Title
A simulation study of packet forwarding methods for ATM data services

Authors
Hong, Duke P.
Suda, Tatsuya

Publication Date
1994-10-28

A Simulation Study of Packet Forwarding Methods for ATM Data Services

Duke P. Hong and Tatsuya Suda

Technical Report No. 94-42

28 October 1994

Department of Information and Computer Science
University of California, Irvine, CA 92717-3425
(phone) 714-824-5474 (fax) 714-824-4056 (email) suda@ics.uci.edu

Abstract

It is crucial for ATM to efficiently support existing data applications. One of the promising methods of supporting data over ATM uses connectionless servers providing a packet switching service in the ATM network. There are two primary methods for forwarding cell based datagrams in a connectionless server: reassembly mode, where the datagrams are reconstructed before forwarding, and streaming mode, where each cell is forwarded as it arrives. The two modes impact the packet loss probability at a connectionless server, the packet delay, the interleaving method used, the AAL used, and the buffer space required. This paper investigates the performance of two forwarding modes. Extensive simulations show that forwarding connectionless packets in streaming mode is shown to result in lower packet loss and end-to-end delay than reassembly mode.

This research is supported in part by grants from the University of California MICRO program, Hughes Aircraft, Nippon Steel Information and Communication Systems Inc. (ENICOM), Hitachi Ltd., Hitachi America, Tokyo Electric Power Company, and Hewlett Packard (HP).

1 Introduction

The deployment of Asynchronous Transfer Mode (ATM) networks promises to open opportunities to provide advanced network applications such as teleconferencing, remote learning, multimedia news retrieval, and so forth. ATM has been designed with these applications in mind. This brings up the familiar "chicken-and-egg" dilemma. The ATM network will not develop without the innovative applications to justify the network's expense, but the applications will not be developed without the networks to support them. For ATM to foster in its infancy, it must be capable of supporting the large base of applications that exists today. These applications are mainly data oriented and are currently supported by IEEE 802 LANs. As evidence, much of the current commercial deployment of ATM is providing high speed backbones for LANs. ATM presents a significant departure from current network architectures presenting some challenges in supporting data traffic. In particular, ATM employs a connection-oriented service using small fixed sized packets (53 bytes cells) whereas IEEE 802 LANs employ a connectionless service with larger variable sized packets (or datagrams). Thus, some of the key issues in providing data service over ATM involve the convergence of the connection-oriented service of ATM with the connectionless applications and the management of multi-cell datagrams.

There are a variety of approaches to providing connectionless data service in ATM networks. A simple approach is to use the connection oriented service of ATM to interconnect pairs of ATM end systems, (*i.e.*, either a host machine with an ATM interface or an interworking unit between ATM and another network). For example, Figure 1(a) illustrates an interworking unit (IWU) in one 802 LAN maintaining a connection to another 802 LAN's interworking unit. Note that in order to send packets to a third LAN, a separate connection is required. This approach is called the *indirect* approach since the network only indirectly achieves connectionless service by using the connections already provided by ATM.

There are two types of connections that can be used in the indirect approach. One possibility is to use permanent virtual connections (PVC) between each pair of ATM end systems that require communication. PVCs are particularly useful in situations where traffic between the two endpoints is known to have regular characteristics or where connection establishment overhead cannot be tolerated. The drawback of PVCs is that they may consume network resources even when they are idle. Alternatively, switched virtual connections (SVC) can be used to remove idle connections. An SVC is established when a packet is destined for an ATM end system for which no connection is open. SVCs are terminated after some period of inactivity on the connection. The use of SVCs reduces the waste of network resources but imposes a delay on packet transmission due to connection setup.

While the indirect approach is appropriate for small ATM environments or with a limited number of ATM end systems involved, in order to efficiently support connectionless service in a public wide area ATM network setting, special attention must be paid to scalability so that as the number of end systems grows, network resources are efficiently utilized. While the approaches described above are good in the interim with small ATM networks, the *direct* approach addresses scalability. The direct approach [1] [2], is based on the deployment of connectionless servers (CLS) throughout the network, forming a virtual connectionless network overlayed on the connection-oriented ATM network. Connectionless servers act as packet switches, routing packets to other connectionless servers until the destination is reached. Connectionless servers are strategically attached to a subset of switches within the network. The topology of the virtual connectionless network may take any number of forms including ring, mesh, and star depending on the desired efficiency and robustness [3]. For example, Figure 1(b) shows how LANs can be interconnected using connectionless servers. With the direct approach, an ATM end systems needs only one open connection to a connectionless server in order to send packets to one or more destination ATM end systems. This method is called the direct approach because the ATM network provides the connectionless service directly to the user; the user is responsible only for sending encapsulated connectionless datagrams into the virtual connectionless network.

The direct approach has several advantages over the indirect approach in a wide area environment. It is inherently scalable. Fewer connections, and thereby fewer network resources, are needed in the direct approach. This also has the effect of multiplexing data traffic from many sources on fewer connections potentially making connection traffic smoother and more predictable. The direct approach relieves the ATM end systems of responsibility for making routing decisions by forcing that responsibility onto connectionless servers. This simplifies the use of connectionless service for the user since little or no special hardware or software support on the user's side is required. Furthermore, this method requires only one connection into the virtual connectionless network for each ATM end system, as opposed to the numerous connections required by the indirect approach. This may aid the end system in terms of resource usage and connection setup overhead.

The direct approach, however, does have its disadvantages. The additional processing time required at connectionless servers increases the end-to-end delay in datagram transmission. The delay is further increased with respect to the indirect approach because packets must be routed through connectionless servers rather than taking the shortest path between the end systems. Congestion at connectionless servers may also increase the end-to-end packet loss probability.

The impact of these disadvantages can be reduced with some forethought. The packet forwarding technique employed at connectionless servers can largely influence the delay. Packet loss due to congestion at connectionless servers and network congestion can be reduced by utilizing some form of bandwidth management on the inter-CLS connections. For instance, the inter-CLS connections may reserve some amount of bandwidth and adjust it slowly as demand changes [4]. Alternatively, unused bandwidth may be borrowed from other inactive connections more dynamically than in the previous approach [5]. Another method would be to use feedback to adjust the output rate dynamically [6].

This paper investigates the performance of connectionless servers, comparing two packet forwarding techniques. The major issues surrounding the architecture and performance of connectionless servers are discussed in Section 2. The simulation model and results of our comparison of forwarding modes are presented in Section 3. Finally, observations are summarized in section 4.

2 Connectionless Server Performance Factors

Connectionless servers form a virtual packet switched network within the ATM connection oriented network. The connectionless service functionality can be added externally to a subset of the ATM switches as shown in Figure 2 or it can be incorporated into the switches. While the former configuration can create some additional switch congestion because packets must be routed through the switch twice (before and after the CLS), it provides greater flexibility in adding, removing and modifying connectionless service functions. As shown in the figure, there may be many connections leading into and out of the CLS. Each connection, implemented using virtual circuits, leads to another connectionless server or to and ATM end system. For brevity, we will refer to the connections leading into and out of the CLS as "*ports*". With N incoming and outgoing ports, the CLS functions as an $N \times N$ packet switch.

