CHAPTER 6

NUMERICAL RESULTS

In this chapter, we present the numerical results from both the queueing models and the simulations, and discuss the implications of such results. Firstly, we compare the numerical results from the $M^{[x]}/D/1/\infty$ queueing model and those from the simulations on the ATM switch, for the case with no discarding. The factors affecting the accuracy of the queueing model are investigated. The queue statistics of the $M^{[x]}/D/1/\infty$ queueing model are then presented for studying the performance of an ATM data switch. Secondly, we compare the numerical results from the $M^{[x]}/D/1$ queueing model with threshold and those from the simulations on the ATM switch when the PDS is operational. The queue statistics of the $M^{[x]}/D/1$ queueing model with threshold are then presented for studying the behaviour of the PDS and selecting the threshold value. Moreover, we use our queueing models to study the goodput characteristics, with different threshold values for the PDS and with different buffer sizes for the Cell Discard Strategy. Finally, we compare the performance of the PDS and that of the Cell Discard Strategy. We also present the results of two simulations to illustrate the effects of buffer overflows above threshold when PDS is operational.

In both the queueing models and the simulations, we employ the Poisson arrival process to understand the behaviour of the systems under consideration. For this purpose, our calculations and simulations are performed over a range of traffic intensities, including gross overload. Different packet-size distributions as listed in Table 7.1 are used. The distributions designated as MM_1 and MM_2 correspond to those which may be regarded as typical for a LAN and MAN respectively and therefore would have relevance to an ATM data network. For comparison, we also considered fixed length packet sizes. The packet sizes designated F_{08} and F_{45} are close to the respective means of MM_1 and MM_2. For the simulations with the distributions F_{08} and MM_1 the run time corresponded to 2,000,000 cell periods, and for F_{45} and MM_2 10,000,000. Instead of running the same simulation for many times to get an averaged result, we run each simulation long enough to produce quasi-static results.

Packet Size Distributions	Size	Probability
F_{08}	8	1
F_{45}	45	1
MM_1	2	0.73
	3	0.03
	11	0.03
	21	0.03
	32	0.18
MM_2	1	0.2
	10	0.1
	30	0.1
	50	0.1
	65	0.2
	80	0.3

Table 7.1: Packet Size Distributions

In all the simulations we performed, ten TEs were generating traffic in each link and the traffic intensity at each input port was kept at a constant 80% of the link capacity while the amount of traffic routed from each input port to the designated output link could produce a loading ranging from 70% to 500%.

6.1 Comparison of Cell Arrivals with Switching and Group Arrivals

In this section, we compare the numerical results from the $M^{[x]}/D/1/\infty$ queueing system and those from the simulations of the ATM switch. Recall that with the $M^{[x]}/D/1/\infty$ queueing system we have group arrivals while in the simulations we have superposition of cell arrivals. In this comparison, no discard strategy is employed and we compare the characteristics of the two arrival patterns in terms of the tail distribution of their respective output buffer occupancy. We observe that the tail distributions associated with the two arrival patterns are close to each other and this relationship is affected by three factors, namely the number of input links to the switch, the traffic loading of an output link and the service rate of an output link.

6.1.1 Number of input links

Firstly, we study the behaviour of the output buffer occupancy with various number of input links (N) and unlimited buffer capacity. For this purpose, we consider three ATM switches with two, five and ten input links respectively. The amount of traffic routed from each input port to the designated output port produces a loading of 80%. The simulations run for different packet size distributions and result for MM_2 is presented in Figure 7.1 together with the corresponding results from the $M^{[x]}/D/1$ queueing model.

Figure 7.1: Number of Links - the effects of varying the number of links (N = 2, 5, 10) in the simulation model while the parameters in the queueing model are unchanged, for the case with the MM_2 distribution and ρ = 0.80.

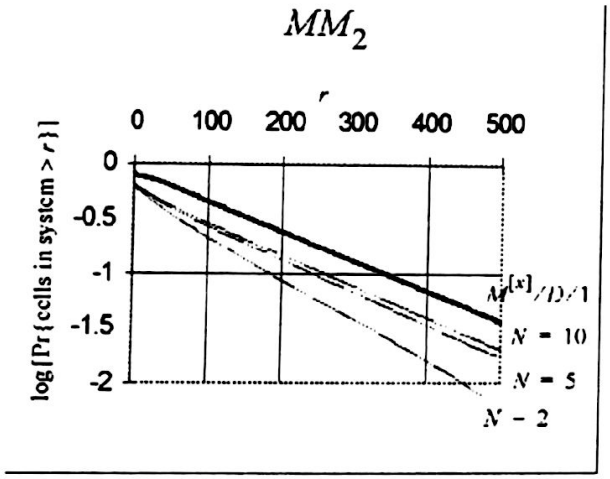

It can be observed from Figure 7.1 that the higher the number of input links, the more likely that the actual system behaves like a $M^{[x]}/D/1$ queueing model and the latter provides a conservative bound for the former. This applies to other packet size distributions as well and the results with ρ = 0.8 and N = 10 are shown in Figure 7.2. In this figure, it seems that the vertical distances between the tow tail distributions is a constant value which varies with different packet size distribution. In other words, the two tail distributions follow the same curvature. For this reason, we set the number of links to be ten in all the simulations used for the modelling of the PDS and the comparisons between the PDS and the Cell Discard Strategy in the subsequence sections.

Figure 7.2: 10 Input Links - for the case with 10 input links and $\rho = 0.80$, the comparison of the queueing model and the simulations with each of the packet size distributions (F_{08}, MM_1, F_{45} and MM_2).

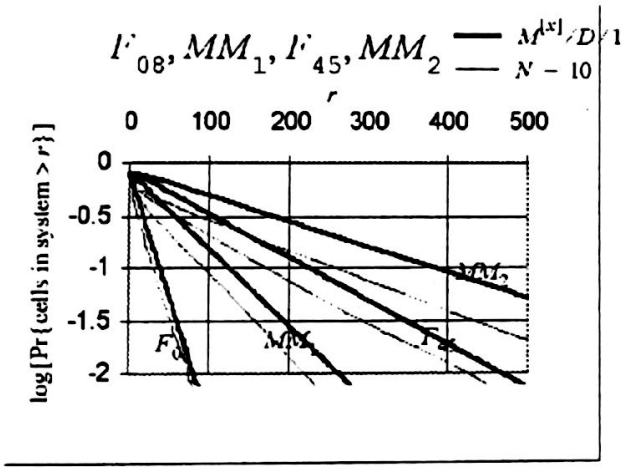

6.1.2 Traffic intensity

Secondly, we study the behaviour of the output buffer occupancy with different levels of traffic intensity at the output port and unlimited buffer capacity. For this purpose, we consider a ten port ATM switch. The amount of traffic routed from each input port to the designated output port produces a loading of 70%, 80% and 90%. The simulations are run for different packet size distributions and results MM_2 are presented in Figure 7.3 together with the corresponding results from the $M^{[x]}/D/1$ queueing model.

