Computer Communication Networks

Final Review

ICEN/ICSI 416 – Fall 2017

Prof. Dola Saha

What is included?

- Network Layer
- Link Layer
- Physical Layer
- Network Security

IP datagram format

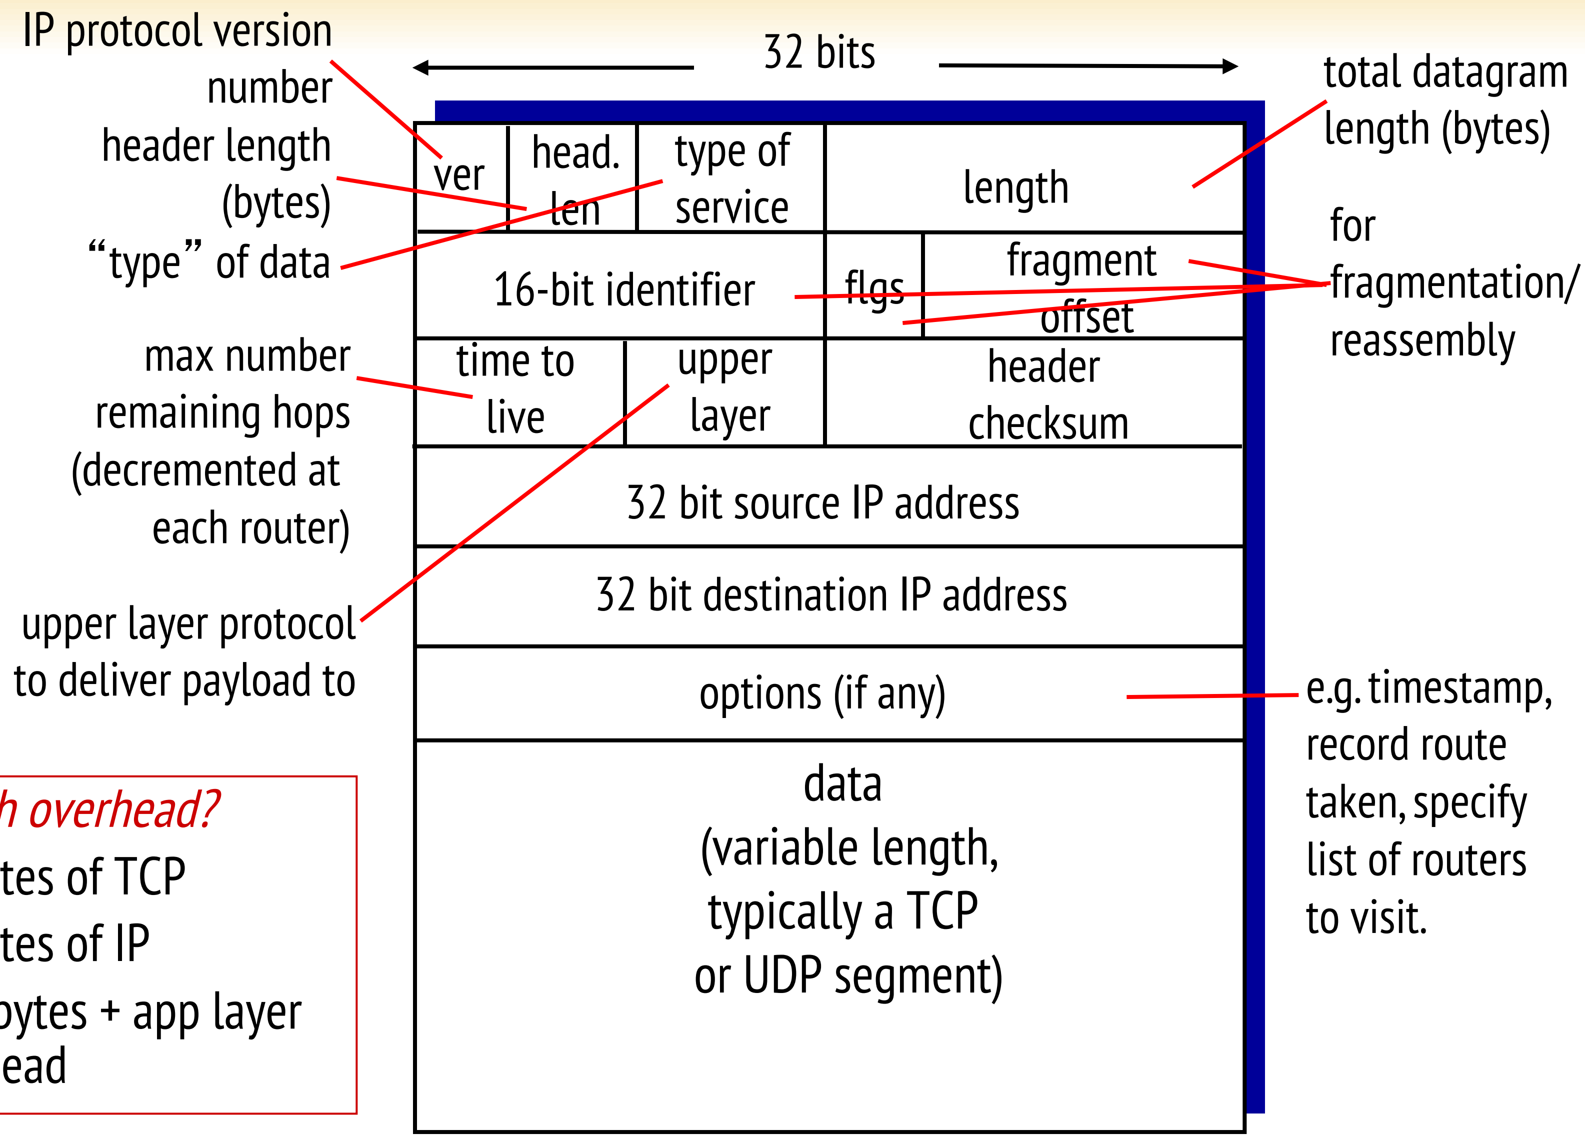

how much overhead?

- 20 bytes of TCP
- 20 bytes of IP
- = 40 bytes + app layer overhead

Router architecture overview

- high-level view of generic router architecture:

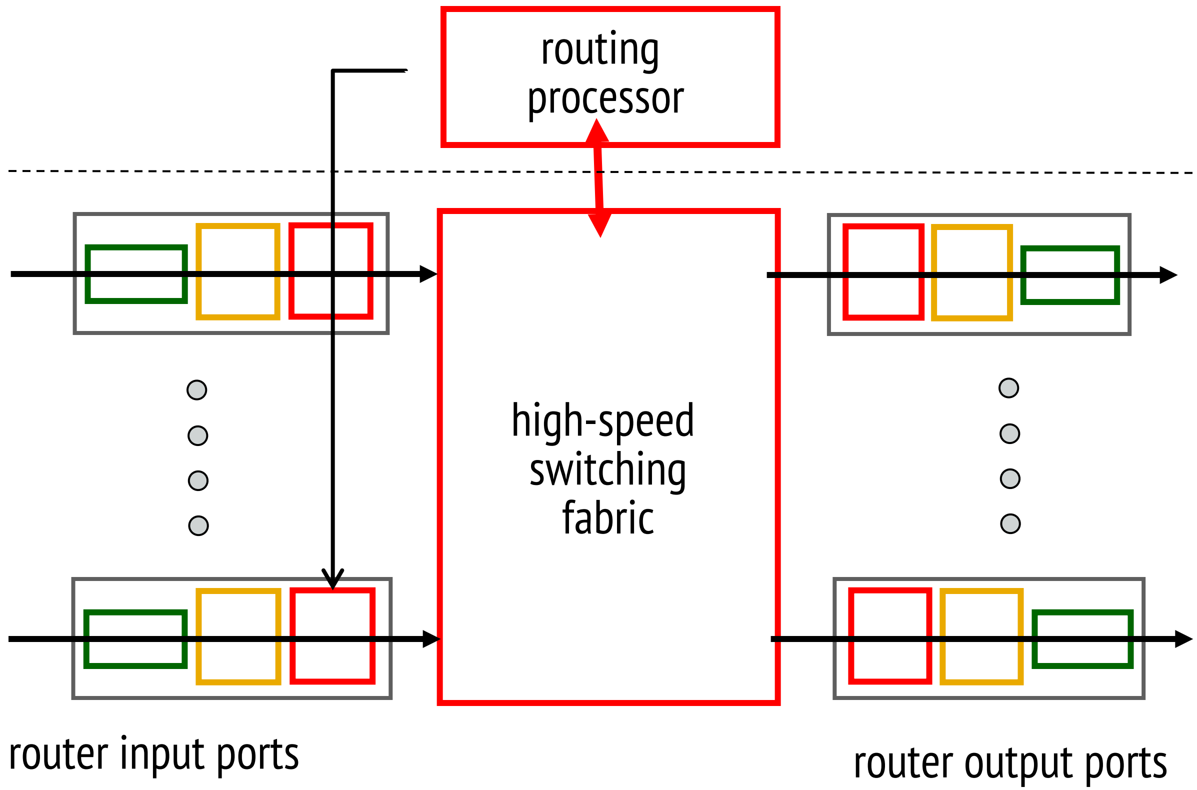

routing, management
control plane (software)
operates in millisecond
time frame

forwarding data plane
(hardware) operttes in
nanosecond timeframe

Longest prefix matching

longest prefix matching

when looking for forwarding table entry for given destination address, use *longest* address prefix that matches destination address.

