The Network Layer: Data Plane

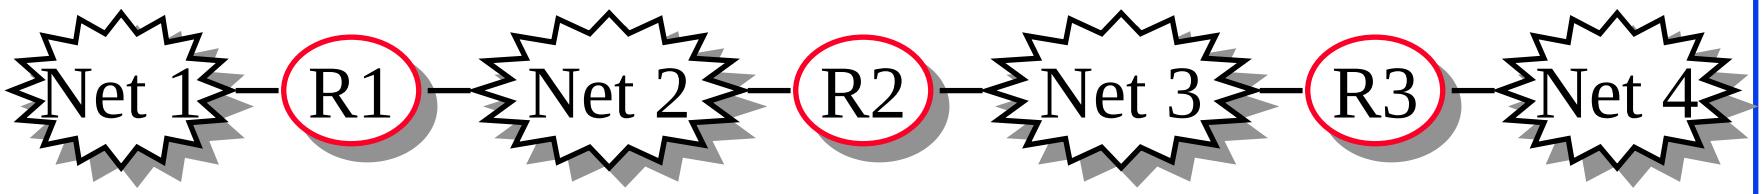

Raj Jain
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@wustl.edu

Audio/Video recordings of this lecture are available on-line at:

http://www.cse.wustl.edu/~jain/cse473-22/

Student Questions

Overview

1. Network Layer Basics
2. What's inside a router?
3. Forwarding Protocols: IPv4, DHCP, NAT, IPv6
4. Software Defined Networking

Note: This class lecture is based on Chapter 4 of the textbook (Kurose and Ross) and the figures provided by the authors.

Student Questions

- Why is the Upper Layer Protocol section of the datagram 8 bits long, are there that many transport protocols?

Network Layer Basics

1. Forwarding and Routing
2. Connection Oriented Networks: ATM Networks
3. Classes of Service
4. Router Components
5. Packet Queuing and Dropping

Student Questions

Forwarding and Routing

- **Forwarding**: Input link to output link via Address prefix lookup in a table.
- **Routing**: Making the Address lookup table
- **Longest Prefix Match**

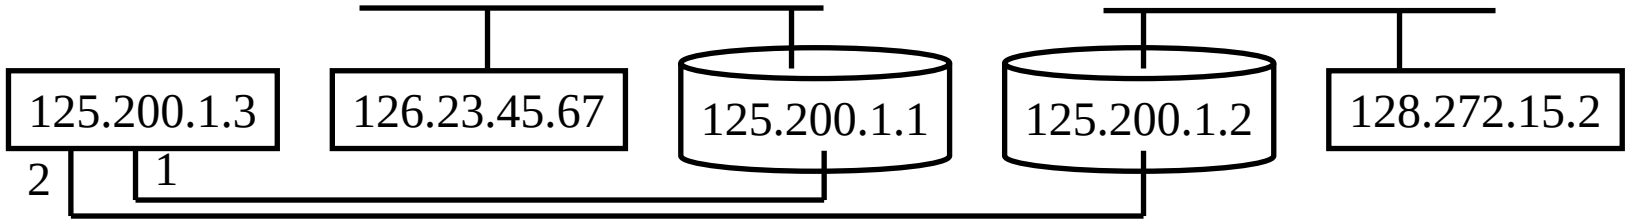

Prefix	Next Router	Interface
126.23.45.67/32	125.200.1.1	1
128.272.15/24	125.200.1.2	2
128.272/16	125.200.1.1	1

Student Questions

- Is there a limit for how long an address table can be?

No. There is no limit.

- Does each router have its own routing table, or can multiple routers share one routing table?

Each router has its own table.

- Is there a default next router to forward a packet to if the destination address does not match any entry in the lookup table?

Yes, it is called "Default Gateway." Actually, the last entry in the table is such that it matches all addresses and the next router is the default gateway.

- Can you explain the 3rd row of this table again?

If the first 16 bits are 128.272, then send it to 125.200.1.1 on interface 1.

- Does the slash in 128.272/16 mean that both 128.272. (...) and 128.16. (...) can be matched here?

/16 means the first 16 bits. 128.272 and 128.16 have only the first 8 bits common. They are different in the 9th bit.

Forwarding and Routing

- **Forwarding**: Input link to output link via Address prefix lookup in a table.
- **Routing**: Making the Address lookup table
- **Longest Prefix Match**

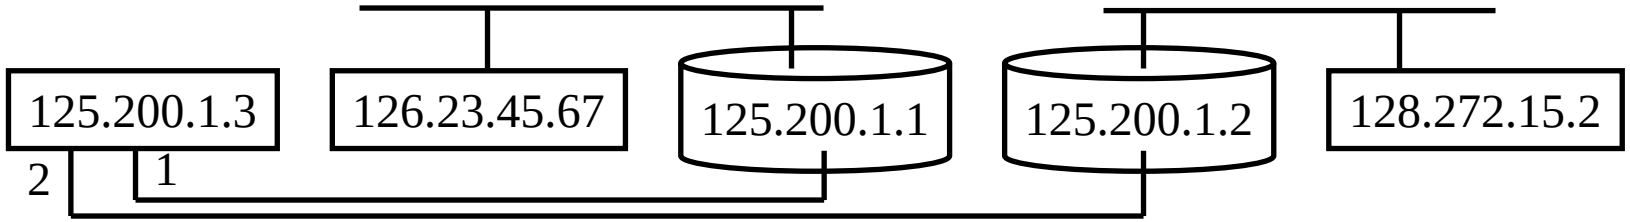

Prefix	Next Router	Interface
126.23.45.67/32	125.200.1.1	1
128.272.15/24	125.200.1.2	2
128.272/16	125.200.1.1	1

Student Questions

- If the prefixes overlap, what will happen to the longest prefix match?

The table entries are sorted by decreasing prefix length. For example, in the table on this slide, the lengths are 32, 24, 16. So the first entry to match is the longest prefix length.

- Is this routing table the standard design/template used in computer networks?

There is no standard. But this shows essential parts.

- Quiz 1 asks: Routing is the function of making the lookup table. The answer is false. Can you explain why?

Please see Piazza post. The answer is wrong.

- Which are considered first for longest Prefix match. From the left or the right?

Prefix is measured from the left.

- Could you explain the graph and the example you gave in the lecture again, since we cannot see your pointer and whiteboard?

Sure.

Network Service Models

- Guaranteed Delivery: No packets lost
- Bounded delay: Maximum delay
- In-Order packet delivery: Some packets may be missing
- Guaranteed minimal throughput
- Guaranteed maximum jitter: Delay variation
- Security Services (optional in most networks)
- ATM offered most of these
- IP offers none of these ⇒ Best effort service (Security is optional)

Optional Homework: R4, R5 in the textbook

Student Questions

- Why did IP become popular without these seemingly essential requirements for a protocol?

It was open-source, non-proprietary with free working implementations. NSF and DARPA made it spread fast in the academic world and then opened it to the outside.

- Does IP have an advantage over other network service models that provide reliability?

Yes. See the answer above.

- I thought we discussed in Ch. 3 that the network layer didn't guarantee delivery? Why does TCP assume unreliability then?

Some network layers do guarantee delivery. IP didn't and so TCP assumed an unreliable network layer.

- Can you explain what bounded delay is again?

Delay should be less than x ms.

What's Inside a Router?

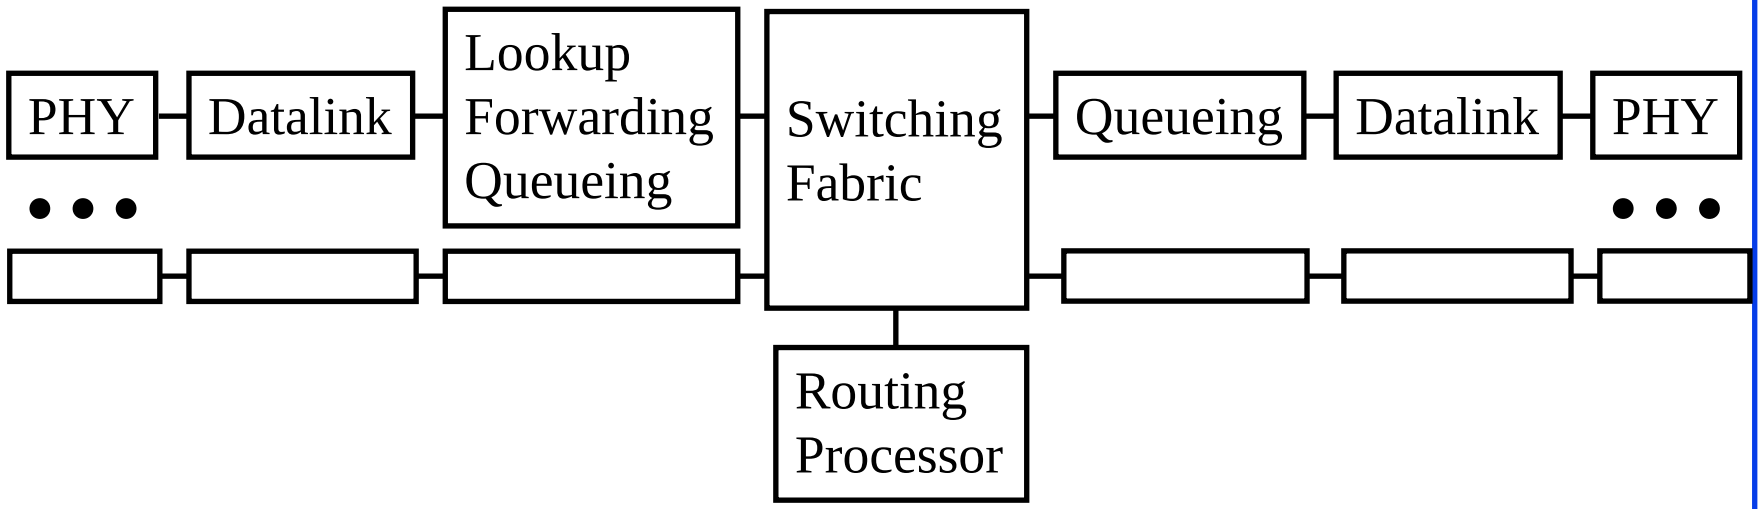

- **Input Ports**: receive packets, lookup address, queue
  Use **Content Addressable Memories** (CAMs) and caching
- **Switch Fabric**: Send from input port to output port
- **Output Ports**: Queuing, transmit packets

Student Questions

- Do this input physical link also serve as the output physical link back to wherever the input came from?

Generally, yes. However, simplex (one-way) links are possible.

- Is there any logic included in the switching fabric, or is it just pipe from input ports to output ports?

There is a lot of logic in the fabric.

- Can I think of forwarding queuing and queuing as buffers? Are they physical components or software/programs?

They are both memory buffers.

- What is PHY?

PHYsical layer

- Could you explain CAMs again?

You can lookup Content addressable memory. by content. It tells you which row matches. Then you can lookup other info in that row. There is no need for searching.

- The router will check the IP header of the passing datagram. Will it process the TCP header or the UDP header?

No, IP processes only IP header.

Types of Switching Fabrics

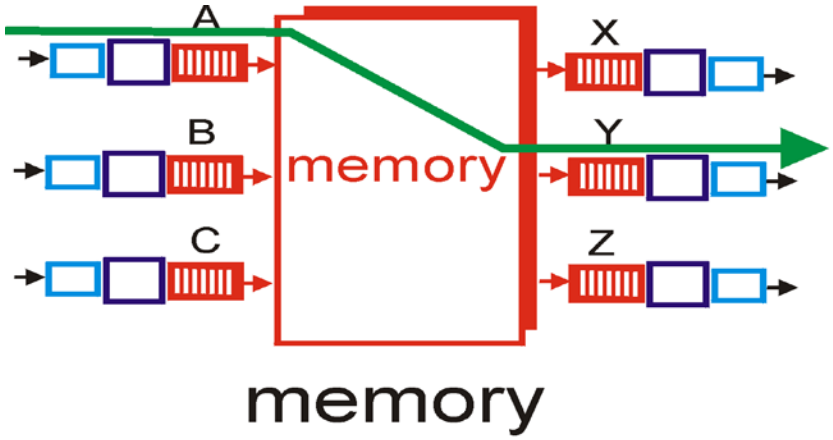

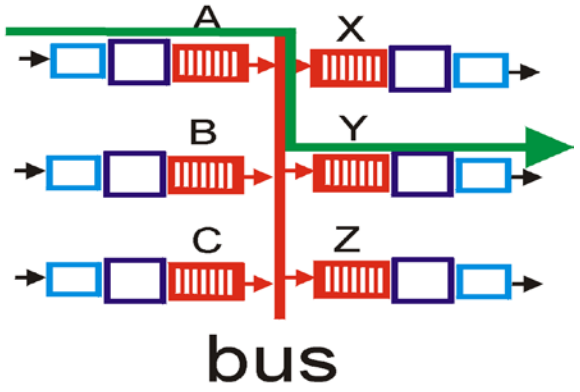

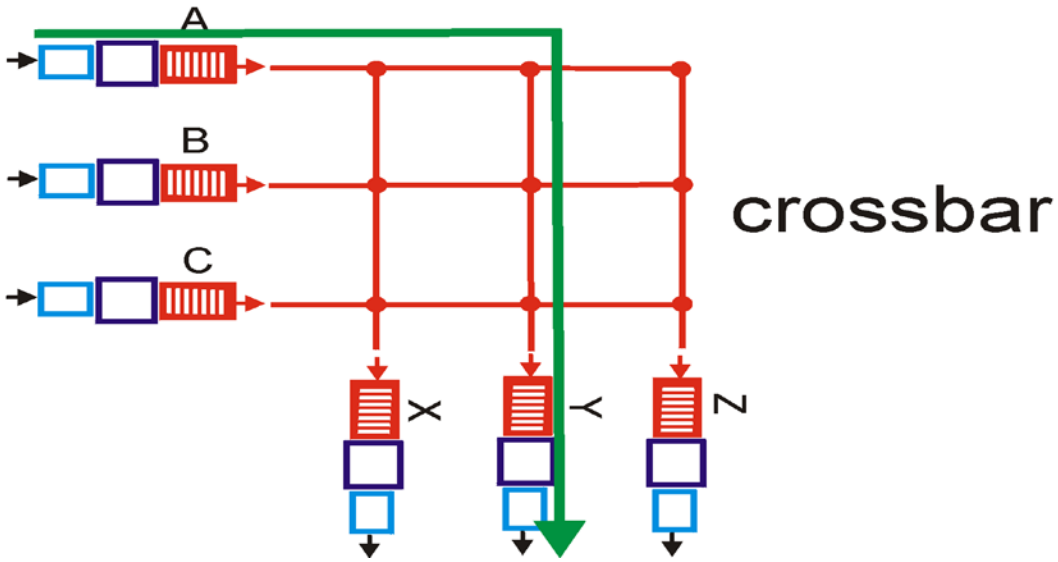

Student Questions

- Is there an industry standard for switching or is it at the discretion of each manufacturer?

