



Packet Spacing: An Enabling Mechanism for Delivering Multimedia Content in Computational Grids*

ANNETTE C. FENG
APU C. KAPADIA

afeng@lanl.gov
akapadia@cs.uiuc.edu

Los Alamos National Laboratory, Los Alamos, NM 87545, and Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801

WU-CHUN FENG

feng@lanl.gov

Los Alamos National Laboratory, Los Alamos, NM 87545, and Department of Computer & Information Science, The Ohio State University, Columbus, OH 43210

GENEVA G. BELFORD

belford@cs.uiuc.edu

Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801

Abstract. Streaming multimedia with UDP has become increasingly popular over distributed systems like the Internet. Scientific applications that stream multimedia include remote computational steering of visualization data and video-on-demand teleconferencing over the Access Grid. However, UDP does not possess a self-regulating, congestion-control mechanism; and most best-effort traffic is served by congestion-controlled TCP. Consequently, UDP steals bandwidth from TCP such that TCP flows starve for network resources. With the volume of Internet traffic continuing to increase, the perpetuation of UDP-based streaming will cause the Internet to collapse as it did in the mid-1980's due to the use of non-congestion-controlled TCP.

To address this problem, we introduce the counter-intuitive notion of inter-packet spacing with control feedback to enable UDP-based applications to perform well in the next-generation Internet and computational grids. When compared with traditional UDP-based streaming, we illustrate that our approach can reduce packet loss over 50% without adversely affecting delivered throughput.

Keywords: network protocol, multimedia, packet spacing, streaming, TCP, UDP, rate-adjusting congestion control, computational grid, Access Grid

1. Introduction

TCP and UDP are the most widely-used transport protocols today, the TCP/IP protocol suite being the *de facto* standard in the Internet-computing environment. TCP enables reliable, bulk-data transfer; however, it is inappropriate for such tasks as live video-on-demand and remote computational steering of visualization data in

* This work was supported by the U.S. Dept. of Energy's Laboratory-Directed Research & Development Program and the U.S. Dept. of Energy's Office of Science through Los Alamos National Laboratory contract W-7405-ENG-36. Any opinions, findings, and conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DOE or Los Alamos National Laboratory. Los Alamos Unclassified Report (LA-UR) 01-0904.

computational grids. Bulk-data transfer requires 100% reliable communication, and hence, TCP. Video-on-demand and remote computational steering generally do not require 100% reliability, therefore, TCP is overkill. For instance, if a video frame is missing a small block of pixels due to a lost packet, the video application is better off displaying the virtually complete frame and moving on to the next frame instead of waiting for the re-transmission of the lost packet (which over the Internet could easily take 100 ms).¹ TCP, in this case, provides too much functionality because its loss detection and re-transmission mechanisms, being tightly integrated with TCP's congestion-control mechanism, are inherent functions of the protocol.

UDP, on the other hand, provides no reliability guarantees. Specifically, it provides best-effort, end-to-end service without performing loss detection and packet re-transmission and without performing congestion control. Because of this, UDP obtains more bandwidth than TCP, albeit at the risk of suffering packet loss and packet re-ordering, problems that can ultimately be addressed by the applications themselves. Therefore, multimedia applications such as RealPlayer [17, 18] and scientific applications such as remote data visualization use UDP in order to improve perceived performance. Because UDP does not self-regulate in response to network congestion, these UDP-based applications gobble up available network resources, stealing bandwidth away from well-behaved applications that use congestion-controlled TCP. An application that blasts UDP packets into the network can readily fill the buffers of an intermediate router, causing severe congestion and packet loss. Since TCP-based applications slow down their sending rates in response to congestion, these applications become starved for network resources as the UDP-based applications continue to blast their packets unchecked into the network and claim the bandwidth being made available to them. Even though sending hosts can inject UDP packets as quickly as they are able, the throughput can suffer dramatically due to heavy packet loss and increased delays as packets spend more time waiting in queues within the network.

A simple observation reveals that adequate throughput can be attained by spacing the packets apart instead of blasting them one right after the other into the network. The next section reveals this insight. The notion of slowing down the sending rate in order to achieve better throughput is certainly counter-intuitive; however, our experiments show the viability and effectiveness of this approach.

1.1. *Insight*

Based on our recent work in network traffic characterization [7, 8, 24], we observed significant packet loss even when the offered load was less than half of the available network bandwidth. An analysis of our *ns* [1] simulations revealed that this behavior was due to simultaneous bursts of traffic coming from client applications and overflowing the buffer space in the bottleneck router. Metaphorically, this could be viewed as what happens at a major highway interchange during rush hour where everyone wants to go home simultaneously at 5:00 p.m., thus "overflowing" the highway interchange. To avoid such a situation, some people self-regulate themselves by heading home at a different time, i.e., spacing themselves out from other people.

