On the Effectiveness of the Longest-Packet-In Selective Packet Discarding Policy

N.L.S Fonseca e M.J. Ferreira

Relatório Técnico IC 98-04

Março de 1998

On the Effectiveness of the Longest-Packet-In Selective Packet Discarding Policy

Nelson L .S. Fonseca and Marcelo J. Ferreira

State University of Campinas
Institute of Computing
P.O. Box 6176
13083-970 Campinas SP
Brazil
Phone/Fax: +55+19+2393115 or 7888442
e-mail: {nfonseca, mjf}@dcc.unicamp.br

Abstract

In this paper, we introduce a novel packet discarding policy called Longest-Packet-In (LPI) which maximizes the cell goodput, the ratio of outgoing good cells to the total number of cells that arrive at a queue. We compare LPI to: Tail Drop, Early Packet Discard, Early Packet Discard with Hysteresis and Fair Early Packet Discard with Hysteresis by investigating the trade-off between maximizing the cell goodput and maximizing the packet goodput. We show that LPI not only provides the highest cell goodput, but also produces the highest packet goodput for a high ratio of the mean packet size to the buffer size and under a short-range dependent process. We evaluate the impact of network load, mean packet size, packet distribution and the Hurst parameter into both the cell goodput and into the packet goodput. Furthermore, we extend our analysis to networks which carry sources with distinct mean packet length.

I) Introduction

In ATM networks, user information is transferred in fixed size cells consisting of 48-byte payload and 5-byte header. As user data are usually larger than a cell payload, a transport level packet is segmented into several ATM cells. The ATM Adaptation Layer (AAL) provides segmentation of variable-size packets into cells at the transmitter side and reassembly function at the receiver side. The ATM standard specifies error control to cell headers, but not to the payload field. Payload error control is handled in an end-to-end basis by the transport layer above the ATM layer. For some applications if one of the cells of a packet is lost, the whole packet has to be retransmitted. Therefore, transmitting cells of a corrupted packet wastes bandwidth, and may increase congestion. In this vein, several packet discarding policies have been defined to preserve the integrity of a higher number of packets [1]-[4].

If on one hand, rate base flow control for available bit rate (ABR) traffic may reduce the need for packet discarding mechanisms, on the other hand, it does not eliminate the need for such mechanisms since buffers using ABR flow control may still face overload periods before rate control mechanisms may react to traffic changes. Actually, as adaptative window mechanisms, rate-flow control mechanisms may need one or more network round trip delay to control congestion. Moreover, rate base flow control will not be universally applied. For instance, it will not be applied to unspecified bit rate (UBR) streams.

Several studies have claimed that different types of network traffic can be accurately modeled by a self-similar process. A self-similar process is able to capture the long-range dependence phenomenon exhibited in this traffic. In addition, series of simulation and analytical studies demonstrated that this phenomenon may have a pervasive effect on queueing performance, *i.e.*, there is a clear evidence that it can potentially cause massive cells loss in ATM multiplexers. Moreover, queueing systems with self-similar input may suffer from the buffer inefficacy phenomenon, *i.e.*, by just increasing the buffer size, we are not able to decrease the buffer overflow probability considerably [5]-[7].

In this paper, we introduce a novel packet discarding policy which maximizes the cell good put, *i.e.*, the ratio of outgoing good cells to the total number of cells that arrive at a queue. In other words, instead of maximizing the number of successfully transmitted packets, we maximize the bandwidth used to carry non-corrupted packets. We compare our policy to existing ones by investigating the trade-off between the cell goodput and the packet goodput (the ratio of good packet on the outgoing flow to the total number of packets which arrive at a queue). We show that our policy can significantly improve the cell goodput and is also able to produce the highest packet goodput for a high ratio of the mean packet size to the buffer size and under a short-range dependent process. We evaluate the impact of network load, mean packet size, packet distribution and the Hurst parameter into both the cell goodput and into the packet goodput. Furthermore, we extend our analysis to networks which carry sources with distinct mean packet length.

This paper is organized as follows. Section II briefly review existing packet discarding policy. Section III introduces the Longest-Packet-In packet discarding policy. Section IV describe the fractal Brownian motion process. Numerical examples are shown in section V and, finally, conclusions are drawn in section VI.

II) Previous Work

For some applications, if one cell of a packet is lost, the whole packet has to be retransmitted. Hence, transmitting cells of a corrupted packet wastes bandwidth, and consequently may increase congestion in already congested nodes. In line with that, several packet discarding policies have been recently defined. In Tail drop (TD), if a cell is lost subsequent cells of the same packet are discarded. In Early-Packet-Discard (EPD), a queue threshold is defined and an incoming packet is discarded if it finds the buffer occupancy above this threshold [1]. Both TD and EPD can significantly improve the goodput when compared to systems with no control at the packet level.

Several studies have investigated these two policies. The seminal work of Ramanow and Floyd concluded that for fixed size packets EPD gives higher goodput than TD does [1]. Lapid et al. [4] analyzed a system with Poisson input and packet size geometrically distributed. They found out that TD performs better than EPD for moderately loaded systems, whereas EPD is preferable for heavily loaded systems. Moreover, they found out that an optimal threshold do exists but is not significant at all in terms of cell goodput. Tsukumatani et. al [2] considered a system loaded with several on-off sources with geometrically distributed periods. They claim that EPD threshold gets smaller as mean packet size increases.

In order to improve EPD for small buffer sizes, Turner defined the Early-Packet-Discard with Hysteresis (EPDH) discipline [3]. In EPDH, a source (connection) can be either active or inactive. Whenever a source is considered inactive all its packets are dropped. EPD-H defines an additional queue level called floor level which should be set to some small number of cells. At the end of a packet transmission, if the queue occupancy exceeds the threshold and the queue level since the most recent threshold crossing, the source is considered inactive. Furthermore, a sources is also considered inactive if it loses a cell. On the other hand, if the queue level is both below threhold and is either below the floor level or below the minimum level since the last threshold crossing, the source becomes active. By observing that sources with higher transmission rates tend to start transmission before the queue level rises above the threshold more often than do sources with lower rates, Turner defined another variant of EPD called Fair Early-Packet-Discard with Hysteresis (F-EPDH). Besides EPDH rules, F-EPDH defines a high and a low threshold. If the current queue level is between the two threshold, and the queue level is failing, a source is marked as active. Conversely, if the queue level is rising, the source is marked as inactive. To the best of our knowledge a numerical evaluation of EPDH and of F-EPDH policies has not been carried out before.

III) The Longest-Packet-In Packet Discarding Policy

Existing packet discarding policies aim at maximizing the number of successfully transmitted packets irrespective of their size. For instance, in Tail Dropping a packet is discarded if one of its cells is lost, however, the discarded packet size may be larger than enqueued packets. In Early Packet Discard a packet may be discarded even if there are available buffer slots to accommodate it. To overcome this drawback, we define a packet discarding policy which maximizes the number of cells in successfully transmitted packet. Since cells are fixed size, maximizing the number of cells in successfully transmitted packets is equivalent to maximize the bandwidth used to transmit integral packets.

The Longest Packet In (LPI) policy always admits an incoming packet to the buffer if there are available buffer slots. Otherwise, it discards one or more enqueued packets if their total size is smaller than the incoming packet size. We assume that consecutive cells of a packet arrive at consecutive slots and that the packet length information is obtained from the first cell of a packet so that the decision of accepting an incoming packet into the buffer is made at the arrival time of its first cell. Any packet with its first cell still in queue is eligible to be dropped, *i.e.* any packet whose cells has neither left the queue nor is in service is eligible for dropping. If a packet is chosen to be discarded, all its enqueued cells are simultaneously dropped and any further cell of this packet is no longer admitted to the buffer. The LPI policy is described below:

If $\left(D = L + \sum_{i \in S} u_i - B + b - max\left(L, \max_{i \in S}(u_i)\right) \le 0\right)$

enqueue the incoming packet

else Solve Problem I:

$$Min \sum_{i \in Q} l_i x_i$$

Subject to:

$$L > \sum_{i \in Q} l_i x_i \ge D$$

$$x_i \in \{0, 1\}, \; i = 1 \ldots |Q|$$

if (Problem I admits a feasible solution $\tilde{x}$)

discard packet i such that $\tilde{x}_i = 1$

else

discard the incoming packet

where:

S - Set of sources transmitting non-corrupted packets at the time of the target (incoming) packet, excluding the incoming packet source;

Q - set of sources with its first cell still in queue;

l_i - size of the i^{th} packet;

u_i - number of cells of the i^{th} packet, $i \in S$, yet to arrive at the beginning of the time slot during which the incoming packet arrives;

B - Buffer size;

b - queue length at the time of the target (incoming) packet arrival;

L - Size of the incoming packet;

$\tilde{x}$ - Problem I feasible solution.

Figure I illustrates the LPI policy.

Note that when we verify whether there is available buffer space to enqueue the incoming packet (if ($D \leq 0$) we consider that $max\left(L, \max_{i \in S}(u_i)\right)$ buffer slots are released during the transmission of active sources. If there is not enough available buffer space, we seek a set of packet with minimum total size (Min $\sum_{i \in Q} l_i x_i$) which if discarded releases the necessary number of buffer slots $\left(\sum_{i \in Q} l_i x_i \geq D\right)$ to admit the incoming packet. Obviously, the total size of dropped packets should be less than the incoming packets $\left(L > \sum_{i \in Q} l_i x_i\right)$, otherwise we decrease the cell goodput.

Problem I is a knapsack problem. We can solve a knapsack problem in *O* (*c x n*) where *c* is the knapsack capacity and *n* is the number of objects. In our problem, the knapsack capacity is the number of buffer slots to be released in order to admit the incoming packet (*D*), and the objects are the packets to be discarded. Note that *D* takes into account the queue length growth up to the transmission time of the last cell of active sources, *i.e.*, it takes into account the number of cells of active sources yet to arrive and the number of buffer slots to be released in a time window which starts at the arrival of the incoming (target) packet and ends at the transmission time of the last cell of the longest non-corrupted packet.

Dropping packets whose first cell are still in queue is a conservative approach, since higher cell goodput could be achieved in a multiplexer in isolation if we choose packets to be discarded among all packets in queue. In fact, choosing a packet among all enqueued packets or selecting a packet among packets whose first cells are still in queue is the trade-off between maximizing the cell goodput at a multiplexer and avoiding clogging downstream multiplexers with cells of dropped packets. From numerical examples shown in section IV, we notice that enqueueing cells of corrupted packets may have a detrimental effect on the cell goodput. Since downstream multiplexers are fed by the output of several multiplexers, the number of useless cells enqueued at downstream queues may increase considerably. Preventing cells of corrupted packets to flow along the network can always be obtained at a higher signaling cost. Nevertheless, such overhead is not desirable at high speeds.

IV) The Fractional Brownian Motion

Several studies have claimed that different types of network traffic can be accurately modeled by self-similar processes. A self-similar process is able to capture the long-range dependence phenomenon exhibited in this traffic. A process with long-range dependences is a process whose auto-correlation function decays very slowly as a function of time. The fractal Brownian motion is a Gaussian self-similar process which accurately captures the self-similar nature of network traffic [8].

The ordinary Brownian motion, $B(t)$, describes the movement of a particle in a liquid subjected to collisions and other forces [9]. It is a real random function with independent Gaussian increments such that

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

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

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 $(t\text{-}s)^{h - 1/2}$ [10].

Definition:

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

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

This equation leads to the ordinary Brownian motion if $H = 1/2$. Its self-similar property is based on the fact that $B_H(\rho s)$ 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):

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

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. The Hurst parameter, H, gives the index of self-similarity of a process. The higher the value of H, the slower the autocorrelation decay is.

Although the fBm accurately describes aggregated network traffic, it does not provide an explanation for the self-similar phenomenon at the source level. Willinger et. al [11]-[12]showed that the aggregation of a large number of individual *on-off* sources with strict alternating *on* and *off* periods generates a fBm if either the distribution of *on* periods or the distribution of *off* periods is a heavy-tail, *i.e.*, if the complementary (or tail) distribution of the duration of either *on* or the complementary distribution of *off* periods can be expressed as:

$$F(x) \sim lx^{\alpha} L(x)$$

where *l* is a constant and *L (x)* is a slowly varying function at infinity. One of the distributions which present such behavior is the Pareto distribution.

