THESIS FOR THE DEGREE OF LICENTIATE OF COMPUTER SCIENCE

A Study of Partially Reliable Transport Protocols for Soft Real-Time Applications

KARL-JOHAN GRINNEMO

Department of Computer Science
KARLSTAD UNIVERSITY
Karlstad, Sweden 2002

A Study of Partially Reliable Transport Protocols for Soft Real-Time Applications

KARL-JOHAN GRINNEMO

Technical Report no. 2002:36

Department of Computer Science
Karlstad University
SE–651 88 Karlstad, Sweden
Phone: +46 (0)54–700 1000

Contact information:
Karl-Johan Grinnemo
Department of Computer Science
Karlstad University
Universitetsg. 2
SE–651 88 Karlstad, Sweden

Phone: +46 (0)54–700 1826
Fax: +46 (0)54–700 14 60
Email: karl-johan.grinnemo@kau.se

Printed in Sweden
Karlstads Universitetstryckeri
Karlstad, Sweden 2002

A Study of Partially Reliable Transport Protocols for Soft Real-Time Applications

KARL-JOHAN GRINNEMO
Department of Computer Science, Karlstad University

Abstract

The profileration of multimedia applications, such as streaming video, teleconferencing, and interactive gaming has created a tremendous challenge for the traditional transport protocols of the Internet – UDP and TCP. Specifically, many multimedia applications are examples of soft real-time applications. They have often relatively stringent requirements in terms of delay and delay jitter, but typically tolerate a limited packet loss rate.

In recognition of the transport service requirements of soft real-time applications, this thesis studies the feasibility of using retransmission based, partially reliable transport protocols for these applications. The thesis studies ways of designing retransmission based, partially reliable transport protocols that are congestion aware and TCP compatible. Furthermore, the transport protocols should provide a service that, in terms of performance metrics such as throughput, delay, and delay jitter, are suitable for soft real-time applications. The thesis work comprises the design, analysis, and evaluation of two retransmission based, partially reliable transport protocols: PRTP and PRTP-ECN. Extensive simulations have been carried out on PRTP as well as PRTP-ECN. These simulations have in part been complemented by some theoretical analysis. The results of the simulations and the analysis suggest that substantial reductions in delay jitter and improvements in throughput can indeed be obtained with both PRTP and PRTP-ECN as compared to TCP. While PRTP reacted too slowly to congestion to be TCP-friendly and altogether fair, PRTP-ECN was found to be both TCP-friendly and reasonably fair.

The thesis work also comprises an extensive survey on retransmission based, partially reliable transport protocols. Based on this survey, we have proposed a taxonomy for these protocols. The taxonomy considers two dimensions of retransmission based, partially reliable transport protocols: the transport service, and the error control scheme.

Keywords: Transport protocol, Internet, partial reliability, retransmission-based, real-time, simulation, taxonomy, survey

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my employer, Ericsson AB in Karlstad. Without their financial support, this work would not have been possible. I would also like to express my sincere gratitude to my advisor, Dr. Anna Brunström, for her guidance and support, which made the completion of this work possible, and which made me grow as a scholar.

Special thanks, goes to Jacob Ängeby at Ericsson AB i Karlstad for his encouragement, and our discussions about research in general, and flow and congestion control in particular. Others at Ericsson AB in Karlstad which I am very much in debt to are Lena Larsson and Staffan Barrefors for recognizing the need for education and long-term research strategies; and Patrik Forslund, Gunnar Lorentzon, Magnus Larsson, and Sören Torstensson for their interest in my research.

Many thanks also goes to my colleagues at the Department of Computer Science at Karlstad University, especially the members of the DISCO (Distributed Systems and Communications) Research Group. You have through numerous discussions both ad-hoc, and at meetings and workgroups, contributed to my research. Furthermore, I would like to send my thanks to Rikard-Ed Svensson at the Department of Electrical Engineering, my "lunch-and coffee-break" partner, for his morale and personal support.

