SCTP: A new networking protocol for super-computing

Mohammed Atiquzzaman
Shaojian Fu
Department of Computer Science
University of Oklahoma.
atiq@ou.edu

Networking Characteristics in Supercomputing

- High bandwidth links between upercomputers;
- Requires low delay in packet delivery;
- Many applications are based on message passing;
- Multimedia applications becoming popular.

SCTP: A new Transport Protocol

- TCP is the most popular transport protocol to connect supercomputers.

- Supercomputer applications can take advantage of many SCTP features:
  - Support for multiple logical streams to improve data transmission throughput;
  - Support for multiple network interfaces to achieve high availability;
  - More secure mechanisms to prevent threats such as Denial of Service (DoS) attack.

What is SCTP?

- Stream Control Transmission Protocol;
- Originally designed to support Telephone signaling messages over IP Networks;
  - Currently supports most of the features of TCP
- Standardized by IETF RFC 2960;
- Reliable transport protocol on top of IP

SCTP and TCP Features

- Both of them are reliable transport protocols;

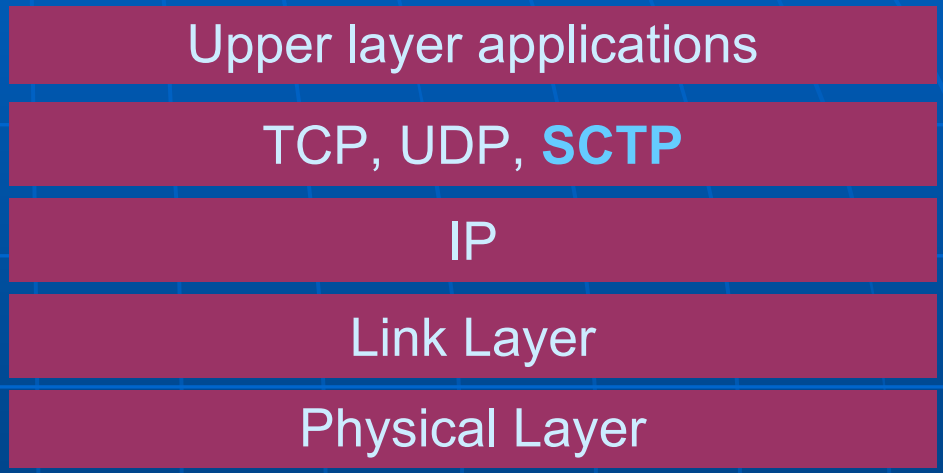

- Similar Congestion Control algorithms

  (slow start, congestion avoidance);

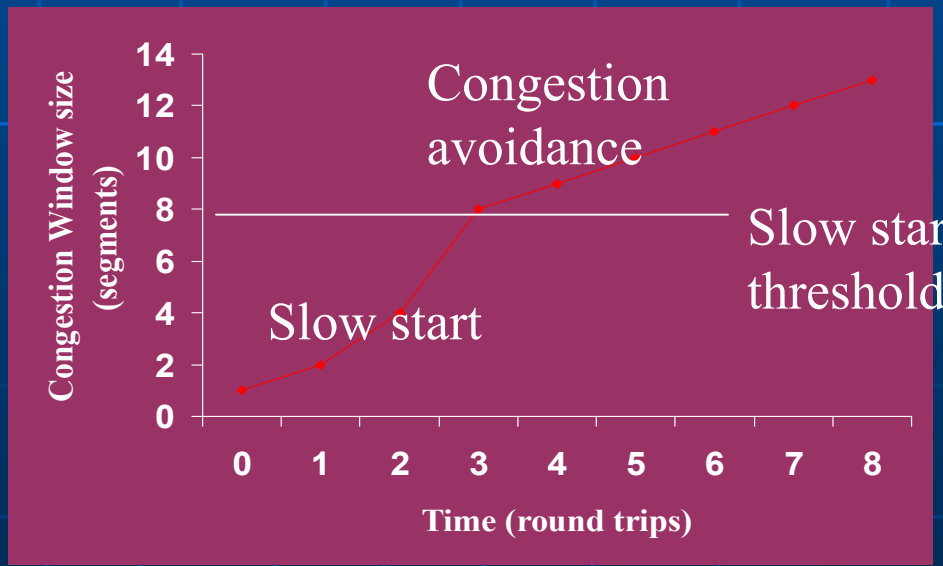

SCTP and TCP Differences

- SCTP has the concept of an association instead of a connection;
- Multiple Logical Streams in a Association;
- Support for Multihoming;
- SCTP is message oriented while TCP is byte stream oriented;
- Protection against DoS attacks;

Where does SCTP fit?

