



Computer Network Courses

Ratan K Guha



Acknowledgement

- Dr. Gerald Marin first taught the undergraduate computer networks course at UCF and gave talk about the course in the summer workshop of 2002. He also developed the networking laboratory.



Computer Network Courses

- CDA 4506 (required for IT majors) Lab
- CDA 4527 Analysis
- CDA 5501 Network Architecture
- CDA 5530 Performance
- CDA 6520 Advanced



History of CDA 4527

- Spring 2002 and Spring 2003 offered the “Network Analysis” course out of Stallings and Trivedi concentrated on working text problems.
- Fall 2004 offered the “Network Analysis” course out of Kurose & Ross and Stallings concentrated on working text problems.



Course Goal

- This course is intended for students in CS and Engineering whose long-term goals may include design and development of networking hardware or software products as well as design and implementation of enterprise or service-provider networks.
- It introduces the concept of good network design using protocol and function layering.
- It interweaves the corresponding control issues throughout including throughput, delay, bandwidth management, congestion control, error control, sliding windows, retransmission strategies, contention resolution.



Textbooks

- A Tanenbaum, Computer Networks – Prentice Hall, Fourth edition, 2003.
- W. Stallings, High-Speed Networks and Internets - Prentice Hall, Second edition, 2002.
- J. Kurose and K. Ross, Computer Networking – Addison Wesley, Third edition, 2005.
- K. Trivedi, Probability and Statistics with Reliability, Queuing and Computer Science Applications – John Wiley, Second edition, 2002.
- W. Stallings, Wireless Communications and Networks - Prentice Hall, Second edition, 2005.



Major Course Constraint

- We (UCF) cannot assume students have had the Network Lab course (or any other networking course).
 - Network Lab course required for IT majors and most do not have calculus.
 - This course requires 2 semesters of calculus plus a course in statistics.
- Must repeat (filtered) some networking material.
 - Some students do take both courses.



Findings

- Many students did not know how to use a simple random variable, X , or compute $E(X)$.
- Virtually no students understood a simple Binomial distribution (certainly had no instinct for when/how to apply).
- Could not begin to explain any problem that might require a continuous probability distribution (such as talking about arrivals having an exponential distribution).
- Students challenged by mathematical notation.



Class Presentation

- 2003 – Stallings first three weeks and then one day Stallings and one day Trivedi in each week
- 2004 – Kurose and Ross as the text book and Stallings as the reference book for coverage of ATM, Powerpoint slides for wireless networks



Topics from Trivedi Includes:

- Introduction to Probability
- Discrete Random Variables
- Continuous Random Variables
- Expectation
- Stochastic Processes
- Queuing Analysis



Student Experience -Marin (Positive)

- 3 or 4 of 30 students were excited to learn how these problems could be treated.
- Approximately 10 students really seemed to have “mastered” most of the material.
- A number of students said it was there most difficult class at the university but they learned a great deal.



Student Experience- Marin (Negative)

- Many students felt the math was too demanding.
- Students with no previous network experience were overwhelmed.
- Students especially did not like the Stallings book because it did not show them how to work key problems.



Student Experience –Guha 2003

- Students liked having one day on computer networks and one day on probability in a week.
- Students liked the discussion of algorithms in chapters 11-13.
- Students did not like any of the text books



Student Experience –Guha 2004

- Students liked the text book.
- Students liked the discussion of fundamental concepts, algorithms and coverage of wireless networks.
- Students did not like the reference book



Introduction and Basics



Topics

- Important networking ideas: 1970's through 1990's.
- Networking Basics
- Connection-oriented vs connectionless
- Protocols and Building Blocks
- OSI and Course Reference Models
- Transmission Delay, Queuing Delay, Propagation Delay
- Error Detection and Correction Codes
- Example Networks



Network Topics Emphasis

- TCP/IP
- LAN - CSMA
- ATM – QoS
- 802.11



TCP/IP

- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Internet Protocol (IP)
- IPv6
- Flow Control and Congestion Control



Delays

- Example: Highway with toll booth
 - Toll Booth – Node
 - Highway – Transmission medium
- Transmission time
- Propagation delay
- Queuing delay



Error Detection and Correction

- Parity bits
- Checksum bits
- CRC codes
- Hamming Code (Single Error Correcting)



Hamming Code (Single Error Correcting)

- m - data; r - check; $m+r+1 \leq 2^r$
- Check bits at power of 2 - 1, 2, 4, 8
- Data bits at 3, 5, 6, 7, 9
- Example 4 data bits, 3 parity bits (Even)
- P1 - 1, 3, 5, 7; P2 - 2, 3, 6, 7
- P3 - 4, 5, 6, 7
- Even Parity Check - C1, C2, C3



