Katarina Asplund

Implementation and Experimental Evaluation of a Partially Reliable Transport Protocol

Karlstad University Studies
2004:11

Katarina Asplund. *Implementation and Experimental Evaluation of a Partially Reliable Transport Protocol.*

Licentiate thesis
March 2004

Karlstad University Studies 2004:11
ISSN 1403-8099
ISBN 91-85019-75-5

Distribution:
Karlstad University
Division for Information Technology
Department of Computer Science
SE-651 88 KARLSTAD
SWEDEN
+46 54-700 10 00

www.kau.se

Printed at Universitetstryckeriet, 2004

Implementation and Experimental Evaluation of a Partially Reliable Transport Protocol

KATARINA ASPLUND
Department of Computer Science, Karlstad University

Abstract

In the last decade, we have seen an explosive growth in the deployment of multimedia applications on the Internet. However, the transport service provided over the Internet is not always feasible for these applications, since the network was originally designed for other types of applications. One way to better accommodate the service requirements of some of these applications is to provide a partially reliable transport service. A partially reliable transport service does not insist on recovering all, but just some of the packet losses, thus providing a lower transport delay than a reliable transport service. The work in this thesis focuses on the design, implementation, and evaluation of a partially reliable transport protocol called PRTP. PRTP has been designed as an extension to TCP in order to show that such a service could be effectively integrated with current protocol standards. An important feature of PRTP is that all modifications for PRTP are restricted to the receiver side, which means that it could be very easily deployed. The thesis presents performance results from various experiments on a Linux implementation of PRTP. The results suggest that transfer times can be decreased significantly when using PRTP as opposed to TCP in networks in which packet loss occurs. Furthermore, the thesis includes a study that investigates how users perceive an application that is based on a partially reliable service. Specifically, how users select the trade-off between image quality and latency when they download Web pages is explored. The results indicate that many of the users in the study could accept less than perfect image quality if the latency could be shortened.

Keywords: Partial reliability, transport protocol, Internet, performance evaluation

Acknowledgments

Firstly I would like to express my heartfelt gratitude to my supervisor Anna Brunstrom who has always supported me in my work and encouraged me when I have had difficulties. I am also grateful to my collegues at the Computer Science Department and especially the members of the DISCO Research group for fruitful discussions and moral support. Furthermore, I would like to thank Bengt and my children Viktor and Malin for their love and support during this time, which has been of immense value to me. Finally, I would like to express my gratitude for the financially support of the work presented in this thesis provided by Ericsson AB, the KK-foundation and CMIT.

List of Appended Papers

This thesis is comprised of the following four papers. References to the papers will be made using the Roman numbers associated with the papers such as Paper I.

I. Sean Schneyer, Johan Garcia, Anna Brunstrom, Katarina Asplund, PRTP: A Partially Reliable Transport Protocol for Multimedia Applications, *Proceedings of ISIMADE 1999*, Baden-Baden, Germany, August 1999.

II. Katarina Asplund, Anna Brunstrom, Partially Reliable Multimedia Transport, *Karlstad University Studies 2004:5*, Karlstad University, Sweden, February 2004. A slightly revised version of this paper will be published as

> Katarina Asplund, Anna Brunstrom, Partially Reliable Multimedia Transport, chapter in "Perspectives on Multimedia", *to be published by John Wiley & Sons in 2004*.

III. Katarina Asplund, Johan Garcia, Anna Brunstrom, Decreasing Transfer Delay Through Partial Reliability, *Proceedings of PROMS 2000*, Cracow, Poland, October, 2000.

IV. Katarina Asplund, Anna Brunstrom, The Trade-off Between Latency and Image Quality on the Web: A User Perspective, *Proceedings of CIIT 2002*, St. Thomas, US Virgin Islands, November, 2002.

Some of the papers have been subjected to some minor editorial changes.

Contents

Introductory Summary [redacted]

1 Introduction

The history of the Internet goes back only some twenty years. The development of this now globe-spanning network was an initiative from the US Department of Defense, which wanted a communication network that was simple, reliable and robust. The Internet architecture and its protocols were of course designed to work well for the applications that were used in those days, such as file transfer, e-mail and DNS. The services provided by the two standard transport protocols of the Internet, the Transmission Control Protocol (TCP) [1] and the User Datagram Protocol (UDP) [2], are especially well suited for these kinds of applications. On the one hand, TCP provides full reliability, covering for the imperfectnesses[1] of the underlying network. The service TCP provides is well suited for file transfer, e-mail and remote login for which full reliability is of the utmost importance. Of course this full reliability comes at a price; since it is realized through retransmissions of lost data, the latency of data transfer is increased. This price is almost always worth paying, however, since reliability has an absolute priority over timeliness for these applications. UDP, on the other hand, provides a kind of ”no frills” service to the application. This means that it does not offer any extra reliability, but neither does it increase the latency. The service that UDP provides is well suited for applications such as DNS and SNMP.