An interesting characteristic of the convergence property is the relationship between the parameter α of the aggregating sources and the parameter *H* of the aggregated process, which is given by $H = (3 - \alpha) / 2$. In the heterogeneous case, each set of *on-off* sources with α_i converges to a fBm with parameter H_i and the highest H_i dominates the aggregated process in the long-run ($t \rightarrow \infty$).

V) Numerical Examples

To assess the effectiveness of the Longest Packet In policy, we compare it to: Tail Drop (TD), Early Packet Discard (EPD), Early Packet Discard with Hysteresis (EPDH) and Fair Early Packet Discard with Hysteresis (FEPDH). We analyze the trade-off between maximizing the number of cells in successfully transmitted packets and maximizing the number of successfully transmitted packets. Furthermore, we numerically evaluate EPDH and FEPDH. To the best of our knowledge such evaluation has not been carried out before.

It is worth noting that previous studies adopted different performance objective to compare discarding policies. Turner [3] and Romanow and Floyd [1] considered the effective throughput which is ratio of good cells on the outgoing flow to the total outgoing cells. Lapid et. al defined goodput as the ratio of outgoing good cells to the total number of cells that arrive at a queue. We define cell goodput as is in Lapid et al [4]. The cell goodput shows how much of the traffic is not wasted. It differs from the effective throughput in the sense that effective throughput does not take into account discarded cells. Effective throughput consider just the outgoing flow. By definition, the Longest-Packet-In policy produces effective throughput 1.0, since LPI allows only good cells in the outgoing flow. Moreover, we also show the packet goodput which is the ratio of good packet on the outgoing flow to the total number of packets which arrive at a queue. The packet goodput gives the percentage of user's transmitted packet which is successfully transmitted. In summary, the cell goodput gives the effectiveness of a policy at the link level, while the packet goodput indicates the efficiency of a policy at the transport level.

Given that a queue with LPI is not amenable to analytical analysis, we use event-driven simulation. The replication method is used to derive 95% confidence intervals. In the simulation experiments, we consider a set of queues. Each queue is under a different discarding policy and its input is the aggregated process of several of *on-off* sources (Figure 2). While in state *on*, a source produces a cell of a packet at every time slot, *i.e.*, a source generates a different packet at every visit to state *on* and the length of the packet is the duration of the *on* period. *Off* periods represent interpacket generation time. To evaluate how different policies perform under diverse traffic patterns, we generate short-range and long-range dependent processes by using exponentially and Pareto distributed *on-off* periods, respectively. Moreover, we investigate whether results for long-range dependent processes vary when inter-packet generation time is exponentially distributed.

The offered load is given by $\rho = N \; x \; \varepsilon$, where N is the number of sources, and $\varepsilon = T_{on} / (T_{on} + T_{off})$ is the activity ratio, *i.e.*, is the stationary probability of a source being in state *on*. T_{on} and T_{off} are the mean duration of *on* and *off* periods, respectively. We vary the offered load by changing either N or ε, and we comment on discrepancy in results whenever they occur.

In the simulation experiments, we investigate the impact of buffer size, mean packet length and threshold value on the cell goodput, and on the packet goodput. We show findings for buffer size 100 and 400, and a threshold value of 70% of the buffer size. We notice that for threshold values over 80%, results given by EDP policies tend to values given by TD. Results for thresholds values in the range between 50% and 70% are quite similar, and we comment any divergency. These findings are consonant with results in [4]. For EPDH we show results for low threshold of 60% of the buffer size and for high threshold of 80%. We show results for streams with mean packet length of 42 ATM cells. Finally, we fix the activity ratio of each source to 0.025.

V.a) On-off sources with geometrically distributed periods

We first consider geometrically distributed *on-off* periods. In Figures 3 and 4, we show the cell goodput and the packet goodput as a function of the network load for mean packet length of 42 ATM cells, and for buffer size 100 and 400, respectively. We vary the load by increasing the number of sources. Each source contributes 0.025 to the offered load which gives an interpacket arrival time of 1638 time slots. As the load increases, the goodput, obviously, decreases. Nonetheless, LPI gives the highest cell goodput. For small buffer sizes (100) and for network load of 1.0, while the cell goodput given by LPI is above 0.78, the cell goodput given by TD and (F)EPDH are below 0.65 (note that the goodput given by EPDH and F-EPDH are indistinguishable). The difference between the cell goodput produced by LPI and the cell goodput produced by (F)EPDH can be as high as 0.2. As we increase the buffer size (to 400), the difference between the cell goodput given by LPI and the cell goodput given by other policies decreases. However, this difference can still be as high as 0.1. In addition, LPI cell goodput has the lowest rate of decrease. For instance, for buffer size 400, LPI cell goodput decreases 0.06 whereas (F)EPDH decreases 0.15.

For buffer size 100 and mean packet length of 42 ATM cells, *i.e.*, for a high ratio of the mean packet size to the buffer size, LPI gives the highest packet goodput. It happens because under LPI buffer slots are occupied only by cells of non-corrupted packets. Conversely, under other policies, buffer slots may be occupied by cells of corrupted packets, and, consequently, enough slots to accommodate an incoming packet may not exist. Furthermore, LPI typically needs to drop only one packet to accommodate an incoming packet. For buffer size 400, TD gives the highest packet goodput. Nevertheless, the cell goodput difference between TD and LPI is at most 0.04.

In Figure 5 and 6, we show the relationship between the goodput and the mean packet size for buffer size 100 and 400, respectively. We fix the number of sources to 32, which gives an offered load of 0.8. Hence, to maintain constant the activity ratio as we vary the mean packet length (duration of *on* periods), we also vary the mean duration of *off* periods. For packet size smaller than 10 ATM cells, the goodput produced by different policies are almost indistinguishable. As the packet size increases, the cell goodput decreases because available buffer space is normally smaller than incoming packets. For small buffer sizes (100), while LPI cell goodput is above 0.7, the cell goodput given by EPDH and F-EDPH policies drop below 0.57, which might be unacceptable in a real network. This illustrates the benefit of occupying the buffer space only with cells from non-corrupted packets. For large buffer size (> 400), the difference between the cell goodput produced by LPI and the goodput produced by other policies decreases. Nonetheless, for large packet size (> 40 ATM cells), LPI is still able to keep the cell goodput at least 0.1 higher than any other policies. The difference between LPI cell goodput and (F)EPDH can be as high as 0.2.

The packet goodput also decrease as the packet size increases. For small buffers, the difference between the packet goodput produced by LPI and the packet goodput produced by other policies can be 0.05 and the difference between LPI goodput and (F)EPDH goodput be as high as 0.15. For large buffers (> 400 buffer slots) and packet size smaller than 50 ATM cells, Tail Drop produces the highest packet goodput, whereas for packets larger than 50 ATM cells, LPI gives the highest packet goodput. For long packets, the difference between LPI packet goodput and (F)EPDH packet goodput can be 0.13. For packets smaller than 50 ATM cells, LPI may drop more than one packet to enqueue an incoming long packet, whereas TD loses at most one packet in overflow situations. As the packet size increases, so does the probability of having long packets in queue, and, consequently, LPI typically drops at most one packet to accept an incoming one. Furthermore, under other policies as packet size increases the buffer is clogged with a higher number of useless cells. It is quite noticeable that the packet goodput produced by EPDH and by F-EPDH sharply decreases for long packets, whereas the packet goodput produced by TD and EPD have smooth decays. In summary, LPI gives the highest cell goodput irrespective of the packet, and produces the highest packet goodput for long packet.

In our simulation experiments, we notice that if instead of increasing the number of sources, we increase the activity ratio, *i.e.*, we keep the number of sources and the mean duration of *on* periods constant and decrease the mean duration of *off* periods, LPI cell goodput and LPI packet goodput roughly stays at the same value, while the cell goodput and the packet goodput produced by other policies increase 0.05. In other words, for a fixed number of sources with higher activity ratio, LPI results are slightly less advantageous than they are for a higher number of sources with lower activity ratio.

To understand the extend to which findings from the previous example are dependable on the ratio of packet size to the buffer size, we evaluate the goodput as a function of this ratio, *i.e.*, we simultaneously increase the mean packet size and the buffer size for a fixed load. From Figure 7, we conclude that both the cell and the packet goodput are almost the same for a fixed ratio of the mean packet size to the buffer size.

To evaluate discarding policies in networks with heterogeneous source, we consider four different sources with distinct mean packet length. We fix the total number of sources to 32, which gives an offered load of 0.8, and we vary the proportion of each type of source. We consider sources with mean packet length of 5, 21, 85 and 170 ATM cells, and we initially distribute the load equally among the four set of sources. We, then, vary the number of sources of a specific set, and distribute the remaining load among the other three sets. In Figure 8, for buffer size 400, we vary the number of sources with mean packet length of 170 ATM cells. In the horizontal axis, we show the proportion of sources of this type. As the number of long packets increases the cell goodput decreases since we increase the chance of an incoming packet to find a low number of available buffer slots. Nevertheless, LPI gives the highest cell goodput. While LPI goodput is above 0.9, (F)EPDH cell goodput is 0.71. LPI also has the lowest cell goodput rate of change. In this specific example, while LPI cell goodput decreases 0.05, the cell goodput produced by other policies decreases at least 0.1. For small buffer size (100), the difference between LPI cell goodput and the cell goodput produced by other policies can be as low as 0.1 (EPD) and as high as 0.3 (EPDH and F-EPDH).

Although TD and EPD packet goodput are higher than LPI packet goodput, the difference between the highest packet goodput (TD packet goodput) and LPI packet goodput is less than 0.05, irrespective of the buffer size. In summary, LPI provides the highest cell goodput, and its packet goodput does not meaningfully differ from the packet goodput given by other policies.

However, for networks with a higher proportion of small packets, we have a considerably different picture. In Figure 9, we increase the proportion of sources with mean packet length of 5 ATM cells and distribute the remaining load among the other policies. As the proportion of small packets increases, the cell goodput also increases, and the difference between cell goodput given by different policies decreases. Once again, LPI packet goodput is lower than both TD and EPD packet goodput, but the difference between TD packet goodput and LPI packet goodput is again less than 0.05. For large buffer size (> 500) the cell goodput given by different policies are indistinguishable.

V.b) On-off Heavy-tail Sources

Self-similar process may have a pervasive effect on queueing performance, *i.e.*, by just increasing the buffer size, we are not able to decrease buffer overflow probability considerably. Therefore, it is of paramount importance to assess the performance of packet discarding policies under self-similar processes. In line with that, we generate a fractal Brownian motion process by aggregating *on-off* heavy-tail sources. We use the Pareto distribution for both the duration of *on* periods and for the duration of *off* periods. In our simulation experiments, we verified that 20 sources is usually enough to generate an aggregated self-similar process. To make sure that the aggregated process was indeed a long-range dependent process, we collected for each simulation experiment the number of cells (bytes) generated at every time slot. We, then, used the variance-time analysis to verify if the Hurst parameter of the aggregated stream is the specified value. In other words, we checked whether the Hurst parameter of the aggregated stream matched $(3 - \alpha) / 2$ where α is the shape of the Pareto distribution.

In Figure 10, we show the cell goodput as a function of the offered load for buffer size 400, mean packet length of 42 ATM cells, and for different values of the Hurst parameter. As the Hurst parameter increases, overload periods get longer, and consequently the cell goodput decreases. LPI gives the highest cell goodput, and EPD cell goodput is close to LPI cell goodput. LPI is able to keep the cell goodput over 0.74, even under a stream with $H = 0.9$. Conversely, the goodput produced by TD and by (F)EPDH drops below 0.6. We notice that long-range dependencies considerably impacts the cell goodput produced by all policies. For instance, for a load of 0.7, LPI cell goodput is 0.95 for $H = 0.7$, whereas it is 0.8 for $H = 0.9$. Furthermore, TD, EPDH and F-EPDH are more sensitive to the Hurst parameter than LPI is. For example, TD cell goodput for an offered load of 0.7 and for $H = 0.7$ is 0.88 while for $H = 0.9$ it is 0.64. As the Hurst parameter increases, overload periods are much longer than the mean *on/off* residence time. Since TD uses any available buffer slot in attempting to accommodate a new packet, it clogs the queue with a higher number of useless cells than it does under a short-range dependent process. Thus, an incoming packet may find less available buffer slots than needed, while buffer slots are occupied by useless cells. EPDH and F-EPDH are more sensitive to the Hurst parameter because during long overload periods, there is a high chance that the queue level occupancy at the end of a packet transmission be above the last threshold crossing. In other words, EPDH and F-EPDH tend to consider a source as inactive unnecessarily. Besides giving the highest cell goodput, for a fixed value of H, LPI has the lowest rate of decrease.