The transmission of datagrams through connectionless servers requires a framing architecture (*e.g.*, Connectionless Interworking Protocol (CLIP) [7]) that absorbs the differences in packet format and size differences. Connectionless servers operate above and within the AAL layer receiving ATM cells and forwarding them based on network layer addresses, such as IP, NSAP, or E.164, of the destination ATM end systems. The cell headers are insufficient for this task since the VPI/VCI only identifies the connections between connectionless servers. Because of the small cell size, placing an ATM end system address into every cell results in very low protocol efficiency. A more efficient approach is to segment the larger datagrams into several cells, placing the ATM end system address into the first cell of a packet (the *beginning-of-message* or BOM cell), and associating the cells together when the datagram is forwarded [8]. It was shown previously that the CLIP header can fit in one cell payload [7]. An encapsulated IP header can fit in the first cell of a segmented datagram as well. Since not all cells contain the destination address, the cells belonging to the same datagram must follow the same path. Thus, the primary task of a connectionless server is to route datagrams based on the destination addresses contained in the BOM cells and ensure that the remaining cells of the datagrams follow the same path as their respective BOM cells.

This remainder of this section will describe aspects of a connectionless server architecture that affect its performance. These are: the interleaving scheme, the forwarding mode, and flow control. Other factors which affect the effectiveness of connectionless servers are the routing algorithm and the network topology used. As these issues have been extensively investigated before, they will not be discussed further.

Interleaving scheme

The interleaving scheme is the method by which cells from two or more packets are multiplexed in a single connection. In packet interleaving, all the cells of one packet are sent contiguously. In cell interleaving, cells from different packets can intermingle. There are compatibility and performance issues raised by the choice in interleaving scheme.

With packet interleaving, packets can be distinguished from each other by defining the packet boundaries using the PTI field in the ATM cell header. With cell interleaving, cells from different packets can only be distinguished by using a multiplexing identifier within the 48-byte payload of the ATM cell. Thus, packet interleaving is compatible with AAL 5 while cell interleaving requires AAL 3/4 which specifies a 10 bit multiplexing identifier within the 48 byte cell payload. As the data network community is moving towards AAL 5, compatibility may become an issue for any scheme using AAL 3/4. There are two methods to resolve the interleaving and compatibility problem. The first is to reassemble the higher layer PDUs from one AAL, then convert and resegment the PDU using the desired AAL. This introduces some delay but otherwise allows

AAL compatibility. The second option is to use a virtual path (VP) as the connection. This allows AAL 5 packets to be sent cell interleaved on the virtual path with different packets identified by the virtual circuit identifiers (VCI).

The most significant issue related to the performance of the two interleaving schemes is how they are affected by consecutive cell loss. It has been documented that ATM exhibits consecutive loss characteristics in which the time period following a cell loss has a high probability of having another cell loss occur [9]. With packet interleaving, a lossy period is likely to affect many cells within the same packet but with cell interleaving, a lossy period can affect cells from many packets. Thus, it is possible that for the same cell loss rate, a packet interleaved scheme will show a smaller packet loss rate than a cell interleaved scheme.

Forwarding mode

The method by which cell segmented datagrams are forwarded can have a significant impact on the performance of a connectionless server. There are two primary methods of forwarding packets at a connectionless server: packet based forwarding and cell based forwarding [10] [11].

In packet based forwarding, incoming cells are reassembled into packets at each connectionless server. The outgoing connection to the next connectionless server is determined for the packet, and the packet is (re)segmented into cells which are delivered to the next hop using the appropriate connection and packet identifiers. This mode of operation is referred to as *reassembly mode* operation. A reassembly mode connectionless server can use either cell or packet interleaving. Since packet interleaving fares better under consecutive loss, we will assume that reassembly mode connectionless servers will use packet interleaving.

In cell based forwarding, packets are not reassembled at connectionless servers. Instead, the packet's constituent cells are forwarded one at a time, as they are received, in a cut-through manner to the next hop. When the first cell of a packet arrives, the routing algorithm uses the destination ATM end system address stored in the BOM to determine the next hop connectionless server or interworking unit. This routing information is stored in a forwarding table so that the connectionless server can simply perform a table lookup and determine the next hop for each of the packet's remaining cells as they arrive. This scheme ensures that cells of the same packet are forwarded on the same routes while allowing cells to be delivered before the entire packet has been received. This mode of operation is referred to as *streaming mode* operation. Since a streaming mode connectionless server must forward cells from packets from two or more input ports on shared output ports, it must use cell interleaving.

The type of forwarding mode used affects the packet delay and loss probability. The reassembly mode introduces a packet reconstruction delay that is not seen using streaming mode connectionless servers. Thus, the end-to-end packet delay increases linearly by at least the packet reconstruction time for each reassembly mode connectionless server. The end-to-end delay using streaming mode will also increase with the number of hops traversed, but may not increase linearly. The interleaving of cells from many packets may increase the transmission times between the BOM and EOM (*end-of-message*) cells of a packet. As more hops are traversed, the delay may increase slower as the multiplexing of packets from incoming ports onto outgoing ports is preceeded by the demultiplexing of packets from incoming ports to outgoing ports.

Packet loss occurs at a connectionless server when the output buffer overflows. Streaming mode servers use buffer space for output rate control to maintain flow control between connectionless servers. The reassembly mode connectionless server uses buffer space for packet reconstruction in addition to output rate control. Therefore, for the same amount of buffer space, buffer utilization is greater for reassembly mode servers than for streaming mode servers, increasing their cell loss rate due to buffer overflow. However, reassembly mode servers have better discarding characteristics than streaming mode. When a reassembly mode connectionless server must drop a cell due to shortage of buffer space or detects a missing cell, the entire packet to which the cell belonged can be discarded. On the other hand, streaming mode servers can only discard the cells arriving after the loss. Thus, useless traffic needlessly will remain in the output buffer and fill connectionless server buffers downstream.

Another consideration is the effect of the interleaving scheme on the effectiveness of packet discarding. Suppose cell i in an L cell packet is dropped due to buffer overflow. The remaining $k = L - i - 1$ cells can be discarded. With packet interleaved input, the next k cells that arrive on that connection can be discarded. This is particularly of importance because the buffer occupancy was at 100% at the time loss occured. The ability to discard the next k arriving cells reduces demand for the buffer space. With cell interleaved input, this reduction in demand is not as great. Although k cells arriving on that connection can be discarded, they will arrive interleaved with other cells that cannot be discarded because they belong to packets which have not yet been damaged. Thus, if m packets were interleaved, then only 1 out of m cells for the next km cells on that connection can be discarded. This maintains high demand for buffer space which could lead to additional packet loss. For these reasons, reassembly mode may show lower overall packet loss under certain circumstances, while streaming mode may show lower packet loss in others.

Flow Control

In virtually all ATM environments, some form of rate control is necessary to prevent congestion and to enforce flow control between connectionless servers. Suppose a connectionless server is capable of processing p cells per second and there are N incoming connections to the server. In order to avoid overloading the server, the cell arrival rate of the incoming ports should not exceed p/N cells per second on average. Assuming all connections have identical traffic loads, this restriction implies that a rate control permitting the transmission of only one cell per N/p seconds be implemented for the connections leading to the connectionless server. There are two classes of rate control devices: closed loop and open loop.

