ELEN 4017

Network Fundamentals

Lecture 22

Purpose of lecture

Chapter 4: Network Layer

- Internet Protocol
  - Format and Fragmentation

IP datagram format

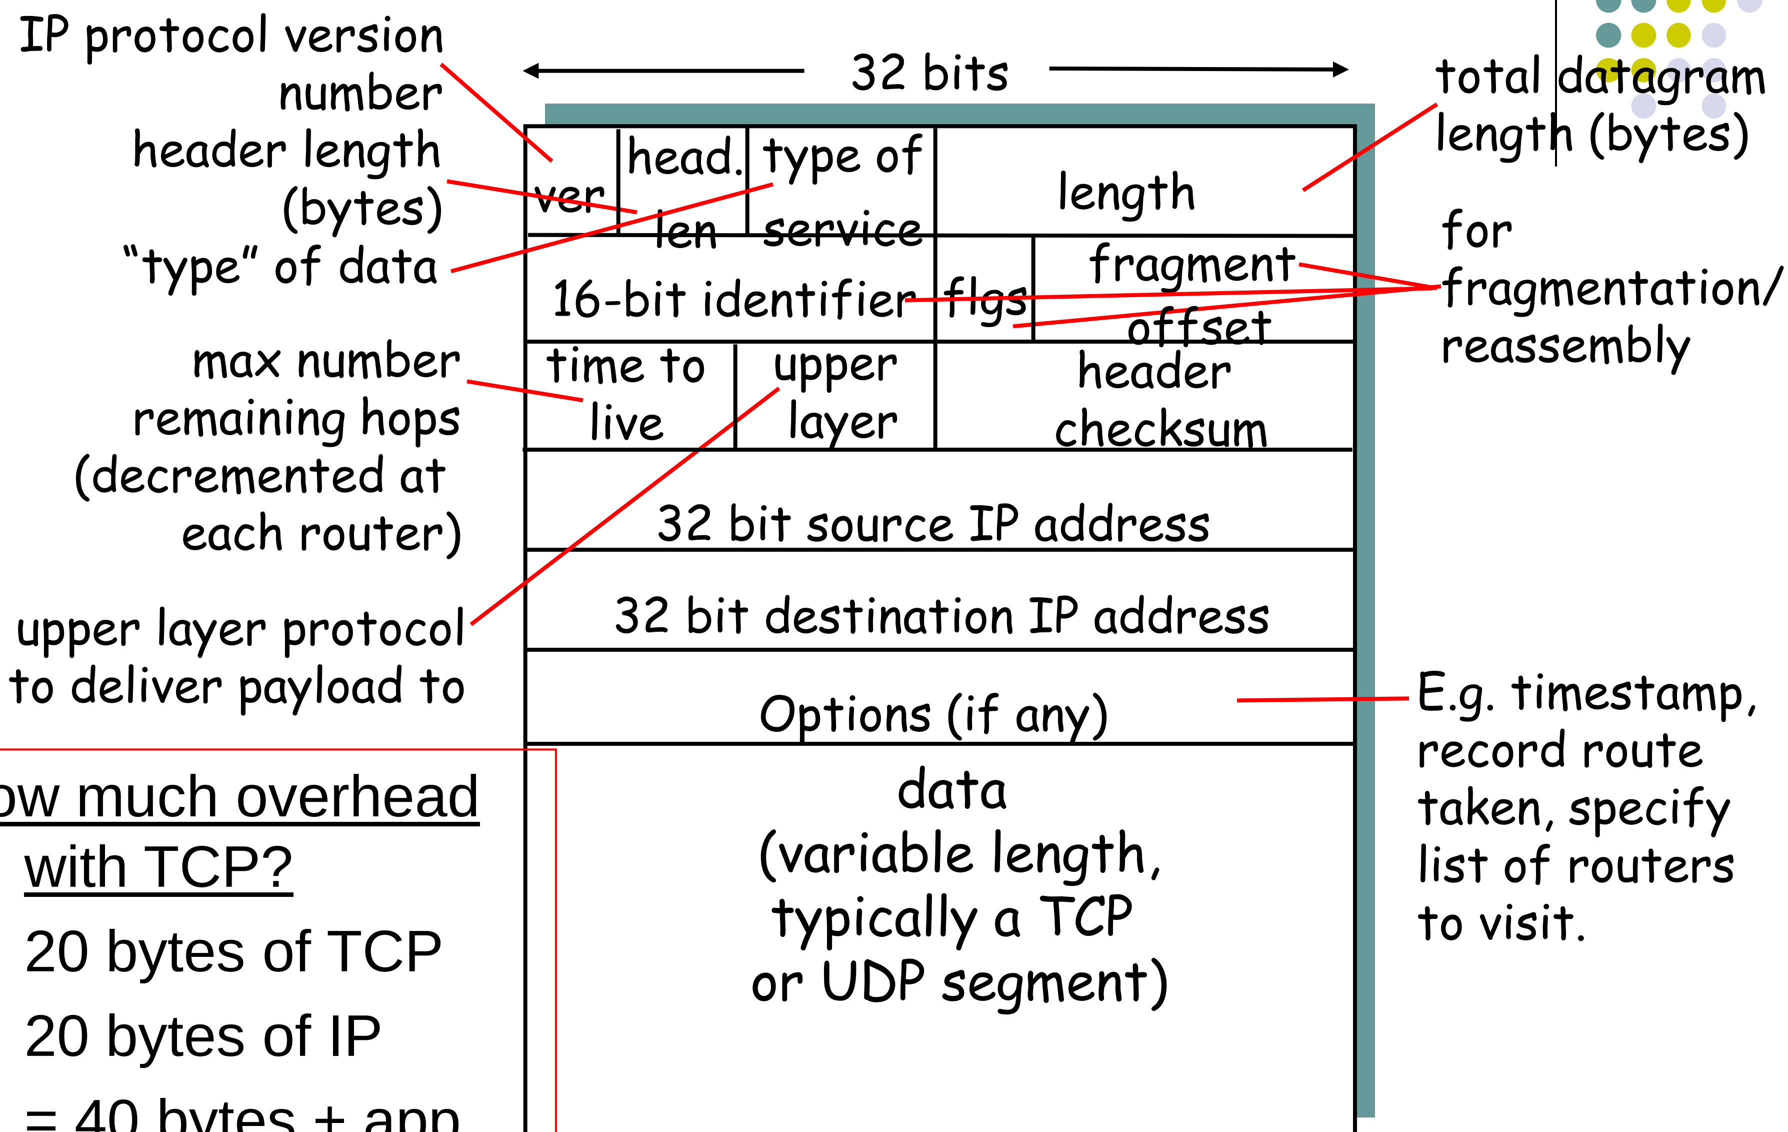

<u>how much overhead with TCP?</u>

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

IP Fragmentation & Reassembly

- network links have MTU (max.transfer size) - largest possible link-level frame.
  - different link types, different MTUs
- large IP datagram divided ("fragmented") within net
  - one datagram becomes several datagrams
  - "reassembled" only at final destination
  - IP header bits used to identify, order related fragments

