Performance of Stream Control Transmission Protocol (SCTP)

Mohammed Atiquzzaman, Ph.D.
School of Computer Science
University of Oklahoma.

Email: atiq@ieee.org
Web: www.cs.ou.edu/~atiq

Tohoku University, Sendai, Japan.
Aug 6, 2002.

Introduction

- TCP is the main transport protocol in the Internet protocol suite
- Original TCP performed poorly in satellite networks
  - errors
  - long propagation delay.
- Many schemes for enhancing TCP for satellite networks.
- IETF developing Stream Control Transmission Protocol (SCTP) for PSTN signaling over IP.

Open question: How good is multistreaming?

Objectives

- Evaluate the performance of SCTP multistreaming.
- Suitability of SCTP for wireless networks
- Performance comparison of SCTP and TCP.

Outline

- SCTP
  - Multistreaming
  - Multihoming
- Performance of multistreaming

Stream Control Transmission Protocol

- SCTP (RFC 2960) being developed by IETF as a transport protocol for PSTN signaling.
  - Reliable: retransmission of lost packets, ack of packets.
  - Non-duplicated service: uses sequence numbers.
  - In-order delivery: re-sequencing at the destination.
- Transport layer protocol which operates on top of an unreliable connectionless network layer such as IP.
  - Transparent to IPv4 or IPv6
- Key features:
  - Multistreaming – multiple streams per association
  - Multihoming – multiple IP addresses per host

SCTP in the protocol stack

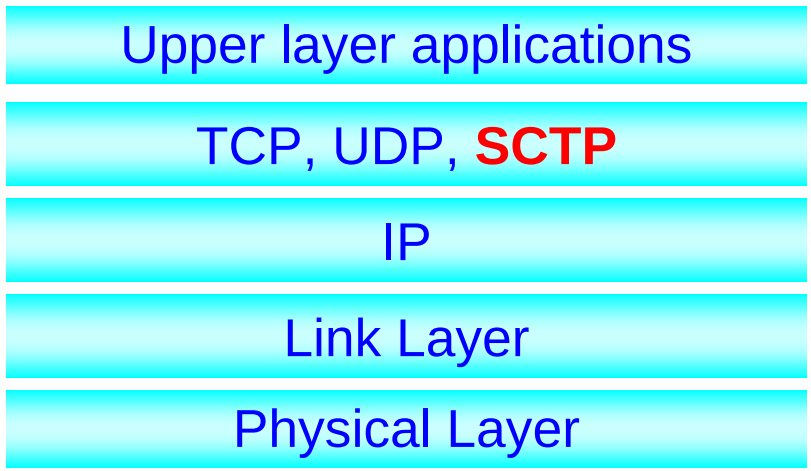

- SCTP accomplishes multistreaming by creating independence between
  - data transmission (uses Transport Sequence Number)
  - data delivery (uses Stream Sequence Number)

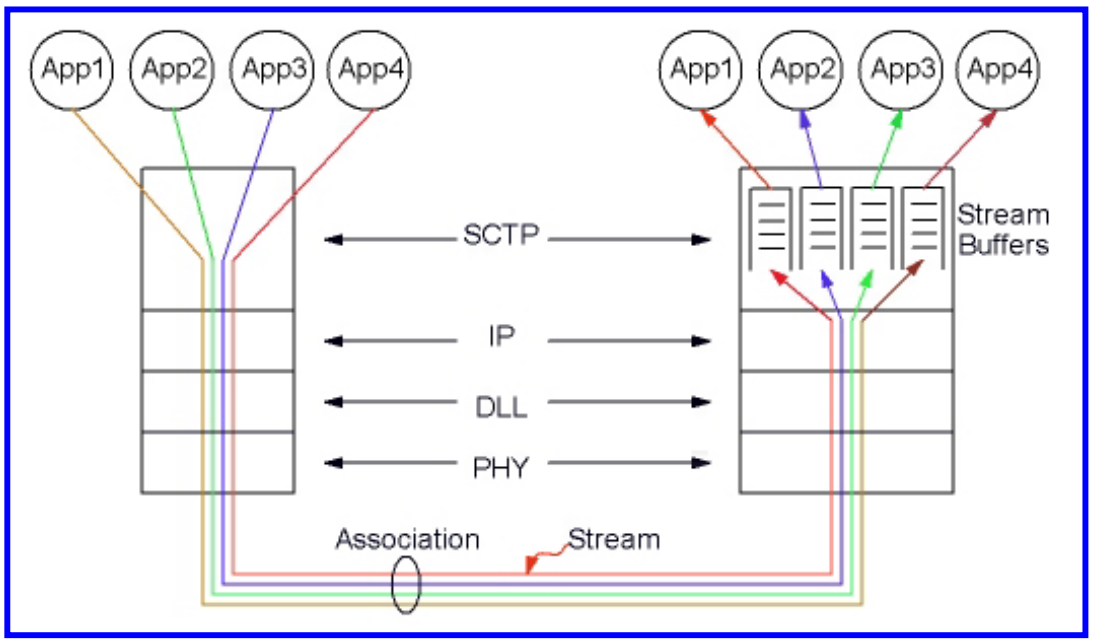

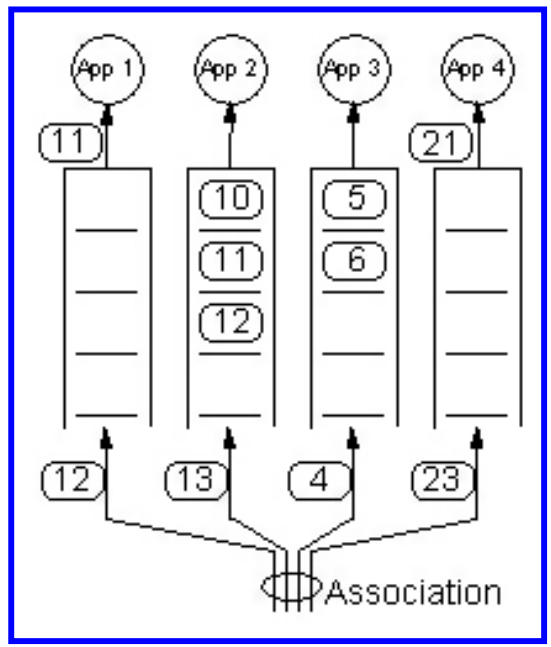

- Supports multiple IP addresses in an *association.*
- Requires multiple Network Interface Cards – already quite common in laptops !!

