffer Occupancy vs Time

In this thesis, we investigate the relationship between the size of the buffer below threshold and the probability of underflow. In our queueing analysis, we will investigate the relationship between the probability of underflow, the packet loss probability and the goodput. The investigation of the buffer size required above the threshold, i.e. the overflow problem, is not analysed in this thesis but is the subject of ongoing work. However, we will run some simulations to show how overflow above threshold will affect the goodput of the PDS and will compare these two simulations results with the those from the queueing analysis of the PDS and the Cell Discard Strategy.

In our analysis of the PDS, we investigate the goodput, associated with the cell flow at an output port of a switch, against offered load. In this context, the offered load to an output port refers to the average rate of cells directed to the output port divided by the maximum possible transmission rate of cells on the link connected to the output port. In this thesis, we describe as underload the case when the offered load is below unity and as overload the case when the offered load is above unity.

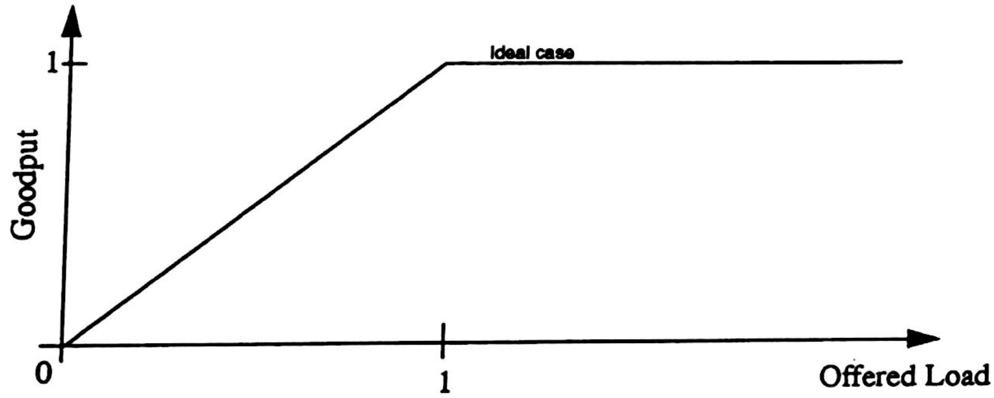

Figure 2.4: Goodput against Offered Load

We show in Figure 2.4 the ideal goodput characteristic with respect to offered load. Recall that goodput refers to the output traffic intensity counting only cells which are part of undestroyed packets. Here goodput characteristic refers to the variation of goodput over a range of offered load. With the ideal case all the cells are transmitted through the switch without loss when the offered load is less than or equal to unity, and the switch output port remains 100% goodput once the offered load is larger than unity. When the goodput is 100%, the output traffic intensity is 100% and all the transmitted cells are part of undestroyed packets, even if the offered load is more than 100%. In other words, with the ideal case the switch always maintains maximum possible output traffic intensity and all the transmitted cells are part of undestroyed packets. In this thesis, we use queueing analysis and simulations to compare the goodput characteristics, associated with the discard strategies, with the ideal goodput characteristics.

2.3 Modelling of the Discard Strategies

In our simulations, we assume the packets are generated at the traffic sources according to a Poisson generating process, and we place no restrictions on the distribution of the packet size. With this generating process, cells will be presented to an input port as a compound Poisson process and then to the respective output buffer as a superposition of the output streams from the respective input ports. Our simulation analysis is capable of examining this model for the case with and without the PDS. However, because the model has to distinguish at the output buffer cells from new packets and cells from packets in transit, the queueing analysis of this system when the PDS is in operation becomes intractable. Moreover, with this model, if the size of the output buffer is limited, it is extremely difficult to relate cell loss to packet loss in the queueing analysis in order to calculate the goodput of the system. Consequently in our analytical modelling, we make the simplifying assumption that packets generated at the various sources arrive at the output buffers directly as a packet. However, we still assume the packets arrive as a Poisson arrival process and place no restrictions on the distribution of the packet size.