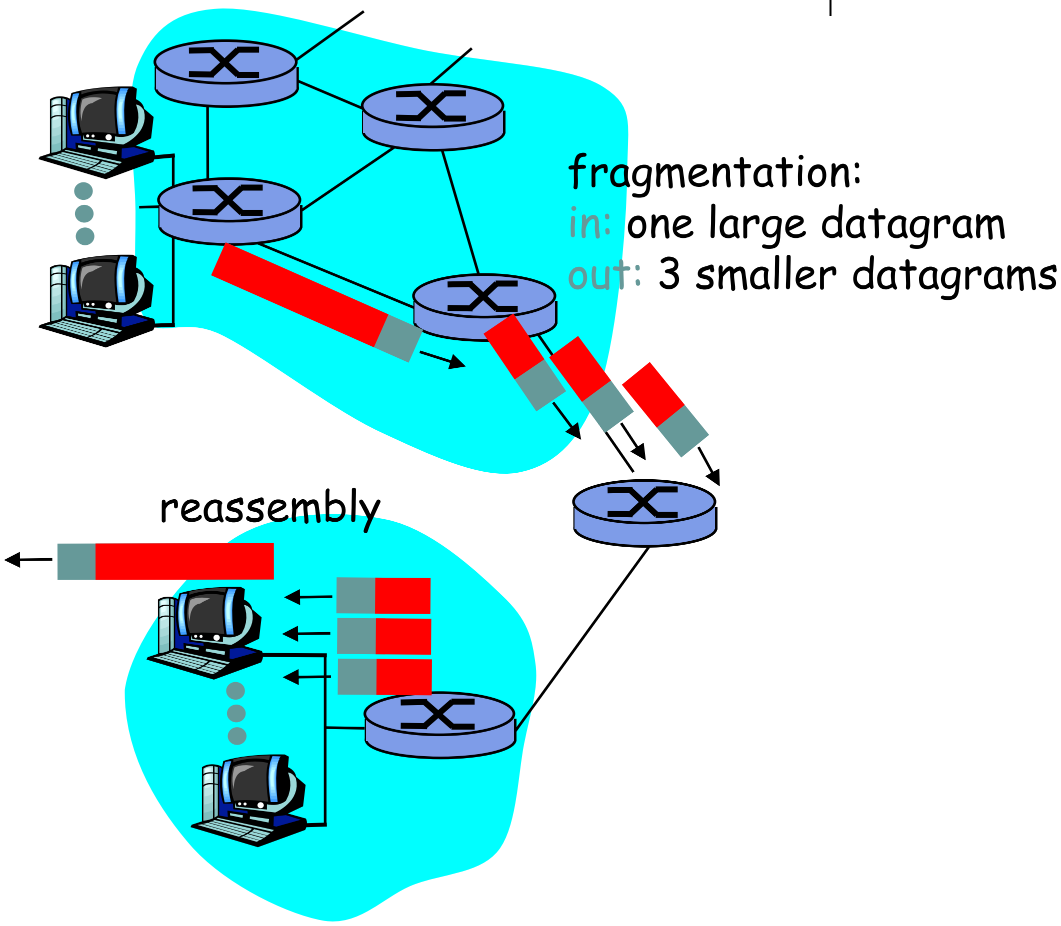

IP Fragmentation and Reassembly

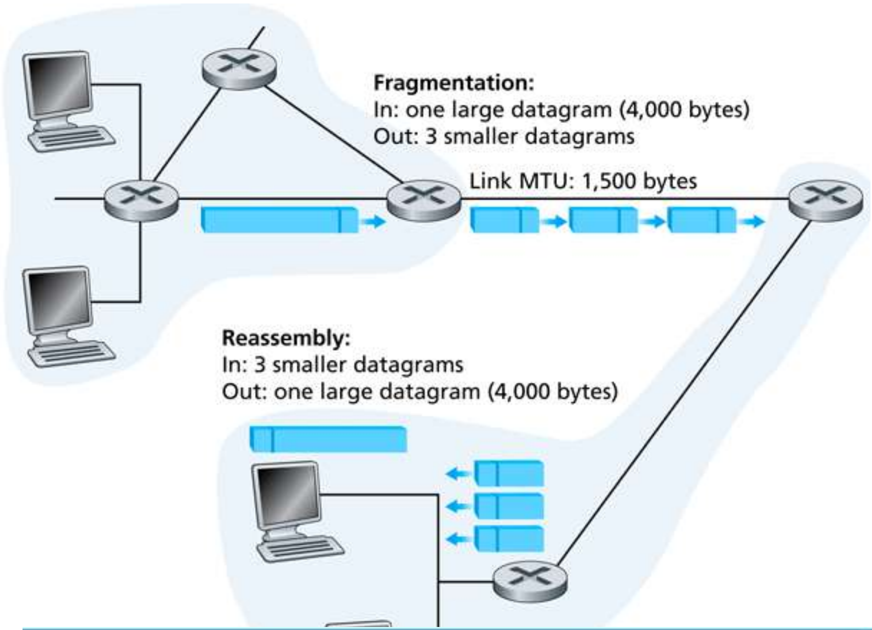

Re-assembly only at end-system

8 byte chunks

Fragment	Bytes	ID	Offset	Flag
1st fragment	1,480 bytes in the data field of the IP datagram	identification = 777	offset = 0 (meaning the data should be inserted beginning at byte 0)	flag = 1 (meaning there is more)
2nd fragment	1,480 bytes of data	identification = 777	offset = 185 (meaning the data should be inserted beginning at byte 1,480. Note that 185 · 8 = 1,480)	flag = 1 (meaning there is more)
3rd fragment	1,020 bytes (= 3,980–1,480–1,480) of data	identification = 777	offset = 370 (meaning the data should be inserted beginning at byte 2,960. Note that 370 · 8 = 2,960)	flag = 0 (meaning this is the last fragment)

Fragmentation vulnerabilities

- Fragmentation places extra burden on routers and end systems.
- Can be used for DoS attacks – send a series of unexpected and bizarre fragments.
- New version of IP (IPv6) does away with fragmentation.

Purpose of lecture

Chapter 4: Network Layer

- Internet Protocol
  - Format and Fragmentation
  - IP addressing

Overview of addressing

- IPv4 is the currently used IP addressing scheme.
- IPv6 is the proposal to replace IPv4, since address space is filling up quickly. Why ?
- IPv6 will increase the IP address space, optimize forwarding.
- Important issue is how do you make the transition whilst maintaining service:
  - Dual stack devices
  - Tunnelling

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 interface
  - IP addresses associated with each interface

