DRONACHARYA
College of Engineering

Computer Science & Engineering

Data Communication and Computer Networks

(MTCSE-101-A)

Stream Control Transmission Protocol (SCTP)

OBJECTIVES:

- To introduce SCTP as a new transport-layer protocol.
- To discuss SCTP services and compare them with TCP.
- To list and explain different packet types used in SCTP and discuss the purpose and of each field in each packet.
- To discuss SCTP association and explain different scenarios such as association establishment, data transfer, association termination, and association abortion.
- To compare and contrast the state transition diagram of SCTP with the corresponding diagram of TCP.
- To explain flow control, error control, and congestion control mechanism in SCTP and compare them with the similar mechanisms in TCP.

Chapter Outline

16-1 INTRODUCTION

Stream Control Transmission Protocol (SCTP) is a new reliable, message-oriented transport-layer protocol. Figure 16.1 shows the relationship of SCTP to the other protocols in the Internet protocol suite. SCTP lies between the application layer and the network layer and serves as the intermediary between the application programs and the network operations.

Figure 16.1 *TCP/IP Protocol suite*

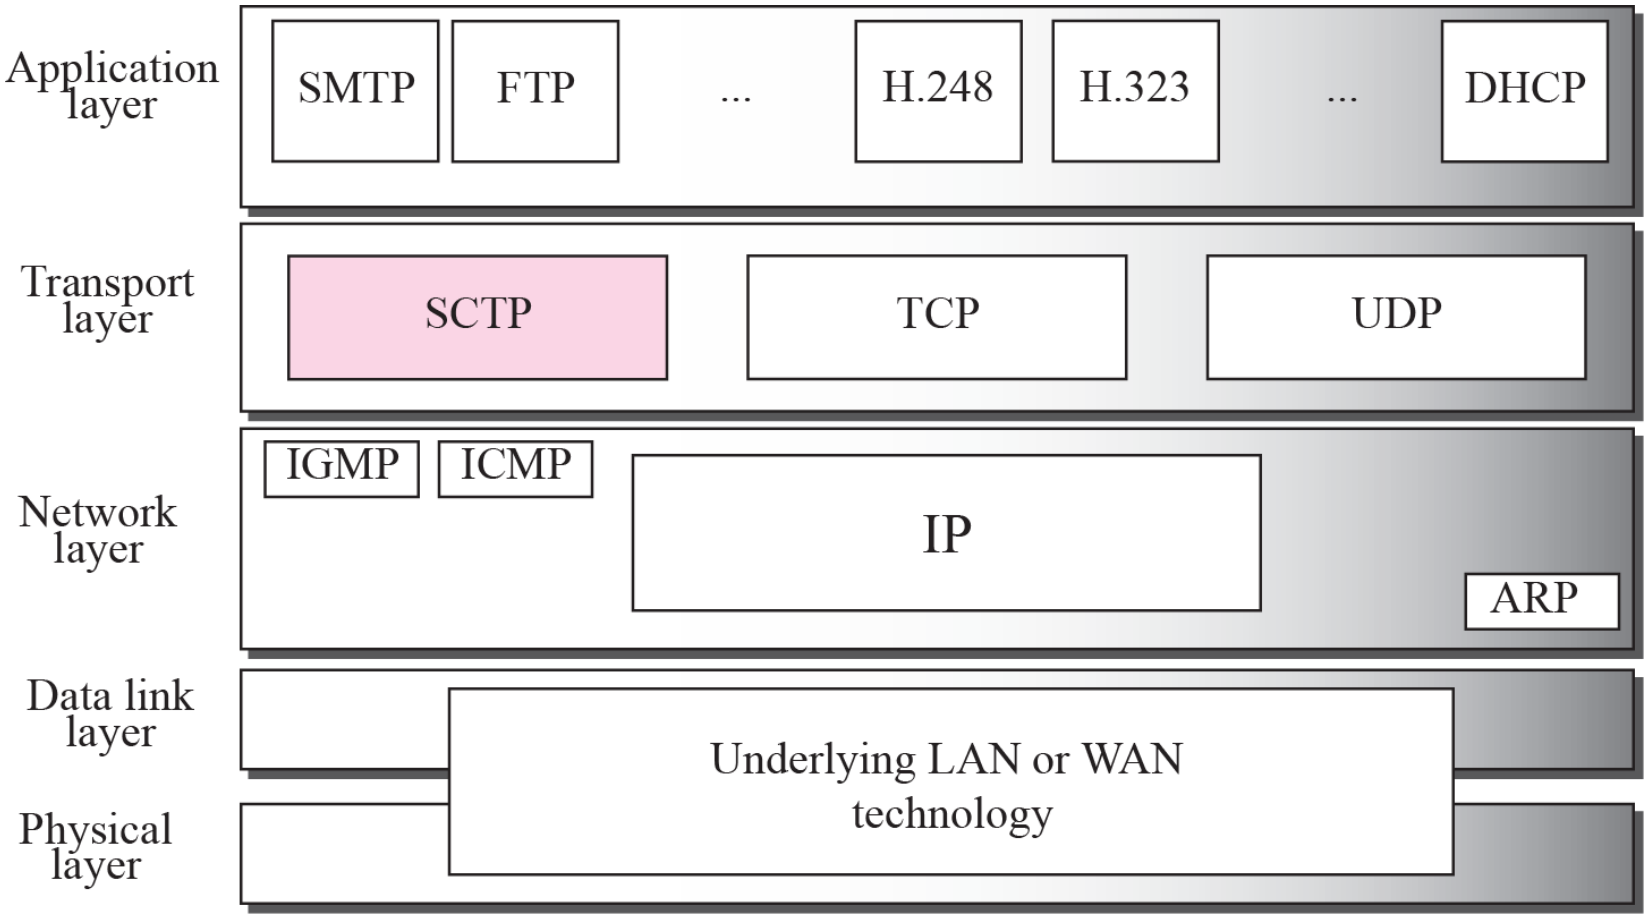

Note

SCTP is a message-oriented, reliable protocol that combines the best features of UDP and TCP.

Comparison

- UDP: Message-oriented, Unreliable
- TCP: Byte-oriented, Reliable
- SCTP
  - Message-oriented, Reliable
  - Other innovative features
    - Association, Data transfer/Delivery
    - Fragmentation, Error/Congestion Control

16-2 SCTP SERVICES

Before discussing the operation of SCTP, let us explain the services offered by SCTP to the application layer processes.

Topics Discussed in the Section

- ✓ **Process-to-Process Communication**
- ✓ **Multiple Streams**
- ✓ **Multihoming**
- ✓ **Full-Duplex Communication**
- ✓ **Connection-Oriented Service**
- ✓ **Reliable Service**

Table 16.1 *Some SCTP applications*

Protocol	*Port Number*	*Description*
IUA	9990	ISDN over IP
M2UA	2904	SS7 telephony signaling
M3UA	2905	SS7 telephony signaling
H.248	2945	Media gateway control
H.323	1718, 1719, 1720, 11720	IP telephony
SIP	5060	IP telephony

Figure 16.2 *Multiple-stream concept*

If one of the streams is blocked, the other streams can still deliver their data.

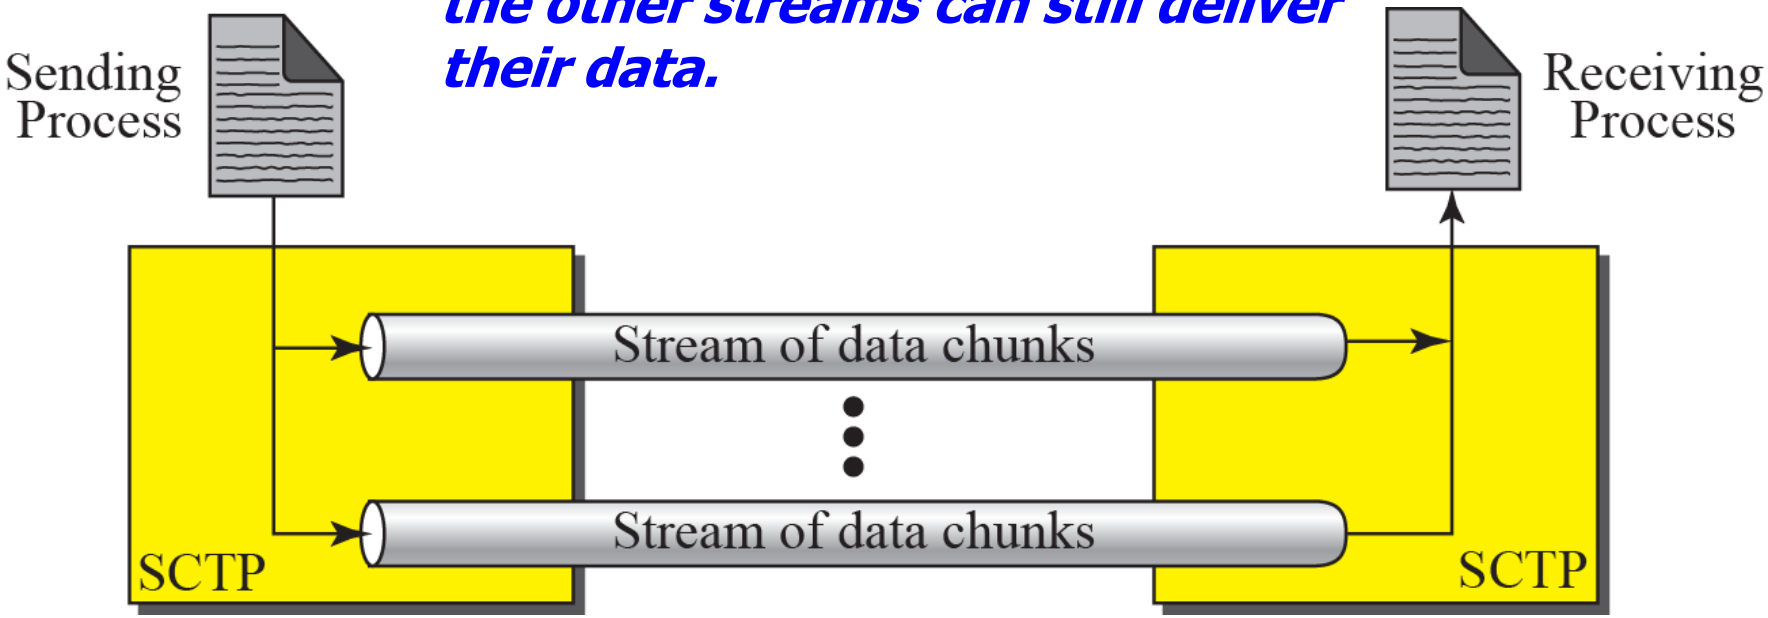

Note

An association in SCTP can involve multiple streams.

Figure 16.3 *Multihoming concept*

At present, SCTP does not allow load sharing between different path.

Currently, it is only for fault-tolerance.

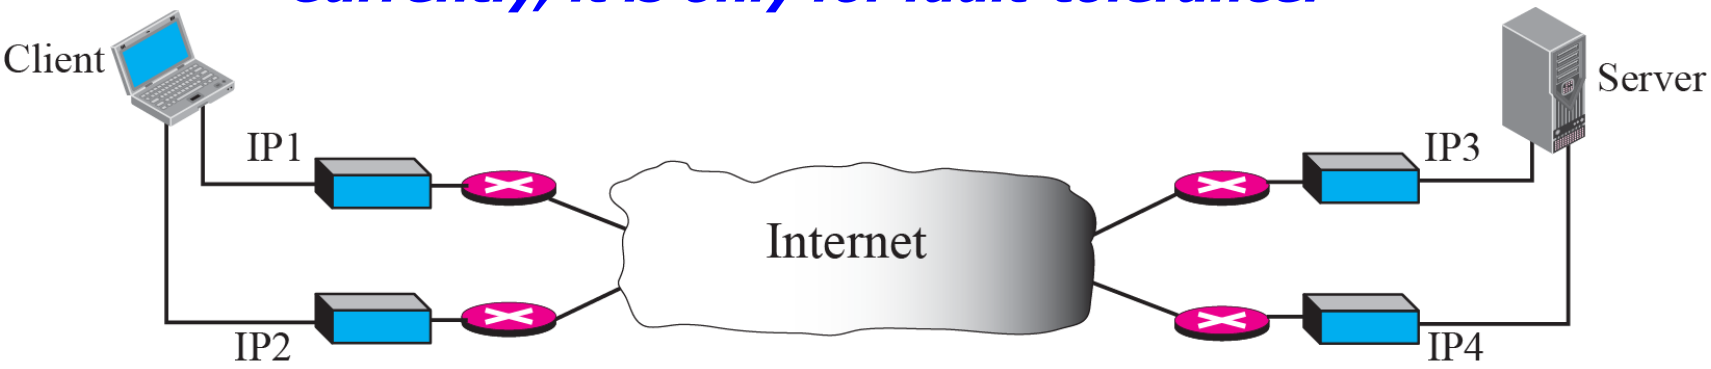

Note

SCTP association allows multiple IP addresses for each end.

16-3 SCTP FEATURES

Let us first discuss the general features of SCTP and then compare them with those of TCP.

Topics Discussed in the Section

- ✓ Transmission Sequence Number (TSN)
- ✓ Stream Identifier (SI)
- ✓ Stream Sequence Number (SSN)
- ✓ Packets
- ✓ Acknowledgment Number
- ✓ Flow Control
- ✓ Error Control
- ✓ Congestion Control

Note

In SCTP, a data chunk is numbered using a TSN.

Note

To distinguish between different streams, SCTP uses an SI.

Note

To distinguish between different data chunks belonging to the same stream, SCTP uses SSNs.

Figure 16.4 *Comparison between a TCP segment and an SCTP packet*

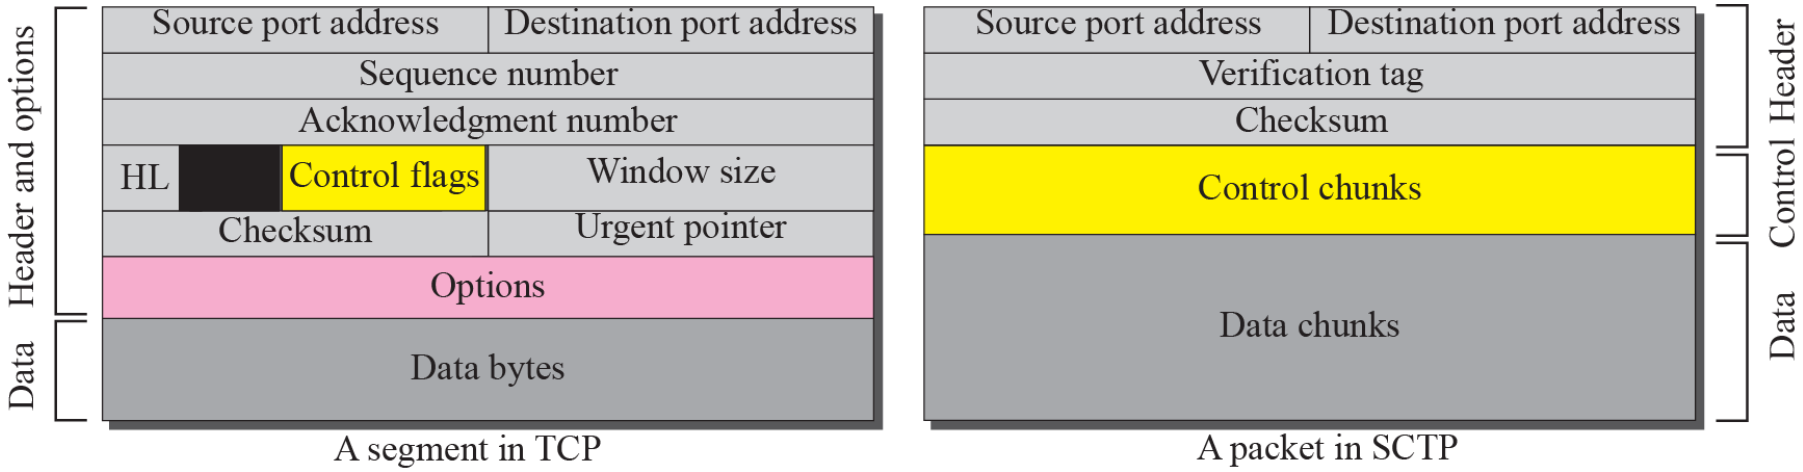

Note

TCP has segments; SCTP has packets.