CRC Codes

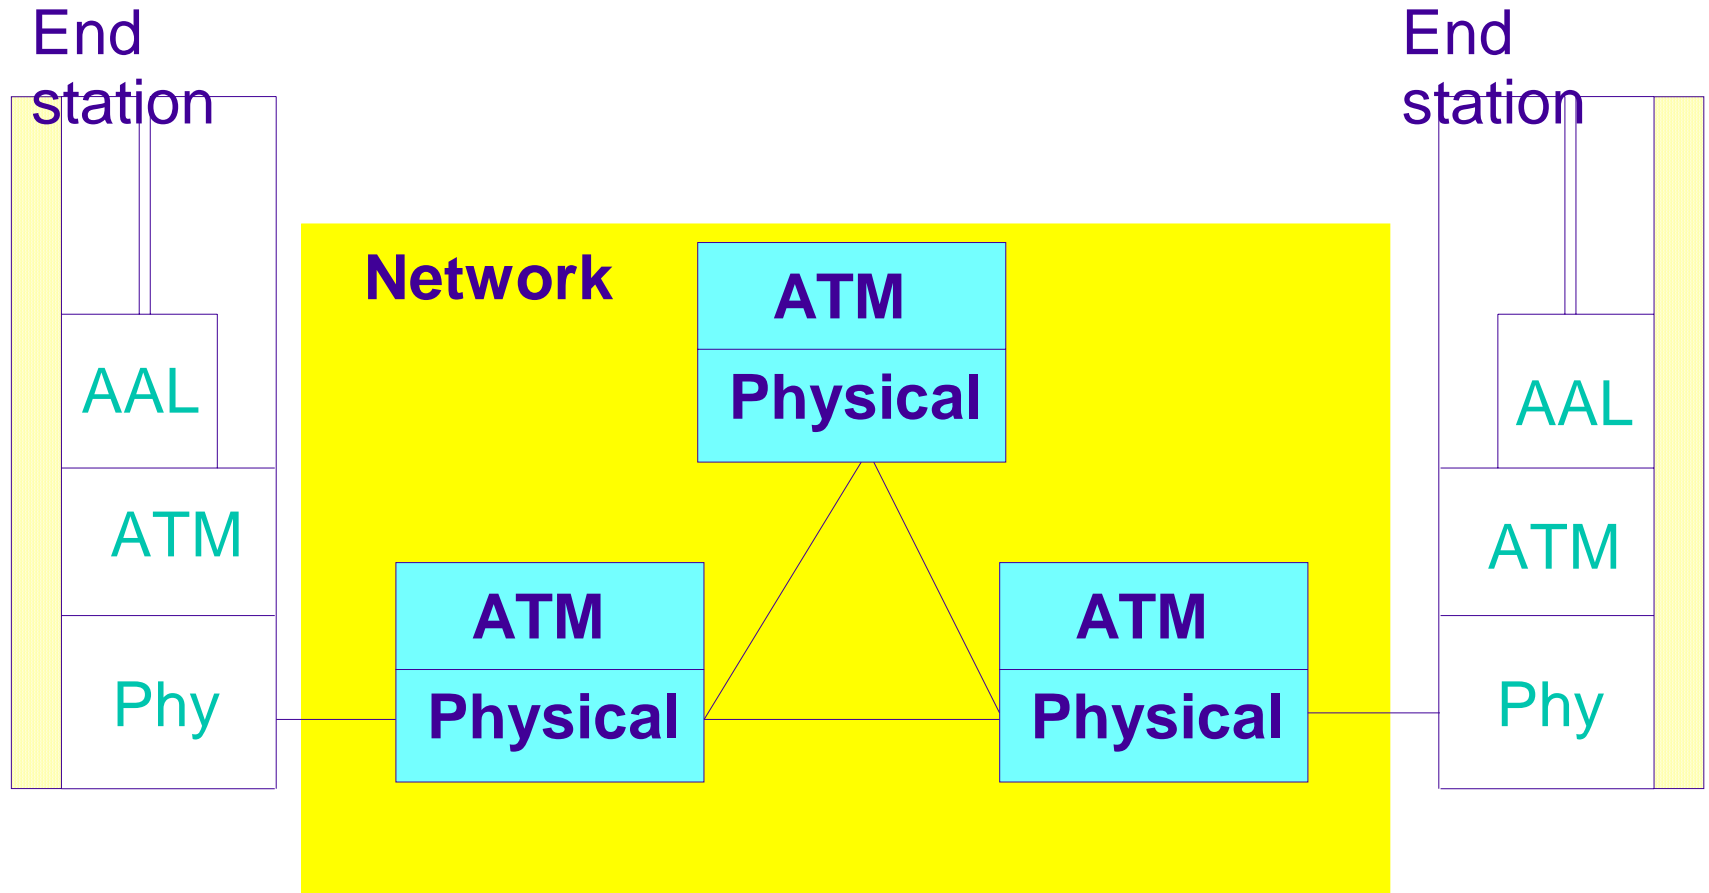
- Error detecting codes
- Message $m(x)$, Generator polynomial $G(x)$
- $x^r m(x)/G(x)$ using modulo 2 division
- Append remainder at the end
- CRC-16, CRC-CCITT
- Catches all single and double errors, all burst errors of length 16 or less, all errors with an odd no. of bits, 99.997% of 17-bit bursts, 99.998% of 18 bit and longer bursts



Checksum codes

- Used in header
- 16 bits as a number
- One's complement arithmetic
- One's complement of the sum as checksum

Introduction to ATM





Includes

- ATM Traffic-Related Attributes
- Traffic Management Framework
- Traffic Control
- ABR Traffic Management
- Emphasized on Algorithms



ATM

- Minimal error and flow control
- Data rates 155.52 Mbps, 622.08 Mbps, Gigabits/s
- ATM cells – 53 Octets; 5 Octet header
- Header Error Control on 32 bit header
- $G(x) = x^8 + x^2 + x + 1$
- Correction mode/Detection mode



Introduction

- Control needed to prevent switch buffer overflow
- High speed and small cell size gives different problems from other networks
- Limited number of overhead bits



Overview

- Congestion problem
- Framework adopted by ITU-T and ATM forum
 - Control schemes for delay sensitive traffic
 - Voice & video
 - Not suited to bursty traffic
 - Traffic control
 - Congestion control
- Bursty traffic
 - Available Bit Rate (ABR)
 - Guaranteed Frame Rate (GFR)