A closed loop rate control device uses feedback to dynamically adjust the output rate. An example of this is the proportional rate control algorithm (PRCA) being adopted by the ATM Forum. In PRCA, the destination sends a positive feedback cell for every m cells received. Congested switches along the path discard these feedback cells and the sources adjust their output rates according to the rates at which their feedback cells are received. This scheme is amenable to local and wide area environments although it exhibits heavy oscillatory behavior in the wide area [6].

An open loop rate control device does not use any feedback from the network. Examples of this type include the well known leaky bucket and constant rate control. The constant rate control simply ensures a minimum cell interdeparture time. This type of rate control is suitable in wide area environments where the propagation delay for feedback messages are too long to be effective.

3 Simulations

Simulation Model

The network simulation model is composed of two $N \times N$ connectionless servers as shown in Figure 3. The first server, *CLS-1*, is fed by N traffic generators, each of which is assumed to maintain a PVC (or a long-lived SVC) to the server. In other words, connection setup is not considered in the simulations. The traffic generators simulate other connectionless servers or interworking units between ATM and 802 LANs. Of the N output PVCs from *CLS-1*, one connection that is observed end-to-end is fed through the network where it experiences network cell loss and delay at ATM switches before reaching *CLS-2*. The remaining $N - 1$ outgoing connections are assumed to be destined for other connectionless servers or interworking units not shown in the figure. *CLS-2* also receives traffic from $N - 1$ additional traffic generators. Using this model, the packet delays and loss probabilities have been studied for both single hop (*IWU* to *CLS-1*) and 2 hop (*IWU* to *CLS-1* to *CLS-2* to *IWU*) scenarios. Both streaming and reassembly mode connectionless servers have been simulated.

The traffic generators shown in Figure 3 can generate either packet interleaved (PI) or cell interleaved (CI) packet streams simulating interworking units and connectionless servers. Interworking units are assumed to generate packet interleaved output. Reassembly mode connectionless servers generate packet interleaved output. Thus, to simulate a network of reassembly mode servers, all input to the connectionless servers are packet interleaved. However, since streaming mode connectionless servers generate cell interleaved output, it is necessary to investigate scenarios where they receive packet and cell interleaved input. In the first scenario, all inputs are packet interleaved, with the exception of the *CLS-1* to *CLS-2* connection, to simulate connectionless servers at the edge of the virtual connectionless network receiving traffic from interworking units. In the second scenario, all inputs are cell interleaved. This simulates connectionless servers deeper within the virtual connectionless network since the cell interleaved inputs are similar to those produced by other streaming mode connectionless servers.

A simple traffic arrival process is chosen for the simulations; for each connection, the cell arrival process is similar to IPP where the ON states represents packets. The packet sizes are randomly generated using a truncated geometric distribution with minimum size L_{min}, maximum size L_{max}, and average size L cells. To simulate random packet generation by interworking units, the time between the arrival of the last cell of one packet and the first cell of the next packet at the connectionless servers for each connection is assumed to follow a geometric distribution with mean $1/\lambda_F$ cell slots. Similarly, the cell interarrival time on a connection is drawn from a geometric distribution with mean $1/\lambda_C$. Since cell interleaved input simulates the output from an $N \times N$ upstream streaming mode connectionless server, up to N packets can be interleaved at any time.

A shared buffer of size B_{CLS} is employed for both streaming and reassembly mode connectionless servers. The destinations of the incoming packets are assumed to be uniformly distributed among the N output ports, at both *CLS-1* and *CLS-2*. The load applied to *CLS-2* is set to be equiv-

Parameter	Value	Description
N	4	Number of input/output ports at CLS
$1/\mu_{Co}$	N cell slots	Minimum cell interdeparture time at output ports
$1/\lambda_C$	N cell slots	Average cell interarrival time from IWUs
$1/\lambda_F$	*variable*	Average interpacket delay from IWUs
L	15 cells	Mean packet size
L_{min}	2 cells	Minimum packet size
L_{max}	210 cells	Maximum packet size
B_{CLS}	250/500/1000 cells	Buffer size at CLS
l_{net}	10^{-5}	Network cell loss probability

Table 1: Simulation Parameters

alent to the load applied to *CLS-1*. Each output port is regulated using a simple open loop constant rate control scheme which restricts the cell interdeparture times to a minimum of $1/\mu_{Co}$. Such rate control prevents traffic overload at the next hop connectionless server.

In addition to cell loss and delay within connectionless servers, there may be cell loss and delay within the network. The virtual connections leading to *CLS-2* are affected by both network cell loss and delay. The network cell loss is assumed to be random and independent with probability l_{net}. The network delay on the connection between *CLS-1* and *CLS-2* is simulated by a single queue with an interfering Poisson cell stream with a cell interarrival time of 1.8 cell slots and a constant service time of 1 cell slot. Table 1 summarizes the simulation parameters. The confidence intervals show the range where 8 of 10 simulation runs resulted.

In most of the simulation results that follow, the output load on each CLS connection is the parameter against which connectionless server characteristics are measured. The load on each CLS output port is defined as the fraction of time an output port is busy with the transmission of a cell or rate control countdown until it is permitted to send the next cell. It therefore equals the number of cells that arrive per unit time (λ) divided by the maximum number of cells that can depart per unit time (μ_{Co}).

$$\rho = \frac{\lambda}{\mu_{Co}} \tag{1}$$

The overall cell arrival rate λ is obtained from the average interpacket delay ($1/\lambda_F$) and average per-packet cell interarrival time ($1/\lambda_C$) as follows:

$$\lambda = \frac{L}{\frac{L-1}{\lambda_C} + \frac{1}{\lambda_F}} \tag{2}$$

Results

Several performance measures were obtained to evaluate the two forwarding modes. The first performance measure is packet loss probability for a single CLS hop shown in Figure 4. For small buffers and light loads, it can be seen that streaming mode shows a substantially lower packet loss probability than reassembly mode. With $\rho = 0.63$, the packet loss probability for streaming mode with packet interleaved input is 2×10^{-5} compared to 1×10^{-3} for reassembly mode. However, with cell interleaved input, the streaming mode showed a much higher packet loss probability of 2×10^{-4}. For larger buffers and higher loads, the difference between reassembly and streaming modes diminish. For $B_{CLS} = 1000$, the packet loss probabilities between the two modes with packet interleaved input is nearly identical. Additionally, with cell interleaved input, the packet loss probabilities for streaming mode is higher than for packet interleaved input reassembly mode with $B_{CLS} \geq 500$ for all loads.

In reassembly mode servers, packets reside in the buffer for longer periods of time as they are reassembled and then buffered for output rate control. This longer buffer occupancy time results in a greater propensity for buffer overflow as other packets arrive, especially for small buffers. Streaming mode servers exhibit a smaller packet loss probability because cells flowing through them are buffered only for rate control purposes. However, with cell interleaved input, packet loss probability with streaming mode is significantly higher as suspected from the discussion in Section 2. The poor discarding characteristics of cell interleaved inputs was shown to have a considerable impact on packet loss probability, increasing the loss probabilities by as much as a factor of 10.