SCTP vs. TCP (1)

- Control information
  - TCP: part of the header
  - SCTP: several types of control chunks
- Data
  - TCP: one entity in a TCP segment
  - SCTP: several data chunks in a packet
- Option
  - TCP: part of the header
  - SCTP: handled by defining new chunk types

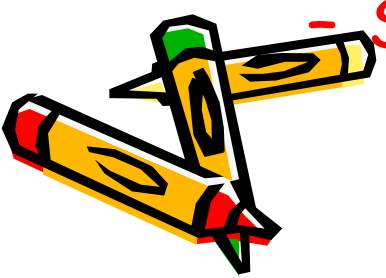

SCTP vs. TCP (2)

- Mandatory part of the header
  - TCP: 20 bytes, SCTP: 12 bytes
  - Reason:
    - TSN in data chunk's header
    - Ack. # and window size are part of control chunk
    - No need for header length field (∵no option)
    - No need for an urgent pointer
- Checksum
  - TCP: 16 bits, SCTP: 32 bit

SCTP vs. TCP (3)

- Association identifier
  - TCP: none, SCTP: verification tag
  - Multihoming in SCTP
- Sequence number
  - TCP: one # in the header
  - SCTP: TSN, SI and SSN define each data chunk
  - SYN and FIN need to consume one seq. #
  - Control chunks never use a TSN, SI, or SSN number

Note

In SCTP, control information and data information are carried in separate chunks.

Figure 16.5 *Packet, data chunks, and streams*

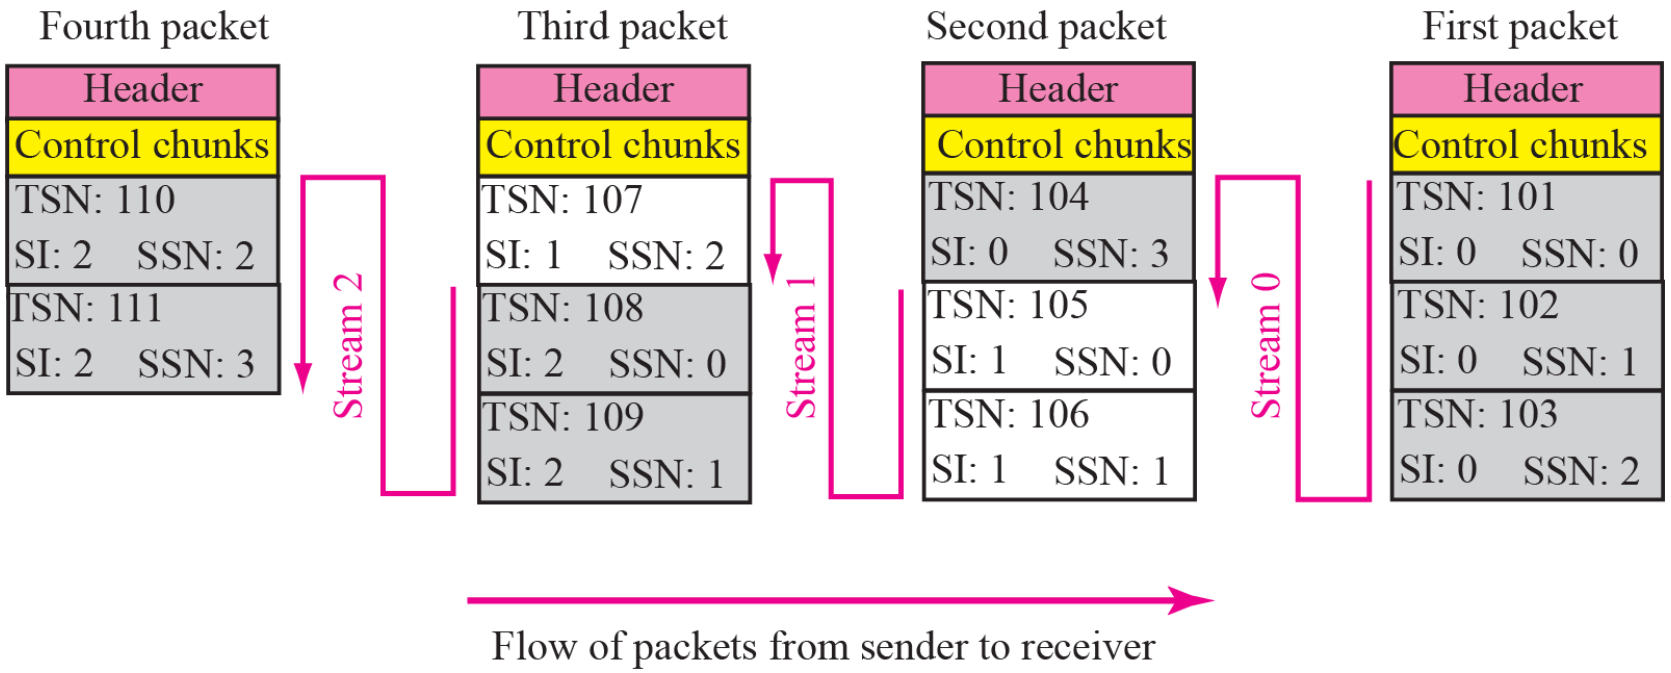

Note

Data chunks are identified by three identifiers: TSN, SI, and SSN. TSN is a cumulative number identifying the association; SI defines the stream; SSN defines the chunk in a stream.

Note

In SCTP, acknowledgment numbers are used to acknowledge only data chunks; control chunks are acknowledged by other control chunks if necessary.

16-4 PACKET FORMAT

In this section, we show the format of a packet and different types of chunks. Most of the information presented in this section will become clear later; this section can be skipped in the first reading or used only as the reference. An SCTP packet has a mandatory general header and a set of blocks called chunks. There are two types of chunks: control chunks and data chunks.

Topics Discussed in the Section

- ✓ General Header
- ✓ Chunks

Figure 16.6 *SCTP packet format*

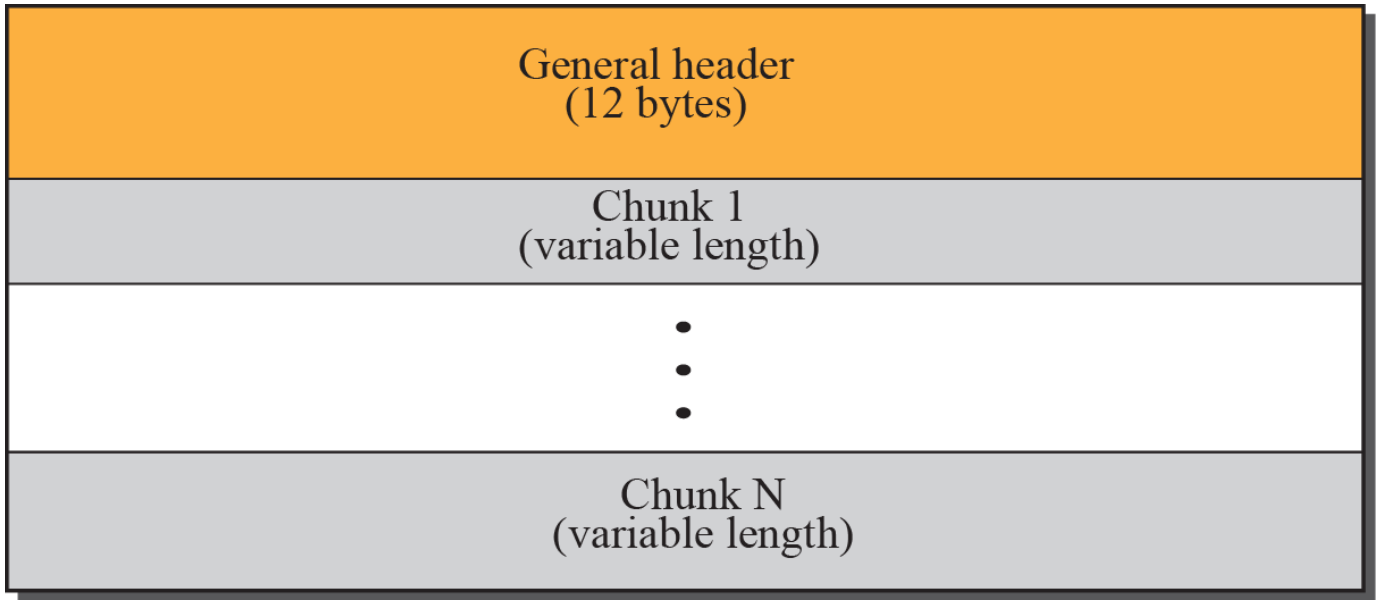

Note

In an SCTP packet, control chunks come before data chunks.

Figure 16.7 *Common layout of a chunk*

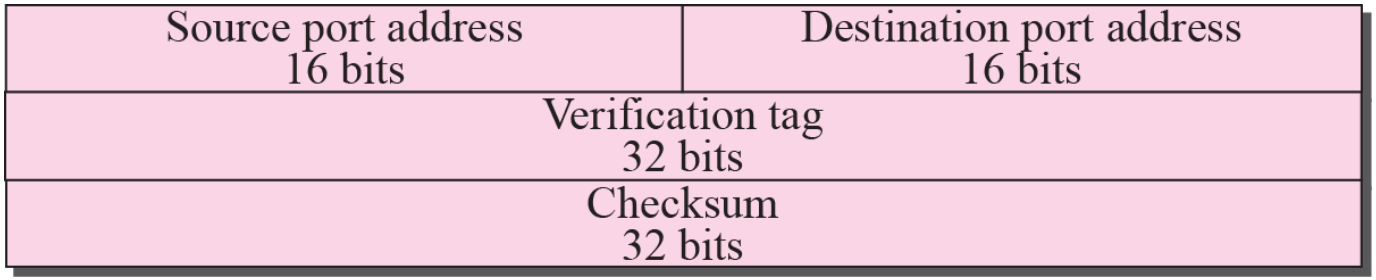

Figure 16.8 *Multiple-stream concept*

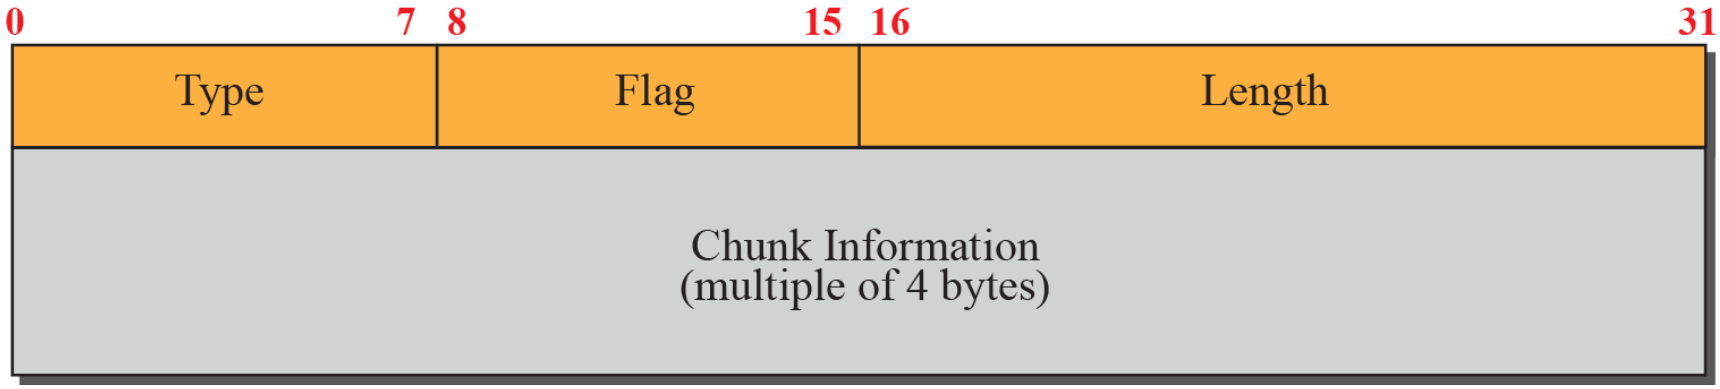

Note

Chunks need to terminate on a 32-bit (4-byte) boundary.

Table 16.2 *Chunks*

Type	*Chunk*	*Description*
0	DATA	User data
1	INIT	Sets up an association
2	INIT ACK	Acknowledges INIT chunk
3	SACK	Selective acknowledgment
4	HEARTBEAT	Probes the peer for liveliness
5	HEARTBEAT ACK	Acknowledges HEARTBEAT chunk
6	ABORT	Abort an association
7	SHUTDOWN	Terminates an association
8	SHUTDOWN ACK	Acknowledges SHUTDOWN chunk
9	ERROR	Reports errors without shutting down
10	COOKIE ECHO	Third packet in association establishment
11	COOKIE ACK	Acknowledges COOKIE ECHO chunk
14	SHUTDOWN COMPLETE	Third packet in association termination
192	FORWARD TSN	For adjusting cumulating TSN

Note

The number of padding bytes is not included in the value of the length field.