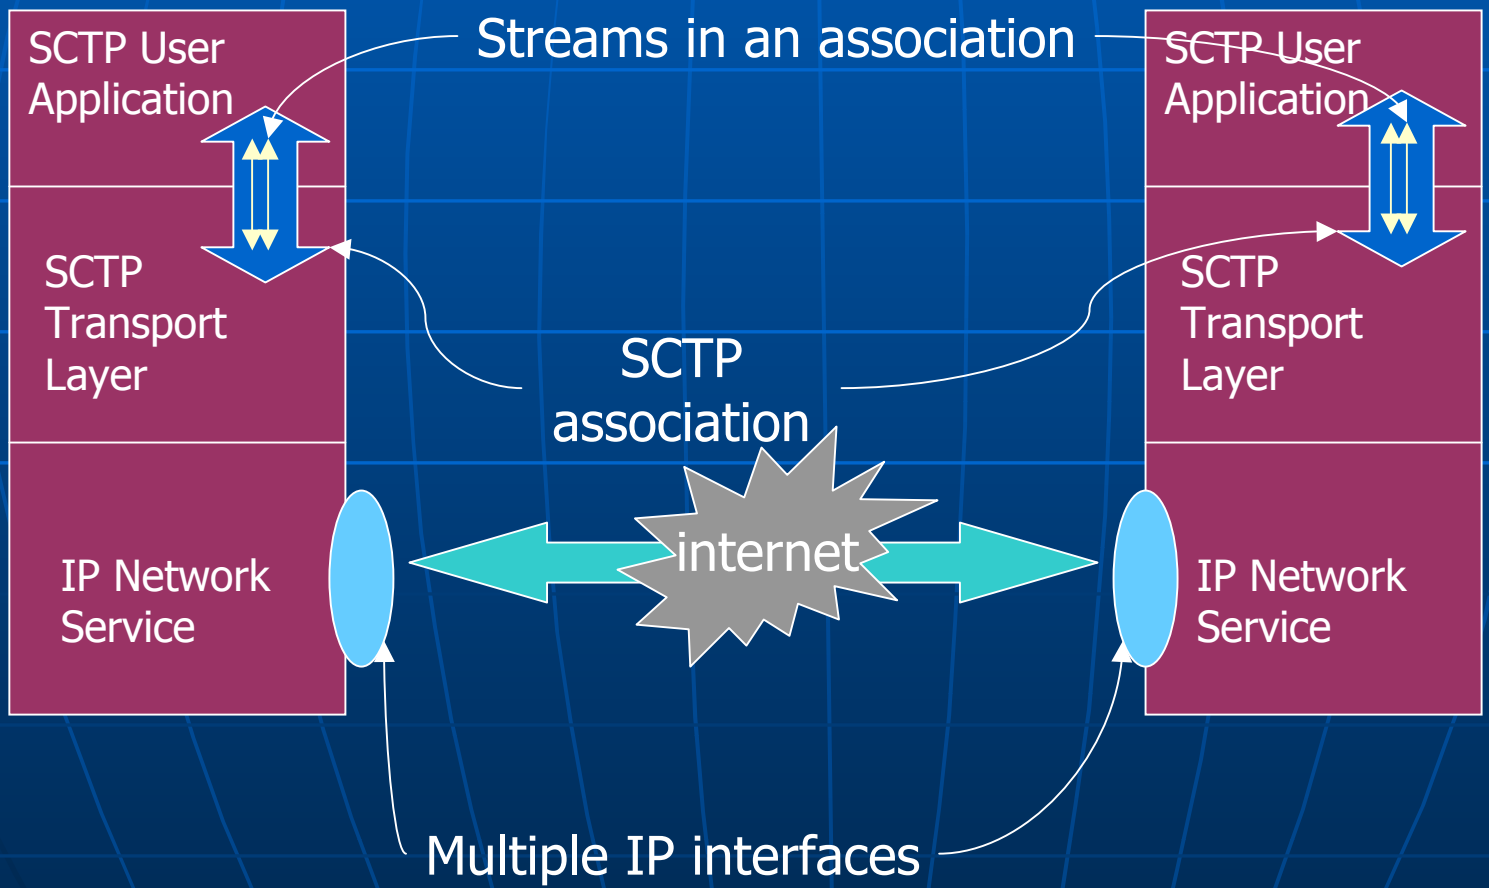

SCTP Multistreaming

- Multiple streams per association;
- Multiple streams prevent head-of-line blocking that occurs in a single stream scenario;
- Stream properties can be individually tailored to fit application requirements.