Last, but certainly not least, I wish to thank my family. My parents Göran and Ingrid Grinnemo for their support: They have always encouraged me in my education. If it had not been for them, I would not have come as far as I have today. Thanks also goes to my brother, Karl-Henrik Grinnemo, and my nephew, Felix Grinnemo.

List of Appended Papers

This thesis consists of an introductory summary and reprints of the following five papers:

Paper I: K-J Grinnemo, A. Brunstrom, and J. Garcia. *A Taxonomy and Survey of Retransmission Based Partially Reliable Transport Protocols*. Karlstad University Studies 2002:34, Karlstad University, Sweden, October 2002.

Paper II: K-J Grinnemo, A. Brunstrom. *A Simulation Based Performance Evaluation of PRTP*. Karlstad University Studies 2002:35, Karlstad University, Sweden, October 2002.

Paper III: K-J Grinnemo, A. Brunstrom. Evaluation of the QoS offered by PRTP-ECN – A TCP Compliant Partially Reliable Transport Protocol. In *9th International Workshop on QoS (IWQoS)*, pages 217–231, Karlsruhe, Germany, June 2001.

Paper IV: K-J Grinnemo, A. Brunstrom. A Simulation Based Performance Analysis of a TCP Extension for Best Effort Multimedia Applications. In *35th Annual Simulation Symposium (ANSS35)*, pages 327–336, San Diego, USA, April 2002.

Paper V: K-J Grinnemo, A. Brunstrom. Enhancing TCP for Applications with Soft Real-Time Constraints. In *The Convergence of Information Technologies and Communications (ITCom)*, Volume 4518, pages 18–31, Denver, USA, August 2001.

Minor editorial changes have been made to some of the papers.

Contents

Introductory Summary

1 Introduction

The origins of the Internet goes back to the experimentation of packet switching conducted by Baran et al. [3] at RAND in the early sixties. The primary objective was to design a network that would withstand a military attack, i.e., would continue to function even if a number of network nodes went down. As a consequence of this, the Internet protocols were mainly designed to provide for a dynamic and robust network, and not, as was the case with the telecommunication network, a particular real-time transport service. In particular, the two standard Internet transport protocols, UDP [22] and TCP [23], were designed to provide an unreliable, best effort transport service and a reliable transport service, respectively.

Traditionally, Internet has been used for applications such as remote login, e-mail, and file transfer, i.e., applications whose service requirements mesh well with the transport services offered by UDP and TCP. However, in the last decade, this has started to change. The technological advances have revolutionized communication systems. The speed and capacity of various components in a communication system, such as transmission media, switches, processors, and memory, have grown linearly, and sometimes even exponentially. This has introduced opportunities for new applications in the Internet such as video broadcasting and conferencing, voice, distributed interactive games, medical imaging, to name a few.

Furthermore, the market calls for a convergence in the communications industry. They argue against a panoply of networks to meet all diverse communication requirements, especially considering the fact that more or less all communication today employ digital techniques. That is, at the heart of most of today's communication networks is a collection of flows of bits, of 1s and 0s, which in many ways are indistinguishable from each other.

Although, discussions about a ubiquitous service platform already took place when the Integrated Services Digital Network (ISDN) architecture was introduced in the eighties, and also when Broadband ISDN (B-ISDN) and Asynchronous Transfer Mode (ATM) seemed to be on everyones agenda in the beginning of the nineties, it is, in at least one way, different this time. Back then, there was a major disagreement between the data- and telecommunication industry: the major proponents of ISDN and B-ISDN came from the telecommunication sector, while the datacommunications sector were more skeptical. However, this time there appears to be a consensus among the two sectors: both camps seem to have agreed that if there is to be a convergence, the platform for it will be the Internet.