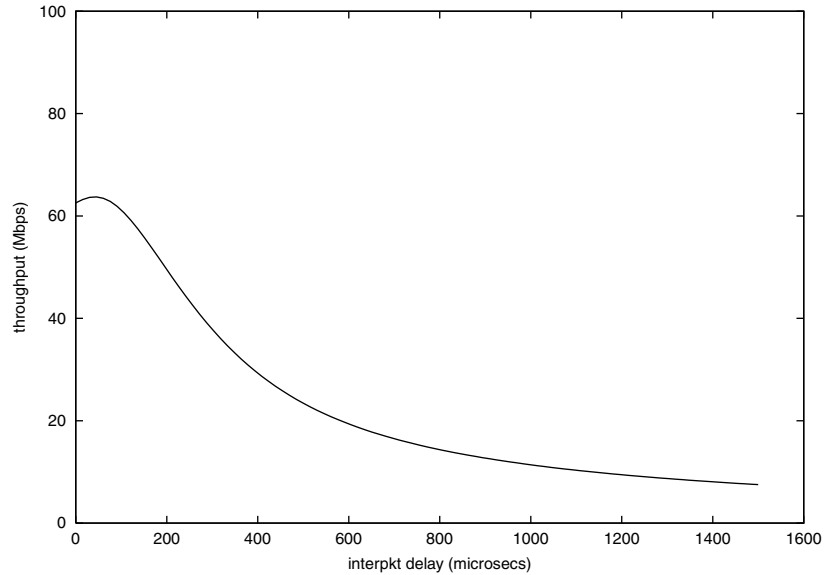


Figure 1. Delivered throughput to the receiver.

If we view vehicles as packets and the highway interchange as a router, then to avoid buffer overflow and enhance throughput, packets should not be blasted onto the network one after another. Instead, packets should be spaced out over time. To test this hypothesis, we ran live wide-area network (WAN) tests between Los Alamos National Laboratory (LANL), University of Illinois at Urbana-Champaign (UIUC), and Ohio State University (OSU). These tests consisted of sending UDP packets between LANL and either UIUC or OSU at different packet-spacing intervals. Figures 1 and 2 show the throughput and packet loss, respectively, of a representative test between LANL and UIUC [6]. When the packet spacing is zero, e.g., today's UDP-based multimedia-streaming applications, the throughput is 62 Mb/s but with a packet loss of almost 90%! With as little as 100 μ s of spacing between packets, the throughput remains the same, but the packet loss drops all the way down to 35%. And when the packet spacing is 50 μ s, the throughput is actually *higher* than when the packets are not spaced as in UDP-based multimedia streaming.

All curves from our other live WAN tests have the same general shape. That is, the throughput initially increases when the amount of packet spacing increases and then decreases exponentially as the amount of spacing increases further. The packet-loss percentage immediately decreases in an exponential manner as packet spacing increases.

1.2. Related work

Many transport protocols for the delivery of multimedia content certainly have been proposed, among them being XTP, RAP, and HPF. The Xpress Transport Protocol (XTP) [4] uses explicit rate control to combat congestion, however, the congestion-control mechanism must be implemented within the network and not simply at the

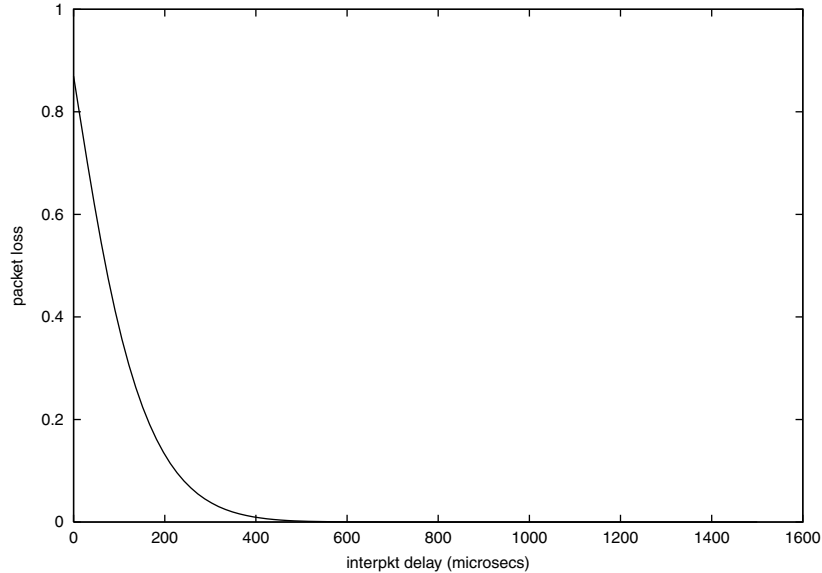


Figure 2. Packet-loss percentage.

edges. Fluctuating round-trip times (RTTs) cause poor performance because of a design feature whereby XTP enters a synchronizing handshake when a timer expires while XTP awaits a response to a request for information on missing data [3]. Furthermore, because it is a complex protocol, XTP is meant to be implemented in VLSI for performance reasons, so software implementations are too slow for multimedia traffic [22].

The Rate Adaptation Protocol (RAP) [19] is a TCP-friendly protocol that employs an “additive increase, multiplicative decrease” (AIMD) algorithm for rate adjustment. RAP is intended for the transmission of delay-sensitive, semi-reliable, rate-based applications which use layered-encoding of their data streams. RAP is therefore not a general solution but specifically targets layered-encoded multimedia content which it uses to adjust its transmission rate by adjusting the number of layers it sends.