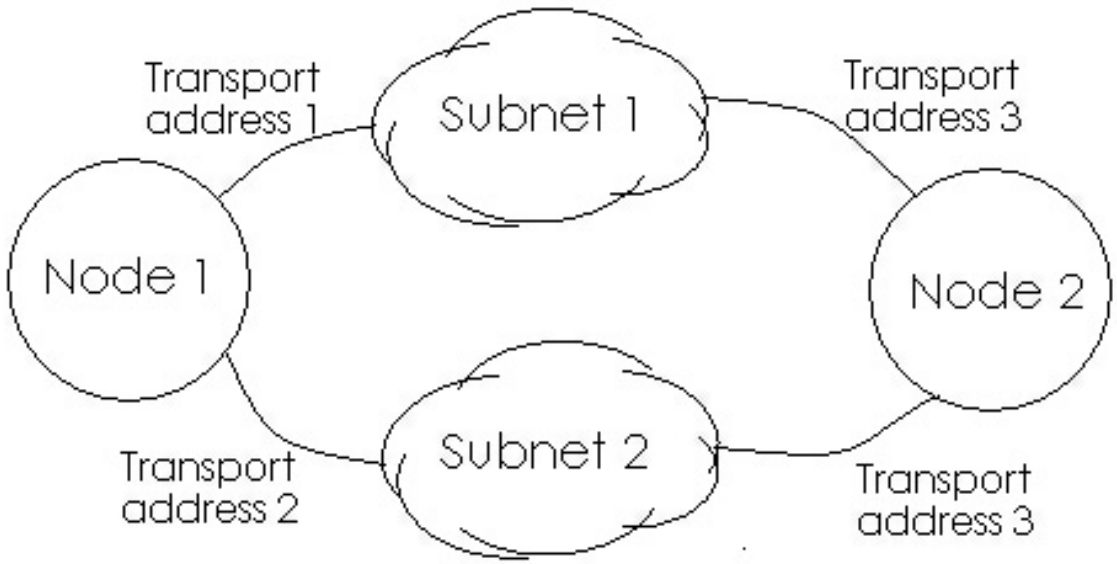

SCTP Packets

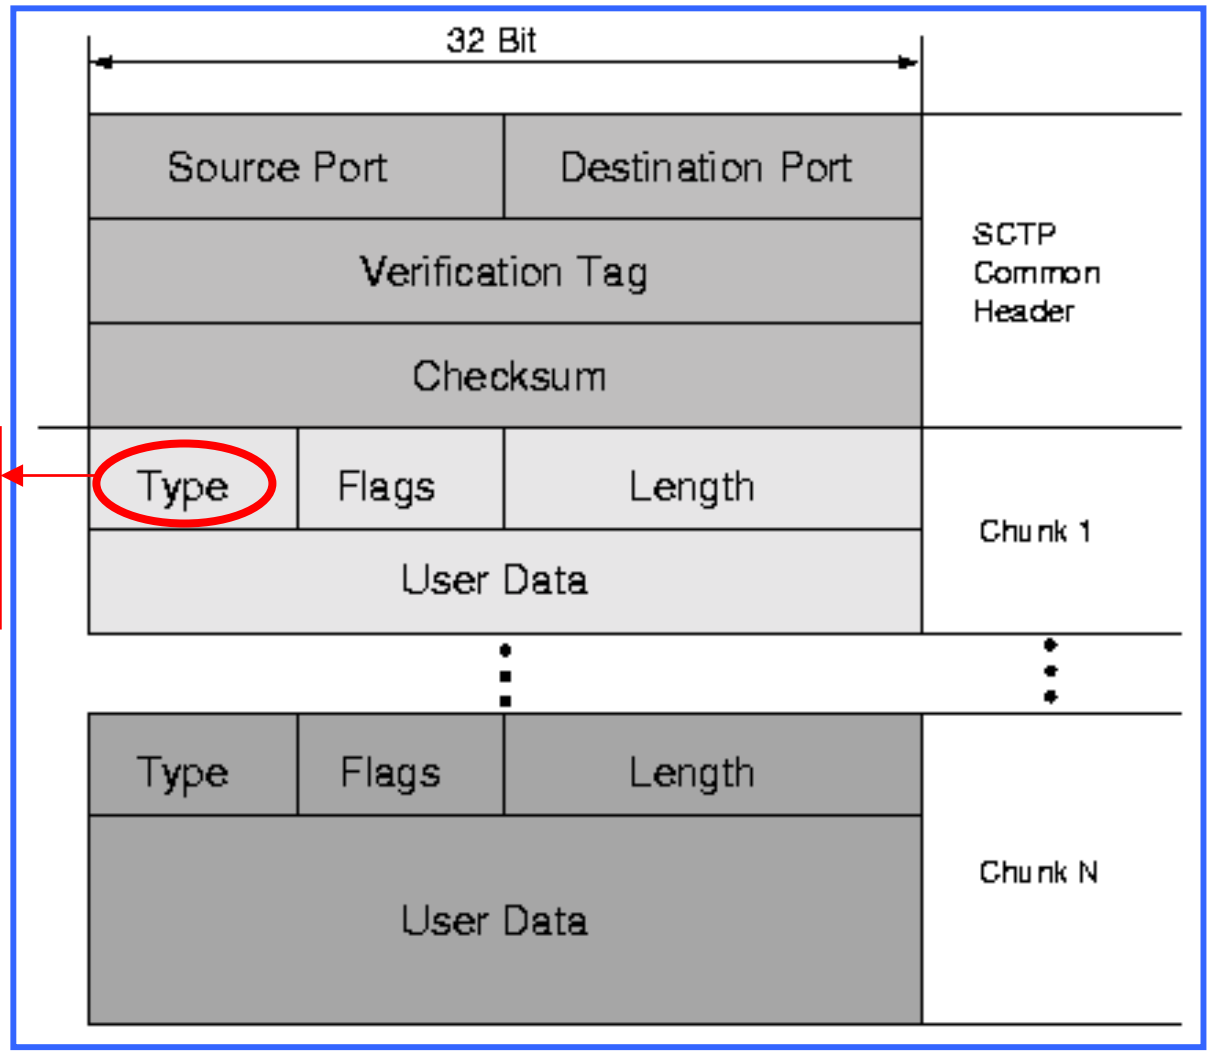

Chunk Type

- Payload
- SACK, etc.

