

IMPROVING MULTIMEDIA PERFORMANCE OVER LOSSY NETWORKS VIA SCTP

Armando L. Caro Jr.
Paul D. Amer

Computer and Information Science Department
University of Delaware, Newark, DE 19716 USA
Email: {acarо,amer}@cis.udel.edu

Phillip T. Conrad
Gerard J. Heinz

Computer and Information Science Department
Temple University, Philadelphia, PA 19122 USA
Email: conrad@joda.cis.temple.edu, gheinz@astro.temple.edu

ABSTRACT

The Stream Control Transmission Protocol (SCTP) is a new Internet standards track transport layer protocol. SCTP was designed to transport PSTN signaling messages over IP networks, but SCTP is also capable of serving as a generalized transport protocol. As such, SCTP provides an alternative to the traditional transport protocols, TCP and UDP, that may be better able to satisfy the requirements of future battlefield networks.

Unlike traditional transport protocols, SCTP allows multiple streams of messages within a single connection (or, in SCTP terminology, a single association). As the results in this paper show, this ability is particularly helpful in reducing latency for streaming multimedia in high loss environments.^a

1 INTRODUCTION

The future battlefield environment will include mobile ad-hoc and wireless sensor nodes which deliver

^aPrepared 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. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes not withstanding any copyright notation thereon. 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.

streaming real time multimedia to end users. Traditional transport protocols are not well suited to the relatively high loss rates of battlefield networks. New protocols which are designed to handle flexible service requirements can offer better QoS tradeoffs for future army networks.

Previous work has shown that traditional transport protocols, such as TCP and UDP, are less robust to packet loss than protocols incorporating partial order and partial reliability [2]. Previously, partial order and partial reliability were only implemented in experimental protocols. However, the telecommunication industry is strongly backing a new internet standards track protocol, SCTP (RFC2690) [4], which incorporates partial order and has extensions for partial reliability in the Internet Draft stage.

Section 2 provides an overview of SCTP. The merits of a partially-ordered service will be covered in Section 3. Section 4 presents results from previous study of partially ordered vs ordered service. These results indicate the nature of the expected performance gains SCTP may be able to provide. Section 5 will end the paper with some concluding remarks and ideas for future work.

2 SCTP OVERVIEW

The Stream Control Transmission Protocol (SCTP) is a reliable transport protocol operating over a connectionless packet-switched network, such as IP. SCTP emerged from the need for telecommunications companies to manage SS7 applications and

services over an IP infrastructure. SS7 is a protocol suite for managing PSTNs and other telecommunication networks. The upper layers of SS7 are designed to operate over a circuit-switched control channel of the industry's TDM phone system network. Therefore, SCTP is oriented towards providing connection-oriented reliable message streams between communication endpoints.

While SCTP was originally designed for signaling transport, and much of the current work on SCTP in the IETF is centered around this application, the protocol designers recognize that the protocol *is capable of broader applications* [4]. Therefore, it is important to compare SCTP's transport services with that of UDP and TCP.

UDP is an unreliable, yet fast connectionless datagram service. Delay sensitive messages find UDP suitable in that regard, but many applications cannot tolerate the lack of ordered delivery, loss recovery, duplicate detection, congestion control, and flow control.

Contrasting UDP, TCP does provide a reliable transport service with all the lacking features of UDP. However, many applications find TCP too restrictive. Some of the drawbacks of TCP which users have wanted to bypass are as follows.

- TCP is byte-stream-oriented, which means that applications are responsible for tracking message boundaries and using the push mechanism to ensure messages are transferred in reasonable time.
- TCP is ordered. Strict order-of-transmission data delivery is a restriction for some applications. Many times unordered or partially ordered data delivery is sufficient for the application's needs. For such applications, TCP causes unnecessary delays due to its head-of-line blocking.
- TCP does not transparently support multihomed hosts.
- TCP is vulnerable to denial of service attacks, which makes it a risky protocol to use in mission critical applications.

- TCP does not give applications control over protocol parameters, but this may be necessary for some applications to have their specific needs met.

As a result of TCP's limitations, customized protocols have been built over UDP to provide the features needed. This creates much unnecessary work in the form of duplicate effort. Many application developers find themselves wasting energy and resources developing new protocols to fit their specific transport needs which could not be met by UDP or TCP.