The Heterogeneous Packet Flows (HPF) [12] protocol supports the delivery of packets having differing QoS requirements within a single stream. Addressing a design flaw of TCP, HPF decouples congestion control from reliability and uses a rate-based, AIMD approach to combat congestion. The problem with the AIMD approach (also used by RAP) is that such an approach will not scale to high-performance (or more precisely, high bandwidth-delay product) networks. For example, when the window size is one, a linear increase is a 100% increase. When the window size is 1000, a linear increase is a mere 0.1%. An absolute linear increase in window size from 500 to 1000 (as during TCP’s congestion-avoidance phase) will take 500 round-trip times to converge! More realistically, the situation is actually much worse. If we assume a typical WAN with a high bandwidth-delay product, i.e., 1 Gb/s WAN \times 100 ms RTT = 100 Mb, then for an uncongested network, the ubiquitously deployed TCP Reno continually increases its window size until it induces

packet loss (i.e., just after 100 Mb) and then chops its window size in half (i.e., 50 Mb). The re-convergence back to the “optimal window size” of 100 Mb using TCP’s absolute linear increase takes much too long and results in lowered network utilization. In this particular case, convergence can take as long as $(100 \text{ Mb} - 50 \text{ Mb}) / (1500 \text{ B/RTT} * 8 \text{ b/B}) = 4,168 \text{ RTTs}$ or $(4,168 \text{ RTTs} * 100 \text{ ms/RTT}) = 416.8 \text{ seconds} = 6.947 \text{ minutes!}$

In 1997, Mahdavi and Floyd [13] informally proposed the notion of equation-based congestion control for unicast applications. While the AIMD algorithm found in TCP backs off by cutting its sending rate in half in response to a single congestion indication, equation-based congestion control uses a control equation that more gradually and smoothly adapts its maximum rate because some real-time applications find that halving the sending rate is unnecessarily severe and can noticeably reduce the user-perceived quality [23]. Although the above work has given rise to a significant amount of research on equation-based and other types of congestion-control mechanisms [10, 16, 19 – 21, 23], we still do not have any deployable congestion-control mechanisms for best-effort streaming multimedia.

Previous work in packet spacing includes [2, 11]. In [11], Jain argues that rate-control protocols for congestion control may not work without the cooperation of intermediate routers because packets may get clumped together at the intermediate routers anyway. This would result in larger bursts at the intermediate routers even though the goal may have been to reduce the burstiness of the traffic. While this may have been true a decade ago, we believe that the boom of the world-wide web and other multimedia applications creates enough interleaving traffic to maintain packet spacing between end hosts. We will substantiate this belief in Section 3.2.3.

Aggarwal et al. [2] study the effect of uniform packet spacing (or “pacing”) over a round-trip time in TCP. While pacing results in better fairness, throughput, and lower drop rates in some cases, the throughput is worse than regular TCP most of the time because a paced-TCP is susceptible to synchronized losses and delays congestion notification. In contrast, we focus on the effects of packet spacing over UDP with control feedback rather than on TCP itself.

In general, our packet-spacing protocol differs from the above work in several ways. First, rather than focusing primarily on being compatible or *fair* with TCP, our rate-adjusting protocol addresses fairness while simultaneously delivering UDP-like bandwidth. Second, we accomplish the above feat by introducing the counterintuitive notion of packet spacing.

2. Approach

Packet spacing refers to the delay introduced between two consecutive packets, as shown in Figure 3. Here, t_s is the amount of spacing between packets, and t_x is the transmission time for each packet. By introducing such a delay, bursts of packets can be spaced out, resulting in fewer packet drops at intermediate routers and potentially *higher* throughput at the end host, as shown back in Figure 1. Thus, packet spacing can potentially be used as a mechanism to assist in congestion avoidance and control.

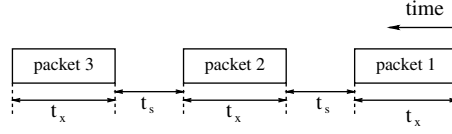


Figure 3. Packet spacing.

Based on Figure 1, the ideal operating region of our packet-spacing mechanism ranges from $50 \mu\text{s}$ to $500 \mu\text{s}$. No packet spacing or packet spacing of less than $50 \mu\text{s}$ results in very high packet loss with less delivered bandwidth than when the packet spacing is $50 \mu\text{s}$.

Depending on the application, the ideal packet-spacing range may be as small as $100 \mu\text{s}$ to $200 \mu\text{s}$ in order to get UDP-like bandwidth but with significantly less packet loss, e.g., at $200 \mu\text{s}$, bandwidth is 50 Mb/s while packet loss is only 10% , or as large as $400 \mu\text{s}$ to $500 \mu\text{s}$ to obtain TCP-like reliability but with higher throughput than TCP. To exploit this counterintuitive finding, we develop an ad-hoc packet-spacing protocol (PSP) to adjust the amount of packet spacing based on feedback from the network.²

2.1. Ad-hoc packet-spacing protocol

In our ad-hoc packet-spacing protocol (PSP),³ the sender initially transmits packets at the highest possible rate, i.e., no inter-packet spacing, and the receiver sends acknowledgments every round-trip time (RTT) for the packets it receives. (This RTT is the base propagation-delay time, not the dynamic RTT. To keep the protocol simple, we did not experiment with dynamic RTTs.)

We calculate the base RTT by performing *ping* during connection set-up.⁴ After the connection is established, the sender conveys the calculated RTT to the receiver by including it within the header of each packet. Note that this is not required after the first acknowledgment is received, but we have left this provision so that dynamic RTTs can be used in the future. Each acknowledgment contains the number of packets that were received in the previous RTT.

When the sender receives such acknowledgments, it compares the number of packets sent, p_{sent} , in the previous RTT to the number of packets received, p_{rcvd} . Based on the values of p_{sent} and p_{rcvd} , the sender adapts its packet spacing ps as shown in Figure 4.

```

if  $p_{sent} > p_{rcvd}$  (i.e., packets were lost) then
  /* sender must reduce its transmission rate */
  if  $ps = 0$  then
     $ps \leftarrow 50 \mu\text{s}$ 
  else
     $ps \leftarrow \min(ps * 2, RTT)$ 
  else /* sender tries to increase its sending rate */
     $ps \leftarrow ps - 2$ 

```

Figure 4. Ad-hoc packet-spacing protocol.