Figure 7.3: Traffic Intensity - the effects of varying the traffic intensity (ρ = 0.70, 0.80, 0.90), for the case with the MM_2 distribution and N = 10

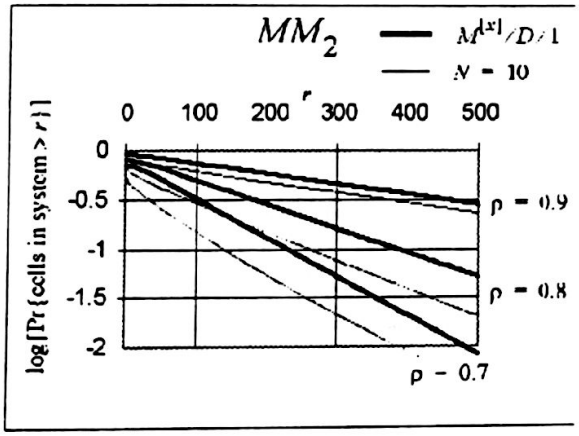

It can be observed from Figure 7.3 that the higher the traffic intensity at the output port, the more likely that the actual system behaves like a $M^{[x]}/D/1$ queueing model and the latter provides a conservative bound for the former. As the service process of both the queueing model and the simulation model are the same, the implication of this relationship is that the two arrival process are getting more similar to each other as the offered load increases.

6.1.3 Rate adaptation

Next, we study the behaviour of an ATM switch output buffer when the service rate is less than that of an individual incoming link. An example is a slow TE connected to a concentrator. MM_2 is used as the packet size distribution for this experiment. The simulations run for the amount of traffic routed from that input port to the designated output to produce loadings of 70% and 80%. In each case, the ratio of incoming link rate to the service rate is varied with ratios 3:1, 4:1 and 12:1. It should be noted that if the service rate is equal to that of an incoming link and traffic is arriving from a single link than buffers are needed only for synchronization. Results are presented in

Figure 7.4 together with the corresponding results from the $M^{[x]}/D/1$ queueing model.

Figure 7.4: Rate Adaptation - the effects of changing the traffic rate ratio (incoming link rate: service rate = 3:1, 4:1, 12:1), for the case with the MM_2 distribution, ρ = 0.80, 0.90 and N = 1.

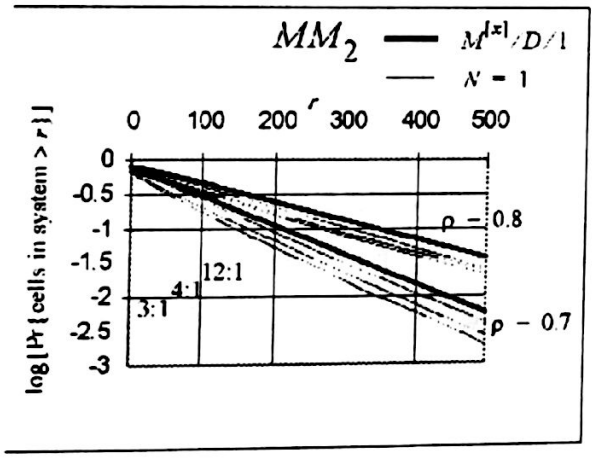

It can be observed from Figure 7.4 that the higher the ratio of incoming link rate to the service rate, the more likely that the actual system behaves like a $M^{[x]}/D/1$ queueing model and the latter provides a conservative bound for the former.

6.1.4 The application of the $M^{[x]}/D/1/\infty$ queueing model

The $M^{[x]}/D/1/\infty$ queueing model can be used to study the behaviour of an ATM data switch. In the comparisons of the two arrival patterns made in the previous sub-sections, the packets in both arrival patterns are generated according to the same process, and the traffic intensity is also the same. In other words, the number of packets directed to the tagged output buffer are the same on average. The difference between the two arrival patterns is how the cells of the packets arrive at the output buffer. We observed that when the level of convergence increases or when the traffic intensity increases, the probability of superposition of cells arrive at the same instant increases. This may explain how the factors affect the relationship between cell

arrivals and group arrivals, however, the detail analysis of this relationship is the subject of ongoing work. In this thesis, this relationship shows that our queueing model is a good approximation of the cell arrival pattern which occurs within an ATM data switch. Hence we make use of this relationship to support our queueing model to be used as a tool for studying the behaviour of an ATM data switch. To study this behaviour, we consider the queue statistics obtained from the $M^{[x]}/D/1/\infty$ queue. For a given packet size distribution, Figure 7.5 shows how the queue occupancy varies with the traffic loading in the range 50 to 90%. We observed that with a higher traffic intensity level, the probability of having a particular buffer occupancy level is higher. This applies to all the packet size distributions considered. On the other hand, Figure 7.6 shows how the queue occupancy varies with the packet size distribution. Note that with the same mean on packet sizes, a higher variability in the packet length distribution results in a higher queue occupancy distribution. This observation holds for the packet size distributions F_{08} and MM_1, and for F_{45} and MM_2.

Figure 7.5: Occupancy distributions - the effects of changing ρ (from 0.5 to 0.9) for the case with each of the four different packet size distributions.