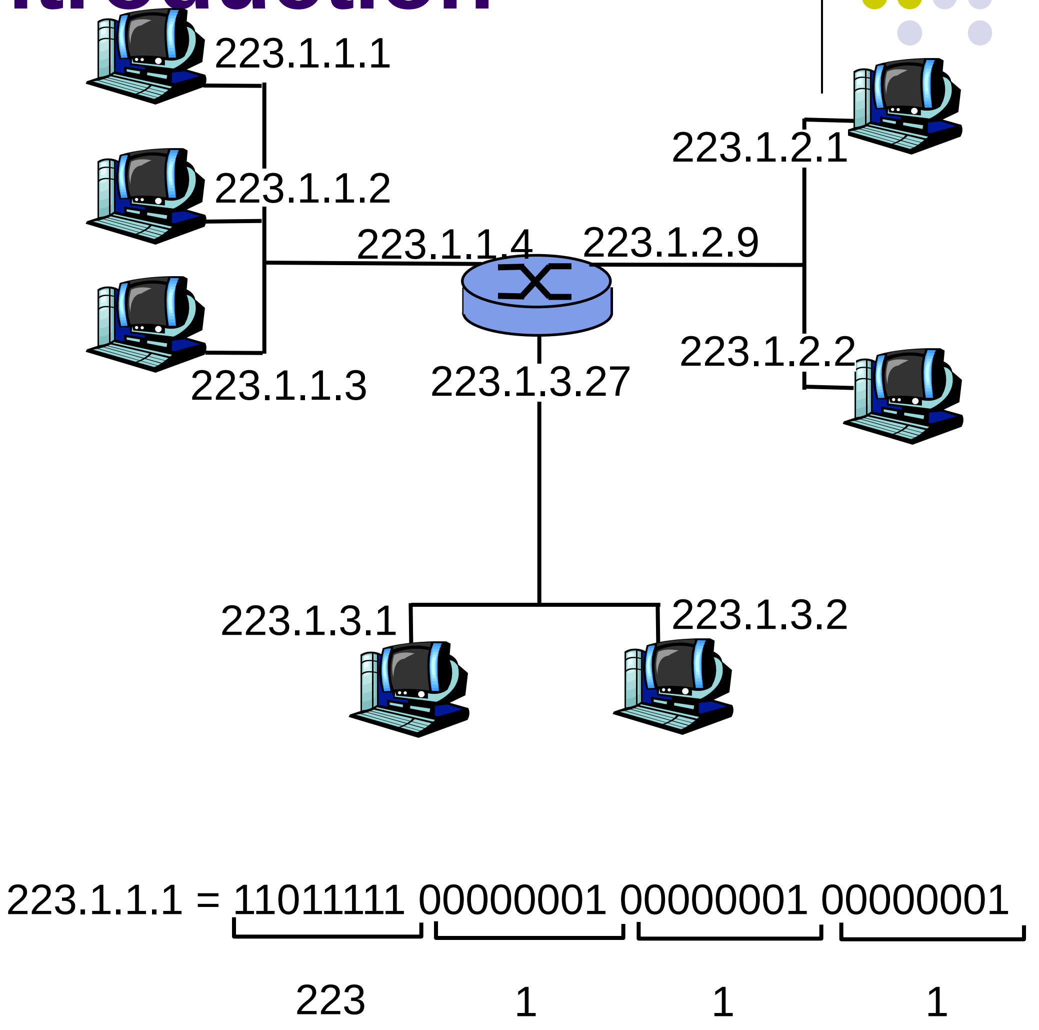

Internet addressing (IPv4)

- An IPv4 address is 32 bits long (~ 4 billion addresses available).
- Stated in **dotted decimal notation**, each **byte** is written in decimal form and separated by a dot e.g. 193.32.216.9. The binary equivalent is
  - 11000001 00100000 11011000 00001001
- Each interface on every host or router in the global Internet must have a **globally unique address.**
- Addresses are not chosen freely but determined by the **subnet** to which interface is connected.

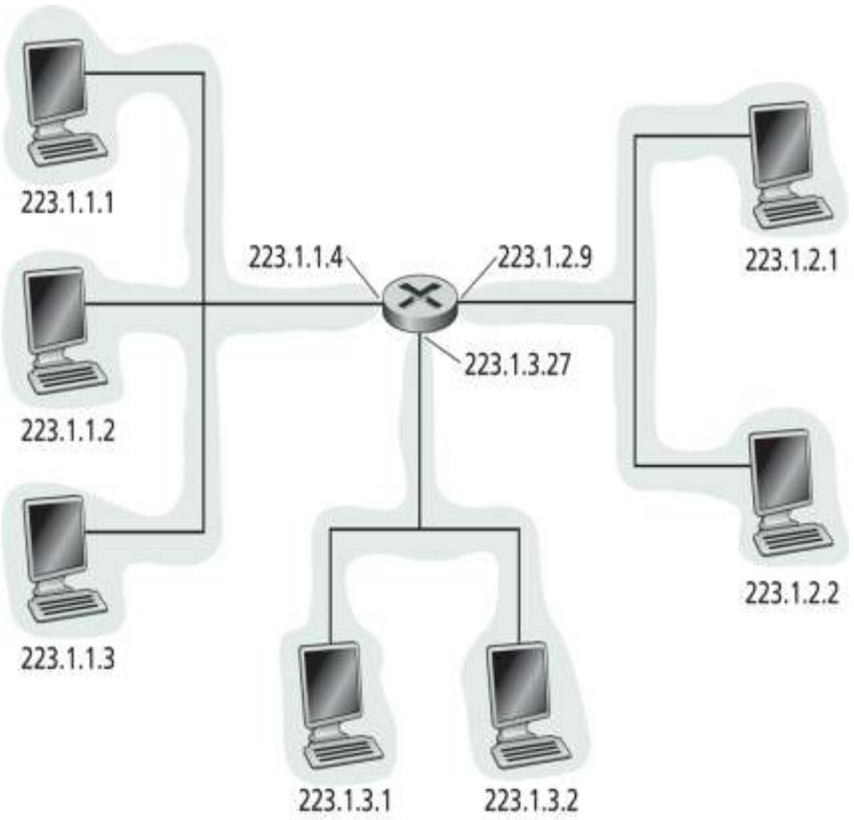

Figure 4.15 ♦ Interface addresses and subnets

Subnets

- IP address:
  - subnet part (high order bits)
  - host part (low order bits)
- *What's a subnet ?*
  - device interfaces with same subnet part of IP address
  - can physically reach each other without intervening router