For buffer size 100, and mean packet length of 42 ATM cells (*i.e.*, for a high ratio of the mean packet size to the buffer size) LPI also has good performance. For a stream with Hurst parameter of 0.9, LPI cell goodput is above 0.8, whereas the cell goodput given by TD, EDPH and (F)EDPH are below 0.6. In other words, even for high values of the Hurst parameter, LPI is able to maintain a reasonable bandwidth utilization, whereas TD, F-EPDH and EPDH utilizes less than 60% of the bandwidth to carry useful information. For small buffers (100), EPD cell goodput is no longer similar to LPI goodput, because in this scenario the threshold mechanism discards packets unnecessarily. Actually, it is 0.1 lower than LPI cell goodput. As H increases, TD goodput gets closer to F-EPDH and EPDH goodput.

If we compare the cell goodput under a short-range dependent process (exponentially distributed *on-off* periods) to the cell goodput under a long-range dependent process, we observe that the cell goodput is lower under a long-range dependent process than it is under a short-range dependent process. For example, for buffer size 400, and offered load 0.8, TD cell goodput is 0.96 under a short-range dependent process while it is 0.62 under a stream with $H = 0.9$. Nevertheless, LPI goodput under a long-range dependent stream decreases less than any other policies. For instance, for a load of 0.7 and for $H = 0.9$, LPI cell goodput drops 0.18 when compared to short range dependent process while (F)EPDH cell goodput drops 0.33.

On the other extreme, LPI packet goodput is highly sensitive to the Hurst parameter and it gives the lowest packet goodput among all policies irrespective of the buffer size (Figure 11). It happens because as overload periods get longer, LPI drops a higher number of small packets to enqueue long incoming packets. Although, LPI produces the lowest packet goodput, it is above 0.775, even for $H = 0.9$. The difference between the highest packet goodput, TD packet goodput, and LPI packet goodput is at most 0.2.

It bears noting that irrespective of the correlation-decay pattern of the input process (short or long range dependent), TD cell goodput is always lower than EPD cell goodput, whereas TD packet goodput is always higher than EPD packet goodput. This trend can be easily understood if one considers the fact that during overload TD tends to clog the queue with a higher number of useless cells. Therefore, at a certain time slot, a long incoming packet may find EPD queue length below the threshold value whereas it may not find enough available buffer slots in TD queue. On the other hand, TD may admit a higher number of small packets when EPD queue length is above the threshold value. However, the total size of a high number of small packets admitted under TD may be lower than the total size of a small number of long packet admitted under EPD.

In Figure 12, we show the cell goodput as a function of the mean packet size for buffer size 400, and for different values of H. As the packet size increases, the cell goodput obviously decreases. Nonetheless, LPI produces the highest cell goodput. The Hurst parameter impacts considerably the cell goodput. Streams with a high mean packet size are more affected than streams with a low mean packet size, *i.e.*, long packets have higher probability of being dropped during overload periods, and this probability increases as overload periods get longer. For instance, LPI cell goodput for packet size smaller than 10 ATM cells and for H = 0.7 is 0.98, while for H = 0.9 it is 0.875. Additionally, for H = 0.9 and for a stream with mean packet size smaller than 10 ATM cells, LPI cell goodput is 0.875 while it is 0.625 for mean packet size of 350 ATM cells. Moreover, LPI is less affected by the Hurst parameter than are other policies. For example, for H = 0.9 and mean packet size less than 10 ATM cells TD cell goodput is 0.75, whereas for mean packet size of 350 cells it is 0.45. Besides giving the highest cell goodput, LPI has the lowest rate of change.

For small buffer size (100), *i.e.*, for a higher ratio of the packet size to the buffer size, the cell goodput is lower than for a lower ratio of the packet size to the buffer size. Nevertheless, LPI is still able to maintain a reasonable goodput while (F)EDPH cell goodput drops to unacceptable values. For instance, for a packet size of 350, LPI cell goodput is 0.65 while EDPH goodput is 0.35.

We notice that the cell goodput is typically lower under a short-range dependent process than it is under a long-range dependent process. For example, for a packet size of 170, buffer size 400 and under a short range dependent process, LPI cell goodput is 0.87 while it is 0.725 under a stream with H = 0.9. Furthermore, under a short-range dependent process the cell goodput for small packet size are almost indistinguishable which does not happen under a long-range dependent process, specially for high H values.

On the other hand, both the Hurst parameter and the mean packet size influence the packet goodput (Figure 13). Although LPI gives the lowest packet goodput for high value of the Hurst parameter, it is always higher than 0.725. While under a short-range dependent process and for long packet, LPI produces the highest packet goodput, under a long-range dependent process, it produces the lowest packet goodput. For sources which generates long packets in average, as overload periods get longer and, the probability of a source producing small packets increases, so does the probability of a long incoming packet finding small packets in queue. Hence, for long overload periods LPI may drop more than one packet to accommodate an incoming packet.

To investigate discarding policies in heterogeneous networks, we also consider sources with mean packet length of 5, 21, 85 and 170 ATM cells. We initially distribute the load equally among all set of sources. We then increase the number of sources of a specific set and distribute the remaining load equally among the other sets. In Figure 14, we vary the number of sources with mean packet length of 170 ATM cells for a fixed load of 0.8 and buffer size 400. LPI produces the highest cell goodput. The Hurst parameter has a great influence on the cell goodput. Nonetheless, LPI cell goodput is above 0.67 even for H = 0.9. Again, keeping high cell goodput while other policies goodput decay considerably illustrates the advantages of maximizing the buffer occupancy only with good cells. On the other extreme, in heterogeneous networks with a high percentage of long packets, LPI typically drops more than one packet to admit an incoming packet, and, therefore it produces the lowest packet goodput. As the Hurst parameter increases, LPI packet goodput drops considerably. While TD and EPD are able to maintain a high packet goodput, LPI packet goodput drops to 0.67 for H = 0.9 (Figure 15). These results differ from results given by short-range dependent process. Under short-range dependent process, LPI produces the highest cell goodput and also preserves a packet goodput value close to TD and EPD packet goodput. When we vary the proportion of packets with mean of 5 ATM cells, LPI cell goodput is over 0.85 even for H = 0.9 and LPI packet goodput differs by at most 0.1 from the highest packet goodput (TD). In other words, for small packet size and under long-range dependent process, results are quite similar to results under short-range dependent process.

V.c) Source with Pareto distributed *on* periods and with geometrically distributed *off* periods.

The convergence theorem which establishes that the aggregation of a large number of on-off heavy tail converges to a fractal Brownian motion does not require that both *on* and *off* periods be heavy-tail. To investigate whether results might be affected by assuming geometrically distributed off periods, we repeated the same experiments described in last section. No significant difference was noted.

VI) Conclusions

In this paper, we introduced a packet discarding policy called Longest-Packet-In which maximizes the number of cells in successfully transmitted packets. We also compared LPI to existing policies by investigating the trade-off between the cell goodput and the packet goodput. We show through extensive simulation results that LPI produces the highest cell goodput among all policies. Furthermore, for a high ratio of the mean packet size to the buffer size and under a short range dependent process, it also gives the highest packet goodput. However, under a long-range dependent process LPI packet goodput is the lowest one. Under a long-range dependent process Tail Drop gives the highest packet goodput. Early Packet Discard with Hysteresis and Fair Early Packet with Hysteresis may give poor performance irrespective of input traffic characteristics.

LPI requires that the packet size be transmitted in the first cell of a packet. Since half of a cell payload is statistically wasted when we chop a variable size packet into a certain number of fixed size cells. Therefore, the unused payload may carry the packet size information implying in a overhead of less than one cell per packet. Nonetheless, during congestion overhead cells may impact network performance. Future research should investigate the effect of overhead cells.

Implementing LPI in real time demands that we solve a knapsack problem which can be done in $O(c \times n)$ where c is the number of buffer slots to be made available for the incoming packet and n is the number of eligible packets to be dropped. Since LPI selects packets to be discarded among packets with their first cell still in queue, the hardware complexity can be reduced. Choosing a packet to be discard among the packets with their first cells still in queue or selecting among all enqueued packets is the trade-off between maximizing the cell goodput in a multiplexer in isolation and clogging downstream multiplexers with useless cells. Further investigations should address this issue.

VII) References

[1] A. Romanow and S. Floyd, “Dynamics of TCP Traffic over ATM Networks”, *IEEE JSAC*, pp. 633-641, May, 1995.

[2] N. Tsukutani, K. Kawahara, T. Takine, H. Sunahara and Y. Oie, “Throughput Analysis of Selective Cell Discard Schemes in Transport layer over ATM Networks”, *In Proc. of IEEE Globecom’97*, pp. 1067-1074, 1997.

[3] J. S. Turner, “Maintaining High Throughput During overload in ATM Switches”, *In Proc. INFOCOM'96*, 287-295, 1996.

[4] Y. Lapid and R. Rom and M. Sidi, “Analysis of Packet Discarding Policies in High-Speed Networks”, *In Proc. of INFOCOM'97*, pp. 1191-1199, 1997.

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

[6] 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.

[7] N.L.S. Fonseca, M.J. Ferreira and G.S. Mayor, “On the effectiveness of Loss-Conserving Disciplines under a Long-range Dependent Process”, *In Proc. of IEEE Globecom’97*, pp. 1102-1107, 1997.

[8] I. Norros, “A Storage Model with Self-Similar Input”, Queueing Systems 16, 1994.

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

[10] B. Mandelbrot and J.Ness, “Fractional Brownian Motions, Fractional Noises and Applications”, *SIAM Review*, October 1968.

[11] W. Willinger and M. S. Taqqu and R. Sherman, Proof of a Fundamental Result in Self-simlar Traffic Modeling, Computer Commun. Review, 1997.

[12] W. Willinger and M. S. Taqqu and R. Sherman and D. V. Wilson, “Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level”, pp. 71-86, Feb 1997.