SCTP attempts to enhance the services of UDP and overcome the limitations of TCP. SCTP is a reliable message-based connection-oriented transport protocol, which according to [4], provides the following services:

- acknowledged error-free non-duplicated transfer of user data,
- congestion avoidance behavior,
- sequenced delivery of user messages within multiple streams ^b (i.e., partially-ordered data delivery), with an option for order-of-arrival delivery of individual user messages,
- data fragmentation to conform to discovered path MTU size,
- optional bundling of multiple user messages into a single SCTP packet,
- network-level fault tolerance through support of multihoming at either or both ends of an association, and
- resistance to flooding and masquerade attacks.

In this paper, we will focus on the feature of partially-ordered data delivery. The next section will explain the benefits of such a service.

^bThe term *stream* is used in SCTP to refer to a sequence of user messages that are to be delivered to the upper-layer protocol in order with respect to other messages within the same stream. This is in contrast to its usage in TCP, where it refers to a sequence of bytes. [4]

3 BENEFITS OF PARTIALLY-ORDERED SERVICE

Since traditional protocols (UDP and TCP) offer only the extremes (unordered and strictly ordered service, respectively), then application developers with needs in between these extremes are faced with a dilemma. If TCP is chosen, unnecessary performance penalties must be paid. On the other hand, if UDP is chosen, then developers must build their own transport protocol over UDP to provide the exact services that are needed.

A flexible transport protocol offering a partially-ordered service is ideal for applications that need flexible control over the ordering of individual elements. Such a service is essential for balancing various QoS parameters, and to avoid having to implement a custom protocol for each new application.

Through analysis and simulation, we investigated the benefits of a partially-ordered service, and found that such a service could provide improvements in throughput, delay, and buffer utilization for a normalized time-scale [7, 8, 9]. Later, the development of an innovative transport protocol that provides a partially-ordered service (POCv2) and an application which uses the protocol (ReMDoR), demonstrated that the theoretical advantages of a partially-ordered service claimed in previous work can be achieved in practice [2].

Early in 2000, serious work began on the new protocol SCTP, which also incorporates features of partial order. Thus, SCTP provides a capability very similar to the stream capability proposed within POCv2. However, SCTP provides a restricted form of partial order service. While SCTP does not allow for complex partial orders, as does POCv2, one can imagine implementing a thin layer on top of SCTP that would result in a protocol equivalent or nearly equivalent to POCv2's partially-ordered service.

4 EMPIRICAL RESULTS

Reference [2] includes extensive results from sixteen experiments comparing the performance of multimedia document retrieval over reliable transport services providing unordered, partially-ordered, and ordered delivery. In this section, we present results from one of these experiments that highlights the

kind of benefits we expect to be obtainable from SCTP.

Experiment R1 from [2] compares ordered/reliable service to partially-ordered/reliable service for retrieval of a document with eight images presented in parallel. Experiment R1 uses the standard GIF compression technique rather than a specific network conscious compression technique [5, 6]. The GIF format requires ordered/reliable delivery for each image, so unordered service cannot be used. However, partially-ordered service can be used because the data for each image can be interleaved in eight parallel streams. This experiment uses the ReMDoR application [1] in addition to the UTL and Lossy Router tools developed by the Protocol Engineering Laboratory at the University of Delaware [3].

The hypothesis for this experiment is that for all loss rates $> 0\%$, partially-ordered/reliable (PO/R) service provides, on average, better progressive display for parallel GIF images than ordered/reliable (O/R) service. For practical reasons, it is necessary to refine this hypothesis somewhat. It would be absurd to evaluate the gain at very low loss rates ($< 1\%$). For very low loss rates, the gain will be so small as to be insignificant for all practical purposes. It would be equally absurd to evaluate PO/R service vs O/R service at loss rates approaching 100%. At these loss rates, the performance of both services would be unacceptable. So instead of trying to characterize the gains of partial order over the entire range of loss rates between 0% and 100%, we focused on a few loss rates, and on the trend in performance as the loss rate increases.