It is at the discretion of each manufacturer.

- Will there be performance differences between different types of switching fabrics?

Yes. Different delays and throughput.

- Switching Fabrics sounds like a process changing the physical circuit, but indeed it is logical switching, right? Just like request forwarding?

It could be either.

- With memory-based switching, does the memory maintain a table to decide how to forward packets? If so, is that table fixed, or can it be modified?

Yes. It can be modified by changing the firmware.

- Can you explain how to run multiple switching fabrics in parallel to scale the router's capacity?

In CSE 570S: Advanced networking, we teach how to make 4x4 switches using 2x2 switches and so on.

Where Does Queuing Occur?

- If switching fabric is slow, packets wait on the input port.
- If switching fabric is fast, packets wait for output port
  ⇒ Queueing (Scheduling) and drop policies
- Queueing: First Come First Served (FCFS), Weighted Fair Queueing

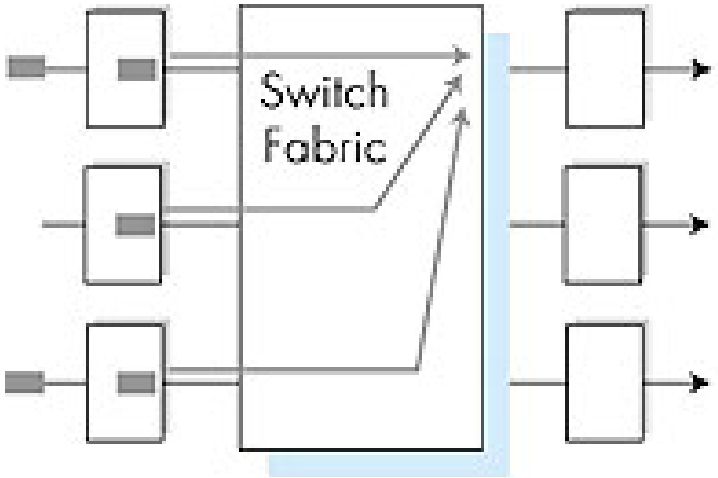

Output Port Contention at Time *t*

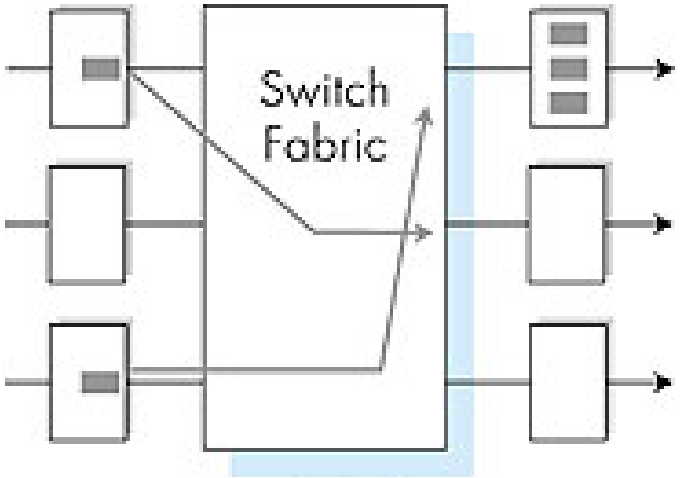

One Packet Time Later

Student Questions

- Why are packets switching in router, instead of byte streaming taught in TCP?

Packets are variable size. So you can send a packet with 1-byte of data. However, TCP header and applications headers will result in larger data sizes.

- If increasing usage over a greater number of ports would help decrease queueing, why aren't more ports in standard use?

They are. Datacenters may use 4096x4096 port switches (as taught in the CSE 570S: Advanced networking)

- How do we indicate priority for queuing?

IP header has a priority field.

Ideal Buffering

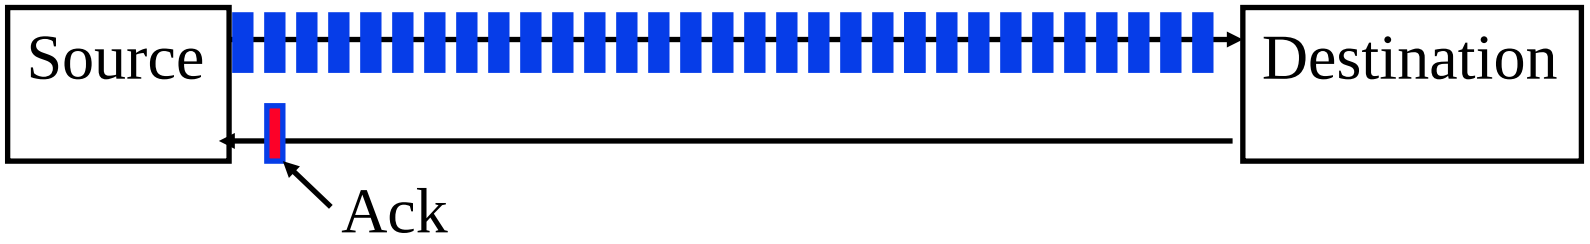

- Flow Control Buffering = RTT*Transmission Rate
- Buffer = RTT*Transmission Rate/√(# of TCP flows)

Student Questions

- Can you clarify what this flow control buffering referring to? Is this the buffer for the entire link and then when you divide by sqrt(# TCP flows) that is the buffer for what? Do input ports have a separate buffer from the entire link?

Buffers are at the destination. The buffers have to be as large as the number of bits on the wire.

- The book says: "router buffers ... for buffer sizing ... the amount of buffering should be equal to the average RTT times the link capacity" Where does this fit in?

Number of bits on the wire
= Length of the link in sec × Bits/sec
=RTT Link× Capacity

- What are TCP flows, the number of packets in one pipeline?

All packets with the same IP address and port number.

- Do we take the floor of the square root of TCP flows?

This is an approximation. So you can use the roof, floor, or rounding to the nearest integer. However, if the answer is less than 1, the roof would be the only possibility.

Packet Dropping Policies

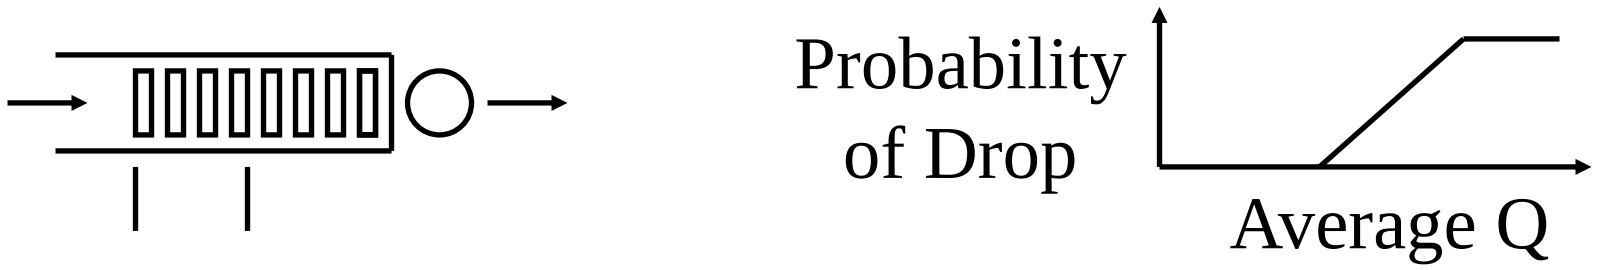

- **Drop-Tail**: Drop the arriving packet
- **Random Early Drop (RED):** Drop arriving packets even before the queue is full
  - Routers measure average queue and drop incoming packet with certain probability

  ⇒ **Active Queue Management** (AQM)

Student Questions

- Are dropped packets managed as packet loss?

Yes.

- How to mark the lost packets? How to recover the dropped packets, like resend the loss packets?

There is no marking. Transport will recover if necessary.

- Which policy is followed today, drop tail or random early drop?

RED is quite common.

- How well does RED work in practice?

RED works quite well.

- Will it be more efficient for Go-back-N if we drop the last packet or drop the smallest packet for the selective repeat? I mean, there could be a way to mark the most efficient packet dropping in the header.

Yes, that marking is implemented as the "Selective Ack." feature introduced later as an enhancement. See RFC 2018.

Head-of-Line Blocking

- Packet at the head of the queue is waiting
  ⇒ Other packets can not be forwarded even if they are going to other destination

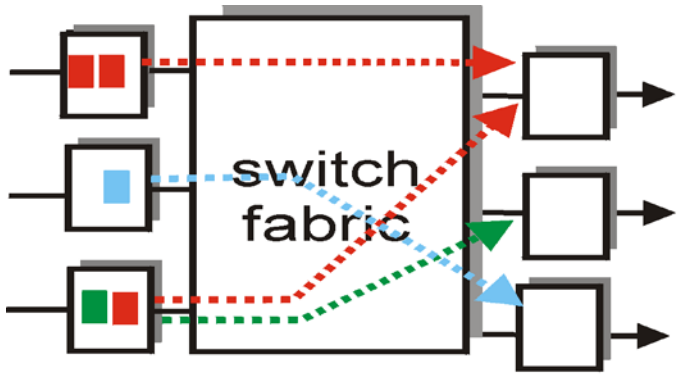

output port contention at time t - only one red packet can be transferred

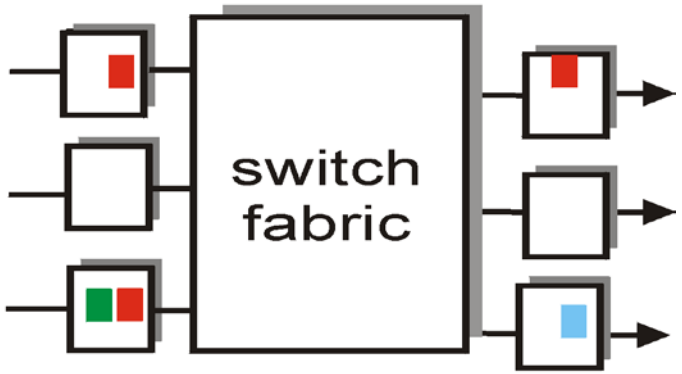

green packet experiences HOL blocking

Student Questions

- Is there a way to alleviate the HOL blocking?

Yes. There are ways. Too much detail for this first course.

- Can you please explain head-of-line blocking again?

Sure.

- Are there any other packet-dropping policies?

Anyone can make his own.

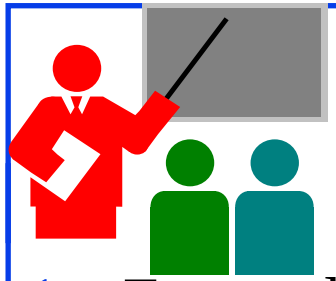

Network Layer Basics: Review

1. Forwarding uses routing table to find output port for datagrams using **longest prefix match**. Routing protocols make the table.
2. IP provides only **best effort** service (KISS).
3. Routers consist of input/output ports, **switching fabric**, and processors.
4. Datagrams may be dropped even if the queues are not full (**Random early drop**).
5. Queueing at input may result in **head of line blocking**.

Student Questions

- What is the advantage of having datagrams dropped randomly?

It makes it fairer. Since the drops are distributed among many flows.

Ref: Read Sections 4.1, 4.2, full, Page 303-330 of the textbook. Try R1 through R16

Forwarding Protocols

1. IPv4 Datagram Format
2. IP Fragmentation and Reassembly
3. IP Addressing
4. Network Address Translation (NAT)
5. Universal Plug and Play
6. Dynamic Host Control Protocol (DHCP)
7. IPv6

Student Questions

IP Datagram Format

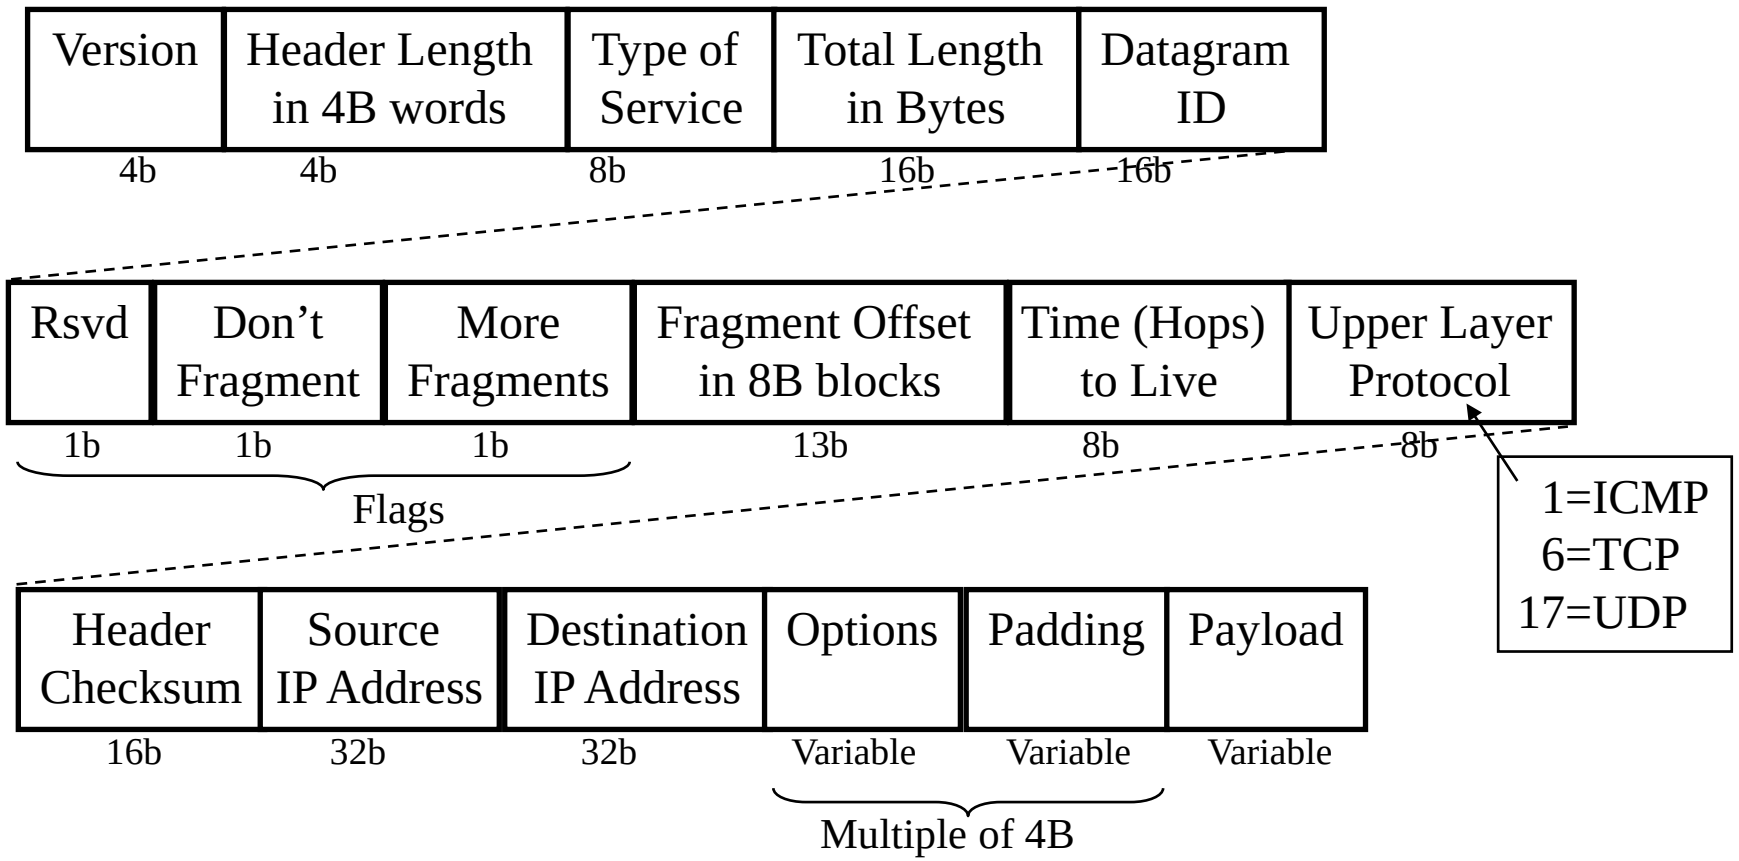

Student Questions

- To clarify, type of service is not used?

