Efficient Multicast in IP/ATM Networks

Walid Mostafa   Mukesh Singhal
Department of Computer and Information Science
The Ohio State University
Columbus, OH 43210
email: (mostafa,singhal)@cis.ohio-state.edu
Phone: 614-292-5839
Fax: 614-292-2991

Abstract

Real-time transport protocols rely on IP to disseminate audio/video packets from a sender to one or more receivers. In IP over ATM broadband networks, studies have shown that the effective bandwidth is reduced by at least 10% and that larger protocol payload sizes yield higher throughput. However, a dropped cell corrupts the IP packet it belongs to and reduces the rate successfully-received data, or goodput. This results in quality of service degradation to multimedia traffic, particularly for intra-frame encoded video raffic that is sensitive to a packet loss. A key observation for multimedia applications is that achieving maximum throughput versus reaching optimum goodput is determined by the end-to-end quality of service requirements. In this report, we derive analytic expressions for the maximum throughput and packet-loss probability which provides us with a better understanding of the goodput of real-time transport protocols. We also derive an analytic expression for optimum PDU that maximizes the goodput with respect to cell-loss probability and number of ATM switches. We show that optimum PDUs yield higher goodput than the default MTU for IP over ATM, particularly as the cell-loss probability and/or the number of ATM switches increase. For reliable multicast, We also derive analytic expressions to the expected total number of transmissions per packet and to the goodput of reliable multicast.

Keywords: Multimedia Dissemination, Reliable Multicast, Transport PDU, IP over ATM.

1 Introduction

The continuous growth in multimedia applications over the Internet, e.g, video conferencing [Jac94] and distributed whiteboards [McC92], and the availability of high-speed broadband ATM networks have led to considerable interests in the interoperability of IP and ATM [Lau94, CSV96]. Real-time transport protocols, e.g., RTP [SCFJ96, Sch96], cater for the needs of audio/video traffic of multimedia applications [HFG96, BFFM96, SH96, Vau96]. These protocols provide high throughput because they neither guarantee delivery or ordering of transport packets nor provide flow control.

Multimedia applications over the Internet disseminate audio/video traffic in real-time and are sensitive to network delays. However, they have various degrees of tolerance to packet loss. For example, the ITU recommendation H.261 [IT93] that uses the temporal redundancy of video to perform compression, is sensitive to packet loss. Thus, in specifying the RTP payload format for H.261 video, Turletti and Muitema [TH96] defined negative acknowledgement control packets to request packet retransmission. This approach mimics reliable transport protocols.

Multimedia transport requires either high throughput of redundant data that tolerate packet loss or lower throughput of compressed data that is sensitive to packet loss. In other words, techniques that reduce multimedia bandwidth requirements increase its quality of service demands. This fundamental observation is the motivation to analyze the tradeoff between the maximum rate of transmitted traffic, *throughput*, at a source and the maximum rate of successfully-received traffic, *goodput*, at one or more destinations.

In this report, we derive analytic expressions for the maximum throughput and packet-loss probability which provide us with a better understanding of the goodput of real-time transport protocols. We also derive an analytic expression for the optimum protocol payload data unit size that maximizes the goodput of data transport with respect to cell-loss probability and number of ATM switches. We show that optimum PDUs yield higher goodput than the default MTU for IP over ATM (9180 bytes), specially as the cell-loss probability and/or the number of ATM switches increase. We also derive similar expressions to the goodput of reliable multicast where a packet is retransmitted until it is correctly received by all receivers.

The rest of the report is organized as follows: Section 2 describes the protocol stack overhead for multimedia transport using RTP and IP over ATM. In Section 3, we derive analytic expressions for maximum throughput, packet loss probability, and optimum PDU using real-time multicast with no retransmissions of lost packets. In Section 4, we derive analytic expressions for the expected total number of retransmissions per lost packet and the goodput for reliable multicast. Section 5 presents numerical results for the analytic expressions derived earlier and compares them with simulation results. It also presents a goodput comparison between optimum PDU and default IP/ATM MTU. A summary is presented in Section 6.

2 Protocol Stack Overhead

In this section, we present the model and assumptions of multimedia transport using IP over ATM networks. The ATM layer reduces the available bandwidth to the higher protocol layer by at least 10% due to ATM cell header overhead. Cavanaugh [Cav94] evaluated the available bandwidth to an application using Internet protocols and ATM over a variety of physical layers. Other studies evaluated the performance of Internet protocols over ATM [RF94, LO95, LFEN96].

The real-time transport protocol, RTP, provides end-to-end network transport services. It supports both point-to-point and point-to-multipoint transport. It does not address resource reservation and does not guarantee quality of service for real-time services. It provides a control protocol to allow monitoring of data delivery and to provide minimal control and identification functionality.

Figure 1 shows the various protocol layers that are involved in multimedia data transmission using RTP. The application data is referred to as RTP payload (PDU). RTP adds an