Due to space limitations, this paper will present only a subset of the experimental results of R1, which we will refer to as Experiment R1.1. Experiment R1.1 illustrates the performance of the PO/R service provided by the R2E protocol versus the O/R service provided by the T2E protocol ^c. The performance comparison is done over a PPP connection of 9.6 kbps with a 10% loss rate. Figure 1 shows an average performance graph for Experiment R1.1. From this graph, we conclude that this set of experimental

^cThere are several reasons explained in [2] why these two particular services were chosen from among the dozens available in UTL.




Figure 1: Performance graph: Experiment R1.1 (9.6kbps PPP link at 10% loss)

data supports the hypothesis stated above.

In addition, the results of the complete Experiment R1 presented in [2] show the following:

- At 0% loss, R2E and T2E have virtually identical performance.
- While both R2E and T2E experience worse performance as the loss rate increases, the performance of R2E degrades more slowly than that of T2E.
- At nearly every point in time, on average, R2E provides more data (show both in bytes and pixels) to the end-user.

To provide an end-user perspective, Figure 2 shows the difference between R2E and T2E performance at a few sample points for 10% loss. As can clearly be seen, at each of these points, partially-ordered service provides better performance than totally-ordered service. While human factor studies (which we suggest as future work) would be necessary to establish this scientifically, we hypothesize that the initial delivery of at least a few pixels will prove to be highly correlated with user satisfaction. Seeing at least some progress provides hope to the user, while seeing a screen that does not change for a long period of time (especially a blank one) can be discouraging.

5 CONCLUSION AND FUTURE WORK

Previously, the merits of partial order were theoretical advantages proven by analytical models and simulation only. The results in this paper show, however, that applications requiring such a transport service do actually achieve the benefits in practice. Until recently, partial order has only been implemented in experimental transport protocols for research purposes, but the telecommunication industry is now interested the new and upcoming transport protocol which supports partial order, namely SCTP.

The PEL researchers of the University of Delaware in collaboration with Temple's NetLab researchers are investigating SCTP. We plan to evaluate the performance of SCTP when used for streaming real-time multimedia and other partial order benefiting applications. Our goal is to develop an integrated multimedia transport protocol to fit the needs of future army networks.

6 REFERENCES

- [1] A. Caro. Remdor 2.0: Remote multimedia document retrieval over partially-ordered, partially-reliable transport protocols, May 1998. BS Thesis, CIS Dept., University of Delaware.
- [2] P. Conrad. Order, reliability, and synchronization in transport layer protocols for multimedia document retrieval, 2000. PhD Dissertation, CIS Dept. University of Delaware.
- [3] P. Conrad, P. Amer, M. Taube, G. Sezen, S. Iren, and A. Caro. Testing environment for innovative transport protocols. In *MILCOM '98*, Bedford, MA, October 1998.
- [4] R. Stewart et al. Stream control transmission protocol. RFC 2960, October 2000.
- [5] S. Iren. Network-conscious image compression, 1999. PhD Dissertation, CIS Dept., University of Delaware.
- [6] S. Iren, P. Amer, A. Caro, G. Sezen, M. Taube, and P. Conrad. Network-conscious compressed image transmission over battlefield networks. In *MILCOM '98*, Bedford, MA, October 1998.
- [7] R. Marasli. Partially ordered and partially reliable transport protocols: Performance analysis, 1997. PhD Dissertation, CIS Dept., University of Delaware.
- [8] R. Marasli, P. Amer, and P. Conrad. Retransmission-based partially reliable services: An analytic model. In *IEEE INFOCOM*, San Fransisco, CA, March 1996.
- [9] R. Marasli, P. Amer, and P. Conrad. An analytic model of partially ordered transport service. *Computer Networks and ISDN Systems*, 29(6):675–699, May 1997.

	Partially-ordered service (R2E)	Ordered service (T2E)
25 seconds	<p>avg. 14447 (13508 pixels shown)</p>	<p>avg. 1530 (1535 pixels shown)</p>
50 seconds	<p>avg. 53767 (53294 pixels shown)</p>	<p>avg. 40437 (41745 pixels shown)</p>
75 seconds	<p>avg. 87345 (86446 pixels shown)</p>	<p>avg. 75665 (76506 pixels shown)</p>

Figure 2: Screenshot: Experiment R1.1 (9.6kbps PPP link at 10% loss)