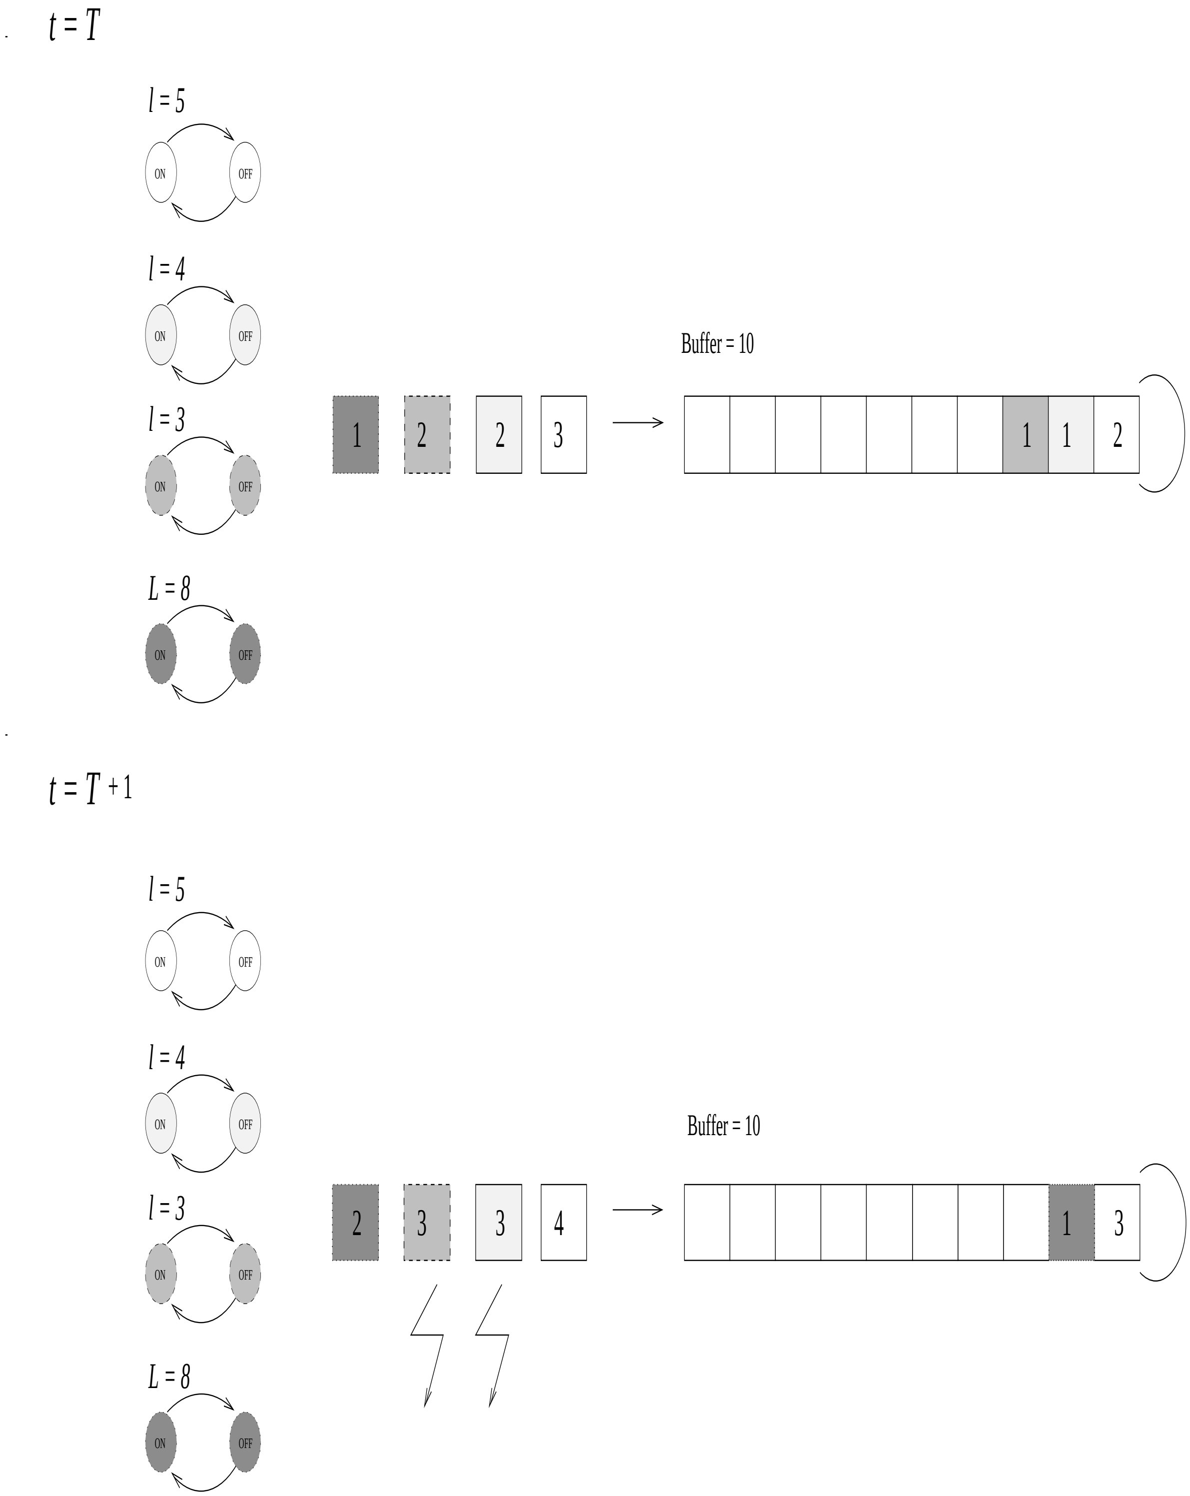

Figure 1: The LPI Packet Discarding Policy. Note that Only Packets with its First Cell in Queue are Eligible to be Dropped.

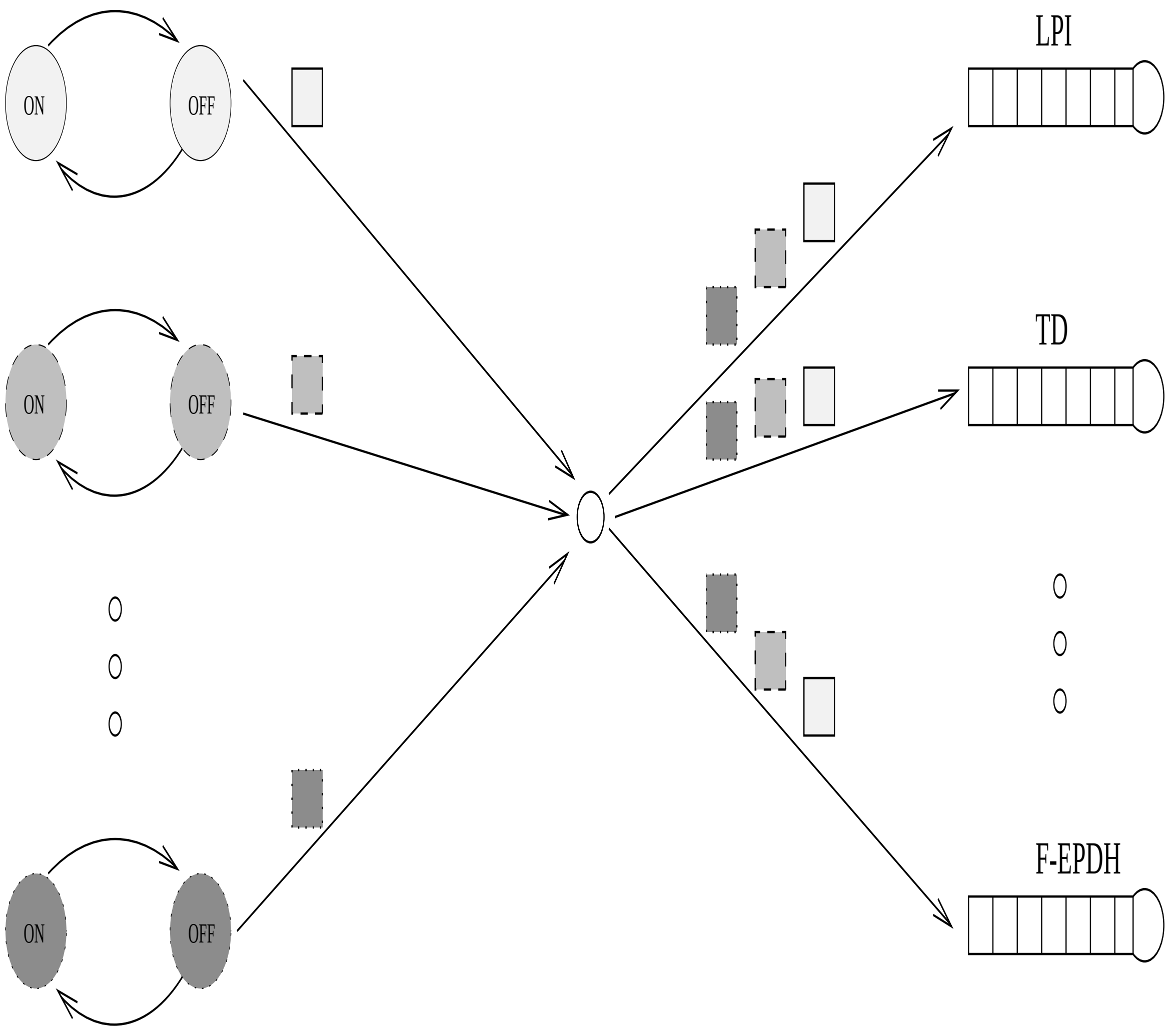

Figure 2: The Simulation Experiments. The Same Stream Produced by Aggregating Several On/Off sources Feeds Different Queues Under Different Discarding Policies.

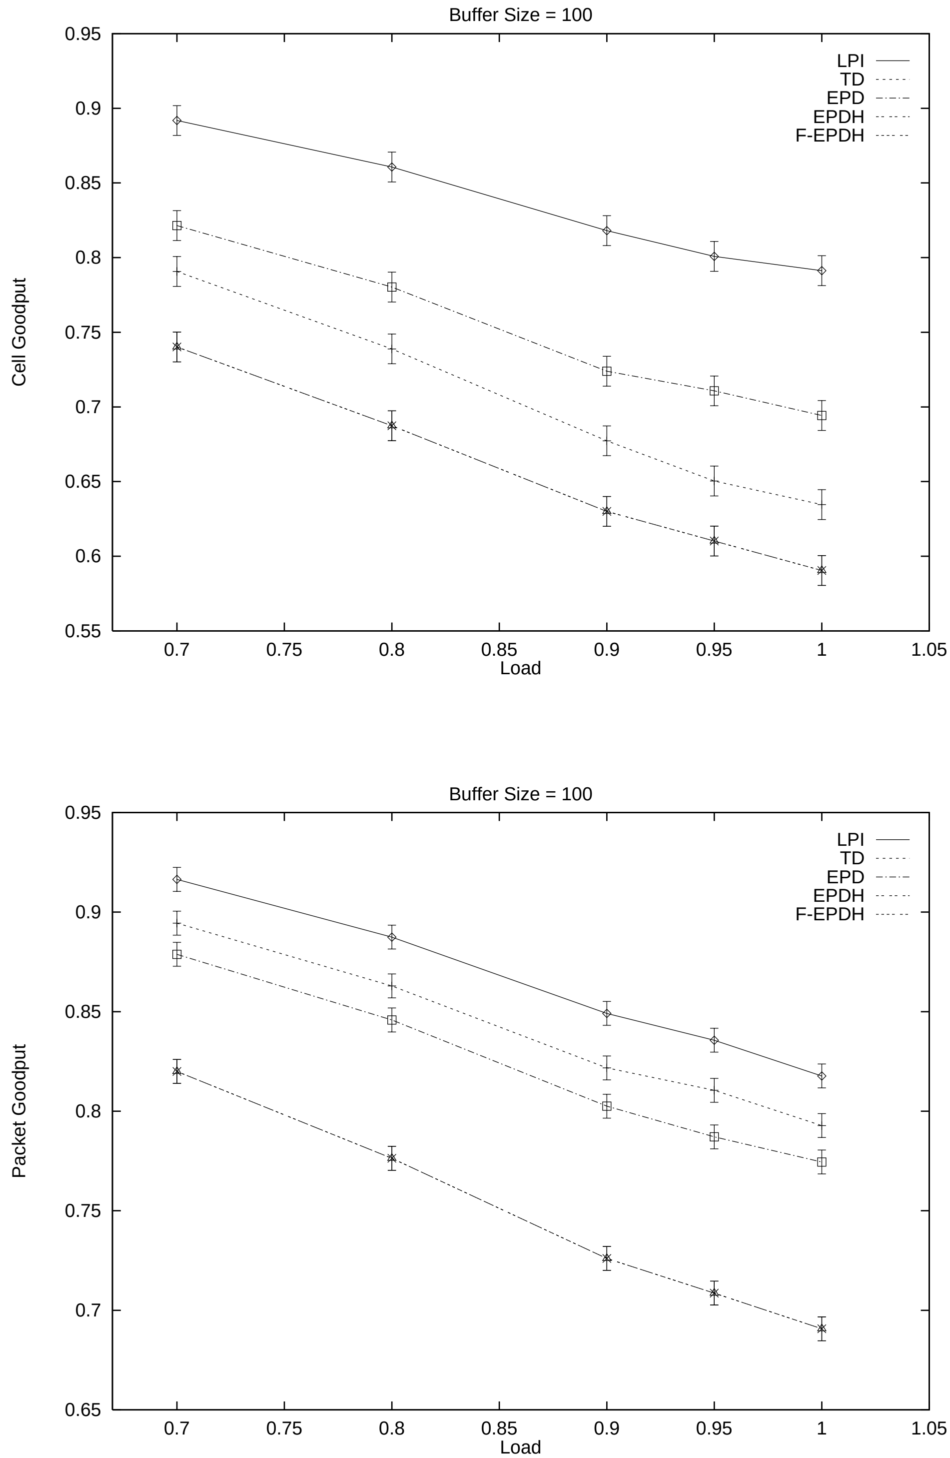

Figure 3: Cell Goodput and Packet Goodput x offered load for mean packet size of 42 ATM cells and buffer size 100.