Because our WAN experiments and simulations showed that the ideal packet spacing occurred between 0 μ s and 2000 μ s, we chose an initial packet spacing of 50 μ s because (1) anything smaller generated significantly higher packet loss with no benefit with respect to throughput and (2) finding the ideal packet spacing within this range quickly would take no more than seven RTTs. Larger spacings can be reached in only a few more RTTs because the packet spacing increases exponentially.

The $ps \leftarrow \min(ps * 2, RTT)$ clause ensures that the maximum packet spacing is one RTT. That is, at least one packet is sent every RTT.

2.2. Damped packet-spacing protocol

Due to the opposing packet-spacing decisions in our ad-hoc PSP, initial tests of PSP resulted in large oscillations around the ideal sending rate. To address this problem, we added the following heuristic to damp the oscillations: *If a loss occurs due to a deliberate decrease in the packet spacing (and consequently, increase in rate), then the sender reverts to the previous packet-spacing value.* Using this heuristic, the sender makes significantly smaller oscillations around the ideal operating point, resulting in a 10% increase in overall throughput. Thus, in this paper, we will present experimental results for the damped PSP rather than the ad-hoc PSP.

2.3. Equation-based packet-spacing protocol

Although the purpose of the ad-hoc and damped PSPs is to demonstrate the benefits of packet spacing, these protocols are point-specific solutions that apply only to the tested topology shown in Figure 5. Furthermore, the halving of the sending rate in the ad-hoc and damped PSPs in response to a single congestion indication may be unnecessarily severe for real-time applications such as video and induces bursty traffic behavior. To address these problems, we discuss a general solution that adapts an equation-based, congestion-control mechanism [10] and incorporates packet spacing.

An application that uses a congestion-control mechanism that is more aggressive than TCP's could unfairly limit TCP traffic from getting its fair share of bandwidth [9] and vice versa. To achieve a proper balance, we select a formulation that is based on the steady-state sending rate of TCP [15]:

$$T = \frac{s}{R\sqrt{\frac{2p}{3}} + t_{RTO}(3\sqrt{\frac{3p}{8}})p(1 + 32p^2)}$$

This equation provides an upper bound on the sending rate T in bytes/sec, as a function of the packet size s , round-trip time R , steady-state loss event rate p , and the TCP retransmit timeout t_{RTO} .

While setting the rate according to the above equation ensures TCP-friendliness between the equation-based packet-spacing protocol and TCP, it will result in bursty

traffic behavior because TCP itself is bursty in nature. Thus, we set the sending rate not only based on the above equation but also on an exponentially weighted moving average of the number of packets between loss events [10] in order to smooth out changes in the sending rate.

3. Experiments

For our WAN simulations, we used *ns-2*, which is a network simulator developed by the VINT group [1]. We refer to senders and receivers as *agents*, which follows from the terminology used by *ns-2*.

Our packet-spacing agent (PSA) implements packet spacing *without* any congestion-control feedback. Because the damped PSP outperforms the ad-hoc PSP, we only implement and experiment with the former in a damped PSA. Lastly, our equation-based PSA implements the equation-based PSP.

3.1. Network topology

Figure 5 shows the network topology that we used in our experiments. The k nodes on the left (n_1, n_2, \dots, n_k) simulate senders on a local-area Ethernet, transmitting via a common gateway router (e.g., LAN/WAN gateway or n_{middle}) to a WAN backbone running at 155 Mb/s or OC-3; this topology models the LAN and WAN at Los Alamos National Laboratory. All the receivers are aggregated into the node n_{sink} . The gateway router has a buffer size of 10 packets, 100-Mb/s Ethernet links with 2-ms delays to the senders, and a 155-Mb/s link with 40-ms delay to the receivers. This delay is typical of the delay found in a transcontinental WAN connection.

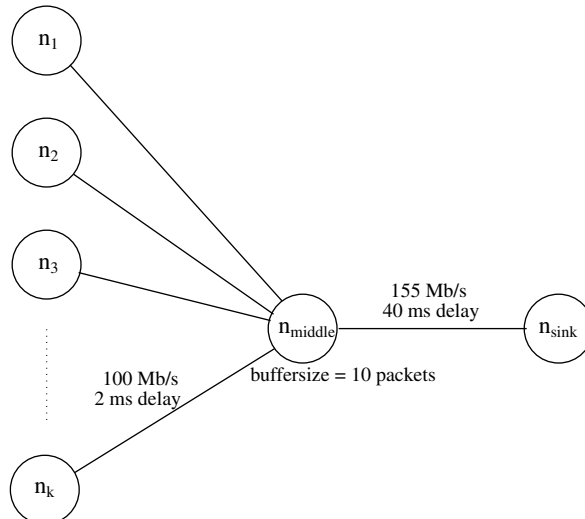


Figure 5. Topology for WAN simulations.