Figure 16.9 *Data chunk*

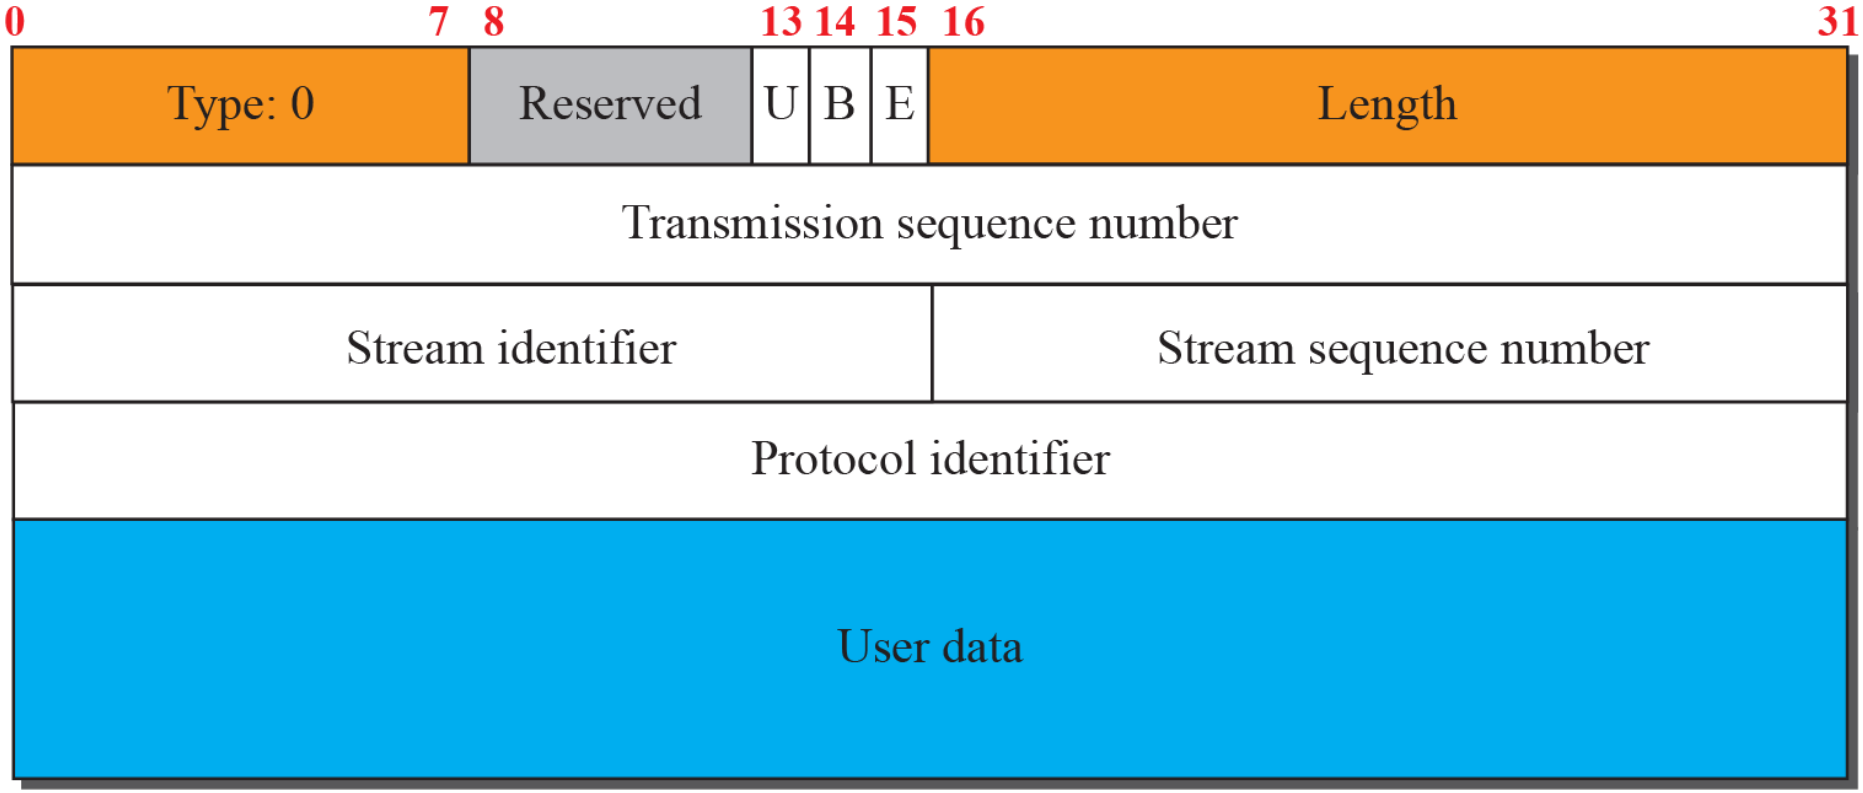

Note

A DATA chunk cannot carry data belonging to more than one message, but a message can be split into several chunks. The data field of the DATA chunk must carry at least one byte of data, which means the value of length field cannot be less than 17.

Figure 16.10 *INIT chunk*

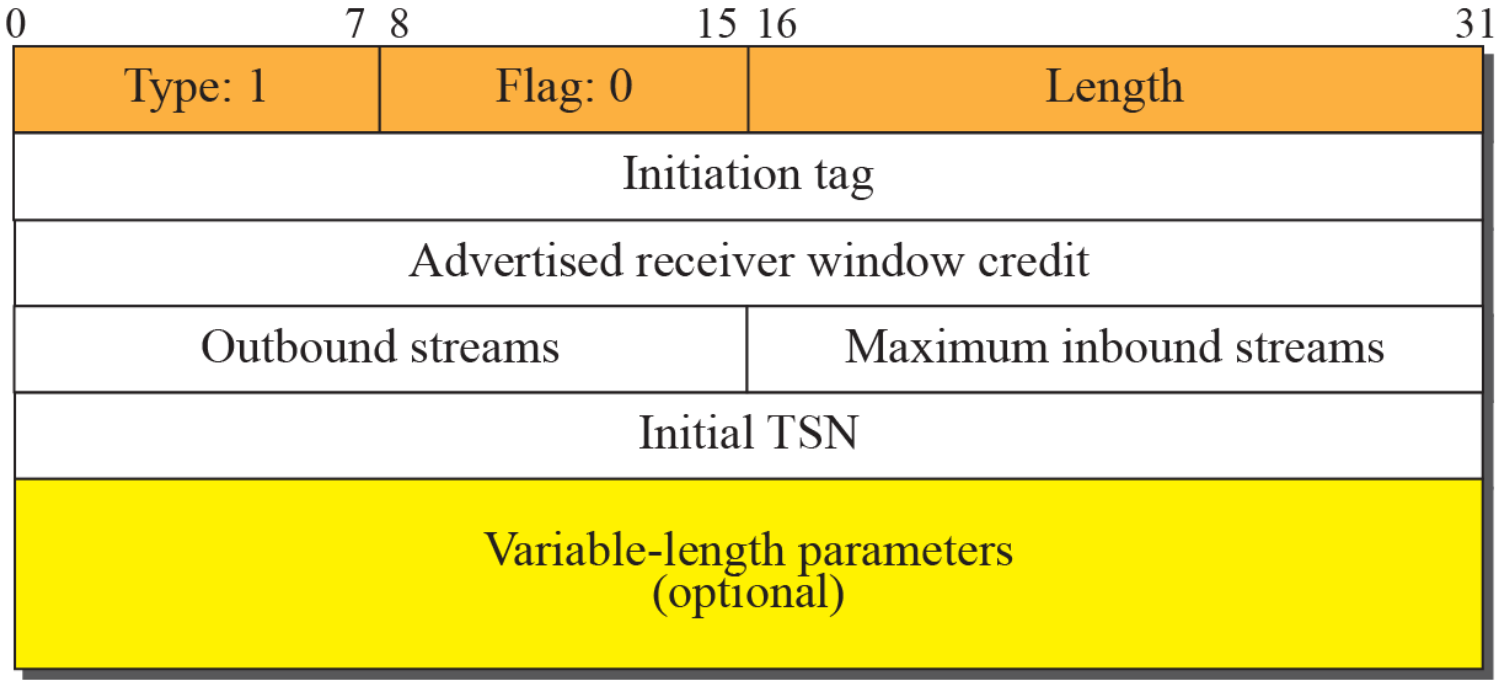

Note

No other chunk can be carried in a packet that carries an INIT chunk.

Figure 16.11 *INIT ACK chunk*

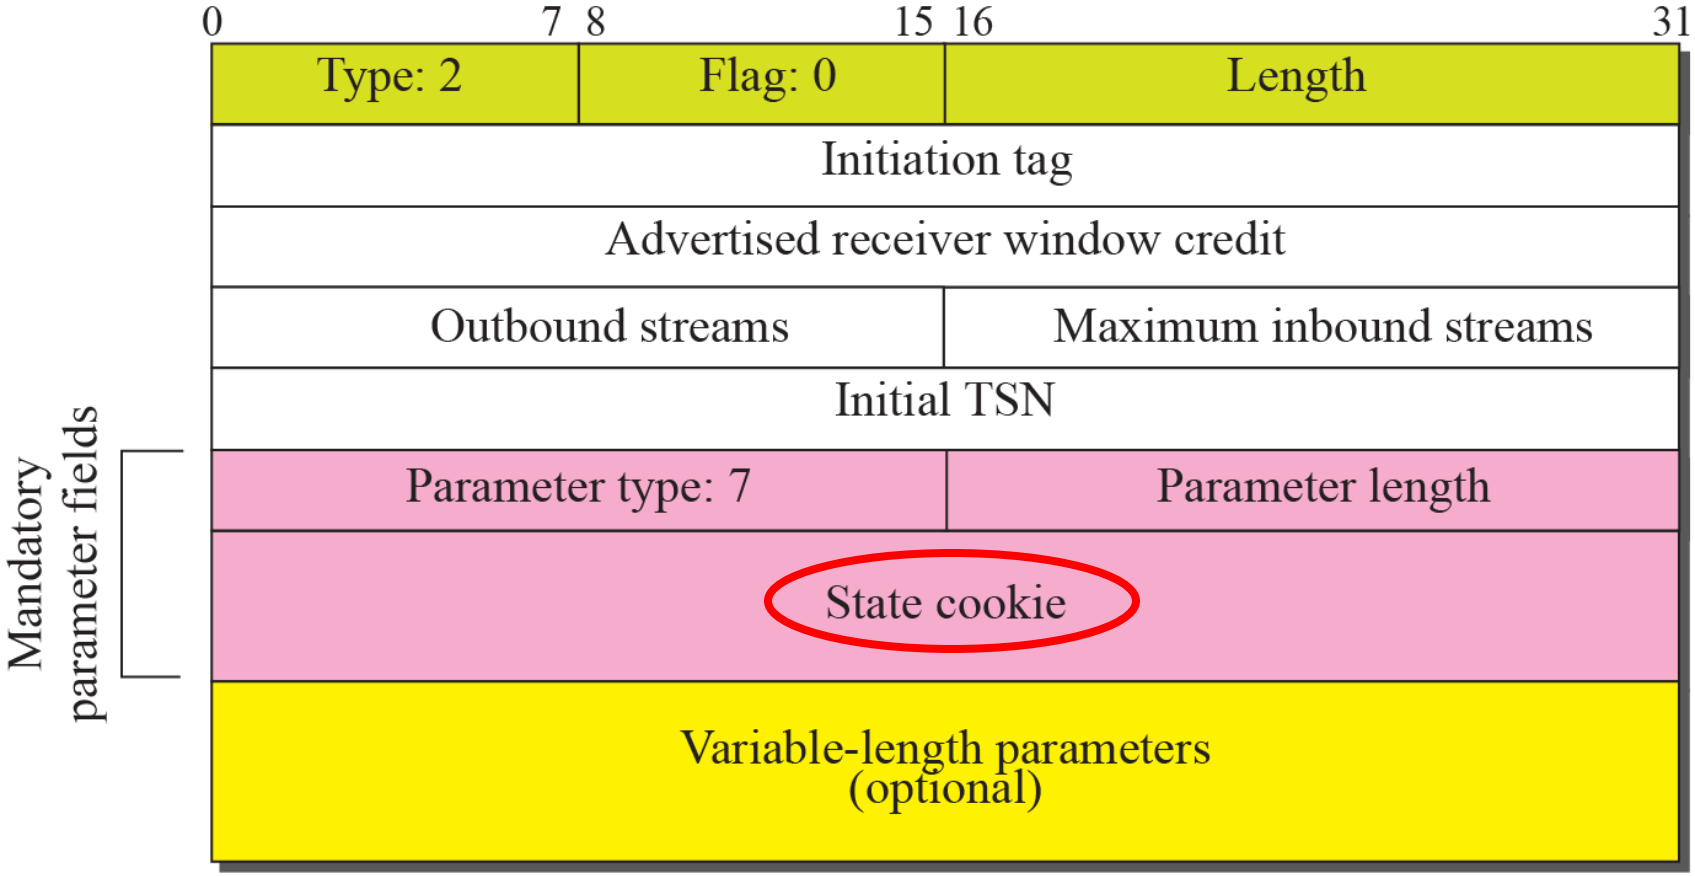

Note

No other chunk can be carried in a packet that carries an INIT ACK chunk.