Figure 5 illustrates the end-to-end packet loss probability of packets that traverse both *CLS-1* and *CLS-2*. This figure shows that although the packet loss rates increase substantially from the single CLS case, the same relationships remain between the loss probabilities for the various forwarding modes as was shown in the previous figure: streaming mode with packet interleaved input has a lower packet loss probability than reassembly mode, but cell interleaved input increases packet loss probability significantly for streaming mode connectionless servers. This observation implies the following conclusion. Recall that reassembly mode servers are capable of discarding entire packets when a single cell is lost, whereas streaming mode servers only discard the remainder of a packet. The packet discarding with reassembly mode would thereby reduce the load on downstream connectionless servers more than streaming mode servers would. Figure 5 shows that this downstream load reduction effect is not significant enough to noticeably improve the end-to-end packet loss probability of reassembly mode in comparison to streaming mode. The single hop packet loss probability is clearly more important in determining end-to-end packet loss than the beneficial discarding characteristics of reassembly mode servers. Thus, packet loss in reassembly mode servers is only influenced to an insignificant extent by improved packet discarding.

The impact of network cell loss is also evident in Figure 5. With a network cell loss rate of $l_{net} = 10^{-5}$, the network cell loss has very little effect on the end-to-end packet loss probabilities when the connectionless server packet loss probabilities are much larger than 10^{-4}. With an average packet size of 15 cells, the packet loss probability due to random cell loss in the network is $15 \times l_{net}$. This is evident in the figure as the minimum packet loss probability is approximately 1.5×10^{-4}. The minimum packet loss probabilities appear to be independent of the forwarding method used. Therefore, both streaming and reassembly modes suffer equally from random network cell loss.

End-to-end delay is another useful measure by which the forwarding modes can be compared. Figure 6 illustrates the end-to-end delay of streaming and reassembly mode servers for various buffer sizes and loads. A packet's end-to-end delay is defined as the time between the arrival of its EOM cell at *CLS-1* and the departure of its EOM from *CLS-2*. The delay for packets traversing reassembly mode servers is consistently greater than that for streaming mode servers with inputs using either interleaving scheme. This relationship is due largely to the packet reassembly time required before transmission.

Another observation obtained from Figure 6 is the narrowing of the delays between streaming and reassembly mode servers as load increases. For low loads or large buffer size where the packet loss probability is relatively low, the difference in delay for the two modes is fairly constant. However, as load increases to a point where packet loss becomes a significant factor, the difference in end-to-end delay becomes very small. This effect may be due to the better discarding characteristics of reassembly mode. Reassembly mode's better discarding characteristics reduce the load on the outputs of the congested connectionless server. This reduction in load improves the delay of packets that have not been discarded, thereby narrowing the gap between streaming and reassembly mode end-to-end delays.

The effect of cell interleaving on packet delay can be seen by examining the output duration of packets. The output duration of a packet is the time between the departures of its first and last cells from a connectionless server. Figure 7 illustrates how the average output duration changes in relation to the input duration (*i.e.*, the time between the arrivals of a packet's first and last cells) at *CLS-1*. Reassembly mode shows little evidence of increased output duration since it utilizes packet interleaving. The output duration is the same from both connectionless servers. On the other hand, streaming mode, which uses cell interleaving, exhibits output durations at *CLS-1* that are 50–70% longer than the input durations. After *CLS-2* the average packet output duration is more than twice as long as the original input duration at *CLS-1*. As expected, the output duration increases as more hops are traversed. However, the rate of growth diminishes at the second CLS. The growth of the output duration from *CLS-1* to *CLS-2*, the "CLS-2/CLS-1" graph, is only about 30%. This diminishing growth can also be seen in the output duration growth for streaming mode servers receiving cell interleaved input. The growth in output duration for cell interleaved input (CI) is considerably less than that for packet interleave input (PI). This confirms our earlier hypothesis of nonlinear growth in end-to-end packet delay.

Finally, the effect of the number of ports (N) on the performance of streaming and reassembly mode operation is interesting to note. Figures 8 and 9 illustrate the effect of increasing N on packet loss and delay at a connectionless server. The load on each output connection is kept constant at ($\rho = 0.75$), while the total buffer space increases proportionally with the number of ports: 250 cells for $N = 4$, 500 cells for $N = 8$, etc. Packet loss decreases, while packet delay increases due to increased buffer sharing. Note however that as N increases, reassembly mode exhibits a slower improvement in packet loss and a faster increase of delays than streaming mode.

In summary, the following observations can be made:

- Streaming mode connectionless servers with packet interleaved input provide better packet loss performance than reassembly mode connectionless servers. However, with cell interleaved input, the packet loss for streaming mode can be worse. This suggests that reassembly mode will perform better overall as more hops are traversed but streaming mode is better for shorter routes.

- The packet delay is consistently smaller using streaming mode than reassembly mode. Delay also grows at a slower pace for routes with more hops (see following observation).

- Although the output duration is longer under cell interleaving than under packet interleaving, this duration grows at a slower rate as more hops are traversed.
- While the better discarding characteristics and consecutive cell loss resilience of reassembly mode make it seem competitive with streaming mode, these characteristics do not significantly impact the loss performance of reassembly mode servers although it slightly improves delay performance.
- Random cell loss in the network does not affect one forwarding mode or interleaving scheme to a greater degree than the other.
- As N increases, the packet loss rate and delay of streaming mode further improves relative to reassembly mode.

4 Conclusion

Two packet forwarding techniques for ATM connectionless servers, streaming and reassembly modes, were studied through simulation. The simulations showed that streaming mode provides significant advantages in terms of packet delay. The packet delay is generally much lower with streaming mode. The simulations also indicated that the growth rate of packet delay for longer routes should also be lower with streaming mode. While streaming mode showed consistently lower packet loss probabilities for the simulation configuration, the high packet loss probability for streaming mode connectionless servers with cell interleaved input indicates that with longer routes the packet loss probability can be higher than for reassembly mode servers. However, the final two figures showed that with more than 4 input and output connections, the loss probabilities for streaming mode are much smaller than for reassembly mode. Although reassembly mode has some advantages over streaming mode, namely better packet discarding and more flexibility in interleaving, it was shown that these factors have very little effect on overall performance.

Extensions of this work will investigate the performance of early packet discarding (where newly arriving packets are discarded if the buffer occupancy is above a threshold), adaptive rate controls between connectionless servers, and the effect of consecutive network cell loss on the packet loss probabilities of cell and packet interleaving.

References

[1] CCITT Study Group XVIII, *CCITT Recommendation I.327: B-ISDN Functional Architecture*, 1991.

[2] CCITT Study Group XVIII, *CCITT Draft Recommendation I.364: Support of Broadband Connectionless Data Service on B-ISDN*, June 1992.