SCTP Multistreaming

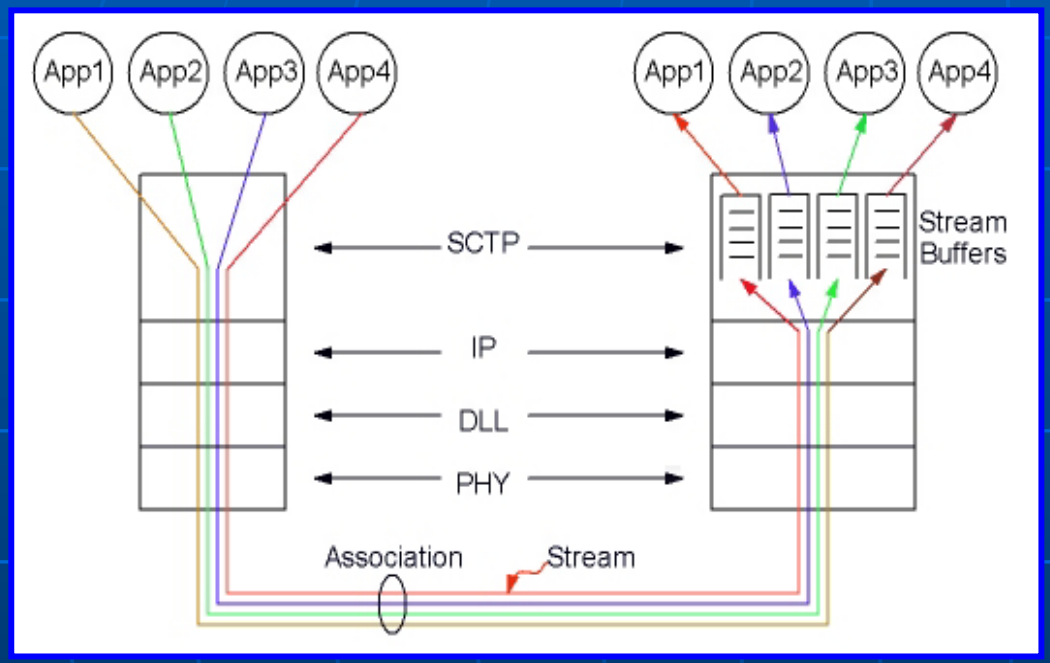

Multiple logical streams

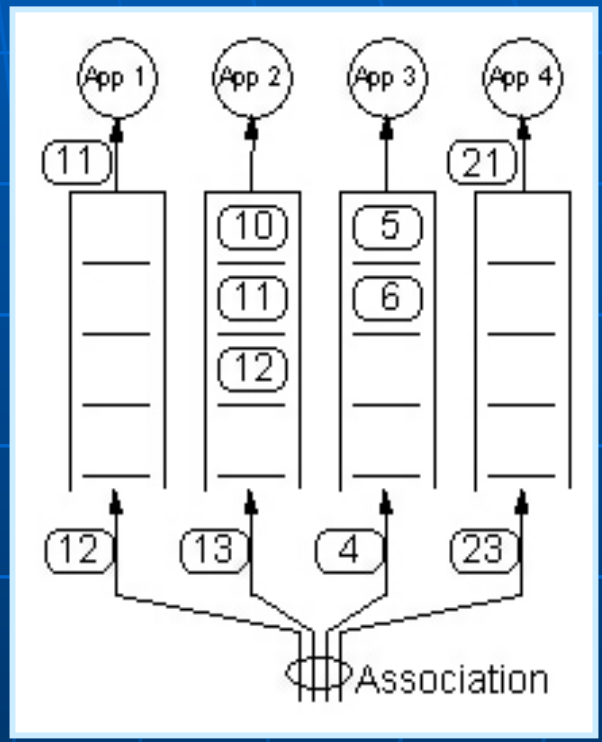

Head-of-line blocking

Benefit of Multistreaming

- Allows application to send multiple objects (text, images, audio) simultaneously. An object can delay the delivery of other objects;
- Improve the data transmission throughput under packet losses;
- Can help supercomputing applications involving multimedia objects.

Multihoming----Supports multiple IP addresses in an *association*.