Chunk Type: Payload

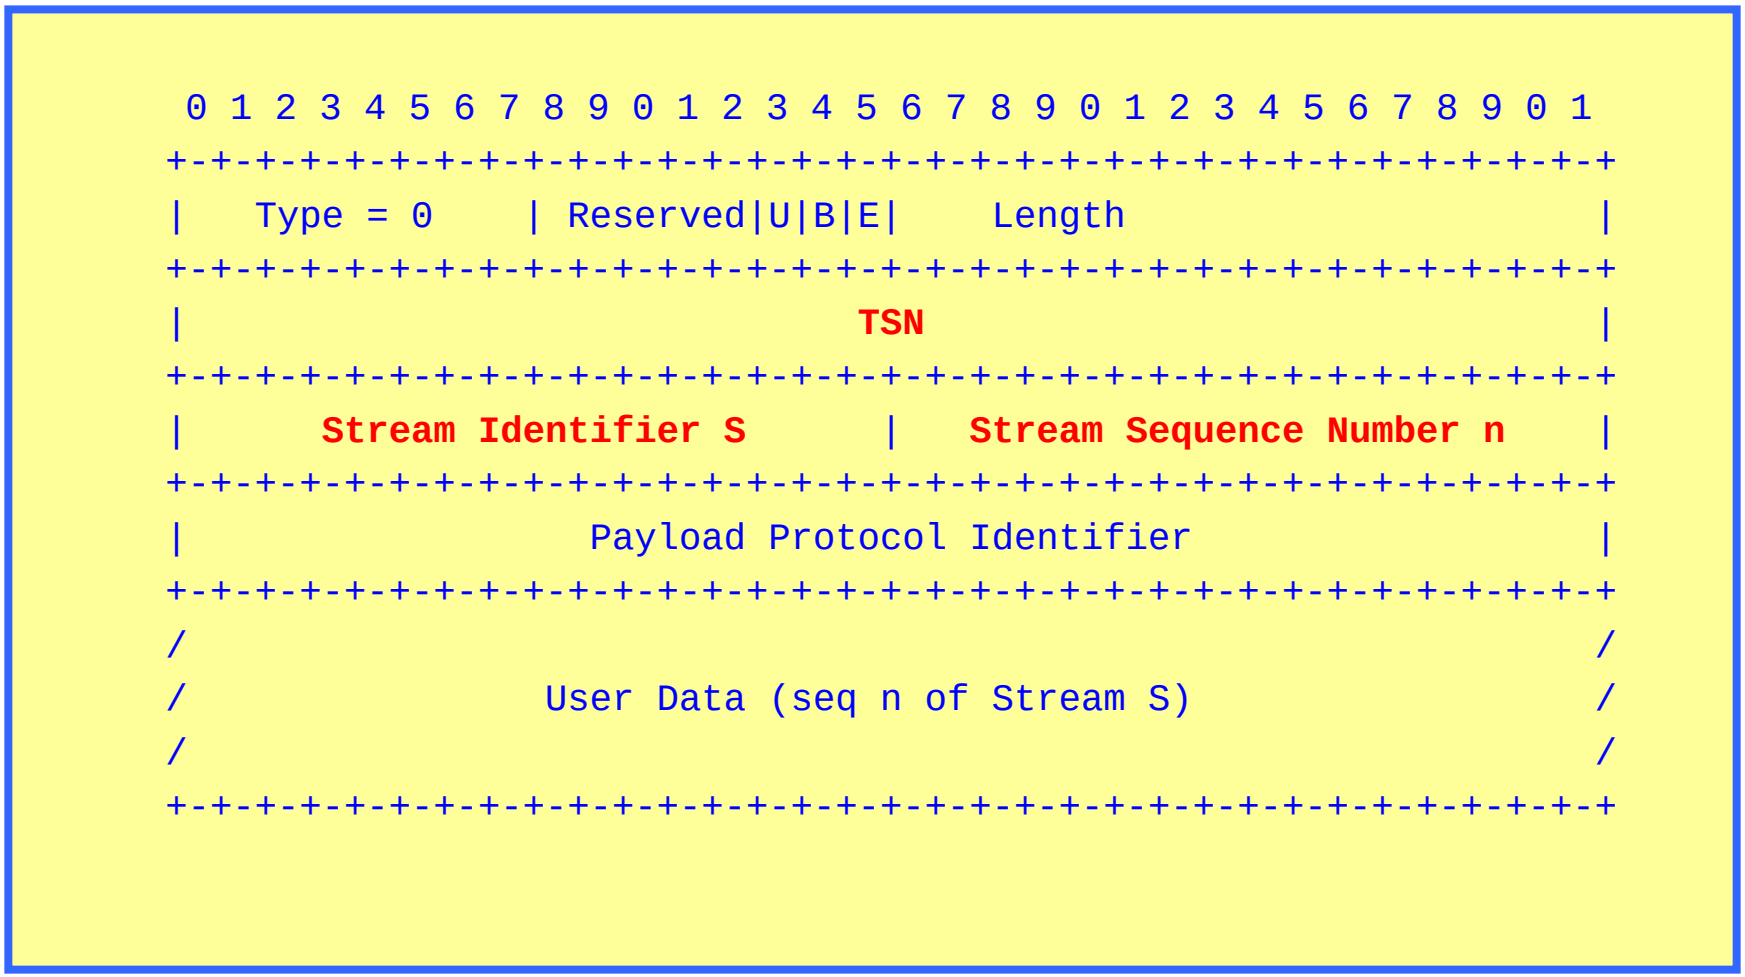

Chunk Type: SACK

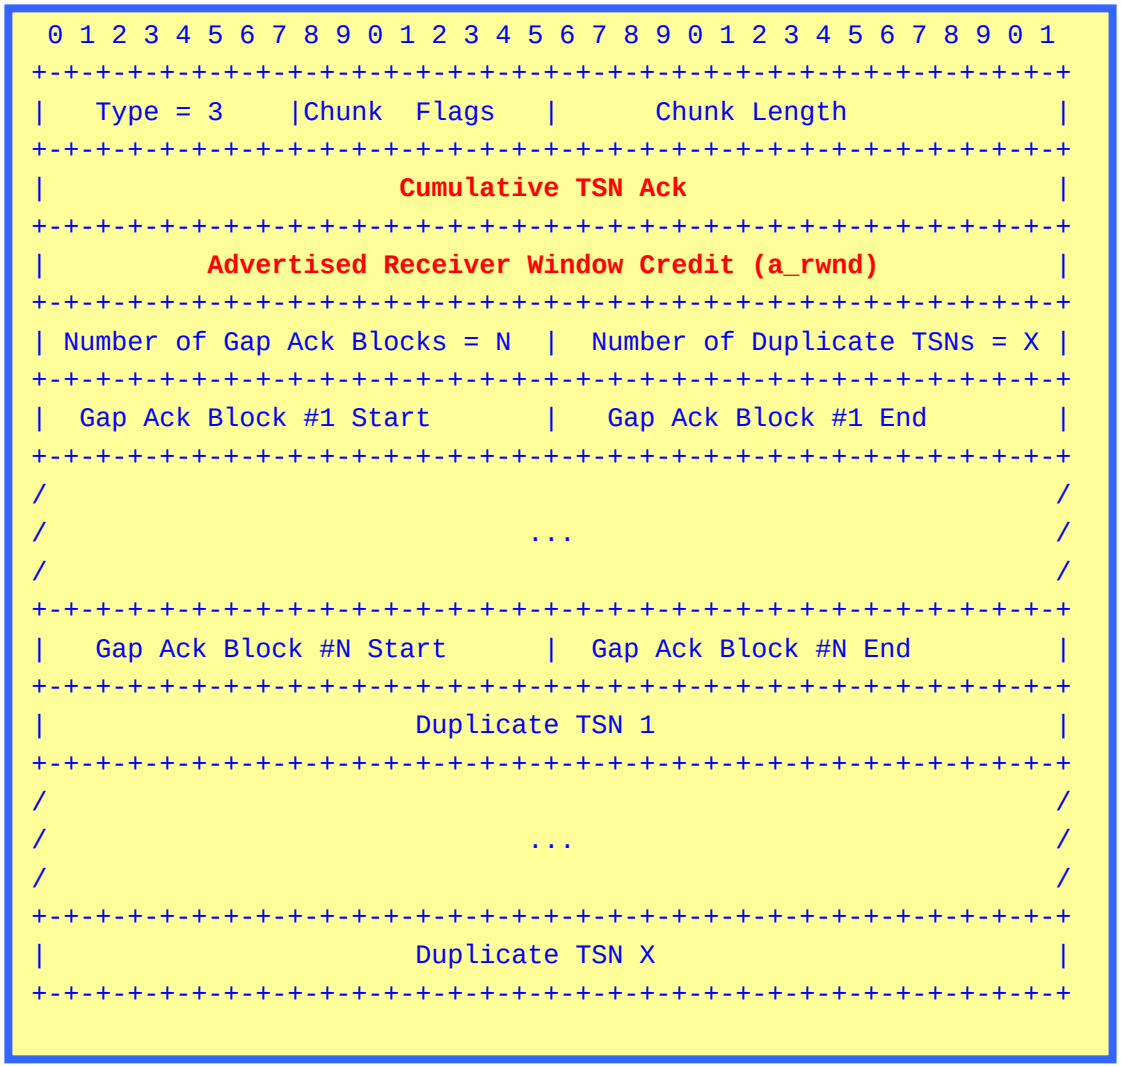

SCTP Congestion Control

- SCTP congestion control is similar to TCP congestion control.
  - Enables seamless introduction of SCTP into IP networks.
- SCTP is rate adaptive similar to TCP.
    - Slow Start, Congestion Avoidance, Fast Retransmit
    - Fast Recovery is implemented, but in a slightly different way than TCP.
- Differences with TCP
  - Number of bytes acknowledged is used to increase *cwnd.*
  - SACK is mandatory
    - Unlimited number of Gap Ack Blocks in SACK
  - No explicit fast recovery phase