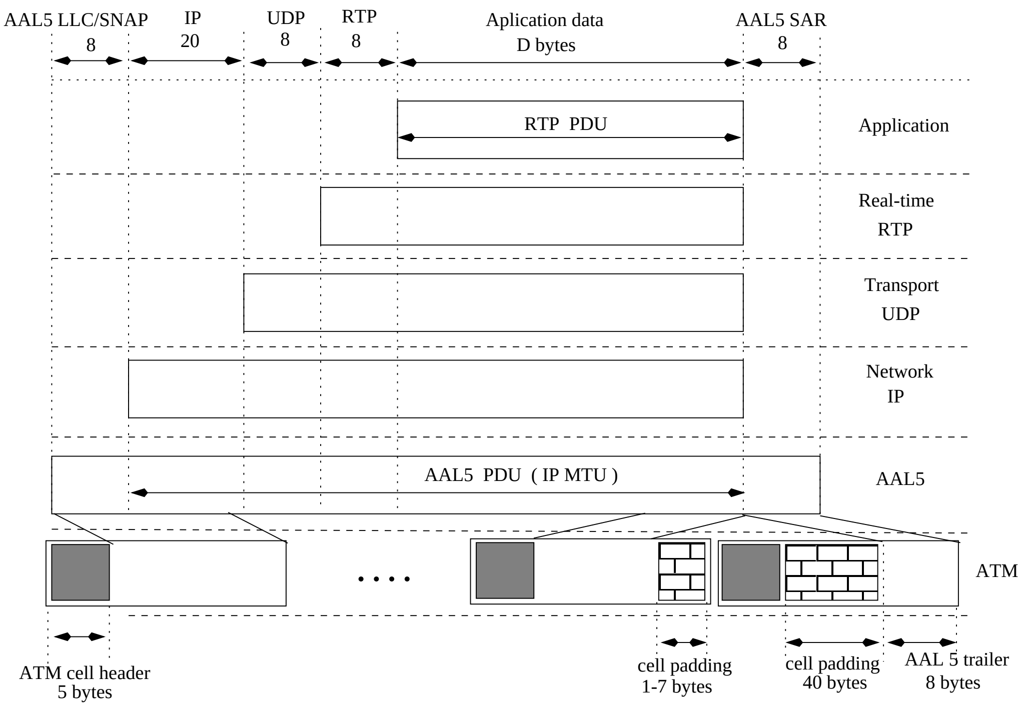

Figure 1: IP/ATM Protocol Stack Overhead

8-byte header. For point-to-point connections, RTP uses UDP which in turn uses IP. UDP header is 8 bytes and IP header is 20 bytes. The constructed IP datagram is referred to as ATM AAL5 payload. Although AAL5 can accept payloads up to 64K bytes, the maximum transmission unit, MTU, of IP over AAL5 is set to 9,180 bytes [Atk94]. AAL5 adds an 8-byte trailer that occupy the last 8 bytes of the last cell. If LLC/SNAP is used, 8 bytes are added as a header.

3 Optimum PDU for Maximum Unreliable Goodput

In this section, we derive analytic expressions to maximum throughput, ρ, and packet loss probability, p. We use these expressions to derive an expression for the goodput α of a transport protocol as a function of cell loss probability, q, number of switches per receiver, s, and number of cells per PDU, c. We then derive the optimum PDU, θ, that maximizes α as a function of q and s.

Figure 1 shows the overhead added to application data to be transmitted as an IP datagram over an ATM network. The overhead bytes account for a total of 52 bytes. However, RTP can use IP directly, e.g., for multicasting, so that may reduce the 8 bytes used by UDP. LLC/SNAP is optional as well. To accommodate these options, we assume that the overhead is 48 bytes or less. This is the equivalent of an ATM cell payload. Thus, we allocate one cell for protocol stack overhead and assume that a PDU is a multiple of 48 bytes. In other words, we assume that a PDU fits into c cells where the size of the PDU equals $48c$. Therefore, the actual throughput available to AAL5 can be expressed as

$$\rho = \frac{48c}{53(c+1)} \tag{1}$$

where ATM layer constructs $(c+1)$ cells with 53-bytes each and pass it down the protocol stack to the network device driver.

We assume a multicast network model where the expected number of ATM switches between a sender and a receiver is s and assume that these switches are independent. A transport packet is considered lost or corrupted if one or more cells are dropped at any of the ATM switches between a sender and a receiver. Assuming that cell loss in ATM switches are independent, the probability that a node receives a cell correctly after passing through s switches is $(1-q)^s$. However, a transport packet is comprised of $(c+1)$ cells. Thus, the probability that a node receives a transport packet is $\prod_{i=1}^{c+1}(1-q)^s = (1-q)^{s(c+1)}$. Therefore, the probability that a packet is lost or is received corrupted at a receiver is

$$p = 1 - (1-q)^{s(c+1)}. \tag{2}$$

We now derive an expression for the goodput α of a transport protocol. We define the goodput as the percentage of transport packets that are received correctly per unit time. Therefore, the goodput can be expressed as follows:

$$\alpha = \rho(1-p). \tag{3}$$

We note that for multiple receivers, the number of receivers that receive a packet correctly is binomially-distributed assuming that their network paths are independent. Therefore, the ratio of the expected number of receivers that receive the packet correctly to the total number of receivers is p. Using Equations (1) and (2), we rewrite Equation (3) as follows:

$$\alpha = \frac{48c(1-q)^{s(c+1)}}{53(c+1)}. \tag{4}$$

We are interested in the value of c that maximizes α. Note that q and s are parameters of the ATM network that a multimedia application has no control over. We derive the partial derivative of α with respect to c, equate it to zero, and then solve the resulting equation for c. Hence,

$$\begin{aligned}\frac{\partial\alpha}{\partial c} &= \frac{48(1-q)^{s(c+1)}}{53(c+1)} - \frac{48c(1-q)^{s(c+1)}}{53(c+1)^2} + \frac{48c(1-q)^{s(c+1)} s ln(1-q)}{53(c+1)} \\ &= \frac{48(1-q)^{s(c+1)}}{53(c+1)}(1 - \frac{c}{(c+1)} + c\ s\ ln(1-q)).\end{aligned} \tag{5}$$

Equating $\frac{\partial\alpha}{\partial c}$ to zero in Equation (5) and substituting θ for c, we get

$$\theta(\theta+1)\ s\ ln(1-q) + 1 = 0 \tag{6}$$

and the larger root can be derived as

$$\theta = \frac{1}{2}(-1 + \sqrt{1 - \frac{4}{s\ ln(1-q)}}). \tag{7}$$

4 Optimum PDU for Maximum Reliable Goodput

In section 3, we analyzed the goodput α of multicast protocols with unreliable delivery, i.e., a packet is sent only once by a sender. The goodput of unreliable multicast is the percentage of packets that are received correctly per unit time. The goodput of reliable multicast however, is defined as the percentage of the bandwidth used for one transmission per packet. The difference between the throughput and the goodput is attributed to lost packet retransmissions.

