

MULTIPLE CLASS SELECTIVE DISCARD MECHANISM FOR THE PROVISION OF DIVERSE LOSS REQUIREMENTS UNDER A LONG-RANGE DEPENDENT PROCESS*

Nelson L. S. Fonseca¹

and

Marcelo de J. Ferreira²

¹Universidade Estadual de Campinas
Instituto de Computação
C.P. 6176 - Campinas - SP - 13083-970
nfonseca@dcc.unicamp.br

²Instituto Luterano de Ensino Superior de Manaus
Departamento de Informática
Av. Solimões, 2 - Manaus - AM - 69077-077
mmjf@internext.com.br

Resumo - A taxa de perda em um multiplexador sob a influência de um processo com dependências de longa duração decai muito lentamente com o aumento do tamanho do buffer e consequentemente prover Qualidade de Serviço diferenciada é uma tarefa desafiante. Neste artigo, investigam-se as vantagens de se adotar mecanismo de descarte seletivo com múltiplas classes quando comparado com mecanismo com dois níveis de prioridade.

Abstract - Providing diverse QoS requirements in multimedia networks is a challenging task. Under a long-range dependent process, the loss rate does not decrease considerably as we increase the buffer size. In this paper, we investigate the advantages of adopting a multi-priority selective discard mechanisms over traditional two-priority mechanisms under a long range dependent process.

Keywords: Self similar network traffic, quality of service, statistical multiplexer, selective discard mechanism.

1. INTRODUCTION

Several studies [1]-[2] have claimed that different types of network traffic, e.g. local area network traffic (LAN), can be accurately modeled by a self-similar process. A self-similar process is able to capture the long-range dependence (LRD) phenomenon exhibited by this traffic. Moreover, series of simulation and analytical studies [3]-[6] demonstrated that this phenomenon might have a pervasive effect on queueing performance, *i.e.*, there is clear evidence that it can potentially cause massive cell losses in ATM networks. In fact, Norros [5] and Duffield [6] showed that the buffer overflow probability for an ATM queueing system with fractional Brownian arrivals follows a Weibull distribution. Furthermore, this queueing system suffers from the buffer inefficiency phenomenon [4], [7], *i.e.*, by just increasing the buffer size, one is not able to decrease the buffer overflow probability considerably.

Different multimedia applications have diverse loss requirements. Coping with different loss requirements is a challenging task. Selective discard is a congestion control mechanism aimed at enabling the network to deal with diverse loss requirements. In a selective discard mechanism

cells are discarded in an overflow situation according to their priority level [8].

Selective discard has been studied in the past few years. However, only recently, has it been investigated under long-range processes. In [9], we analyzed selective discard with two priority levels under a long-range dependent process. We found out that the complete sharing with push-out buffer policy is clearly worth adopting, while complete sharing with guaranteed queue minimum buffer policy may be not. It was also evident that the choice of push-out policies has a significant impact on loss performances.

A multiple class selective discard mechanism furnishes more than two priority levels. Besides providing distinct loss rate, a multiple class mechanism demands less buffer space to support diverse QoS than does a two-priority mechanism. In other words, with a higher number of priority classes we do not need to guarantee loss rate bounds lower than the required. Additionally, we may carry higher loads than does a two-priority mechanism.

At the time of the adoption of selective discard in the ATM standard research on long-range dependent traffic was still in its infancy. Multiple-class selective discard mechanism under a short-range dependent process seems to be unattractive since the buffer gain is not significant when compared to a two-priority scheme [10]. However, under a long-range dependent process a two-priority mechanism demands huge buffer sizes to provide the same loss requirements supported by a multi-priority system [9]. In this paper, we show the advantage of adopting a multi-priority level mechanism over a traditional two-level mechanism under a long-range dependent process [11]. In addition, we highlight the difference between selective discard under a short-range dependent and selective discard under a long-range dependent process.

This paper is organized as follows: Section 2 introduces the Fractal Brownian Motion. Section 3 describes the buffer inefficiency phenomenon. Sections 4 and 5 describe selective discard mechanism and selective discard under a long range dependent process, respectively. Section 6 discusses the effectiveness of a multiple priority scheme. Finally, conclusions are drawn in section 7.