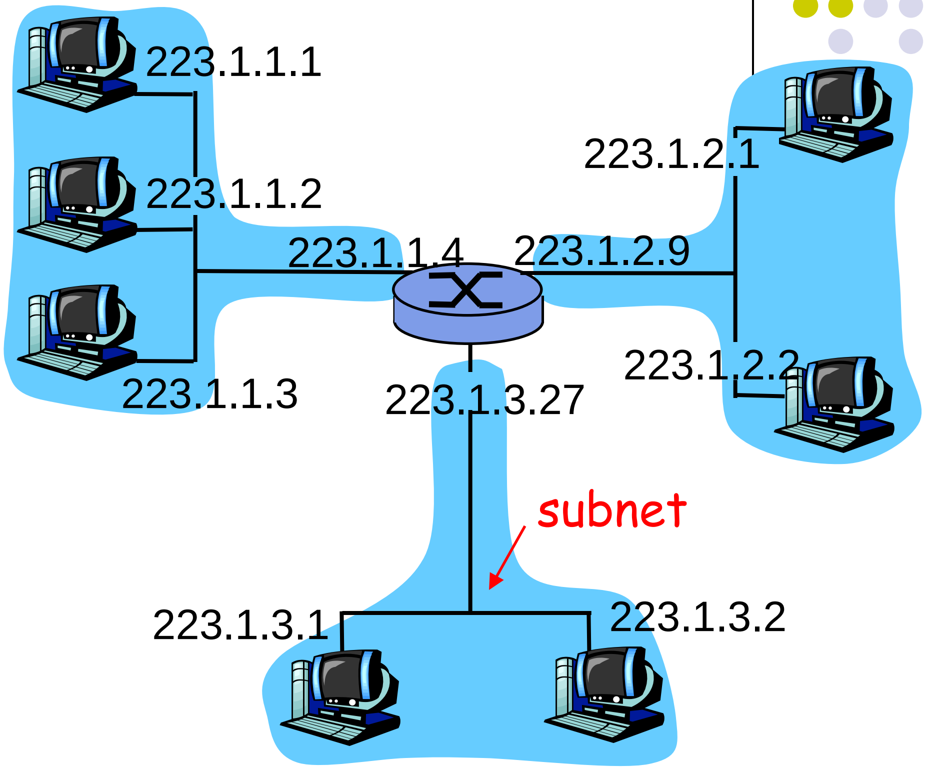

network consisting of 3 subnets

Subnets

Recipe

- To determine the subnets, detach each interface from its host or router, creating islands of isolated networks. Each isolated network is called a subnet.

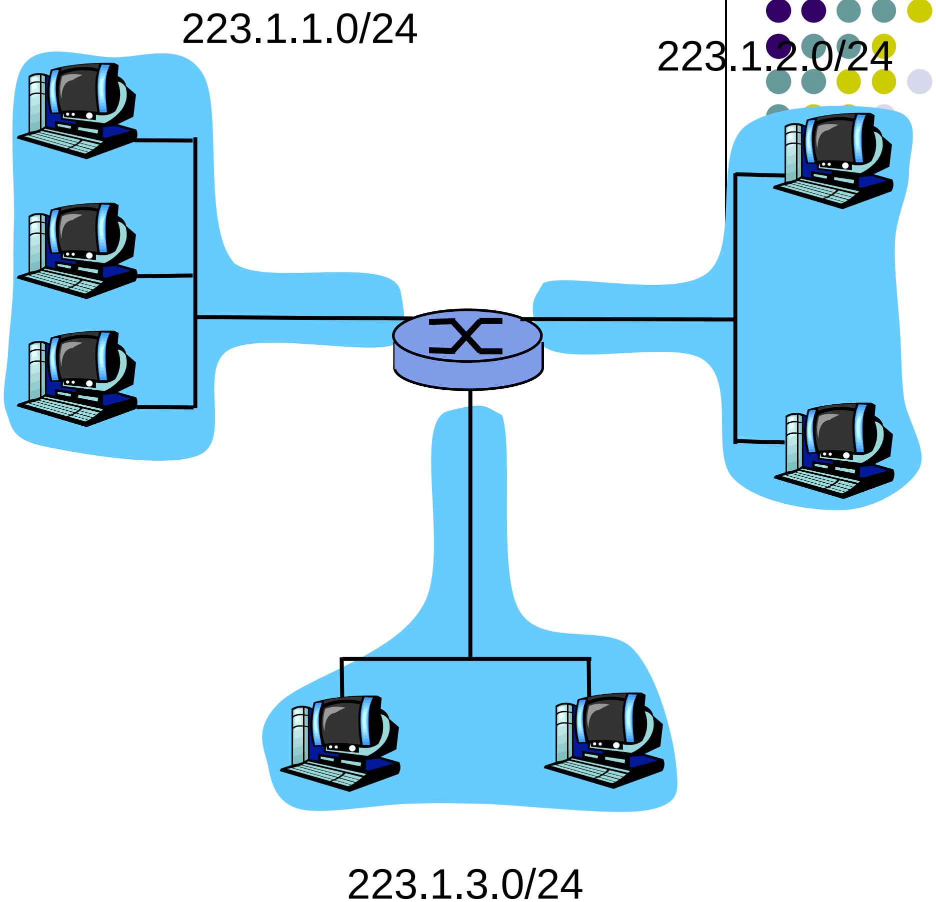

Subnet mask: /24

Subnets

How many?