In this section, we derive analytic expressions for the goodput of multicast protocols that guarantee reliable delivery of packets form one sender to R receivers. We use the same assumptions as stated in Section 3. Therefore, the probability p that a packet is lost or is received corrupted at a receiver is that expressed in Equation (2).

We now derive an expression for reliable multicast goodput $\alpha^{'}$. Assume that the transport PDU size is $48c$ bytes. Thus, for a sender to transmit such a transport packet once, the ATM layer at the sender's host constructs $(c+1)$ cells with 53-bytes each and pass them down the protocol stack to the network device driver. Let M be the total number of transmissions required for a packet to be received correctly by all receivers. Therefore, the ATM layer of a sender passes $53(c+1)E[M]$ bytes to the device driver in oreder to reliably multicast $48c$ bytes. Hence, the goodput of reliable multicast protocols can be expressed as

$$\alpha^{'} = \frac{48c}{53(c+1)E[M]}. \tag{8}$$

Using Equation (1), we rewrite $\alpha^{'}$ as

$$\alpha^{'} = \rho/E[M]. \tag{9}$$

The only unknown in Equation 8 is $E[M]$. Using the result obtained in [PTK94], we have:

$$E[M] = \sum_{i=1}^{R} \left(\begin{array}{c} R \\ i \end{array} \right) (-1)^{i+1} \frac{1}{(1-p^i)}. \tag{10}$$

Thus, we can rewrite Equation (8) as

$$\alpha^{'} = \frac{48c}{53(c+1)\sum_{i=1}^{R} \left(\begin{array}{c} R \\ i \end{array} \right) (-1)^{i+1} \frac{1}{(1-p^i)}}. \tag{11}$$

The expression for $\alpha^{'}$ is much more complex than that of α in Equation (3). Therefore, obtaining a derivative of $\alpha^{'}$ with respect to c, $\frac{\partial \alpha^{'}}{\partial c}$, to derive the maximum PDU may be very complex and hence, not useful. In Section 5, we present a numerical evaluation of Equation (11) to show the local optimal values of c that maximizes the goodput $\alpha^{'}$.

5 Analytic and Simulation Results

In this section, we present numerical results for optimum PDU sizes using both unreliable and reliable multicast.

5.1 Unreliable Multicast

We show the impact of the real-time protocol PDU size c on the maximum packet transport throughput ρ available to ATM AAL 5 layer, on the transport packet loss probability p and on the packet transport goodput α We also show the optimum PDU size θ for different packet loss probabilities and for different number of ATM switches per receiver. We then compare the goodput using θ with the goodput using the default IP/ATM MTU for different values of p and s. Finally we show that our simulation model provides goodput asymptotic to the analytic one and that it yields an optimum PDU consistent with the analytically-derived one.

Maximum Throughput

Figure 2 shows how rapidly the maximum throughput ρ available to AAL5 layer increases before it asymptotically reaches a constant value as the PDU size increases. The asymptotic value $\frac{48}{53}$ is attributed to the ATM layer overhead of 5 bytes out of every 53-byte cell. The rapid increase in ρ is due to the ratio $\frac{c}{c+1}$ as c increases for smaller values of c. It asymptotically reaches a value of 1 as $c >> 1$. This is a key observation because it shows that the larger the PDU, the higher the throughput. We also observe that the asymptotic value of the maximum throughput is reached around the 200-300 mark. This corresponds to a range of 9600-14400 bytes, which is in the range of the default IP/ATM MTU 9180 bytes. We note that we use PDU sizes that are multiple of 48. Otherwise, we would have obtained a sawtooth curve as the padding in the last one or two ATM cells ranges from 1 to 47 bytes.

Packet Loss

Figure 3 shows that the transport packet loss probability increases exponentially either as the PDU size increases or as the the number of ATM switches per receiver increases. This is because increasing c implies that an additional cell has to traverse s switches, thus, increasing the probability of dropping the packet by a factor of $(1-q)^s$. A similar argument applies to increasing s. This is the other key observation because it shows that the larger the PDU size, the higher the probability of corrupted delivery of packets.

Goodput

Figure 4 shows the impact of increasing the PDU size, c, on packet transport goodput, α. We observe that α increases as c increases for lower values of c similar to the behavior of ρ. This is because for low PDU values, the goodput behavior is dominated by the ATM layer overhead while packet loss probability is very low. The goodput reaches its peak at an optimum PDU before it starts dropping as the PDU increases. This is because after a point, packet loss probability, p, becomes more dominant. Thus, the faster p dominates the goodput, the smaller the peak/optimum PDU size. This is true as s and/or q increase.

Optimum PDU

Figure 5 shows the optimum PDU, θ, for different values of cell loss probability, q, and number of ATM switches per receiver, s. We observe that θ is over 300 cells $(14,480$ bytes) when $q < 10^{-6}$ and $s < 10$. However, θ reaches the single digits, i.e., < 480 bytes, for very high cell loss probabilities, e.g., 10^{-3}.

Comparison with Default MTU

Figure 6 shows the ratio of the goodput using the optimum PDU size, θ, and the goodput using the default IP/ATM MTU for different values of cell loss probability, q, and number of ATM switches per receiver, s. We observe that both PDUs yield similar goodput for low values of q because the packet loss probability p is low. However, the ratio increases exponentially past the 10^{-4} cell loss probability. This is because using the default IP/ATM MTU, which is equivalent to 192 cells, suffers higher packet loss as q increases.