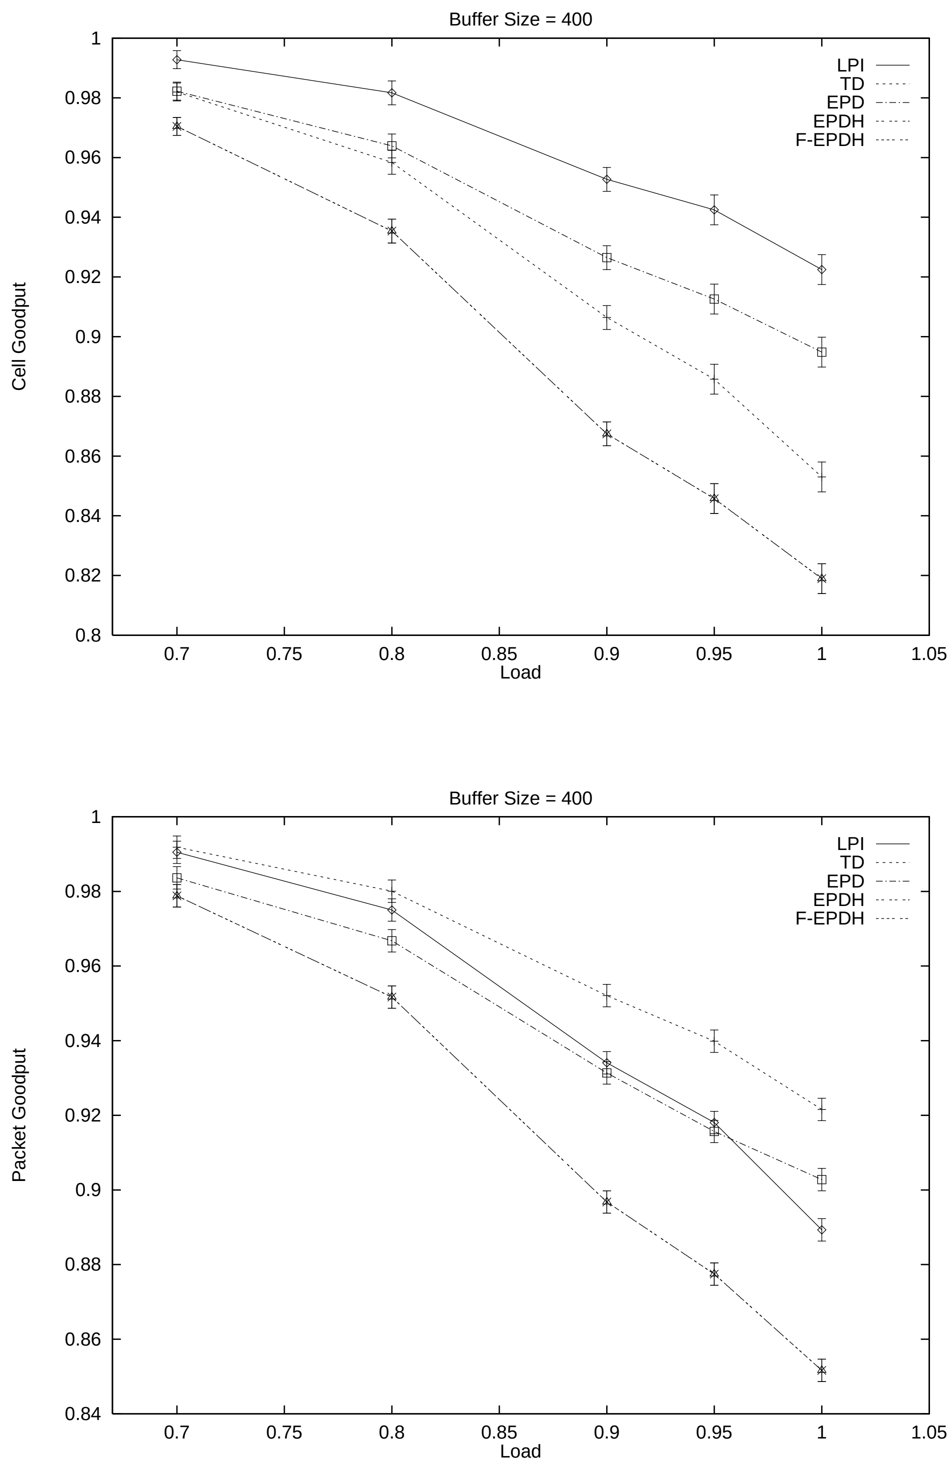

Figure 4: Cell Goodput and Packet Goodput x offered load for mean packet size of 42 ATM cells and buffer size 400.

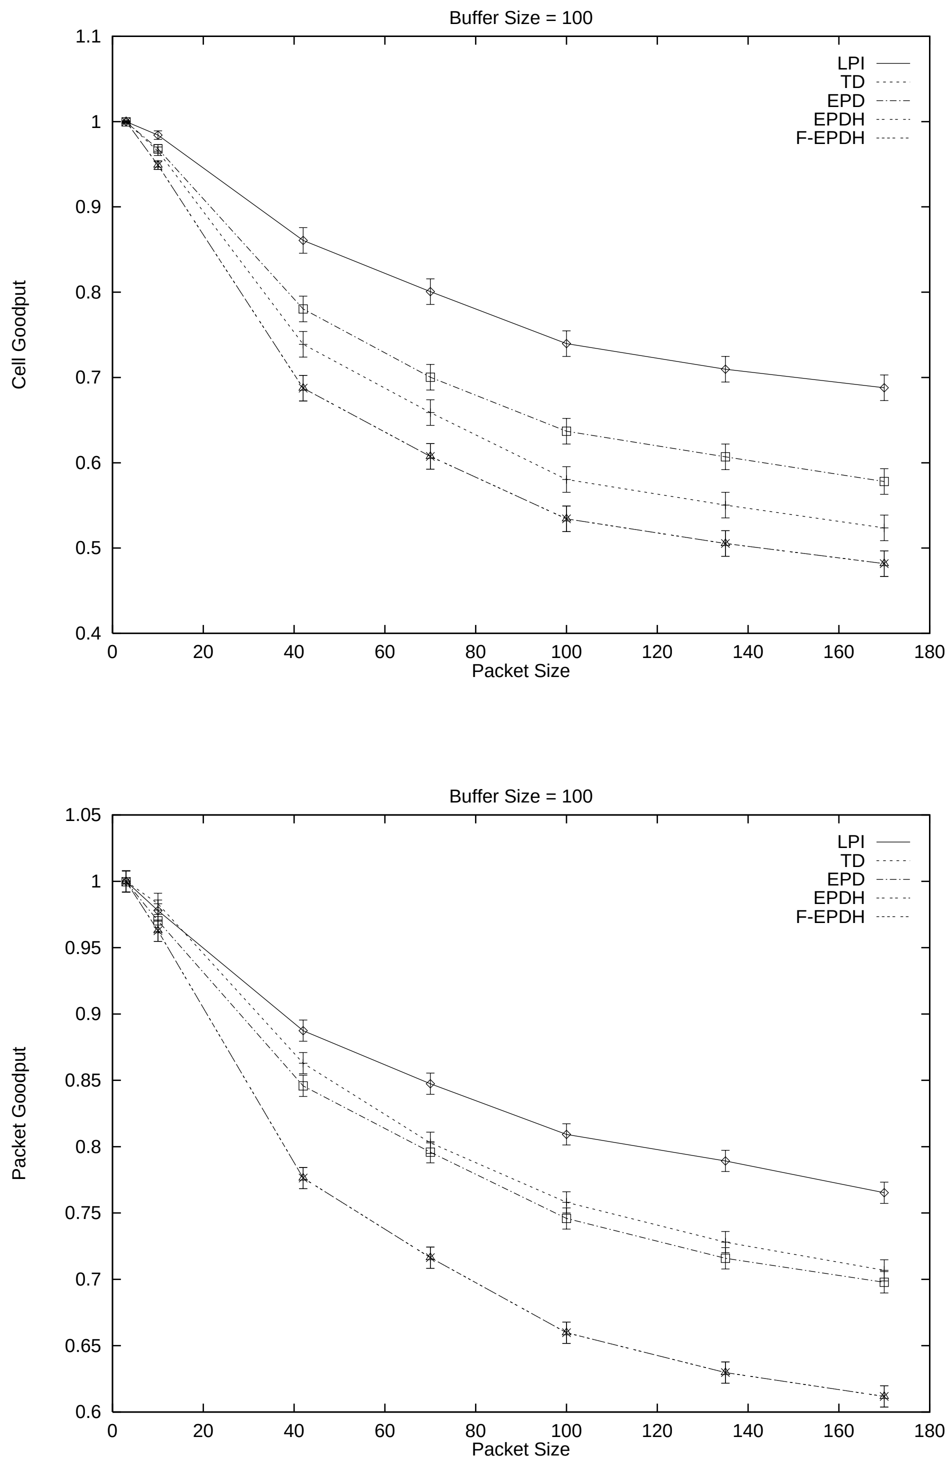

Figure 5: Cell Goodput and Packet Goodput x Mean Packet Size of for an offered load of 0.8 and buffer size 100.