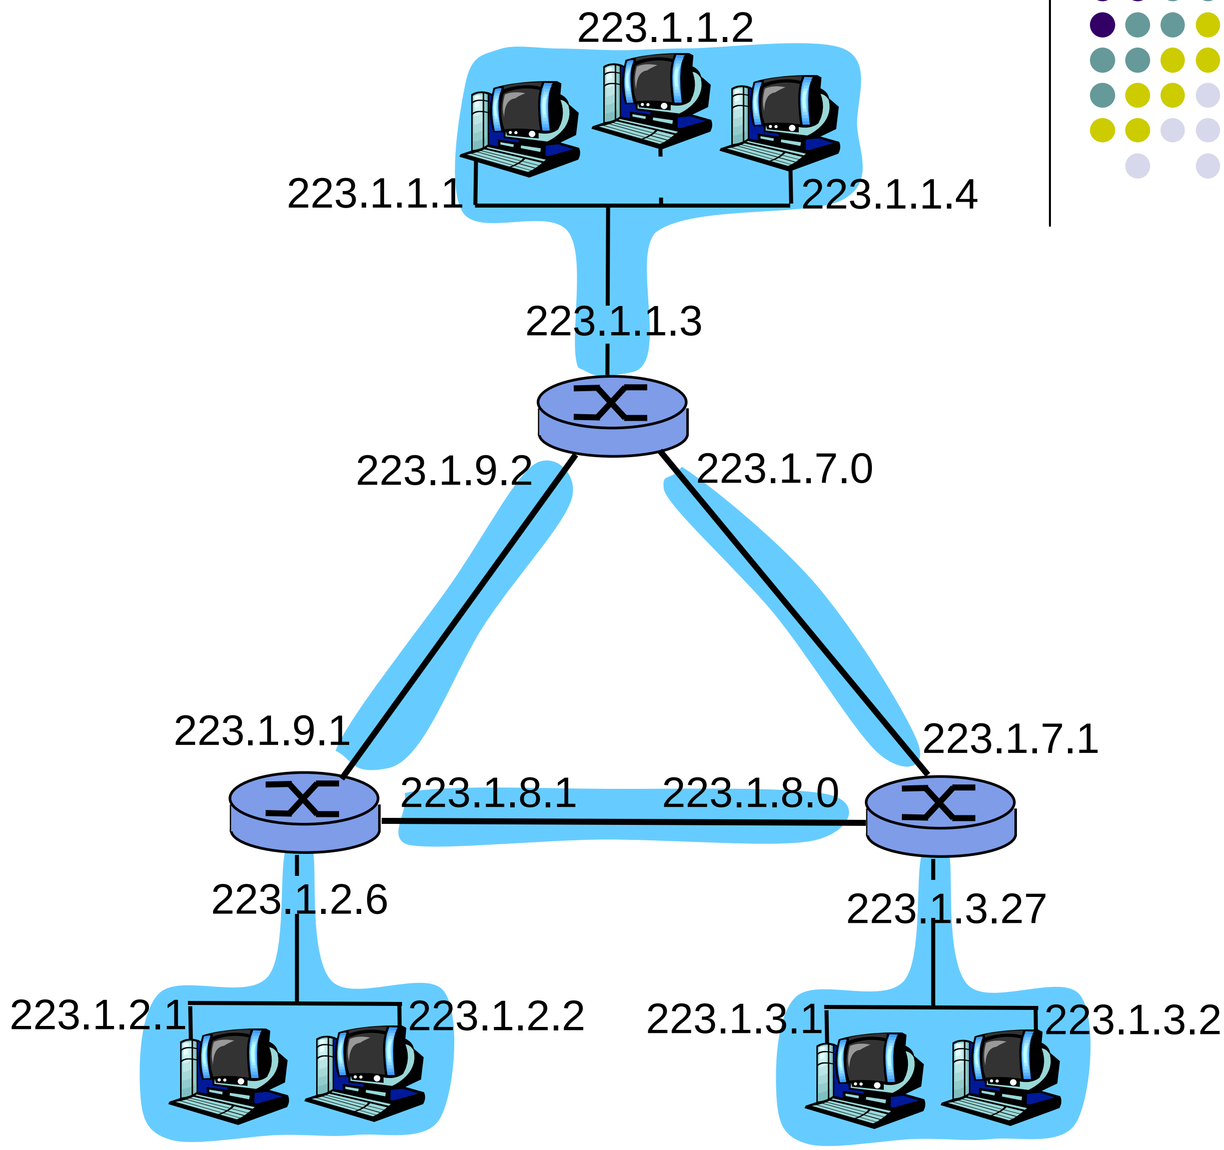

Address assignment - CIDR

- **C**lassless **I**nter-**D**omain **R**outing (CIDR)
- 32 bit IP address is divided into parts
  - a.b.c.d/x where x indicates number of bits in first part of address.
  - **x most significant bits** indicate the network part (called prefix).
- Suppose that Wits has been assigned the prefix **a.b.c.d/21**. Within Wits each school could be assigned a lower level subnet. Thus EIE could have the mask **a.b.c.d/24.**
- Before CIDR, **classful addressing** was used, for which the network portion of an IP address were constrained to be in multiples of bytes.
  - Class A had prefix of 8 bits
  - Class B had prefix of 16 bits
  - Class C had prefix of 24 bits
- Why do you think this proved to be unsuitable?
- Within a subnet **2 addresses are reserved**:
  - a.b.c.**0** is the address of the subnet
  - a.b.c.**255** is the broadcast address for that subnet.

IP addresses: how to get one?

Q: How does a *host* get IP address?

- hard-coded by system admin in a file
  - Windows: control-panel->network->configuration->tcp/ip->properties
  - UNIX: /etc/rc.config
- DHCP: Dynamic Host Configuration Protocol: dynamically get address from as server
  - “plug-and-play”

DHCP: Dynamic Host Configuration Protocol

<u>Goal:</u> allow host to *dynamically* obtain its IP address from network server when it joins network

Can renew its lease on address in use

Allows reuse of addresses (only hold address while connected and "on")

Support for mobile users who want to join network (more shortly)

DHCP overview:

- host broadcasts "DHCP discover" msg
- DHCP server responds with "DHCP offer" msg
- host requests IP address: "DHCP request" msg
- DHCP server sends address: "DHCP ack" msg