Figure 16.12 *COOKIE ECHO chunk*

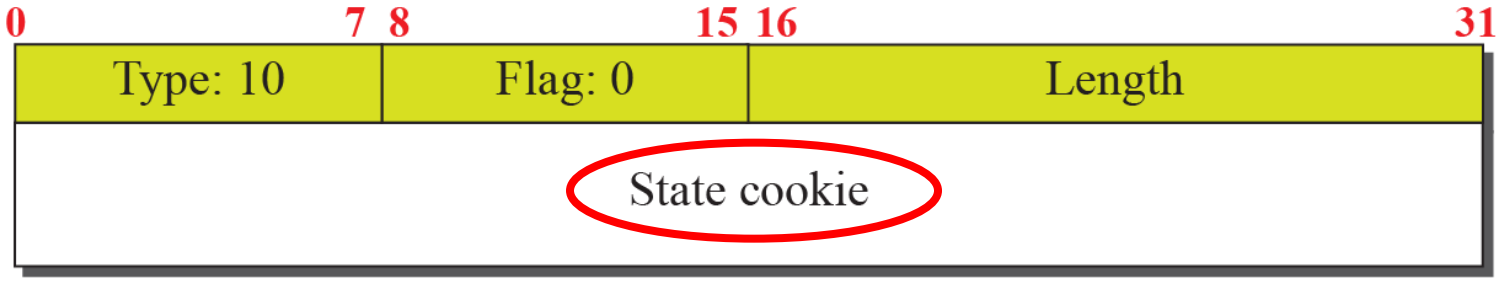

Figure 16.13 *COOKIE ACK*

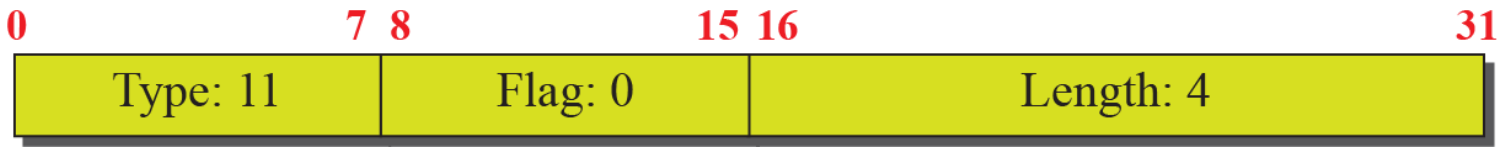

Figure 16.14 *SACK chunk*

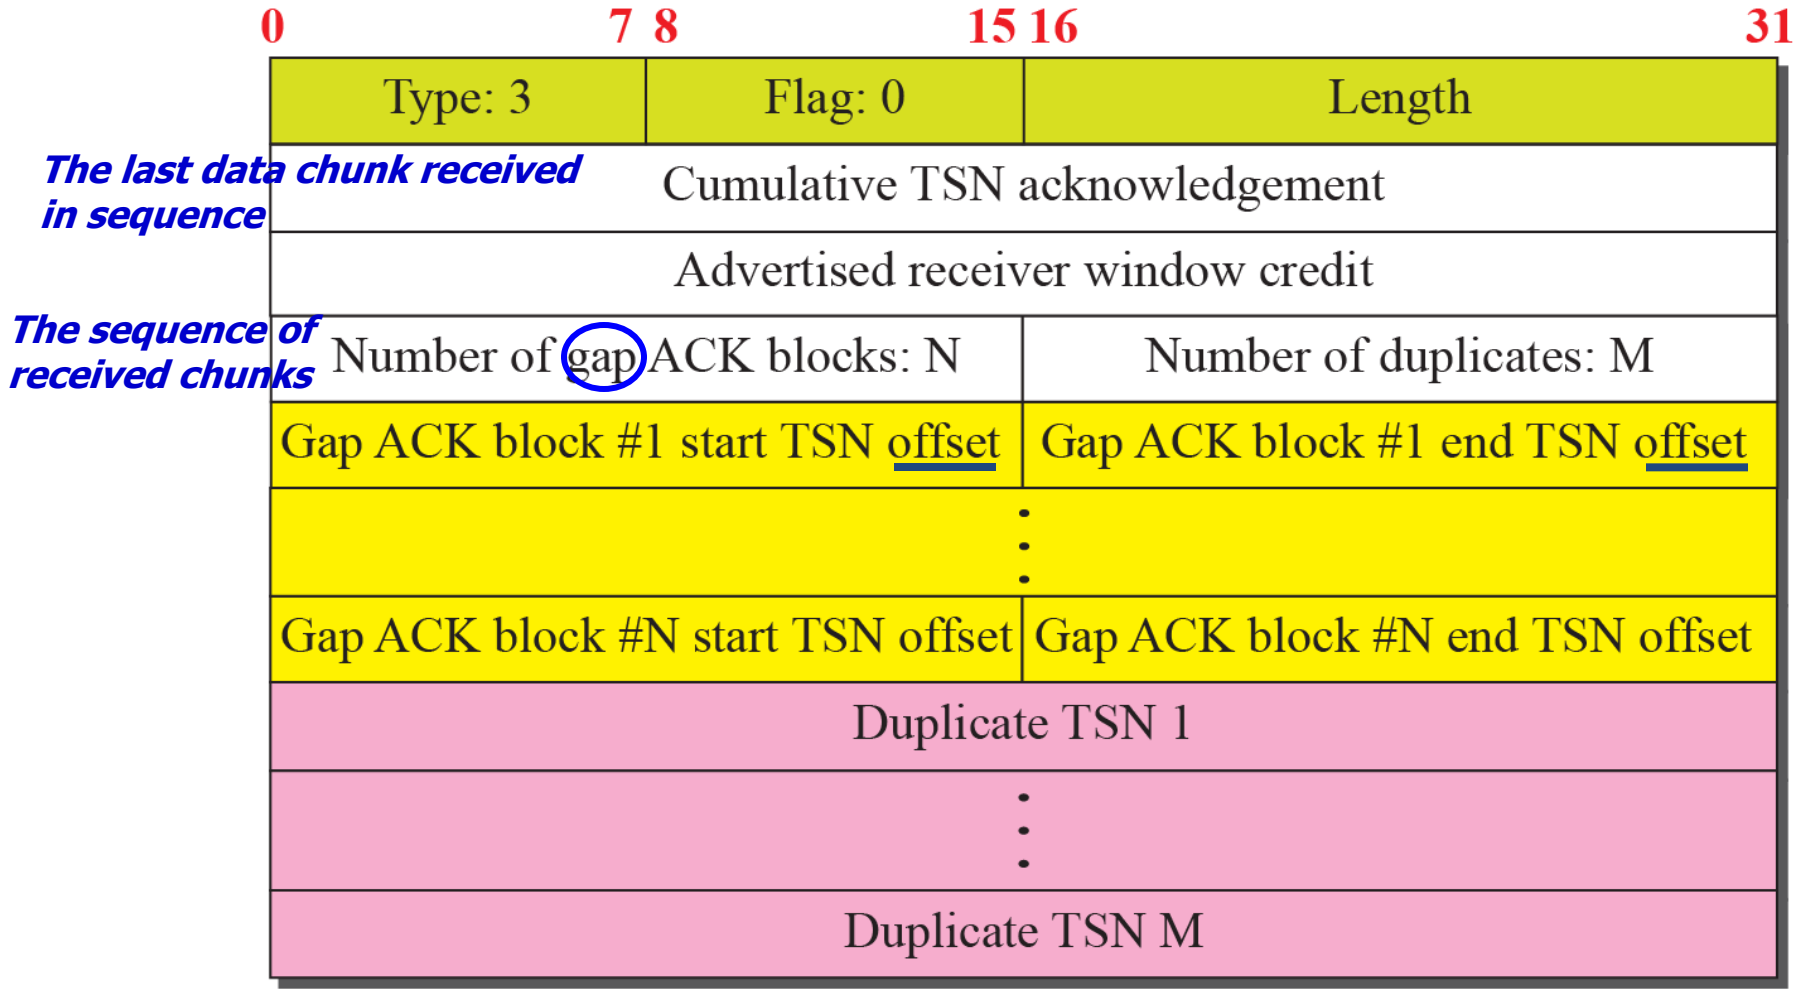

Figure 16.15 *HEARTBEAT and HEARTBEAT ACK chunk*

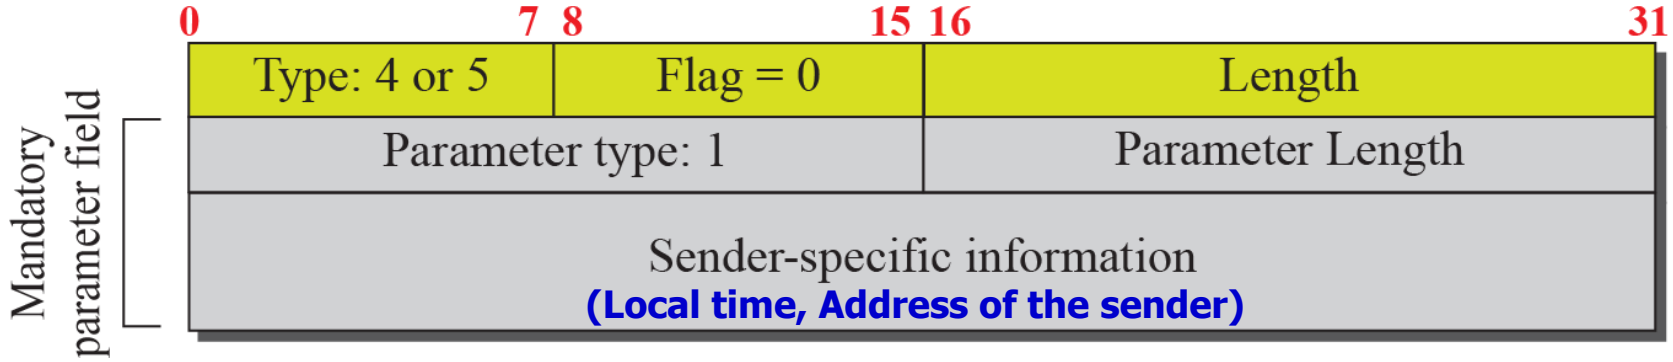

Used to periodically probe the condition of an association

Figure 16.16 *SHUTDOWN chunks*

Figure 16.17 *ERROR chunk*

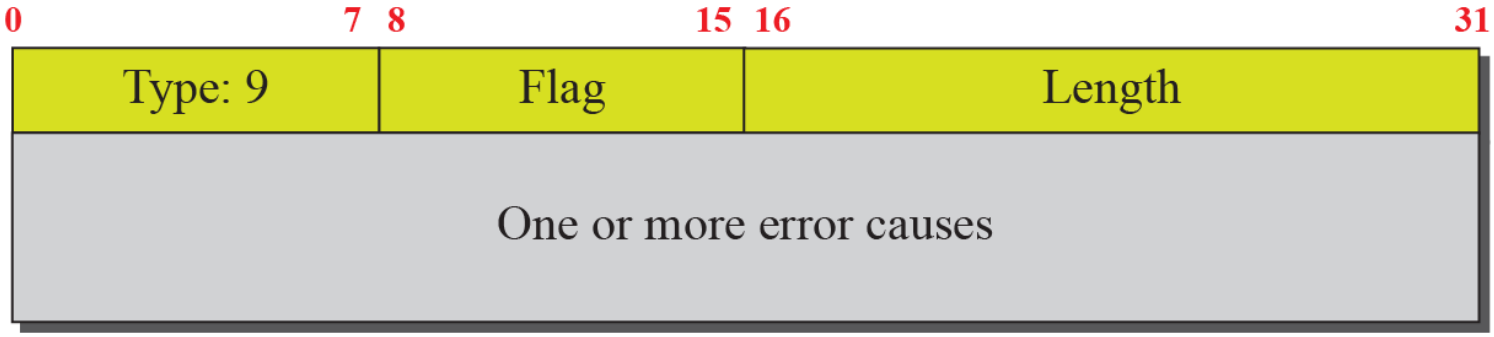

Sent when an end point finds some error in a received packet
But, which packet is with the error?

Table 16.3 *Errors*

Code	*Description*
1	Invalid stream identifier
2	Missing mandatory parameter
3	State cookie error
4	Out of resource
5	Unresolvable address
6	Unrecognized chunk type
7	Invalid mandatory parameters
8	Unrecognized parameter
9	No user data
10	Cookie received while shutting down

Figure 16.18 *ABORT chunk*

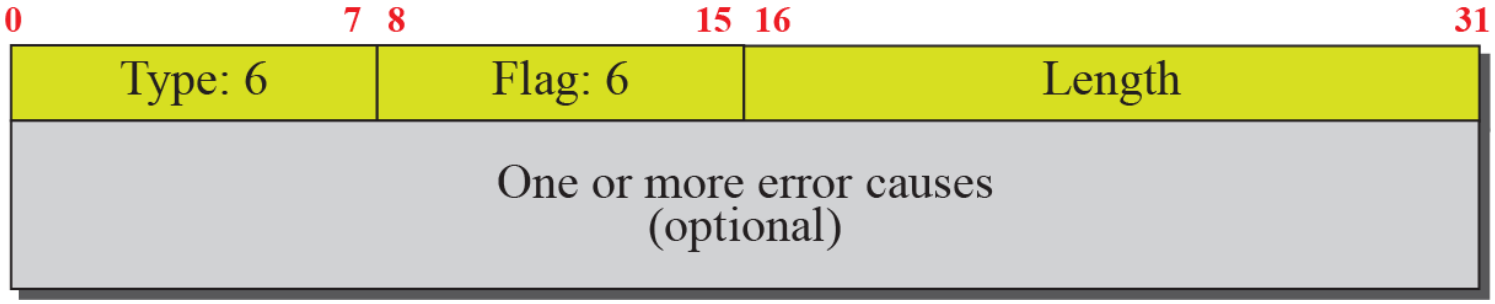

Forward TSN Chunk

- Recently added to the standard (RFC 3758)
- Used to inform the receiver to adjust its cumulative TSN
- It provides partial reliable service

16-5 AN SCTP ASSOCIATION

SCTP, like TCP, is a connection-oriented protocol. However, a connection in SCTP is called an association to emphasize multihoming.

Topics Discussed in the Section

- ✓ **Association Establishment**
- ✓ **Data Transfer**
- ✓ **Association Termination**
- ✓ **Association Abortion**

Note

A connection in SCTP is called an association.

Figure 16.19 *Four-way handshaking*

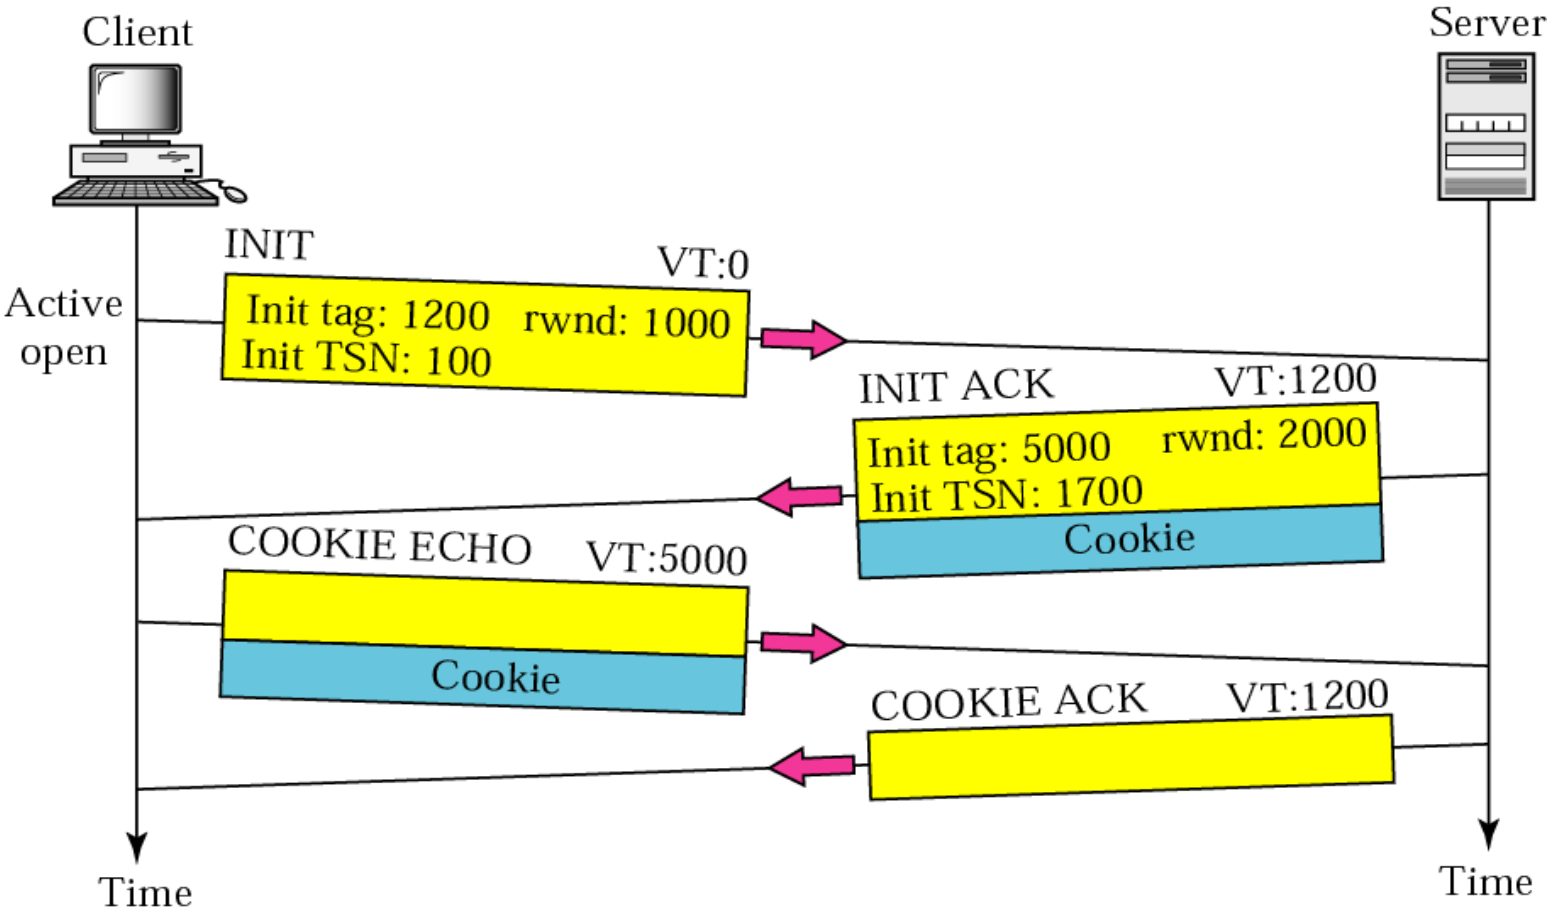

Verification Tag

- In TCP, a connection is identified by a combination of IP addresses and port numbers
  - A blind attacker can send segments to a TCP server using randomly chosen source and destination port numbers
  - Delayed segment from a previous connection can show up in a new connection that uses the same source and destination port addresses (incarnation)

TIME-WAIT timer

- Two verification tags, one for each direction, identify an association

Cookie (1)

- In TCP
  - Each time the server receives a SYN segment, it sets up a TCB and allocates other resources
- In SCTP
  - Postpone the allocation of resources until the reception of the third packet, when the IP address of the sender is verified

Cookie (2)

- In SCTP
  - The information received in the first packet must somehow be saved until the third packet arrives
  - Solution: to pack the information and send it back to the client (cookie)
  - The above strategy works if no entity can "eat" a cookie "baked" by the server
  - To guarantee this, the server creates a digest from the information using its own secret key

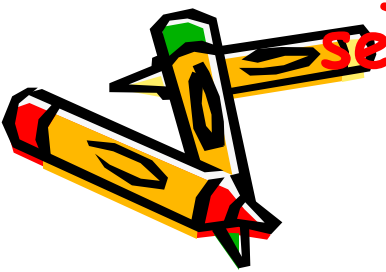

Note

No other chunk is allowed in a packet carrying an INIT or INIT ACK chunk. A COOKIE ECHO or a COOKIE ACK chunk can carry data chunks.

Note

In SCTP, only data chunks consume TSNs; data chunks are the only chunks that are acknowledged.