[3] K. Yamazaki, Y. Wakahara and Y. Ikeda, "Networks and switching for B-ISDN connectionless communications — Issues on interworking of two connectionless services, network topologies and connectionless message switching method," *IEICE Trans. Commun.*, vol. E76-B, no. 3, pp. 229–236, March 1993.

[4] A. Biocca, G. Freschi, A. Forcina and R. Melen, "Architectural issues in the interoperability between MANs and the ATM network," In *Proceedings of ISS '90*, vol. 2, pp. 23–28, May 1990.

[5] M. Gerla, T.-Y. Charles Tai and G. Gallassi, "Internetting LAN's and MAN's to B-ISDN's connectionless traffic support," *IEEE J. Select. Areas Commun.*, vol. 11, no. 8, pp. 1145–1159, Oct. 1993.

[6] Andrew W. Barnhart, "Baseline Performance Using PRCA Rate-Control," ATM Forum 94-0597 - 94-0599, July, 1994.

[7] D.F. Box, D.P. Hong and T. Suda, "Architecture and design of connectionless data service for a public ATM network," In *Proceedings of INFOCOM '93*, pp. 722–731, March 1993.

[8] Shun-ichi Iisaku and Masami Ishikura. "ATM network architecture for supporting the connectionless service," In *INFOCOM 90*, pp. 796-802, 1990.

[9] T. Takine, T. Suda, and T. Hasegawa, "Cell Loss and Output Process Analyses of a Finite-Buffer Discrete-Time ATM Queueing System with Correlated Arrivals," In *Proceedings of INFOCOM '93*, pp. 1259–1269, March 1992.

[10] T. Kawasaki, S. Kuroyanagi, R. Takechi and K. Hajikano, "A study on high speed data communication system using ATM," In *Proceedings of GLOBECOM '91*, pp. 2105–2109, Dec. 1991.

[11] G. Fioretto, T. Demaria, R. Vaglio, A. Forcina, and T. Moro, "Connectionless Service Handling Within B-ISDN," In *Proceedings of GLOBECOM '91*, pp. 217–222, Dec. 1991.