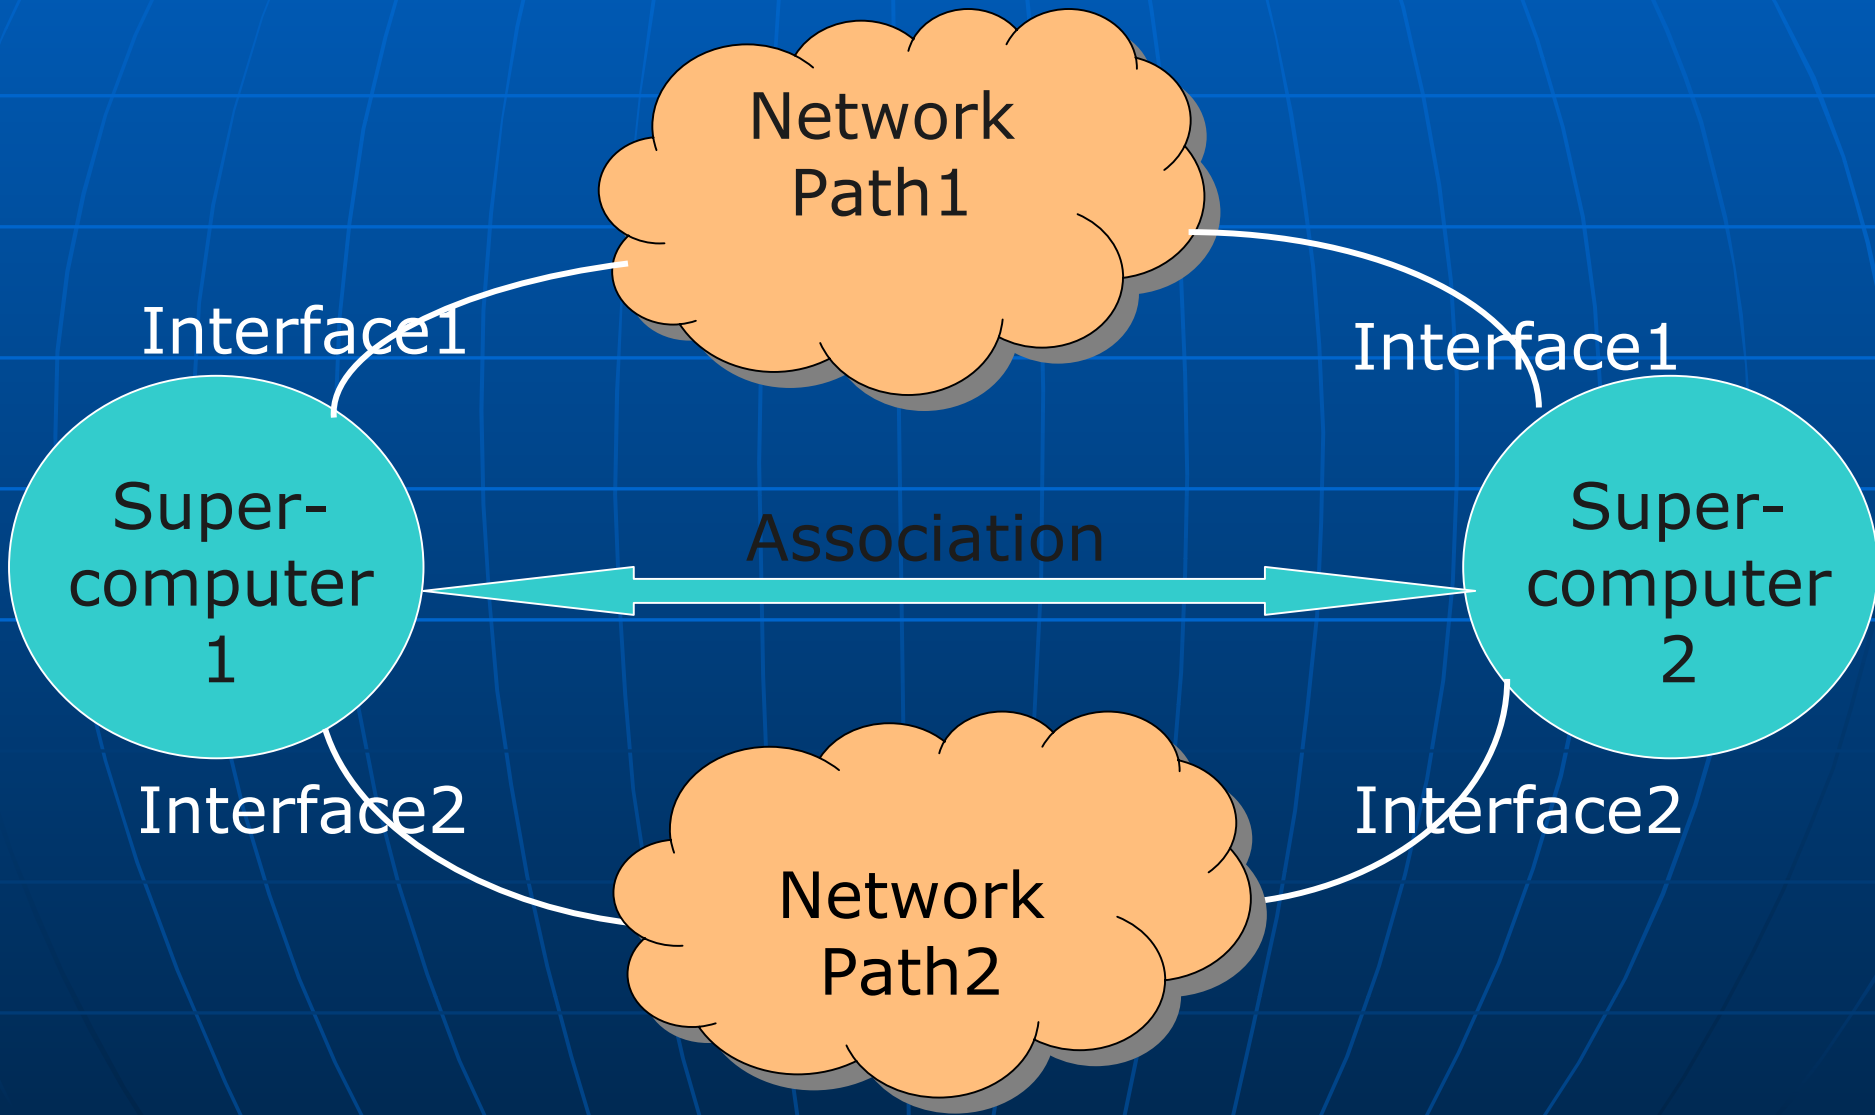

Message-oriented feature

- Communication between supercomputer is often message-oriented;

- TCP's stream-oriented nature is often an inconvenience for applications based on message passing;

- SCTP is message oriented:
  SCTP saves the programmers from doing framing at the application layer manually.

High bandwidth delay links

- Supercomputers are usually connected by high bandwidth-delay product links;
- For optimal performance keep link full
  - maximum TCP window size is 64KB
  - a 192KB bandwidth-delay product line is empty more than 60% of the time