Simulation Results

Figure 7 shows the transport goodput using both analytic and simulation models when q is 10^{-3} and s takes the values 2, 5, 10, respectively. We used the analytically-driven optimum PDU θ for a given (q, s) to derive our simulation running the model at PDUs smaller and larger than θ. The results of the simulation model exhibited the convex shape around the optimum PDU and is asymptotic to the analytic results.

5.2 Reliable Multicast

For reliable multicast, we show the impact of cell loss and number of switches on the number of transmissions required to multicast a packets correctly to all R receivers.

Number of Transmissions/Packet

Figure 8 shows how the expected total number of transmissions per packet increases as the transport PDU size increases. This is because the more cells in a transport packet, the higher the probability that such a packet is not received at one or more receivers and, consequently, the higher the number of retransmissions required for this packet. We also observe that the rate of increase in number of transmissions becomes more significant as cell loss probability increases. For example, when multicasting to 500 receivers, the expected total number of transmissions of a 528-bytes packet is 1.26, 1.35, and 1.59 for cell loss probabilities of 10^{-6}, 10^{-5}, and 10^{-4}, respectively. The corresponding numbers for a 9168-bytes packet are 1.37, 1.70, and 2.90, respectively. For a 65,520-bytes packet, the expected total number of transmissions are 1.62, 2.64, and 6.73, respectively.

Goodput

Figure 9 shows the impact of increasing the PDU size c on reliable multicast goodput $\alpha^{'}$. Like the case of unreliable multicast, $\alpha^{'}$ increases as c increases for lower values of c until it reaches a peak and then decreases as c increases. This is because the peak point represents the change from an ATM-overhead-dominated goodput to a packet-loss-dominated one.

Optimum PDU

Figure 10 shows that the numerically-obtained local optimum transport PDU sizes decrease as the number of receivers increases. This is due to the increased expected number of transmissions per packet as explained in Figure 8. We observe that the optimum PDU size decreases with a high rate for low number of receivers before it asymptotically reaches a constant value. This is because the ATM-layer overhead becomes dominant at small values of transport PDU sizes.

Figure 11 shows that the numerically-obtained local optimum transport PDU sizes decrease as the number of ATM switches between the sender and a receiver increase. This is attributed to the increased probability of cell loss and consequently, packet loss, which result in an increased number of transmissions per packet. We observe that the optimum PDU size decreases with a high rate for low number of switches before it asymptotically reaches a constant value. This is similar to its behavior with the increase in number of receivers as explained in Figure 10.

6 Summary

Multimedia transport requires either high throughput of redundant data that tolerate packet loss or lower throughput of compressed data that is sensitive to packet loss. Thus, techniques that reduce multimedia bandwidth requirements increase its quality of service demands. In this report, we highlighted the tradeoff between the maximum throughput at the source and the maximum goodput at the destination for both unreliable real-time multimedia traffic dissemination and reliable data dissemination.

We derived analytic expressions for both quantities and showed their relationship to transport packet loss probability. We also derived an analytic expression for the optimum protocol payload data unit that a multimedia application can use to maximize its goodput and showed that it has better goodput than the default IP/ATM MTU, particularly for high cell loss probability. For reliable multicast in IP/ATM networks, we derived analytic expressions for the expected number of transmissions per packet and the protocol goodput. We showed that the effective bandwidth increases until it reaches a maximum value and then decreases as the transport PDU size increases. We showed that the optimum transport PDU size decreases as the number of receivers or the number of ATM switches per receiver increases.

References

[Atk94] R. Atkinson. "Default IP MTU for use over ATM AAL5 ". *Internet Request for Comment (RFC 1626)*, May 1994.

[BFFM96] L. Berc, W. Fenner, R. Frederick, and S. McCanne. "RTP Payload Format for JPEG-compressed Video ". *Internet Request for Comment (RFC 2035)*, October 1996.

[Cav94] John David Cavanaugh. "Protocol Overhead in IP/ATM Networks". Technical report, Minnesota Supercomputer Center, August 1994.

[CSV96] R. Cole, D. Suhr, and C. Villamizar. "IP over ATM: A Framework Document ". *Internet Request for Comment (RFC 1932)*, April 1996.

[HFG96] D. Hoffman, G. Fernando, and V. Goyal. "RTP Payload Format for MPEG1/MPEG2 Video ". *Internet Request for Comment (RFC 2038)*, October 1996.

[IT93] ITU-T. "Video Codec for Audiovisual Services at p 64 kbit/s ". *(International Telecommunication Union - Telecommunication Standardisation Sector) Recommendation H.261*, 1993.

[Jac94] Van Jacobson. "Multimedia Conferencing on the Internet". In *SIGCOMM 1994 Tutorial Notes*. ACM, August 1994.

[Lau94] M. Laubach. "Classical IP and ARP over ATM ". *Internet Request for Comment (RFC 1577)*, January 1994.

[LFEN96] G. Y. Lazarou, V. S. Frost, J. B. Evans, and D. Niehaus. "Using Measurements to Validate Simulation Models of TCP/IP over High Speed ATM Wide Area Networks". In *International Conference on Communications*, page S11, Dallas, Texas, June 1996.

[LO95] Gunnar Lindberg and Robert Olsson. “Report from COAST Phase 1: Assorted Experiments with TCP/IP over ATM Wide Area Links”. Technical report, SUNET, Gothenburg and Uppsala, Sweden, September 1995.

[McC92] S. McCanne. “A Distributed Whiteboard for Network Conferencing”. Technical report, Technical report, Real Time Systems Group, Lawrence Berkeley Labo ratory, 1992. DSG Working Paper.

[PTK94] S. Pingali, D. Towsley, and J.F. Kurose. “A Comparison of Sender-Initiated and Receiver-Initiated Reliable Multicast Protocols”. In *SIGMETRICS '94 Conference*. ACM, May 1994.