It was not used for long time. Several proposal have recently been made to use it. So it is used now.

- Payload is what everything passed from the top layer? *Yes.*
- What type of service is used now?

Details of this 8-bit field are discussed in CSE570S: Advanced Networking. Two of those bits are used for ECN discussed in Chapter 3.

- What proposals use a type of service?

See the previous question.

- Is the pseudo-header used in computing TCP checksum stored in this IP datagram?

TCP Checksum covers TCP header, TCP payload, IP addresses, and Protocol fields. IP Checksum covers only the IP header.

- Why do we do separate checksums for IP and TCP? Couldn't you just include the payload in the IP checksum so that it would only need to be calculated once?

See the previous question.

IP Datagram Format

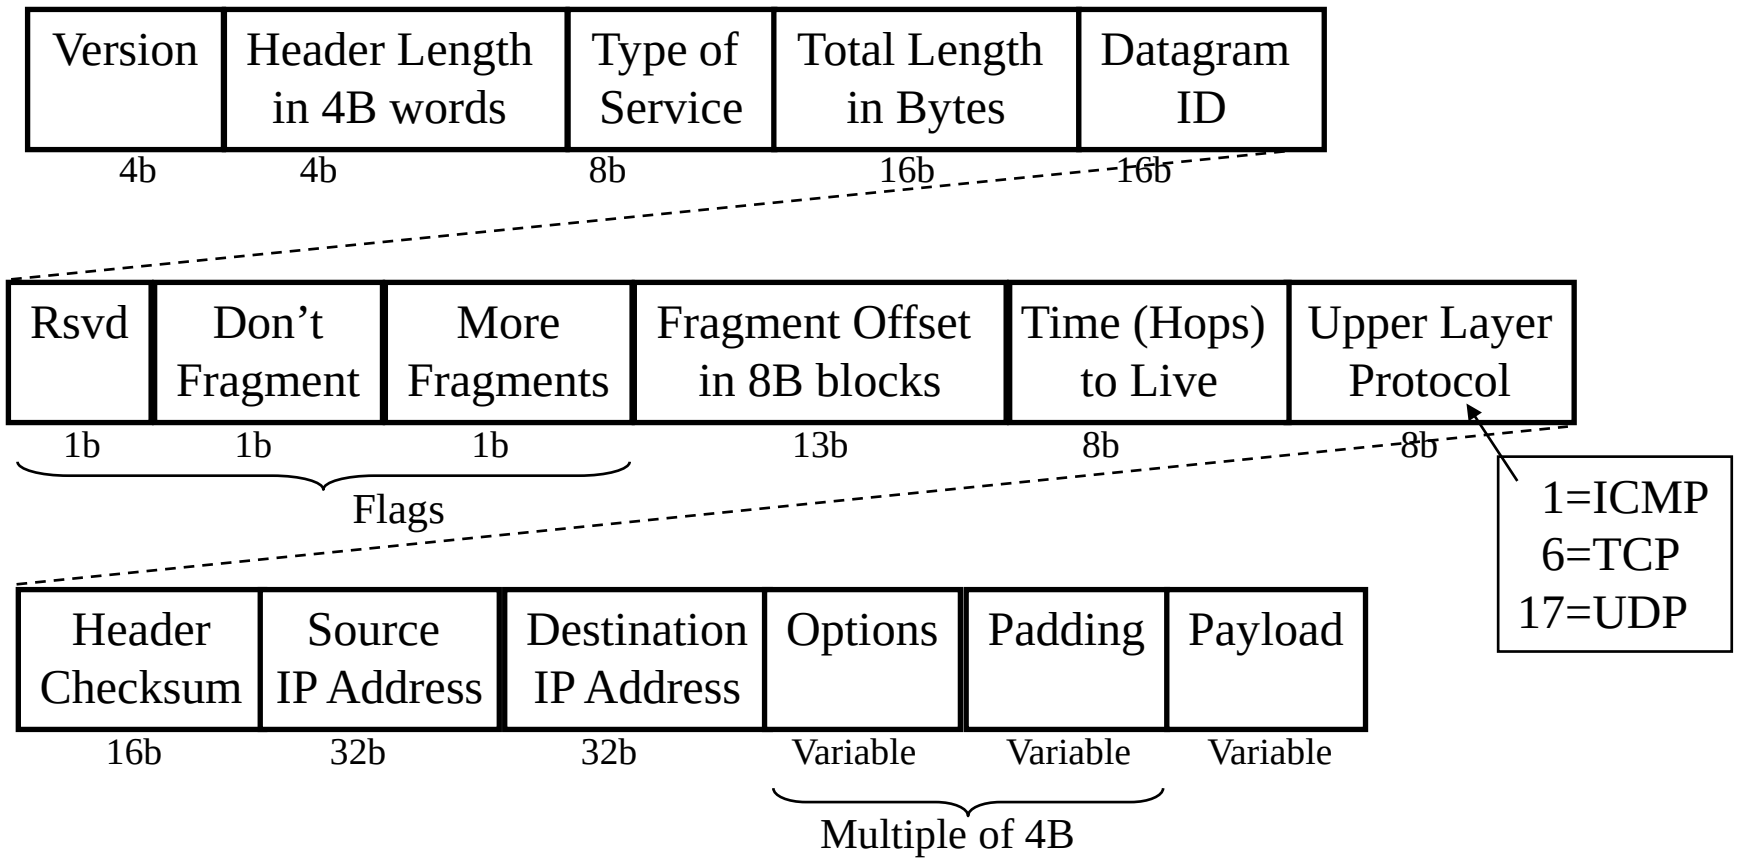

Student Questions

- To clarify, type of service is not used?

It was not used for long time. Several proposal have recently been made to use it. So it is used now.

- For the IP Datagram format, why do the options and padding section have to be multiples of 4 bytes?

Every layer in TCP/IP uses its own unit of length. IP specifies header length in units of 4-byte words.

IP Fragmentation Fields

- Header length: in units of 32-bit words
- Data Unit Identifier (ID)
  - Sending host puts an identification number in each datagram
- Total length: Length of user data plus header in bytes
- Fragment Offset - Position of fragment in original datagram
  - In multiples of 8 byte blocks
- *More fragments* flag
  - Indicates that this is not the last fragment
- Datagrams can be fragmented/refragmented at any router
- Datagrams are reassembled only at the destination host

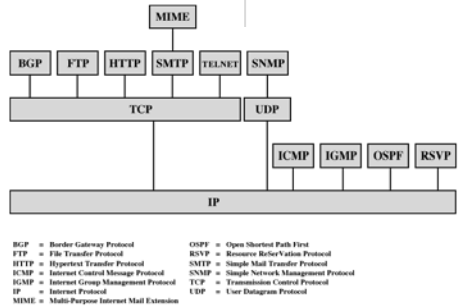

Student Questions

- What are some examples of other Upper Protocol Layer numbers? How many are there?

See Slide 1-44 (Figure above)

- Does total length include the other layers?

Higher layer headers are simply data for IP. Lower layers, it does not know.

- Is IP ID the same as TCP's unique sequence number?

No. TCP uses byte numbers. IP uses packet number.

- For header length, is it a 4-Byte word or a 4-bit word? It's a little confusing.

4-byte words.

- Why do we need offset?

To be able to put the fragment back in its position.

- TCP can ensure that routers do not fragment segments. How does it get the minimum MTU on the path and calculate the MSS?

One way would be to send a test datagram with the no-fragment bit set and see if it is dropped. The dropping router will send a reason using ICMP.

IP Fragmentation Fields

- Header length: in units of 32-bit words
- Data Unit Identifier (ID)
  - Sending host puts an identification number in each datagram
- Total length: Length of user data plus header in bytes
- Fragment Offset - Position of fragment in original datagram
  - In multiples of 8 byte blocks
- *More fragments* flag
  - Indicates that this is not the last fragment
- Datagrams can be fragmented/refragmented at any router
- Datagrams are reassembled only at the destination host

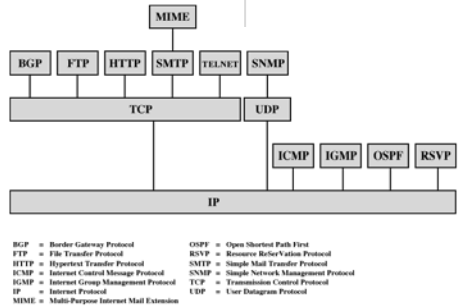

Student Questions

- What are some examples of other Upper Protocol Layer numbers? How many are there?

See Slide 1-44 (Figure above)

- Does total length include the other layers?

Higher layer headers are simply data for IP. Lower layers, it does not know.

- Are these fields the same between IPv4 and IPv6?

No. IPv4 and IPv6 are not compatible.

- Does the total length only represent the length of this layer?

Header plus data at this layer. Data includes all upper-layer headers and data.

IP Fragmentation and Reassembly

Example

- 4000 byte datagram
- Maximum Transmission Unit (MTU) = 1500 bytes

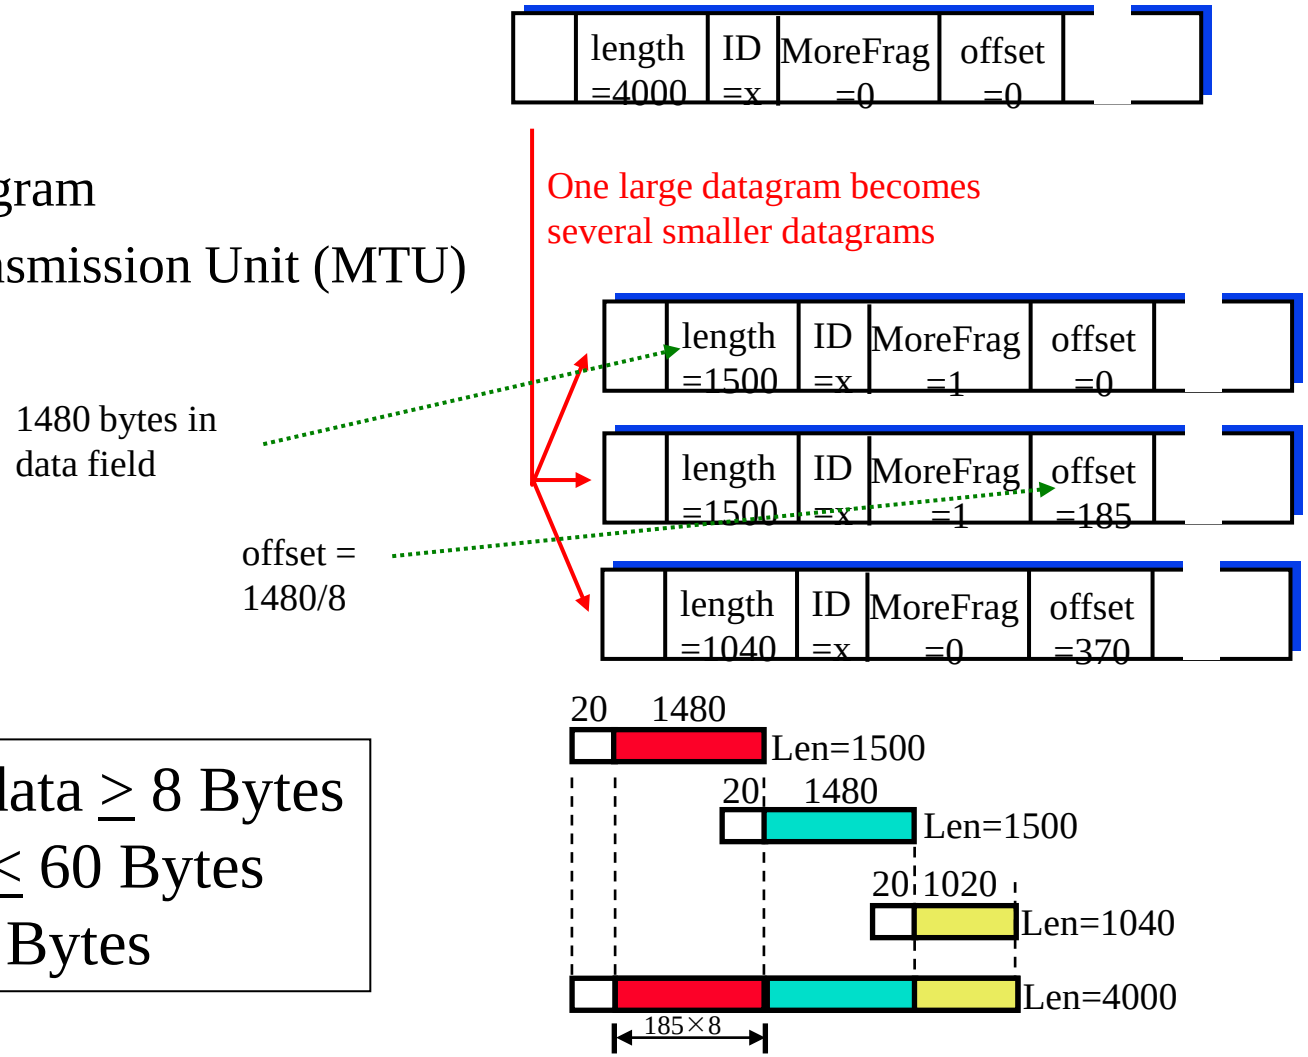

Fragment data ≥ 8 Bytes
IP Header ≤ 60 Bytes
MTU ≥ 68 Bytes

Student Questions

- Would you please explain how we get the header is 20 bytes here?