*This work was partially sponsored by CNPq, CNPq/PROTEM-III Almaden, CAPES and CENAPAD/SP.

2. THE FRACTIONAL BROWNIAN MOTION

The ordinary Brownian motion, $B(t)$, describes the movement of a particle in a liquid subjected to collisions and other forces [12]. It is a real random function with independent Gaussian increments such that

$$E[B(t+s) - B(t)] = 0$$

$$Var[B(t+s) - B(t)] = \sigma^2|s|$$

Mandelbrot defines fractional Brownian motion (fBm) as being the moving average of $dB(t)$ in which past increments of $B(t)$ are weighted by the kernel $(t-s)^{H-1/2}$ [13].

Definition: Let H be such that $0 < H < 1$. The fBm is defined as the Weyl's fractional integral of $B(t)$.

$$B_H(t) = \frac{1}{\Gamma(H+1/2)} \int_{-\infty}^0 ((t-s)^{H-1/2} - (-s)^{H-1/2}) dB(s) + \int_0^t (t-s)^{H-1/2} dB(s).$$

This equation leads to the ordinary Brownian motion if $H = 1/2$. Its self-similar property is based on the fact that $B_H(rs)$ is identical in distribution to $\rho^H * B_H(s)$. The increments of the fBm, Y_j form a stationary sequence called fractional Gaussian noise (fGn):

$$Y_j = B_H(j+1) - B_H(j), j = \dots, -1, 0, 1, \dots$$

We should note that these increments are not independent unless you have pure Brownian motion, i.e., $H = 1/2$. Moreover, Hurst law states that $Var[B_H(t+s) - B_H(t)] = \sigma s^{2H}$, i.e., a fBm arrival model is also able to capture the inherent high variability exhibited by real network traffic.

3. THE BUFFER INEFFICACY PHENOMENON

The buffer inefficacy phenomenon is the queueing phenomenon in which by just increasing the buffer size, we are not able to decrease the buffer overflow probability considerably. This phenomenon has been reported earlier by several other studies [16]-[18]. In this section, we present a very intuitive explanation for it and show that it is of particular importance when the traffic source exhibits long-range dependencies.

We model an ATM node as a deterministic queueing system with constant departure rate given by c and finite buffer size given by b . The input traffic is given by the stochastic process $A(t)$ with mean input rate $\bar{a} < c$. It defines the aggregate number of cell arrivals up to time t , $t \geq 0$. Assume that the buffer overflow occurs at time t so that we can write $A(t) = ct + b$. Moreover, $A(t)/t \geq c + b/t$.

By the law of large numbers, the average arrival rate $A(t)/t$ converges to its mean \bar{a} . Therefore, the probability that it exceeds the term $(c + b/t)$ decreases with t : $P(A(t)/t \geq c + b/t) = \Psi(t)$

In other words, $\Psi(t)$ is a decreasing function with time. The buffer inefficacy phenomenon occurs, if the buffer overflow probability given by decays slowly with t , i.e. $\Psi(t)$ is nonnegligible for large t . In this case, since t is large, the term (b/t) is negligible. Therefore, even if we increase the buffer size, we are not able to increase the term (b/t) significantly in order to decrease the cell loss probability. Intuitively, this phenomenon occurs if the arrival process is able to transmit at high rates for very long periods of time, i.e. if it converges slowly to its mean. We show that a LRD source can transmit at high rates for very long periods of time. Following Norros' work [5], assume that the arrival process $A_H(t)$ is a fractional Brownian motion (fBm) process given by $A_H(t) = \bar{a}t + \sigma Z(t)$ where $\bar{a} > 0$ is the mean input rate, $\sigma > 0$ is the standard deviation, $H \in [1/2, 1)$ is the self-similar (Hurst) parameter and $Z(t)$ is a normalized fractional Brownian motion. When $H = 1/2$, we have the special case of the ordinary Brownian motion. The probability that over a time interval of length t the source can overcome the potential service ct and further exceed a buffer level b is given by:

$$P(A_H(t) \geq ct + b) = P(\bar{a}t + \sigma Z(t) > ct + b) = P(Z(t) > \frac{t(c - \bar{a}) + b}{\sigma})$$