Figure 16.20 *Simple data transfer*

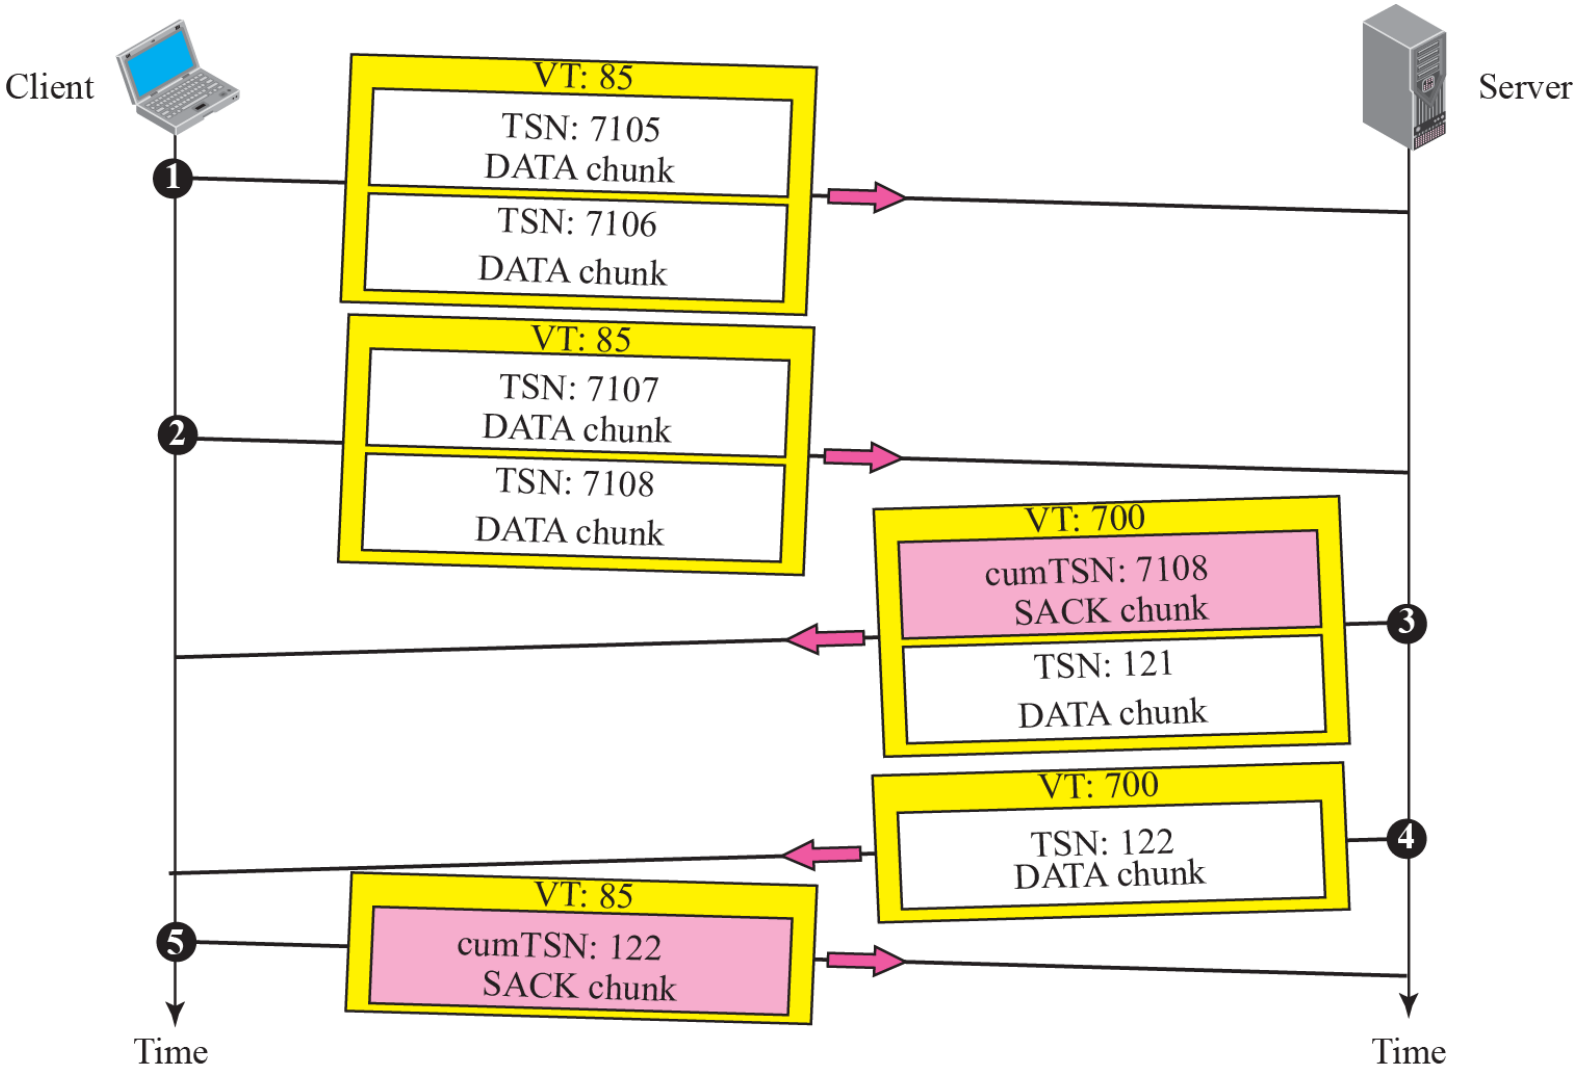

Note

The acknowledgment in SCTP defines the cumulative TSN, the TSN of the last data chunk received in order.

Multi-homing Data Transfer

- Primary address
  - The rest are alternative addresses
  - Defined during association establishment
  - Determined by the other end
  - The process can always override the primary address (explicitly)
  - SACK is sent to the address from which the corresponding SCTP packet originated

Multi-stream Delivery

- Interesting feature in SCTP
  - Distinction between data transfer and data delivery
  - Data transfer: TSN (error/flow control)
  - Data delivery: SI, SSN
- **Data delivery** (in each stream)
  - Ordered (default)
  - Unordered: using the U flag, do not consume SSNs (U flag with fragmentation?)

Fragmentation

- IP fragmentation vs. SCTP
  - SCTP preserves the boundaries of the msg from process to process when creating a DATA chunk from a message if the size of the msg does not exceed the MTU of the path
- SCTP fragmentation
  - Each fragment carries a different TSN
  - All header chunks carries the same SI, SSN, payload protocol ID, and U flag
  - Combination of B and E flag: 11,10,00,01

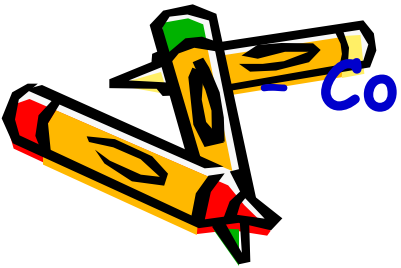

Figure 16.21 *Association termination*

SCTP does not allow a "half-closed" association

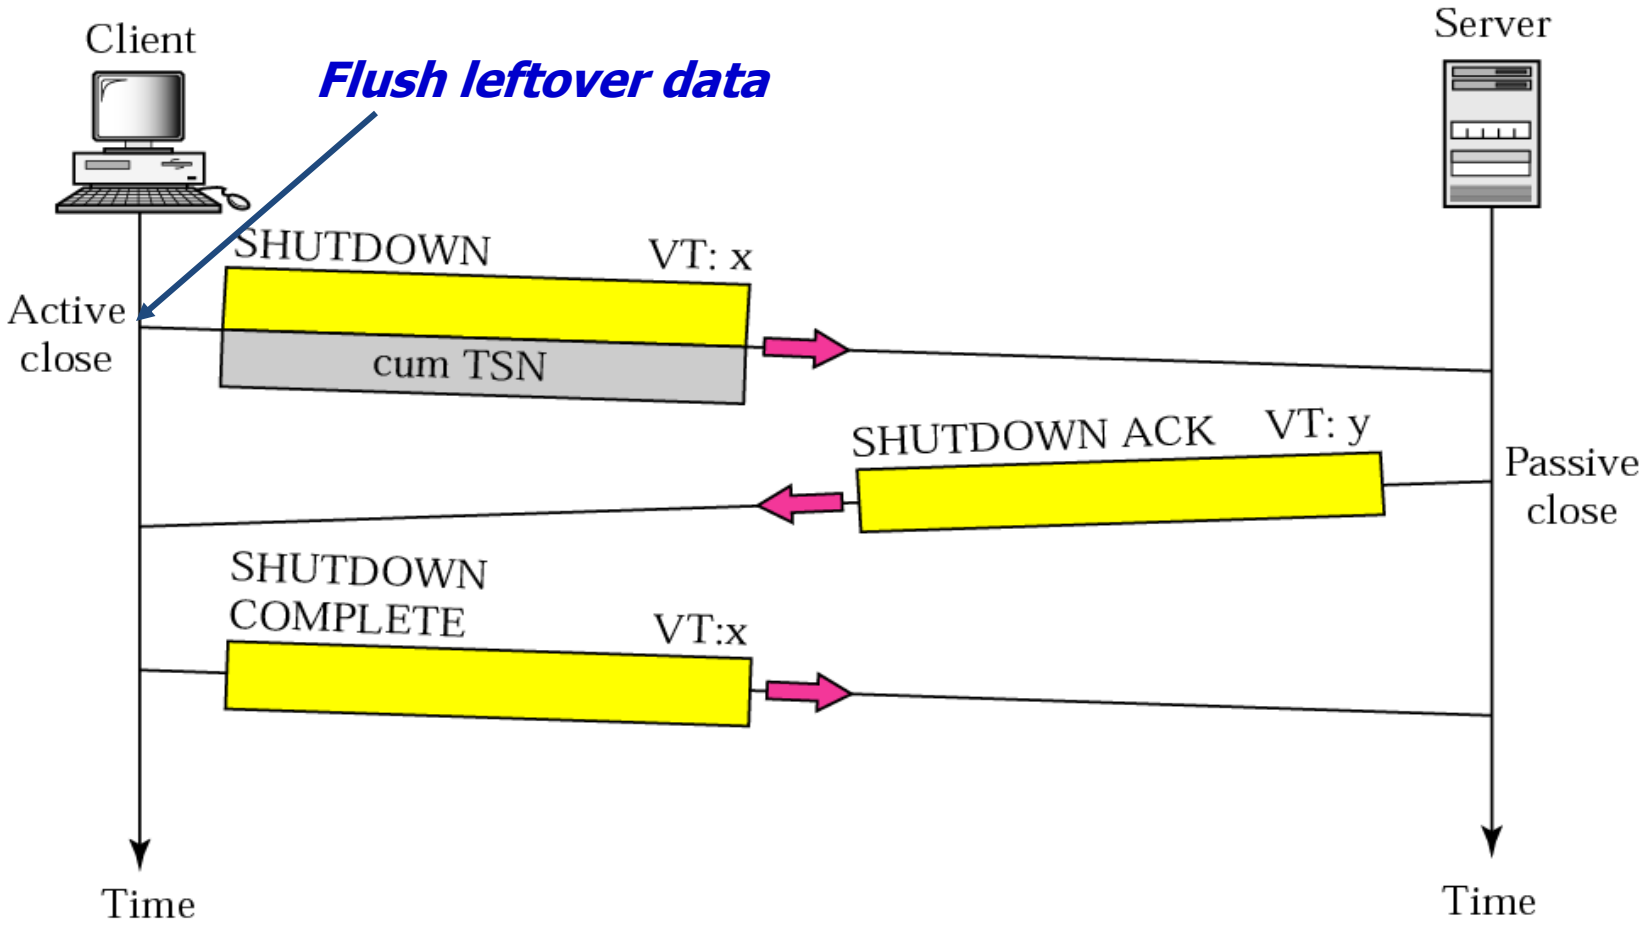

Figure 16.22 *Association abortion*

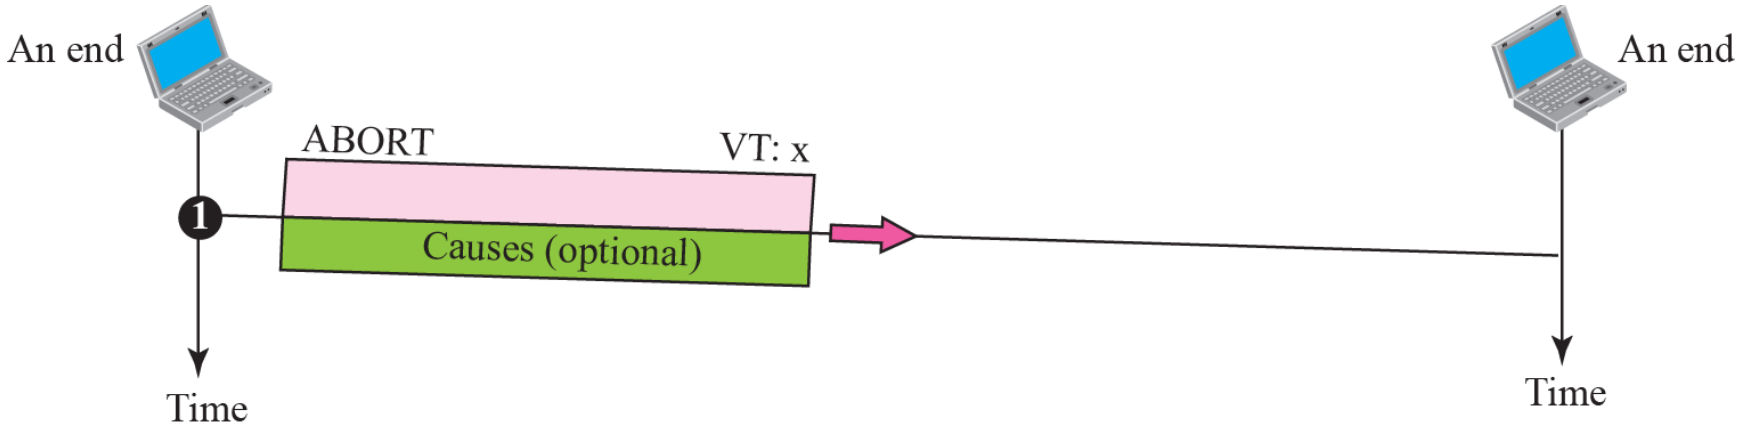

16-6 STATE TRANSITION DIAGRAM

To keep track of all the different events happening during association establishment, association termination, and data transfer, the SCTP software, like TCP, is implemented as a finite state machine. Figure 16.23 shows the state transition diagram for both client and server.