[RF94] A. Romanow and S. Floyd. “The Dynamics of TCP Traffic over ATM Networks”. In *Proceedings of ACM SIGCOMM ‘94*, pages 79–88, London, UK, September 1994.

[SCFJ96] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. “RTP: A Transport Protocol for Real-Time Applications ”. *Internet Request for Comment (RFC 1889)*, January 1996.

[Sch96] H. Schulzrinne. “RTP Profile for Audio and Video Conferences with Minimal Control ”. *Internet Request for Comment (RFC 1890)*, January 1996.

[SH96] M. Speer and D. Hoffman. “RTP Payload Format of Sun's CellB Video Encoding ”. *Internet Request for Comment (RFC 2029)*, October 1996.

[TH96] T. Turletti and C. Huitema. “RTP Payload Format for H.261 Video Streams ”. *Internet Request for Comment (RFC 2032)*, October 1996.

[Vau96] G. Vaudreuil. “Voice Profile for Internet Mail ”. *Internet Request for Comment (RFC 1911)*, February 1996.

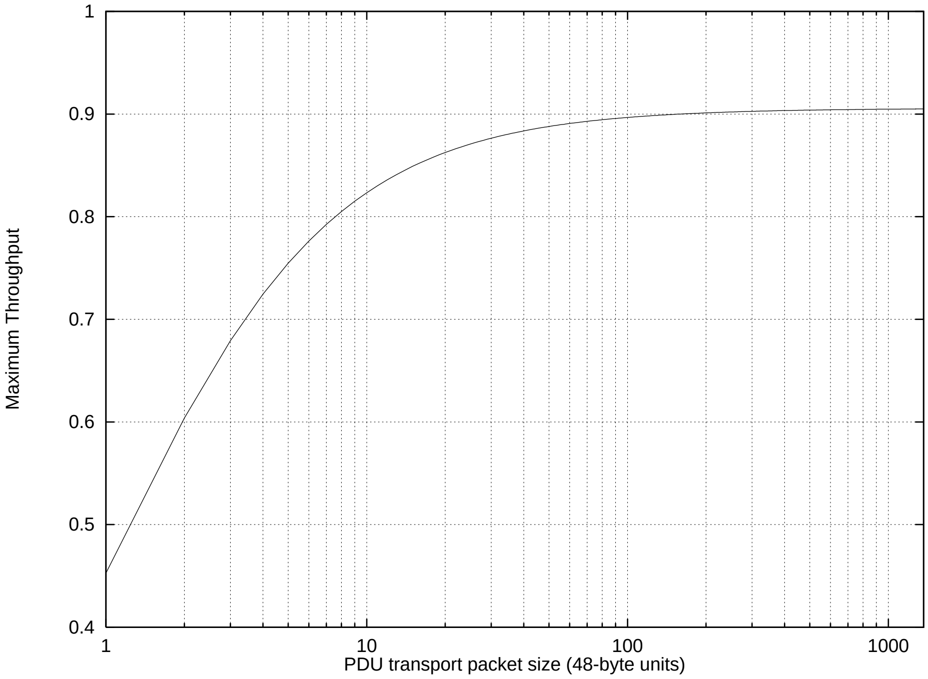

Figure 2: Maximum Effective Bandwidth vs. Protocol PDU

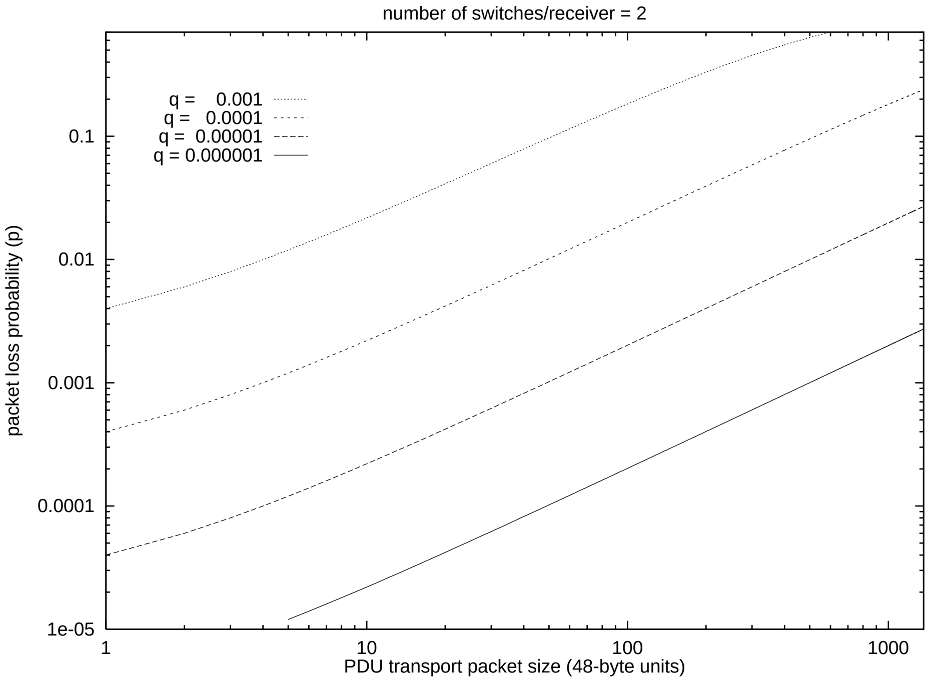

Figure 3: Transport Packet Loss Probability vs. Protocol PDU

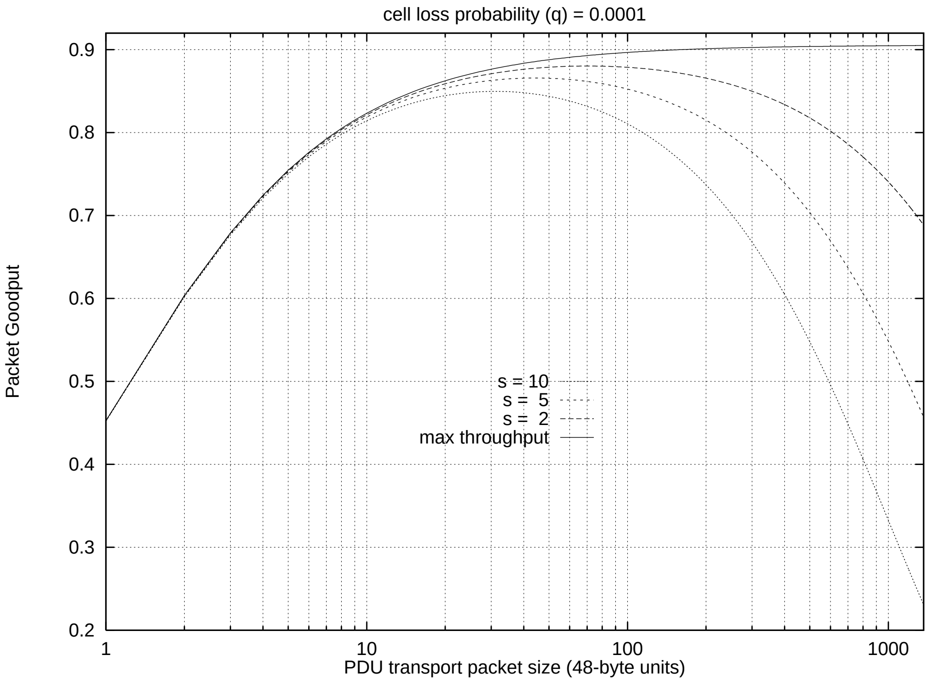

Figure 4: Transport Packet Goodput vs. Protocol PDU

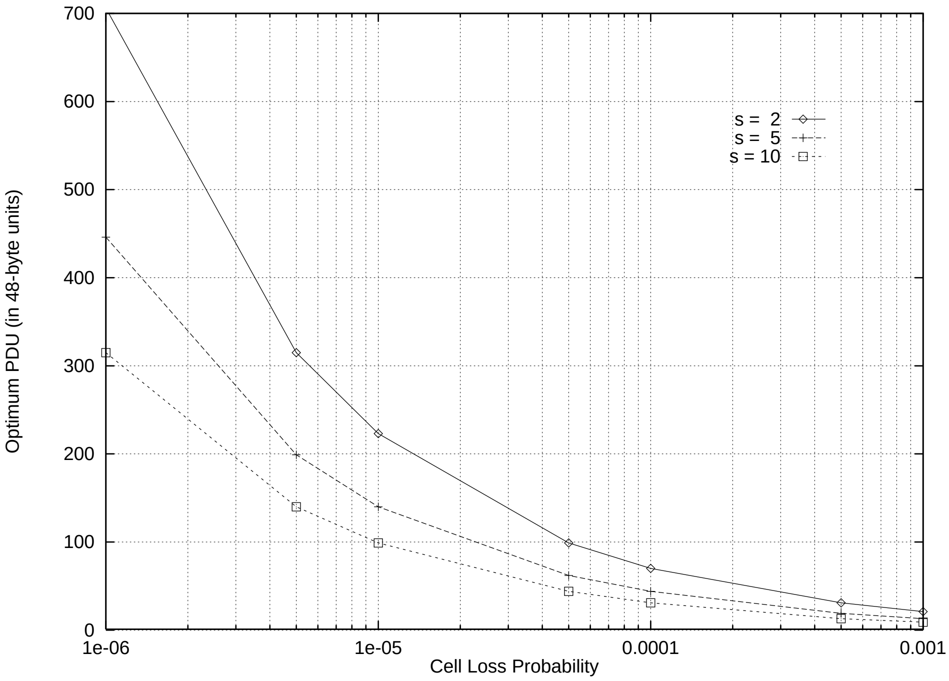

Figure 5: Optimum PDU vs. ATM Cell Loss Probability

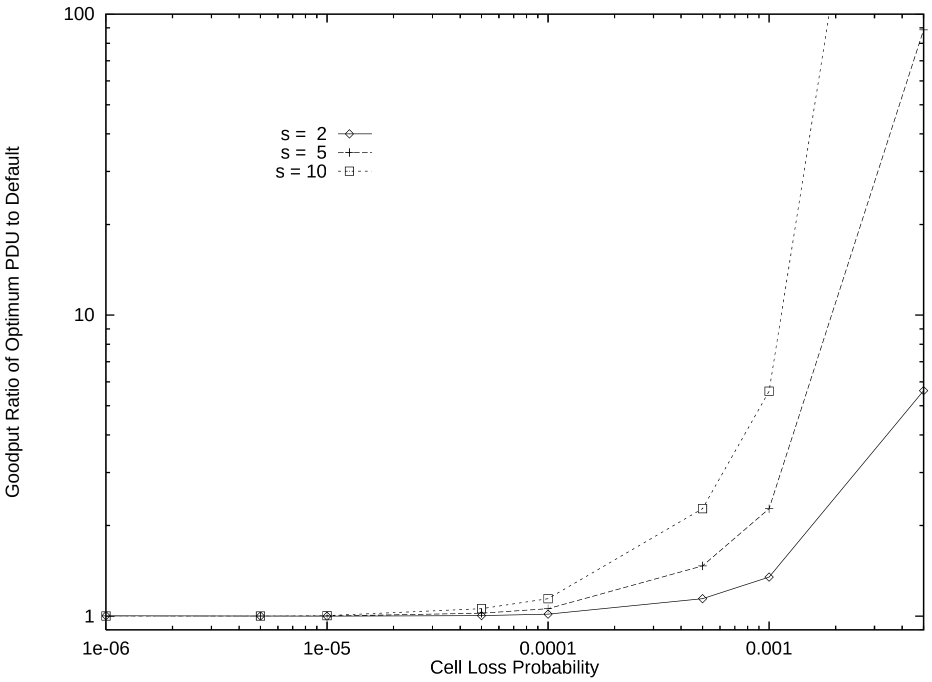

Figure 6: Goodput Ratio of Optimum PDU to Default IP/ATM (9180 bytes)