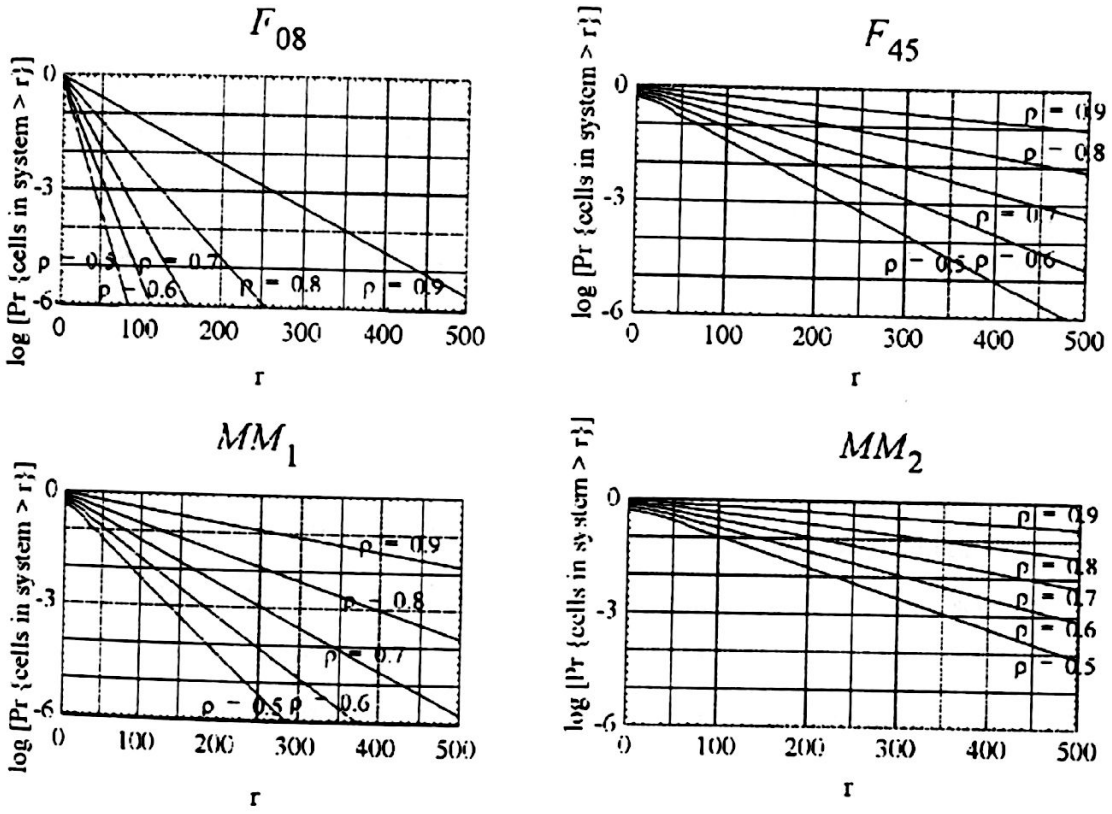

Figure 7.6: Occupancy distributions - the effects of using different packet size distributions, for the case with ρ = 0.80

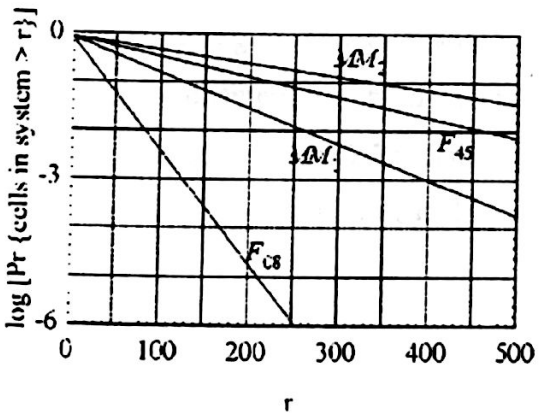

6.2 Modelling of the Packet discard strategy

In this section, we study the performance of an ATM switch operating with the Packet Discard Strategy. Firstly, we compare the packet loss probabilities obtained from the $M^{[x]}/D/1$ queueing model with threshold and those obtained from the simulations. This comparison will show that our queueing model is a good approximation of the cell arrival pattern even when the PDS is operational. Secondly, we study the goodput performance of the PDS using both the queueing model and simulations. Finally we use the $M^{[x]}/D/1$ queueing model with threshold to study the performance of the PDS by observing the packet loss probabilities and goodput over a range of threshold values, traffic intensities and offered load. Recall from chapter two that with the PDS, we have the overflow and underflow problem. In this thesis we concentrate on investigating the overflow problem. Hence in this section we set the proportion of the buffer above threshold to be large enough to accommodate without loss all the packets in transit when the threshold is exceeded. This applies to both the queueing model and the simulations.

To study the change in the packet loss probabilities, we varies the specified threshold value from 0 to 500 cells. The amount of traffic routed from each input port to the designated output port produces a loading of 70%, 80% and 90%. For a given packet size distribution, simulations run for different threshold values and the results have been interpolated. Results are presented in Figure 7.7 together with the corresponding results from the $M^{[x]}/D/1$ queueing model with threshold, respectively for all the packet size distributions considered.

Figure 7.7: PDS Performance: Loss Probabilities vs Threshold with $\rho = 0.70, 0.80, 0.90$ and $N = 10$

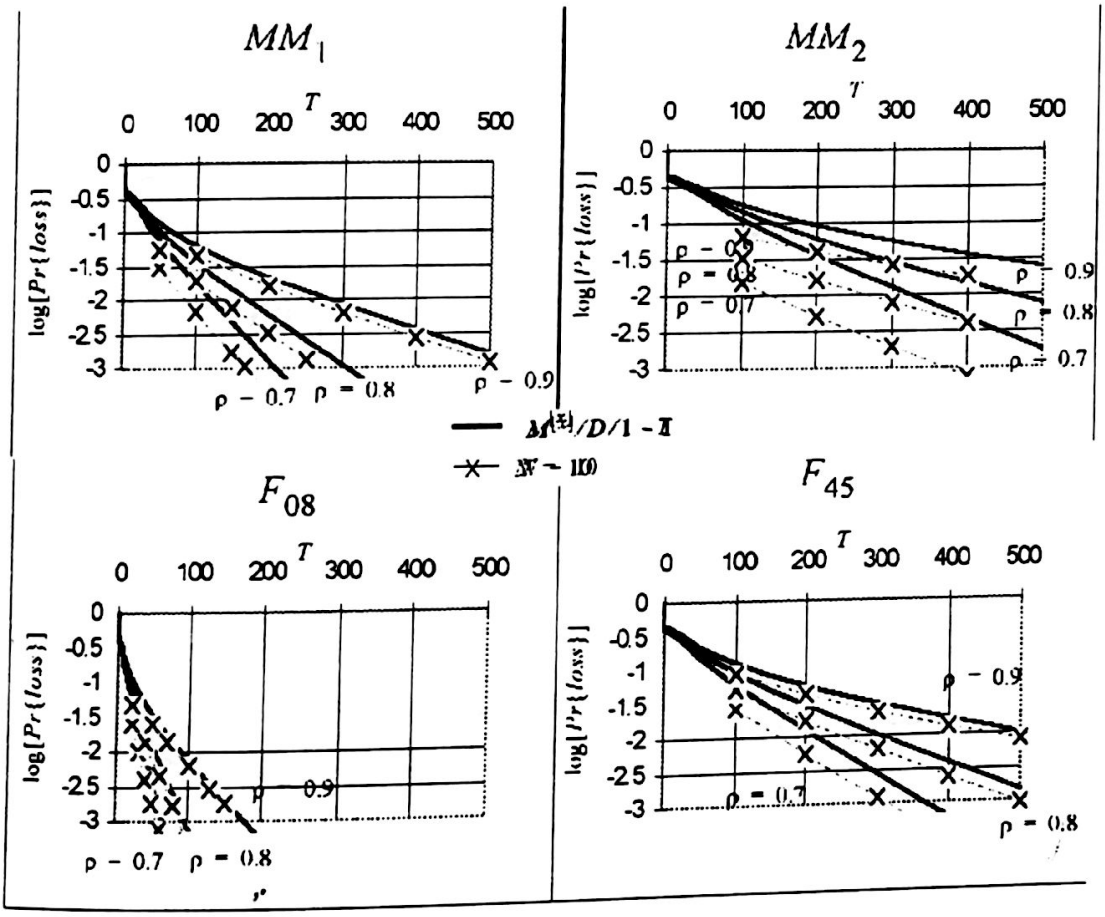

It can be observed from Figure 7.7 that the $M^{[x]}/D/1$ queueing model with threshold closely matches the results of the actual system and are conservative. This implies that even with the PDS in operation, there exists a similar relationship between group arrivals and the cell arrivals. Here, the similarity appears in the packet loss probability while in the case without the PDS, the similarity appears in the buffer occupancy distribution. We also observed from the figures that once again the vertical distance between the two curves seems to be a constant value which varies with different packet size distributions, however, the detail analysis of this relationship is the subject of ongoing work. These observations support our $M^{[x]}/D/1$ queueing model with threshold to be used as the guideline for design.