See Slide 4-14. Assume options + pad length=0

- Why is the MoreFrag header necessary?

To indicate whether this is the last fragment.

- Should the length of all of the 3 fragments add up to 4040?

In general, lengths are related as follows: $\sum(Fragment_i\text{-header})=total\ datagram\text{-}header$

- If only the whole length of 4000 bytes and memory capacity of 1500 bytes is given, how should we work out the length for the header?

Assume a 20-byte header unless specified otherwise.

- In IP fragmentation, how did you come to the offset being 370 for the last segment?

Offset is measured in units of 8B.
The last segment is 1480+1480=2960B away from the beginning ⇒ Offset=2960/8=370

IP Fragmentation and Reassembly

Example

- 4000 byte datagram
- Maximum Transmission Unit (MTU) = 1500 bytes

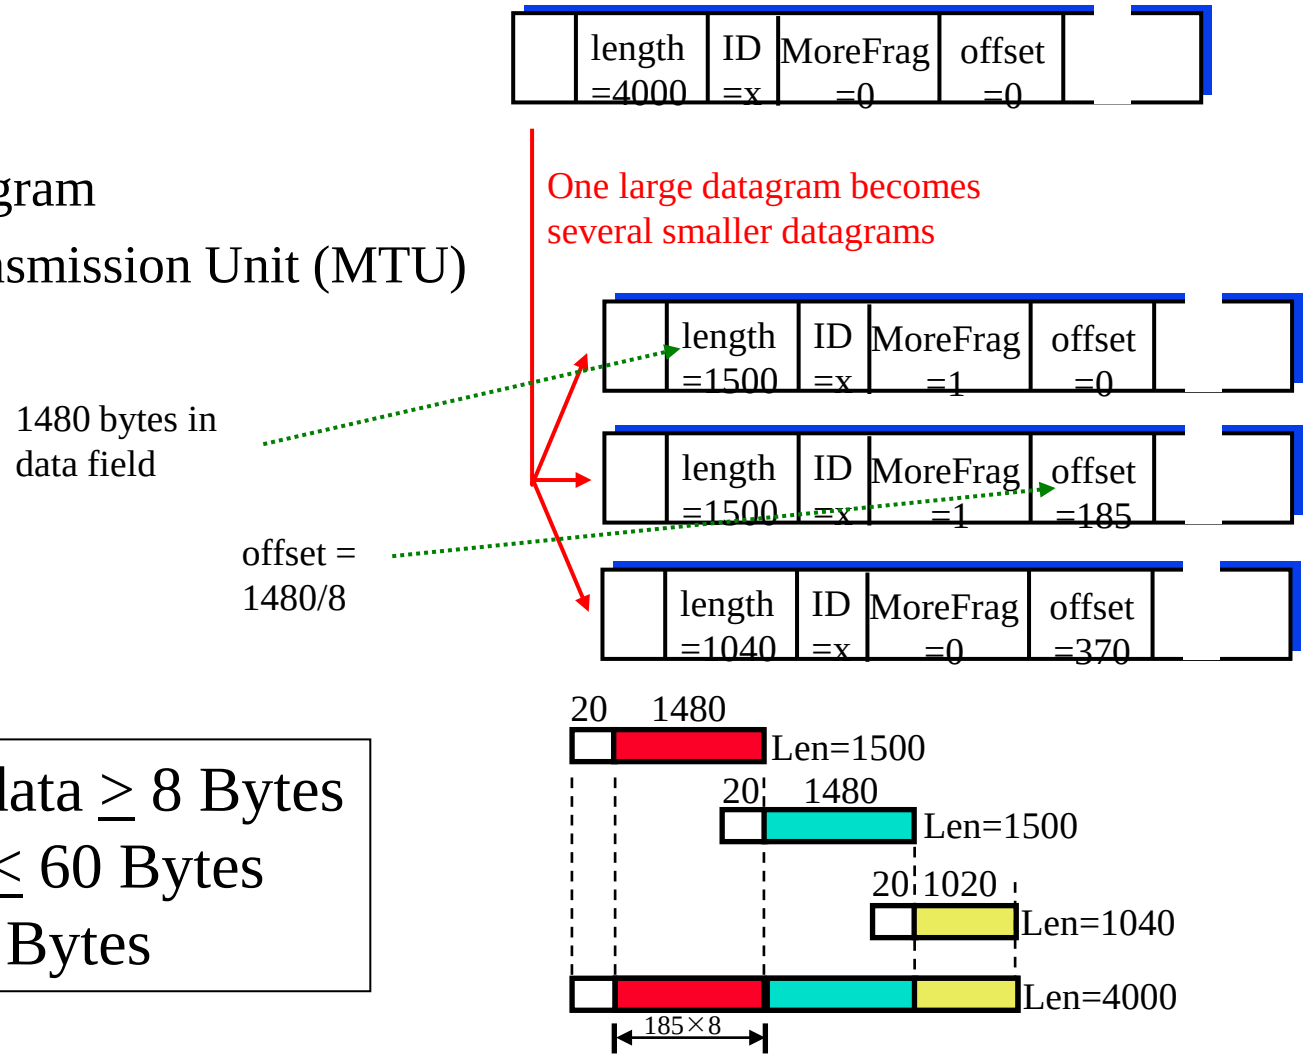

Fragment data ≥ 8 Bytes
IP Header ≤ 60 Bytes
MTU ≥ 68 Bytes

Student Questions

- Why is fragmentation not the responsibility of TCP?

Paths can change. TCP is only end-to-end. IP determines the path. Fragmentation depends on the path.

- Is it necessary to do fragmentation when every router has enough memory and MTU is large enough?

No, not unless the segment is too large. Datalink layers also put limits on frame sizes.

- Would the ID for each fragment be similar to the TCP sequence number, or would the ID of all the fragments be the same as the ID number of the initial fragment?

ID here is the packet number. All fragments will have the same ID.

- When a host receives the datagram, how does it know more segments are coming?

MoreFrag flag

- Is this graph showing that the last 20 bytes of the first segment will be the header of the second segment?

No. There is no dotted line connecting the beginning of 20B to the previous fragment.

Homework 4A: Fragmentation

- [8 points] Consider sending a 2000-byte datagram into a link that has an MTU of 600 bytes. Suppose the original datagram is stamped with the identification number 422. How many fragments are generated? What are the values in the various fields in the IP datagram(s) generated related to fragmentation?

Student Questions

IP Address Classes

- Class A:

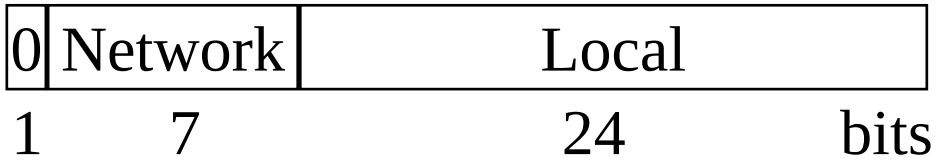

- Class B:

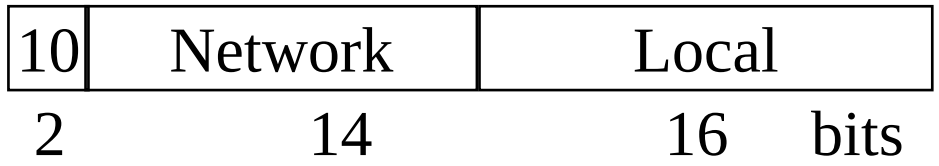

- Class C:

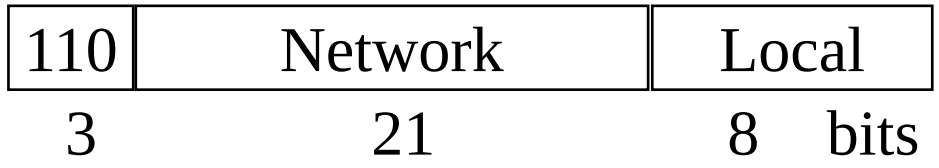

- Class D:

- Class E:

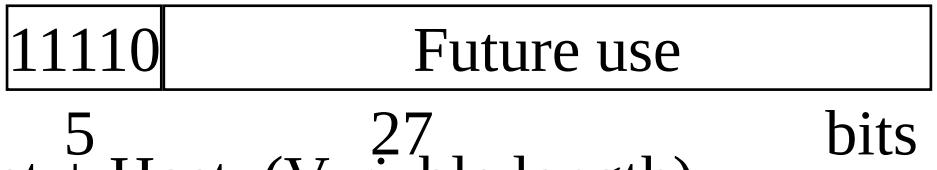

- Local = Subnet + Host (Variable length)

Student Questions

- Since the number of bits in the beginning to specify class is not fixed, how do we know when these bits end and the bits belonging to the "Network" section start? (Do we keep reading 1s until we see a 0)?

Yes, classes were coded to end in 0.

- Does this explain why my laptop and smartphone have different IP addresses using the same Wi-Fi?

Every node in the network has its own IP address.

- What are some possible future uses of class E?

Classes are no longer used.

- Could you go over the differences between network bits and local bits?

Network bits are same for all hosts on the network. Only local bits distinguish one node from another. Similar to the last name in a family. Local bits are like the first names.

IP Addressing

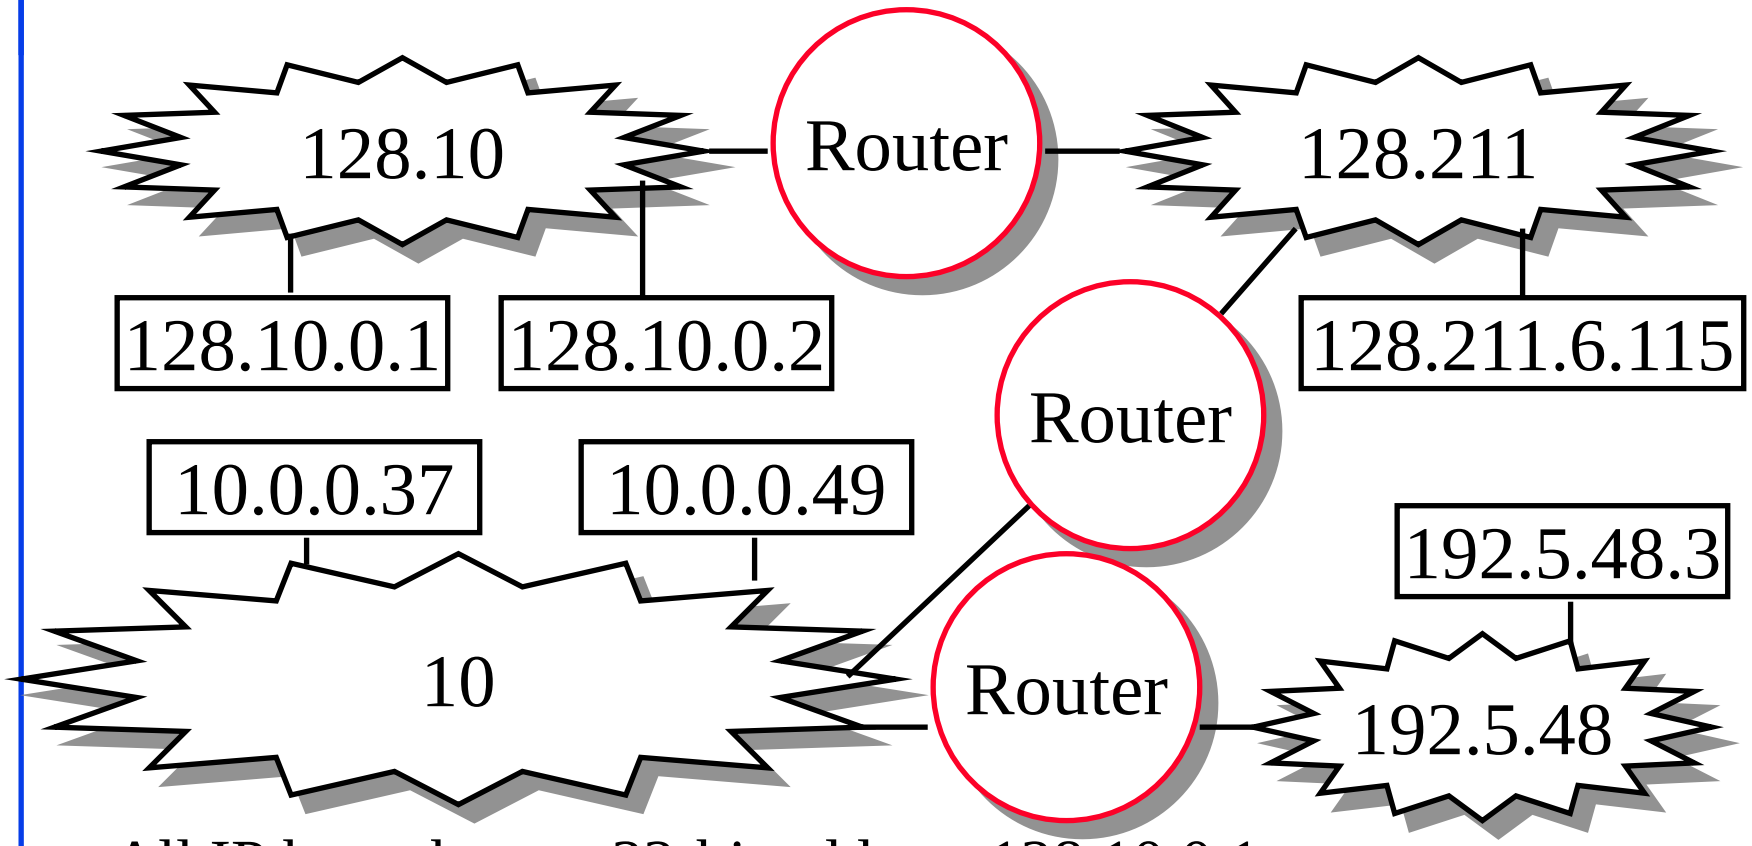

- All IP hosts have a 32-bit address.128.10.0.1 =1000 0000 0000 1010 0000 0000 0000 0001
- All hosts on a network have the same network prefix

Student Questions

- Do routers have IP addresses?

Yes. You need them to program routers.

- The IP address on host devices will change when switching networks?

Yes.