Topics Discussed in the Section

- ✓ **Scenarios**

Figure 16.23 *State transition diagram*

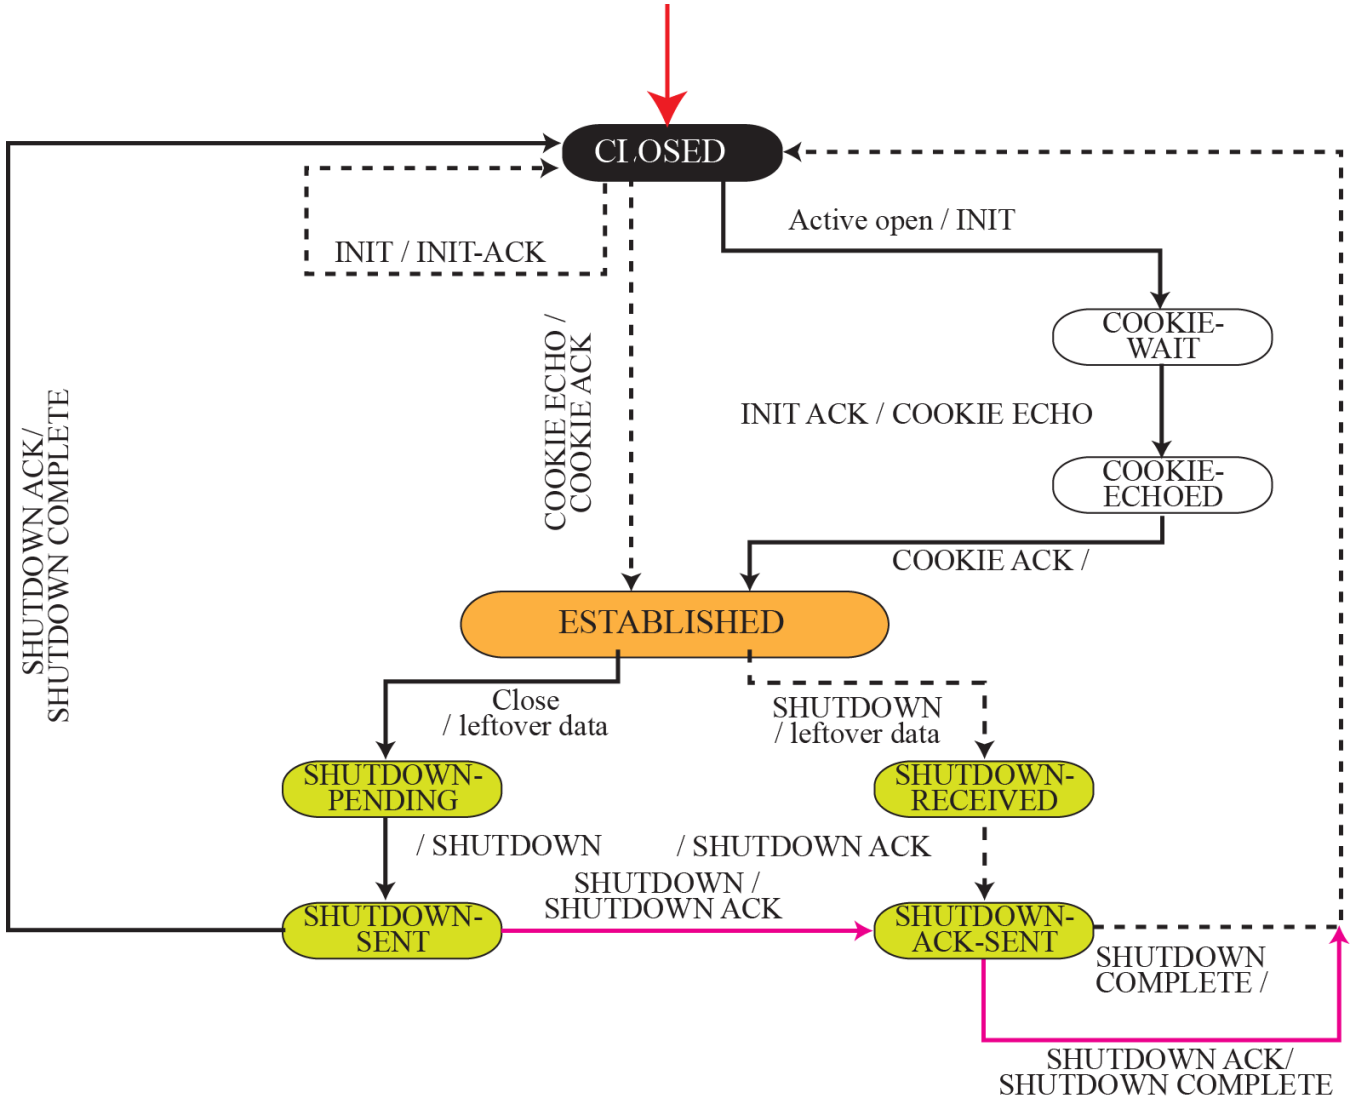

Table 16.4 *States for SCTP*

State	*Description*
CLOSED	No connection
COOKIE-WAIT	Waiting for a cookie
COOKIE-ECHOED	Waiting for cookie acknowledgment
ESTABLISHED	Connection is established; data are being transferred
SHUTDOWN-PENDING	Sending data after receiving *close*
SHUTDOWN-SENT	Waiting for SHUTDOWN acknowledgment
SHUTDOWN-RECEIVED	Sending data after receiving SHUTDOWN
SHUTDOWN-ACK-SENT	Waiting for termination completion

Figure 16.24 *A common scenario of state*

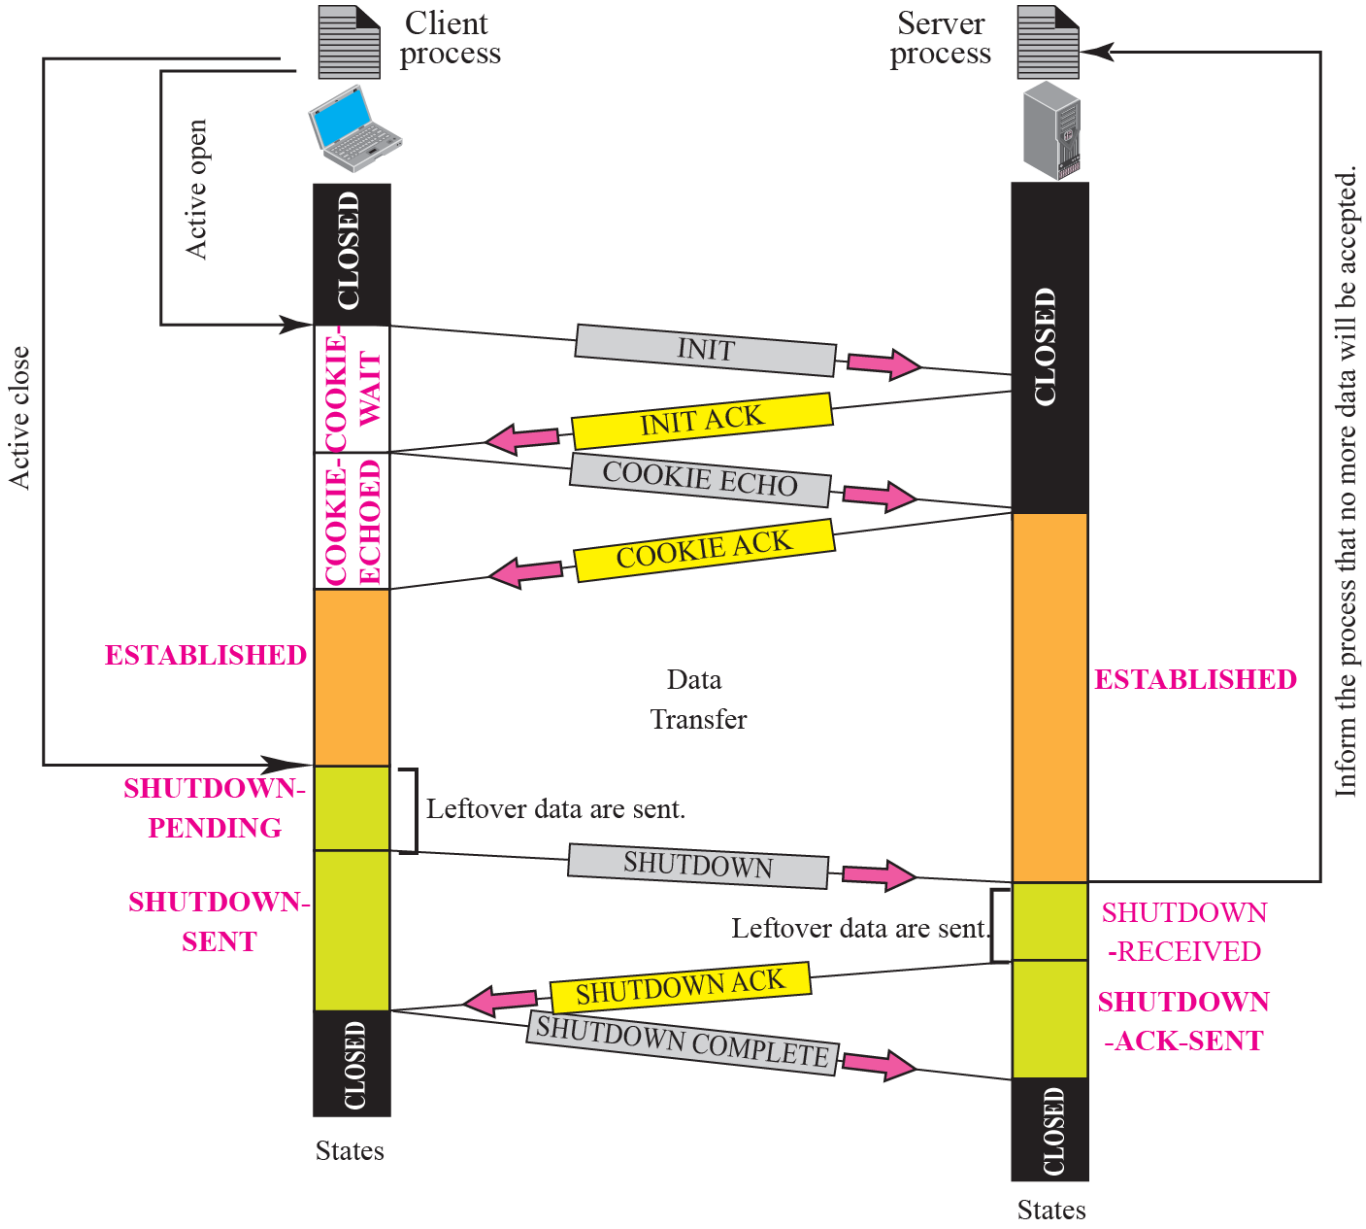

Figure 16.25 *Simultaneous open*

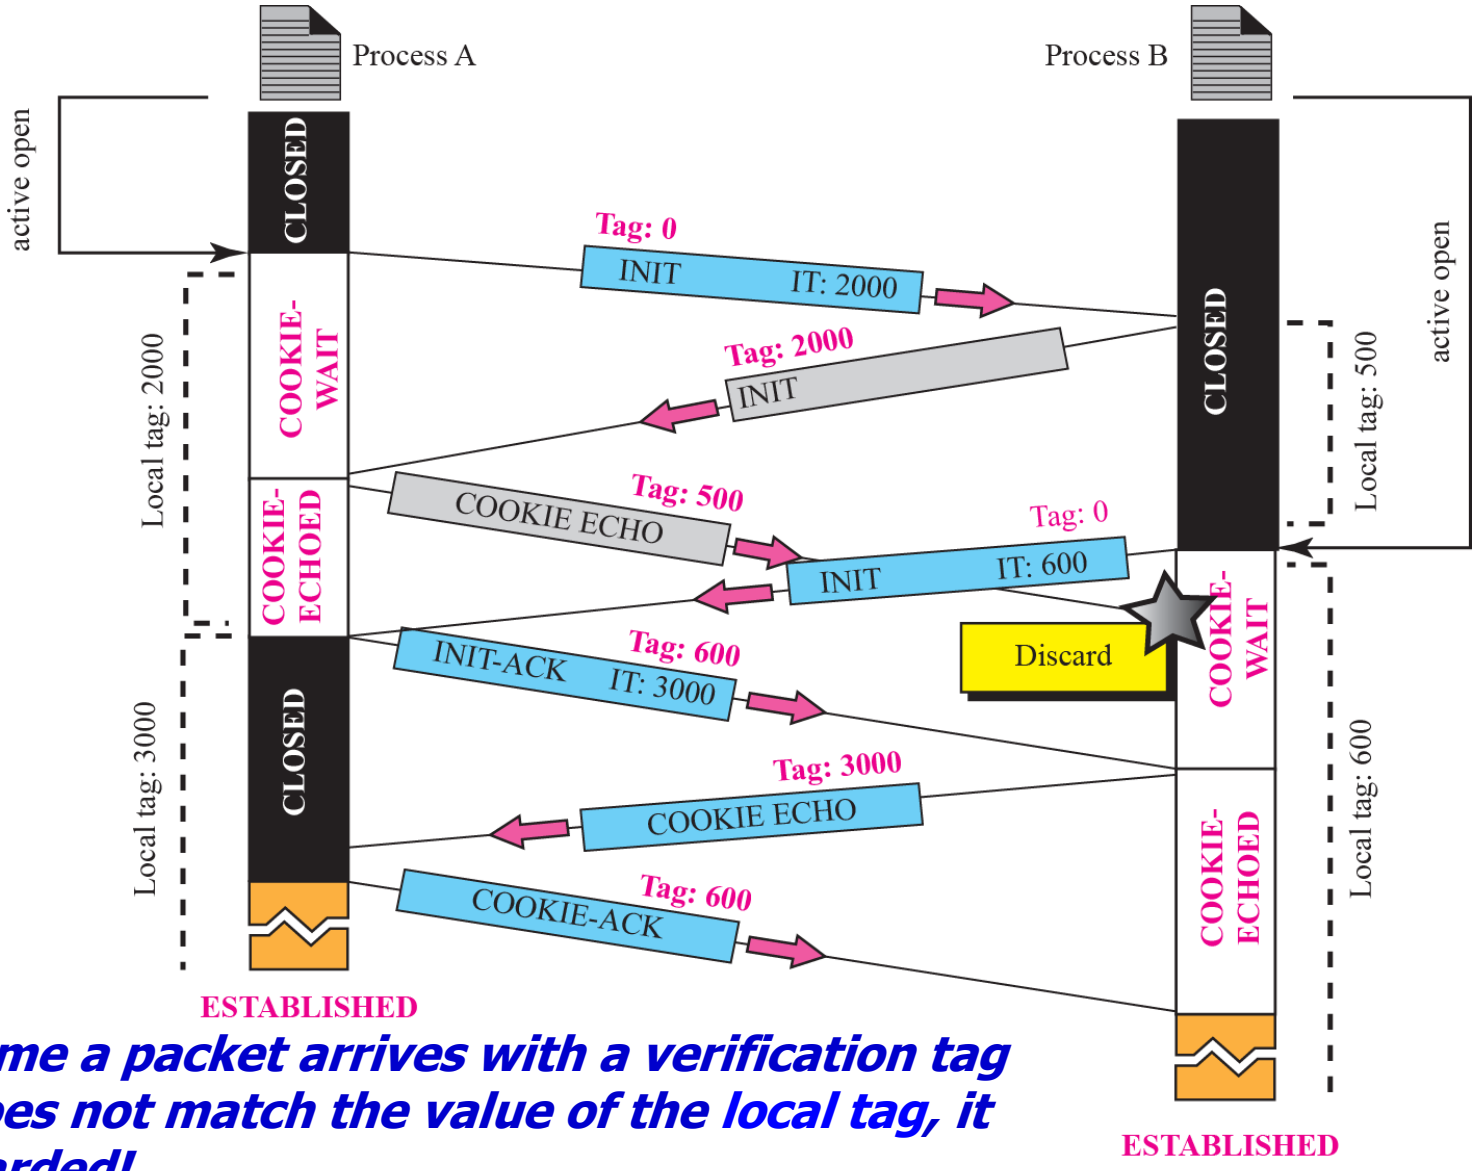

Each time a packet arrives with a verification tag that does not match the value of the local tag, it is discarded!