ATM Traffic-Related Attributes

- Six service categories
 - Constant bit rate (CBR)
 - Real time variable bit rate (rt-VBR)
 - Non-real-time variable bit rate (nrt-VBR)
 - Unspecified bit rate (UBR)
 - Available bit rate (ABR)
 - Guaranteed frame rate (GFR)
- Characterized by ATM attributes in four categories
 - Traffic descriptors
 - QoS parameters
 - Congestion
 - Other



Traffic Control

- Resource management using virtual paths
- Connection admission control
- User parameter control
 - Peak cell rate algorithm
 - Virtual scheduling algorithm
 - Leaky bucket algorithm
 - Sustainable cell rate algorithm
- Selective cell discard
- Traffic shaping
- Explicit forward congestion indication



Some slides of CDA 4527

Chapters 2 and 3 - Stallings



Protocols and the TCP/IP Suite

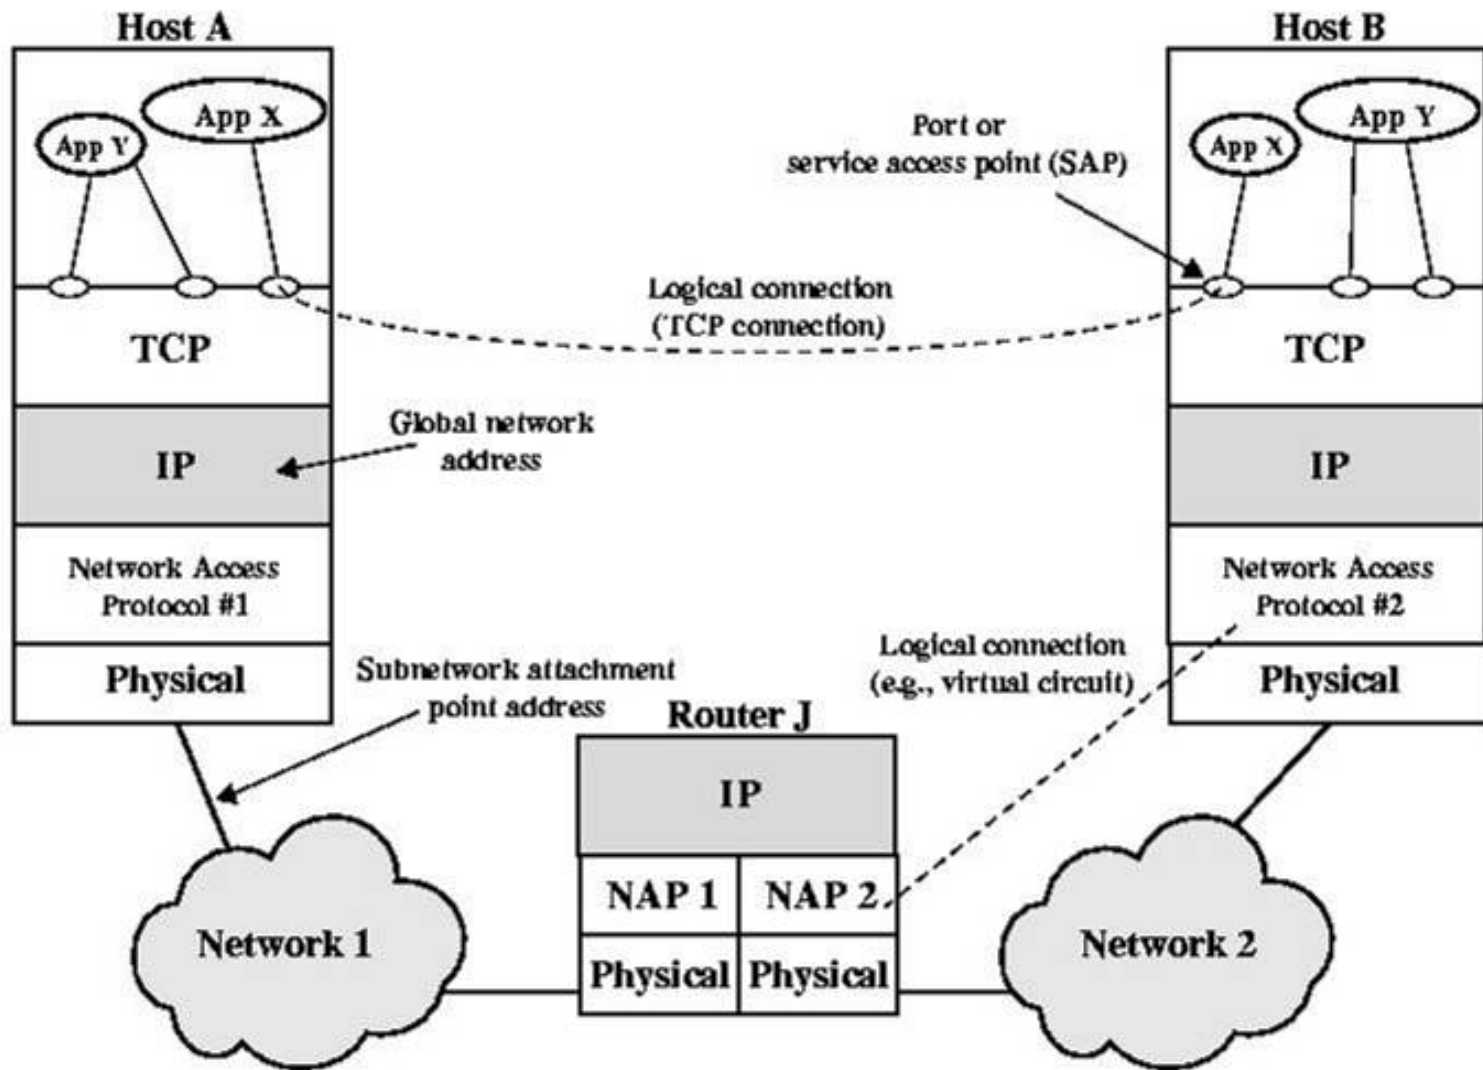


Figure 2.3 TCP/IP Concepts




Figure 2.4 Protocol Data Units (PDUs) in the TCP/IP Architecture



TCP and UDP

- TCP:
 - connection-oriented
 - Reliable packet delivery in sequence
- UDP:
 - connectionless (datagram)
 - Unreliable packet delivery
 - Packets may arrive out of sequence or duplicated

AND SO ON...