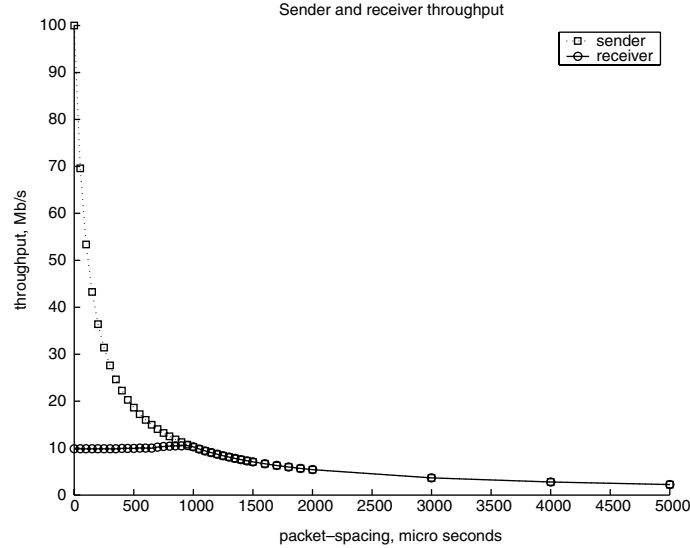


Figure 6. Throughput for one of the 16 PSAs.

3.2. PSA simulations

Here we study the behavior of competing PSAs and PSAs competing with TCP agents. Like Mo et al. [14] who compare TCP Reno and TCP Vegas using infinite file transfers, we use infinite file transfers for the TCP connections as well. (For the figures in this section, each data point in the simulation graphs represents the result of a 500-s simulation for a particular packet-spacing interval.)

3.2.1. PSAs competing. In this set of experiments, we ran simulations with 2, 4, 8, and 16 PSAs competing against each other, respectively. Figures 6 and 7 show the results for the last case. The resulting behavior is similar to what we observed in the actual WAN experiments (i.e., Figures 1 and 2). (Note that all the 16 competing PSAs showed similar behavior.)

In Figures 6 and 7, the region of interest occurs between 0 μ s and 1000 μ s. With a packet spacing of 0 μ s, the sender throughput is 100 Mb/s while the receiver-realized throughput is only a measly 10 Mb/s with a packet loss of 90%! As packet spacing increases, the packet-loss percentage drops sharply, and the throughput at the receiver actually *increases* to its maximum point at 1000 μ s of inter-packet spacing. This phenomenon is similar to what we found with our live WAN tests in Figures 1 and 2.

3.2.2. PSAs competing with TCP agents. In these experiments, we ran simulations with 1, 2, 4, 8, and 16 sender/receiver TCP pairs and an equal number of PSA pairs, respectively. Figures 8 and 9 show the behavior of one particular PSA competing with 15 other PSAs and 16 TCP connections. All other simulations resulted in similar behavior. Again, we see that the behavior is strikingly similar to that seen in

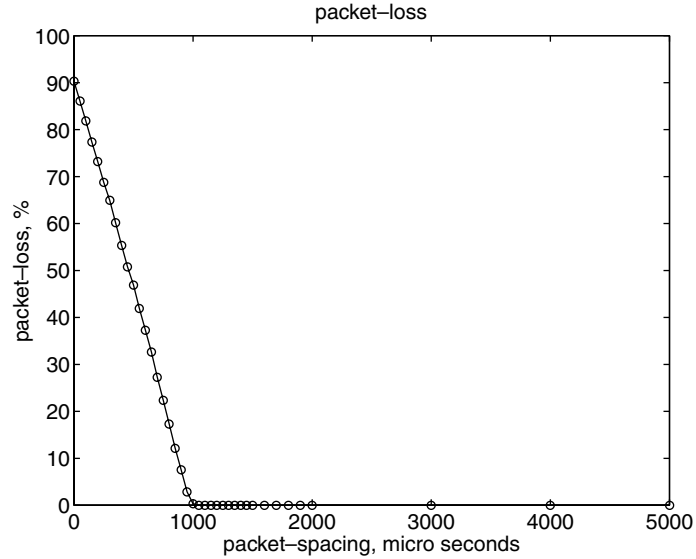


Figure 7. Packet loss for one of 16 PSAs.

the actual WAN experiments. The optimal performance of the PSAs with respect to throughput and packet loss occurs at $1000 \mu\text{s}$ to $1050 \mu\text{s}$, i.e., throughput is 11 Mb/s while packet loss is 0%.

Figures 10 and 11 show the throughput and packet loss, respectively, of a representative TCP connection with its buffers tuned to the bandwidth-delay product.

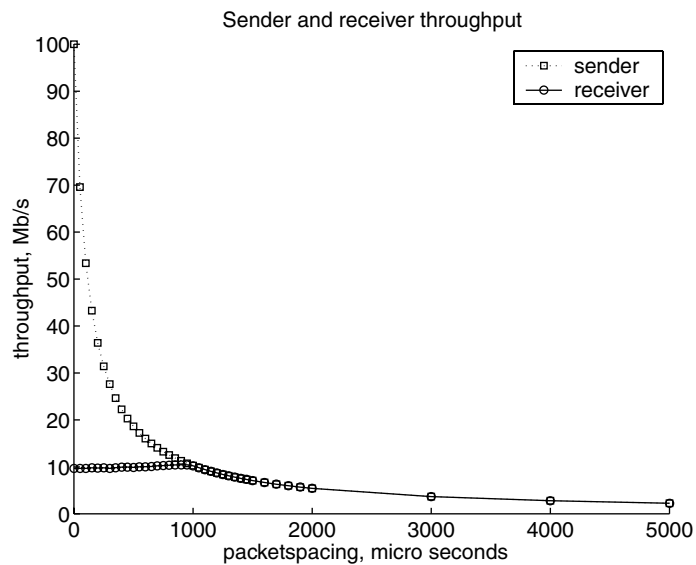


Figure 8. Throughput for one PSA of 16 PSAs and 16 TCP connections.