It is widely agreed that one of the major hurdles to overcome in order for Internet to become a ubiquitous service platform is to provide acceptable services for real-time applications. That is, to design services that are able to meet the stringent performance requirements of real-time applications in terms of throughput, delay, delay jitter, and reliability. This problem has essentially been addressed in two ways: through the use of integrated services (IS) architectures such as IntServ [6] and DiffServ [7], and through novel transport protocols or transport protocol frameworks such as RTP [25].

From a real-time service perspective, the two approaches target somewhat different types of applications. The principal idea behind the IS architectures is to provide something like virtual circuits in the Internet. They could be seen as primarily targeting isochronous applications, i.e., applications which only tolerate very small deviations from their service requirements. In contrast, the second approach mainly involves changing the transport protocols used by the end nodes. It entails no or only a few changes to the Internet network infrastructure. Most notably, the second approach does not include any service provisioning in the interior nodes in the Internet, and is therefore not suitable for isochronous real-time applications. However, since the second approach is less costly and easier to deploy, it is still considered an attractive solution for soft real-time applications, i.e., applications that are more lenient against violations of the service requirements.

An example of the second approach is the so called partially reliable transport protocols. These protocols are characterized by providing a transport service which in terms of reliability places itself in between the services provided by UDP and TCP, i.e., a service which is more reliable than UDP but less reliable than TCP. In other words a transport service which tolerates a limited amount of packet loss. Specifically, these protocols fill the gap between the unreliable and completely reliable transport protocols by permitting an application to specify a controlled level of loss, and then enhance the service provided by the network level only enough to guarantee this loss level. The reason partially reliable protocols are considered attractive for soft real-time applications is that they explicitly provide for the applications to make trade-offs between the performance parameters reliability, delay, and delay jitter.

Essentially, the class of partially reliable transport protocols breaks down into two subclasses: open-loop partially reliable transport protocols and closed-loop partially reliable transport protocols. The subclass of open-loop protocols comprises those protocols which do not employ feedback from the network or end-node(s) when they perform error recovery. Typically open-loop protocols use Forward Error Correction [16, 18] (FEC) or channel coding [15] techniques to recover from corrupt and lost packets. Contrary to open-loop protocols, closed-loop protocols do employ feedback from the network when they perform error recovery. While open-loop partially reliable transport protocols have been extensively studied in the last 20 years [4, 5], the interest for closed-loop protocols only goes back to the beginning of the 1990s [10, 20].

A contributing factor to the previous uninterest for closed-loop partially reliable transport protocols was that they were not considered feasible for soft real-time multimedia applications. However, this view was revised when Dempsey et al. [10], Papadopoulos and Parulkar [20], and others demonstrated the feasibility of using closed-loop partially reliable transport protocols also for this kind of applications. In particular, they demonstrated that retransmission based, partially reliable transport protocols, i.e., those closed-loop protocols which perform error recovery through retransmissions, are in fact suitable for some classes of continuous media applications.

Our research is a continuation of the research on retransmission based, partially reliable transport protocols started by Dempsey et al. In line with their work, we study retransmission based, partially reliable transport protocols in connection with soft real-time applications, and then, with an emphasis on multimedia communication. Since we do not consider it feasible, at least not in the short term, to introduce an altogether new transport protocol in the Internet, our research is mainly concerned with extensions for partial reliability to existing transport protocols. In that way, we are able to introduce partially reliable transport services into the Internet, without imposing any changes to the existing Internet infrastructure.

Although many other researchers have made significant contributions to this research area [1, 9, 11, 12, 17, 21], we only know of one research group [8] that have considered retransmission based, partially reliable transport protocols in terms of extensions to the standard transport protocols. In fact, the majority of researchers seem to have been far more concerned with finding the most performance efficient retransmission based, partially reliable transport protocol than looking for solutions that are compatible with the existing Internet.