Routers

- Provide link between networks
- Accommodate network differences:
 - Addressing schemes
 - Maximum packet sizes
 - Hardware and software interfaces
 - Network reliability





Figure 2.7 Configuration for TCP/IP Example

1. Preparing the data. The application protocol prepares a block of data for transmission. For example, an email message (SMTP), a file (FTP), or a block of user input (TELNET).

2. Using a common syntax. If necessary, the data are converted to a form expected by the destination. This may include a different character code, the use of encryption, and/or compression.

3. Segmenting the data. TCP may break the data block into a number of segments, keeping track of their sequence. Each TCP segment includes a header containing a sequence number and a frame check sequence to detect errors.

4. Duplicating segments. A copy is made of each TCP segment, in case the loss or damage of a segment necessitates retransmission. When an acknowledgment is received from the other TCP entity, a segment is erased.

5. Fragmenting the segments. IP may break a TCP segment into a number of datagrams to meet size requirements of the intervening networks. Each datagram includes a header containing a destination address, a frame check sequence, and other control information.

6. Framing. An ATM header is added to each IP datagram to form an ATM cell. The header contains a connection identifier and a header error control field.

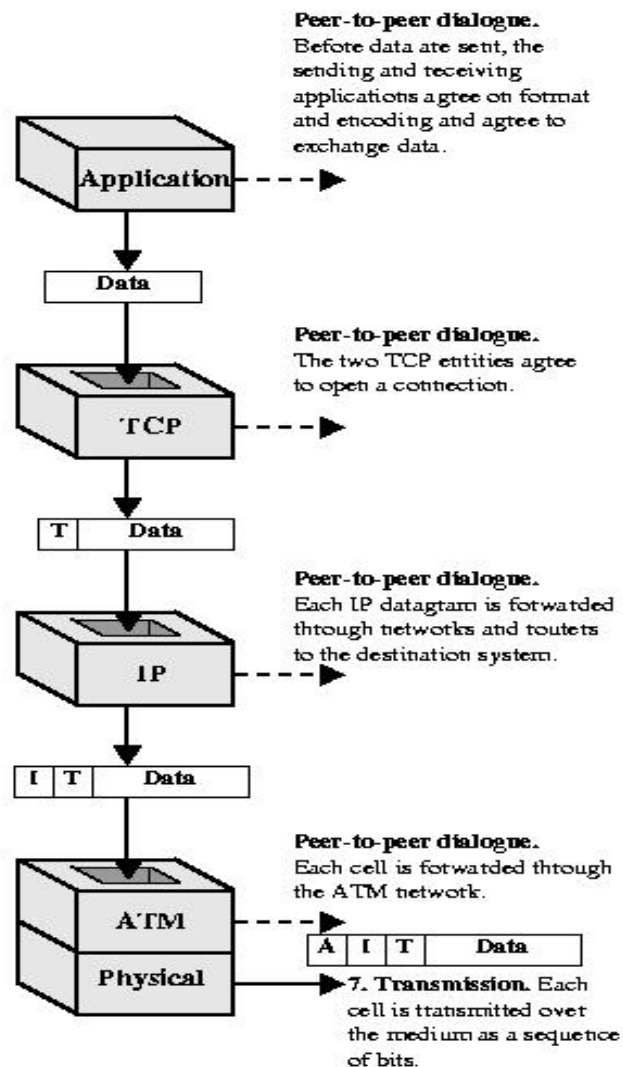


Figure 2.8 Operation of TCP/IP: Action at Sender

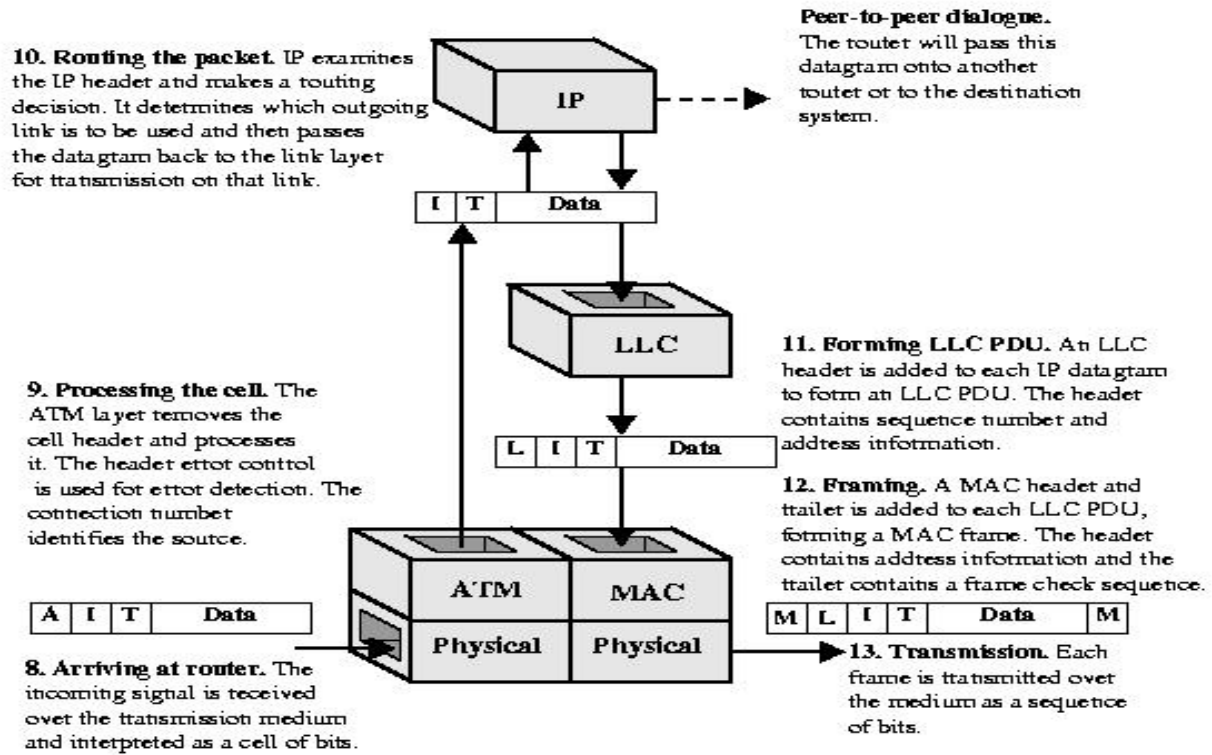




Figure 2.9 Operation of TCP/IP: Action at Router



20. Delivering the data. The application performs any needed transformations, including decompression and decryption, and directs the data to the appropriate file or other destination.

19. Reassembling user data. If TCP has broken the user data into multiple segments, these are reassembled and the block is passed up to the application.

18. Processing the TCP segment. TCP removes the header. It checks the frame check sequence and acknowledges if there is a match and discards for mismatch. Flow control is also performed.

17. Processing the IP datagram. IP removes the header. The frame check sequence and other control information are processed.

16. Processing the LLC PDU. The LLC layer removes the header and processes it. The sequence number is used for flow and error control.

15. Processing the frame. The MAC layer removes the header and trailer and processes them. The frame check sequence is used for error detection.

14. Arriving at destination. The incoming signal is received over the transmission medium and interpreted as a frame of bits.





Figure 2.10 Operation of TCP/IP: Action at Receiver



Chapter 4



Frame Relay



Introduction

- Circuit Switching Networks
- Packet-Switching Networks
 - Switching Technique