By the self-similarity property $Z(t) = t^H Z(1)$ we have:

$$P(Z(1) > \frac{t(c - \bar{a}) + b}{\sigma t^H}) = \Phi(\frac{t(c - \bar{a}) + b}{\sigma t^H})$$

where $\Phi(y) = P(Z(1) > y)$ is the residual distribution function of the standard Gaussian distribution. In fact, using the approximation:

$$\Phi(y) \approx (2\pi)^{-1/2} (1+y)^{-1} \exp(-y^2/2) \approx \exp(-y^2/2)$$

we obtain:

$$P(A_H(t) > ct + b) = \Phi(\frac{t(c - \bar{a}) + b}{\sigma t^H}) \approx \exp(-\frac{1}{2}g(t)^2) = \exp(-\frac{1}{2}(\frac{t(c - \bar{a}) + b}{\sigma t^H})^2) \quad (1)$$

We compute Equation 1 for two sources with same mean, standard deviation and Hurst parameter $H = 0.50$ and $H = 0.85$ respectively. We choose the link bandwidth so that the link utilization is 50%. Figure 1. shows the results. We can see that the probability of buffer overflow for the LRD source decays very slowly with time. Therefore, increasing the buffer size is not enough to accommodate the strong low-frequency component of this source in order to avoid cell losses. On the other hand, Ψ decays very fast in the case of uncorrelated arrivals (Brownian motion) [17][18].

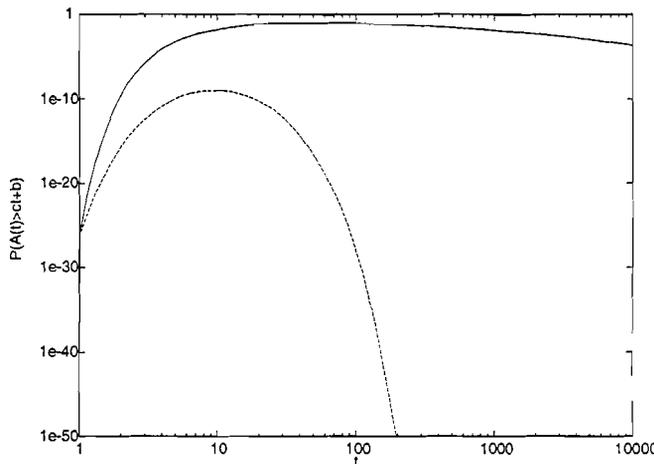


Figure 1. The Buffer Inefficacy Phenomenon.

4. SELECTIVE DISCARD MECHANISM

In a selective discard mechanism cells are discarded in overflow situations according to their priority level. A selective discard mechanism is completely specified by a buffer policy and by a push-out policy. While a buffer policy defines which buffer slot can be occupied by which cell, a push-out policy chooses a cell to be discarded among the cells with lowest priority.

Although the loss rate (the ratio between the number of lost cells and the total number of transmitted cells) is a meaningful and measurable parameter, it is an average value which does not entirely describe the loss process entirely. The number of cells consecutively lost (loss gap) gives a more detailed description of the loss process. For a certain value of loss rate, cells may be lost in several different ways. For instance, for a loss rate of 0.1, we may lose one cell out of every ten cells or we may lose one tenth of the total number of cells in a row. Depending on the signal recovery procedure at the receiver side, the length of the loss gap may have different impact on user's perceived QoS.

Our investigation considers the Complete Sharing (CS) with Push out buffer policy. In a CS policy if a high priority cell finds the buffer full it drops an enqueued low priority cell. CS is a loss conserving discipline. In a loss conserving discipline, (fixed size) cells are lost only in overflow situations. In other words, a loss-conserving discipline always admits a cell into the buffer if there is available space. Loss-conserving queues are of special interest because they minimize the overall cell loss, and consequently maximizes the throughput.