Destination Address Range	Link Interface
11001000 00010111 00010*** *********	0
11001000 00010111 00011000 *********	1
11001000 00010111 00011*** *********	2
otherwise	3

examples:

DA: 11001000 00010111 00010110 10100001 which interface?

DA: 11001000 00010111 00011000 10101010 which interface?

Input port queuing

- fabric slower than input ports combined -> queueing may occur at input queues
  - *queueing delay and loss due to input buffer overflow!*
- Head-of-the-Line (HOL) blocking: queued datagram at front of queue prevents others in queue from moving forward

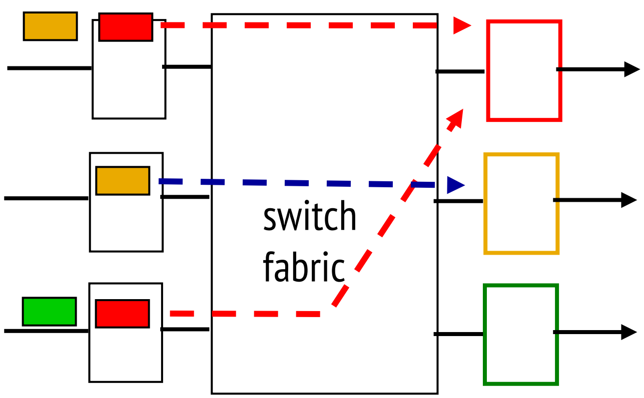

output port contention:
only one red datagram can be transferred.
lower red packet is blocked

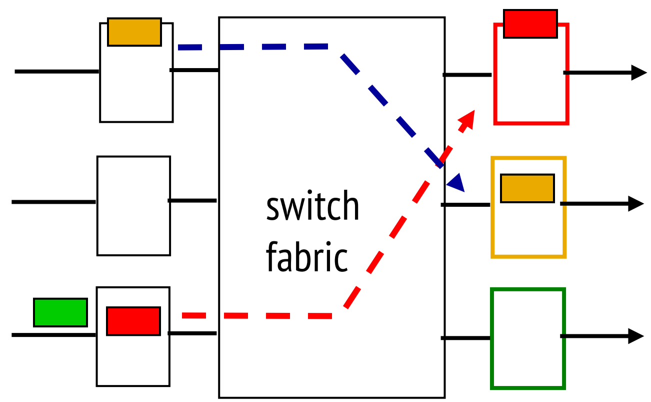

one packet time later:
green packet experiences HOL blocking

Output port queueing

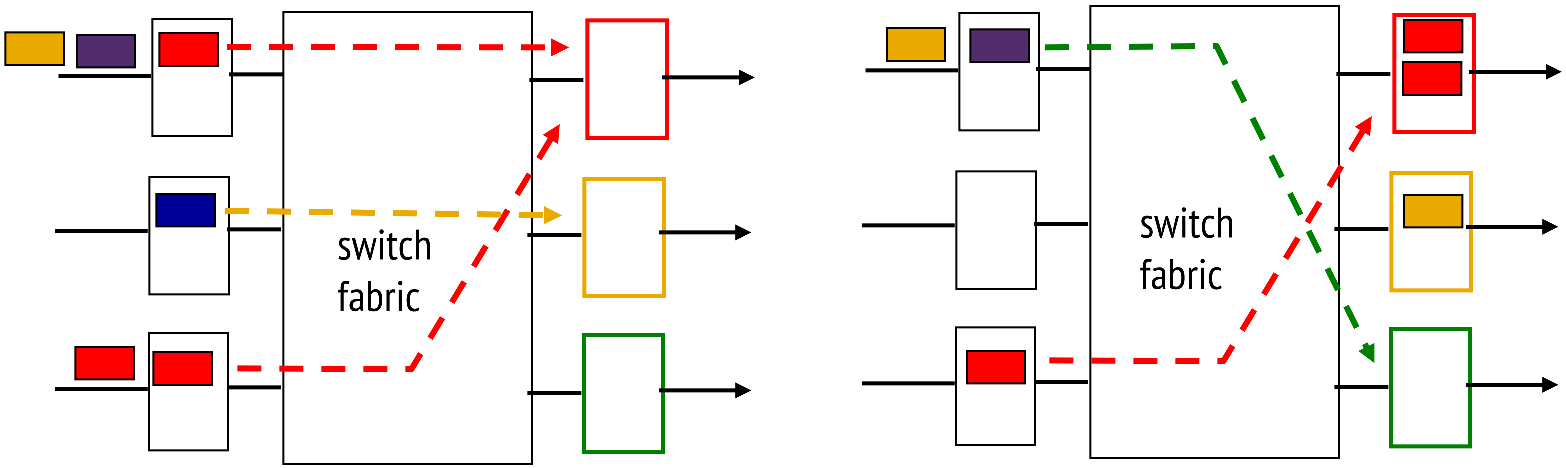

at *t*, packets more from input to output

one packet time later

- buffering when arrival rate via switch exceeds output line speed
- *queueing (delay) and loss due to output port buffer overflow!*

Scheduling mechanisms

- *scheduling:* choose next packet to send on link
- *FIFO (first in first out) scheduling:* send in order of arrival to queue
  - real-world example?
  - *discard policy:* if packet arrives to full queue: who to discard?
    - *tail drop:* drop arriving packet
    - *priority:* drop/remove on priority basis
    - *random:* drop/remove randomly

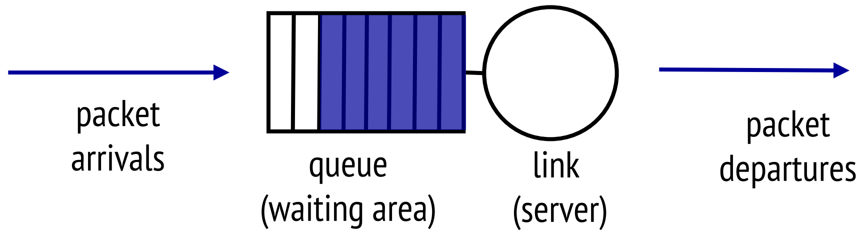

Scheduling policies: priority

priority scheduling: send highest priority queued packet

- multiple *classes*, with different priorities
  - class may depend on marking or other header info, e.g. IP source/dest, port numbers, etc.
  - real world example?

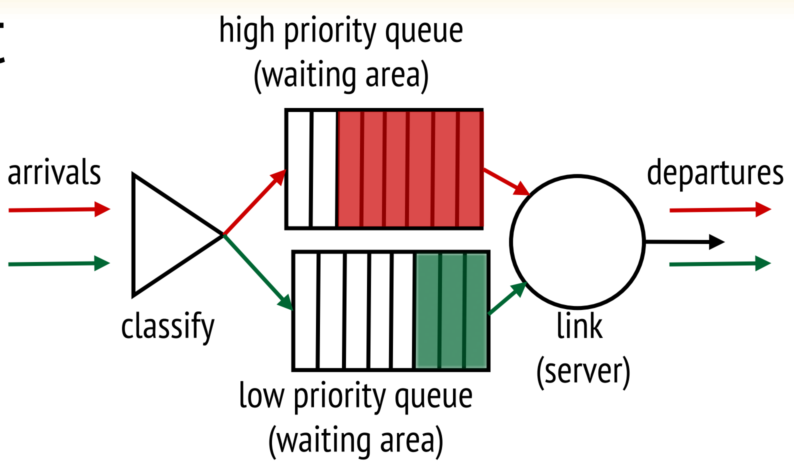

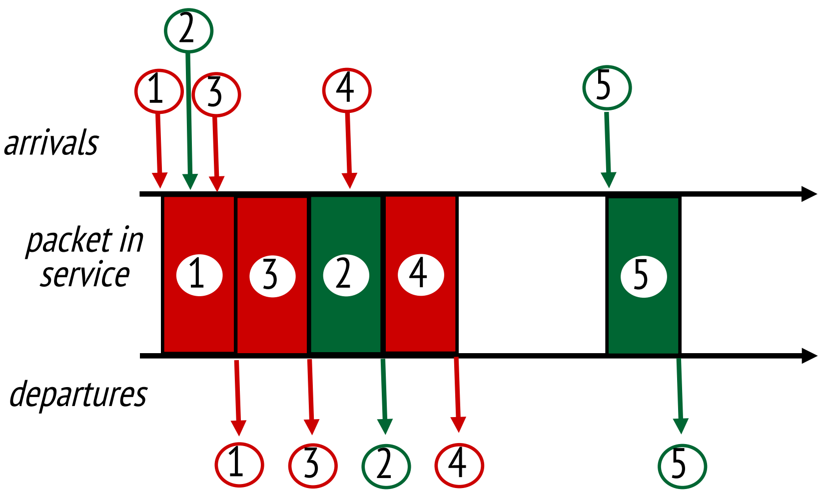

Scheduling policies: still more

Round Robin (RR) scheduling:

- multiple classes
- cyclically scan class queues, sending one complete packet from each class (if available)
- real world example?