Performance/Advantages of

SCTP Multistreaming

ns-2 Simulation Setup

- *Ftp* traffic.
- Packets of *fixed length* (one MTU).
- Upper layer at destination is always ready to accept data.
- Association consists of a number of streams

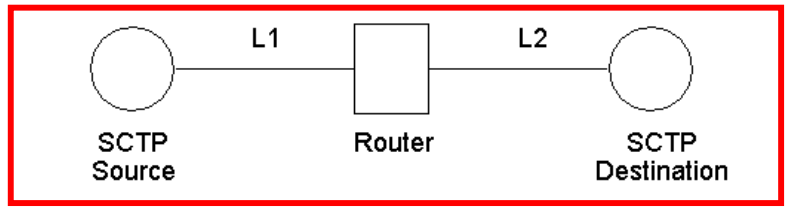

Link Delay: L1+L2 = 260 msec
Receiver buffer size = *B*

Performance Metrics

- Goodput
- Optimal receiver buffer size
  - as a function of packet **error probability (e)**.

Packet Plot: No loss

s = 4, e = 0, B = 15K

- No packet loss → no blocking at receiver.
- *cwnd* increases until *B.*
- Goodput limited to *B/MTU* packets per RTT.

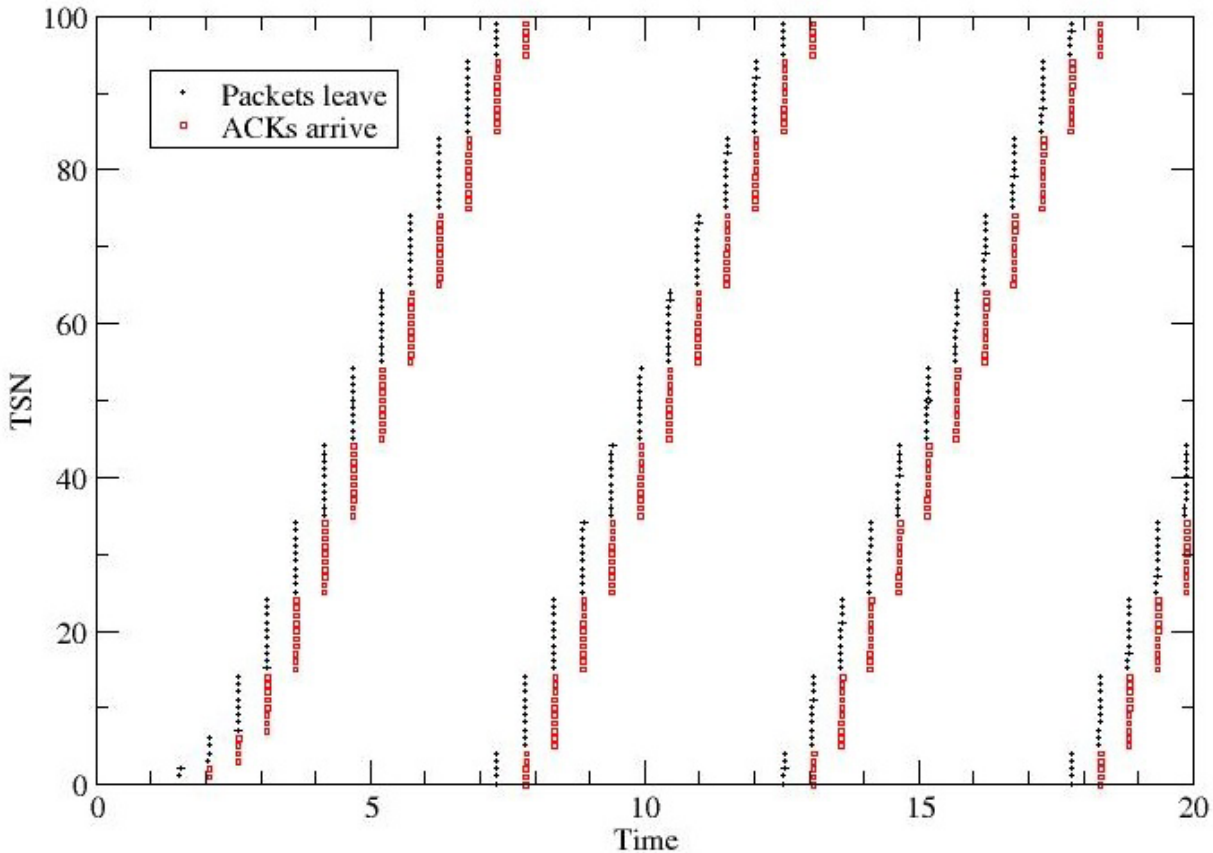

Goodput: No loss

- Since goodput is limited to *B/MTU* packets every RTT; it increases linearly with *B.*

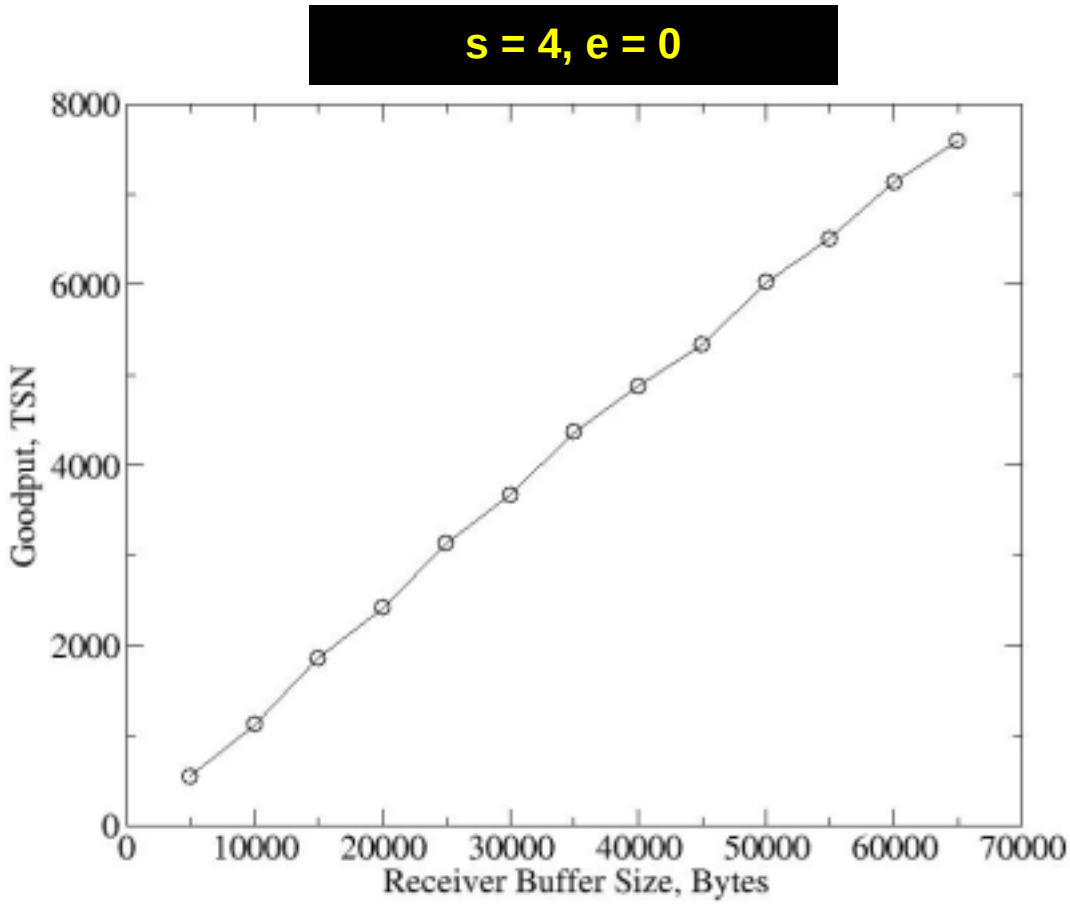

Packet plot: Congestion Control limited

- Poor goodput when receiver buffer is not a constraint
  - Long delays in Retx while waiting for DupAcks