2 Research Statement

The principal objective of our research is to design transport protocols suitable for soft real-time applications, especially multimedia applications, in the Internet. As part of our research, we study the feasibility of using retransmission based, partially reliable transport protocols for this kind of applications. Specifically, we are studying ways of designing retransmission based, partially reliable transport protocols that

- provide a transport service that, in terms of performance metrics such as throughput, delay, and delay jitter, are suitable for soft real-time applications;
- are congestion aware, and ideally react to congestion in a TCP-friendly and fair manner;
- are compatible with the existing Internet transport protocols, and only involve small modifications to the existing Internet infrastructure.

3 Scope and Contributions of the Thesis

Our research has primarily involved the design, analysis, and evaluation of two particular retransmission based partially reliable transport protocols: PRTP (Partially Reliable Transport Protocol) and PRTP-ECN (Partially Reliable Transport Protocol using Explicit Congestion Notification). These protocols are compatible with TCP, and primarily differ from TCP in the way they make retransmissions: the application atop

PRTP/PRTP-ECN specifies its minimum acceptable reliability level and PRTP/PRTP-ECN only makes the number of retransmissions sufficient to uphold the specified reliability level. We have performed extensive simulation experiments on both PRTP and PRTP-ECN, and compared their performance with TCP in terms of performance metrics of major importance to soft real-time applications.

In the simulation experiments performed on PRTP, we evaluated the performance of PRTP compared to TCP for long-lived as well as short-lived connections, i.e., the stationary as well as the transient behavior of PRTP were studied. The stationary analysis comprised three performance metrics, average interarrival jitter, average throughput, and average fairness, while we in the transient analysis only studied throughput. Furthermore, in the stationary analysis, we also tested if PRTP exhibited a TCP-friendly behavior.

The results of the simulation experiments showed that PRTP indeed is likely to give substantial improvement in average throughput both for long-lived and short-lived connections, and that significant reduction in average interarrival jitter can be obtained for long-lived connections. However, the fairness and TCP-friendliness tests conducted on PRTP in the stationary analysis gave at hand that PRTP is not altogether fair and TCP-friendly. As a consequence of this, PRTP was modified to react more appropriately to incipient congestion. In particular PRTP was modified to use parts of the explicit congestion notification (ECN) mechanism proposed by IETF [24] to better signal congestion. The modified version of PRTP was called PRTP-ECN.

The same simulation experiments as were performed on PRTP were also performed on PRTP-ECN. While the improvement in average throughput and reduction in average interarrival jitter were not as large for PRTP-ECN as for PRTP, they were still substantial. Furthermore, in contrast to PRTP, PRTP-ECN showed a TCP-friendly behavior and was reasonably fair against contending TCP flows.

The simulation experiments on PRTP-ECN were complemented with a theoretical analysis. Both the stationary and the transient behavior of PRTP-ECN were theoretically analyzed. Specifically, we derived analytical expressions for the packet loss tolerance of PRTP-ECN at startup and in steady-state. We also analytically established the tolerance of PRTP-ECN for packet loss bursts.

In addition to the work on PRTP and PRTP-ECN, we have done an extensive survey of retransmission based, partially reliable transport protocols. On the basis of this survey, we have proposed a taxonomy. Apart from serving as a framework for classification, the taxonomy articulates the principal components of retransmission based, partially reliable transport protocols; introduces a uniform terminology; and emphasizes those aspects of retransmission based, partially reliable transport protocols that need further research.

To summarize, the main contributions of our research are as follows:

- We have demonstrated the feasibility of designing TCP compatible, retransmission based, partially reliable transport protocols that are congestion aware, and that, in terms of throughput and interarrival jitter, are more suitable for soft real-time applications than TCP.

- We have made predictions, both analytically and through simulations, of the performance improvements that can be obtained by using retransmission based, partially reliable transport protocols for soft real-time applications in the Internet.

- We have presented a survey and a taxonomy of retransmission based partially reliable transport protocols.

4 Thesis Outline

This thesis is a compilation of five papers. The five papers are listed below, together with short synopses about their content.