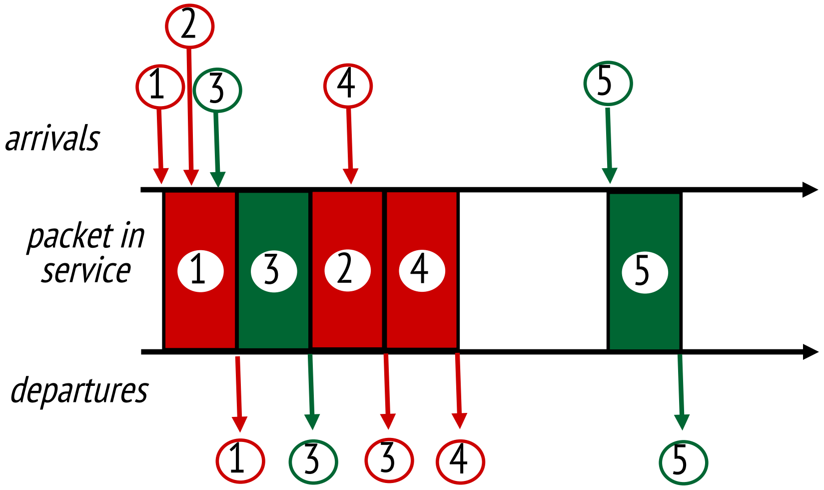

Scheduling policies: still more

Weighted Fair Queuing (WFQ):

- generalized Round Robin
- each class gets weighted amount of service in each cycle
- real-world example?

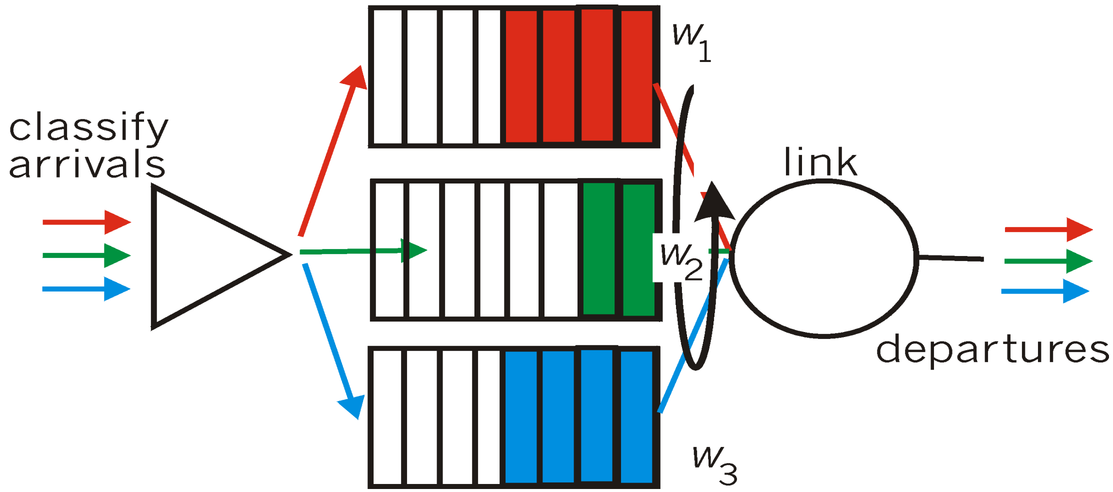

IP fragmentation, reassembly

example:

- 4000 byte datagram
- MTU = 1500 bytes

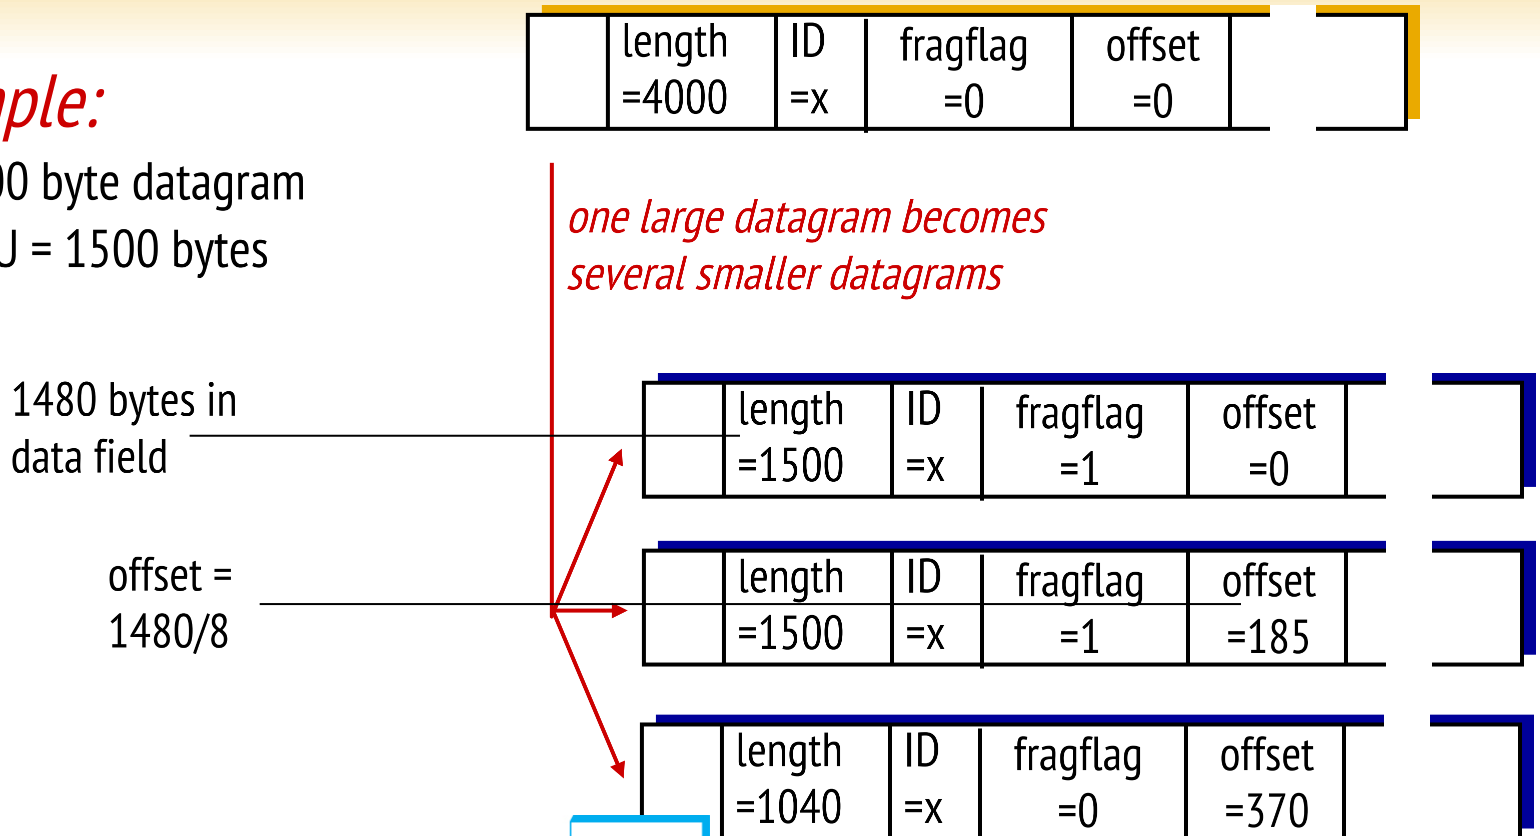

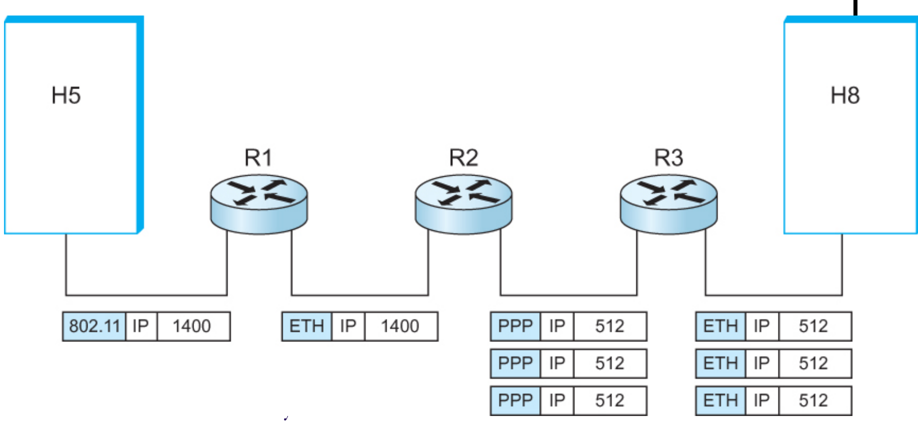

IP addressing: introduction

- *IP address:* 32-bit identifier for host, router *interface*
- *interface:* connection between host/router and physical link
  - router's typically have multiple interfaces
  - host typically has one or two interfaces (e.g., wired Ethernet, wireless 802.11)
- *IP addresses associated with each interface*