Figure 16.26 *Simultaneous close*

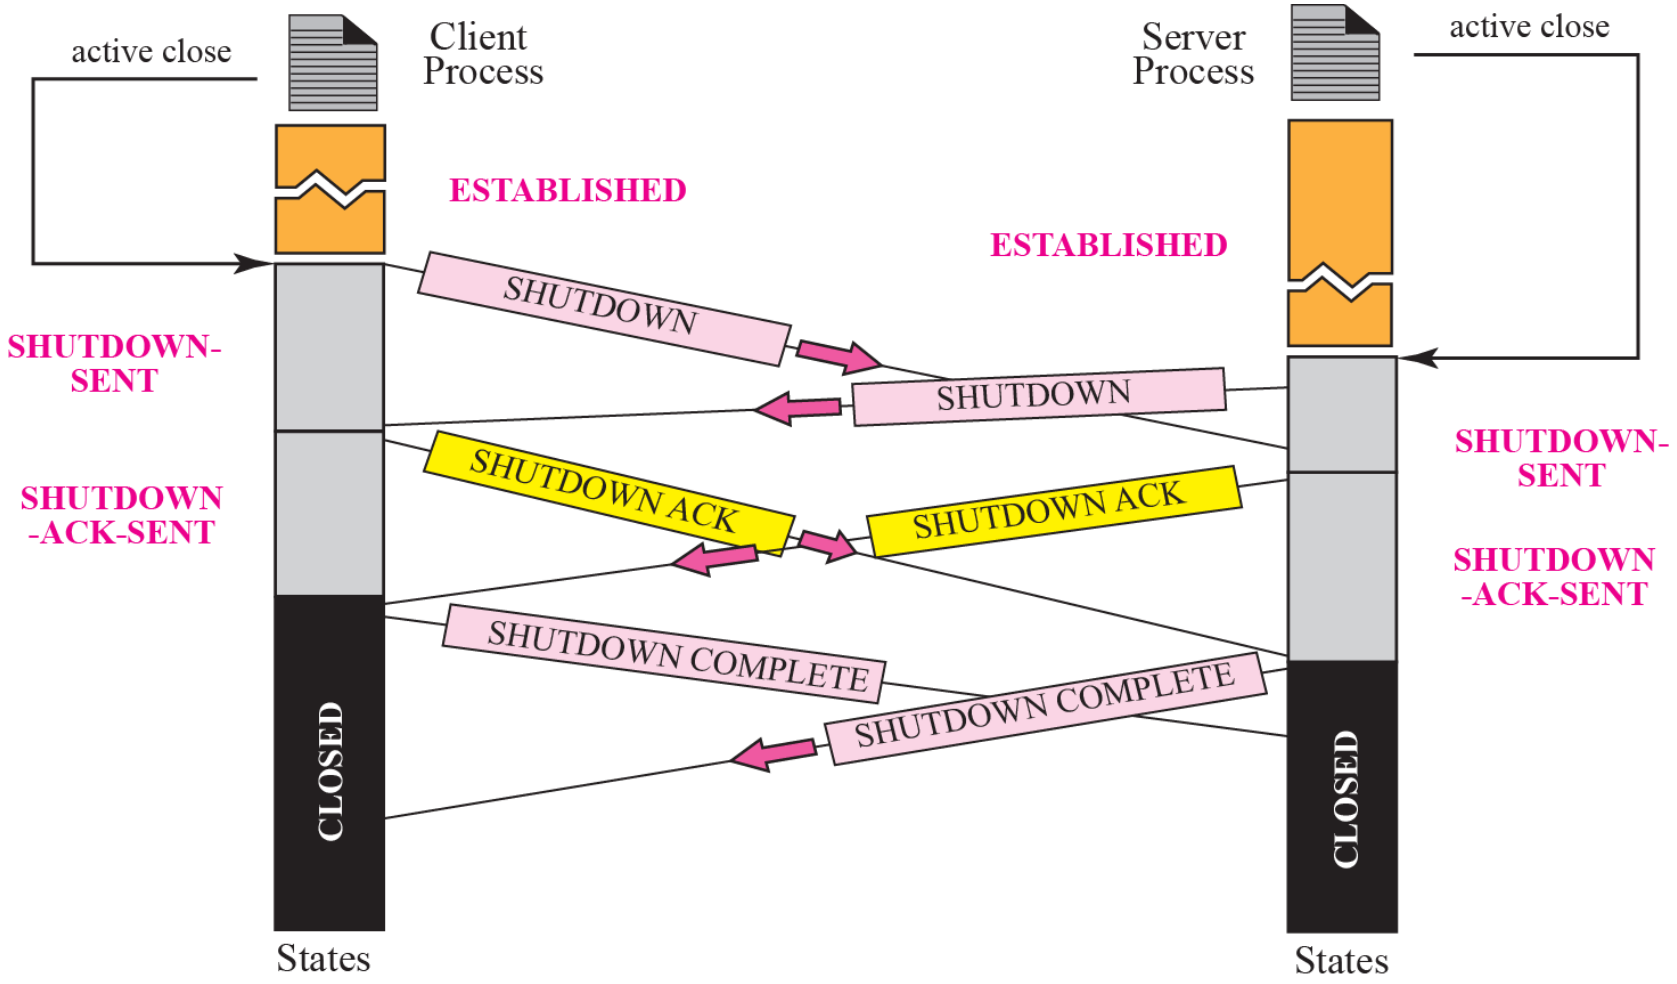

16-7 FLOW CONTROL

Flow control in SCTP is similar to that in TCP. In TCP, we need to deal with only one unit of data, the byte. In SCTP, we need to handle two units of data, the byte and the chunk. The values of rwnd and *cwnd* are expressed in bytes; the values of TSN and acknowledgments are expressed in chunks.

Topics Discussed in the Section

- ✓ **Receiver Site**
- ✓ **Sender Site**
- ✓ **A Scenario**

Figure 16.27 *Flow control, receiver site*

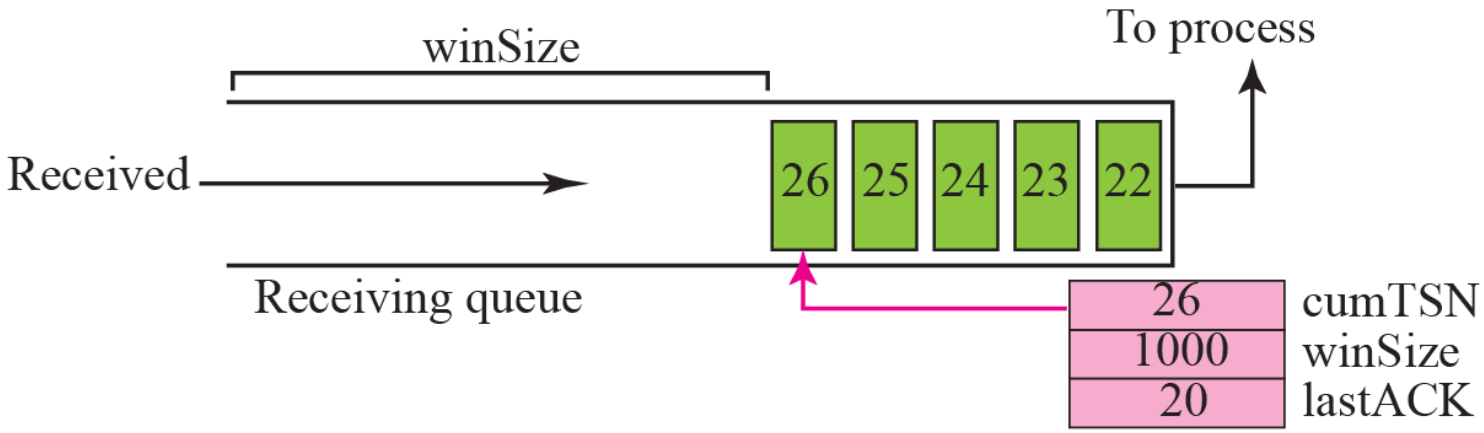

rwnd, cwnd: in bytes

TSN and Acknowledgement : in chunks

Figure 16.28 *Flow control, sender site*

1. A chunk pointed to by curTSN can be sent if the size of the data is less than or equal to the quantity (rwnd-inTransit)

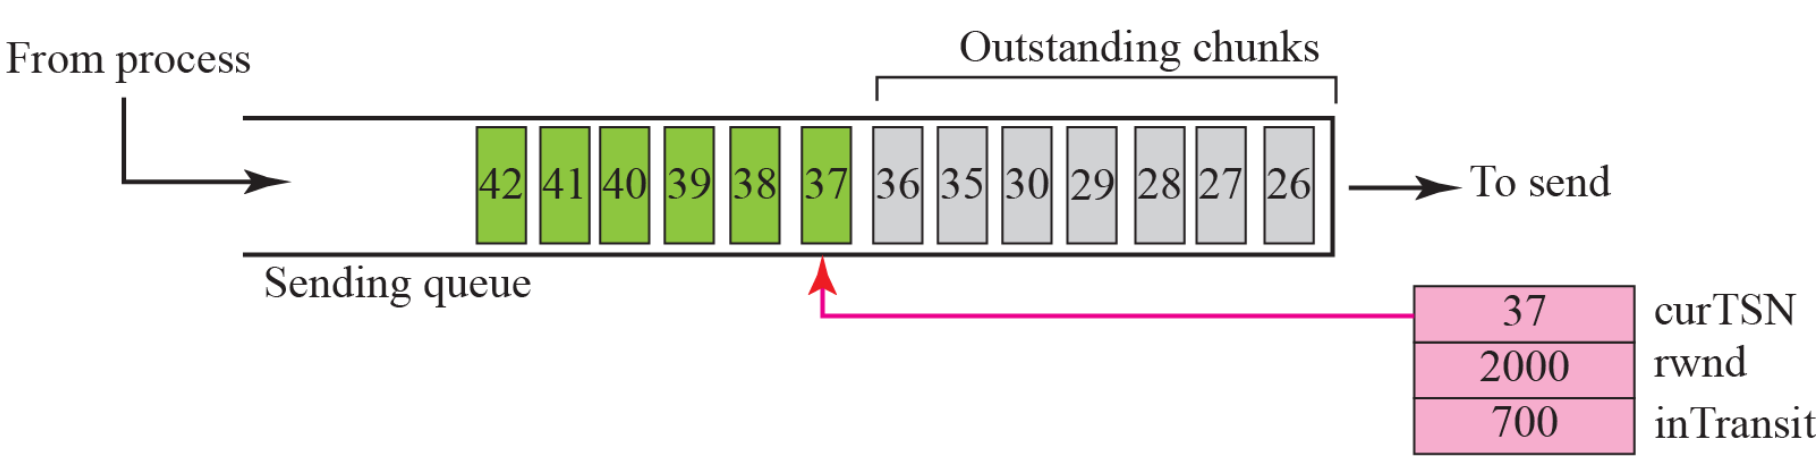

2. When a SACK is received, the chunks with a TSN less than or equal to the cumulative TSN in the SACK are removed from the queue and discarded. The values of rwnd and inTransit are updated properly

Figure 16.29 *Flow control scenario*

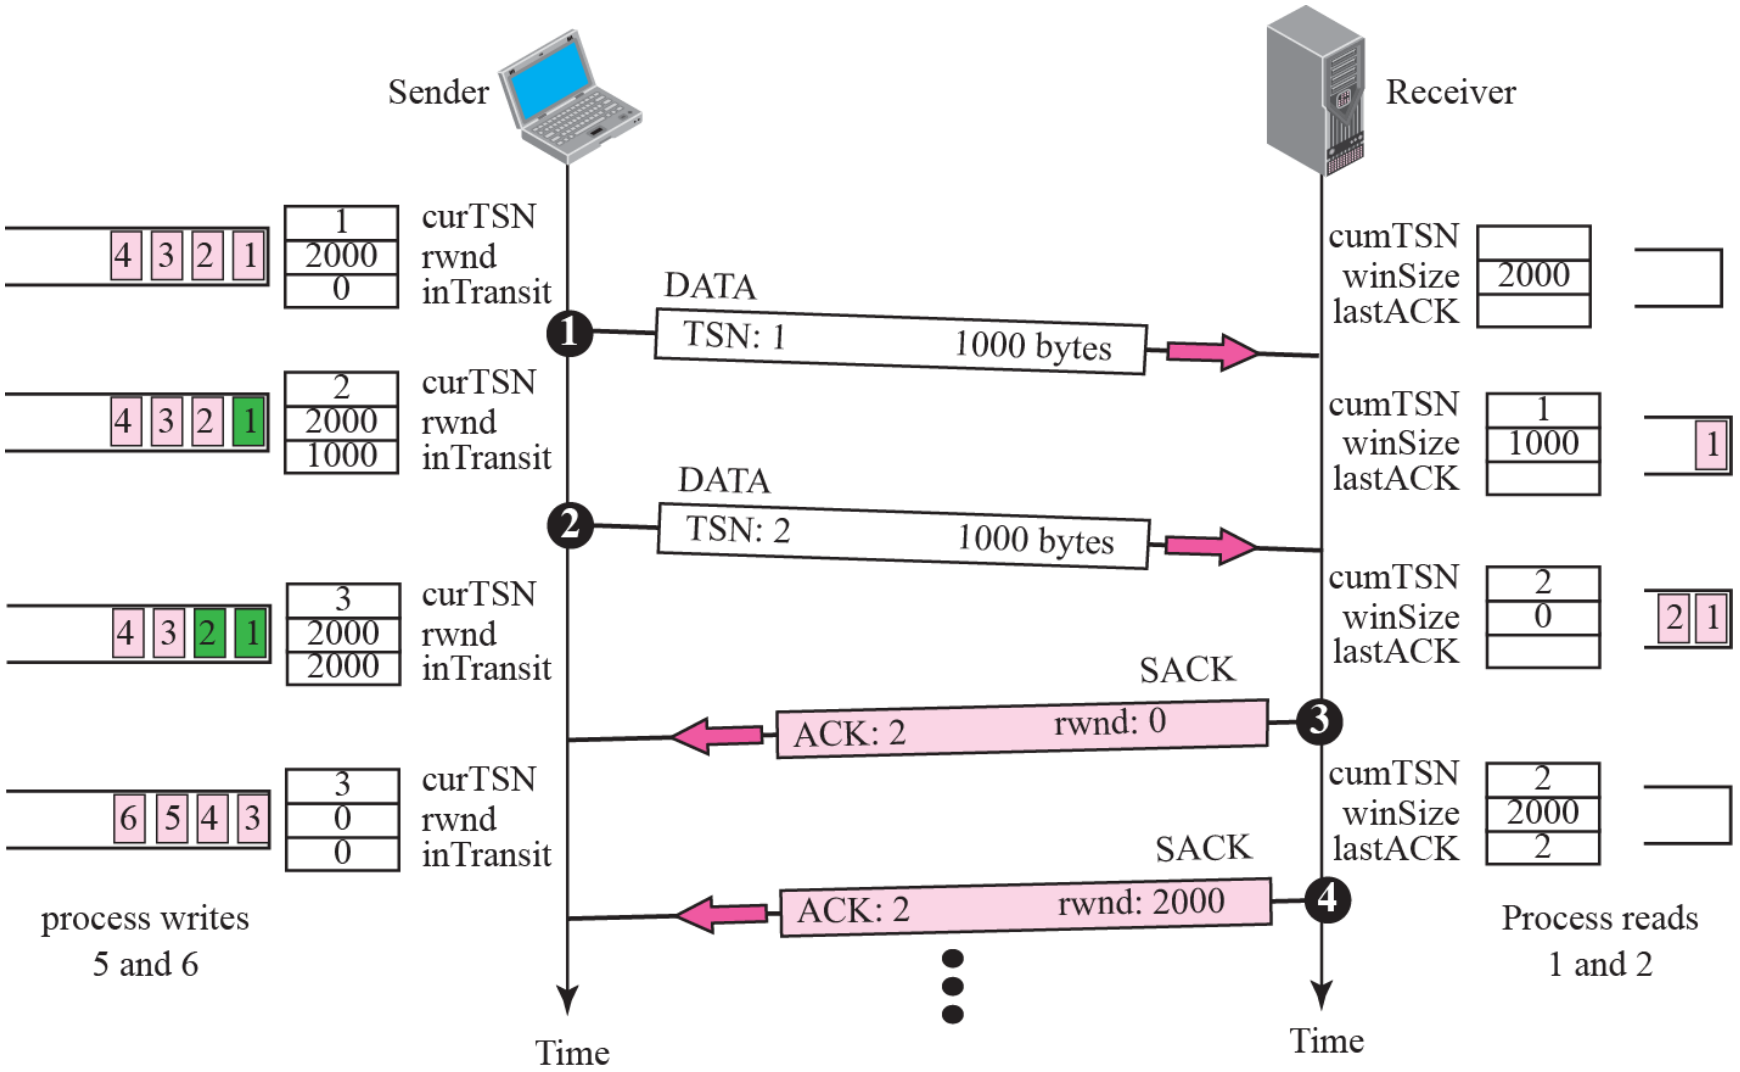

16-8 ERROR CONTROL

SCTP, like TCP, is a reliable transport-layer protocol. It uses a SACK chunk to report the state of the receiver buffer to the sender. Each implementation uses a different set of entities and timers for the receiver and sender sites. We use a very simple design to convey the concept to the reader.