 - Routing
 - Datagram
 - Virtual Circuit
- Frame Relay Networks
 - Architecture
 - User Data Transfer
 - Call Control

Frame Switching: Making a network look like a link.





Two Possible Approaches

- Terminate the link control protocol (much like router) and use network headers to route across the network.
 - Problem: Must understand various layer 3's
- Put our own network headers onto whatever is sent by Box A (another envelope), route through the network, then strip off headers and send to Box B.
 - Frame Switching does this.



Frame Switching Problems

- Link control protocols have very short time-outs and network delay is irregular.
- Many link control protocols use polling; sending polls across a network increases network overhead.
- The network's internal routing protocol must use a known address structure and these internal addresses must somehow be associated with external box addresses.
- Network will not be able to recognize priorities or TOS bits in network layer header.



Frame Relay Concept

- The concept of frame relay is simple. A network is interposed between devices communicating on a link, but the devices are not aware that this has happened. Because of the problems already discussed, this is not simple in implementation. However, it is a simple concept.



Link-Level Flow and Error Control Includes

- Flow Control
- Link Control Mechanisms
 - Stop and Wait
 - Sliding-Window Techniques
 - Go-back-N ARQ
 - Selective-Reject ARQ
- ARQ Performance
- HDLC, CRC



Introduction

- The need for flow and error control
- Link control mechanisms
- Performance of ARQ (Automatic Repeat Request)



Flow Control and Error Control

- Fundamental mechanisms that determine performance
- Can be implemented at different levels:
link, network, or application
- Difficult to model performance
- Simplest case: point-to-point link
 - Constant propagation
 - Constant data rate
 - Probabilistic error rate
 - Traffic characteristics



Flow Control

- Limits the amount or rate of data that is sent
- Reasons:
 - Source may send PDUs faster than destination can process headers
 - Higher-level protocol user at destination may be slow in retrieving data
 - Destination may need to limit incoming flow to match outgoing flow for retransmission



Flow Control at Multiple Protocol Layers

- X.25 virtual circuits (level 3) multiplexed over a data link using LAPB (X.25 level 2)
- Multiple TCP connections over HDLC link
- Flow control at higher level applied to each logical connection independently
- Flow control at lower level applied to total traffic




Figure 11.1 Flow Control at Multiple Protocol Layers



Flow Control Scope

- Hop Scope
 - Between intermediate systems that are directly connected
- Network interface
 - Between end system and network
- Entry-to-exit
 - Between entry to network and exit from network
- End-to-end
 - Between end user systems