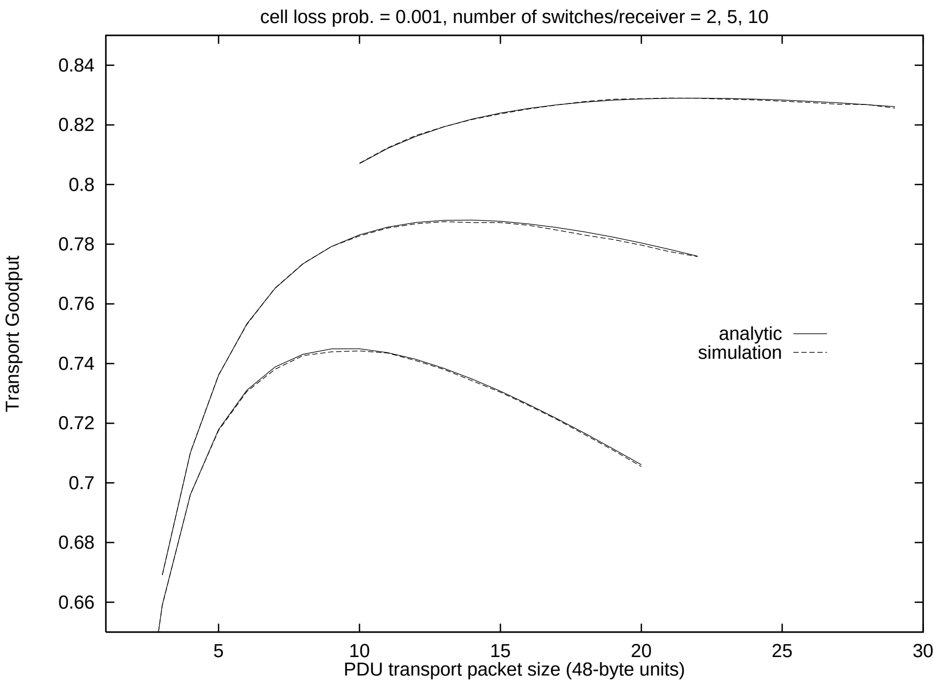

Figure 7: Goodput and Optimum PDU using Analysis and Simulation

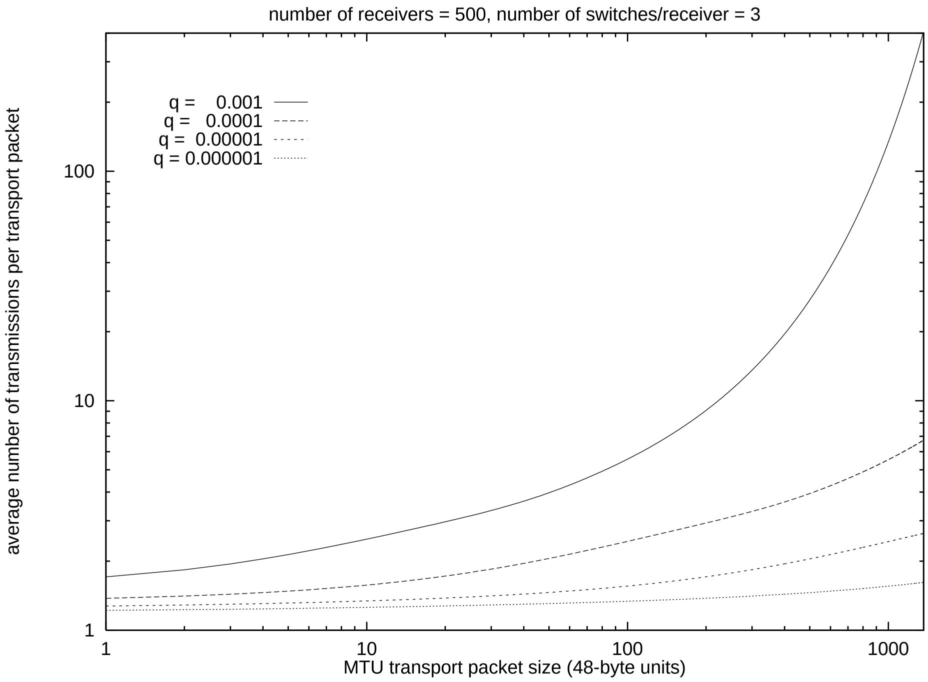

Figure 8: Packet Transmissions vs. Protocol PDU

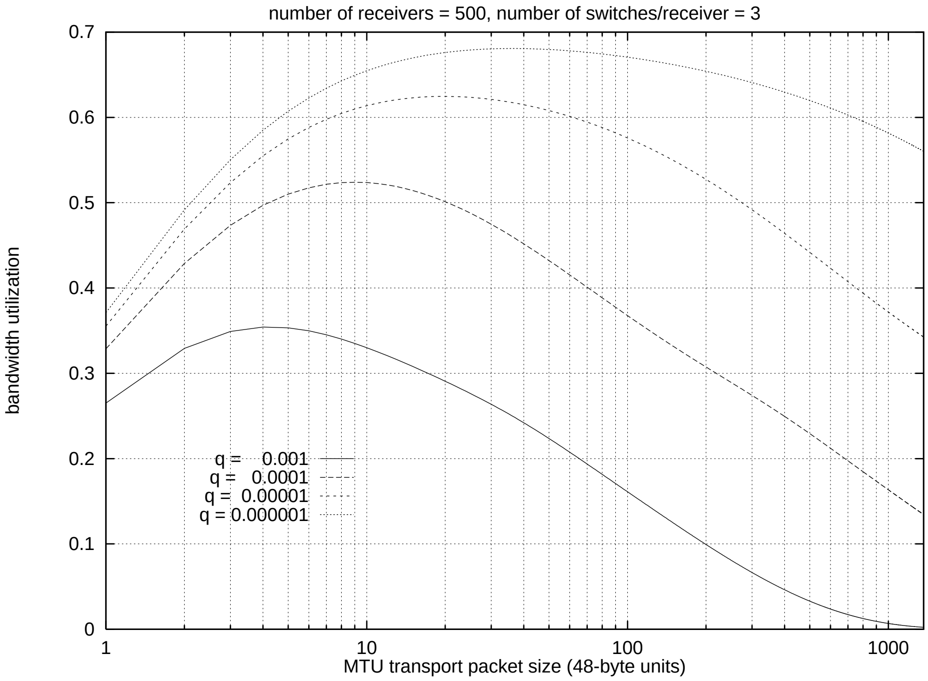