We now study the change in goodput characteristics as the traffic intensity is varied. For a given packet size distribution, simulations run for different traffic intensities with different threshold values, and the results have been interpolated. Results are presented in Figure 7.8 together with the corresponding results from the $M^{[x]}/D/1$ queueing model with threshold.

Figure 7.8: PDS Performance: Goodput vs Offered Load for N = 10

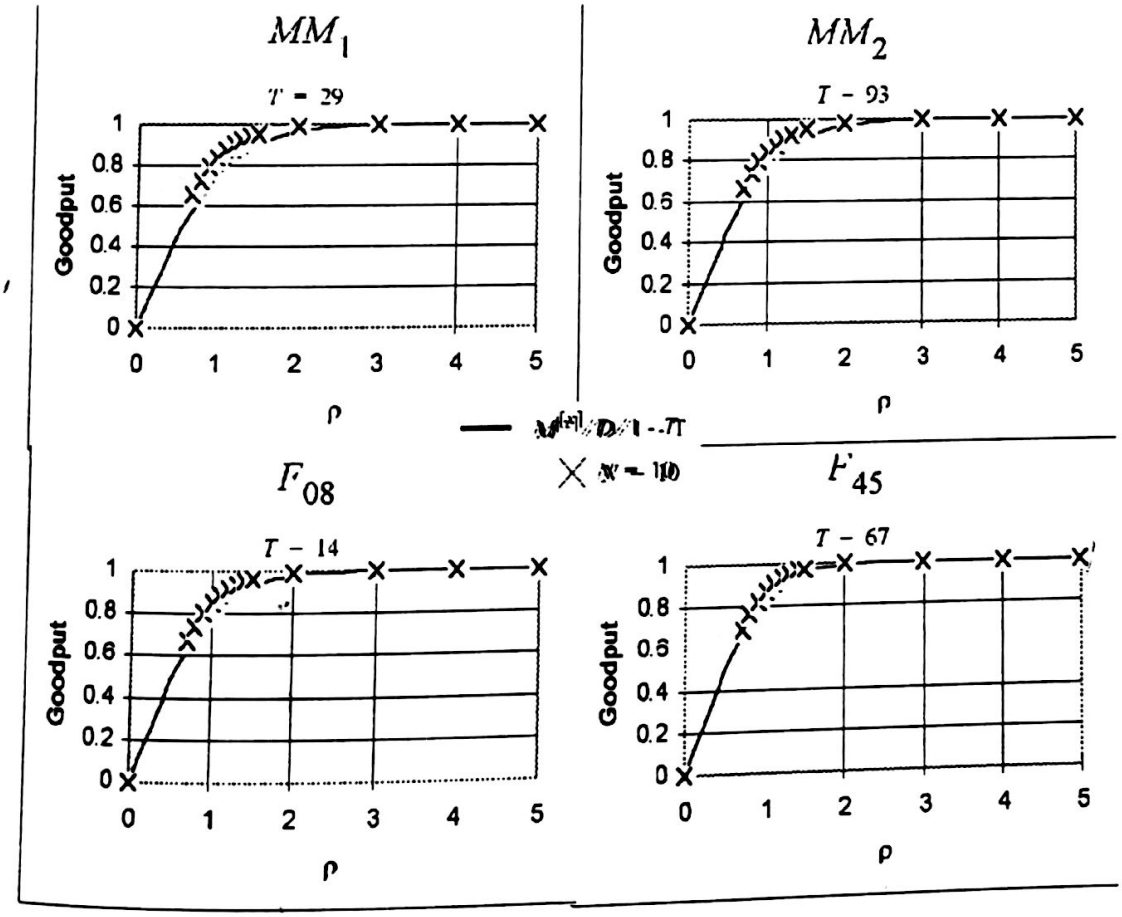

It can be observed from Figure 7.8 that once again the $M^{[x]}/D/1$ queueing model with threshold closely matches the results of the actual system and are conservative. For all the packet size distributions we considered, both the analytical and simulation results show similar patterns. The performance of the system with the PDS is very close to the ideal goodput of a system. The system has sustain goodput even in the

overload region. Having such a good performance, the Packet Discard Strategy is yet simple and can be easily implemented at the ATM layer. Choosing PDS as the basic congestion control strategy for the UBR traffic should largely enhance the performance of the ATM data network.

6.2.1 The queue statistics of the $M^{[x]}/D/1$ queueing model with threshold

To facilitate the use of the $M^{[x]}/D/1$ queueing model with threshold as the guideline for the design of the PDS, we investigate the queue statistics of the $M^{[x]}/D/1$ queueing model with threshold. For a given packet size distribution, Figure 7.9, Figure 7.10, Figure 7.11 and Figure 7.12 show the queue statistics of the $M^{[x]}/D/1$ queueing model with threshold. In each of these figures, Graph (a) shows how the probability of packet loss varies with the threshold value, for traffic loading ranging from 50 to 90%. These graphs can be used to estimate the threshold value as to be described in the following sub-section. Each Graph (b) shows how the goodput varies with traffic loading, for different threshold values. The implications of these graphs will be discussed in the next section. Each Graph (c) shows how the utilization factor varies with the threshold value, for traffic loading ranging from 50 to 90%. With these graphs, we observe how fast does an output port achieve maximum possible goodput as the threshold value increases. When the system is underloaded, the maximum possible goodput is equal to the offered load, and when the system is overloaded, the maximum possible goodput is equal to unity. Each Graph (d) shows how the probability of packet loss varies with traffic loading, for threshold values ranging from 100 to 500. With these graphs, it can been seen that once the system is overloaded, a percentage of packets must be discarded regardless of the threshold values. However, the packet loss probability can be minimized with a larger threshold value.