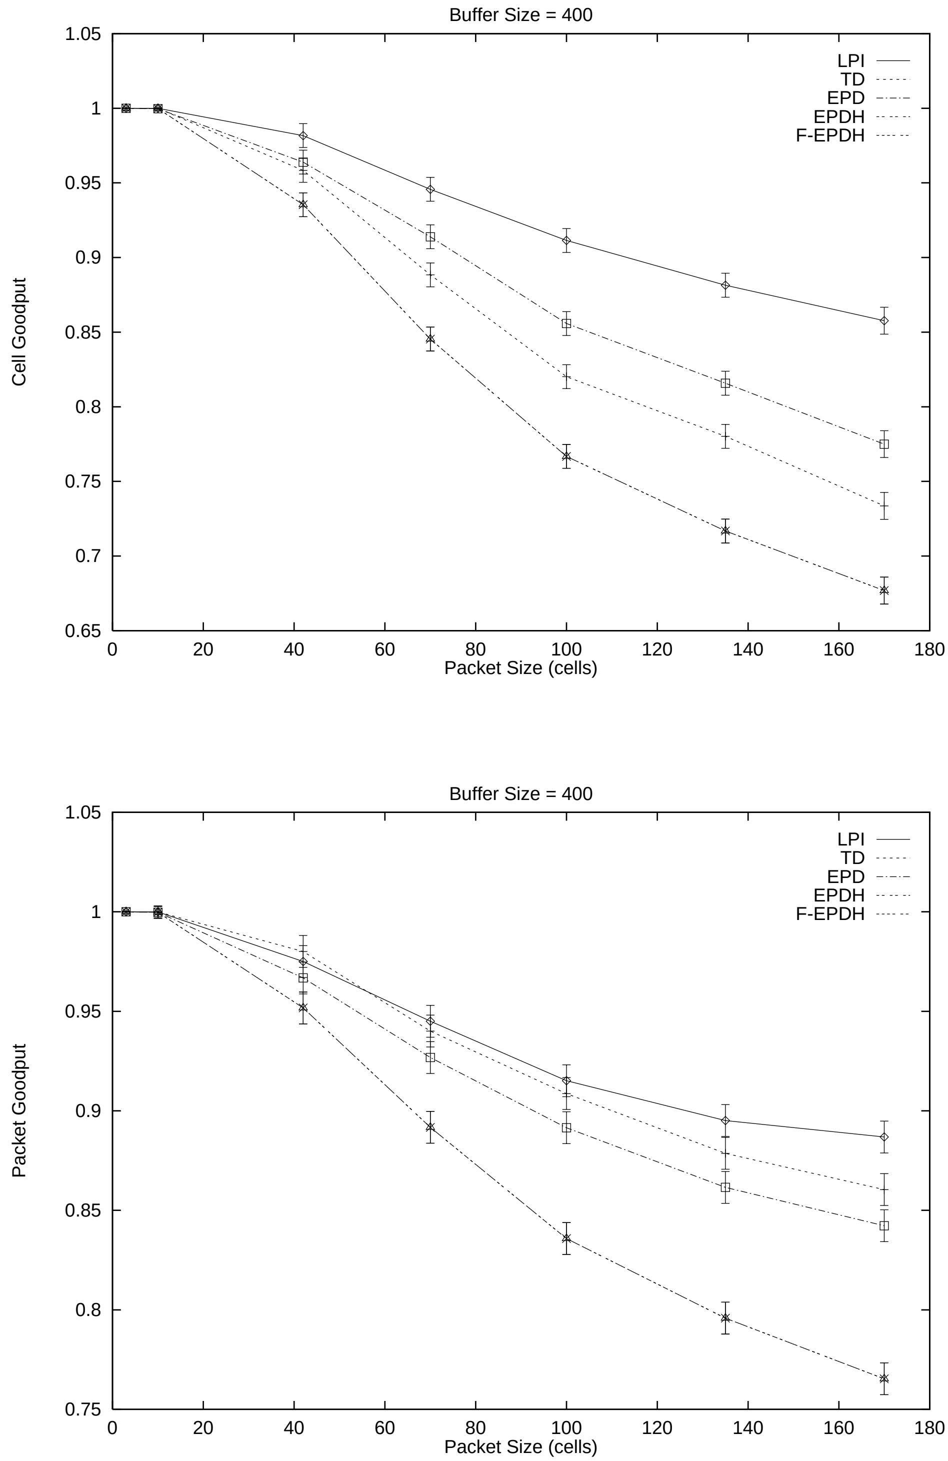

Figure 6: Cell Goodput and Packet Goodput x Mean Packet Size for an offered load of 0.8 and buffer size 400.

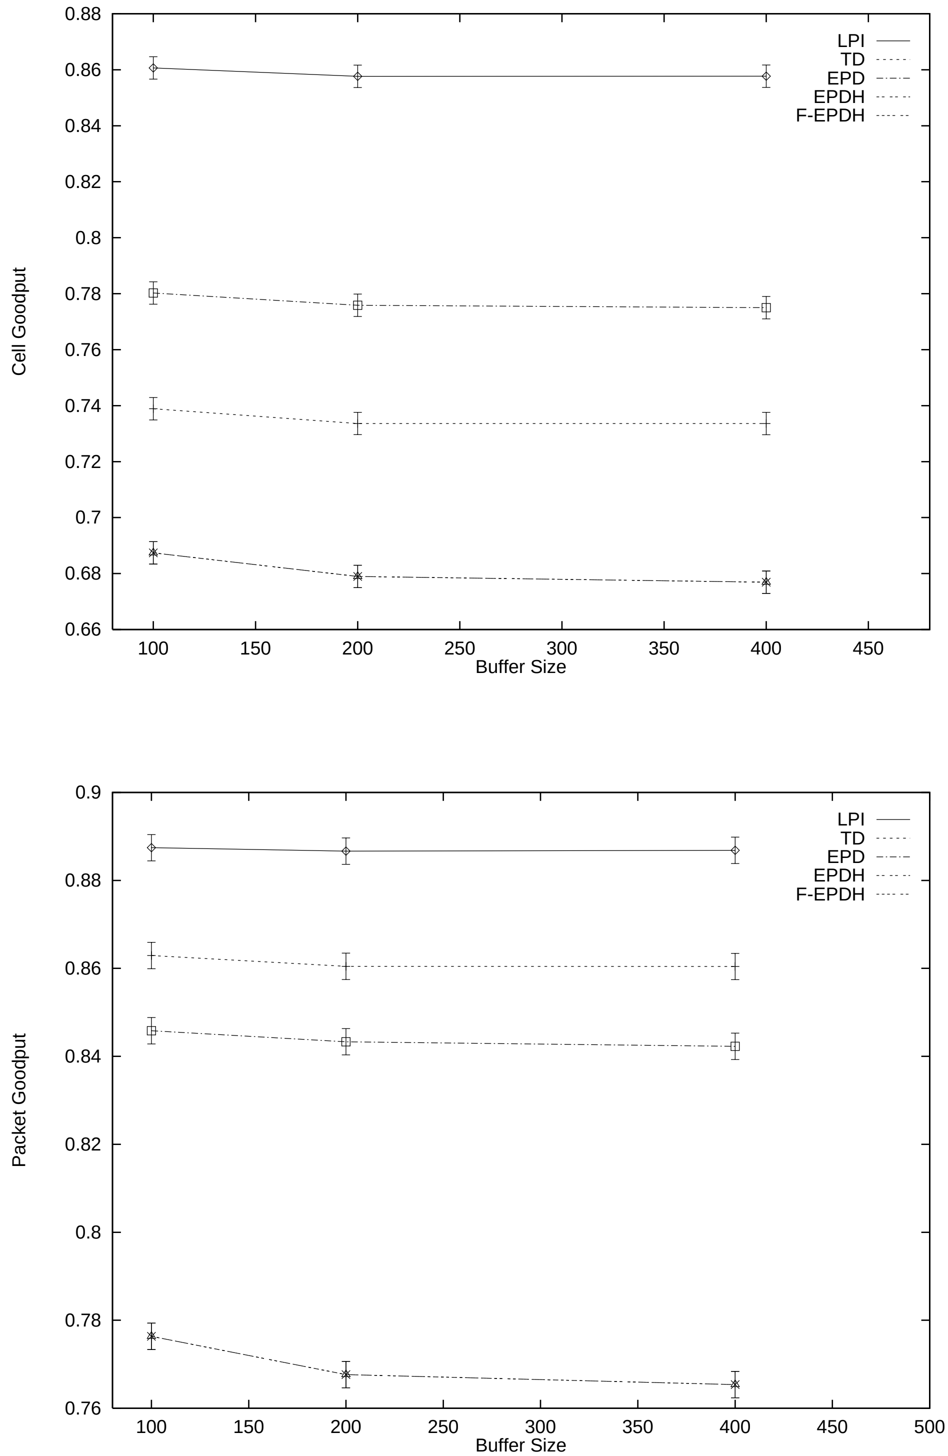

Figure 7: Cell Goodput and Packet Goodput x Buffer Size for a fixed ratio of packet size to buffer size and an offered load of 0.8.

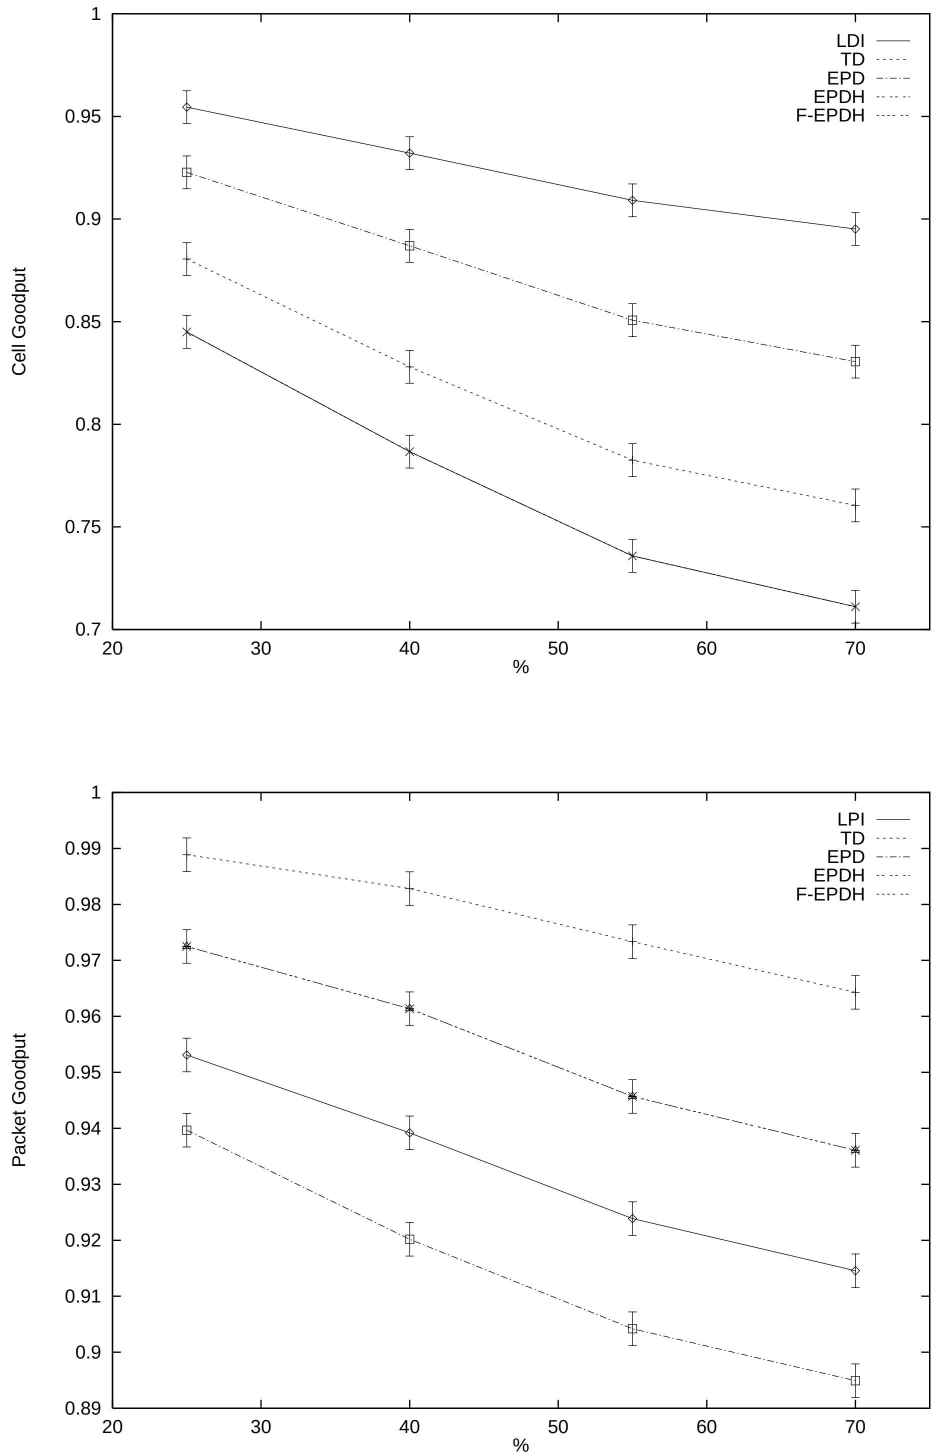

Figure 8: Cell Goodput and Packet Goodput x Proportion of Cells with Mean Size of 170 ATM Cells in a Network with 32 Sources and buffer size 400. Packets may have mean size of 5, 21, 85 and 170 ATM.