Paper I: A Taxonomy and Survey of Retransmission Based Partially Reliable Transport Protocols

A taxonomy for retransmission based, partially reliable transport protocols is presented. The taxonomy consists of two classification schemes. The first classification scheme classifies retransmission based, partially reliable transport protocols with respect to their offered reliability service, and the second classification scheme with respect to their error control scheme.

Furthermore, a survey of retransmission based, partially reliable transport protocols is presented. The survey primarily focuses on the services offered by the protocols; how they have been realized; and how they map into the taxonomy.

Taken together, the taxonomy and survey not only provides the foundation for retransmission based, partially reliable transport protocols, they also serve as an introduction to this class of transport protocols.

Paper II: A Simulation Based Performance Evaluation of PRTP

As mentioned in Section 3, our research has primarily concerned the design, analysis, and evaluation of two particular retransmission based, partially reliable transport protocols: PRTP and PRTP-ECN. This paper introduces PRTP. The design and the core principles of the protocol are described. However, the largest portion of the paper is devoted to an extensive simulation study of PRTP.

The simulation study of PRTP comprised three simulation experiments. In the first simulation experiment our simulation model of PRTP was validated against a prototype of PRTP that we have implemented in Linux [2]. The second simulation experiment evaluated the stationary performance of PRTP compared to TCP. In particular, we considered the performance of PRTP for long-lived connections in terms of average interarrival jitter, average throughput, and average fairness. We also investigated whether or not PRTP is TCP-friendly. Finally, in the third simulation experiment the transient performance of PRTP compared to TCP was evaluated. Specifically, the third simulation experiment considered the throughput performance of PRTP in a typical Web browsing scenario. Since, the throughput obtained in Web browsing is very much dependent on the type of Internet connection, three types of connections were studied: fixed, modem, and GSM.

Paper III: Evaluation of the QoS Offered by PRTP-ECN – A TCP Compliant Partially Reliable Transport Protocol

The simulations presented in Paper II found PRTP to be TCP-unfriendly and not altogether fair. To address this, PRTP-ECN was conceived. This paper considers PRTP-ECN: its design and the principal ideas behind the protocol.

The stationary performance of PRTP-ECN was evaluated using the same simulation testbed as was used in the stationary analysis of PRTP (see Paper II). This paper gives a detailed description of the stationary analysis of PRTP-ECN. Specifically, it evaluates the stationary performance of PRTP-ECN compared to TCP in terms of average interarrival jitter, average throughput, average goodput, average fairness, and TCP-friendliness.

Paper IV: A Simulation Based Performance Analysis of a TCP Extension for Best Effort Multimedia Applications

In the same way as Paper III, this paper considers the stationary performance analysis of PRTP-ECN. However, in this paper the focus is on the statistical design and analysis of the simulation experiment.

The simulation experiment was designed as a series of factorial experiments, one for each studied performance metric. This paper elaborates on the underlying effects model. Examples of issues discussed are model fitting, e.g., variance stabilizing transforms, and the statistical hypotheses tested. In addition to the performance metrics discussed in Paper III, this paper also considers the link utilization of PRTP-ECN as compared to TCP.

Paper V: Enhancing TCP for Applications with Soft Real-Time Constraints

While Papers III and IV only evaluated the performance of PRTP-ECN through simulations, the major contribution of this paper is a theoretical analysis of the transient and stationary behavior of PRTP-ECN. In particular, this paper presents analytical expressions for the packet loss tolerance of PRTP-ECN at startup; explicit formulae for the upper an lower bounds of the stationary packet loss tolerance of PRTP; and finally an expression for the maximum packet loss burst tolerated by PRTP-ECN in stationary state.

Although, the central theme of this paper is a theoretical evaluation of PRTP-ECN, it also provides a summary of the stationary performance analysis of PRTP-ECN as described in Paper III.