Subnetting

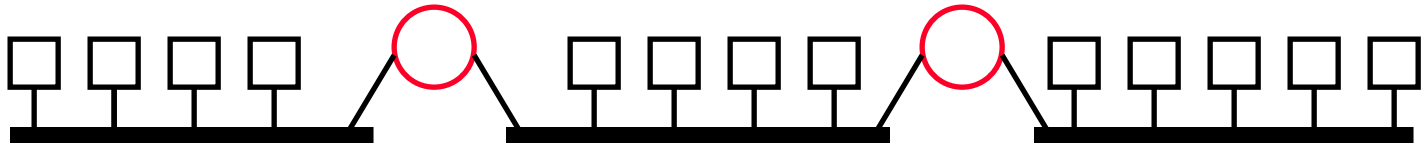

- All hosts on a subnetwork have the same prefix. Position of the prefix is indicated by a "subnet mask"
- Example: First 23 bits = subnet

Address:	10010100 10101000 00010000 11110001
Mask:	11111111 11111111 11111110 00000000
.AND.	10010100 10101000 00010000 00000000

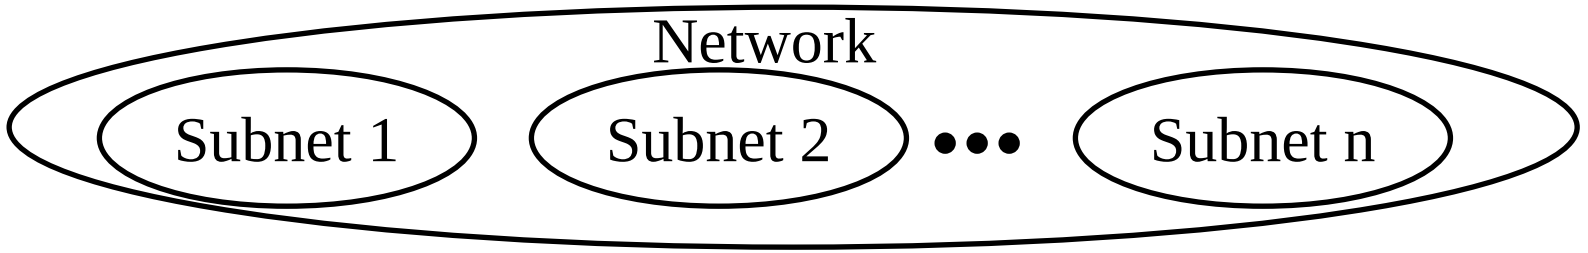

Student Questions

- What prevents two systems from having (or claiming they have) the same IP address?

IP addresses are allocated by the network manager or by DHCP. They cannot be the same.

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
  - All 1's in the host part is used for subnet broadcast
  - All 0's in the host part was meant as "subnet address" but not really used for anything. Some implementation allow it to be used as host address. Some don't. Better to avoid it.

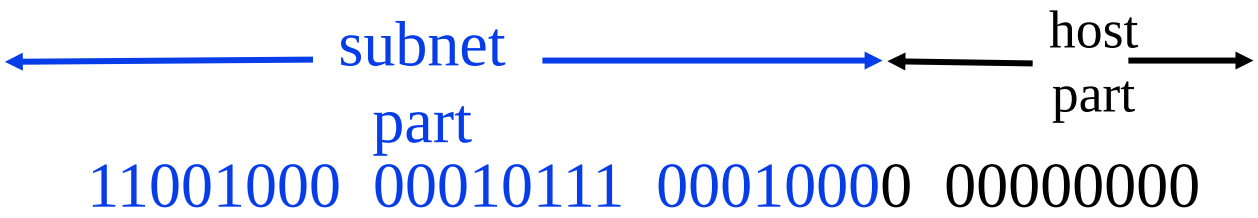

200.23.16.0/23

Student Questions

- How does it determine what size mask is necessary for the subnet?

The network manager determines depending on how many nodes, she needs to support on the network.

Homework 4B: Subnets

- [18 points] Consider a router that interconnects 3 subnets: Subnet 1, Subnet 2, and Subnet 3. Suppose all of the interfaces in each of these three subnets are required to have the prefix 223.1.17/24. Also suppose that Subnet 1 is required to support up to 51 interfaces, Subnet 2 is to support up to 96 interfaces, and Subnet 3 is to support up to 26 interfaces. Provide three network address prefixes (of the form a.b.c.d/x) that satisfy these constraints. **Use adjacent allocations**. For each subnet, also list the subnet mask to be used in the hosts.

Student Questions

- Since the subnets are required to have the prefix 223.1.17/24, wouldn't they be required to have the same mask as their required prefix, if not then why does the prefix even include a mask?

Each subnet can have a longer prefix than 24 bits, leaving only as many bits as required for the hosts. For example, subnet 1 can have 6-bits for hosts and 26 bits for the subnet prefix.

Forwarding an IP Datagram

- Delivers **datagram**s to destination network (subnet)
- Routers maintain a "routing table" of "next hops"
- Next Hop field does not appear in the datagram

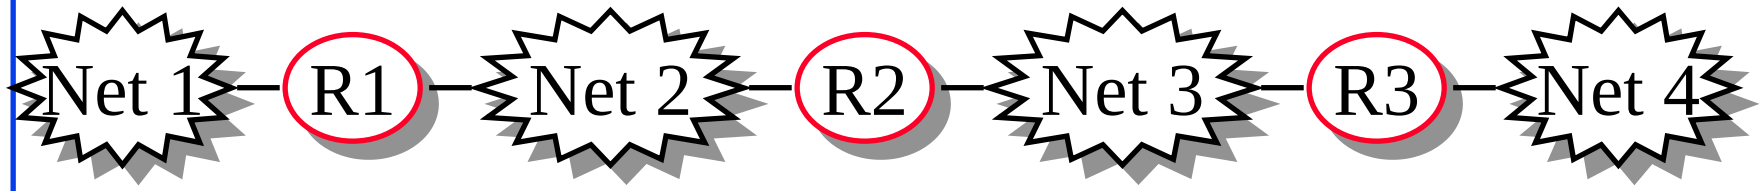

Table at R2:

Destination	Next Hop
Net 1	Forward to R1
Net 2	Deliver Direct
Net 3	Deliver Direct
Net 4	Forward to R3

Student Questions

- What is the length of the IP datagram header? Does it vary?

See Slide 4-14

- Do the other layers' headers need to get duplicated for each fragment?

IP only cares about its headers. Its header gets duplicated. Other layers are part of data.

- What is the purpose of having various networks in the same path?

Different companies cover different areas.

- So the network delivers datagrams to the closest router if it does not have next hop?

No. If it does not have the next hop, it delivers it to the "Default Gateway"

- How often does routing tables are updated in a router?

Whenever a new link is added, a link breaks, or the router boots up.

Route Aggregation

- Can combine two or more prefixes into a shorter prefix
- ISPs-R-Us has a more specific route to organization 1

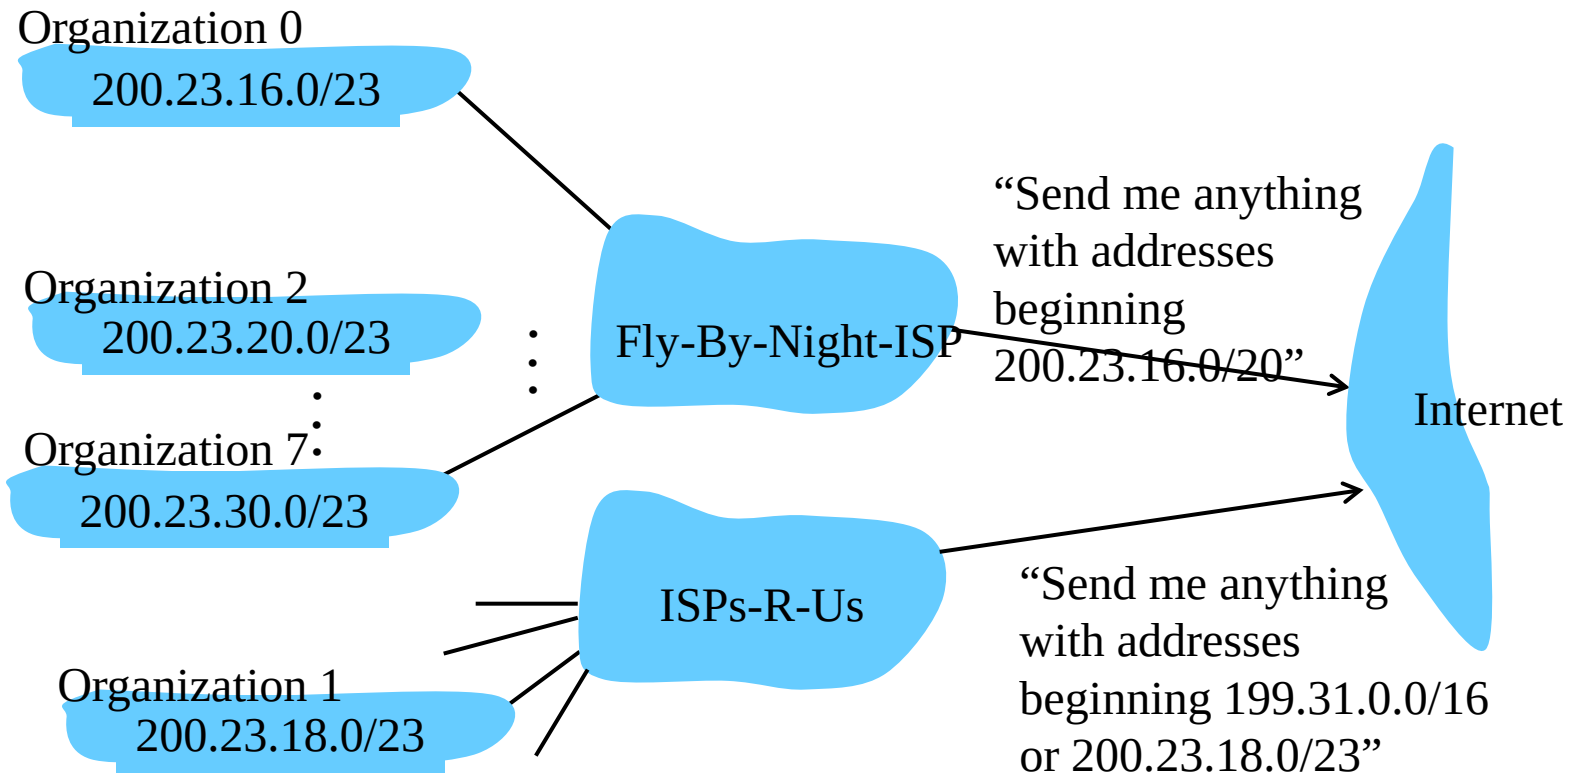

Student Questions

- What is the difference between Fly-By-Night-ISP vs ISPs-R-Us?

No difference. Just funny names.

"Route Print" Command in Windows

MAC: netstat -rn

```
===========================================================================
Interface List
0x1 ........................... MS TCP Loopback interface
0x2 ...00 16 eb 05 af c0 ...... Intel(R) WiFi Link 5350 - Packet Scheduler Miniport
0x3 ...00 1f 16 15 7c 41 ...... Intel(R) 82567LM Gigabit Network Connection - Packet Scheduler Miniport
0x40005 ...00 05 9a 3c 78 00 ...... Cisco Systems VPN Adapter - Packet Scheduler Miniport
===========================================================================
===========================================================================
Active Routes:
Network Destination        Netmask          Gateway       Interface  Metric
          0.0.0.0          0.0.0.0      192.168.0.1   192.168.0.108              10
          0.0.0.0          0.0.0.0      192.168.0.1   192.168.0.106              10
        127.0.0.0        255.0.0.0        127.0.0.1       127.0.0.1               1
    169.254.0.0        255.255.0.0    192.168.0.106   192.168.0.106              20
    192.168.0.0    255.255.255.0    192.168.0.106   192.168.0.106              10
    192.168.0.0    255.255.255.0    192.168.0.108   192.168.0.108              10
  192.168.0.106  255.255.255.255        127.0.0.1       127.0.0.1              10
  192.168.0.108  255.255.255.255        127.0.0.1       127.0.0.1              10
  192.168.0.255  255.255.255.255    192.168.0.106   192.168.0.106              10
  192.168.0.255  255.255.255.255    192.168.0.108   192.168.0.108              10
        224.0.0.0        240.0.0.0    192.168.0.106   192.168.0.106              10
        224.0.0.0        240.0.0.0    192.168.0.108   192.168.0.108              10
  255.255.255.255  255.255.255.255    192.168.0.106   192.168.0.106      1
  255.255.255.255  255.255.255.255    192.168.0.106         40005        1
  255.255.255.255  255.255.255.255    192.168.0.108   192.168.0.108      1
Default Gateway:       192.168.0.1
===========================================================================
Persistent Routes:
  None
```

Adr & mask = Dest
⇒ Match

Longest Prefix match is used

Metric: Lower is better

Note: 127.0.0.1 = Local Host, 224.x.y.z = Multicast on local LAN

Student Questions

- Do packets sent to 127.0.0.1 ever actually leave the computer onto the network before returning or is it all internal?

Internal loopback.

- What is the difference between the interface and the gateway? What is network destination vs. gateway? How do you know which interface is specified by the given address under that field?
- *Interface=Adapter*
  Gateway=Router
  Net. Destination=Dest Adr

- What exactly is the metric?

Lower is better

- What is the difference between gateway and router?

None in practice.

Lab 4A: Routing Table

- [8 Points] Use "Route Help" in Windows (or man route in MAC) to learn the route command
- Ping www.google.com to find its address
- Make sure that you have two active interfaces preferably connected to different routers. For example, create a 2nd interface by connecting a smart phone hot spot via USB. Or by connecting to a router in our lab during TA hours
- Print route table
- Trace route to www.google.com using tracert
- Modify the routing table so that the other interface will be used.
- Note the command you used to modify the routing table
- Print the new routing table
- Trace route to the same numeric address for www.google.com as before . Submit underlined items.

Student Questions

- Don't have a phone hotspot? Could I just use a non-washu VPN?

Not sure if traceroute will work with VPN. Did you try and did it work?