A push-out policy selects a cell to be dropped among cells with lowest priority. The most common policies are [14]: Last-In-First-Drop (LIFD), First-In-First-Drop (FIFD), Random selection (RAND) and Modified-FIFD (M-FIFD) [15]. The modified-FIFD policy always drops the oldest low priority cells to make room for an arriving cell irrespective of its class. Discarding low priority cells at different positions define different queue distributions. If cells close to the head of the queue are discarded, we increase the residence time of low priority cells. Thus, we augment the likelihood

of a high priority cell to find a low priority cell in queue to be dropped. Accordingly, the First-In-First-Drop policy gives lower high priority loss rate than does the Last-In-First-Drop policy. Conversely, the Last-In-First Drop increases the chances that a low priority cell close to the head of the queue be transmitted. It consequently, may break a long loss gap into two smaller gaps. The Modified-First-In-First-Drop maximizes the likelihood of a high priority cell to find a low priority cell by always dropping an enqueued low priority cell irrespective of the priority level of the arriving cell.

An in-depth view of selective discard mechanism may be found in [8].

5. SELECTIVE DISCARD UNDER A LONG-RANGE DEPENDENT PROCESS

In [9], we investigated the extent to which selective discard under a long-range dependent process can provide differentiated QoS. We analyzed complete sharing with push out and complete sharing with guaranteed queue minimum (CSGQM) buffer policies. Similarly to CS, in a CSGQM, low priority cells are discarded during buffer overflows, however, we guaranteed a minimum queue length to the low priority cells.

We verified that under CS and for Hurst parameter lower than 0.75, the high priority loss rate decreases considerably as the buffer size increases. Although for higher values of the Hurst parameter, the high priority loss rate decreases very slowly, CS produces differentiated per class loss rates. Conversely, for guaranteed queue length higher than 20% of the total buffer size, the per class priority loss rate given by CSGQM becomes indistinguishable.

The choice of push-out policy significantly impacts loss performances. M-FIFD may produce high priority loss rates three orders of magnitude lower than LIFD. On the other hand, LIFD may give maximum loss gap size which are typically half of the maximum loss gap size given by M-FIFD. We also notice that the per class loss rate and loss gap size are highly sensitive to the Hurst parameter, as well as to the variance.

6. THE EFFECTIVENESS OF MULTIPLE CLASS SELECTIVE DISCARD MECHANISM

To assess the effectiveness of multi-class selective discard mechanism we simulate a Complete Sharing queue with constant service time fed by a Fractal Brownian Motion. In order to compare the behavior of multiple class selective discard mechanism under a long-range dependent process to the behavior under a short-range dependent process, we also simulate a queue fed by a two state Discrete Time Batch Markovian Arrival Process (D-BMAP). A D-BMAP is a short-range dependent process with batch arrivals in which the distribution of the arriving batch depends on the state of an underlying Markov Chain. To be more specific, we use a two-sate D-BMAP which correlation coefficient is given by $C_c(n) = \frac{\sigma^2 \rho}{1 + \sigma^2 \rho} * (2\alpha - 1)^n$ where ρ is the arrival rate, σ is

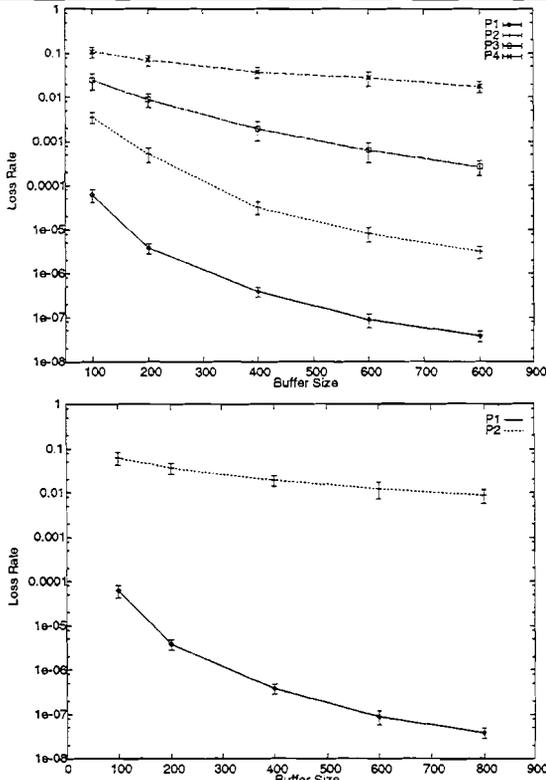


Figure 2. Loss rate \times buffer size for $\rho = 0.8$, $\sigma = 1.0$, $P_1 = 0.6$, $P_2 = 0.1$, $P_3 = 0.1$ and $P_4 = 0.2$. In a queue with two priority levels all classes except the highest one are mapped into the lowest priority class.

the standard deviation of the number of arrivals and a is the transition probability of embedded Markov Chain [19].