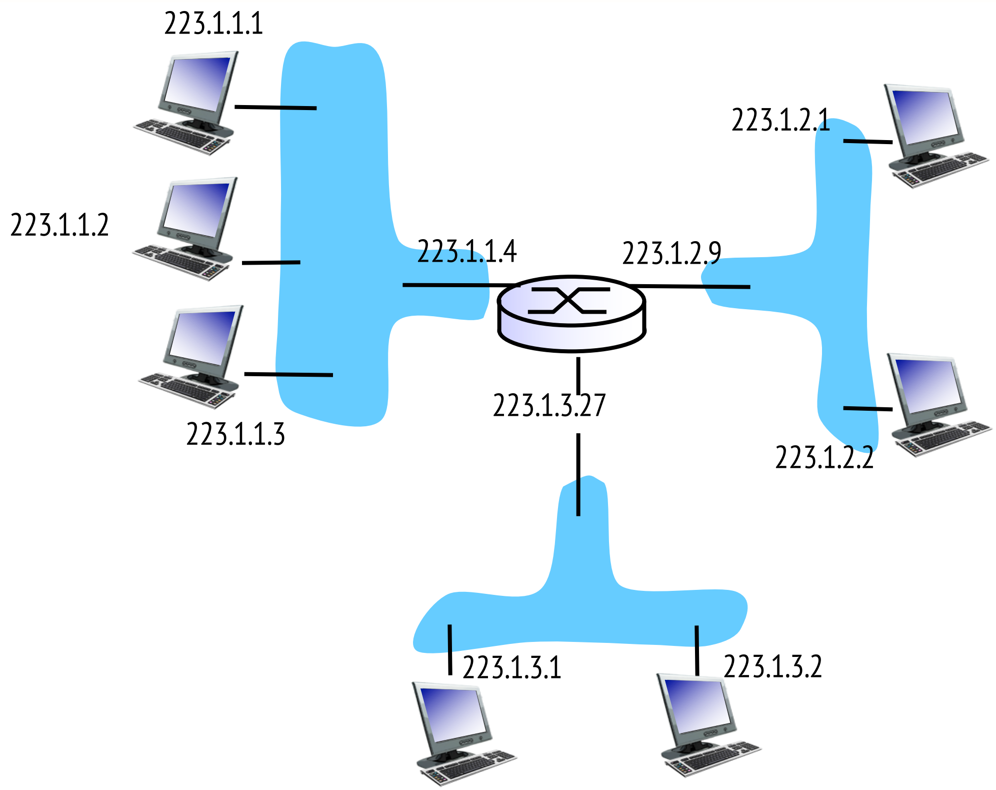

223.1.1.1 =	11011111	00000001	00000001	00000001
	223	1	1	1

IP addressing: CIDR

CIDR: Classless InterDomain Routing

- subnet portion of address of arbitrary length
- address format: a.b.c.d/x, where x is # bits in subnet portion of address