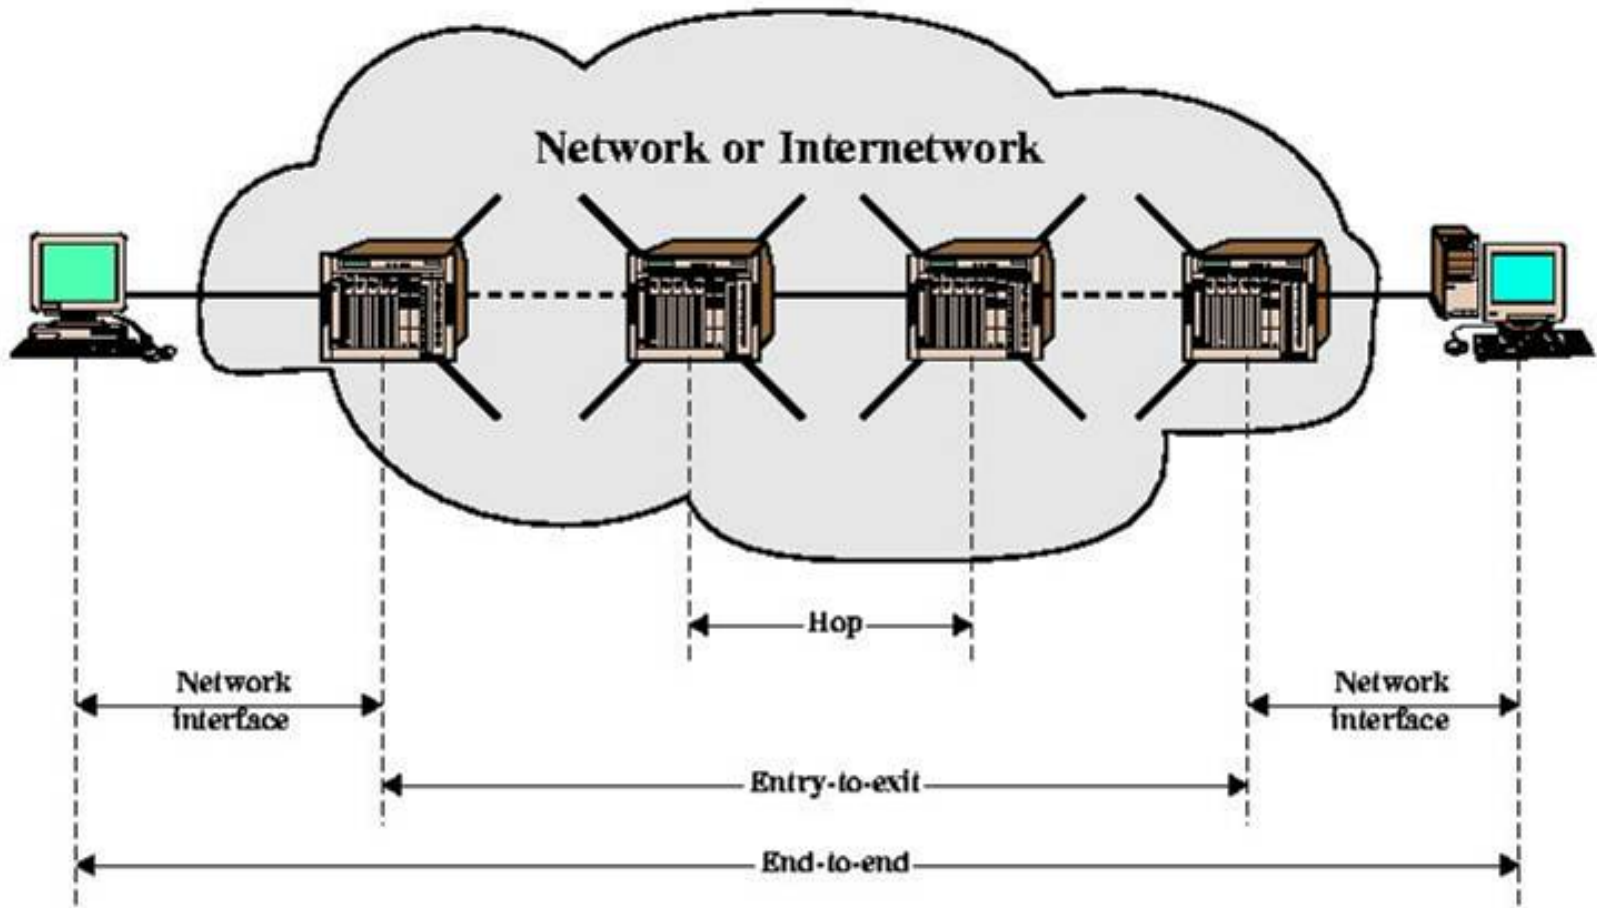


Figure 11.2 Flow Control Scope



Error Control

- Used to recover lost or damaged PDUs
- Involves error detection and PDU retransmission
- Implemented together with flow control in a single mechanism
- Performed at various protocol levels



Link Control Mechanisms

3 techniques at link level:

- Stop-and-wait
- Go-back-N
- Selective-reject

Latter 2 are special cases of sliding-window

Assume 2 end systems connected by direct link



Sequence of Frames

Source breaks up message into sequence of frames

- Buffer size of receiver may be limited
- Longer transmission are more likely to have an error
- On a shared medium, avoids one station monopolizing medium



Stop and Wait

- Source transmits frame
- After reception, destination indicates willingness to accept another frame in acknowledgement
- Source must wait for acknowledgement before sending another frame
- 2 kinds of errors:
 - Damaged frame at destination
 - Damaged acknowledgement at source



ARQ

- Automatic Repeat Request
- Uses:
 - Error detection
 - Timers
 - Acknowledgements
 - Retransmissions



TCP Traffic Control Includes

- TCP Flow Control
- TCP Congestion Control
 - Window Management



TCP Flow Control

- Uses a form of sliding window
- Differs from mechanism used in LLC, HDLC, X.25, and others:
 - Decouples acknowledgement of received data units from granting permission to send more
- TCP's flow control is known as a credit allocation scheme:
 - Each transmitted octet is considered to have a sequence number