Topics Discussed in the Section

- ✓ **Receiver Site**
- ✓ **Sender Site**
- ✓ **Sending Data Chunks**
- ✓ **Generating SANK Chunks**

Figure 16.30 *Error-control receiver site*

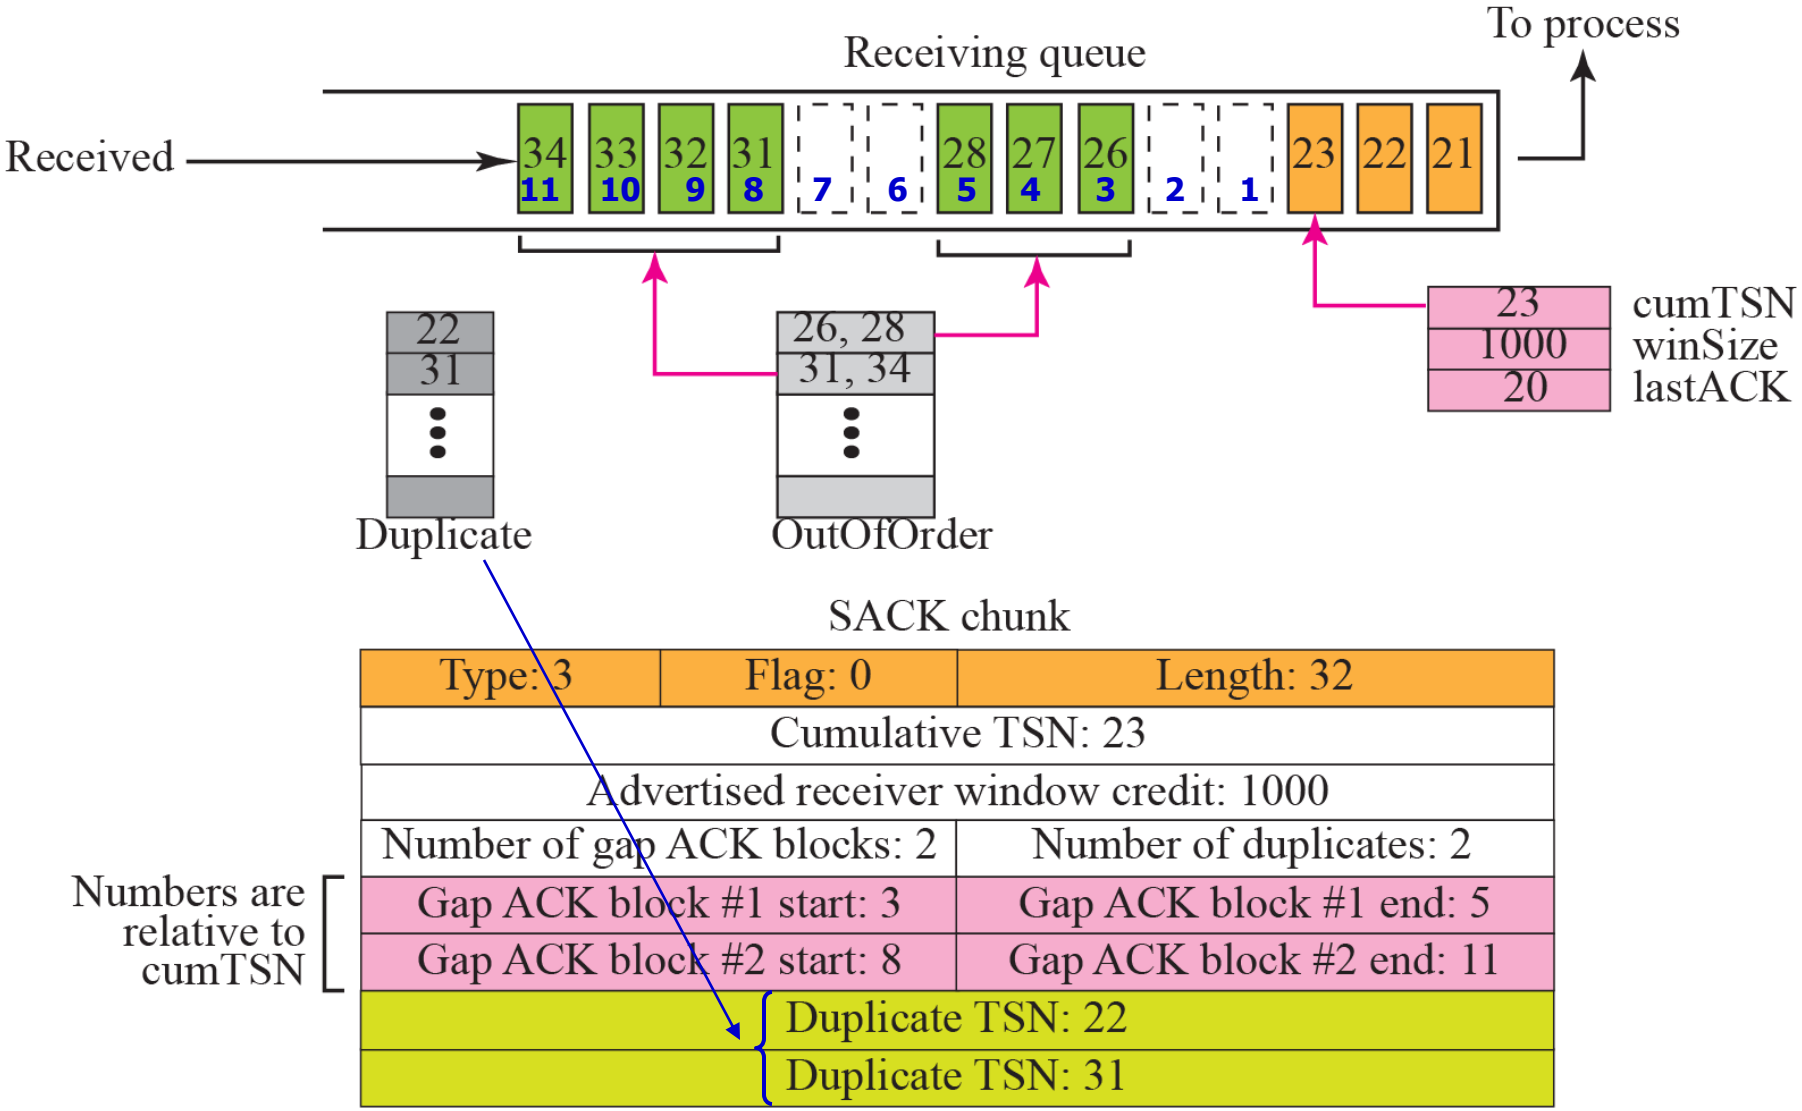

Figure 16.31 *Error control, sender site*

Assume 100 bytes per chunk

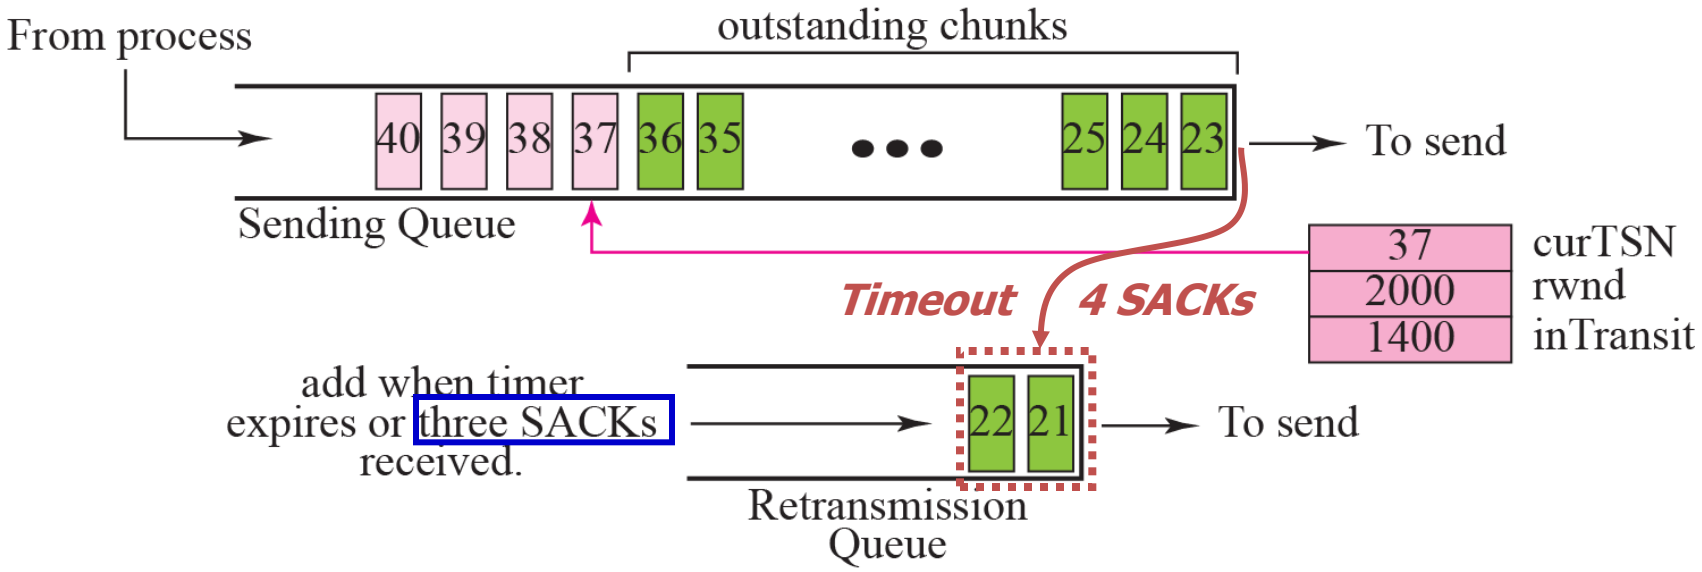

The chunks in the retransmission queue have priority

Figure 16.32 *New state at the sender site after receiving a SACK chunk*

1. ***Chunks 26-28, 31-34 are removed.***
2. ***The value of rwnd is changed to 1000 as advertised in the SACK chunk.***

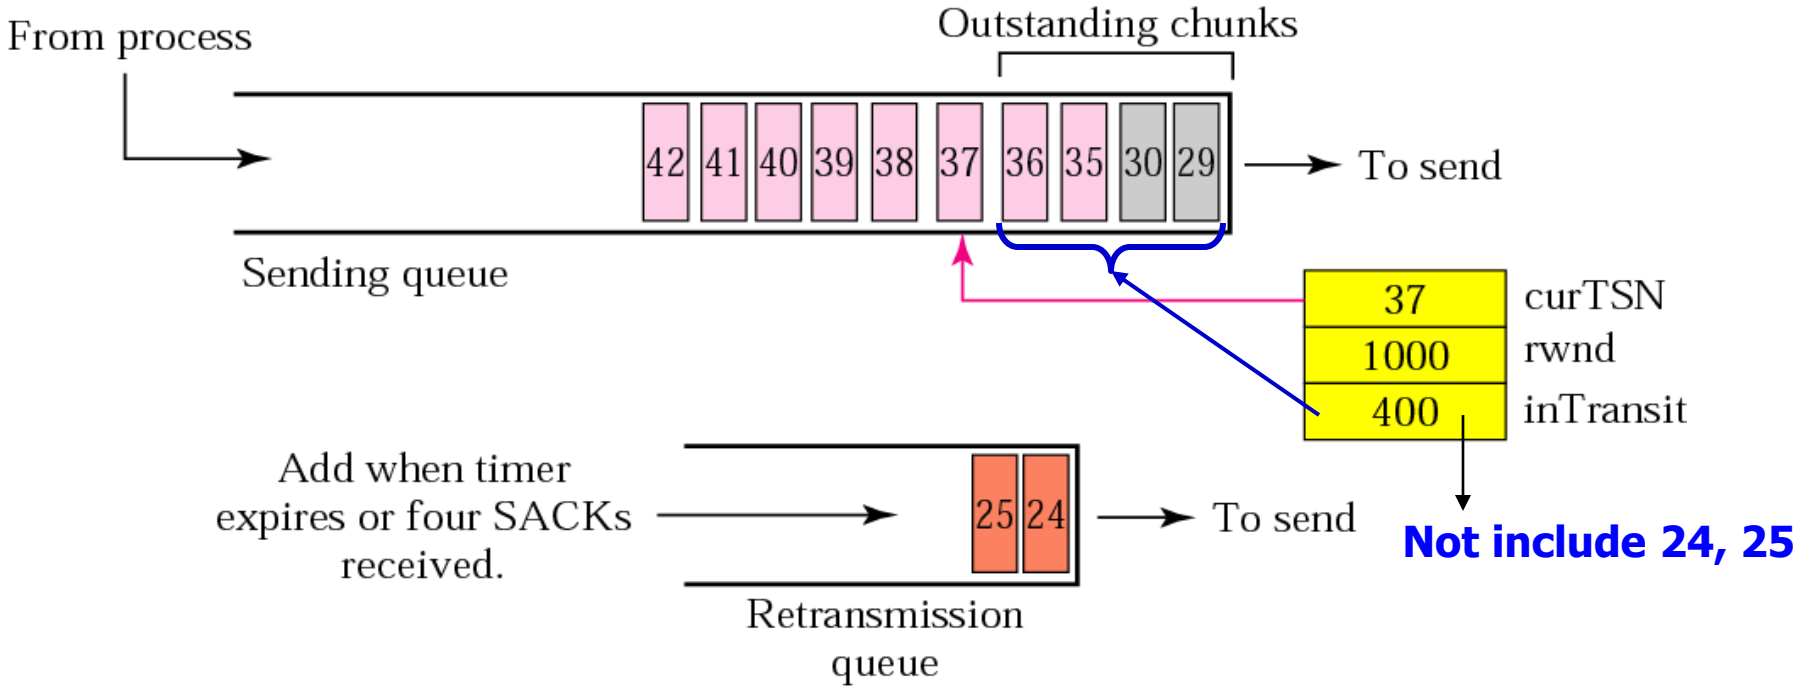

3. ***Also assume timer for chunks 24, 25 has expired. New TO is set according to exponential backoff rule in Chapter 12.***
4. ***4 chunks are now in transit, so inTransit becomes 400.***

Generating SACK Chunks

- Piggybacking
- Delay sending of SACK no more than 500ms
- An end must send at least one SACK for every other packet it receives
- Send a SACK immediately when
  - a packet arrives with out-of-order data chunks
  - a packet arrives with duplicate data chunks and no new data chunks

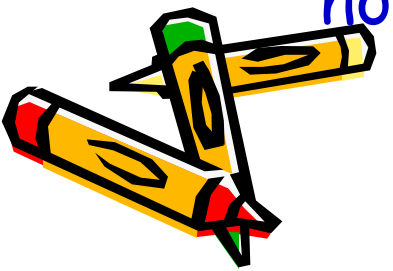

16-9 CONGESTION CONTROL

SCTP, like TCP, is a transport layer protocol with packets subject to congestion in the network. The SCTP designers have used the same strategies we described for congestion control in Chapter 15 for TCP. SCTP has slow start, congestion avoidance, and congestion detection phases. Like TCP, SCTP also uses fast retransmission and fast recovery.

Topics Discussed in the Section

- ✓ **Congestion Control and Multihoming**
- ✓ **Explicit Congestion Notification**

Need to have different values of cwnd for each IP address

It is a process that enables a receiver to explicitly inform the sender of any congestion experienced in the network. E.g. the receiver encounters many delayed or lost packets.