subnet part ← → | host part ← →

11001000 00010111 0001000 0 00000000

200.23.16.0/23

DHCP client-server scenario

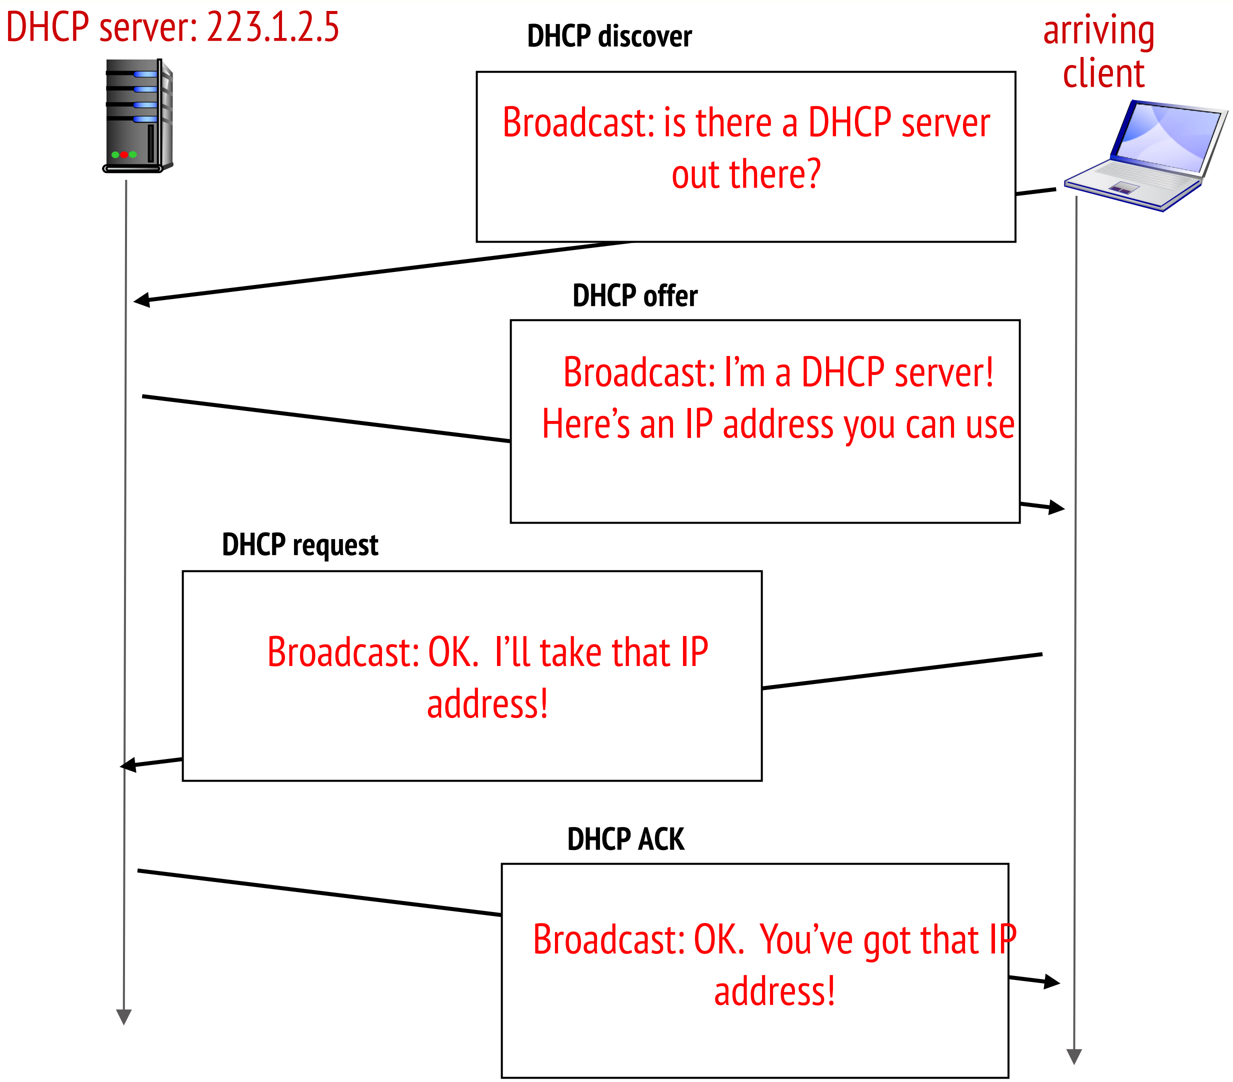

NAT: network address translation

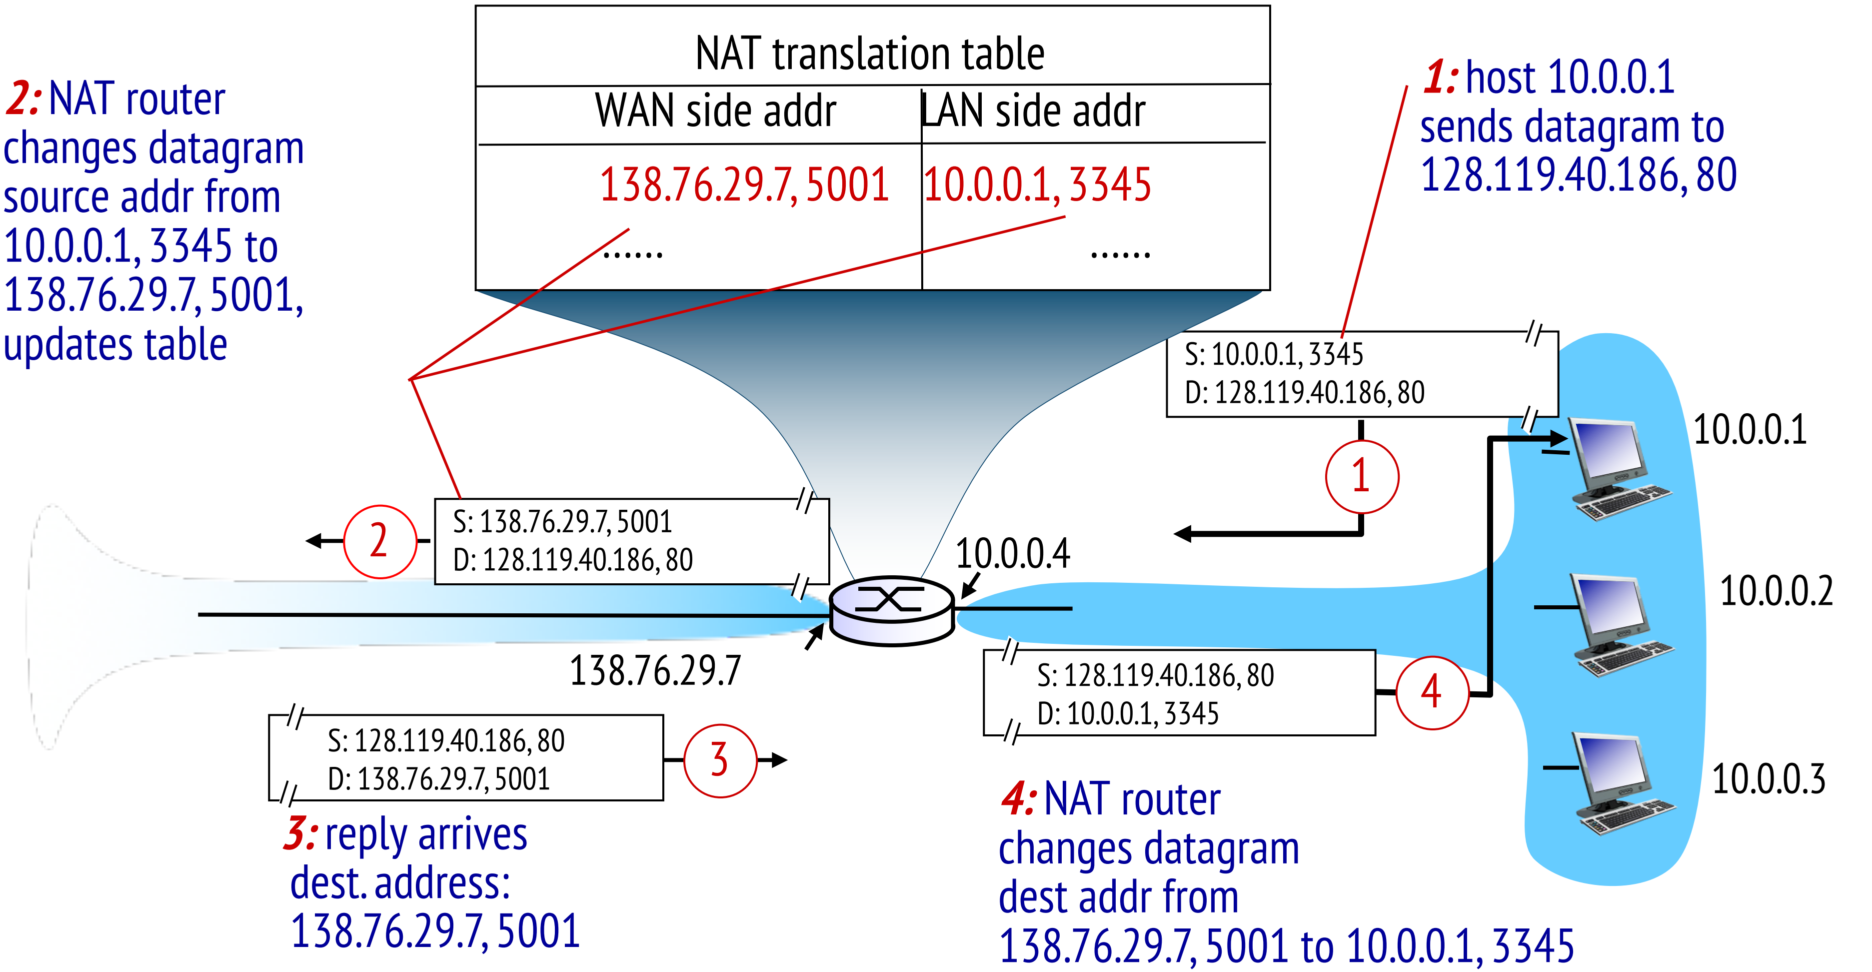

Internet Control Message Protocol (ICMP)

- Defines a collection of error messages that are sent back to the source host whenever a router or host is unable to process an IP datagram successfully
  - Destination host unreachable due to link /node failure
  - Reassembly process failed
  - TTL had reached 0 (so datagrams don't cycle forever)
  - IP header checksum failed

- ICMP-Redirect
  - From router to a source host
  - With a better route information

Traceroute and ICMP

- source sends series of UDP segments to destination
  - first set has TTL =1
  - second set has TTL=2, etc.
  - unlikely port number
- when datagram in *n*th set arrives to nth router:
  - router discards datagram and sends source ICMP message (type 11, code 0)
  - ICMP message include name of router & IP address

when ICMP message arrives, source records RTTs

stopping criteria:

- UDP segment eventually arrives at destination host
- destination returns ICMP "port unreachable" message (type 3, code 3)
- source stops

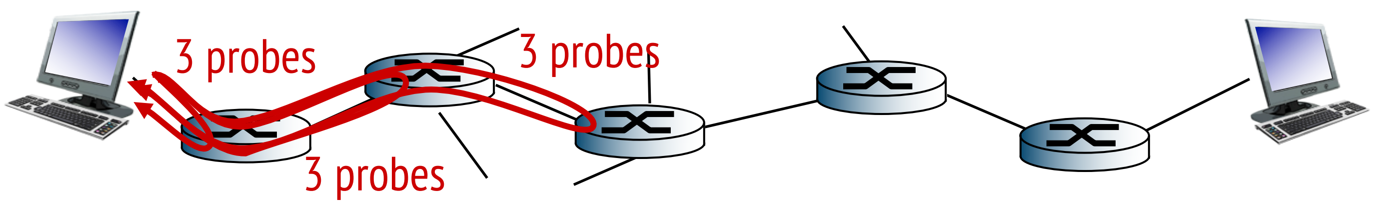

Dijkstra's algorithm: example

Step	N'	D(v) p(v)	D(w) p(w)	D(x) p(x)	D(y) p(y)	D(z) p(z)
0	u	7,u	3,u	5,u	∞	∞
1	uw	6,w		5,u	11,w	∞
2	uwx	6,w			11,w	14,x
3	uwxv				10,v	14,x
4	uwxvy					12,y
5	uwxvyz					

notes:

- construct shortest path tree by tracing predecessor nodes
- ties can exist (can be broken arbitrarily)

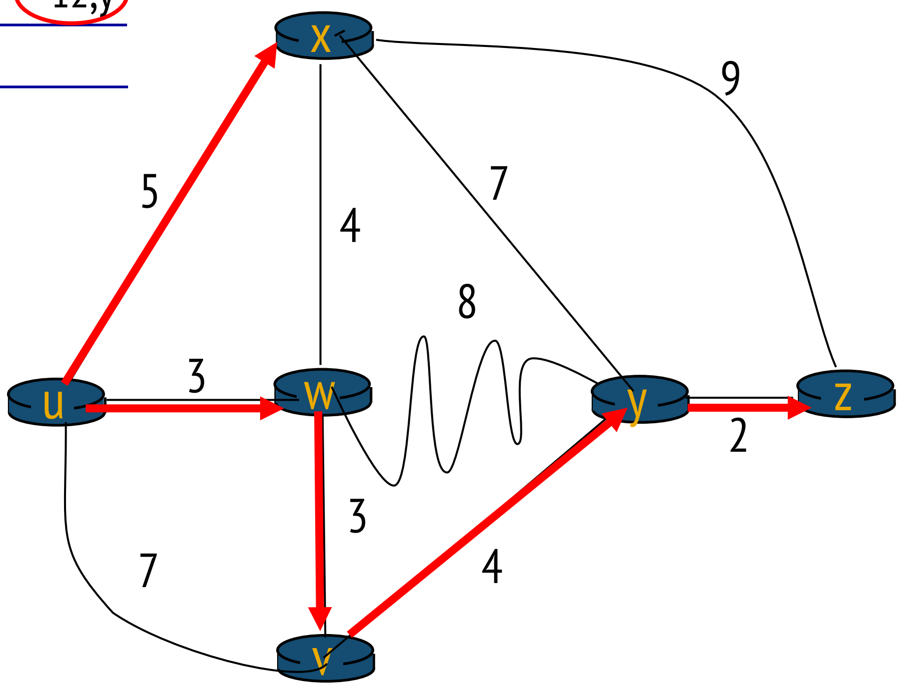

$$D_x(y) = \min\{c(x,y) + D_y(y), c(x,z) + D_z(y)\} = \min\{2+0\ , 7+1\} = 2$$

$$D_x(z) = \min\{c(x,y) + D_y(z), c(x,z) + D_z(z)\} = \min\{2+1\ , 7+0\} = 3$$

node x table

from \ cost to	x	y	z
x	0	2	7
y	∞	∞	∞
z	∞	∞	∞

node y table

from \ cost to	x	y	z
x	∞	∞	∞
y	2	0	1
z	∞	∞	∞

node z table

from \ cost to	x	y	z
x	∞	∞	∞
y	∞	∞	∞
z	7	1	0

from \ cost to	x	y	z
x	0	2	3
y	2	0	1
z	7	1	0

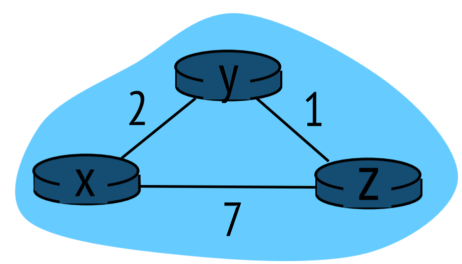

time

Comparison of LS and DV algorithms

message complexity

- ***LS:*** with n nodes, E links, O(nE) msgs sent
- ***DV:*** exchange between neighbors only
  - convergence time varies

speed of convergence

- ***LS:*** $O(n^2)$ algorithm requires O(nE) msgs
  - may have oscillations
- ***DV:*** convergence time varies
  - may be routing loops
  - count-to-infinity problem

robustness: what happens if router malfunctions?

LS:

- node can advertise incorrect *link* cost
- each node computes only its *own* table

DV:

- DV node can advertise incorrect *path* cost
- each node's table used by others
  - error propagates thru network

What's unique about MANET?

