


2.8 TCP CONGESTION AVOIDANCE

- Congestion avoidance mechanisms *prevent* congestion before it actually occurs.
- These mechanisms predict when congestion is about to happen and then to reduce the rate at which hosts send data just before packets start being discarded.
- TCP *creates* loss of packets in order to determine bandwidth of the connection.
- Routers *help* the end nodes by intimating when congestion is likely to occur.
- Congestion-avoidance mechanisms are:
 1. DEC bit - Destination Experiencing Congestion Bit
 2. RED - Random Early Detection

Dec Bit - Destination Experiencing Congestion Bit

- The first mechanism developed for use on the Digital Network Architecture (DNA).
- The idea is to evenly split the responsibility for congestion control between the routers and the end nodes.
- Each router monitors the load it is experiencing and explicitly notifies the end nodes when congestion is about to occur.
- This notification is implemented by setting a binary congestion bit in the packets that flow through the router; hence the name DECbit.
- The destination host then copies this congestion bit into the ACK it sends back to the source.
- The Source checks *how many* ACK has DEC bit set for previous window packets.
- If less than 50% of ACK have DEC bit set, then source *increases* its congestion window by 1 packet
- Otherwise, *decreases* the congestion window by 87.5%. ➤ Finally, the source adjusts its sending rate so as to avoid congestion.



- Average queue length is measured over a time interval that includes the ***last busy + last idle cycle + current busy cycle***.
- It calculates the average queue length by *dividing* the curve area with time interval.

Red - Random Early Detection

- The second mechanism of congestion avoidance is called as *Random Early Detection (RED)*.
- Each router is programmed to monitor its own queue length, and when it detects that there is congestion, it notifies the source to adjust its congestion window.
 - RED differs from the DEC bit scheme by two ways:
 - a. In DECbit, explicit notification about congestion is sent to source, whereas RED implicitly notifies the source by dropping a few packets.
 - b. DECbit may lead to tail drop policy, whereas RED drops packet based on drop probability in a random manner. Drop each arriving packet with some ***drop probability*** whenever the queue length exceeds some ***drop level***. This idea is called ***early random drop***.

2.9 STREAM CONTROL TRANSMISSION PROTOCOL (SCTP)

- Stream Control Transmission Protocol (SCTP) is a reliable, message-oriented transport layer protocol.
- SCTP has mixed features of TCP and UDP.
- SCTP maintains the message boundaries and detects the lost data, duplicate data as well as out-of-order data.
- SCTP provides the Congestion control as well as Flow control.
- SCTP is especially designed for internet applications as well as multimedia communication.


SCTP SERVICES

Process-to-Process Communication

- SCTP provides process-to-process communication.

Multiple Streams

- SCTP allows multistream service in each connection, which is called *association* in SCTP terminology.
- If one of the streams is blocked, the other streams can still deliver their data.

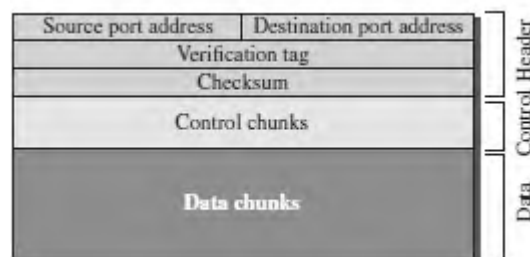


Multihoming

- An SCTP association supports multihoming service.
- The sending and receiving host can define multiple IP addresses in each end for an association.
- In this fault-tolerant approach, when one path fails, another interface can be used for data delivery without interruption.



SCTP PACKET FORMAT



An SCTP packet has a mandatory general header and a set of blocks called chunks.

General Header

- The *general header* (packet header) defines the end points of each association to which the packet belongs
- It guarantees that the packet belongs to a particular association
- It also preserves the integrity of the contents of the packet including the header itself.
- There are four fields in the general header.

Source port This field identifies the sending port.

Destination port This field identifies the receiving port that hosts use to route the packet to the appropriate endpoint/application.

Verification tag A 32-bit random value created during initialization to distinguish stale packets from a previous connection.

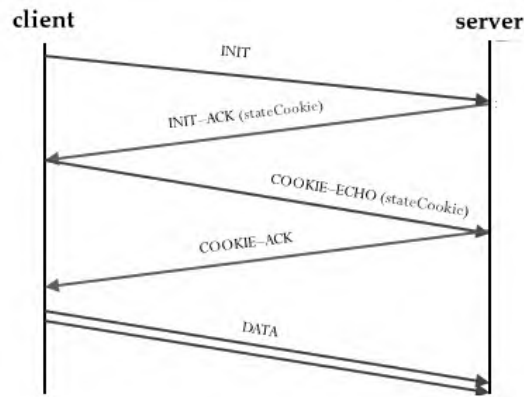
Checksum The next field is a checksum. The size of the checksum is 32 bits. SCTP uses CRC-32 Checksum.

SCTP ASSOCIATION

- SCTP is a connection-oriented protocol.
- A connection in SCTP is called an *association* to emphasize multihoming.
- SCTP Associations consists of three phases:
 1. Association Establishment
 2. Data Transfer
 3. Association Termination

Association Establishment

- Association establishment in SCTP requires a four-way handshake.
- In this procedure, a client process wants to establish an association with a server process using SCTP as the transport-layer protocol.
- The SCTP server needs to be prepared to receive any association (passive open).
- Association establishment, however, is initiated by the client (active open).



- The client sends the first packet, which contains an INIT chunk.
- The server sends the second packet, which contains an INIT ACK chunk. The INIT ACK also sends a cookie that defines the state of the server at this moment.
- The client sends the third packet, which includes a COOKIE ECHO chunk. This is a very simple chunk that echoes, without change, the cookie sent by the server. SCTP allows the inclusion of data chunks in this packet.

Data Transfer

- The whole purpose of an association is to transfer data between two ends.
- After the association is established, bidirectional data transfer can take place.
- The client and the server can both send data.
- SCTP supports piggybacking.

Association Termination

- In SCTP, either of the two parties involved in exchanging data (client or server) can close the connection.
- SCTP does not allow a “half closed” association. If one end closes the association, the other end must stop sending new data.
- If any data are left over in the queue of the recipient of the termination request, they are sent and the association is closed.
- Association termination uses three packets.

