

# Transport QoS over Unreliable Networks: No Guarantees, No Free Lunch!

*P. Conrad, P. Amer, E. Golden, S. Iren, R. Marasli, A. Caro*  
*Computer and Info. Science Dept, Univ. of Delaware*  
*Newark, DE 19716 USA*  
*Email: {pconrad, amer, golden, iren, acar0}@cis.udel.edu*  
*R. Marasli, FORE Systems, Pittsburgh PA, USA, rmarasli@fore.com*

## Abstract

We describe an approach to transport QoS in unreliable networks that focuses on tradeoffs rather than guarantees. In particular, we investigate tradeoffs between qualitative QoS parameters such as order and reliability, and quantitative parameters such as delay and throughput.

## 1 INTRODUCTION

A traditional approach to quality of service (QoS) is to provide applications with guarantees for *quantitative* QoS parameters such as delay and throughput. However, when the underlying network is inherently unreliable, it may be difficult or even impossible to provide guarantees. Therefore, we focus on *QoS tradeoffs* between *qualitative* parameters such as order and reliability, and quantitative parameters such as delay and throughput. A major theme of our work is that there is “no free lunch”; improvements in reliability and order come at a cost of worsening delay and throughput. Our research has two goals: (1) to understand and characterize the nature of these tradeoffs via analysis, simulation and experimentation, and (2) to implement a transport protocol that provides the service user with as much flexibility as possible in selecting the best tradeoff for a particular application. Towards this end, we have investigated *partially ordered and partially reliable (PO/PR) transport service* which provides flexible QoS tradeoffs to application designers and users.

## 2 PARTIALLY ORDERED/PARTIALLY RELIABLE TRANSPORT SERVICE

In a packet-switched network, the transport layer is the lowest layer responsible for end-to-end QoS. Transport layer functions include: (1) recovery from data loss, (2) detecting and removing duplicates, (3) resequencing out-of-order

data, and (4) flow control/congestion avoidance. We view the level of service provided in each of these functional areas as a qualitative QoS parameter.

PO/PR transport service fills the gap between the extremes in service represented by TCP (ordered/reliable/flow-controlled) and UDP (unordered/unreliable/not-flow-controlled). PO/PR service allows a user to avoid the unnecessary degradation in QoS parameters such as delay and throughput that may result if a service is used that is more strict than necessary (e.g., TCP.) Instead, PO/PR service can provide flow control, and the precise degree of order and reliability that an application requires—no more and no less.

Our approach to providing PO/PR service is to implement it as a set of library routines (built over UDP) that applications can utilize to gain flexible control over the ordering and reliability of individual objects. This allows applications to achieve an appropriate balance among various QoS parameters without having to reinvent the transport-layer wheel with every application. Such an approach is consistent with Application Level Framing as proposed in Clark & Tennenhouse (1990). Our results so far include analytic and simulation results characterizing the QoS tradeoffs offered by PO/PR service (Section 3) and the development of software components for experimentation with PO/PR service (Section 4).

### 3 ANALYSIS AND SIMULATION

Analytic and simulation work has yielded several important results about partially reliable and partially ordered services (Marasli 1997). The key results related to partially ordered service are: (1) PO service provides buffer utilization and delay improvements over ordered service, particularly as the loss rate increases, or as the order requirements of applications decrease. Order, however, makes no difference to throughput, as long as the transport sender buffer size is less than or equal to the transport receiver buffer size (Marasli, Amer & Conrad 1997b); (2) by carefully choosing the transmission order of independent objects (linear extension selection), additional buffer utilization and delay improvements can be obtained (Marasli et al. 1997b, Marasli, Amer & Conrad 1996a).

The key results related to partially reliable service are: (1) for lossy networks, using a reliable transport service when only a partially reliable service is needed can cause considerable worsening in throughput and delay (Marasli, Amer & Conrad 1996b); (2) both sender-based and receiver-based reliability schemes for providing partial reliability achieve almost identical reliability and delay. On the other hand, a sender-based approach provides better throughput than a receiver-based approach at higher ack loss rates (Marasli, Amer & Conrad 1997a).