- Moving nodes → ever changing topology
- Wireless links
  - → various and volatile link quality
- Pervasive (cheap) devices
  - → Power constraints
- Security
  - Confidentiality, other attacks

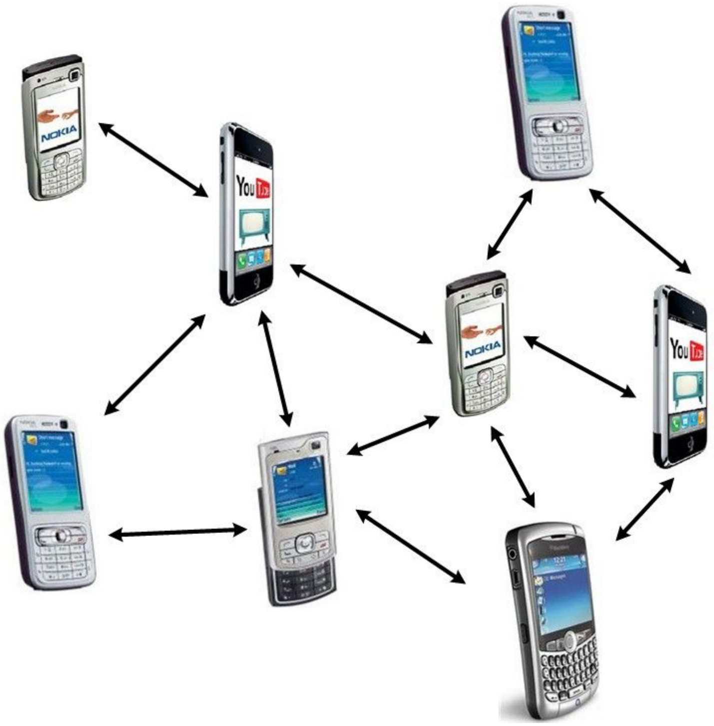

Routing Protocols

- Reactive (On-demand) protocols
  - Discover routes when needed
  - Source-initiated route discovery
- Proactive protocols
  - Traditional distributed shortest-path protocols
  - Based on periodic updates. High routing overhead
- Tradeoff
  - State maintenance traffic vs. route discovery traffic
  - Route via maintained route vs. delay for route discovery

Channel partitioning MAC protocols: TDMA

TDMA: time division multiple access

- access to channel in "rounds"
- each station gets fixed length slot (length = packet transmission time) in each round
- unused slots go idle
- example: 6-station LAN, 1,3,4 have packets to send, slots 2,5,6 idle

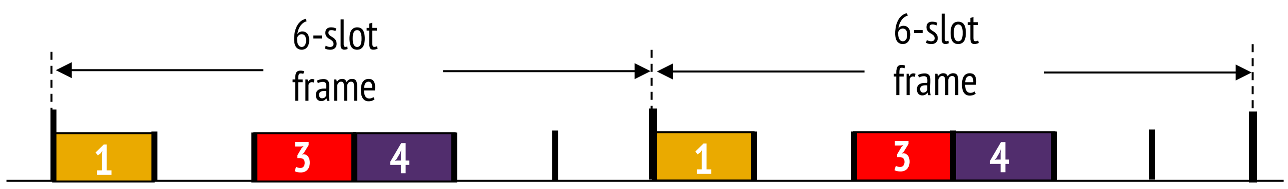

Channel partitioning MAC protocols: FDMA

FDMA: frequency division multiple access

- channel spectrum divided into frequency bands
- each station assigned fixed frequency band
- unused transmission time in frequency bands go idle
- example: 6-station LAN, 1,3,4 have packet to send, frequency bands 2,5,6 idle

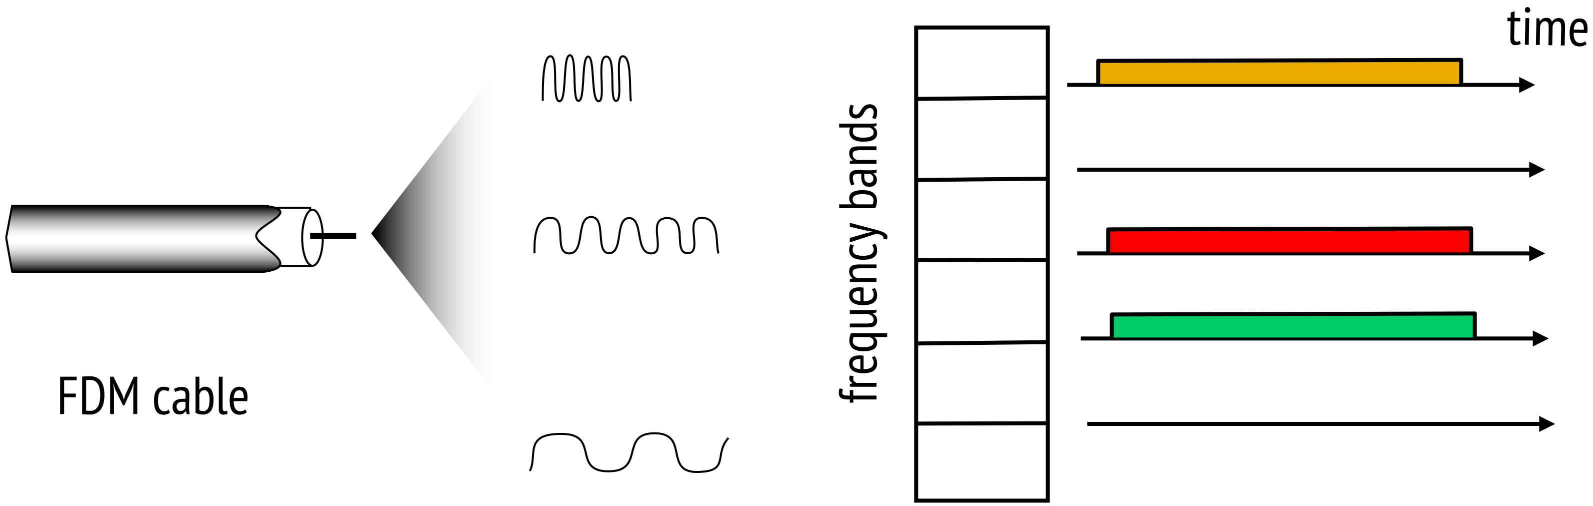

CRC Example

- want:
  - D.2^r XOR R = nG
- equivalently:
  - D.2^r = nG XOR R
- equivalently:
  - if we divide D.2r by G, we want remainder R to satisfy:

$$R = remainder \frac{D.2^r}{G}$$

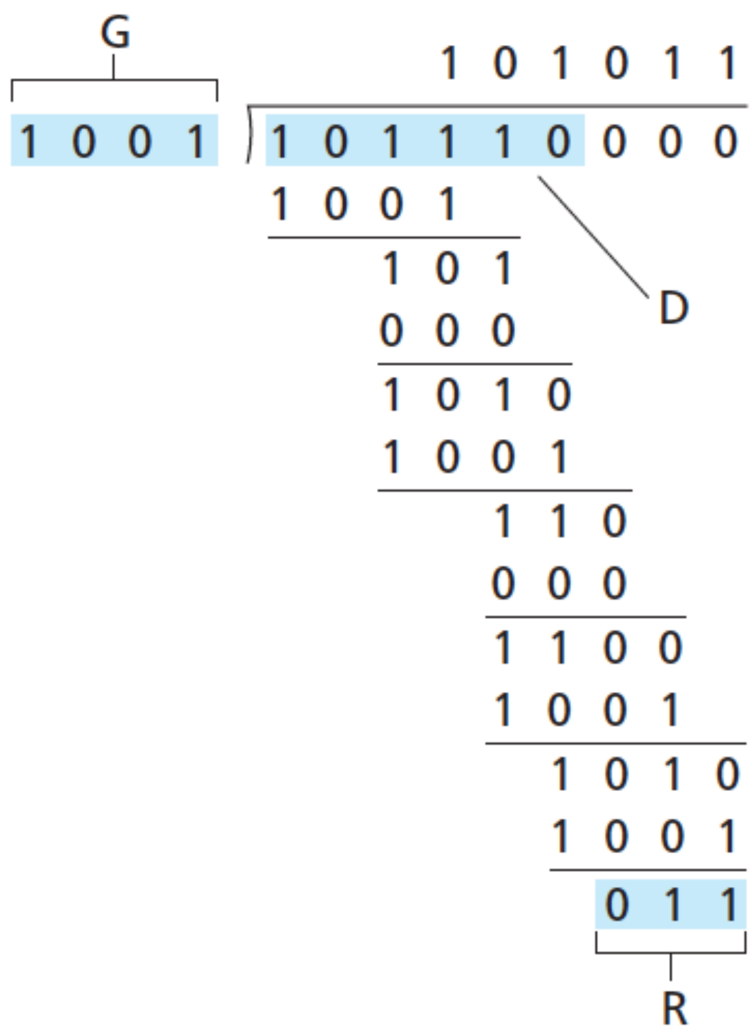

CSMA/CD (collision detection)