Figure 7.9: Queue statistics for F_{08}

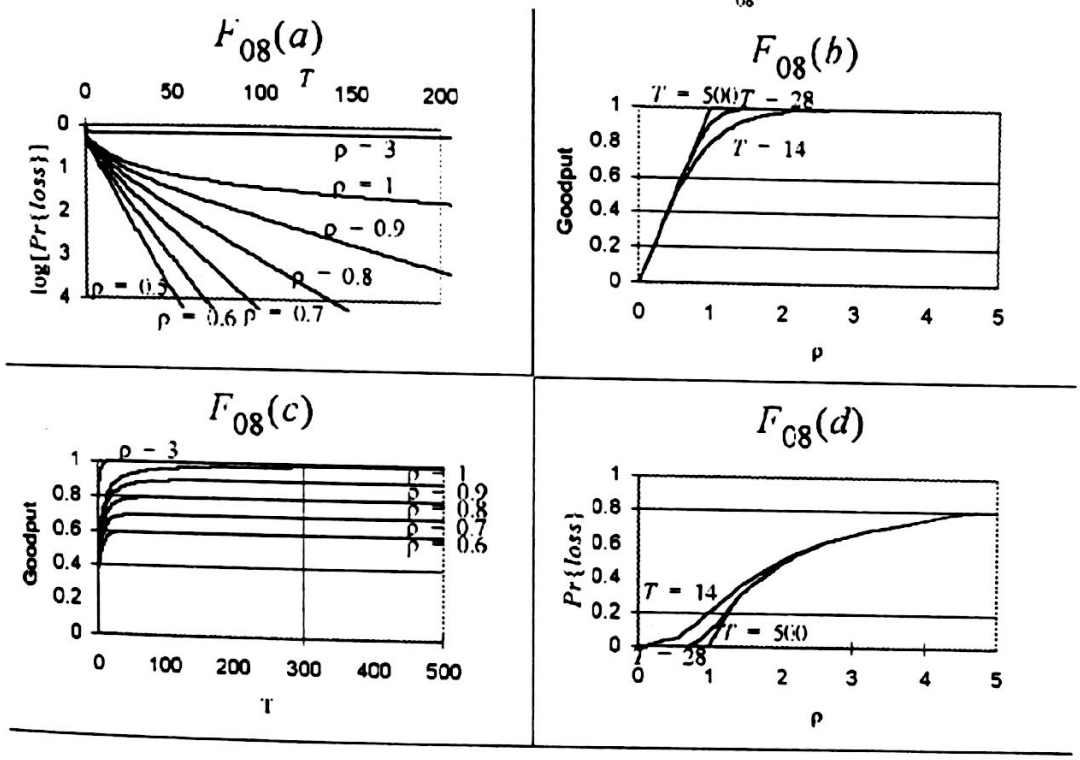

Figure 7.10: Queue statistics for MM_1

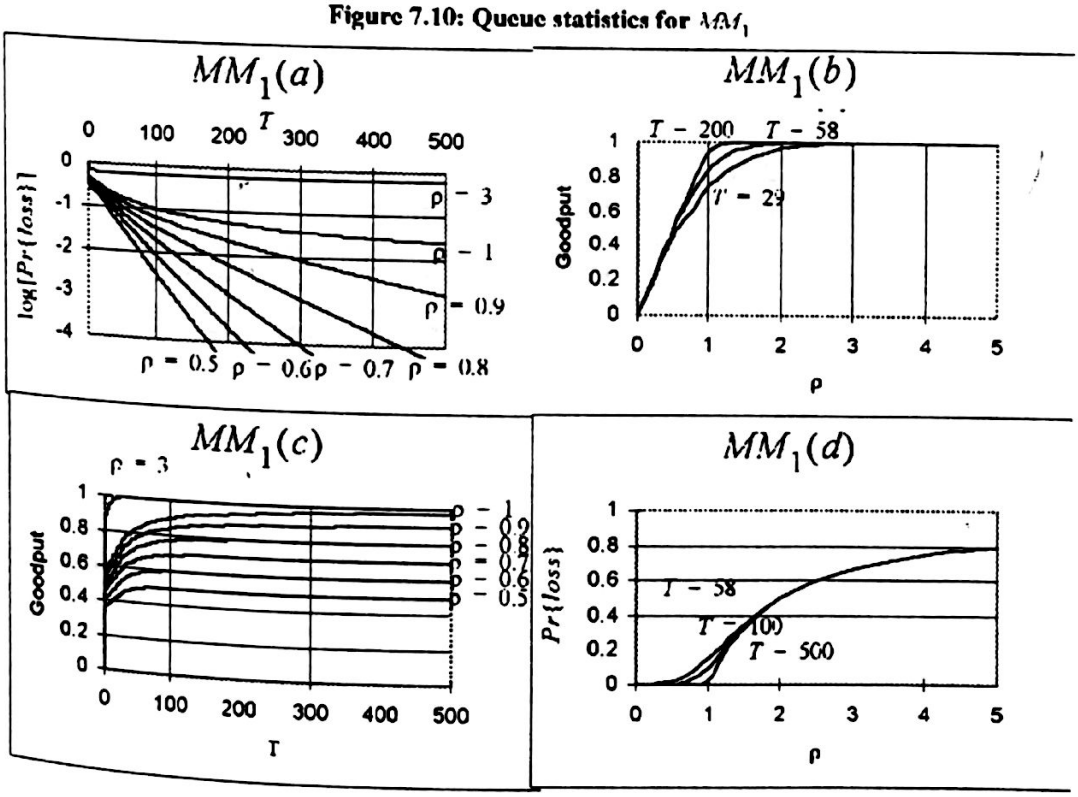

Figure 7.11: Queue statistics for F_{45}

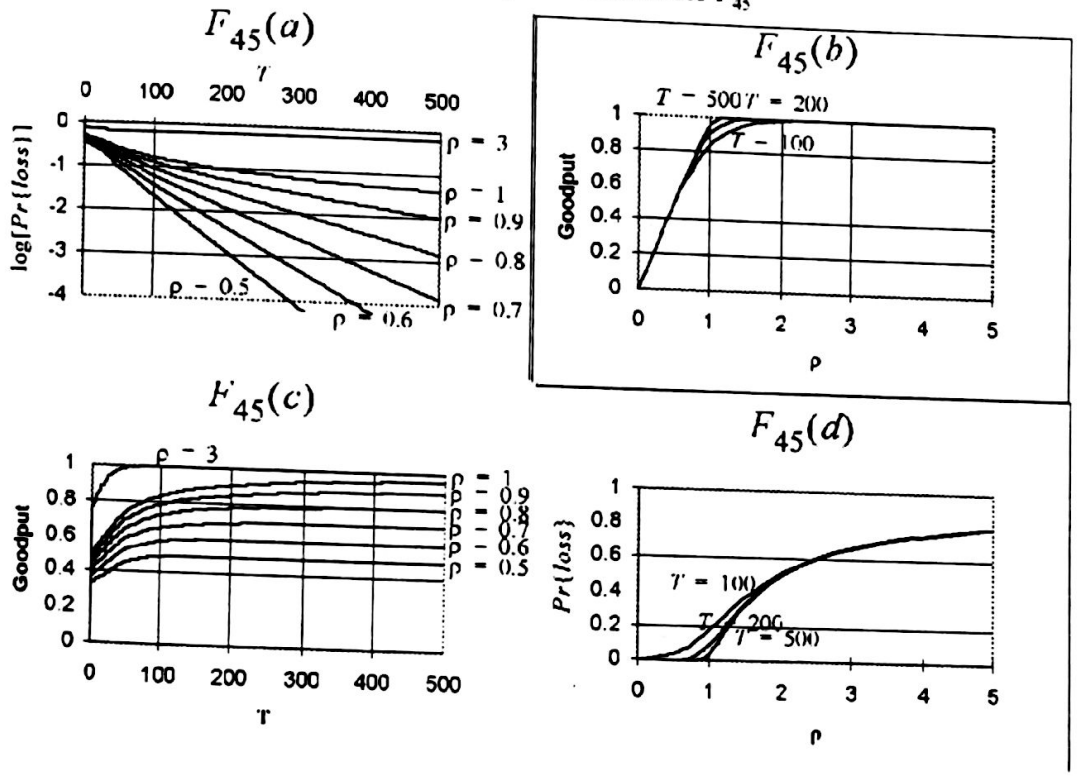

Figure 7.12: Queue statistics for MM_2 -

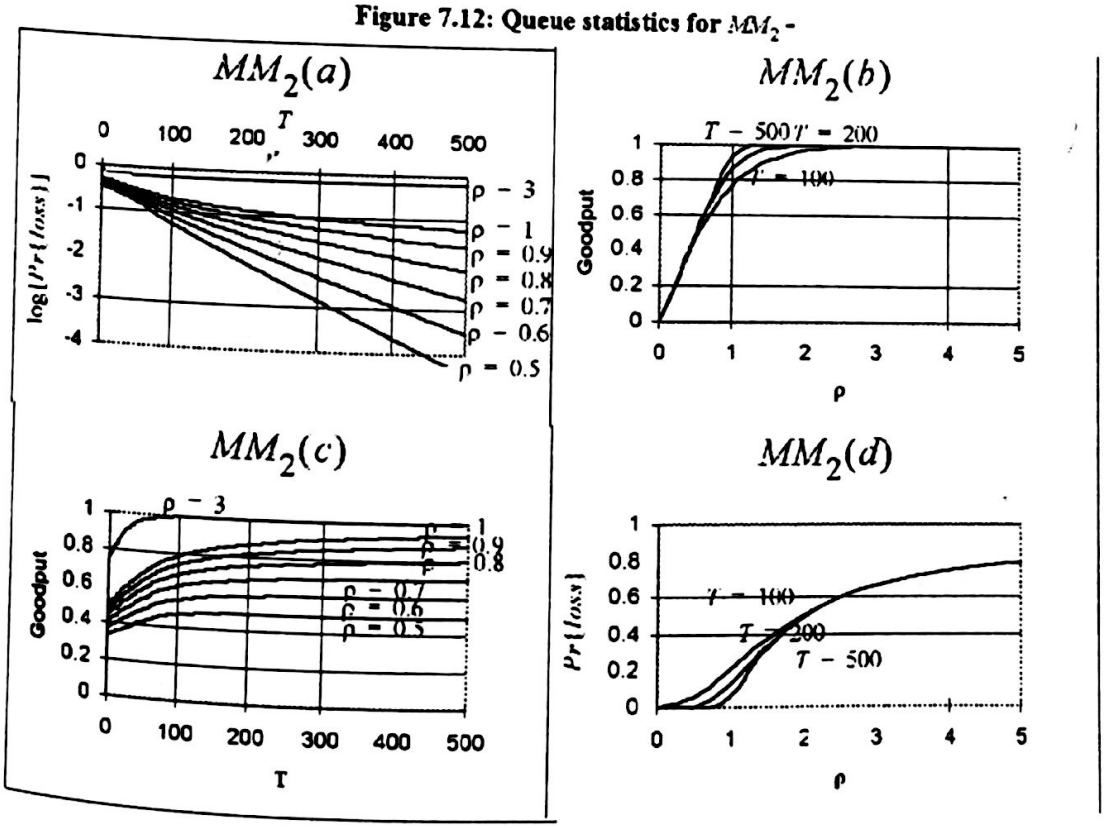

With the analysis of the $M^{[x]}/D/1/\infty$ queueing model in the previous section, we observed that with the same packet mean length, higher variability in the packet length distribution results in higher packet loss probability. This observation indeed also holds in the case with threshold as shown in Figure 7.13.

Figure 7.13: Probability of packet loss vs Threshold with ρ = 0.80

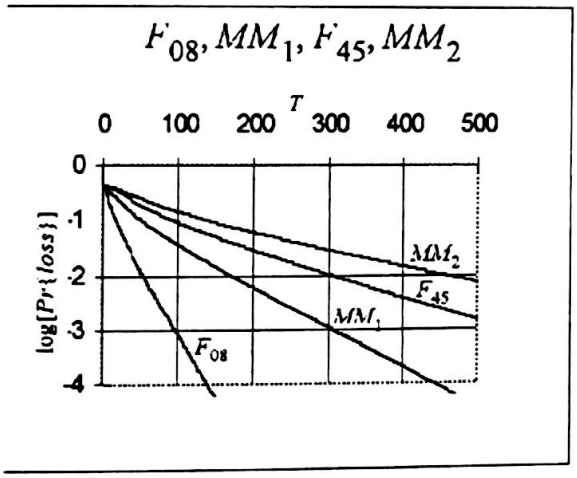

6.2.2 The selection of the threshold

With the similarities between the results from the queueing models and simulations, the results from the queueing model provide a satisfactory basis for design. With a particular traffic intensity, we select the threshold value required to achieve a tolerable packet loss probability for Markovian traffic pattern. Each Graph (a) in Figure 7.9 to Figure 7.12 show the packet loss probabilities against threshold values, for a range of traffic intensity and for different packet size distributions. These graphs provide a guideline for the selection of threshold for the PDS. For example, with MAN like packet size distribution, 80% traffic intensity and 10^{-2} packet loss probability, our estimated threshold value is 500 cells. With LAN like packet size distribution and the same traffic intensity, having such a threshold can achieve a much better packet loss probability of 3.1×10^{-5}. Note that this loss probability is calculated based on the discard of packets when the specified threshold is exceeded, the problem of buffer overflows is discussed in the next section.

6.3 The Packet Discard Strategy and the Cell Discard Strategy

Based on the $M^{[x]}/D/1$ queueing system, in this section we present the performance of an ATM switch operating with the PDS and that with the Cell Discard Strategy. Graph (b)s of Figure 7.9 to Figure 7.12 in the last section show the goodput against offered load for different threshold levels. It is observed that, for the Markovian traffic pattern, quite moderate threshold levels are all that are required to produce a performance near to the ideal even for the large packet distributions. Similar performance is observed for the goodput of the system with different threshold values. As the threshold increases, the goodput characteristic calculated from the $M^{[x]}/D/1$ queueing model with threshold gets closer to the ideal goodput. On the other hand with the Cell Discard Strategy, any increase in the buffer size can only serve to improve the goodput where the offered load is around unity. This is illustrated in Figure 7.14.