However, in the last decade, applications with rather different requirements than those of the traditional applications mentioned above have emerged. One example is multimedia applications, i.e. applications that involve transmission of either video, audio or image data (or some combination thereof). Common to these applications is that they are more sensitive to delay and jitter than traditional applications but instead have a limited tolerance for loss. Neither TCP nor UDP provides an adequate service to many of these applications since TCP often introduces too much latency and jitter and UDP is too unreliable.

The fact that the range of services offered by the transport layer is limited (to say the least), makes it difficult for the Internet to accommodate the service requirements of these, and possibly future, applications. Several ways to address this situation have been proposed. Some of these proposals involve extending the Internet architecture in order to offer other service classes than the single class of best-effort service offered today. Two examples of such proposals are Integrated Services [3] and Differentiated Services [4].

Another way to deal with this situation is to build new end-to-end protocols that, on top of the best-effort Internet, would provide a more flexible service. Such a service would be more adaptable to the differing service requirements of applications. One example of a flexible application protocol which is used by many multimedia applications today, is the Real-Time Transport Protocol[2] (RTP) [21]. Using RTP, applications can obtain

[1] The Internet offers a single class of best-effort service, i.e. the network offers no assurance about when, or even if, packets will be delivered.

[2] Although the name suggests that RTP is a transport layer protocol, it does not provide full transport

a very flexible service, but RTP also has the disadvantage of having to be integrated in nearly each different application. A flexible transport layer protocol, however, can be used directly by a range of different applications without any integration. Best, of course, would be to have flexibility both in the transport and in the application layer. Nevertheless, as the resources of the Internet are limited, flexibility means that something must be traded to obtain something else. For example, to obtain less latency and jitter, you can trade some of the reliability.

In this thesis, we investigate the benefits of a flexible transport service and how users perceive the service that a system that supports such a service provides. Although there exist several conceivable Quality of Service (QoS) parameters, we concentrate in this work on reliability. A service that offers flexible reliability is commonly called a partially reliable service, i.e. a service that does not insist on recovering all but only some of the packet losses. This idea is not new. On the contrary, several transport protocols that provide partial reliability exist. These protocols can be classified as either open-loop or closed-loop protocols. Open-loop protocols involve error- correcting methods such as Forward Error Correction (FEC) [5, 6, 7], which do not require a back channel. In contrast, closed-loop protocols allow the receiver to give feedback on which packets have arrived and which have not. These protocols are often called retransmission-based protocols, as retransmissions are most often the answer to lost packets. Several retransmission-based protocols have been proposed that in different ways provide a partially reliable service [8, 9, 10, 11, 12]. A taxonomy and survey of existing retransmission-based partially reliable transport protocols are presented in [13].

The work in this thesis focuses on the design, implementation and evaluation of such a protocol, called the Partially Reliable Transport Protocol (PRTP). In comparison with other proposed partially reliable transport protocols, PRTP is designed as an extension to current protocol standards and at the same time offers a fully receiver-based approach, which means that it can be very easily deployed. We have implemented the PRTP protocol in the Linux operating system and present in this thesis performance results of various experiments on this Linux implementation. In addition, a study that investigates user perception of a system that provides a flexible reliability service is included.

Our work has been part of a project carried out by the DISCO Research Group at Karlstad University. The aim of this project has been to build a system that provides a flexible QoS and to investigate the performance and usefulness of the system. The system developed consists partly of PRTP and partly of an example application used to test and demonstrate the protocol. The example application is a Web proxy that transcodes JPEG images for partial reliability. Using this system as a test-bed, the trade-off between image quality and latency when transferring Web pages can be studied. A description of the JPEG transcoder and quantitative measurements of the system performance are found in a thesis by Garcia [15]. Besides the performance studies presented in this thesis, simulation studies of PRTP have also been made; these are presented in a thesis by Grinnemo

functionality but is designed to run on top of another transport protocol, typically UDP. Therefore, in this respect RTP can be seen as an application layer protocol.

[14].

2 Research Goals

The Internet architecture and its protocols are currently not well suited for multimedia traffic, such as transfer of audio, video and image data. Our research is aimed mainly at developing and investigating protocols that are better suited to the demands these applications have on data transfer. As part of this overall goal, this thesis studies the performance and usefulness of a transport protocol that provides a partially reliable service. We especially investigate the trade-off between transfer time and level of reliability, where both quantitative performance measurements and studies of user perception are included.