One clear advantage of having a multi-priority level scheme is the ability to provide differentiated loss services. Under a short-range dependent process this advantage is not so striking as it is under a long-range dependent process since the loss rate considerably decreases as the buffer size increases. Thus, to provide low loss rates to all service classes under a short-range dependent process we just need to have large buffers. Figure 2. illustrates this flexibility. In Figure 2..a and in Figure 2..b we show the loss rate as a function of the buffer size in a queue with four-priority levels and in a queue with two priority levels, respectively. P_i denotes the fraction of traffic carried with priority i , where the classes are numbered according to decreasing order of priority. In this example, the highest priority class of the multi-priority queue is mapped into the high priority class of the two priority level queue, whereas the other three classes are mapped into the low priority class. It bears noting that in a queue with two priority levels, we have to transport a fraction of the whole traffic with a loss rate much higher than we would do it in a four priority level queue.

However, as the buffer size increases so does the loss rate difference in a queue with four priority level. For instance, for large buffer (buffer size 800) the loss rates differ by two orders of magnitude, whereas for small buffer size (100) these loss rates differ by less than an order of magnitude. Note that in a queue with two priority levels and buffer size 800 we

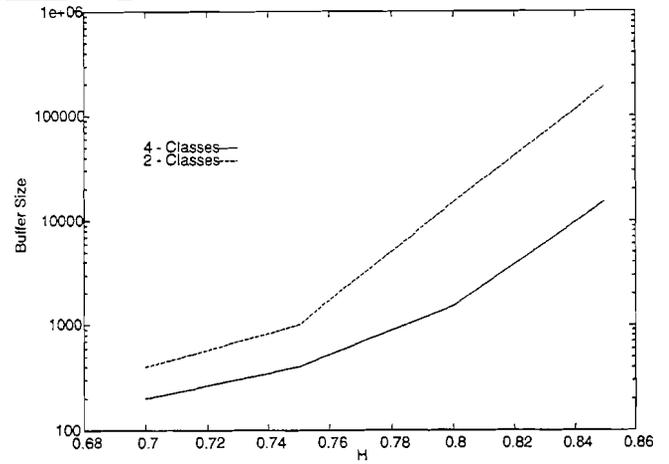


Figure 3. Buffer requirements for a four service classes (applications) with requirements 10^{-9} , 10^{-7} , 10^{-5} and 10^{-3} . In a queue with four priority levels each service class is mapped into a priority level. In a queue with two priority levels the two most stringent classes are mapped into the high priority class and the other two are mapped into the low priority class. The traffic parameters are $\rho = 0.8$, $\sigma = 1.0$, $P_1 = 0.6$, $P_2 = 0.05$, $P_3 = 0.15$ and $P_4 = 0.2$.

would have to carry all classes except the highest priority one with loss rate of 10^{-1} which is much higher than the loss rate for classes 2 and 3 in the four-priority queue.

It seems quite evident that with a higher number of priority levels we have a finer granularity of guaranteed loss rates. However, the great advantage behind the ability to furnish distinct loss bearers is the buffer savings of a multi-priority level mechanism when compared to the buffer requirement of a two-priority mechanism. In a two priority mechanism we typically need to map applications with diverse QoS into a single priority level. Therefore, to satisfy the loss requirement of the most stringent application mapped into a certain priority class we have to increase the buffer size so that the loss rate of this priority level satisfies the application loss requirement. Figure 3. illustrates a scenario in which there are four service classes (set of applications) with loss requirements 10^{-9} , 10^{-7} , 10^{-5} and 10^{-3} . We show the buffer requirements of a queue with four priority levels and the buffer requirement of a queue with two-priority levels to satisfy the mentioned loss requirements. In a queue with four-priority levels the four service classes are carried in distinct priority levels, whereas in a queue with two-priority levels the two most stringent applications are carried as the high priority class and two less stringent classes are carried as the low priority class. It can be observed that as the Hurst parameter increases so does the difference between the buffer requirement of a two-priority level queue and the buffer requirement of a four priority level queue. For a Hurst parameter of 0.85 the buffer requirement difference can be of one order of magnitude. Furthermore, such difference also depends on the offered load. In Figure 4. we can see that when the offered load increases from 0.8 to 0.9 for a stream with $H = 0.8$ the buffer requirement differs by two orders of magnitude!