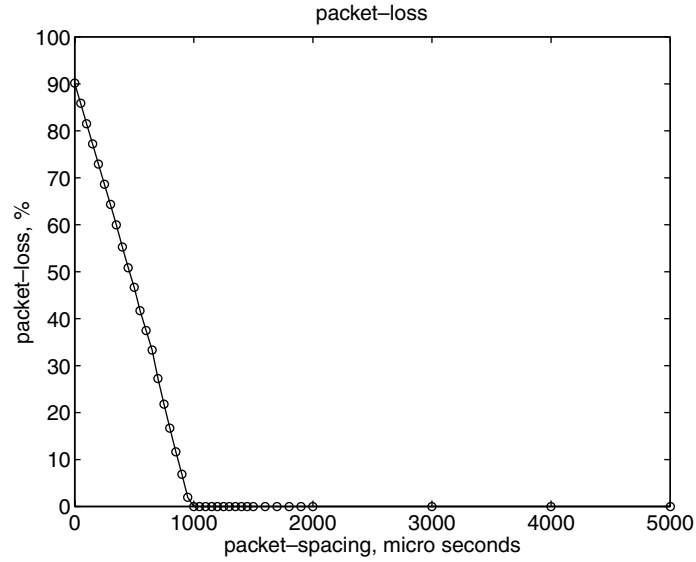


Figure 9. Packet loss for one PSA of 16 PSAs and 16 TCP connections.

These figures show that with sufficient spacing by the PSAs, a TCP connection can consume its share of available bandwidth. For example, Figures 8 and 10 illustrate that with 3000 μ s of packet spacing, each PSA receiver and TCP receiver achieves 5 Mb/s of throughput.

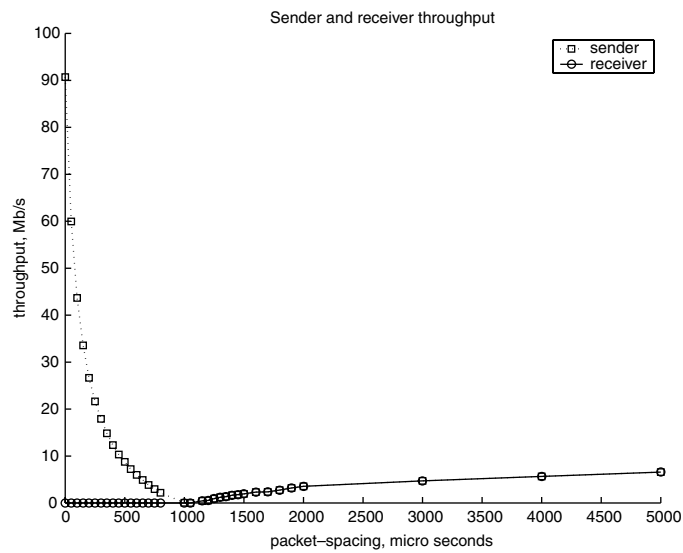


Figure 10. Throughput for one TCP of 16 TCP connections and 16 PSAs.

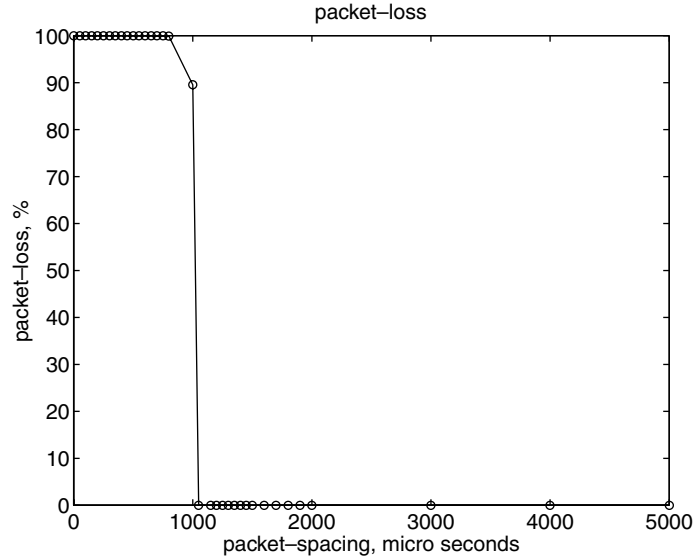


Figure 11. Packet loss for one TCP of 16 TCP connections and 16 PSAs.

3.2.3. PSA spacing at the receiver. Assuming that applications create enough interleaving traffic to maintain packet spacing between communicating hosts rather than getting clumped as claimed by [11], we recorded the inter-arrival time of packets at one PSA receiver, using the same experimental set-up as described in Section 3.2.2. The sending PSAs used a spacing of $1500 \mu\text{s}$; the resulting inter-packet spacings at the receiver averaged $1540.6 \mu\text{s}$ with a standard deviation of $64.75 \mu\text{s}$.

3.3. Adaptive PSA simulations

First, we demonstrate how the damped PSAs try to find the ideal packet spacing under varying network conditions. Next, we show how the equation-based PSAs accomplish the same objectives as the damped PSAs, e.g., TCP-friendliness or fairness, while simultaneously providing a smoother sending rate than the damped PSAs. This smoother sending rate provides better support for time-sensitive applications such as video and audio streaming.