    +
  - drop in *cwnd* due to Retx

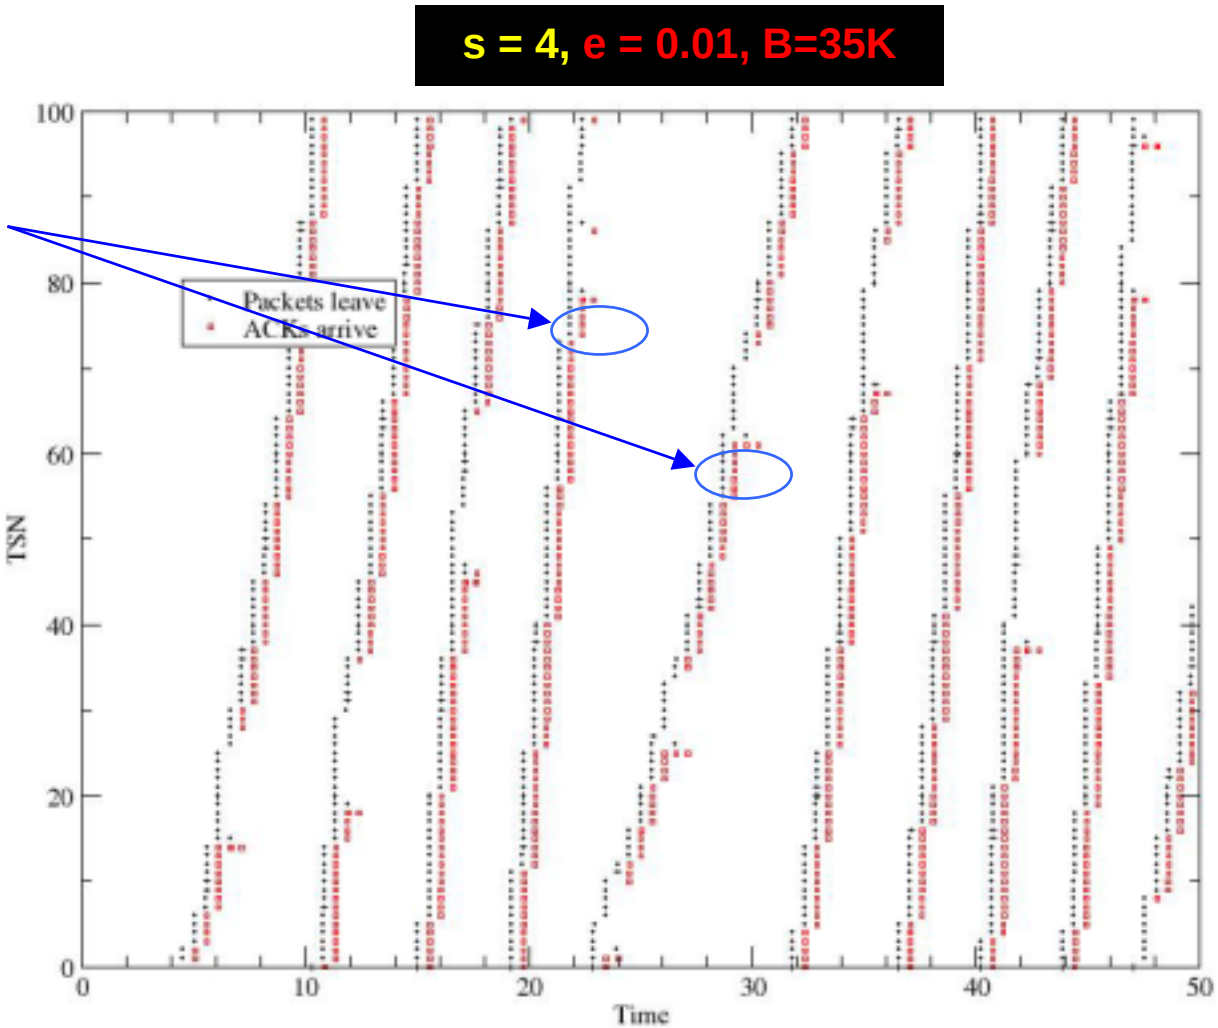

cwnd and *a_rwnd* with errors: Receiver buffer limited

- *B*=15K → throughput constrained by receiver buffer

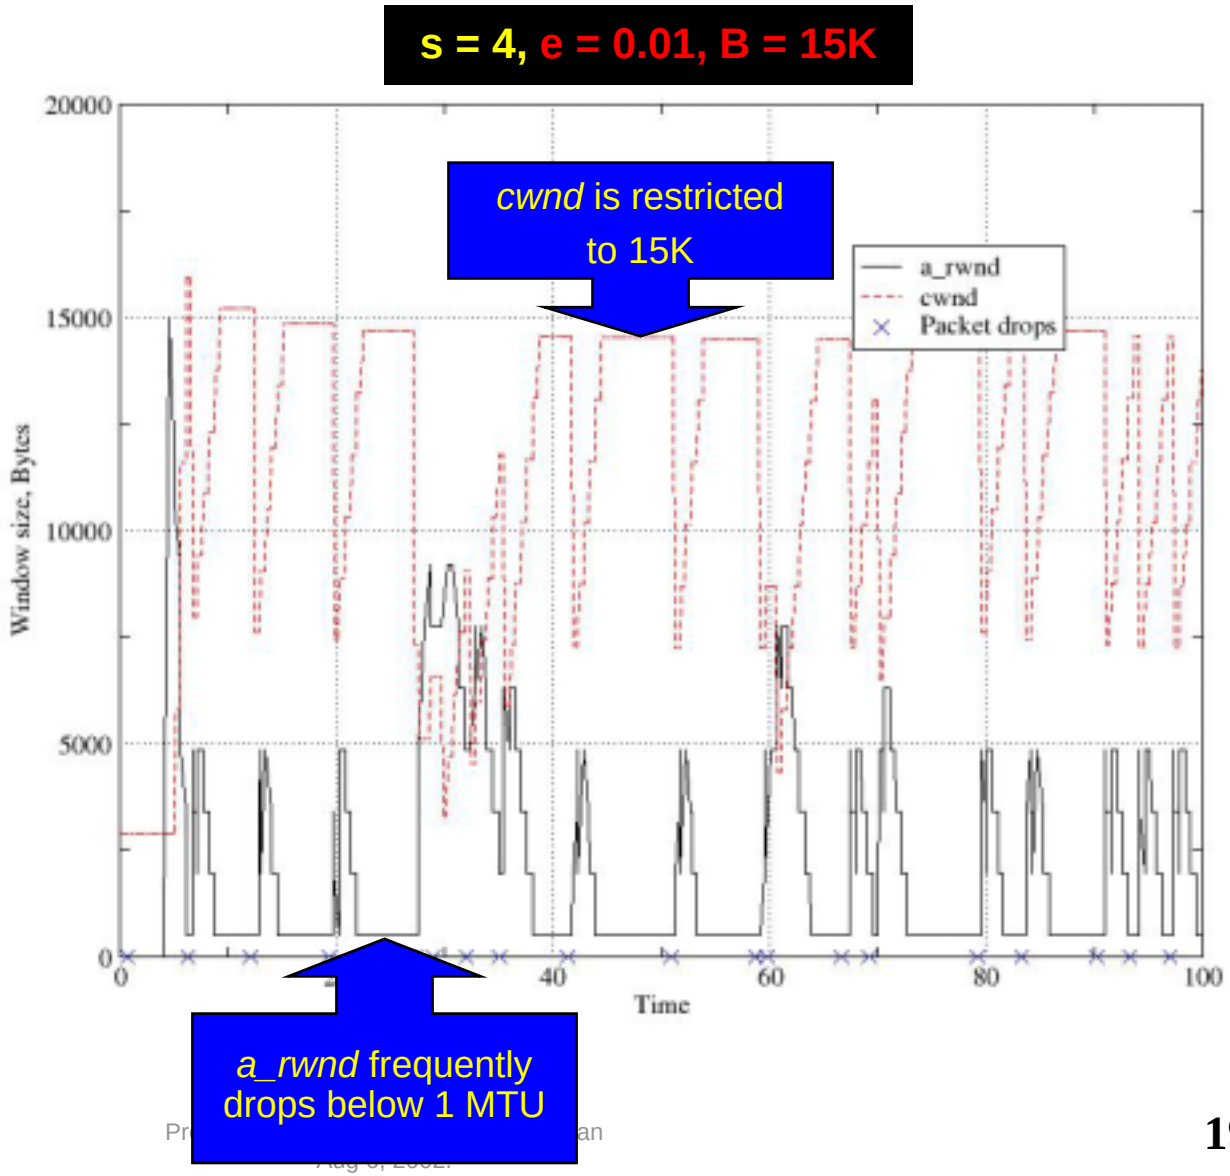

cwnd and *a_rwnd* with errors: Congestion control limited

- *B*=35K → throughput constrained by congestion control of SCTP