Lab 4A Hints

- A host with two interfaces going to the same router:

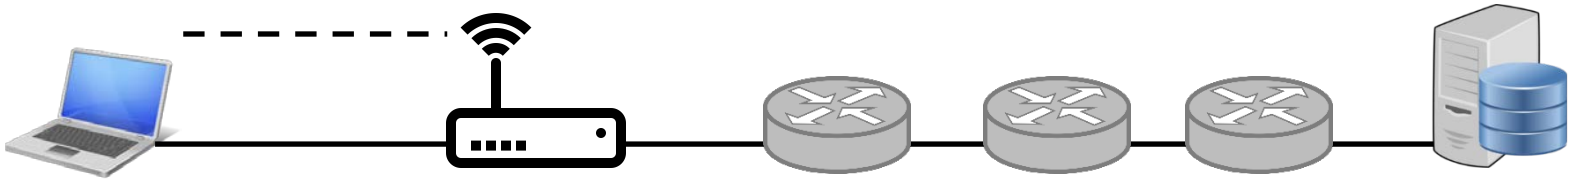

- Trace route result will not change even if you change the interface.

```
IPv4 Route Table
===========================================================================
Active Routes:
Network Destination        Netmask          Gateway       Interface  Metric
          0.0.0.0          0.0.0.0      192.168.0.1    192.168.0.152     55
          0.0.0.0          0.0.0.0      192.168.0.1    192.168.0.151     25
```

Student Questions

Lab 4A Hints (Cont)

- If you have two routers, you can see the effect in trace route. One way to get two routers is to use your cell phone hot spot:

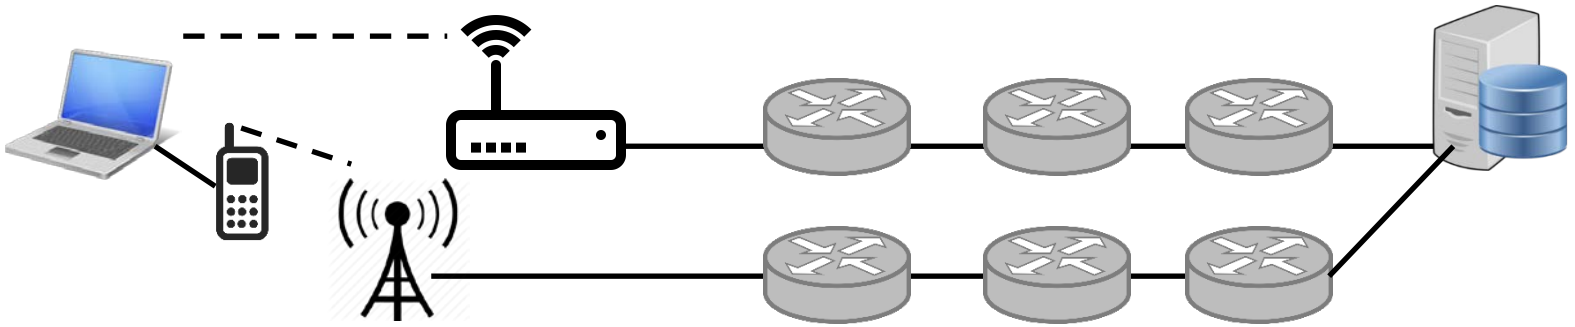

```
IPv4 Route Table
===========================================================================
Active Routes:
Network Destination        Netmask          Gateway       Interface  Metric
          0.0.0.0          0.0.0.0      192.168.0.1    192.168.0.151     25
          0.0.0.0          0.0.0.0      172.20.10.1     172.20.10.2     35
```

- WiFi on phone should be disabled to ensure that it does not forward the traffic to the same home router.

Student Questions

Lab 4A Hints (Cont)

- Another way to get two routers is to use another router. We have placed an extra router in our lab.

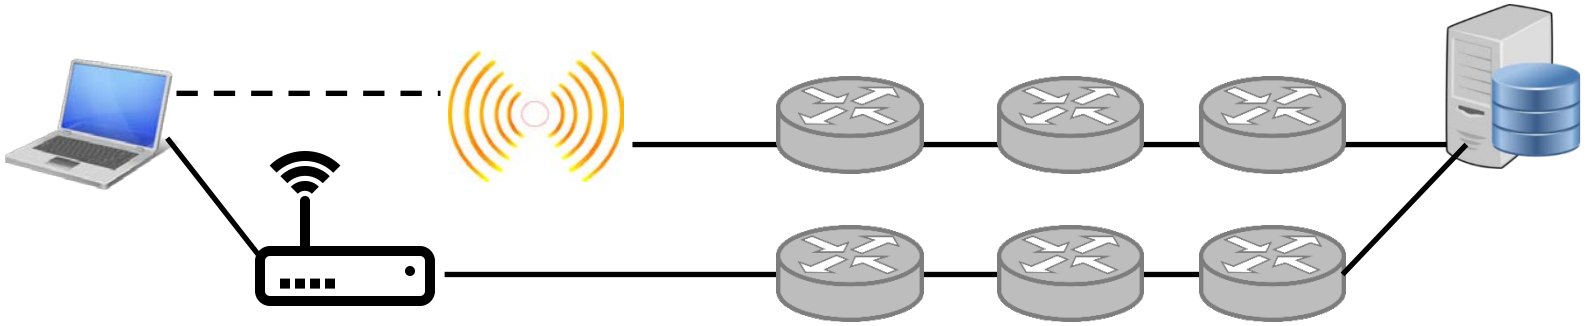

```
IPv4 Route Table
===========================================================================
Active Routes:
Network Destination        Netmask          Gateway       Interface  Metric
          0.0.0.0          0.0.0.0      192.168.0.1    192.168.0.151     25
          0.0.0.0          0.0.0.0      172.20.10.1     172.20.10.2     35
```

Student Questions

Lab 4A Hints (Cont)

- WWW.google.com may have different IP addresses on different networks and so trace route to the same numeric address.
- WUSTL VPN rejects all traffic not going to WUSTL.
  So it can not be used as the 2nd interface.
- The new metric assigned by the route command may not be what you specified. So always check using route print.

Student Questions

Lab 4A Hints (Cont)

A. Use "route help" to learn the route command
- **Windows:** route help
- **Linux:** route help
- **MAC:**
  - man netstat
  - man route

B. Ping www.google.com to find its address
  - ping www.google.com

C. Print the new routing table
- **Windows:**
  - route print
- **Linux:**
  - route
- **MAC:**
  - netstat -nr

D. Modify routing tables
- **Windows:**
  - route add/delete/change
- **Linux:**
  - route add/del
- **MAC:**
  - sudo route –nv add

E. Verify using tracert
- **Windows:**
  - tracert
- **Linux:**
  - traceroute
- **MAC:**
  - traceroute

Student Questions

Private Addresses

- Any organization can use these inside their network
  Can't go on the internet. [RFC 1918]
- 10.0.0.0 - 10.255.255.255 (10/8 prefix)
- 172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
- 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

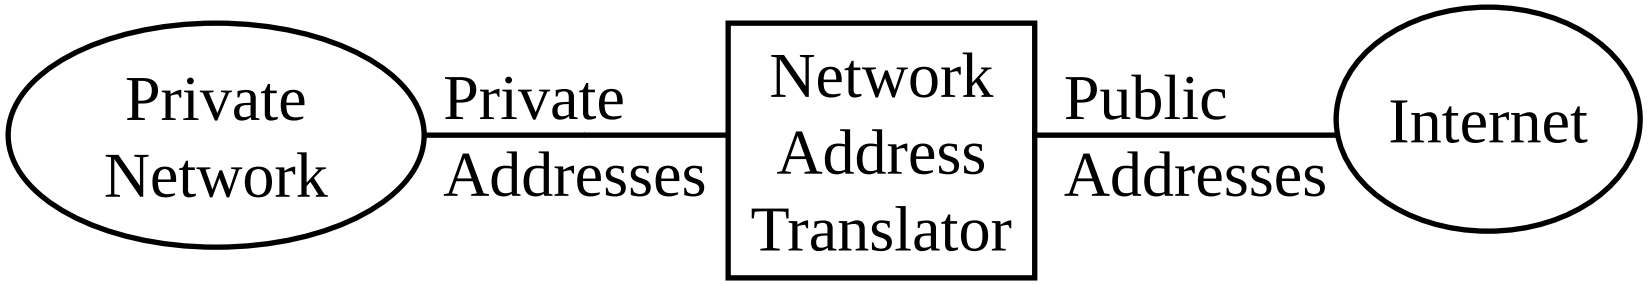

Student Questions

- So does this mean that when we're using VPN service, we get our own private addresses and VPN server acts as network address translator?

Yes, this is one option. Another option is to keep the Internet traffic off the VPN. Only WUSTL traffic going to VPN server.

- What happens if you send a request out of your private network and the destination address specified in the request is a private address?

Both source and destination in public Internet have to be public addresses.

- Can IP address of two private hosts be same, if they share same public address?

Their private addresses need to be different otherwise NAT will have difficulty forwarding the returning responses.

Network Address Translation (NAT)

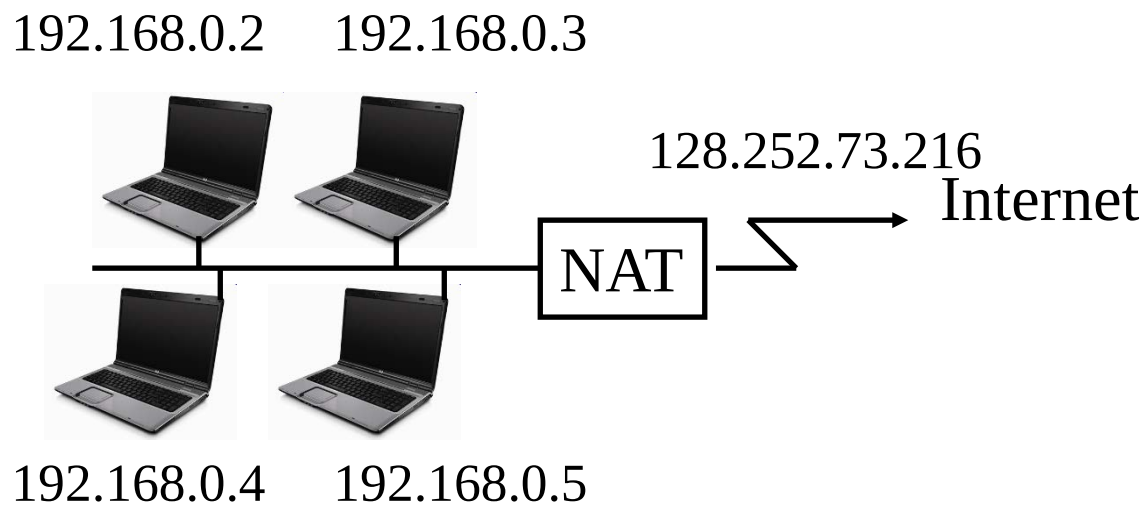

- Private IP addresses 192.168.x.x
- Can be used by anyone inside their networks
- Cannot be used on the public Internet
- NAT overwrites source addresses on all outgoing packets and overwrites destination addresses on all incoming packets
- Only outgoing connections are possible

Student Questions

- Is incoming UDP traffic forwarded differently by NAT?

No

- Does each subnet usually have a DHCP? Does DHCP assign private or public addresses?

Yes, but you can use statically assigned addresses and will not need a DHCP server. DHCP can assign whatever address range you give. Most companies don't have that many public addresses. Some companies do. E.g., WUSTL.

- How do hosts get a more permanent IP address, for example a web server shouldn't be constantly changing IPs.

You can build the address in the server itself. Or ask your DHCP server (router) to assign it a fixed address.

- Can you only have one web server on port 80 within a single private network that you can access from public internet?

Yes. Other servers will have to use other port numbers.

Universal Plug and Play

- NAT needs to be manually programmed to forward external requests
- UPnP allows hosts to request port forwarding
- Both hosts and NAT should be UPnP aware
- Host requests forwarding all port xx messages to it
- NAT returns the public address and the port #.
- Host can then announce the address and port # outside
- Outside hosts can then reach the internal host (server)

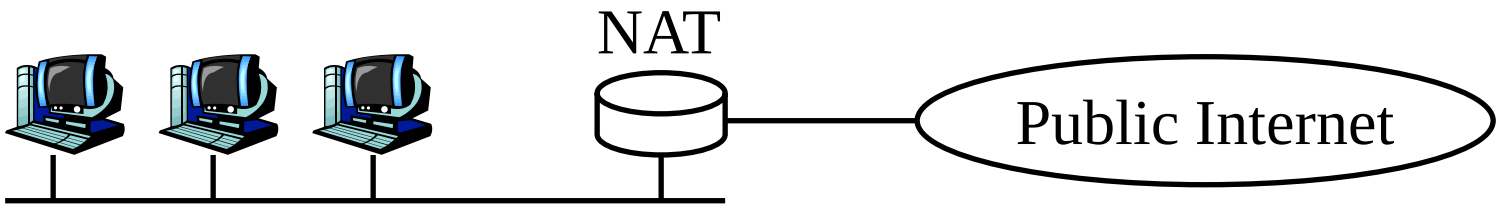

Student Questions

- The ports referred to here are different from the ports of an individual host, correct?

No. These are the TCP ports used by individual hosts.

Homework 4C: NAT

- [20 points] Consider a home network of 3 computers connected to the Internet via a NAT router. Suppose the ISP assigns the router the address 24.34.112.234 and that the network address of the home network is 192.168.1.0/29.
- A. Assign addresses to all interfaces in the home network starting with the lowest possible address.
- B. What is the subnet mask for the home computers?
- C. Suppose each host has two ongoing TCP connections, all to port 80 at host 128.119.40.86. Provide the six corresponding entries in the NAT translation table. Both NAT and computers use source ports starting at 4000.

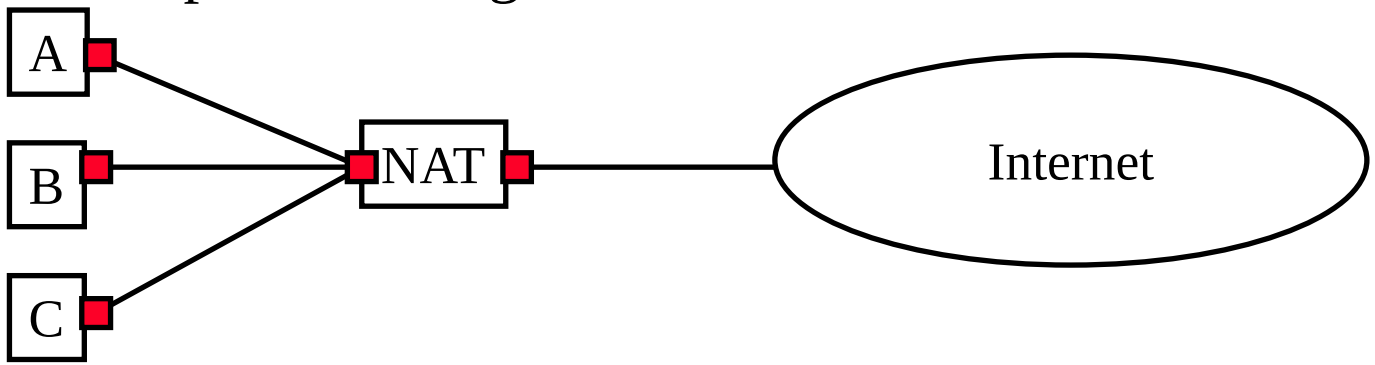

Student Questions

- What exactly is the NAT table translating? is it port number to private IP?

{Private IP- Internal Port} to {Public IP, Public Port} and vice-versa.