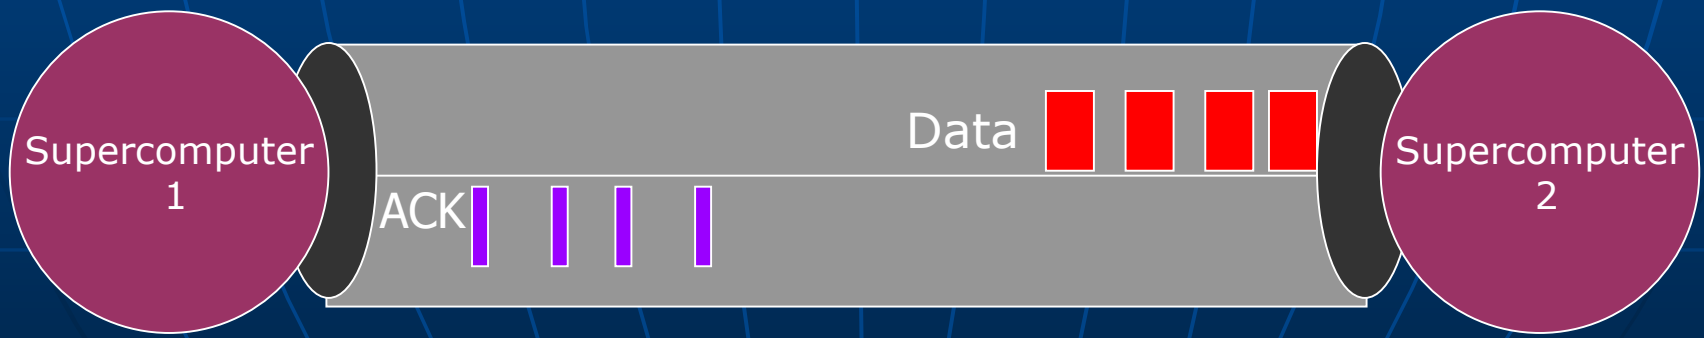

SCTP Large window support

- TCP supports windows up to 2^{16} bytes;
  TCP requires window scaling option (RFC 1323) to support large windows;
- SCTP has a natural support for large windows up to 2^{32} bytes to fill out the pipe.

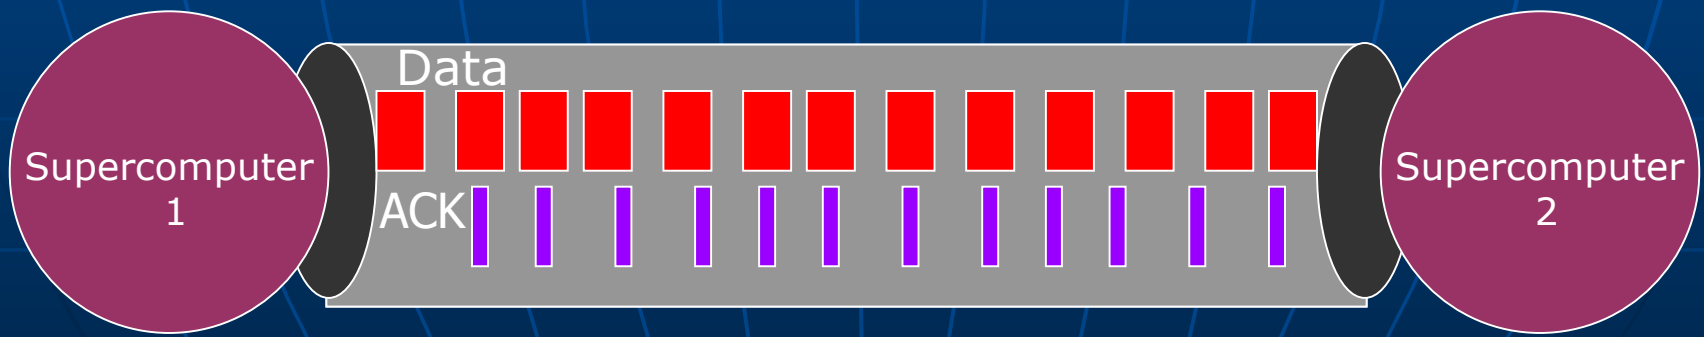

Current SCTP implementations

Linux
(http://sourceforge.net/projects/lksctp/)

FreeBSD/NetBSD/OpenBSD
(http://www.sctp.org)

Solaris
(http://playground.sun.com/sctp/)

Summary

SCTP is now an Internet standard track protocol (RFC 2960)

Advantages over TCP:

- Multistreaming;
- Multihoming;
- Message oriented;
- Unordered data delivery;
- Large window;
- Secure mechanisms.