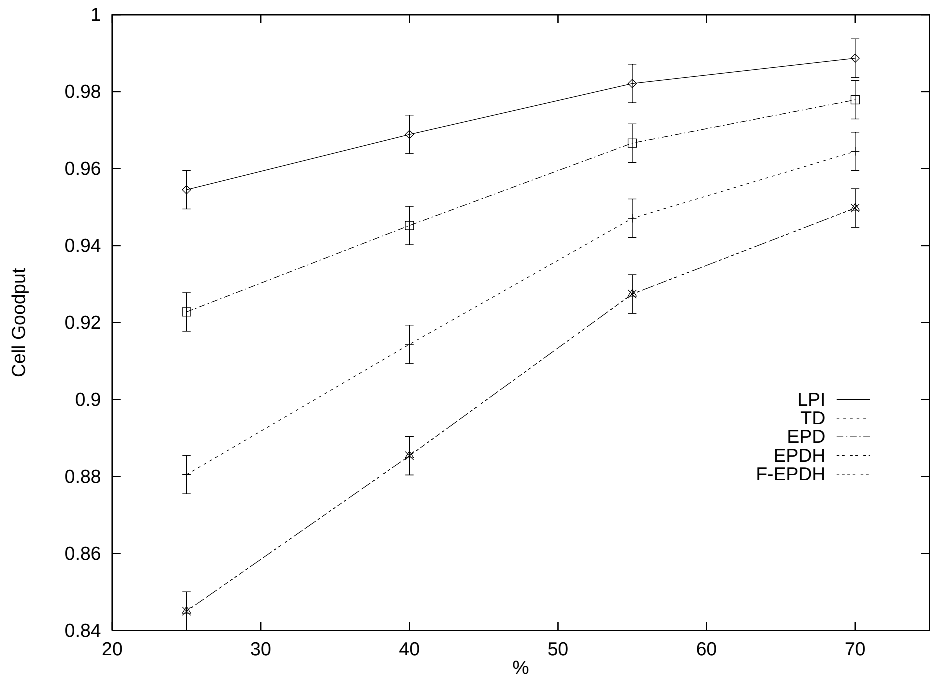

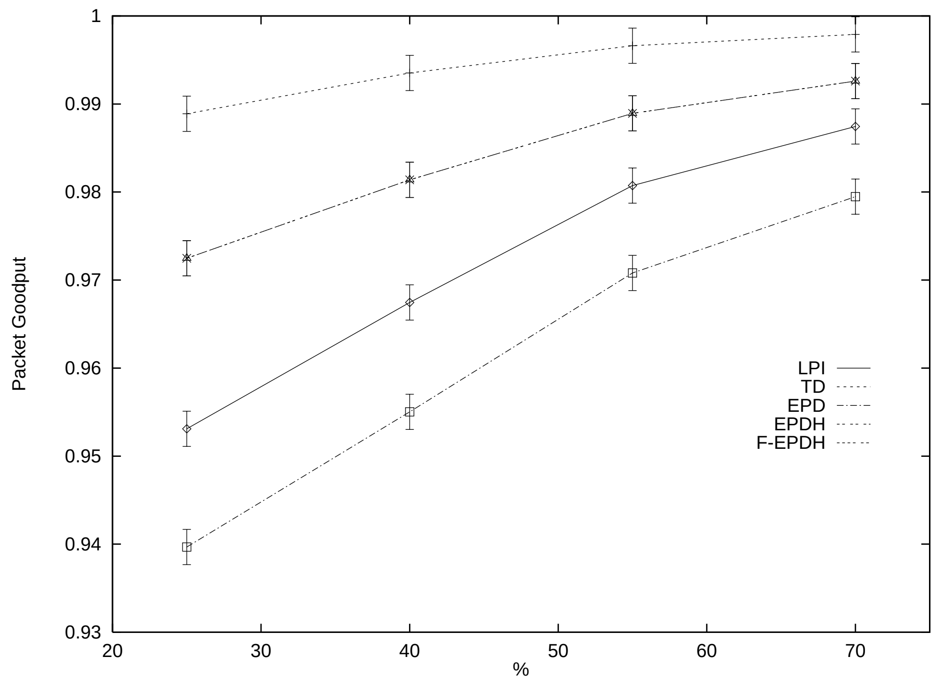

Figure 9: Cell Goodput and Packet Goodput x Proportion of Cells with Mean Size of 5 ATM Cells in a Network with 32 Sources, and buffer size 400
Packets may have mean size of 5, 21, 85 and 170 ATM.

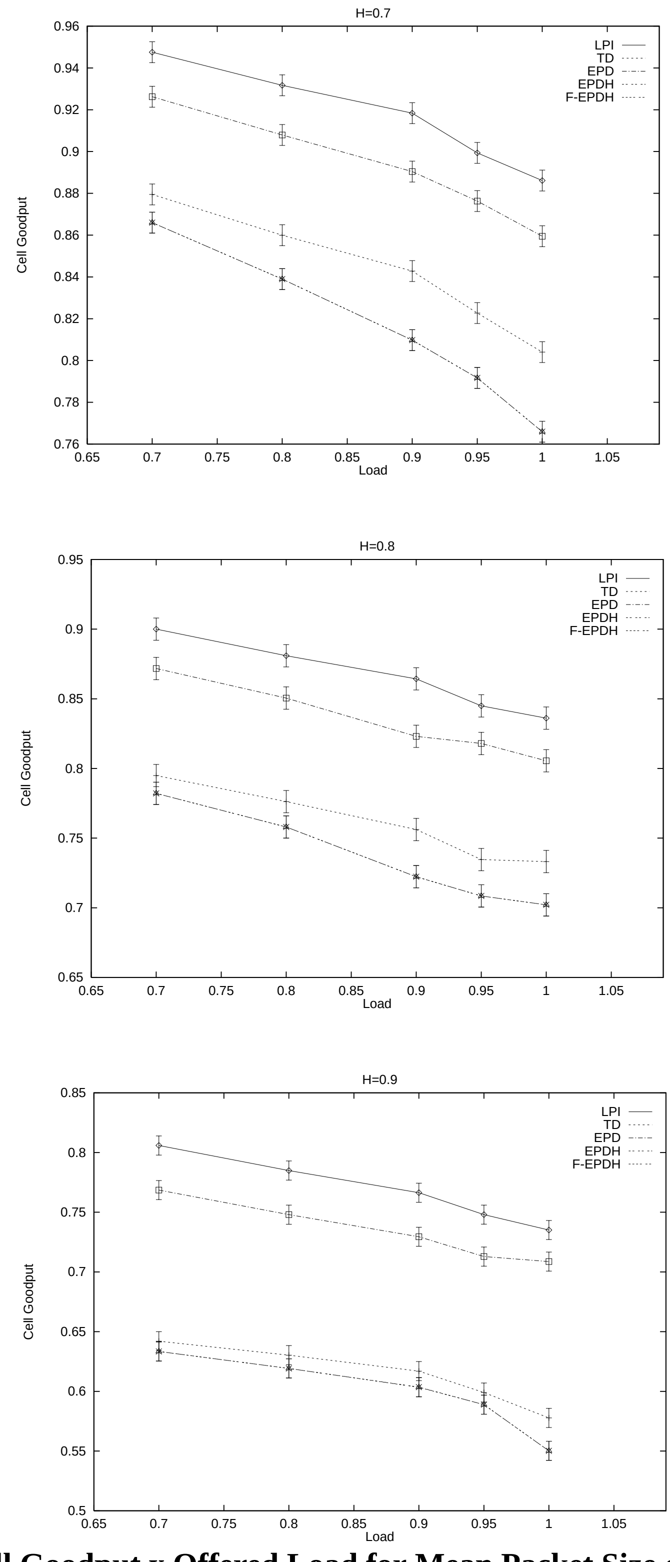

Figure 10: Cell Goodput x Offered Load for Mean Packet Size of 42 ATM Cells Buffer Size 400 and for different values of the Hurst Parameter.

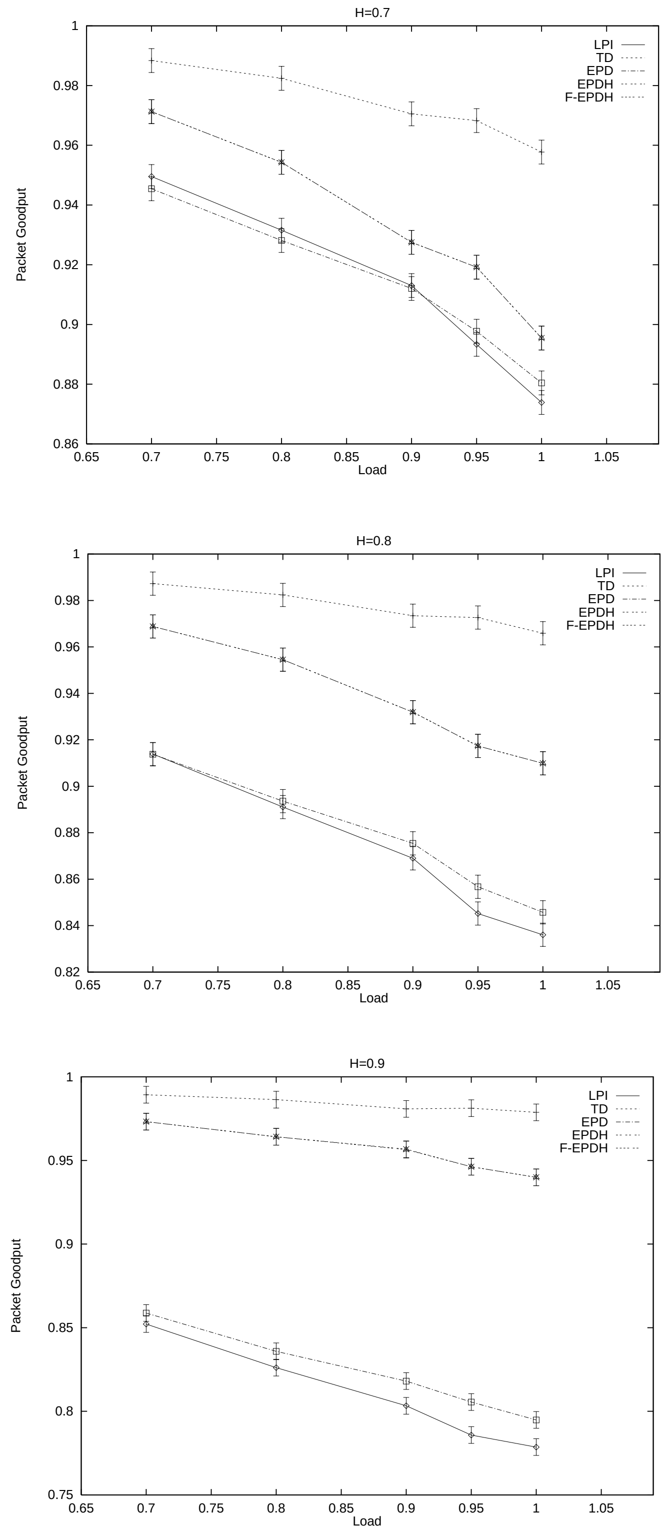

Figure 11: Packet Goodput x Offered Load for Mean Packet Size of 42 ATM Cells Buffer Size 400 and for different values of the Hurst Parameter.

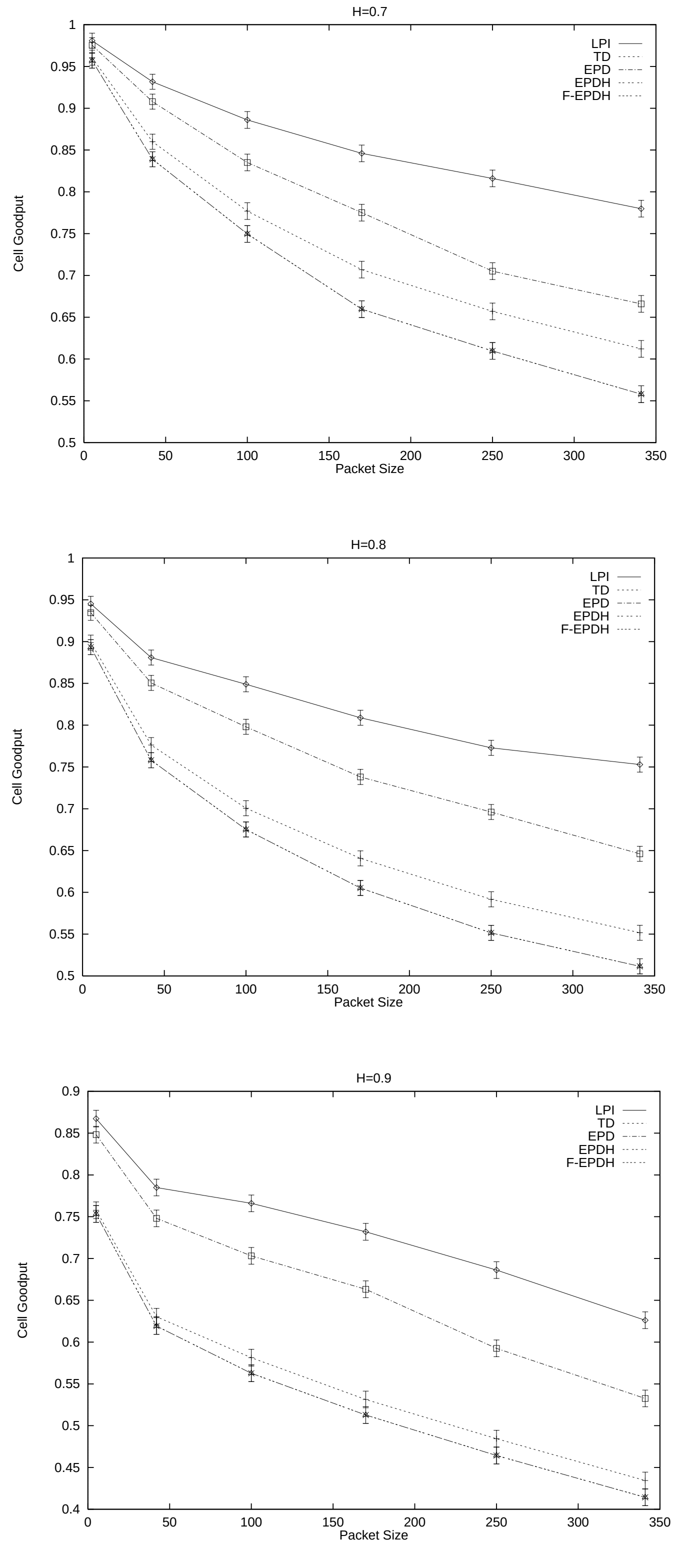

Figure 12: Cell Goodput x Packet Size for an Offered Load of 0.8 and Buffer Size 400 and for Different Values of the Hurst Parameter.