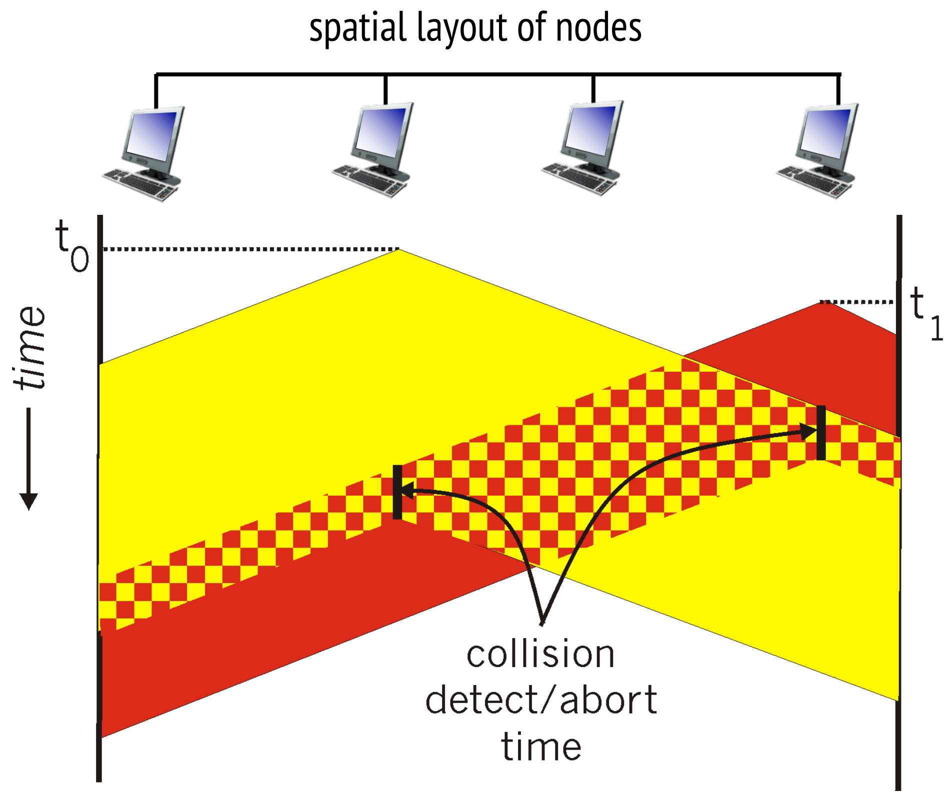

Ethernet CSMA/CD algorithm

1. NIC receives datagram from network layer, creates frame
2. If NIC senses channel idle, starts frame transmission. If NIC senses channel busy, waits until channel idle, then transmits.
3. If NIC transmits entire frame without detecting another transmission, NIC is done with frame !
4. If NIC detects another transmission while transmitting, aborts and sends jam signal
5. After aborting, NIC enters *binary (exponential) backoff:*
   - after *m*th collision, NIC chooses *K* at random from *{0,1,2, …, 2^m-1}*. NIC waits K·512 bit times, returns to Step 2
   - longer backoff interval with more collisions

Popular Interconnection Devices

	Hub	Switch	Router
Traffic Isolation	No	Yes	Yes
Plug and Play	Yes	Yes	No
Optimal Routing	No	No	Yes

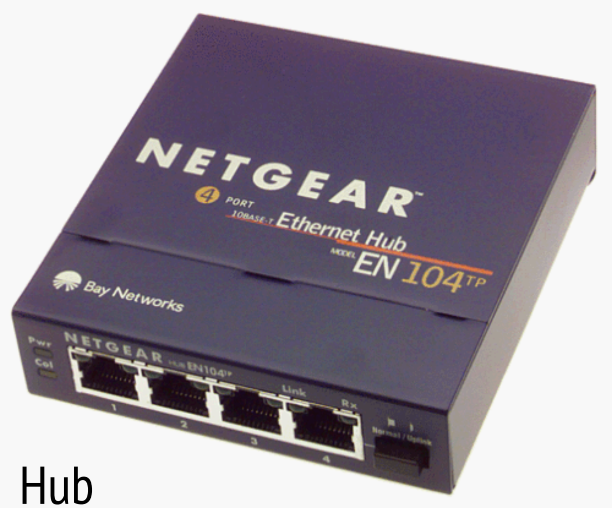

Hub

Switch

Router

Maximum Data Rate of a Channel

- Nyquist's theorem (1924) relates the data rate to the bandwidth (B) and number of signal levels (V):

  Max. data rate = $2B \log_2 V$ bits/sec

- Shannon's theorem (1948) relates the data rate to the bandwidth (B) and signal strength (S) relative to the noise (N):

  Max. data rate = $B \log_2(1 + S/N)$ bits/sec

- Signal to Noise Ratio:

  SNR = $10 \log_{10}(S/N)$ dB

  dB = decibels ➔ deci = 10; 'bel' chosen after Alexander Graham Bell

Baseband Transmission

- Line codes send <u>symbols</u> that represent one or more bits
  - NRZ is the simplest, literal line code (+1V="1", -1V="0")
  - Other codes tradeoff bandwidth and signal transitions

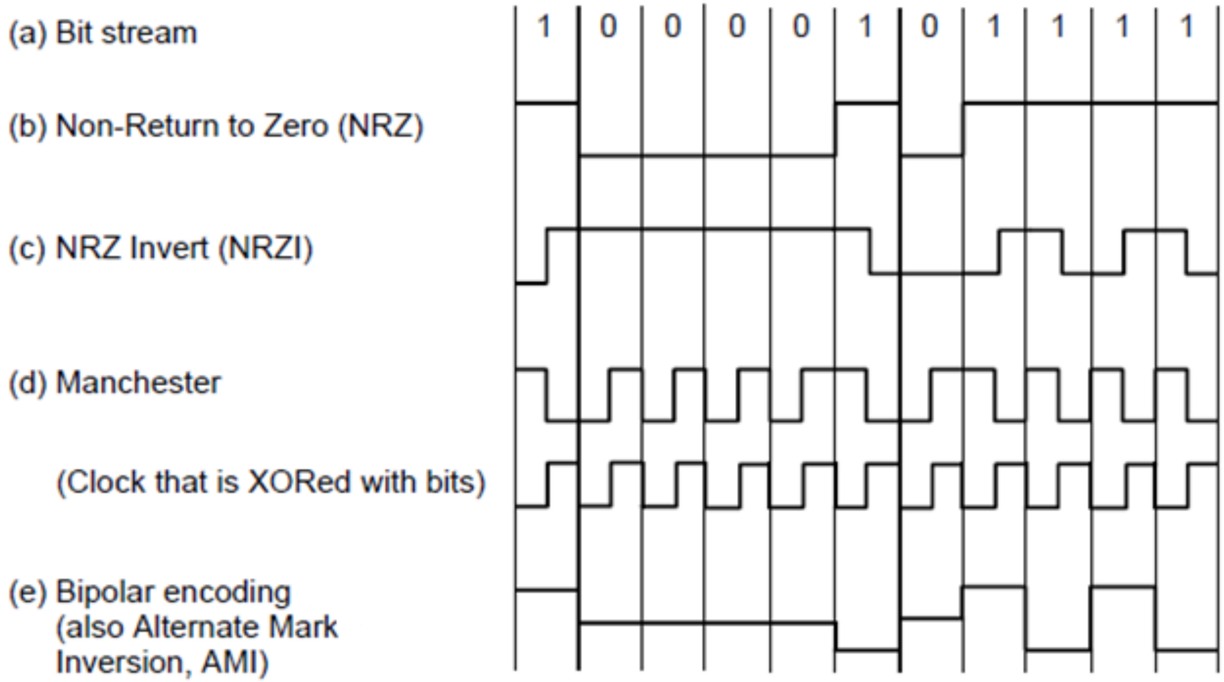

Four different line codes

Clock Recovery

- To decode the symbols, signals need sufficient transitions
  - Otherwise long runs of 0s (or 1s) are confusing, e.g.:

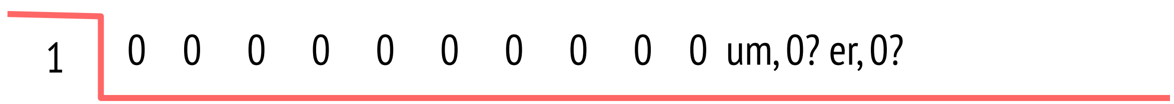

- Strategies:
  - Manchester coding, mixes clock signal in every symbol
  - 4B/5B maps 4 data bits to 5 coded bits with 1s and 0s:

Data	Code	Data	Code	Data	Code	Data	Code
0000	11110	0100	01010	1000	10010	1100	11010
0001	01001	0101	01011	1001	10011	1101	11011
0010	10100	0110	01110	1010	10110	1110	11100
0011	10101	0111	01111	1011	10111	1111	11101

  - Scrambler XORs tx/rx data with pseudorandom bits

Modulation

- *Constellation diagrams* are a shorthand to capture the amplitude and phase modulations of symbols:

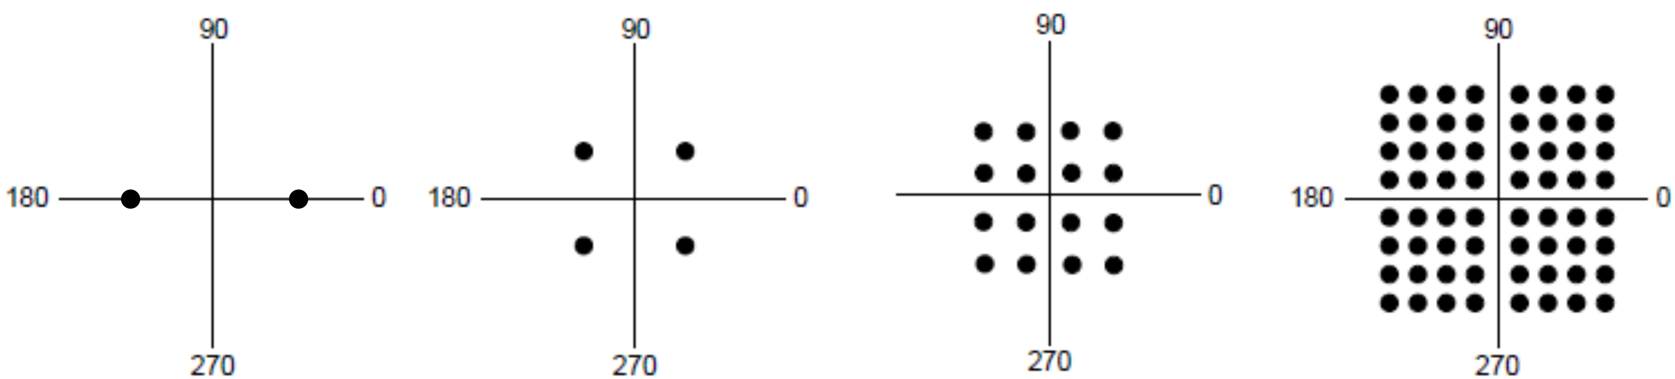

BPSK	QPSK	QAM-16	QAM-64
2 symbols	4 symbols	16 symbols	64 symbols
1 bit/symbol	2 bits/symbol	4 bits/symbol	6 bits/symbol