5 Future Work

As stated in Section 2, the primary objective of our research is to study Internet transport protocols that accommodate the service requirements of soft real-time applications. We have focused on retransmission based, partially reliable transport protocols. Thus far, we have, with PRTP and PRTP-ECN, demonstrated the feasibility of designing retransmission based, partially reliable transport protocols which like TCP are congestion aware, but which, in terms of throughput and jitter, are more suitable for soft real-time applications than TCP. However, both PRTP and PRTP-ECN builds upon TCP, an, in many respects, less than ideal platform for transport protocols that target soft real-time applications. Notably, TCP is a byte-oriented protocol, while many soft real-time applications, especially video and image applications, would benefit from a message-oriented protocol; TCP in itself provides no support for extensions for partial reliability; and furthermore TCP enforces an ordered delivery of packets which is unnecessarily stringent for many soft real-time applications.

In 2000, a new transport protocol was published by IETF – SCTP [27]. Originally, the rationale behind the development of SCTP was that a transport protocol was needed for transportation of signaling information in telecommunication networks. However, this narrow view of SCTP was later revised. Today, SCTP is considered a general purpose transport protocol, on equal footing with transport protocols such as UDP and TCP.

Contrary to TCP, SCTP is a message-oriented transport protocol and consequently keeps message boundaries; a framework for implementing partially reliable transport services on top of SCTP has already been suggested as an Internet draft [26]. Furthermore, SCTP provides for partially ordered message delivery, another service that, in addition to partial reliability, has been proven useful for soft real-time applications [9]. Taken together, we believe these features make SCTP a better platform than TCP for retransmission based, partially reliable transport protocols that target soft real-time applications.

As mentioned, the SCTP transport protocol was primarily designed to be used by signaling protocols in telecommunication networks. Specifically, SCTP was designed to be used by the application and user parts of the SS7 [13] (Signaling System #7) protocol stack. A common denominator of the majority of signaling protocols that will use SCTP is that they have soft real-time requirements. For example, the MTP-L3 [14] layer in the SS7 protocol stack requires response times within 500 to 1200 ms [19]. Failure to meet these response times will result in the initiation of error procedures for specific timers (e.g., timer T4 of ITU-T Q.704 [14]). Furthermore, MTP-L3 has very stringent demands on message loss and sequence errors.

Signaling traffic represents a new type of soft real-time application in the Internet, an application which up to now has attracted very little attention. Therefore, another strong candidate for future research is to study the performance of Internet based signaling traffic. We are particularly interested in studying the performance of the flow and congestion control of SCTP for signaling traffic. SCTP has to a large extent inherited its flow and congestion control mechanisms from TCP. However, the flow and congestion control mechanisms of TCP are most suitable for long-lived connections and less suitable for short-lived ones. The principal objective with our research would be to study extensions to SCTP that improve the performance of SCTP for short-lived signaling connections, while at the same time let SCTP remain a general purpose transport protocol.

References

[1] P. D. Amer, C. Chassot, T. Connolly, M. Diaz, and P. T. Conrad. Partial order transport service for multimedia and other applications. *ACM/IEEE Transactions on Networking*, 2(5), October 1994.

[2] K. Asplund, J. Garcia, A. Brunstrom, and S. Schneyer. Decreasing transfer delay through partial reliability. In *Protocols for Multimedia Systems (PROMS)*, Cracow, Poland, October 2000.

[3] P. Baran. On distributed communications. Memorandum 1–11, RAND, 1981.

[4] G. Barberis and D. Pazzaglia. Analysis and design of a packet-voice receiver. *IEEE Transactions on Communications*, 28(2):152–156, February 1981.

[5] V. Bhargava. Forward error correction schemes for digital communications. *IEEE Communications Magazine*, 21:11–19, January 1983.

[6] R. Braden, D. Clark, and S. Shenker. Integrated services in the internet architecture. RFC 1633, IETF, June 1994.