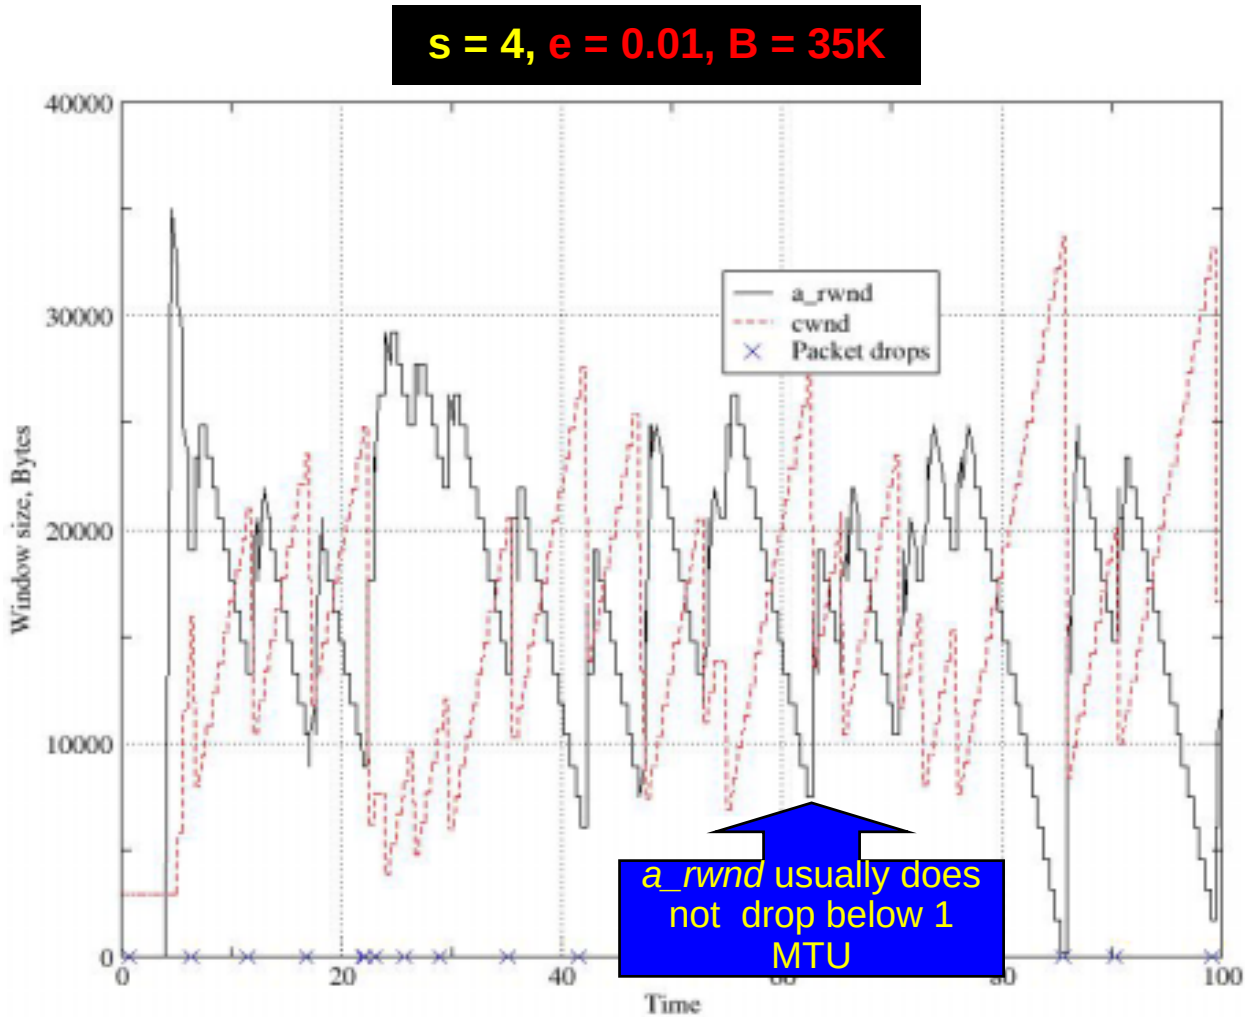

Goodput with errors

s = 4, e > 0

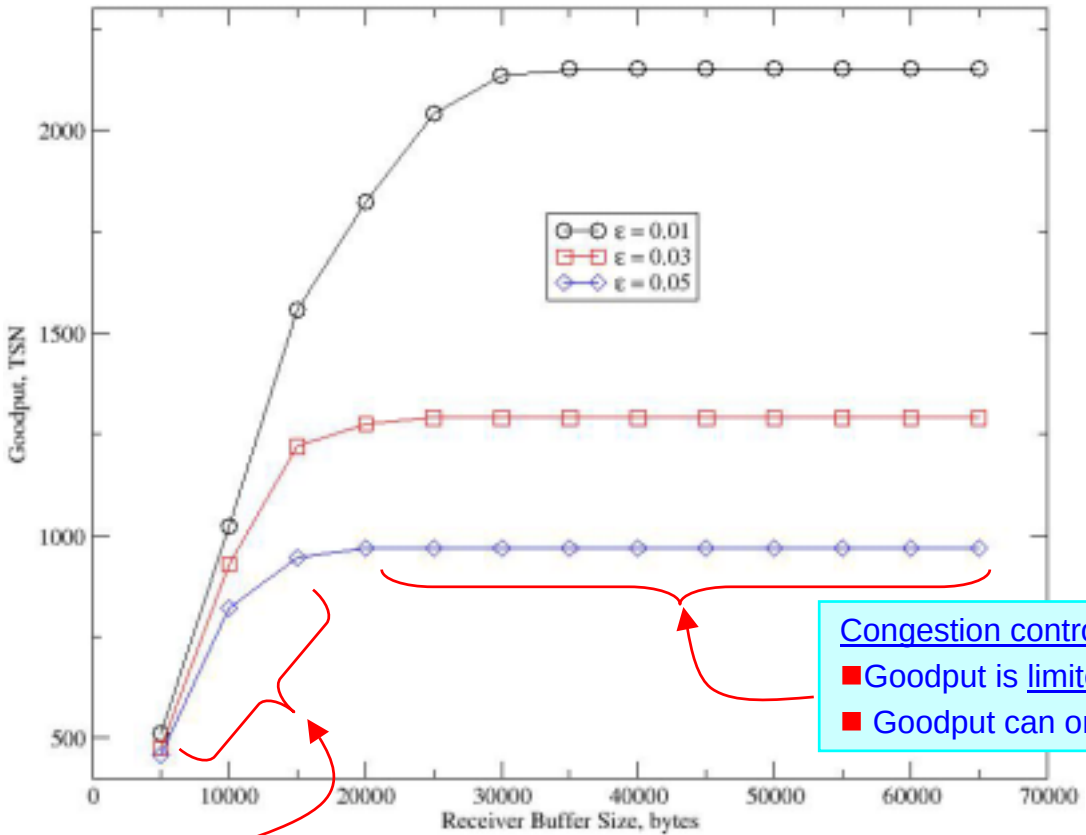

Congestion control limited:

- Goodput is limited by the congestion control of SCTP.
- Goodput can only be increased by lowering the error rate

Receiver buffer limited:

- Goodput increases as *B* increases when the goodput is constrained by B (*a_rwnd* frequently drops below 1 MTU)