Figure 9: Reliable Goodput vs. Protocol PDU

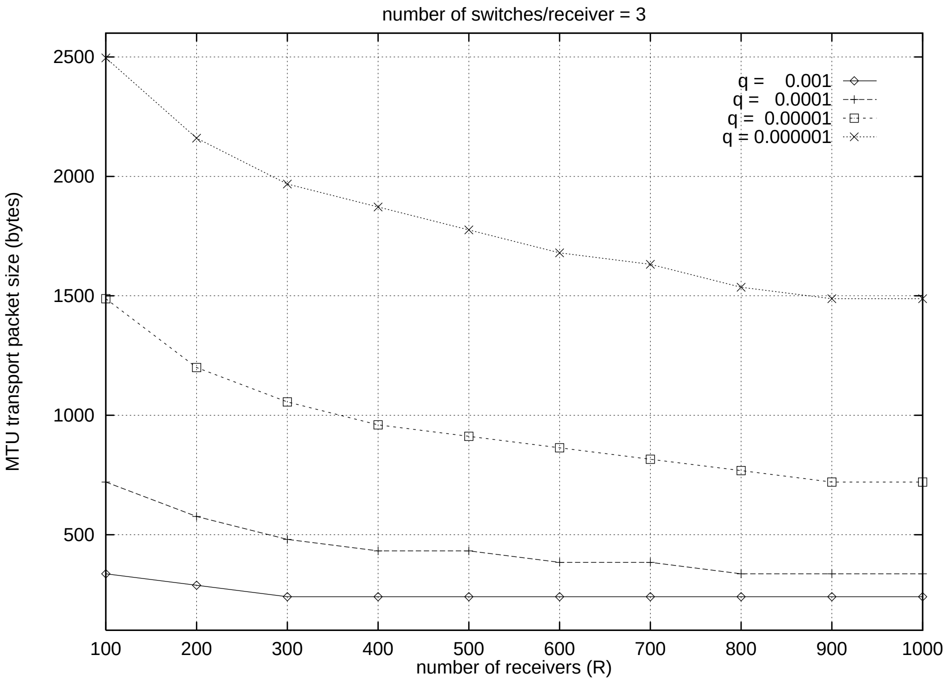

Figure 10: Optimum PDU vs. Number of Receivers

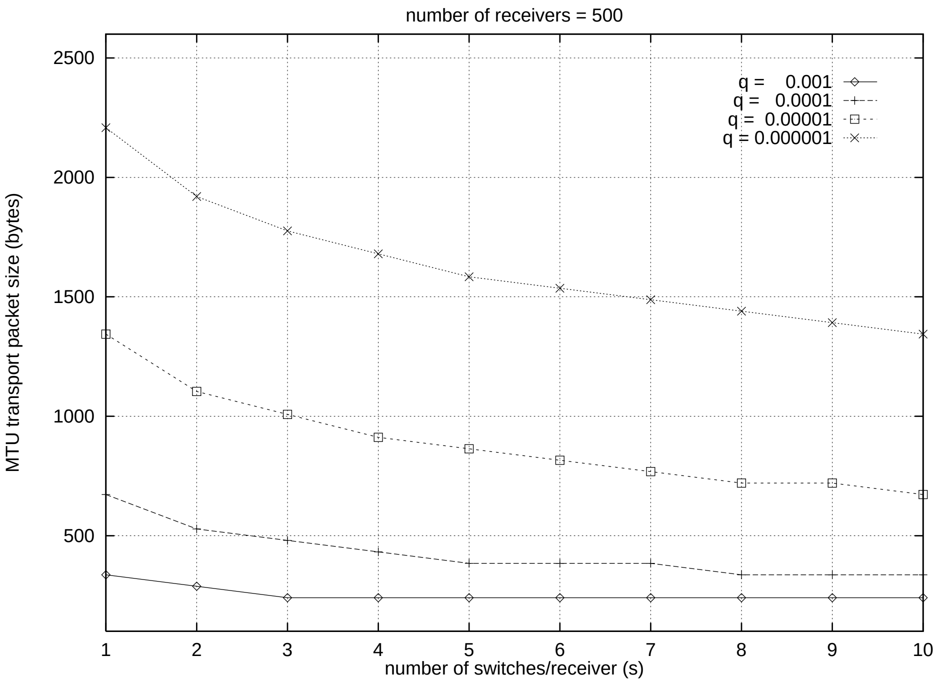

Figure 11: Optimum PDU vs. Number of Switches per Receiver