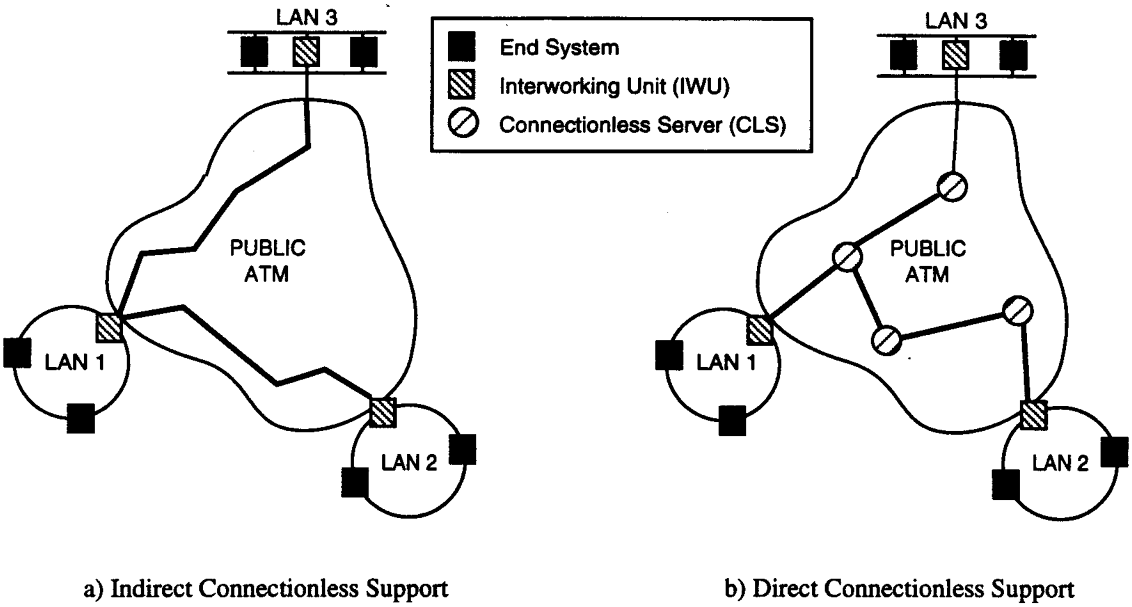

a) Indirect Connectionless Support

b) Direct Connectionless Support

Figure 1: Direct and Indirect Connectionless Service Architectures

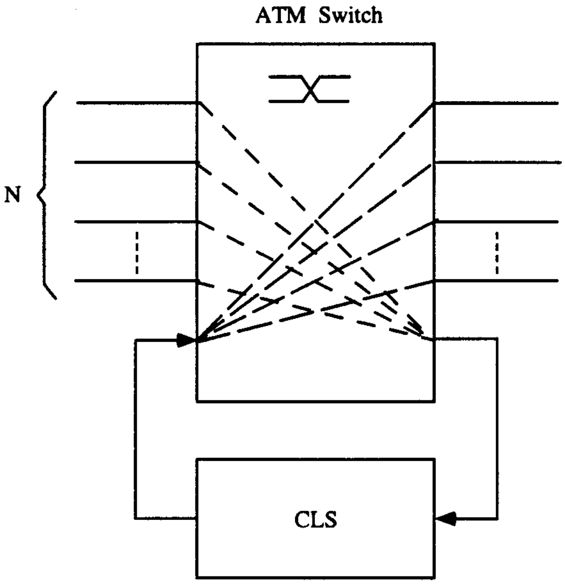

Figure 2: Switch-to-CLS Connection

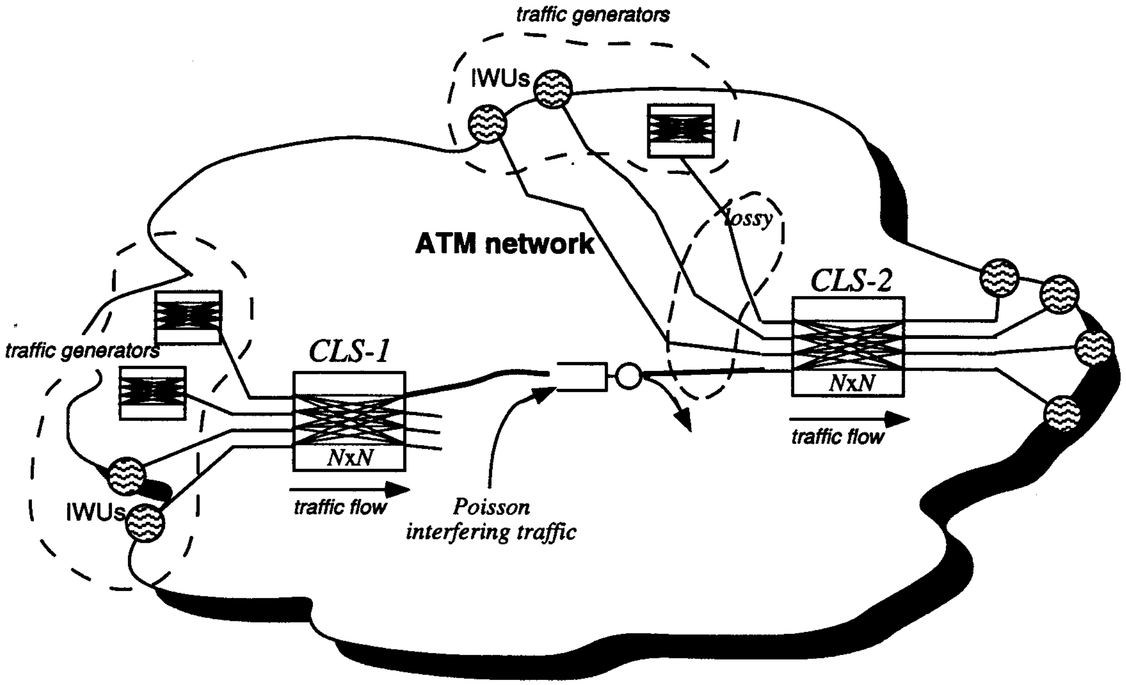