Another advantage of a multi-priority mechanism over a two-priority mechanism is that we can have higher loads in a multi-priority system than we have in a two-priority system. In fact, we can increase the offered load by increasing the load of applications with less stringent QoS. For instance, in a two-priority queue with buffer size 500 and offered load of 0.65 ($H = 0.75$, $P_1 = 0.7$ and $P_2 = 0.3$), if we introduce two more priority levels, we are able to increase the offered load up to 0.9 by distributing the additional load among the two new lowest priority classes.

As mentioned before multiple class selective discard is more attractive under a long-range dependent process than it is under a short-range dependent process. The reason is that under a short-range dependent process we can satisfy diverse loss requirements with moderate buffer size. Figure 5.a illustrates the loss rate as a function of the buffer size for a four priority level queue under a short-range dependent process and Figure 5.b shows both the buffer requirements of a queue with two priority levels and the buffer requirement of a queue with four priority levels. In this scenario the three lowest priority classes of a queue with four priority levels are mapped into the low priority class of a queue with two priority levels. It is observable that low loss rates with small buffer sizes (< 50 buffer slots) can be achieved and the difference of buffer requirement between a queue with four levels and a queue with two levels is less than an order of magnitude.

Previous studies point out that the difference between the high priority loss rate produced by the Modified-First-In-First-Drop and the loss rate produced by the LIFD policy may be of two orders of magnitude in a queue with two priority levels. In this paper, we verified that the difference between loss rate of the highest priority class produced by different pushout policies in a multi-priority queue can be of three orders of magnitude. Furthermore, for low values of the Hurst parameters the difference decreases with the priority level. However, for high values of the Hurst parameter the loss difference can be noted only for the highest priority class.

Another major distinction between a queue with two-priority levels and a queue with four priority levels is that in a queue with two priority levels the low priority loss gap produced by M-FIFD can be three times longer than the loss gap produced by LIFD. Such difference does not exist in a multi-priority queue in which the loss gap of the lowest priority class is typically smaller than the loss gap of the low priority class in a queue with two-priority levels. Moreover, we observe that in a multi-priority queue we can satisfy a finer granularity of low priority loss gap requirements.

In summary, in a queue with two priority levels, the adoption of different push-out policies implies in a trade-off between offering low high priority loss rates and offering small low priority loss gaps. This trade-off does not exist in a queue with multi-priority levels. In a queue with multi-priority levels the Modified-First-In-First-Drop seems to be the most attractive policy since it minimizes the highest priority loss rate and produces low priority loss gaps of the same order of the low priority loss gaps produced by Last-In-First-Drop policy.

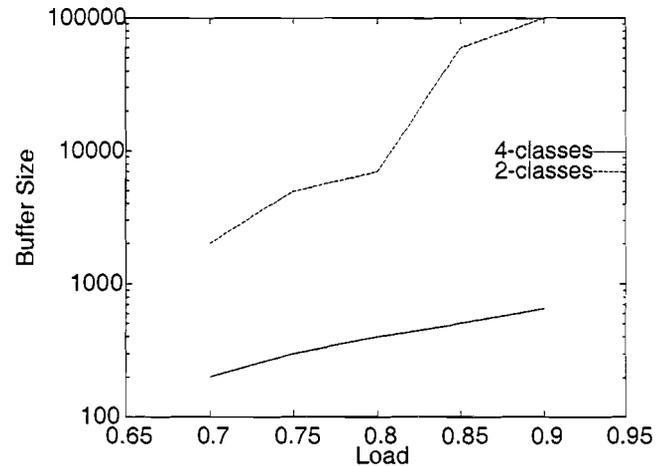


Figure 4. Buffer Size x Offered Load. The traffic parameters are $\rho = 0.8$, $H = 0.8$, $\sigma = 1.0$, $P_1 = 0.6$, $P_2 = 0.05$, $P_3 = 0.15$ and $P_4 = 0.2$.