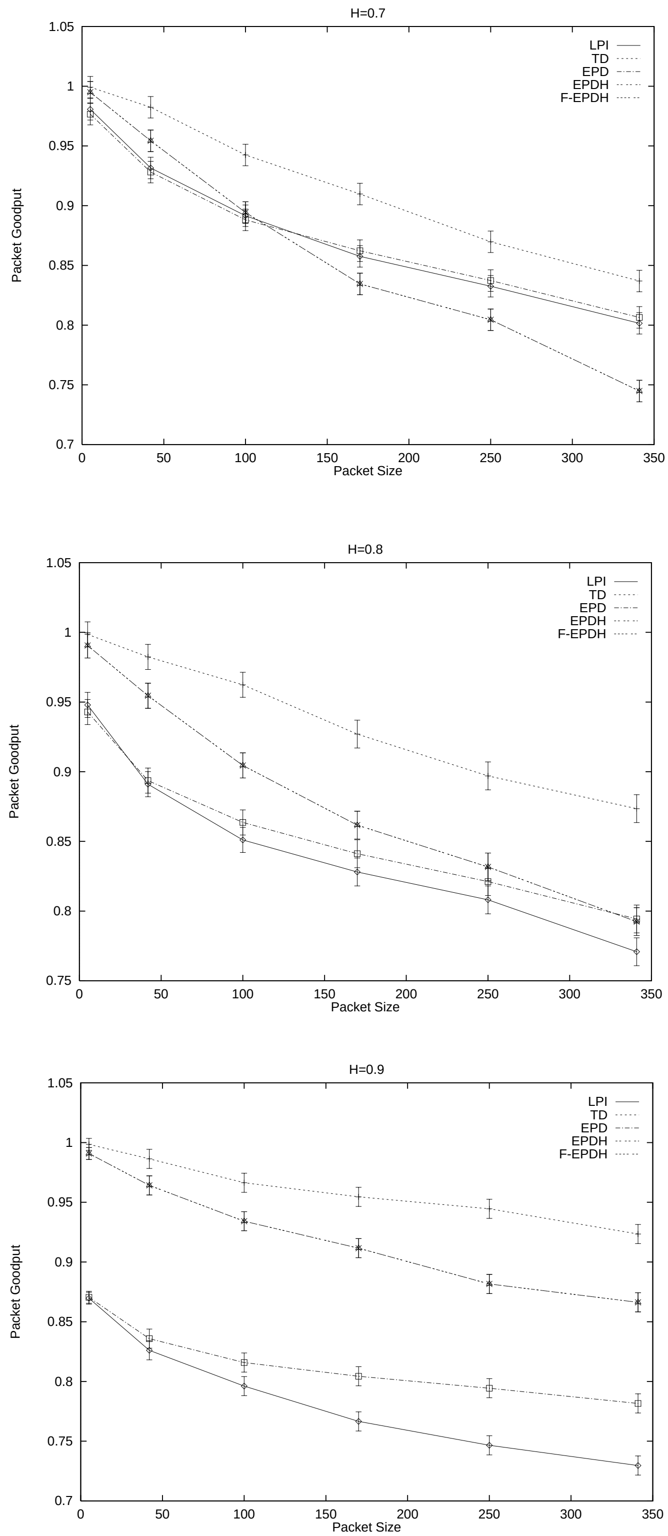

Figure 13: Packet Goodput x Packet Size for an Offered Load of 0.8 and Buffer Size 400 and for Different Values of the Hurst Parameter.

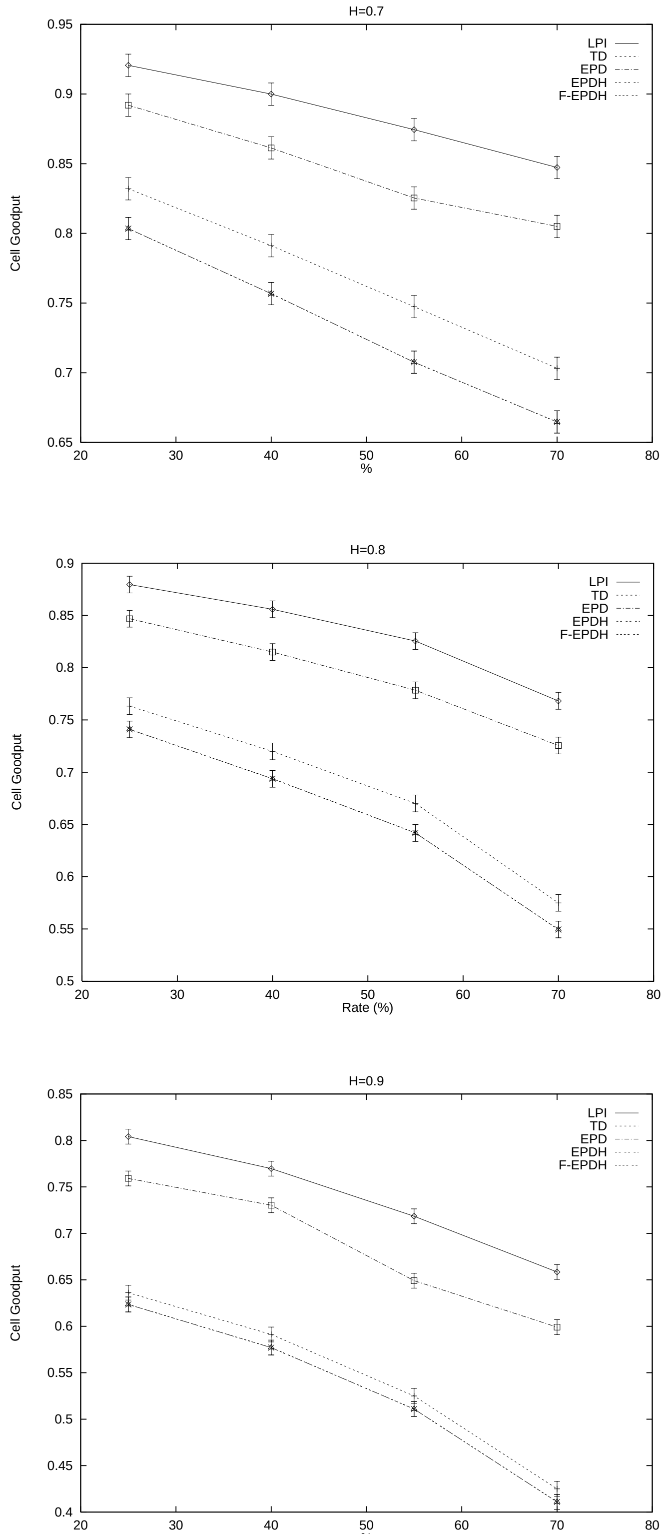

Figure 14: Cell Goodput x Proportion of Cells with Mean Size of 170 ATM Cells for Different Values of the Hurst parameter in a Network with 32 Sources. Packets may have mean size of 5, 21, 85 and 170 ATM.

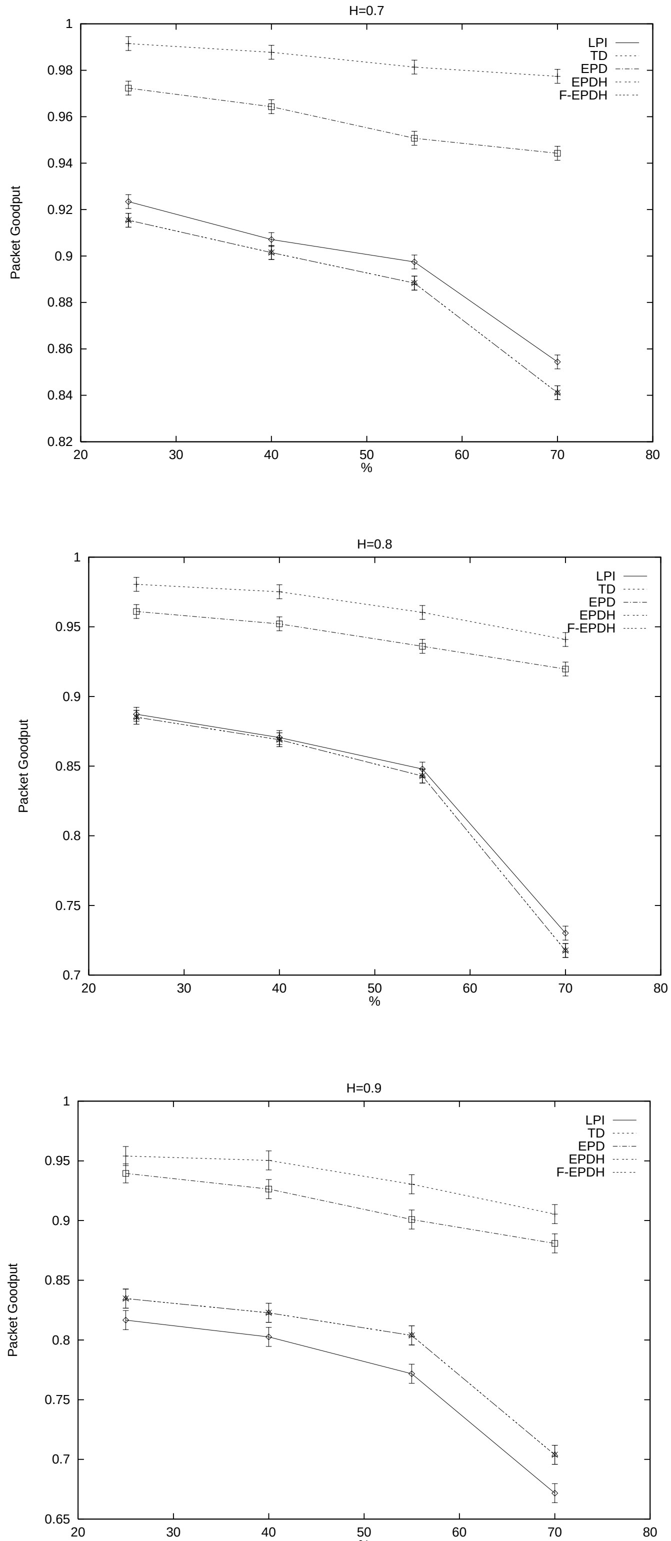

Figure 15: Packet Goodput x Proportion of Cells with Mean Size of 170 ATM Cells for Different Values of the Hurst parameter in a Network with 32 Sources. Packets may have mean size of 5, 21, 85 and 170 ATM.