## 4 EXPERIMENTATION

Our empirical research into PO/PR service compares the performance of PO/PR protocols with that of traditional protocols, under a variety of network conditions. We have developed two client-server multimedia applications to test our notions about QoS tradeoffs:

- A system for retrieving compressed images across a lossy, packet-switched network, allowing experimentation with various combinations of image compression techniques and transport protocols. This demonstration simulates a hypothetical military communications system for transmitting images of wounded soldiers for telemedicine or images of equipment such as tanks, airplanes, etc. for intelligence gathering (Amer, Conrad, Golden, Iren & Caro 1997).
- A multimedia document retrieval system allowing authors to specify synchronization requirements, and varying degrees of reliability for multimedia objects. The basic model is similar to that of the World Wide Web; documents are available on a server, and are retrieved with a browser. However, unlike Web documents, these documents are temporal; they have a time dimension requiring synchronization of elements such as audio, video, still-images, text, pauses and interactions (Conrad, Golden, Amer & Marasli 1996).

We have also built prototype implementations of two PO/PR protocols:

- *k*-XP provides an unordered, partially-reliable transport service (Amer et al. 1997).
- *Partial Order Connection, version 2* (POCv2) provides a partially-ordered/partially-reliable transport service, and features for coarse-grained multimedia synchronization (Conrad et al. 1996).

In support of future work, we are currently developing:

- A library of transport functionality called the *Universal Transport Library (UTL)*. The UTL provides a generalized application programming interface to a number of different transport services with varying degrees of order and reliability. This allows a programmer to write applications that can be run over different transport services simply by varying a single parameter at connection establishment time (Conrad 1997).
- New versions of the compressed image and multimedia document retrieval systems based on UTL. With these applications, we can investigate the impact of QoS tradeoffs on the application as perceived by a real user, in addition to precise quantitative measurements of these tradeoffs under real network conditions (Conrad 1997, Iren in progress).

**Acknowledgments/Disclaimers:** Prepared through collaborative participation in the Advanced Telecommunications Information Distribution Research Program (ATIRP) Consortium sponsored by the U.S. Army Research Laboratory under the Federated Laboratory Program, Cooperative Agreement DAAL01-96-2-0002. This work also supported, in part, by the National Science Foundation (NCR-9314056) and by the US Army Research Office (ARO) (DAAL03-92-G-0070, DAAH04-94-G-0093). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied of the Army Research Laboratory or the U.S. Government.

## REFERENCES

- Amer, P., Conrad, P., Golden, E., Iren, S. & Caro, A. (1997), Partially ordered, partially reliable transport service for multimedia applications, *in* ‘Advanced Telecommunications/Information Distribution Research Program Annual Conference’, College Park, MD.
- Clark, D. D. & Tennenhouse, D. L. (1990), Architectural considerations for a new generation of protocols, *in* ‘SIGCOMM 1990’, pp. 200–208. *Computer Communications Review*, Vol. 20(4), Sept. 1990.
- Conrad, P. T. (1997), Order, Reliability, and Synchronization in Transport Layer Protocols for Multimedia Document Retrieval, PhD thesis, Computer and Information Sciences, University of Delaware. (in progress).
- Conrad, P. T., Golden, E., Amer, P. D. & Marasli, R. (1996), A multimedia document retrieval system using partially-ordered/partially-reliable transport service, *in* ‘Multimedia Computing and Networking 1996 (MMCN96; sponsored by SPIE/IS&T)’, San Jose, CA, USA.
- Iren, S. (in progress), ‘Network-conscious image compression’. PhD Dissertation, CIS Dept., University of Delaware.
- Marasli, R. (1997), Partially Ordered and Partially Reliable Transport Protocols: Performance Analysis, PhD thesis, University of Delaware.
- Marasli, R., Amer, P. & Conrad, P. (1997a), Partially reliable transport service, *in* ‘Proc Second IEEE on Computers and Communications (ISCC’97)’, Alexandria, Egypt. to appear.
- Marasli, R., Amer, P. D. & Conrad, P. T. (1996a), Optimizing partially ordered transport services for multimedia applications, *in* ‘Third International Conference on Multimedia Modeling’, Toulouse, France.
- Marasli, R., Amer, P. D. & Conrad, P. T. (1996b), Retransmission-based partially reliable transport service: An analytic model, *in* ‘IEEE INFOCOM’96’, IEEE, San Francisco, CA, pp. 621–629.
- Marasli, R., Amer, P. D. & Conrad, P. T. (1997b), ‘An analytic study of partially ordered transport services’, *Computer Networks and ISDN Systems*. (To appear).