Figure 7.14: Performance without PDS: Goodput of $M^{[x]}/D/1/B$.

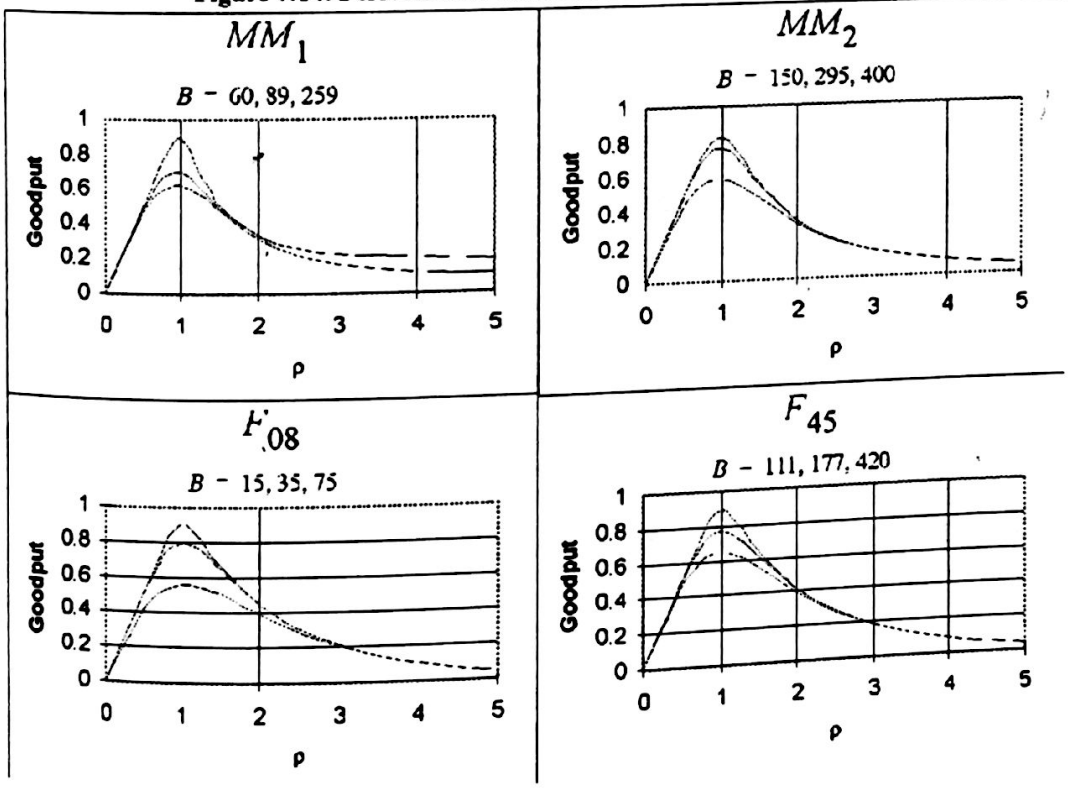

To understand the goodput characteristics associated with the Cell Discard Strategy, we analyse the consequence of increasing the buffer size and the offered load. With Cell Discard Strategy, a larger buffer accommodates more cells and will hence produce a better goodput performance. When the system is underloaded, increasing the offered load increases the number of incoming cells and hence increases the goodput performance. However, even when the system is underloaded, there is still a possibility of buffer overflows because the cells come in a bursty nature rather than a continuous one, hence cell loss multiplication occurs and goodput performance degrades. This possibility of buffer overflows increases as the offered load increases, and once the system is overloaded, part of the incoming cells must be discarded as pointed out in the Graph(d)s in the last section, hence buffer overflows and cell loss multiplication must occur, and goodput performance must degrade. As the offered load increases above unity, the number of packets that must be discarded increases and hence the goodput performance degrades further. These consequences may explain why the maximum goodput performance is around unity. Note that the goodput performance does not drop to zero even under sustained overload, however, most of the packets that get through without being corrupted should be small packets. This conjecture is supported by the goodput characteristic associated with the MM_1 packet size distribution. We observe in the overloaded region, the goodput performance associated with the MM_1 is the best among the others. Recall from the beginning of this chapter that with the MM_1 packet size distribution, 73% of the packets generated contain only two cells. The probability of having two related cells going through the switch without being discarded is higher than other number of related cells. Hence the goodput characteristic associated with the MM_1 packet size distribution is the best among the others in the overload region and our conjecture holds.

We now compare the analytical results on the performance of the PDS and that of the Cell Discard Strategy. Note that the performance of the PDS also depends on the

buffer size above threshold. If it is not large enough to accommodate all the packets in transit, overflows and hence cell loss multiplication will occur as pointed out in Chapter 2. Therefore in this comparison, we set the threshold T for the PDS to be 100 cells and the buffer size B for both discard strategy to be 400 cells. That is, the buffer size above threshold is three times the one below, which we refer as having a buffer size ratio of 1:3. For each of the packet size distributions, our simulations show that with such a buffer size ratio for the PDS, overflow does not occur within the respective run time periods as specified in the beginning of this chapter. We believe that with such a buffer size ratio for the PDS, the probability of cell loss due to buffer overflows could be neglected, but it is the subject of ongoing work.

Results for the comparison are presented in Figure 7.15 which shows that the goodput performance obtained with the PDS is seen to be very nearly ideal providing 100% goodput in the overload region. By comparison the goodput performance of the system with the Cell Discard Strategy degrades quite rapidly with increasing overload or congestion. Even though the goodput of the $M^{[x]}/D/1/B$ queue doesn't drop to zero immediately once it is overloaded, it should be noted that most of the packets that go though without being corrupted are small packets. On the other hand, the systems with PDS doesn't have bias towards small packets at all. It maintains on the output the same packet size distribution as the packets come. These results are reasonable as with the PDS, cells belonging to corrupted packets are prevented from entering the buffer and going through the output port. However, with the Cell Discard Strategy, cells from corrupted packets are admitted into the buffer and bandwidth of the output port is hence wasted.

From Figure 7.15 we observe that with the same threshold value and buffer size, the performance with packet size distributions F_{08} and MM_1 is better than that with F_{45}

and MM_2, where F_{08} and MM_1 contain relatively smaller packets. This implies larger threshold values are required for packet size distributions with larger packets.