DHCP

- **D**ynamic **H**ost **C**ontrol **P**rotocol
- Allows hosts to get an IP address automatically from a server
- Do not need to program each host manually
- Each allocation has a limited “lease” time
- Can reuse a limited number of addresses
- Hosts broadcast “Is there a DHCP Server Here?”
  Sent to 255.255.255.255
- DHCP servers respond
- RFC 2132 defines DHCP options: DHCP Message type option is used to convey the type of the DHCP message. The code for this option is 53, and its length is 1. Legal values for this option are:

Value	Message Type	Value	Message Type
1	DHCP DISCOVER	5	DHCP ACK
2	DHCP OFFER	6	DHCP NAK
3	DHCP REQUEST	7	DHCP RELEASE
4	DHCP DECLINE	8	DHCP INFORM

Ref: https://datatracker.ietf.org/doc/html/rfc2132

Student Questions

- Is there a standard for "lease" time assigned to an allocation?

No, there is no standard. My router has 1 hour lease time. Some routers have 24 hour lease time.

DHCP Example

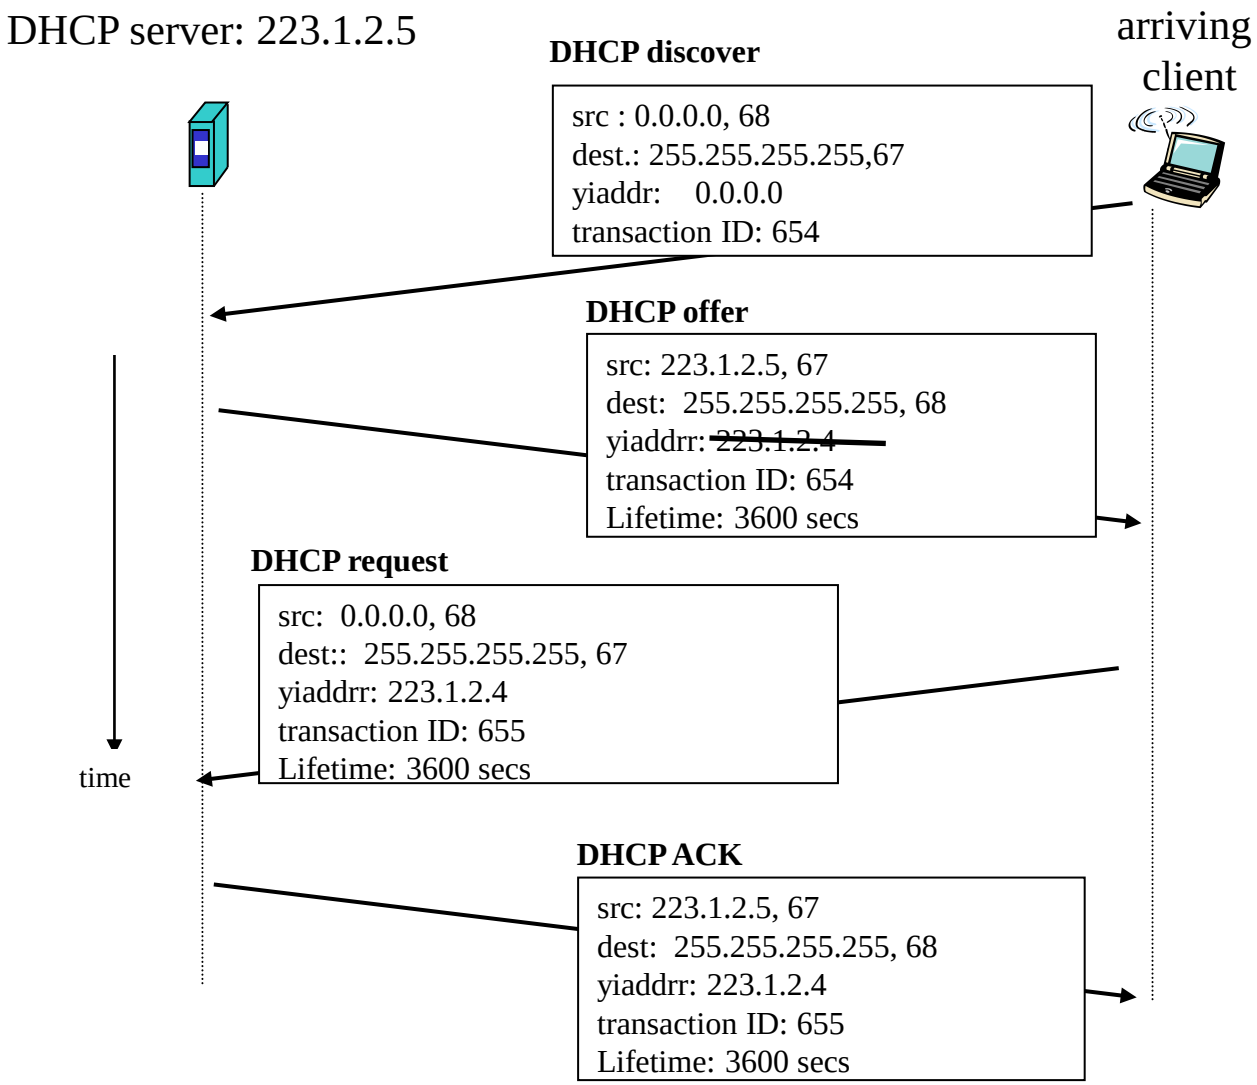

Student Questions

- When the node moves to a new subnet, can it maintain the application's TCP connection?

No. But there are "Mobile IP" extensions that allow this. To be discussed in Chapter 7.

- If the subnet mask is a.b.c/24, Is the maximum number of nodes that the organization can have is 255 end systems? Can a 256th user join, or is the ISP, at this point will request a new subnet mask and reassign all users' IPs?

Yes, the organization is limited to 254 end-nodes. One address will be required for the router. The DHCP server will reject any more address requests. The user will then have to reprogram the router with a new mask. When the lease expires on the existing devices, they will get a new address and a new mask.

- DNS is also related to the IP of a host. How will DHCP and DNS work when a new user connects to a network?

Local addresses are resolved by local DNS, which resides on the same device as DHCP.

Lab 4B: DHCP

- [15 points] Download the Wireshark traces from http://gaia.cs.umass.edu/wireshark-labs/wireshark-traces.zip
- Open *dhcp-ethereal-trace-1* in Wireshark. Select **View → Expand All**. Answer the following questions:
  1. Examine Frame 2 marked DHCP.
     A. What transport protocol and destination port # is used by DHCP?
     B. What are the source and destination IP addresses for this frame and why?
     C. What is the **Code-Length-Type** for the DHCP Discover option?
  2. Examine Frames 4, 5, 6 to find Code-Length-Type for:
     A. DHCP Offer
     B. DHCP Request
     C. DHCP Ack

Student Questions

Lab 4B: DHCP (Cont)

3. Examine Frame 4:
   A. What IP address was assigned by the DHCP server?
   B. What IP address is this frame addressed to and why?
   C. What other information was provided by the DHCP server?
      1. Subnet Mask:
      2. Default Gateway:
      3. DNS1:
      4. DNS2:
      5. Domain Name:
      6. Lease Time:
4. Examine Frame 5 and find what preferred IP address was requested by the client?

Student Questions

IPv6

- Shortage of IPv4 addresses ⇒ Need larger addresses
- IPv6 was designed with 128-bit addresses
- $2^{128} = 3.4\times10^{38}$ addresses
  ⇒ 665×10^{21} addresses per sq. m of earth surface
- If assigned at the rate of 10^6/μs, it would take 20 years
- **Dot-Decimal**: 127.23.45.88
- **Colon-Hex:** FEDC:0000:0000:0000:3243:0000:0000:ABCD
  - Can skip leading zeros of each word
  - Can skip one sequence of zero words, e.g.,
    FEDC::3243:0000:0000:ABCD
    ::3243:0000:0000:ABCD
  - Can leave the last 32 bits in dot-decimal, e.g., ::127.23.45.88
  - Can specify a prefix by /length, e.g., 2345:BA23:0007::/50

Student Questions

- Could you reexplain "::" and skipping a set of zeros

There should be 8 words in the address. If there is a "::" anywhere in the address. You simply put zeros there to the bring the total number of words to 8.

- IPv6 uses only Colon-Hex notation to represent the addresses, then what does Dot-Decimal represent?

Dot-Decimal is used in IPv4.

- Is it allowed to use ":: "multiple times in the address? I feel like that would confuse me.

No. You can not have more than one::

- Can we have two "::s" in the IP address?

No.

- For the notation 2345:BA23:0007::/50, does the "::" here stand for skipped zero? So only 48 bits were explicitly written, and "::" here will be interpreted as two zeros.

Yes.

IPv6

- Shortage of IPv4 addresses ⇒ Need larger addresses
- IPv6 was designed with 128-bit addresses
- $2^{128} = 3.4\times10^{38}$ addresses
  ⇒ 665×10^{21} addresses per sq. m of earth surface
- If assigned at the rate of 10^6/μs, it would take 20 years
- **Dot-Decimal**: 127.23.45.88
- **Colon-Hex:** FEDC:0000:0000:0000:3243:0000:0000:ABCD
  - Can skip leading zeros of each word
  - Can skip one sequence of zero words, e.g.,
    FEDC::3243:0000:0000:ABCD
    ::3243:0000:0000:ABCD
  - Can leave the last 32 bits in dot-decimal, e.g., ::127.23.45.88
  - Can specify a prefix by /length, e.g., 2345:BA23:0007::/50

Student Questions

- Was 128 bit just arbitrarily chosen for IPv6, or were there some criteria behind its selection?

A lot of discussion around how many addresses we need. All values including 64, 256, 128 were discussed. To much vs too little.

- What happened to IPv5? Was version 5 just not accepted?
- *This version was never implemented or used.*
- Is IPv4 more advantageous since it is the default supported protocol by all the routers?

Old is always more prevalent than new. Users need compelling reasons to change.

- Why did IPv5 fail?

IPv6 came in and solved more problems.

- Is the mask size still in decimal?

There is no need for masks. A prefix is sufficient.

- What is Colon-Hex?

Hexadecimal numbers separated by colons.

IPv6 Header

- IPv6:

Version (4b)	Traffic Class (8b)	Flow Label (20b)
Payload Length (16b)	Next Header (8b)	Hop Limit (8b)
Source Address (128b)		
Destination Address (128b)		

- IPv4:

Version	IHL	Type of Service	Total Length
Identification		Flags	Fragment Offset
Time to Live	Protocol	Header Checksum	
Source Address			
Destination Address			
Options			Padding

Student Questions

- Are all of the fields from IPv4 no longer needed for IPv6? There is a lot more space in IPv6 than IPv4 so things could be condensed into one.

All of the things were moved to optional extension headers.

- Why does IPv4 has no hop limit?

In IPv4, it is called Time to Live (TTL).

- Could you explain more about how ipv6 could have multiple headers? What would be in the next header field to represent the next header?

The next header tells precisely what the next header is. It could be another IPv6 extension header or an upper-layer protocol.

- What would go in the "next header" in IPv6? Is this feature similar to fragmenting in IPv4?

Nothing to do with fragmentation. More is coming in the next slides.

- The payload length is 16 bits number means that the payload cannot exceed 2^{16}Bytes=2^6 kB=64 kB?

Yes.

IPv6 Header

- IPv6:

Version (4b)	Traffic Class (8b)	Flow Label (20b)	
Payload Length (16b)		Next Header (8b)	Hop Limit (8b)
Source Address (128b)			
Destination Address (128b)			

- IPv4:

Version	IHL	Type of Service	Total Length
Identification		Flags	Fragment Offset
Time to Live	Protocol	Header Checksum	
Source Address			
Destination Address			
Options			Padding

Student Questions

- Do routers today assign hosts two addresses, one IPv4 and one IPv6, to enable them to support different services ?

Yes. IPv6-enabled routers are generally dual-stacked.

- Does the version for IPv6 only specify the version is either 4 or 6 (the same as IPv4)?

4-bits could be used to indicate 16 versions. However, there are only two in use.

IPv6 vs. IPv4

- 1995 vs. 1975
- IPv6 only twice the size of IPv4 header
- Only version number has same position and meaning as in IPv4
- Removed: header length, type of service, identification, flags, fragment offset, header checksum ⇒ No fragmentation
- Datagram length replaced by payload length
- Protocol type replaced by next header
- Time to live replaced by hop limit
- Added: Priority and flow label
- All fixed size fields.
- No optional fields. Replaced by extension headers.
- 8-bit hop limit = 255 hops max (Limits looping)
- Next Header = 6 (TCP), 17 (UDP)

Student Questions

- Wouldn't removing fragmentation in IPv6 cause major issues with congestion on the network?

Fragmentation is required because some routers have small memory. Congestion is caused if the link or router processing capacity is lower than the load. Compute and storage are different issues.

- Is hop looping possible when using IPv4?

No. TTL is exactly equivalent to Hop Limit. Just a different name.

- If there is no more fragmentation, how is the problem of small router memory handled? will it just change what path is taken so that no low memory routers are encountered?

Every IPv6 router is required to accommodate the largest datagram size.

- Has IPv6 changed in the years since it was created? Are there more IP versions that have been created since then?

Both IPv4 and IPv6 are continuously being enhanced by the IETF. There are no new versions.

IPv6 vs. IPv4

- 1995 vs. 1975
- IPv6 only twice the size of IPv4 header
- Only version number has same position and meaning as in IPv4
- Removed: header length, type of service, identification, flags, fragment offset, header checksum ⇒ No fragmentation
- Datagram length replaced by payload length
- Protocol type replaced by next header
- Time to live replaced by hop limit
- Added: Priority and flow label
- All fixed size fields.
- No optional fields. Replaced by extension headers.
- 8-bit hop limit = 255 hops max (Limits looping)
- Next Header = 6 (TCP), 17 (UDP)

Student Questions

- When I opened Network Sharing through IPV6 on my laptop, I could share the internet with an embedded system through a USB port. Is this happened by using IPV6?

Maybe the embedded system was using IPv6 only. Rare.

- Why is the word "Only" used when explaining that the IPv6 header is twice the size?

As claimed by IPv6 designers to show that it is not bad. Advantages vs. disadvantage tradeoff.