BPSK/QPSK varies only phase

QAM varies amplitude and phase

Gray Coding

- Gray-coding assigns bits to symbols so that small symbol errors cause few bit errors:

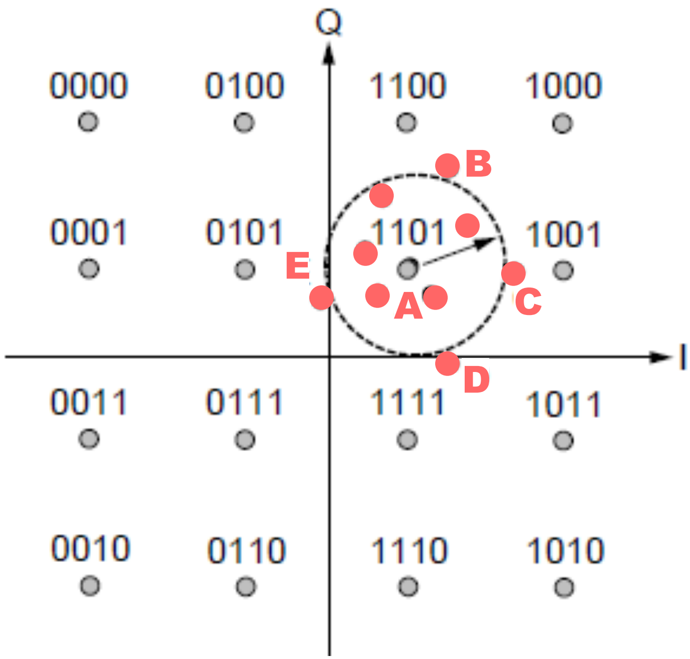

When 1101 is sent:

Point	Decodes as	Bit errors
A	1101	0
B	1100	1
C	1001	1
D	1111	1
E	0101	1

Code Division Multiple Access (CDMA)

- CDMA shares the channel by giving users a code
  - Codes are orthogonal; can be sent at the same time
  - Widely used as part of 3G networks
  - Gold code (GPS Signals), Walsh-Hadamard code, Zadoff-chu sequence

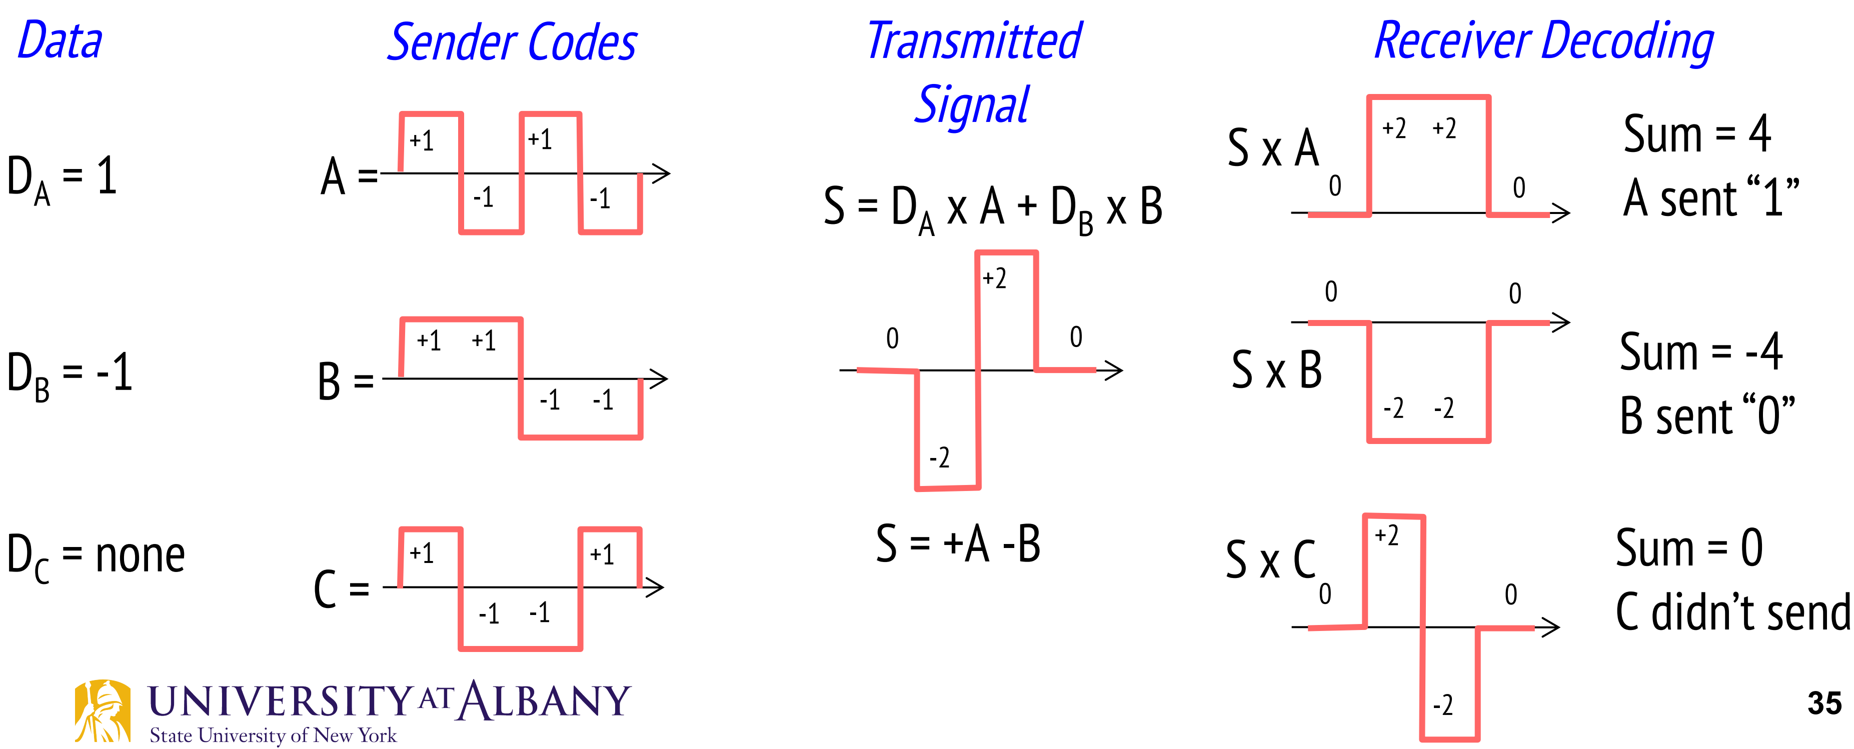

What is network security?

- *confidentiality:* only sender, intended receiver should "understand" message contents
  - Method – encrypt at sender, decrypt at receiver
  - A protocol that prevents an adversary from understanding the message contents is said to provide *confidentiality*.
  - Concealing the quantity or destination of communication is called *traffic confidentiality*.
- *message integrity:* sender, receiver want to ensure message not altered (in transit, or afterwards) without detection
  - A protocol that detects message tampering provides *data integrity.*
  - The adversary could alternatively transmit an extra copy of your message in a *replay attack.*
  - A protocol that detects message tampering provides *originality.*
  - A protocol that detects delaying tactics provides *timeliness.*

What is network security?

- *authentication:* sender, receiver want to confirm identity of each other
  - A protocol that ensures that you really are talking to whom you think you're talking is said to provide *authentication*.
  - Example: DNS Attack [correct URL gets converted to malicious IP]
- *access and availability*: services must be accessible and available to users
  - A protocol that ensures a degree of access is called *availability*.
  - Denial of Service (DoS) Attack
  - Example: SYN Flood attack (Client not transmitting 3rd message in TCP 3-way handshake, thus consuming server's resource)
  - Example: Ping Flood (attacker transmits ICMP Echo Request packets)

Simple encryption scheme

substitution cipher: substituting one thing for another

- *monoalphabetic* cipher: substitute one letter for another

plaintext: abcdefghijklmnopqrstuvwxyz

ciphertext: mnbvcxzasdfghjklpoiuytrewq

e.g.:

Plaintext: bob. i love you. alice

ciphertext: nkn. s gktc wky. mgsbc

Encryption key: mapping from set of 26 letters to set of 26 letters

Polyalphabetic Cipher

Plaintext letter:	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v	w	x	y	z
$C_1(k = 5)$:	f	g	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v	w	x	y	z	a	b	c	d	e
$C_2(k = 19)$:	t	u	v	w	x	y	z	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q	r	s

- n substitution ciphers, $C_1, C_2, \ldots, C_n$
- cycling pattern:
  - e.g., n=4 [C_1-C_4], k=key length=5: C_1,C_3,C_4,C_3,C_2; C_1,C_3,C_4,C_3,C_2; ..
- for each new plaintext symbol, use subsequent substitution pattern in cyclic pattern
  - dog: d from C_1, o from C_3, g from C_4

  Encryption key: n substitution ciphers, and cyclic pattern

  - key need not be just n-bit pattern

Good Luck!!!

Please provide your feedback in online course evaluation.