[7] M. Carlson, W. Weiss, S. Blake, Z. Wang, D. Black, and E. Davies. An architecture for differentiated services. RFC 2475, IETF, December 1998.

[8] T. Connolly, P. D. Amer, and P. T. Conrad. An extension to TCP: Partial order service. RFC 1693, IETF, November 1994.

[9] P. T. Conrad, E. Golden, P. D. Amer, and R. Marasli. A multimedia document retrieval system using partially-ordered/partially-reliable transport service. In *Multimedia Computing and Networking*, San Jose, USA, January 1996.

[10] B. Dempsey, T. Strayer, and A. Weaver. Adaptive error control for multimedia data transfers. In *International Workshop on Advanced Communications and Applications for High-Speed Networks (IWACA)*, pages 279–289, Munich, Germany, March 1992.

[11] M. Diaz, A. Lopez, C. Chassot, and P. D. Amer. Partial order connections: A new concept for high speed and multimedia services and protocols. *Annals of Telecomunications*, 49(5–6):270–281, 1994.

[12] F. Gong and G. Parulkar. An application-oriented error control scheme for high-speed networks. Technical Report WUCS-92-37, Washington University, November 1992.

[13] ITU-T. Introduction to CCITT Signalling System No. 7. Recommendation, ITU-T, April 1993.

[14] ITU-T. Signalling network functions and messages. Recommendation, ITU-T, July 1996.

[15] N. Jayant and P. Noll. *Digital Coding of Waveforms*. Prentice-Hall, Englewood Cliffs, New Jersey, 1984.

[16] G. Karlsson and M. Vetterli. Packet video and its integration into the network. *IEEE Journal on Selected Areas in Communications*, 7(3):739–751, June 1989.

[17] B. Mukherjee and T. Brecht. Time-lined TCP for the TCP-friendly delivery of streaming media. In *IEEE International Conference on Network Protocols (ICNP)*, pages 165–176, Osaka, Japan, November 2000.

[18] H. Ohta and T. Kitami. A cell loss recovery method using FEC in ATM networks. *IEEE Journal on Selected Areas in Communications*, 9(9):1471–1483, December 1991.

[19] L. Ong, I. Rytina, M. Garcia, H. Schwarzbauer, L. Coene, H. Lin, I. Juhasz, M. Holdrege, and C. Sharp. Framework architecture for signalling transport. RFC 2719, IETF, October 1999.

[20] C. Papadopoulos and G. Parulkar. Retransmission-based error control for continuous media applications. In *6th International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, pages 5–12, Zushi, Japan, April 1996.

[21] M. Piecuch, K. French, G. Oprica, and M. Claypool. A selective retransmission protocol for multimedia on the Internet. In *SPIE Multimedia Systems and Applications*, Boston, MA, USA, November 2000.

[22] J. Postel. User datagram protocol. RFC 768, IETF, August 1980.

[23] J. Postel. Transmission control protocol. RFC 793, IETF, September 1981.

[24] K. Ramakrishnan and S. Floyd. A proposal to add explicit congestion notification (ECN) to IP. RFC 2481, IETF, January 1999.

[25] H. Schulzrinne. RTP: A transport protocol for real-time applications. RFC 1889, IETF, January 1996.

[26] R. Stewart, M. Ramalho, Q. Xie, M. Tuexen, and P. T. Conrad. SCTP partial reliability extension. Internet draft, IETF, May 2002. Work in Progress.

[27] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson. Stream control transmission protocol. RFC 2960, IETF, October 2000.

[41] W-T Tan and A. Zakhor. Real-time internet video using error resilient scalable compression and TCP-friendly transport protocol. *IEEE Transactions on Multimedia*, 1(2):172–186, June 1999.

[42] K. Thompson, G. Miller, and R. Wilder. Wide-area Internet traffic patterns and characteristics (extended version). *IEEE Network*, pages 10–23, November 1997.