TCP Header Fields for Flow Control

- Sequence number (SN) of first octet in data segment
- Acknowledgement number (AN)
- Window (W)
- Acknowledgement contains $AN = i$, $W = j$:
 - Octets through $SN = i - 1$ acknowledged
 - Permission is granted to send $W = j$ more octets,
i.e., octets i through $i + j - 1$

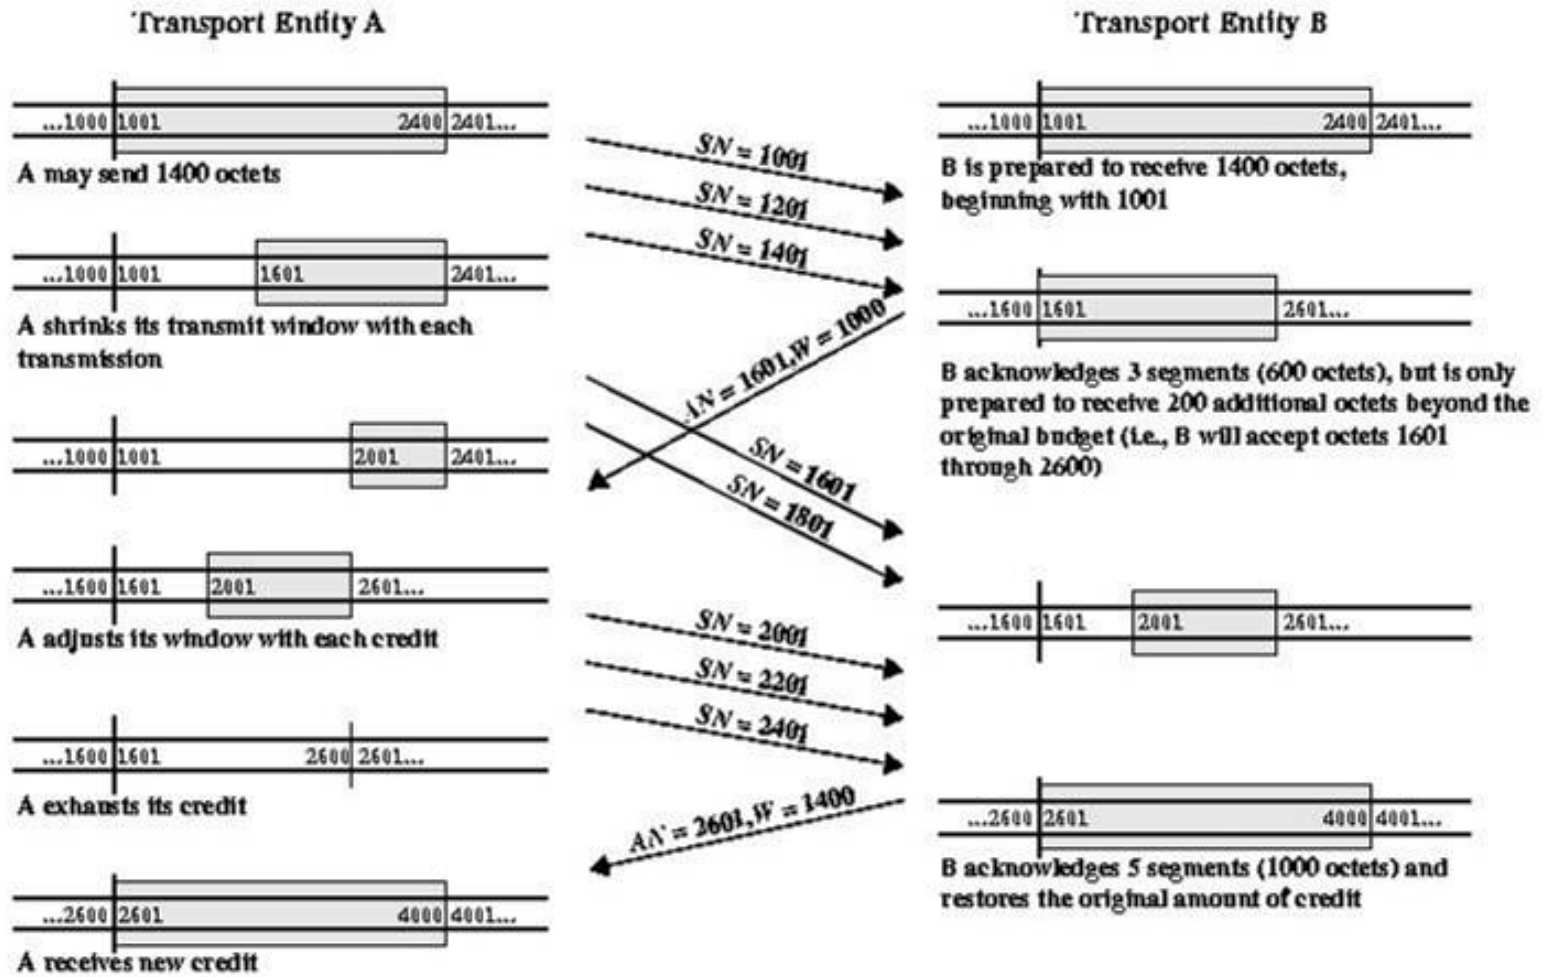


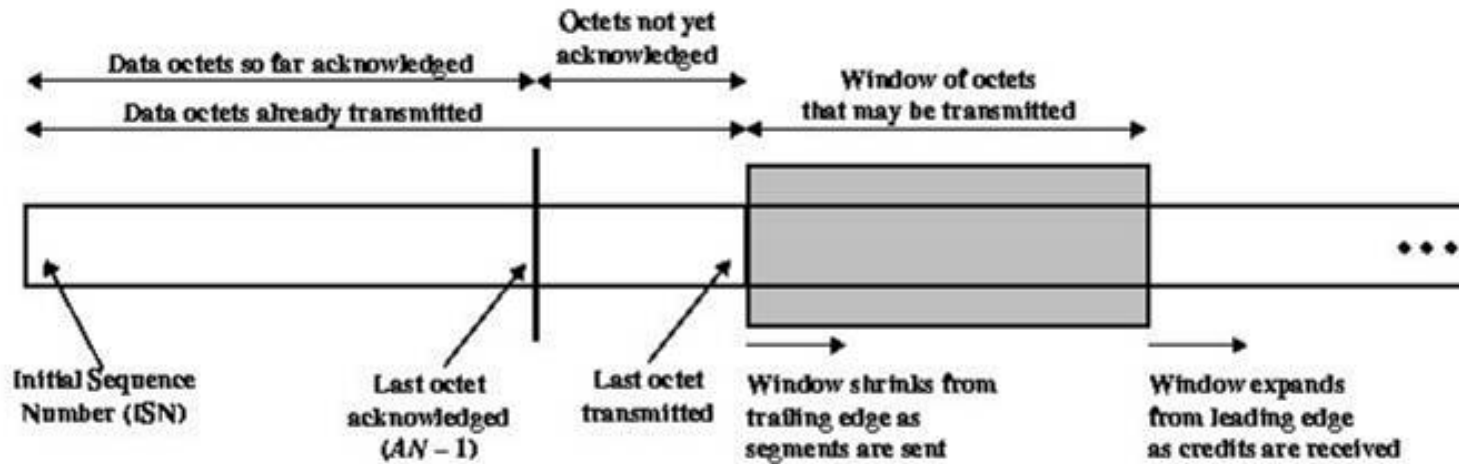
Figure 12.1 Example of TCP Credit Allocation Mechanism



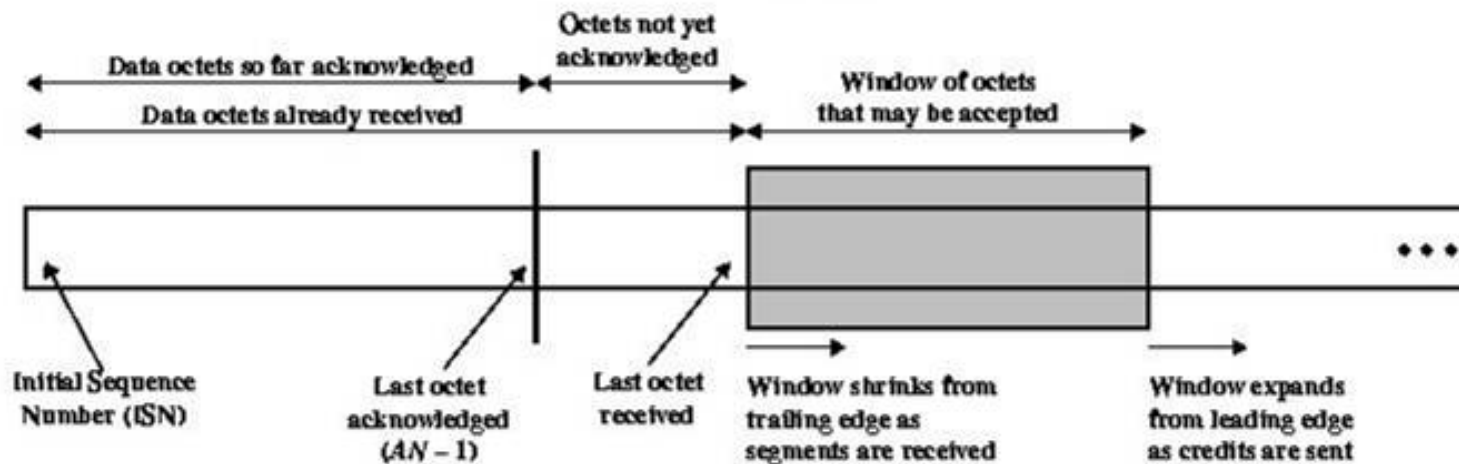
Credit Allocation is Flexible

Suppose last message B issued was $AN = i$, $W = j$

- To increase credit to k ($k > j$) when no new data, B issues $AN = i$, $W = k$
- To acknowledge segment containing m octets ($m < j$), B issues $AN = i + m$, $W = j - m$



(a) Send sequence space



(b) Receive sequence space

Figure 12.2 Sending and Receiving Flow Control Perspectives



Credit Policy

- Receiver needs a policy for how much credit to give sender
- Conservative approach: grant credit up to limit of available buffer space
- May limit throughput in long-delay situations
- Optimistic approach: grant credit based on expectation of freeing space before data arrives



TCP Congestion Control

- Dynamic routing can alleviate congestion by spreading load more evenly
- But only effective for unbalanced loads and brief surges in traffic
- Congestion can only be controlled by limiting total amount of data entering network
- ICMP source Quench message is crude and not effective
- RSVP may help but not widely implemented



TCP Congestion Control is Difficult

- IP is connectionless and stateless, with no provision for detecting or controlling congestion
- TCP only provides end-to-end flow control
- No cooperative, distributed algorithm to bind together various TCP entities



TCP Flow and Congestion Control

- The rate at which a TCP entity can transmit is determined by rate of incoming ACKs to previous segments with new credit
- Rate of Ack arrival determined by round-trip path between source and destination
- Bottleneck may be destination or internet
- Sender cannot tell which
- Only the internet bottleneck can be due to congestion



Window Management

- Slow start
- Dynamic window sizing on congestion
- Fast retransmit
- Fast recovery
- Limited transmit

Traffic and Congestion Control in ATM Networks