3.3.1. Damped PSAs competing. When two damped PSAs compete for bandwidth, each damped PSA makes small oscillations around the ideal sending rate of 550 packets/RTT. Thus, each damped PSA gets its fair share of bandwidth. And even when the damped PSAs are started at different times, the sending rates quickly converge to 550 packets/RTT.

3.3.2. Damped PSAs competing with TCP. In this simulation, we ran 10 TCP connections with infinite file transfers in the background and then added two damped

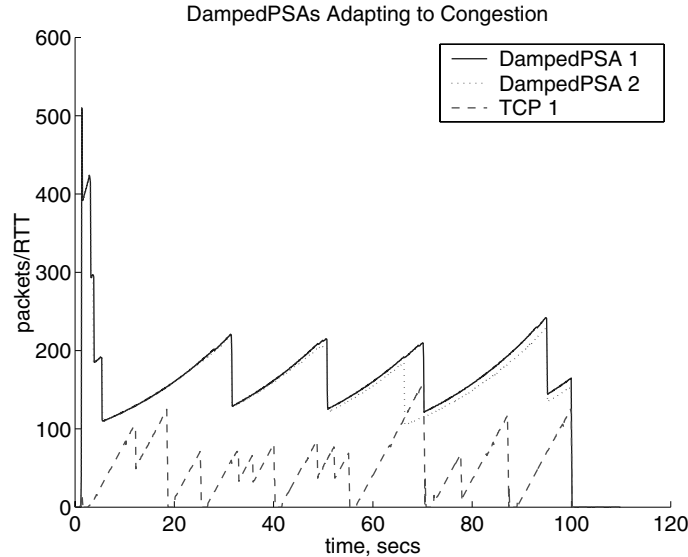


Figure 12. Two damped PSAs competing with ten TCPs.

PSAs to compete for bandwidth. Figure 12 shows that the damped PSAs respond readily to congestion. And again, both damped PSAs have very similar sending rates.

However, the damped PSAs suffer from two problems. First, due to the “additive increase/multiplicative decrease” nature of our damped PSP, each damped PSA sees a significant rate change (i.e., halving of its sending rate) when a single congestion-indication event occurs. And unfortunately, for real-time applications such as video streaming, such a drastic change in sending rate is unnecessarily severe as it can noticeably reduce the user-perceived quality [23]. Second, while the damped PSAs are certainly “TCP-friendlier” than UDP packet blasting, the damped PSAs manage to claim approximately twice as much bandwidth as a representative TCP connection.

3.3.3. Equation-based PSAs competing with TCP. Similar to the previous section, we ran 10 TCP connections with infinite file transfers in the background and then added two equation-based PSAs to compete for bandwidth. In stark contrast to Figure 12, Figure 13 demonstrates that the equation-based PSAs produce a much smoother and fairer sending rate that is appropriate for multimedia applications. Furthermore, the packet-spaced sending rate of the equation-based PSAs is more TCP-friendly than the that of the damped PSAs or standard UDP packet blasting.

4. Conclusion

Perhaps the most interesting result in this paper is that a receiver’s realizable throughput actually *increases* (up to a point) even when the sender’s transmission rate decreases. This result has dramatic implications on many of today’s multimedia

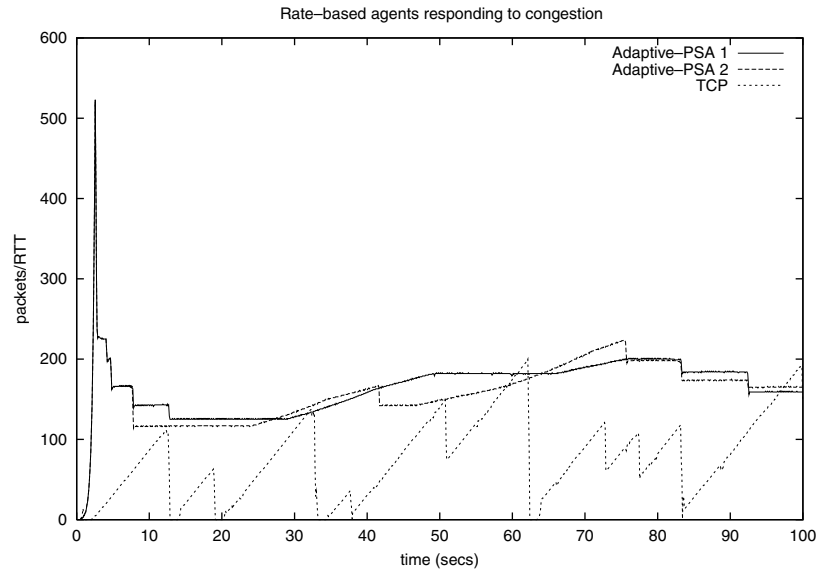


Figure 13. Two equation-based PSAs competing with ten TCPs.

applications that blast packets onto the network as fast as possible, i.e., no packet spacing. By slowing down the introduction of packets into the network, congestion is alleviated at the intermediate routers; this, in turn, results in a net increase in throughput. Thus, this work provides an incentive for multimedia providers not to blast UDP packets indiscriminately into the network. In addition, it provides motivation for the deployment of a packet-spaced protocol that can deliver UDP-like performance yet still be responsive to competing connections, particularly for applications with multimedia streaming such as the Access Grid [5].

Our damped packet-spacing protocol (PSP) sends data near its “optimal” sending rate by using a simple feedback mechanism that reports packet loss every RTT. This mechanism in turn controls the amount of packet spacing. These initial results demonstrate that by introducing packet spacing to a multimedia stream, packet loss can be reduced dramatically while still maintaining decent (and relatively stable) throughput for multimedia applications.