3 Research Overview

In order to investigate the benefits of a partially reliable transport service we have developed the partially reliable transport protocol PRTP. To be able to focus on issues directly related to partial reliability we chose initially to modify TCP, where partial reliability is only an option and normal TCP is the default. This design choice has the advantage that our new PRTP protocol is part of the current Internet infrastructure. PRTP allows an application to specify the reliability level dynamically during a connection's lifetime and also ensures that this reliability level is maintained. This means that PRTP provides a deterministic reliability guarantee to the application. All modifications for PRTP are restricted to the receiver side of the protocol. This approach has two important advantages. Firstly, it means that PRTP can be very easily deployed and, secondly, it means that it can interoperate with regular TCP senders.

We have evaluated the performance of PRTP under a variety of network conditions. Experiments have been done both with and without background traffic and we have measured the performance for both long and short file transfers. As the main purpose of the experiments was to investigate whether we could obtain lower delay in return for decreased reliability, the performance metric we considered was transfer time. The results indicate that PRTP can considerably decrease transfer times, as compared to TCP, in networks in which packet loss occurs. This is especially true if the loss tolerance is rather high (5-10%).

However, some of this performance improvement comes from the fact that PRTP has a more aggressive congestion control than TCP. This means that PRTP is not altogether TCP-friendly [16], but takes more than its fair share of the bandwidth. In a simulation study [17], Grinnemo et al investigated the use of explicit congestion notification (ECN) [18] to make PRTP TCP-friendly. This study shows that this ECN-enabled version of PRTP is indeed TCP-friendly, although the performance improvement in relation to TCP is smaller than for PRTP without ECN.

Although our experiments show that PRTP provides a very flexible trade-off between transfer time and reliability in many situations, the results of the experiments do not say anything about how beneficial such a service would be for end users. For this reason, we studied how users select the trade-off between image quality and delay when downloading different Web pages. The study explores how users perceive the type of service a system based on the PRTP protocol can provide. Investigating how users perceive a new network service is important since research on user QoS has shown that users' perception of the service they receive does not always reflect the objective service provided by the network [19]. The results show that many of the users in the study preferred faster download time to perfect image quality. The extent of the acceptable image quality deterioration, however, was highly dependent on the context and on the individual user.

4 Research Contributions

The main contributions of the work presented in this thesis are the following:

- We implemented the partially reliable transport protocol PRTP in the Linux kernel. Since PRTP is designed as an extension to TCP and all modifications are restricted to the receiver side, PRTP can be very easily deployed.
- We show some of the benefits of the partially reliable service provided by PRTP, especially how transfer times can be decreased when a reliability level of less than 100% is acceptable.
- We investigate how users could benefit from such a partially reliable service that PRTP provides. The study explores how users perceive the trade-off between image quality and latency when downloading Web pages. The results show that many users could accept less than perfect image quality if the latency could be shortened, and therefore indeed would benefit from a flexible Web service.

5 Thesis Outline

This thesis consists of four papers. These papers are listed below, together with a short summary of each paper.

Paper I: PRTP: A Partially Reliable Transport Protocol for Multimedia Applications

In this paper, our ideas about the design of a transport protocol that provides a partially reliable service are presented for the first time. A partially reliable service can be described according to a number of different characteristics. Some of the most important are the type of reliability guarantee, the variability of the reliability level and the

range of reliability. Various design choices made for the PRTP protocol are presented, and a description is given of the application test-bed built to test and demonstrate the PRTP protocol. The application is a Web proxy that transcodes JPEG images for partial reliability.

Paper II: Partially Reliable Multimedia Transport

This technical report describes PRTP in a rather broad context. In particular, background material on functions in TCP that are relevant for PRTP is presented. The report presents results from a number of experiments that investigate different aspects of PRTP. A summary of the experimental study on the steady state behavior of PRTP over an emulated network that is presented in more detail in paper III is first presented. The transient behavior of PRTP, i.e. PRTP's performance for short file transfers, is then investigated and some initial results are presented. The report also describes tests that evaluate the steady state behavior of PRTP over a real transatlantic connection, and thus how PRTP interacts with regular Internet traffic.

Paper III: Decreasing Transfer Delay through Partial Reliability

This paper describes the actual design and Linux implementation of PRTP in detail. In addition, the paper presents performance results in an experimental study that investigates the steady state behavior of PRTP over an emulated network. Specifically, transfer times for PRTP in networks with different one-way packet delays and different packet loss probabilities are presented. As PRTP is designed as an extension to TCP, the performance of PRTP is related to that of TCP. The results indicate that PRTP can decrease transfer times significantly as compared to TCP for loss tolerant applications. The tradeoff between transfer times and reliability can also be done very flexibly.

Paper IV: The Trade-off between Latency and Image Quality: A User Perspective

The previous papers describe the design and implementation of the PRTP protocol and give an overview of the Web browsing application. They also present a number of experiments that measure the objective (quantitative) service provided by the PRTP protocol. In contrast, this paper explores how users perceive the type of service a system based on partial reliability can provide. Specifically, how users select the tradeoff between image quality and latency when downloading Web pages is studied. The results indicate that a small impairment of the image quality in return for decreased download time is preferrable for many of the users in the study.