Figure 3: Network Simulation Model

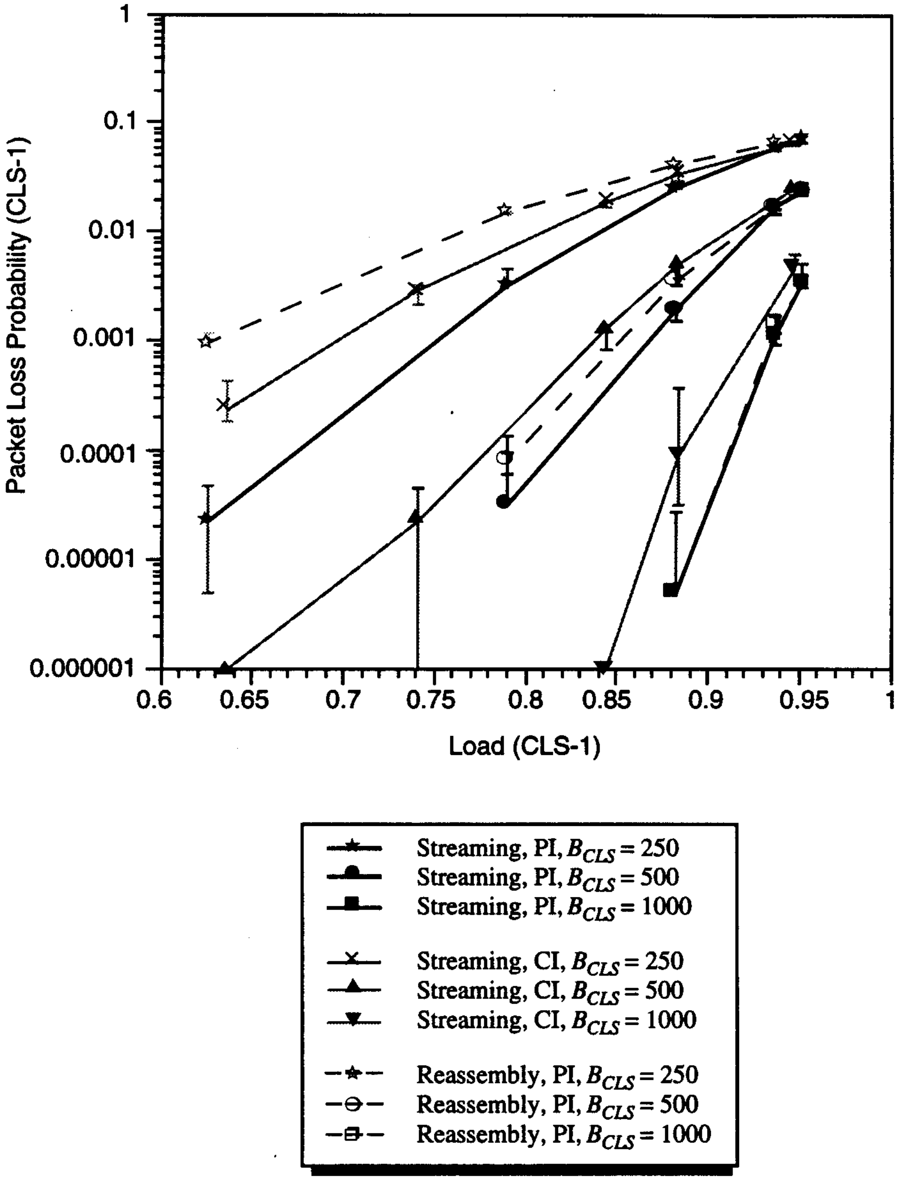

Figure 4: Packet Loss Probability at *CLS-1*

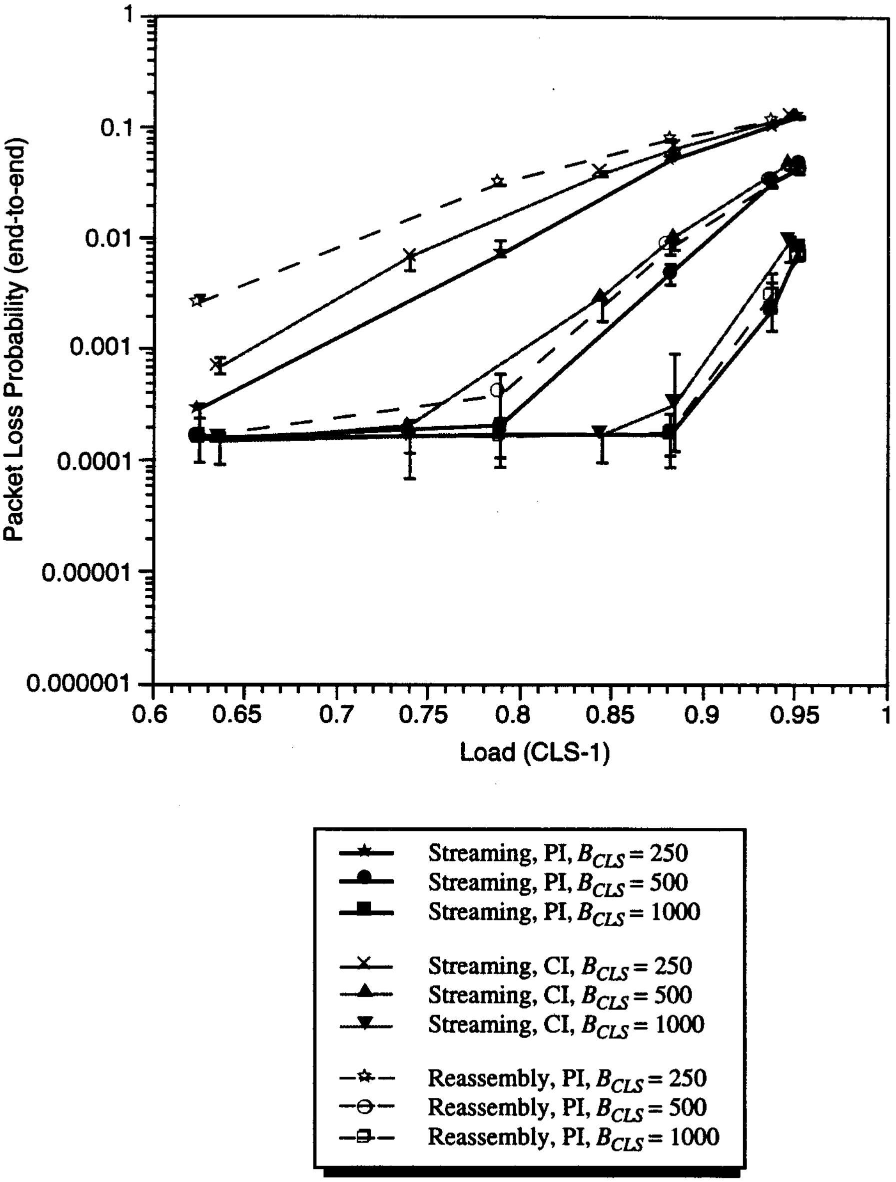

Figure 5: *End-to-End* Packet Loss Probability

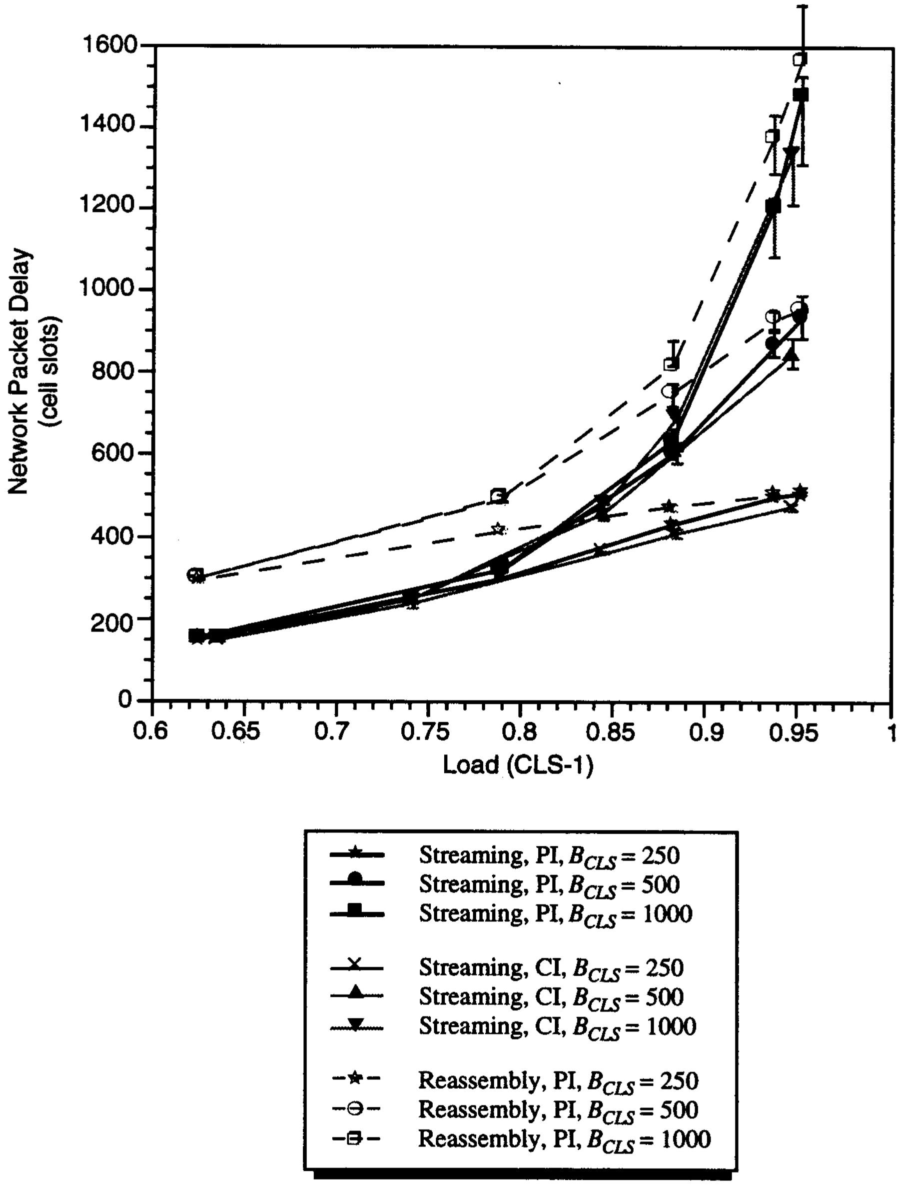

Figure 6: *End-to-End* Packet Delay