DHCP client-server scenario

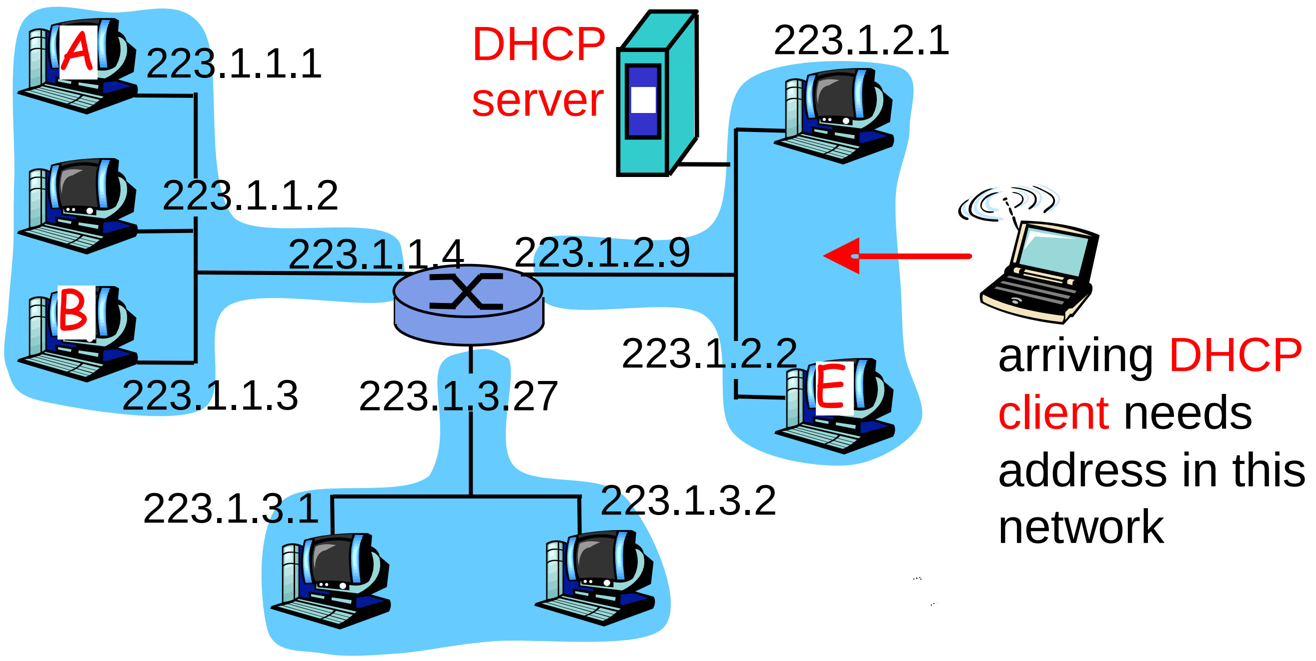

DHCP client-server scenario

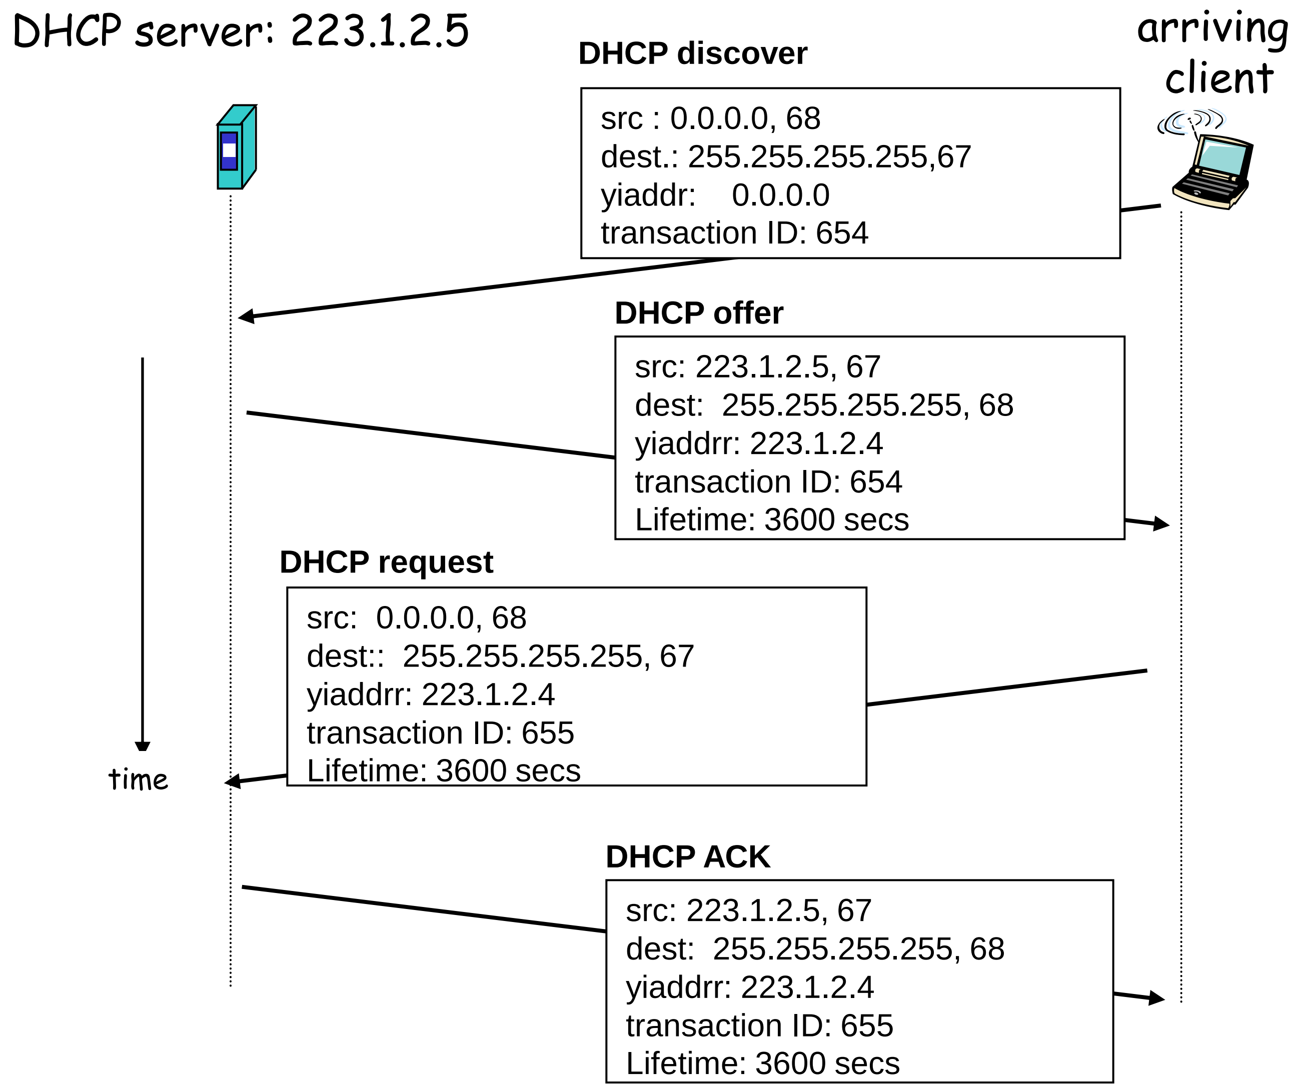

IP addresses: how to get one?

Q: How does *network* get subnet part of IP addr?

A: gets allocated portion of its provider ISP's address space

ISP's block	11001000 00010111 00010000 00000000	200.23.16.0/20
Organization 0	11001000 00010111 00010000 00000000	200.23.16.0/23
Organization 1	11001000 00010111 00010010 00000000	200.23.18.0/23
Organization 2	11001000 00010111 00010100 00000000	200.23.20.0/23
...	….. ….	….
Organization 7	11001000 00010111 00011110 00000000	200.23.30.0/23