7. CONCLUSIONS

In this paper we investigated the advantages of adopting a multiple-class mechanism over a two-priority mechanism. We showed that in a multi-priority queue, we can carry the same load with less buffer. In addition, we can carry higher loads than in a two-priority queue by increasing the load of less stringent applications.

If on the one hand, M-FIFD produces the lowest highest priority loss rate among all policies, on the other hand, the maximum loss gap length given by LIFD is in the same order of the maximum loss gap given other policies. Therefore, M-FIFD seems to be the most appropriate push-out policy to be adopted in an ATM multiplexer. Insofar as results given by FIFD are close to results given by M-FIFD and FIFD is less complex to implement (it requires less buffer shifting than M-FIFD), FIFD is a more attractive choice.

While multiple-priority mechanisms does not seem to be not so attractive under short range dependent process they are very appealing under long-range dependent process, since they offer differentiated service even when massive loss occurs.

REFERENCES

- [1] W. Leland, M. Taqqu, W. Willinger and D. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)", *IEEE/ACM Transaction on Networking*, vol 2, no 1, pp. 1-15, February 1994.
- [2] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", In *Proc. of ACM SIGCOMM*, 1994.
- [3] G. Mayor and J. Silvester, "A Trace-Driven Simulation of an ATM Queueing System with Real Network Traffic", *Proc. of IEEE ICCCN*, 1996.

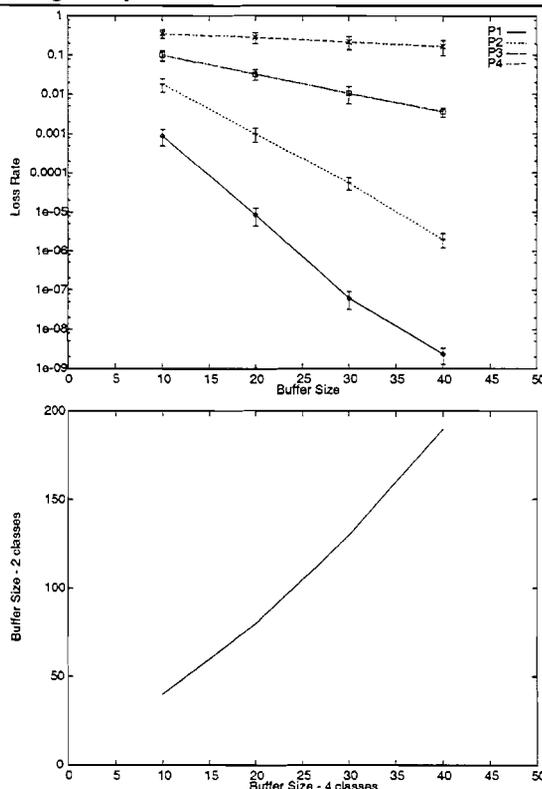


Figure 5. The bottom Figure (5.b) shows the buffer requirement for selective discard with 2 and 4 priority classes to support the loss rate shown in the top Figure (5.a). The traffic parameters are $\rho = 0.8$, $\sigma = 1.0$, $H = 0.75$ and $\alpha = 0.99$.

[4] G. Mayor and J. Silvester, "An ATM Queueing System with a Fractional Brownian Noise Arrival Process", Proc. of IEEE ICC, 1996.

[5] I. Norros, "A Storage Model with Self-Similar Input", Queueing Systems 16, 1994.

[6] N. Duffield, J. T. Lewis, N. O'Connell, R. Russel, F. Toomey, "Predicting Quality of Service with Long-range Fluctuations", In Proc of IEEE ICC'95, pp. 473-477, 1995.

[7] H. D. Sheng and S. Q. Li, "Spectral Analysis of Packet Loss Rate at a Statistical Multiplexer for Multimedia Services", IEEE/ACM Transactions on Networking, January 1994.