To address the burstiness of the damped PSP and its topology-specific congestion-control algorithm, the equation-based PSP provides a more general mechanism for congestion control. The equation-based PSP achieves fairness with TCP connections and results in very low packet loss and a smoother rate that is more appropriate for multimedia applications.

Future work includes examining the performance of the equation-based PSP with different types of application traffic and over a live WAN. Of particular interest are those applications that generate data in short bursts with relatively large intervals between bursts. Based on the experimental results presented here, we expect that the packet loss that would normally be induced by these bursts to be greatly reduced.

Notes

1. If the required frame rate is 30 frames per second, then the interframe delay is only 33 ms. Therefore, a re-transmission delay of 100 ms over the wide-area network is clearly unacceptable.
2. We note that at the present time, the feedback is only used for adjusting the packet spacing and that no retransmissions are done at this time.
3. We refer to this protocol as being “ad-hoc” because it is a point-specific solution that only applies to the tested topology and serves only to demonstrate the benefits of packet spacing. It is not a general solution.
4. A more sophisticated mechanism could be developed to get a better estimate of the RTT. However, for the purposes of our experiments, we only needed a value that was reasonable enough to provide timely feedback.

References

1. ‘UCB/LBNL/VINT network simulator—ns (version 2)’. <http://www.isi.edu/nsnam/vint/index.html>.
2. A. Aggarwal, S. Savage, and T. Anderson. Understanding the performance of TCP pacing. In *Proceedings of IEEE INFOCOM*, 2000.
3. J. W. Atwood and G. C. K. Chung. Error control in the Xpress transfer protocol. In *Proceedings of 18th Conference on Local Computer Networks*, pp. 423–431, 1993.
4. Y. Baguette and A. Danthine. Comparison of TP4, TCP and XTP—Part 2: Data transfer mechanisms. *ETT*, 3(5), 1992.
5. L. Childers, T. Disz, R. Olson, M. E. Papka, R. Stevens, and T. Udeshi. Access grid: Immersive group-to-group collaborative visualization. In *Proceedings of the 4th International Immersive Projection Technology Workshop*, 2000.
6. A. Feng, RAPID: Rate-adjusting protocol for internet delivery. Thesis proposal (in progress), University of Illinois at Urbana-Champaign, 2001.
7. W. Feng and P. Tinnakornsriruphap. The adverse impact of the TCP congestion-control mechanism in heterogeneous computing systems. In *Proceedings of ICPP’00*, 2000a.
8. W. Feng and P. Tinnakornsriruphap. The failure of TCP in distributed computational grids. In *Proceedings of SC’00*, 2000b.
9. S. Floyd and K. Fall. Promoting the use of end-to-end congestion control. *IEEE/ACM Transactions on Networking*, 1999.
10. S. Floyd, M. Handley, J. Padhye, and J. Widmer. Equation-based congestion control for unicast applications. In *Proc. of SIGCOMM*, 2000.
11. R. Jain. Myths about congestion management in high speed networks. *Internetworking: Research and Experience*, 29(2):101–113, 1992.
12. J. Li, S. Ha, and V. Bharghavan. HPF: A transport protocol for supporting heregeneous packet flows in the Internet, 1999.
13. J. Mahdavi and S. Floyd. TCP-friendly unicast rate-based flow control. Technical report, Note sent to end2end-interest mailing list, 1997.
14. J. Mo, J. Walrand, and V. Anantharam. Analysis and comparison of TCP Reno and Vegas. In *Proceedings of IEEE INFOCOM*, 1999.
15. J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP throughput: A simple model and its empirical validation. In *Proceedings of SIGCOMM’98*, 1998.
16. J. Padhye, J. Kurose, D. Towsley, and R. Koodli. A model based TCP-friendly rate control protocol. In *Proceedings of NOSSDAV’99*, 1999.
17. RealNetworks Inc. RBN™ (Real Broadcast Network) White Paper. Available at <http://www.real.com/solutions/-rbn/whitepaper.html>, 1999a.
18. RealNetworks Inc. RealVideo Technical White Paper. Available at <http://www.real.com/devzone/library/whitepapers/-overview.html>, 1999b.
19. R. Rejaie, M. Handley, and D. Estrin. An end-to-end rate-based congestion control mechanism for real-time streams in the Internet. In *Proceedings of INFOCOM’99*, 1999.

20. I. Rhee, V. Ozdemir, and Y. Yi. TEAR: TCP emulation at receivers—Flow control for multimedia streaming. Technical report, North Carolina State University, 2000.
21. D. Sisalem and H. Schulzrinne. The loss-delay based adjustment algorithm: A TCP-friendly adaptation scheme. In *Proceedings of NOSSDAV'98*, 1998.
22. W. T. Strayer, S. Gray, and R. E. C. Jr. An object-oriented implementation of the Xpress transfer protocol. In *Multimedia: Advanced Teleservices and High-Speed Communication Architectures, 2nd International Workshop (IWACA'94)*, pp. 387–400, 1994.
23. D. Tan and A. Zakhor. Real-time internet video using error resilient scalable compression and TCP-friendly transport protocol. *IEEE Transactions on Multimedia*, 1999.
24. P. Tinnakornsisuphap, W. Feng, and I. Philp. On the burstiness of the TCP congestion-control mechanism in a distributed computing system. In *Proceedings of ICDCS'00*, 2000.