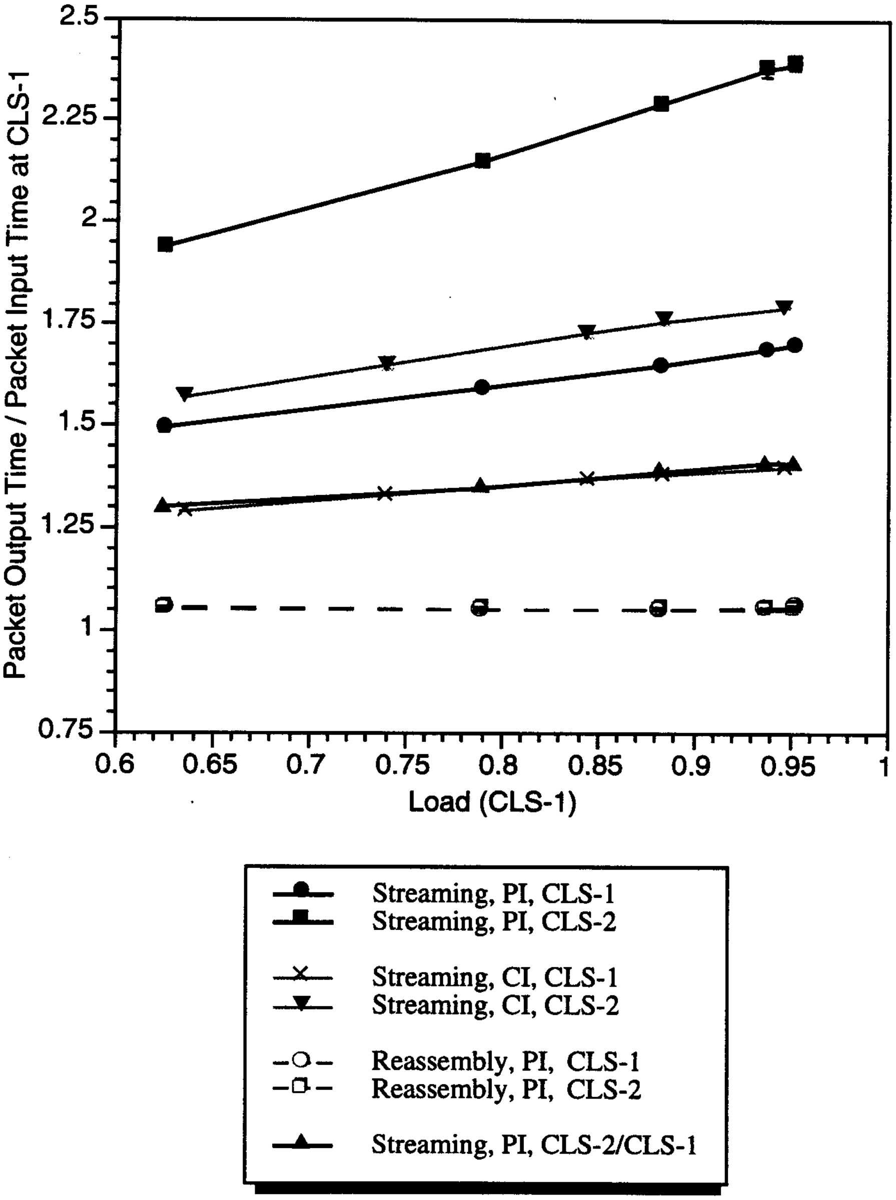

Figure 7: Packet Input Time and Output Times at *CLS-1* and *CLS-2*

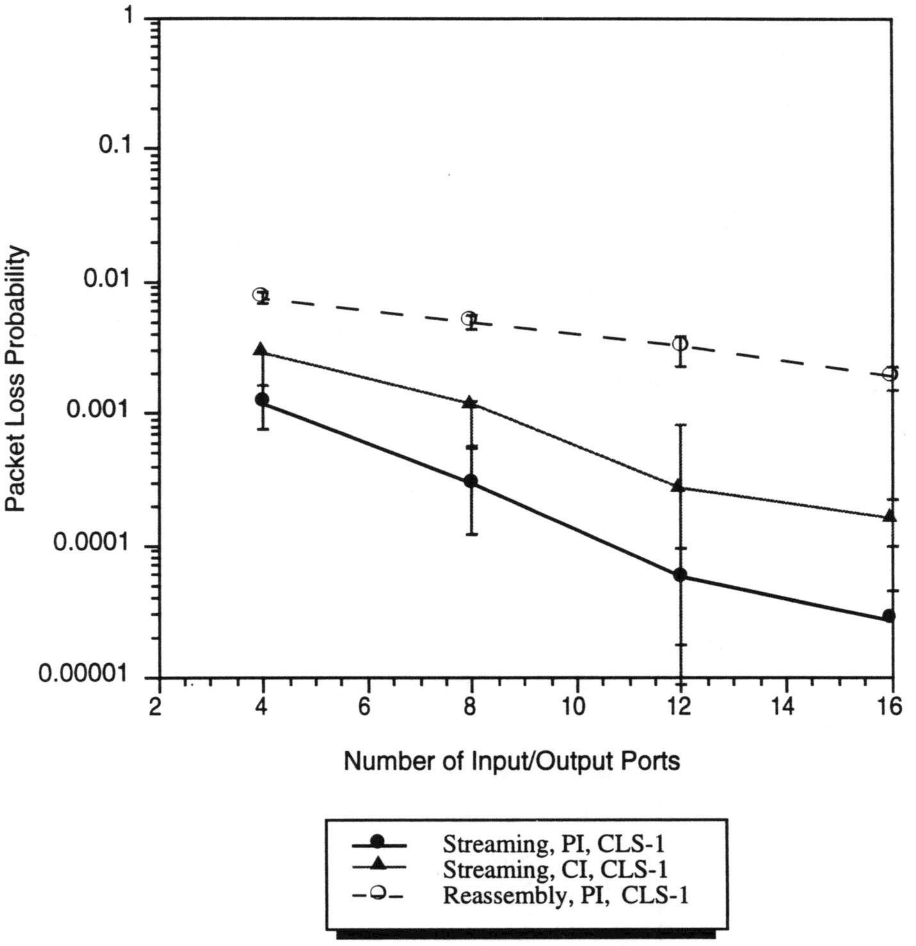

Figure 8: Packet Loss Probability vs. Number of Ports, Single CLS

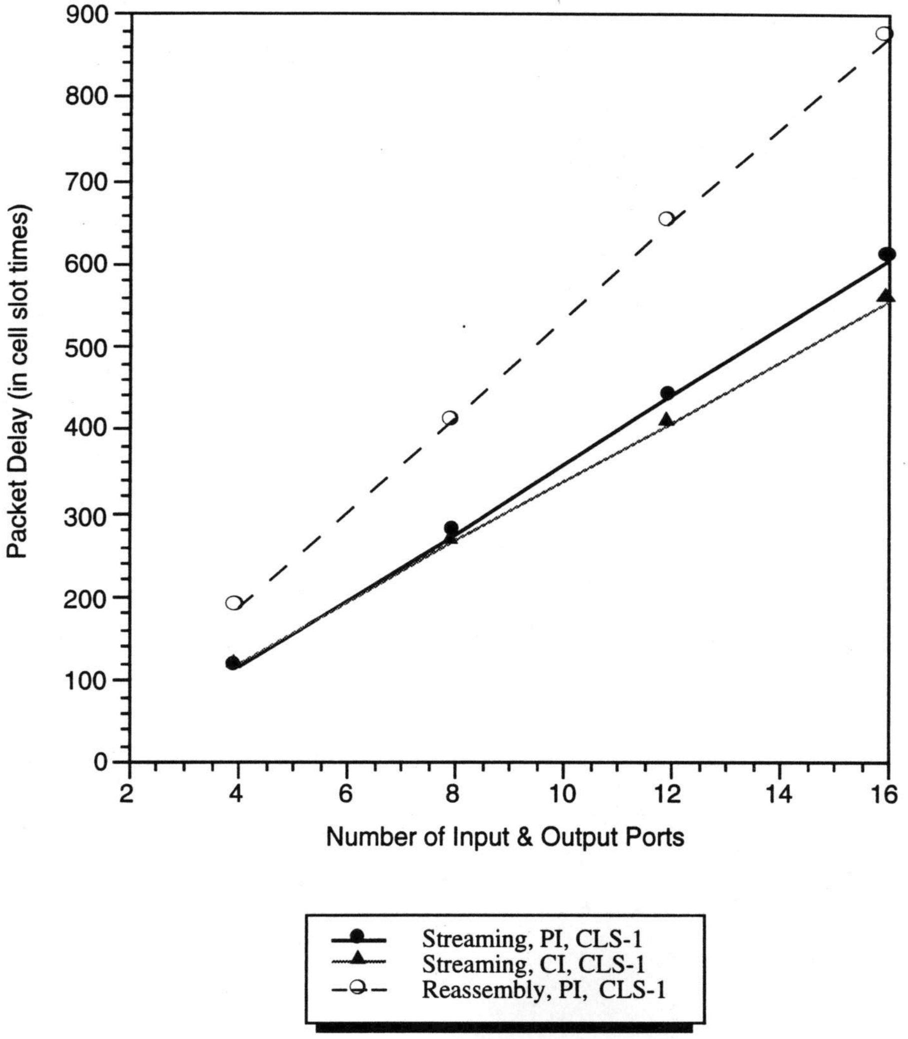

Figure 9: Packet Delay vs. Number of Ports, Single CLS