[8] R. Choen, N.L.S. Fonseca and M. Zukerman, "Traffic Management and Network Dimensioning", Multimedia Communications Networks: Technologies and Services, M Tatipamula and B. Kasnabish Editores, Artech House, 1998 .

[9] N.L.S. Fonseca, M.J. Ferreira and G.S. Mayor, "On the Effectiveness of Loss-Conserving Discipline Under a Long-Range Dependent Process", in Proc of Globecom'97, pp. 1102-1107, 1997

[10] N. L. S. Fonseca e J. A. Silvester, "Estimating the Loss Probability in a Multiplexer Loaded with Multi-priority

MMPP Streams", in Proc of IEEE ICC'93 , pp. 1037-1041, 1993.

[11] N.L.S. Fonseca e M.J. Ferreira, "Multiple Class Selective Discard Under a Long-Range Dependent Process", Proc of IEEE International Telecommunications Symposium 98, pp. 95-101, 1998

[12] A. Papoulis, "Probability, Random Variables and Stochastic Processes", A. Papoulis, 3rd edition, McGraw-Hill, Inc., 1978.

[13] B. Mandelbrot and J.Ness, "Fractional Brownian Motions, Fractional Noises and Applications", SIAM Review, October 1968.

[14] N. L. S. Fonseca e J. A. Silvester, "A Comparison of Push-out Policies in an ATM Multiplexer", In Proc. of IEEE Pacific Rim Conf on Comp. Comm and Signal Proc., pp.. 338-341, 1993

[15] L. Tassiulas, Y.C. Hung, and S.S. Panwar. "Optimal Buffer Control during Congestion in a ATM Network Node. IEEE/ACM Transactions on Networking", 2(4):374-386, Aug 1994.

[16] G. Mayor and J.Silvester, "Time Scale Anaysis of an ATM Queueing System With Long-Range Dependent Traffic", in Proc. of INFOCOM'97, 1997.

[17] R. G. Addie, M. Zukerman and T. Neame, "Fractal Traffic: Measurements, Modelling and Performance Evaluation", In Proc. of IEEE INFOCOM 1995.

[18] B. Ryu and A. Elwalid, "The Importance of Long-Range Dependence of VBR Video Traffic in ATM Traffic Engineering: Myths and Realities", In Proc. of ACM SIGCOMM, 1996.

[19] T. Takine, T. Suda and T. Hasegawa, "Cell loss and output process analyses of finite buffer discrete time queueing system with correlated arrivals", in Proc. of IEEE INFOCOM, pp 1259-1268, 1993

Nelson Fonseca received his Electrical Engineer (1984) and MSc in Computer Science (1987) degrees from the Pontifical Catholic University of Rio de Janeiro, Brazil, and the MSc (1993) and PhD (1994) degrees in Computer Engineering from University of Southern California. Since 1995, he has been an Assistant Professor in the Institute of Computing of the State University of Campinas, Campinas, Brazil. He is the recipient of the USC International Book award and of the Brazilian Computing Society First Thesis and Dissertations Award. Dr. Fonseca is listed in Marqui's Who is Who in the World and Who's Who in Science and Technology. He is the editor-in-chief of IEEE Communications Newsletter, an editor for Computer Networks (former Computer Networks and ISDN), technical editor of IEEE Communications Interactive and Associated technical editor for the IEEE Communications Magazine. He has served as guest editor for special issues of IEEE JSAC, IEEE Communications Magazine and the Journal

of the Brazilian Computer Society. He has also served as technical committee member of IEEE Globecom'97 and SBT/IEEE ITS'98. Dr. Fonseca was the chairman of IEEE CAMAD'98. His main interests are ATM traffic control, and video services.

Marcelo de Jesus Ferreira obtained his Bachelor degree in Computer Science from the University of Amazonas (1995) and the MSc degree in Computer Science from the State University of Campinas. Since 1998 he has been affiliated with the Brazilian Lutheran University, Manaus, Brazil. His Master Dissertation was awarded in the 1998 latin American Master Thesis Contest sponsored by UNESCO and by CLEI. His main interest is ATM traffic control.