SCTP MULTISTREAMING: DESIGN IDEAS ON PREFERENTIAL TREATMENT AMONG STREAMS +

Janardhan R. Iyengar

CIS Department, University of Delaware
Newark, DE - 19716
iyengar@cis.udel.edu

Sunil Samtani

Telcordia Technologies Inc.
Morristown, NJ – 07960
ssamtani@research.telcordia.com

1. BACKGROUND

With the introduction of multistreaming in the Stream Control Transmission Protocol (SCTP), Future Combat Systems (FCS) networks are presented with a new transport layer mechanism, for transmitting multimedia, which is markedly superior to transmission over UDP or TCP. SCTP provides for logical demarcation of data within an application transfer through multistreaming. Streams were originally designed to prevent head-of-line blocking at the receiver. Such head-of-line blocking can be observed during an application transfer of multiple, independent objects through a single TCP connection. One could envision using multiple TCP connections for transferring the different application objects, but such a design can have deleterious effects on the network [Balakrishnan et al., 1999].

Conceptually, multistreaming provides an aggregation mechanism for transferring different objects belonging to the same logical application session, such as a multimedia session. Preliminary work has shown the performance benefits of transferring video data over SCTP as against over TCP [Balk et al., 2002]. Since all these different objects follow the same physical path through the network from the server to the client, the round-trip estimates and the end-to-end available bandwidth estimate (the congestion window, *cwnd*) as probed by the congestion control algorithms are shared among the different streams. Such sharing of congestion information has been shown to have significant benefits [Balakrishnan et al., 1999].

2. PREFERENTIAL TREATMENT AMONG STREAMS

As shown in Figure 1, the "X" mark identifies a target in a reconnaissance mission. The application can partition the map into smaller pieces (e.g., region inside the turquoise square), and transmit the partitioned data in multiple streams back to command and control. We wish to be able to mark the streams with different ToS (Type of Service) bits or DSCP (Differentiated Services Code Point) markings so that the network can treat these streams differently. The stream that transports the "X" mark should be marked with the highest priority so the packets experience minimal latency and loss. Losses experienced by other map pieces should not affect this piece. In this example, the target information is more relevant to the user when received quickly, and the user need not receive all information to make a critical decision. As illustrated above, we wish to provide the service of being able to send data with request for preferential treatment from the network for different parts of a transfer (different streams) to the application. With all of the benefits offered by SCTP multistreaming, its current design falls short of being able to provide such prioritization among streams. Preliminary results and related work [Akella et al., 2001] have shown that the benefits of preferential treatment among streams are lost when such *false sharing* of congestion information exists. A closer inspection shows that the overall reduced throughput is due to a variety of reasons:

Figure 1: Illustration of Multistreaming: Reconnaissance Mapping

(1) The stream receiving a lower level of service from the network may experience more losses. These losses subsequently influence the entire transmission since the

+ Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

cwnd is common, thus reducing overall throughput. (2) Unnecessary Fast Retransmissions: Reordering introduced due to the possibly different forward path delays between the different ToS flows could result in spurious fast retransmissions. Such spurious retransmissions would cause the sender to wrongly infer congestion, thus causing unnecessary reduction of the *cwnd*. (3) Errors in roundtrip time estimate: Since the streams with different ToS marking may experience different delays at the routers.

3. CURRENT WORK

As of today, SCTP is the only transport that provides logical demarcation of data within an association or end-to-end connection. SCTP does so through the provision for multistreaming, but fails to provide ToS marking per stream. No good design or API exists today which an application can use to ask the network for preferential treatment of logically separate parts of an application transfer. We are currently working on a design that extends SCTP to handle such preferential treatment among streams. Our design goals include being able to adapt to highly dynamic network environments, such as the network bottleneck for all the flows being the same ToS unaware router, as may well be the case in FCS networks.