6 Future Work

The work in this thesis has been an attempt to come a few steps closer to offer a more flexible QoS to applications on the Internet. In pursuing this work further, several possibilities exist. Firstly, our transport protocol PRTP could be further developed in a

number of aspects. The most obvious aspect is the addition of ECN in the Linux implementation, in order to make PRTP TCP-friendly.
A second possibility is to use another transport protocol than TCP as a platform on which to build flexible services. One strong candidate for a new platform is the Stream Control Transmission Protocol (SCTP) [20], a transport protocol that was rather newly standardized by the IETF. SCTP has a number of interesting features that could be explored in order to provide flexibility. Furthermore, SCTP already provides partially ordered message delivery, and it is suggested that it also should provide partial reliability [22].
Another important area for future work is to study how users perceive the flexible services that are developed; it is not always a one-to-one relationship between quantitative measurements of the service and a users' qualitative judgment of the same service. There is also the question of how a flexible service should be specified. If the service should be specified at the user level, it is a non-trivial task to design a good, understandable user interface.

References

[1] J. Postel, "RFC 793: Transmission control protocol," Sept. 1981.

[2] J. Postel, "RFC 768: User datagram protocol," Aug. 1980.

[3] R. Braden, D. Clark, and S. Shenker, "RFC 1633: Integrated services in the Internet architecture: an overview," June 1994.

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

[5] J.-C. Bolot and T. Turletti, "Adaptive error control for packet video in the Internet," *Proceedings of CIP '96, Lausanne, Switzerland*, September 1996.

[6] J.-C. Bolot and A. Vega-Garcia, "Control mechanisms for packet audio in the Internet," *Proceedings of IEEE INFOCOM'96, San Fransisco, CA*, March 1996.

[7] Y. Rasheed and A. Leon-Garcia, "Towards a generic real-time transport and adaptation protocol for ip networks," *Computer Communications*, vol. 23, pp. 1448–1458, August 2000.

[8] P. D. Amer, C.Chassot, T. Connolly, M. Diaz, and P. Conrad, "Partial order transport service for multimedia and other applications," *IEEE/ACM Transactions on Networking*, vol. 2, pp. 440–456, October 1994.

[9] M. Diaz, A. Lopez, C. Chassot, and P. Amer, "Partial order connections: a new concept for high speed and multimedia services and protocols," *Annales des telecoms*, vol. 49, no. 5-6, 1994.

[10] B. J. Dempsey, "Retransmission-based error control for continuous media," *PhD thesis, University of Virginia*, 1994.

[11] P. T. Conrad, E. Golden, P. D. Amer, and R. Marasli, "A multimedia document retrieval system using partially-ordered/partially-reliable transport service," *Proceedings of Multimedia Computing and Networking*, January 1996.

[12] M. Piecuch, K. French, G. Oprica, and M. Claypool, "A selective retransmission protocol for multimedia on the Internet," *Proceedings of SPIE Multimedia Systems and Applications*, pp. 79–90, November 2000.

[13] 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, Oct 2002.

[14] K.-J. Grinnemo and A. Brunstrom, "A study of partially reliable transport protocols for soft real-time applications," *Licentiate thesis, Karlstad University Studies 2002:36*, Karlstad University, Sweden, Nov 2002.

[15] J. Garcia, "Application and transport layer flexibility: An image transfer example," *Licentiate thesis, Karlstad University Studies 2002:11*, Karlstad University, Sweden, 2002.

[16] S. Floyd and K. Fall, "Promoting the use of end-to-end congestion control in the Internet," *IEEE/ACM Transactions on Networking*, vol. 7, pp. 458–472, August 1999.

[17] K.-J. Grinnemo and A. Brunstrom, "Evaluation of the QoS offered by PRTP-ECN - a TCP-compliant partially reliable protocol," *Proceedings of IWQoS 2001, Karlsruhe, Germany*, pp. 217–231, June 2001.

[18] K. Ramakrishnan, S. Floyd, and D. Black, " RFC 3168: The addition of explicit congestion notification (ECN) to IP,"September 2001.

[19] A. Bouch and M. A. Sasse, "Network quality of service: What do users need?," *Proceedings of the 4th International distributed conference (IDC'99)*, pp. 78–90, September 1999.

[20] R. Stewart, "RFC 2960: Stream Control Transmission Protocol," October 2000.

[21] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RFC 3550: RTP: A transport protocol for real-time applications, July 2003.

[22] R. Stewart, M. Ramalho, Q. Xie, M. Tuexen, and P. Conrad. SCTP partial reliability extension. *Work in progress, draft-ietf-tsvwg-prsctp-02.txt*, November 2003.