Figure 7.15: Performance with and without PDS: A Comparison of Goodput

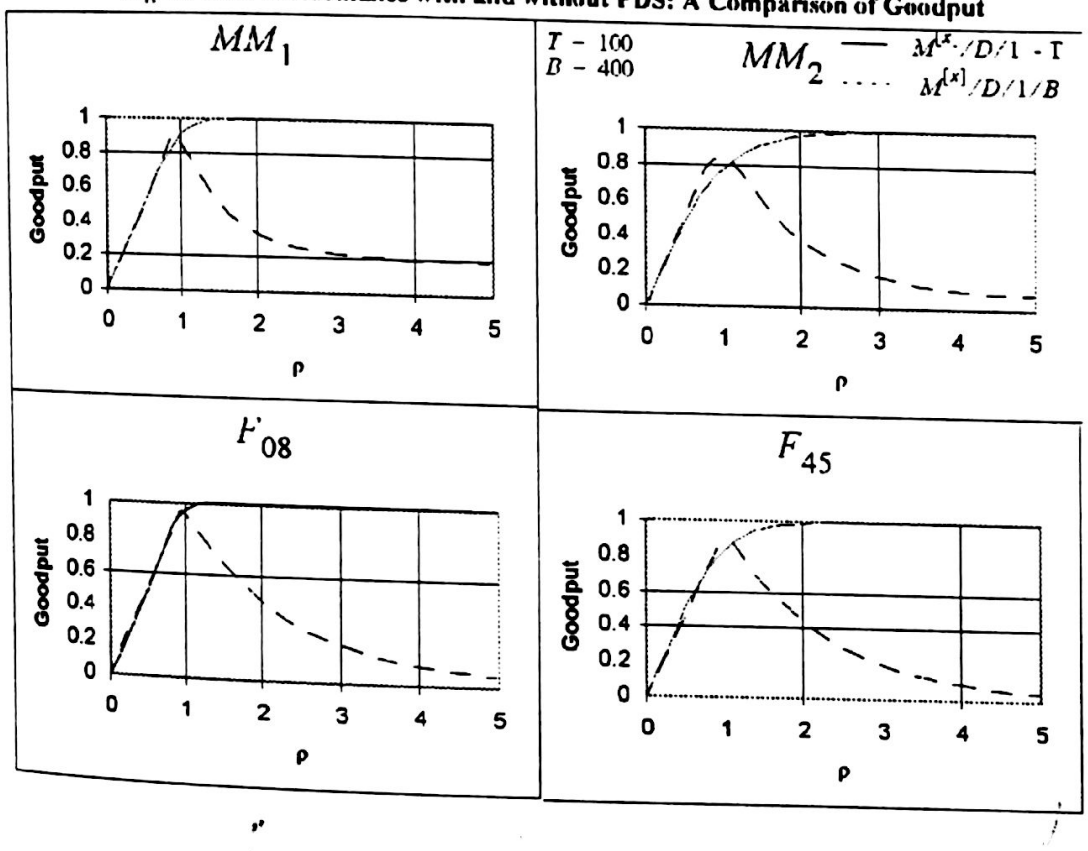

Next, we demonstrate the effect of buffer overflows with the PDS. We observed in this chapter that for the PDS, the buffer occupancy distributions below threshold obtained from the analytical model are related to those from the simulations, however, the distributions above threshold are not related. In our queueing analysis, we derived in (84) that in order to eliminate buffer overflows, the buffer size required above threshold should be $G_{max} - 1$ where G_{max} is the maximum packet size of a packet size distribution. However with cell arrivals, this is not the case. To illustrate this point and the effects of overflows with PDS, we set the buffer size to be $T + G_{max} - 1$ for each of the packet size distributions with their respective G_{max}.

The comparison of the PDS and Cell Discard Strategy for the MM_2 packet size distribution is shown in Figure 7.16.

Figure 7.16: MAN traffic - the comparison of PDS and CDS with T=216 and B=295 and MM_2

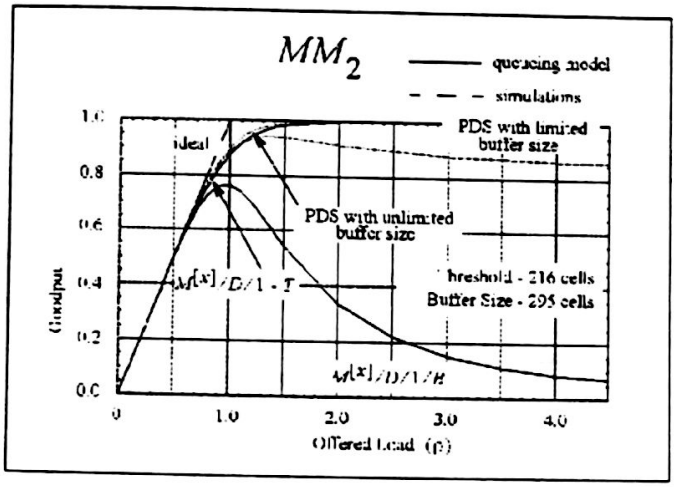

With MM_2, the maximum packet size is 80 and we set the threshold to be 216. This made the buffer size above threshold to be 0.37 times the one below. It can be observed that with such a low buffer size ratio, simulation shows that the performance with the PDS does degrade due to buffer overflows above threshold, however, such performance is still far better than the one with Cell Discard Strategy, especially in the overload region.

Figure 7.17: LAN traffic - the comparison of PDS and CDS with T=58 and B=89 and MM_1

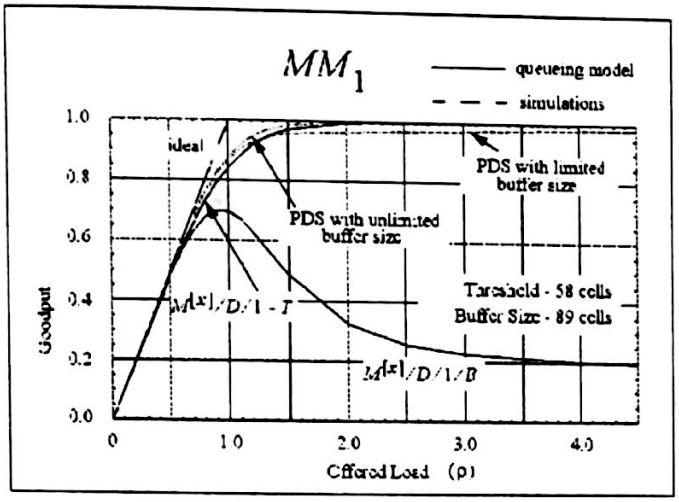

Figure 7.17 shows the same comparison for the MM_1 packet size distribution. Here, the buffer size above threshold is 0.56 times the one below. The performance with the PDS does degrade but much less serious than the one with the MM_2 packet size distribution. This suggests that apart from the buffer size ratio, the packet size distribution may also be a key to the dimensioning of the buffer size with Packet Discard Strategy, but is the subject of ongoing work. The comparisons immediately above are also reported in our paper [22] together with another packet discard strategy based on buffer allocation.