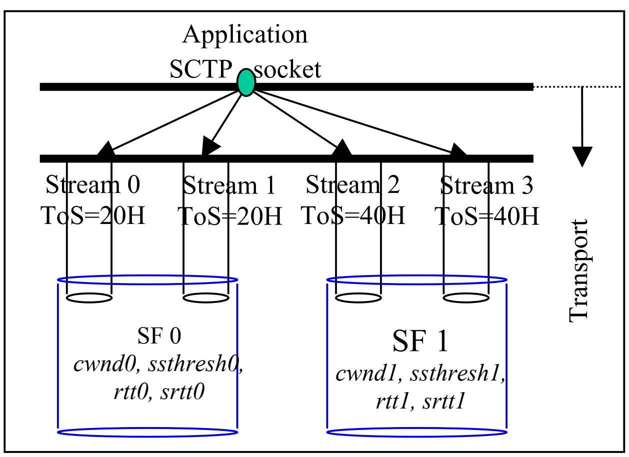

Figure 2: Sub-association Flow Schematic

The design shown in Figure 2 tries to separate the data flow within an association into separate *Sub-association Flows (SF)*, each SF having its own set of congestion control parameters. We are implementing this design using an SCTP stack from Siemens and the University of Essen. As shown in Figures 3 and 4, preliminary results from experiments are encouraging. Figure 3 shows transfer latency with unmodified SCTP and Figure 4 with our design modifications to SCTP. One of the ToS flows in both experiments was subjected to 5% loss while the other was subjected to no loss.

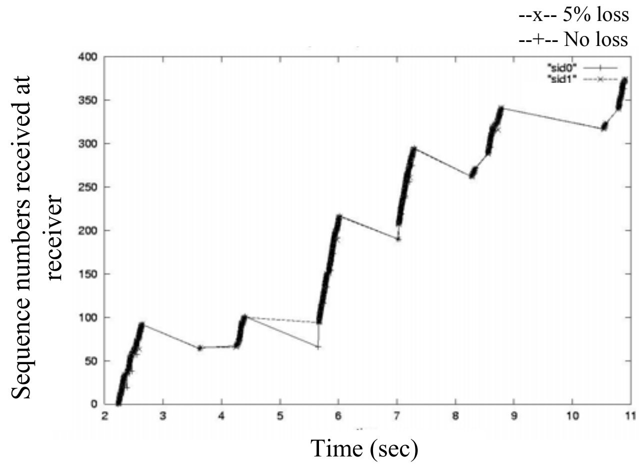

Figure 3: Transfer latency with unmodified SCTP, and per-stream ToS marking

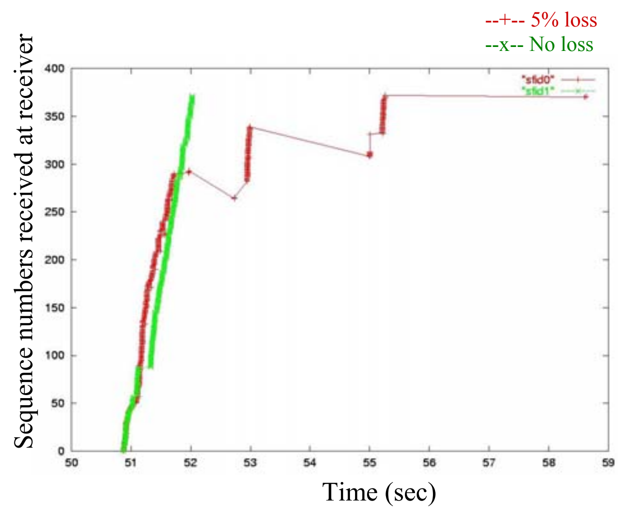

Figure 4: Transfer latency with modified SCTP, and per-stream ToS marking

REFERENCES

[Balakrishnan et al., 1999] H. Balakrishnan, H. Rahul, S. Seshan. An Integrated Congestion Management Architecture for Internet Hosts. In ACM SIGCOMM, Cambridge, MA, September 1999.

[Balk et al., 2002] A. Balk, M. Sigler, M. Gerla, M. Y. Sanadidi. Investigation of MPEG-4 Video Streaming over SCTP. Proc. SCI2002, Orlando, July 2002.

[Akella et al., 2001] A. Akella et al. The Impact of False Sharing on Shared Congestion Management. CMU SCS Technical Report CMU-CS-01-135.

DISCLAIMER