2.3.1 The Analytical Model

The $M/D/1$ queue with group arrivals, denoted as $M^{[x]}/D/1$, is used to model the performance of an ATM switch. The operation with the Packet Discard Strategy is investigated using the $M^{[x]}/D/1$ queue with threshold. It models an output port where new frames are discarded when the buffer fill is at or above the threshold. The operation with Cell Discard Strategy is investigated using the $M^{[x]}/D/1/B$ queue where B denotes the size of the output buffer. With this model, cells are discarded whenever the buffer is full. Using this simplified arrival process, we are able to calculate the worst case performance of the PDS and with the Cell Discard Strategy. Hence, we have group arrivals in our queueing models and cell arrivals in our simulations. Note that with group arrivals cells belonging to the same packet are not interleaved. On the other hand with cell arrivals, cells belonging to a segmented packet are interleaved with cells belonging to other segmented packets. In our queueing models, the output buffer fill is modelled as the queue length, and the corresponding output port is modelled as the server of queueing system. Our simulations show that our simplified queueing model, that assumes group arrivals, yields performance results which are conservative. The analysis of these queues are presented in the following chapters.

2.3.2 The Simulation Model

The switch model that we used as the basis for our simulations is shown in Figure 2.5. This switch has N input links feeding into the N input ports with each input and output port having a bandwidth R. However, the rate at the output port can be varied if necessary. For the sake of clarity, only one output port is considered. Cells from segmented packets can arrive at any one of the N input ports. An incoming cell is immediately switched to the decision box of the designated output port. If the cell is accepted, then it is placed into the output buffer to be transmitted onto the outgoing link.

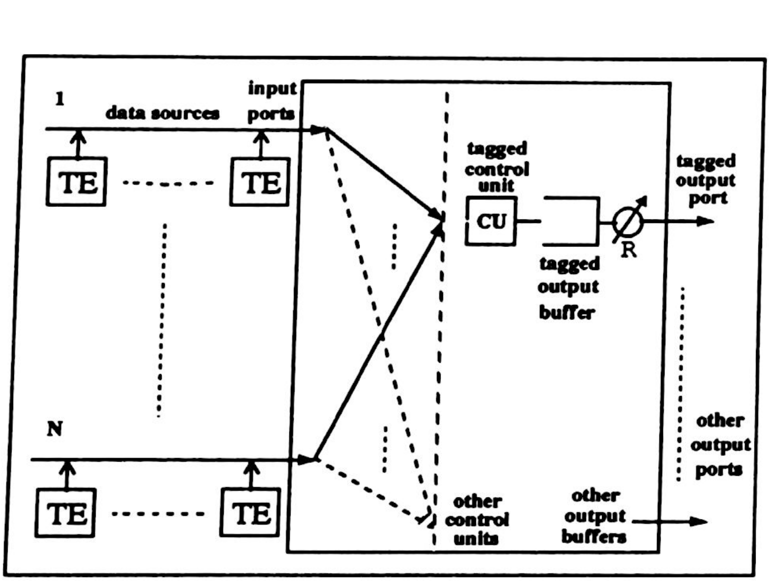

Figure 2.5: The Simulation Model

Each TE shown in Figure 2.5 acts as a traffic generator. The traffic generated by TEs that are connected to a given link is multiplexed according to some MAC protocol before reaching the ATM switch. A typical protocol that may be considered is DQDB. This results in an interleaving of packets arriving at the designated output port. Furthermore, the input to the ATM switch output buffers is always a superposition of various streams from different sources.

There are two major differences between the analytical model and the simulation model. Firstly, the analytical model assumes group arrivals while the simulation model assumes cell arrivals with interlaving of packets and superposition of cells. The second one is that with the analytical model, the number of cells arrive within a service period is not limited while with the simulation model, the number of cells arrive within a service period is N if the output link is not slower than the input link.