Goodput vs. Buffer size: One and four streams

- For small *B*
  Mutlistreaming results in “less” HOL blocking
  - goodput of 4-streams is higher than 1-stream.

- For large *B*

  Goodput is limited by congestion control

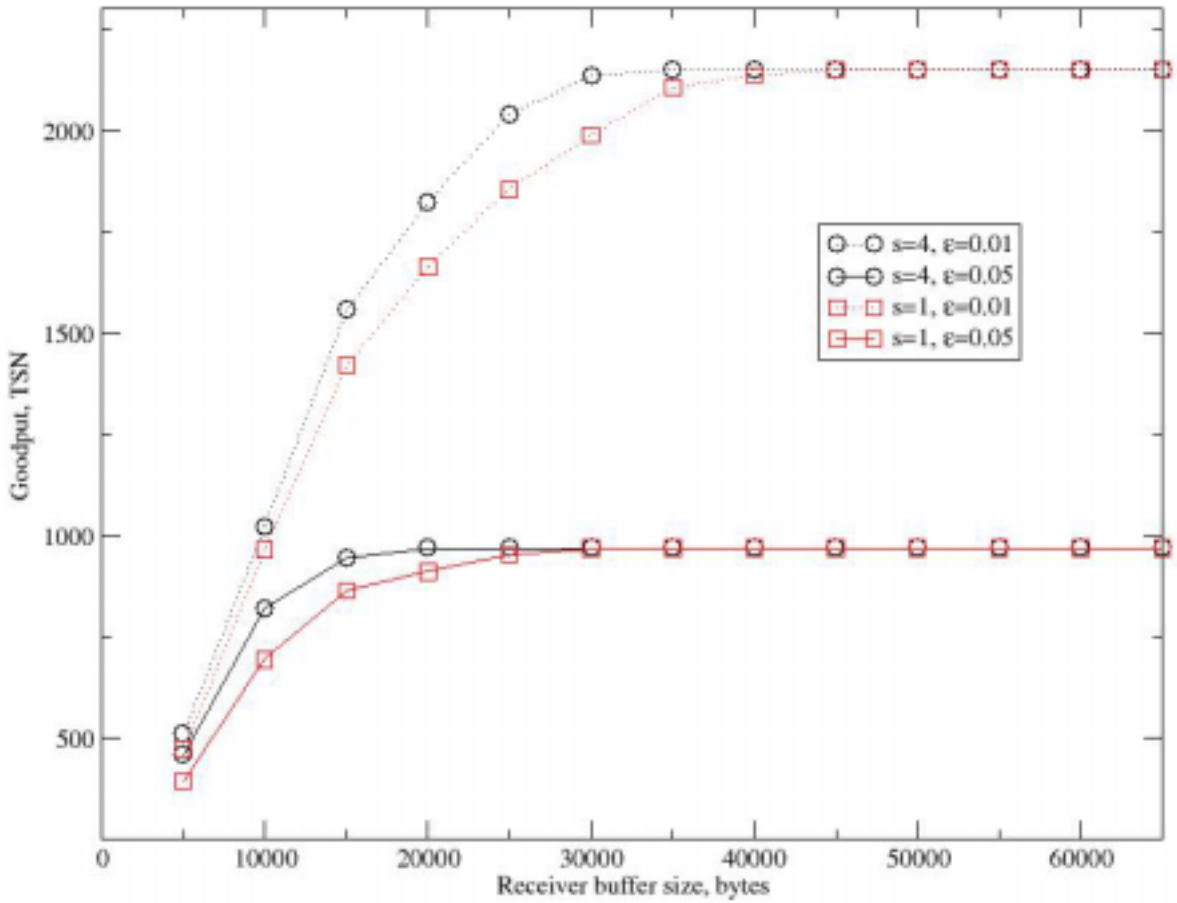

Multistreaming increases goodput for small receiver buffers.

Advantage of Multistreaming: High Throughput

B = 15K

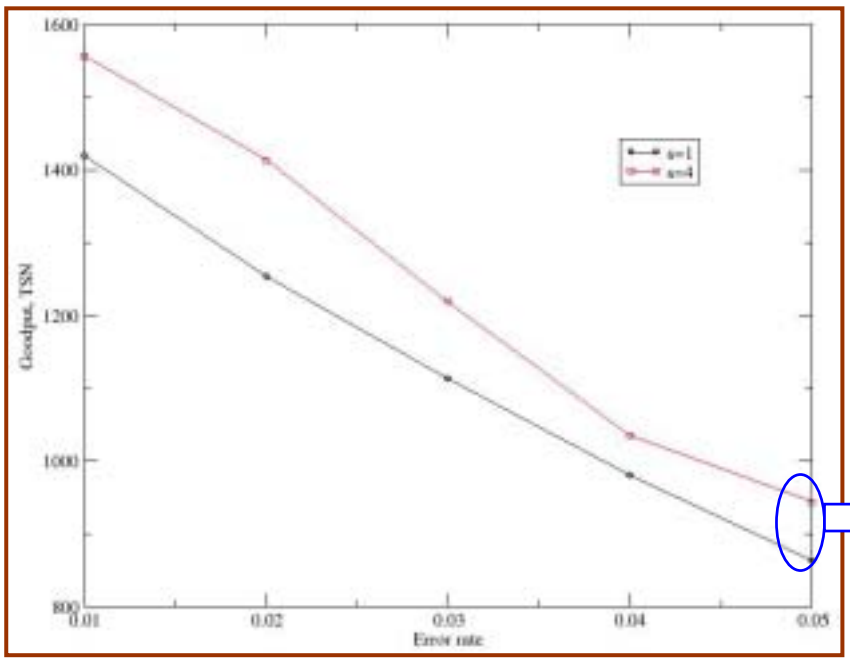

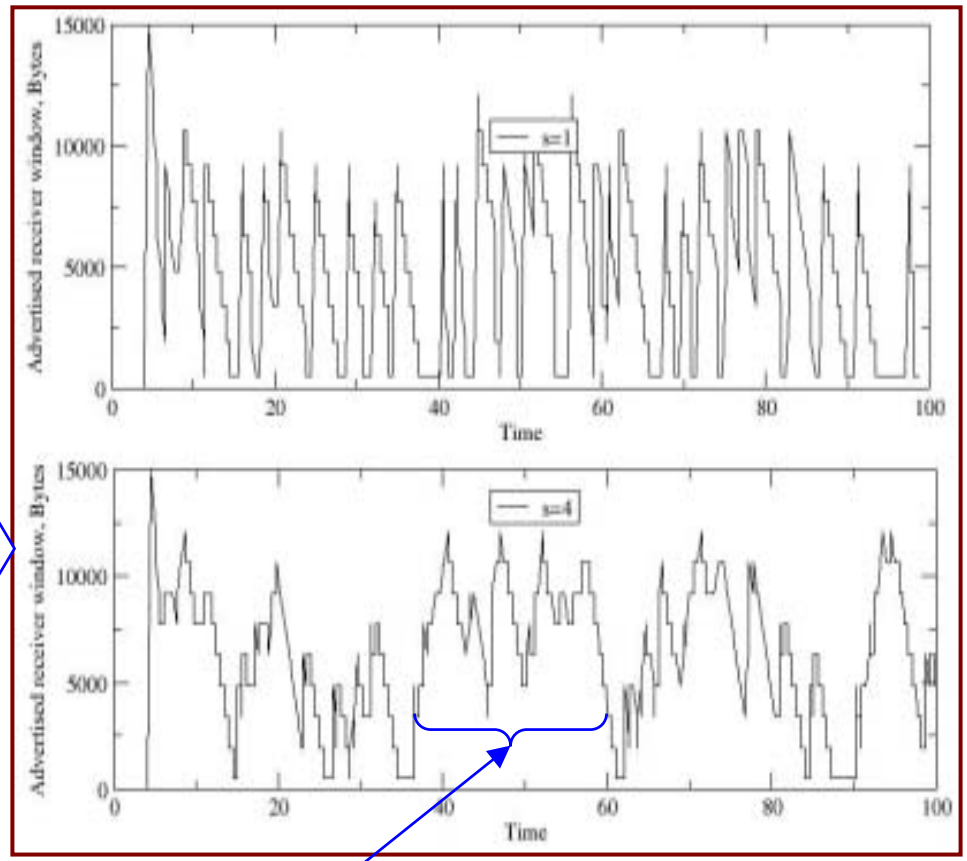

- "Small" Buffer size of 15K shows advantage of multistreaming.

HOL blocking is eliminated as evidenced by the fact that *a_rwnd* is, not a limiting factor

Optimal Buffer Size

Optimal Receiver Buffer Size

Buffer size beyond which *a_rwnd* never falls below 1 MTU

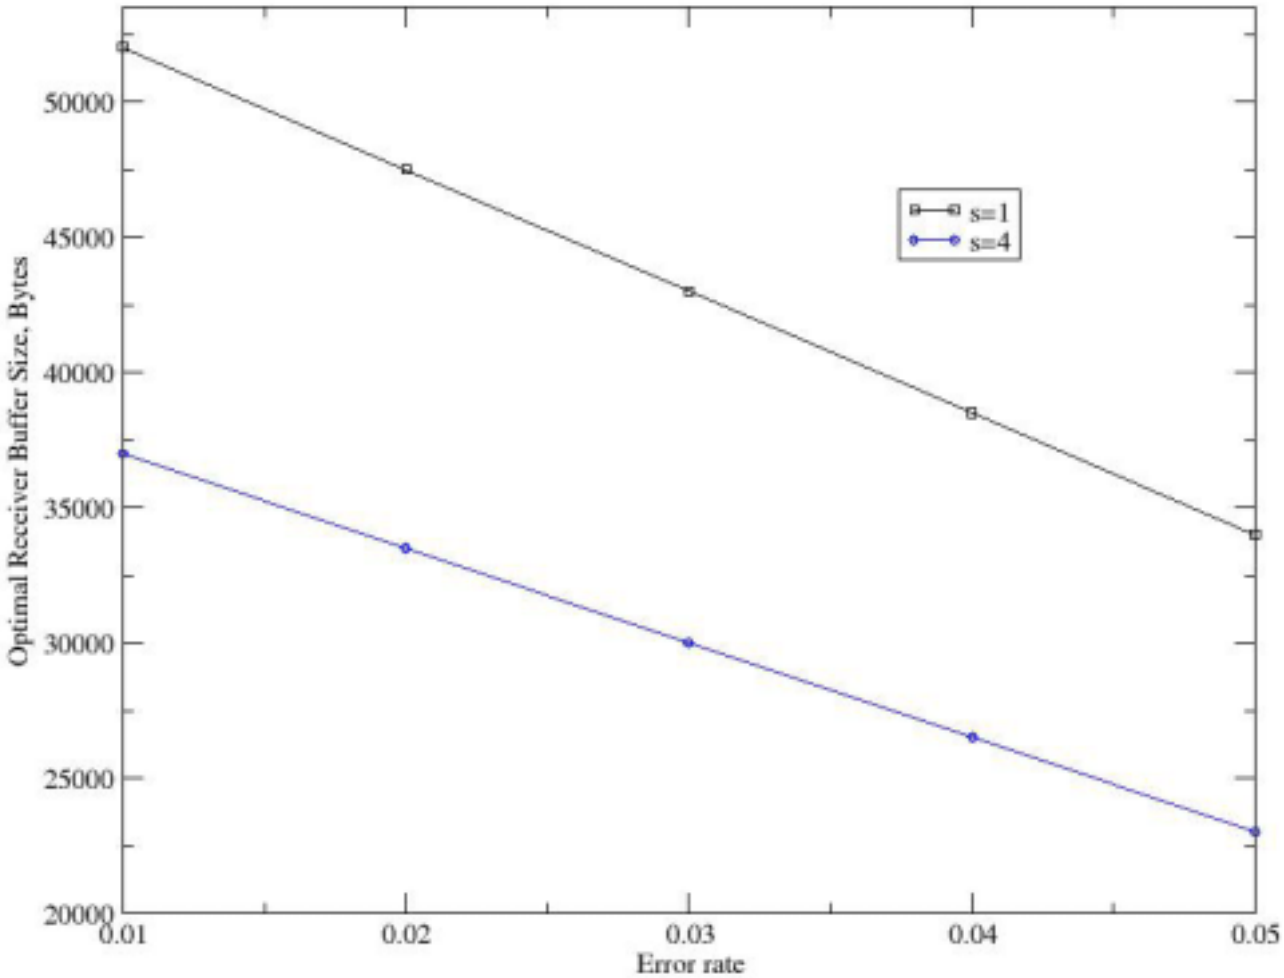

Multistreaming reduces receiver buffer requirements.

Conclusions

- Performance of SCTP multistreaming has been studied
- Multistreaming increases transport level performance for small receiver buffer size.
  - Wireless handheld devices, ex. Mobile and satellite networks.

Further information

- Acknowledgements
  - National Aeronatics and Space Administration (NASA)

- Further Information

Dr. Mohammed Atiquzzaman

atiq@ou.edu, (405) 325 8077

- These slides are available at

www.cs.ou.edu/~atiq