IPv4 to IPv6 Transition

- **Dual Stack**: Each IPv6 router also implements IPv4
  IPv6 is used only if source host, destination host, and all routers on the path are IPv6 aware.
- **Tunneling**:The last IPv6 router puts the entire IPv6 datagram in a new IPv4 datagram addressed to the next IPv6 router
  = **Encapsulation**

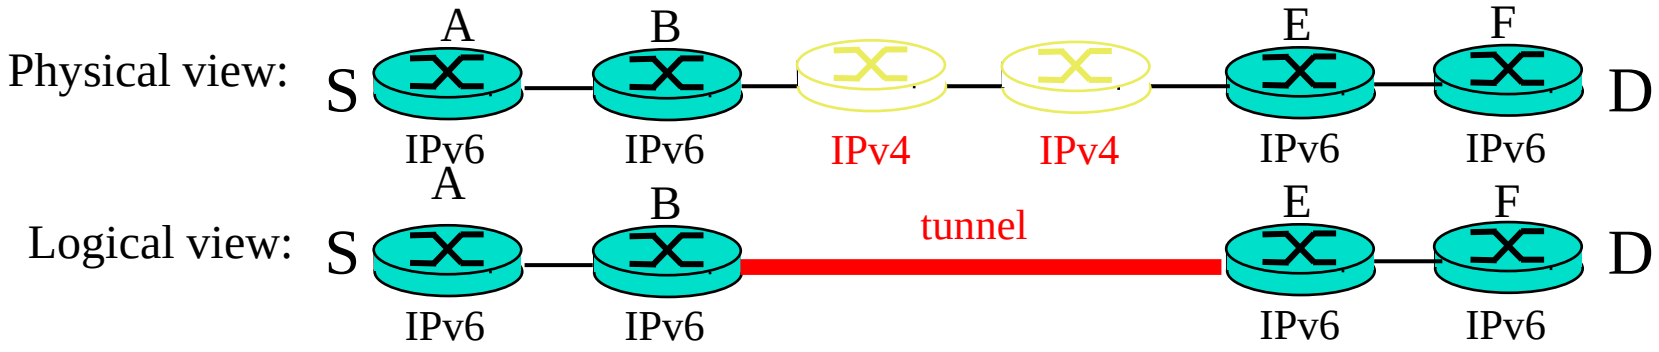

IPv4 From:B to E	IPv6 From:S to D	Data

Student Questions

- Are all new routers required to be dual-stack?

Not required. But, yes, more and more routers are now both IPv4 and IPv6 capable.

- When I display the routing tables on my device, I have two sections: "Internet" and "Internet6". I assume "Internet6" is referring to IPv6, so is this transition currently underway?

Yes.

- Does this mean that the IPv4 header needs have some field to indicate it is encapsulating an IPv6 datagram to allow this datagram to be decapsulated? And how does the last IPv6 router know what's the source IPv4 given that the current datagram only contains IPv6 addresses?

Yes. All of this exists. To be taught in CSE570S: Recent Advances in Networking.

IPv4 to IPv6 Transition

- **Dual Stack**: Each IPv6 router also implements IPv4
  IPv6 is used only if source host, destination host, and all routers on the path are IPv6 aware.
- **Tunneling**:The last IPv6 router puts the entire IPv6 datagram in a new IPv4 datagram addressed to the next IPv6 router
  = **Encapsulation**

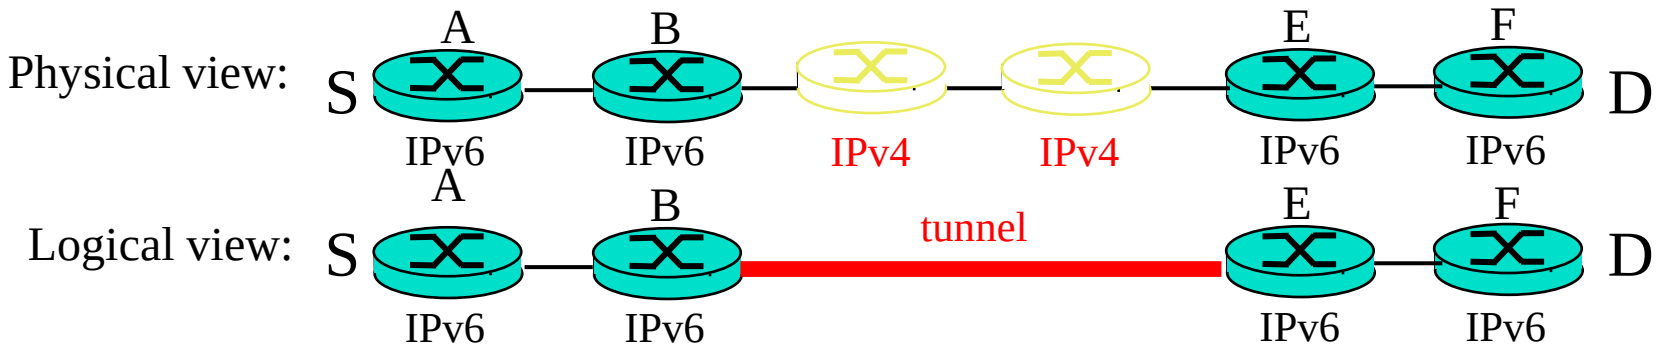

IPv4 From:B to E	IPv6 From:S to D	Data

Student Questions

- How do you know if every router along the path is IPv6 capable if you have not yet traversed the path?

A dual-stacked router knows which neighbor speaks what. Encapsulation takes place only if the shortest path next neighbor speaks a different version.

- If methods of transitioning from IPv4 to IPv6 and visa versa exist, why does the switch to IPv6 need to happen at all?

It is easier if everyone spoke one language. Translation adds overhead.

- Is it worth it to use IPv6 if you're going to have to tunnel anyway? Are there still advantages if everyone doesn't adopt it?

See the answer to the previous question.

- Can you go over tunneling again?

Sure.

- Are there any non-dual-stack routers?

All IPv4 routers are non-dual-stack.

- To be backward compatible, wouldn't there be increasingly more overhead as the technology improves?

Not necessarily.

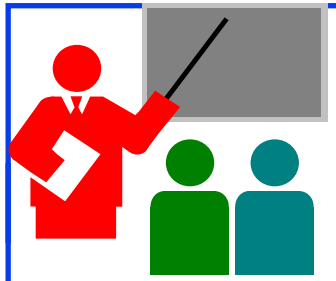

Forwarding Protocols: Review

1. IPv4 uses 32 bit addresses consisting of **subnet + host**
2. **Private addresses** can be reused
   ⇒ Helped solve the address shortage to a great extent
3. **DHCP** is used to automatically allocate addresses to hosts
4. IPv6 uses **128 bit addresses**. Requires dual stack or **tunneling** to coexist with IPv4.

Student Questions

- Are there any other advantages to IPV6, besides increased address space?

Yes. There are. To be discussed in more advanced classes.

- Will IPv6 be updated to any new next version?

The same versions are being updated continuously.

Ref: Read Section 4.3 of the textbook. Try R17 through R29.

Generalized Forwarding and SDN

- Planes of Networking
- Data vs. Control Logic
- OpenFlow Protocol

Student Questions

- How many regions are currently supporting IPv6?

Rest of the world (ROW).

Planes of Networking

- **Data Plane**: All activities involving as well as resulting from data packets sent by the end user, e.g.,
  - Forwarding
  - Fragmentation and reassembly
  - Replication for multicasting
- **Control Plane**: All activities that are necessary to perform data plane activities but do not involve end-user data packets
  - Making routing tables
  - Setting packet handling policies (e.g., security)
  - Base station beacons announcing availability of services

Ref: Open Data Center Alliance Usage Model: Software Defined Networking Rev 1.0,"
http://www.opendatacenteralliance.org/docs/Software_Defined_Networking_Master_Usage_Model_Rev1.0.pdf

Student Questions

- Data plane is in the router, right?

Yes.

Planes of Networking (Cont)

- **Management Plane**: All activities related to provisioning and monitoring of the networks
  - Fault, Configuration, Accounting, Performance and Security (**FCAPS**).
  - Instantiate new devices and protocols (Turn devices on/off)
  - Optional ⇒ May be handled manually for small networks.
- **Services Plane**: Middlebox services to improve performance or security, e.g.,
  - Load Balancers, Proxy Service, Intrusion Detection, Firewalls, SSL Off-loaders
  - Optional ⇒ Not required for small networks

Student Questions

- What do you mean by counters that the management plane keeps track of?

For example, the number of packets received, dropped, or found in error.

Data vs. Control Logic

- Data plane runs at line rate,
  e.g., 100 Gbps for 100 Gbps Ethernet ⇒ Fast Path
  ⇒ Typically implemented using special hardware,
  e.g., Ternary Content Addressable Memories (TCAMs)
- Some exceptional data plane activities are handled by the CPU in the switch ⇒ Slow path
  e.g., Broadcast, Unknown, and Multicast (BUM) traffic
- All control activities are generally handled by CPU

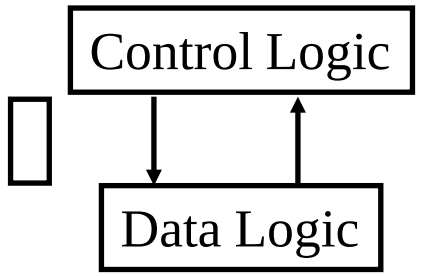

Student Questions

OpenFlow: Key Ideas

1. Separation of control and data planes
2. Centralization of control
3. Flow based control

Student Questions

- Who were the major entities behind OpenFlow?

OpenFlow originated from the Ph.D. thesis of Martin Casado under Prof. Nick McKeown of Stanford University.

Ref: N. McKeown, et al., ``OpenFlow: Enabling Innovation in Campus Networks," ACM SIGCOMM CCR, Vol. 38, No. 2, April 2008, pp. 69-74.

Separation of Control and Data Plane

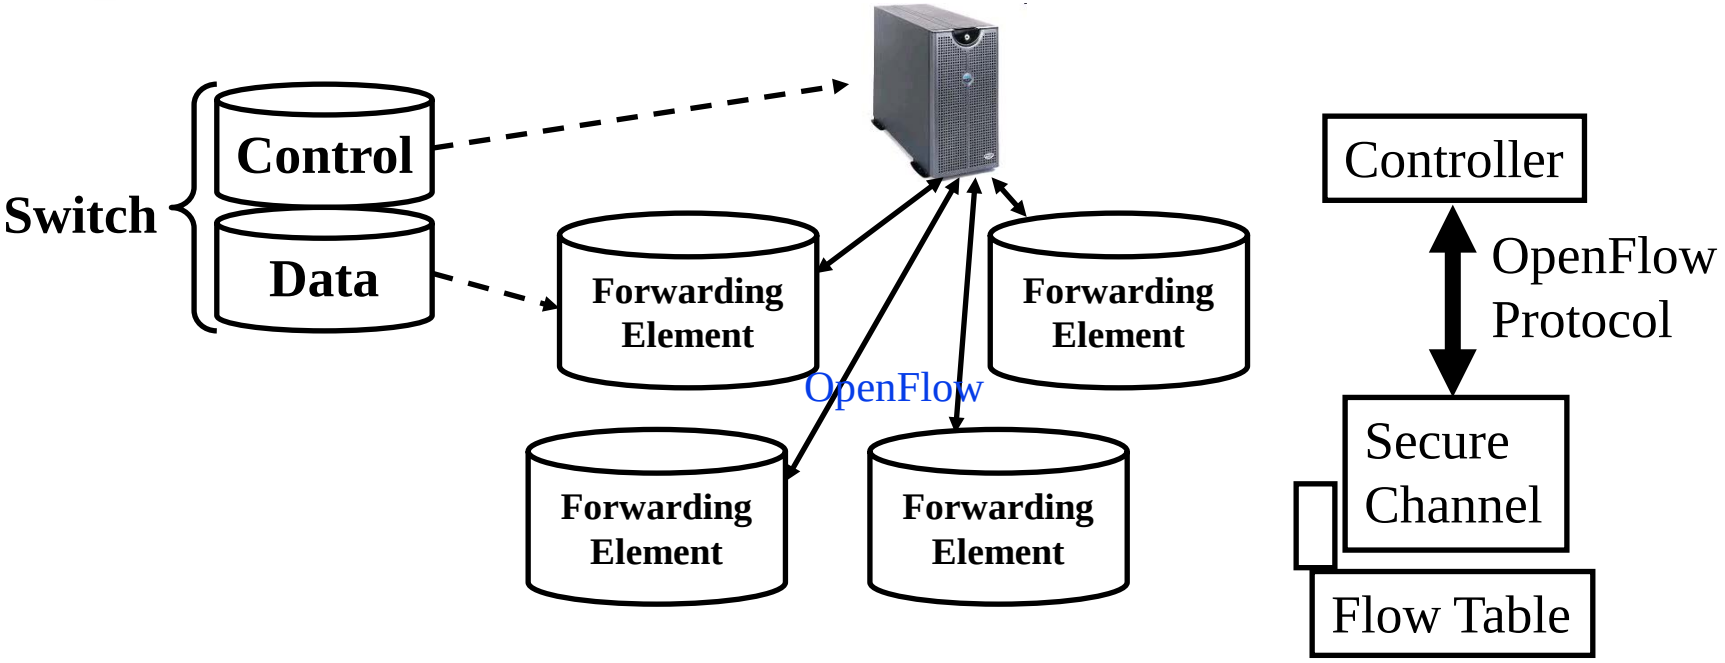

- Control logic is moved to a controller
- Switches only have forwarding elements
- One expensive controller with a lot of cheap switches
- OpenFlow is the protocol to send/receive forwarding rules from controller to switches

Student Questions

OpenFlow V1.0

- On packet arrival, match the header fields with flow entries in a table, if any entry matches, perform indicated actions, and update the counters indicated in that entry

Flow Table:

Header Fields	Actions	Counters
Header Fields	Actions	Counters
…	…	…
Header Fields	Actions	Counters

Ingress Port	Ether Source	Ether Dest	VLAN ID	VLAN Priority	IP Src	IP Dst	IP Proto	IP ToS	Src L4 Port	Dst L4 Port

Student Questions

- Are most routers using OpenFlow protocol to control the traffic these days?

No.

Ref: http://archive.openflow.org/documents/openflow-spec-v1.0.0.pdf

Flow Table Example

Port	Src MAC	Dst MAC	VLAN ID	Priority	EtherType	Src IP	Dst IP	IP Proto	IP ToS	Src L4 Port ICMP Type	Dst L4 Port ICMP Code	Action	Counter
*	*	0A:C8:*	*	*	*	*	*	*	*	*	*	Port 1	102
*	*	*	*	*	*	*	192.168.*.*	*	*	*	*	Port 2	202
*	*	*	*	*	*	*	*	*	*	21	21	Drop	420
*	*	*	*	*	*	*	*	0x806	*	*	*	Local	444
*	*	*	*	*	*	*	*	0x1*	*	*	*	Controller	1

- Idle timeout: Remove entry if no packets received for this time
- Hard timeout: Remove entry after this time
- If both are set, the entry is removed if either one expires.

Ref: S. Azodolmolky, "Software Defined Networking with OpenFlow," Packt Publishing, October 2013, 152 pp., ISBN:978-1-84969-872-6 (Safari Book)

Student Questions

- Do the table entries actually use glob-style expressions?

No. Glob is for ascii strings. Most of these are binary strings. So marking and matching is common.

Are these counter fields denoted by the counter value (like an ID) or is the counter value the actual value being passed back of these instances?

Counters are actual counts of those rows being matched and those actions taken.

- How the Counter is incremented. Can you give some examples?

N←N+1

Here, N is the number of packets received.

Matching

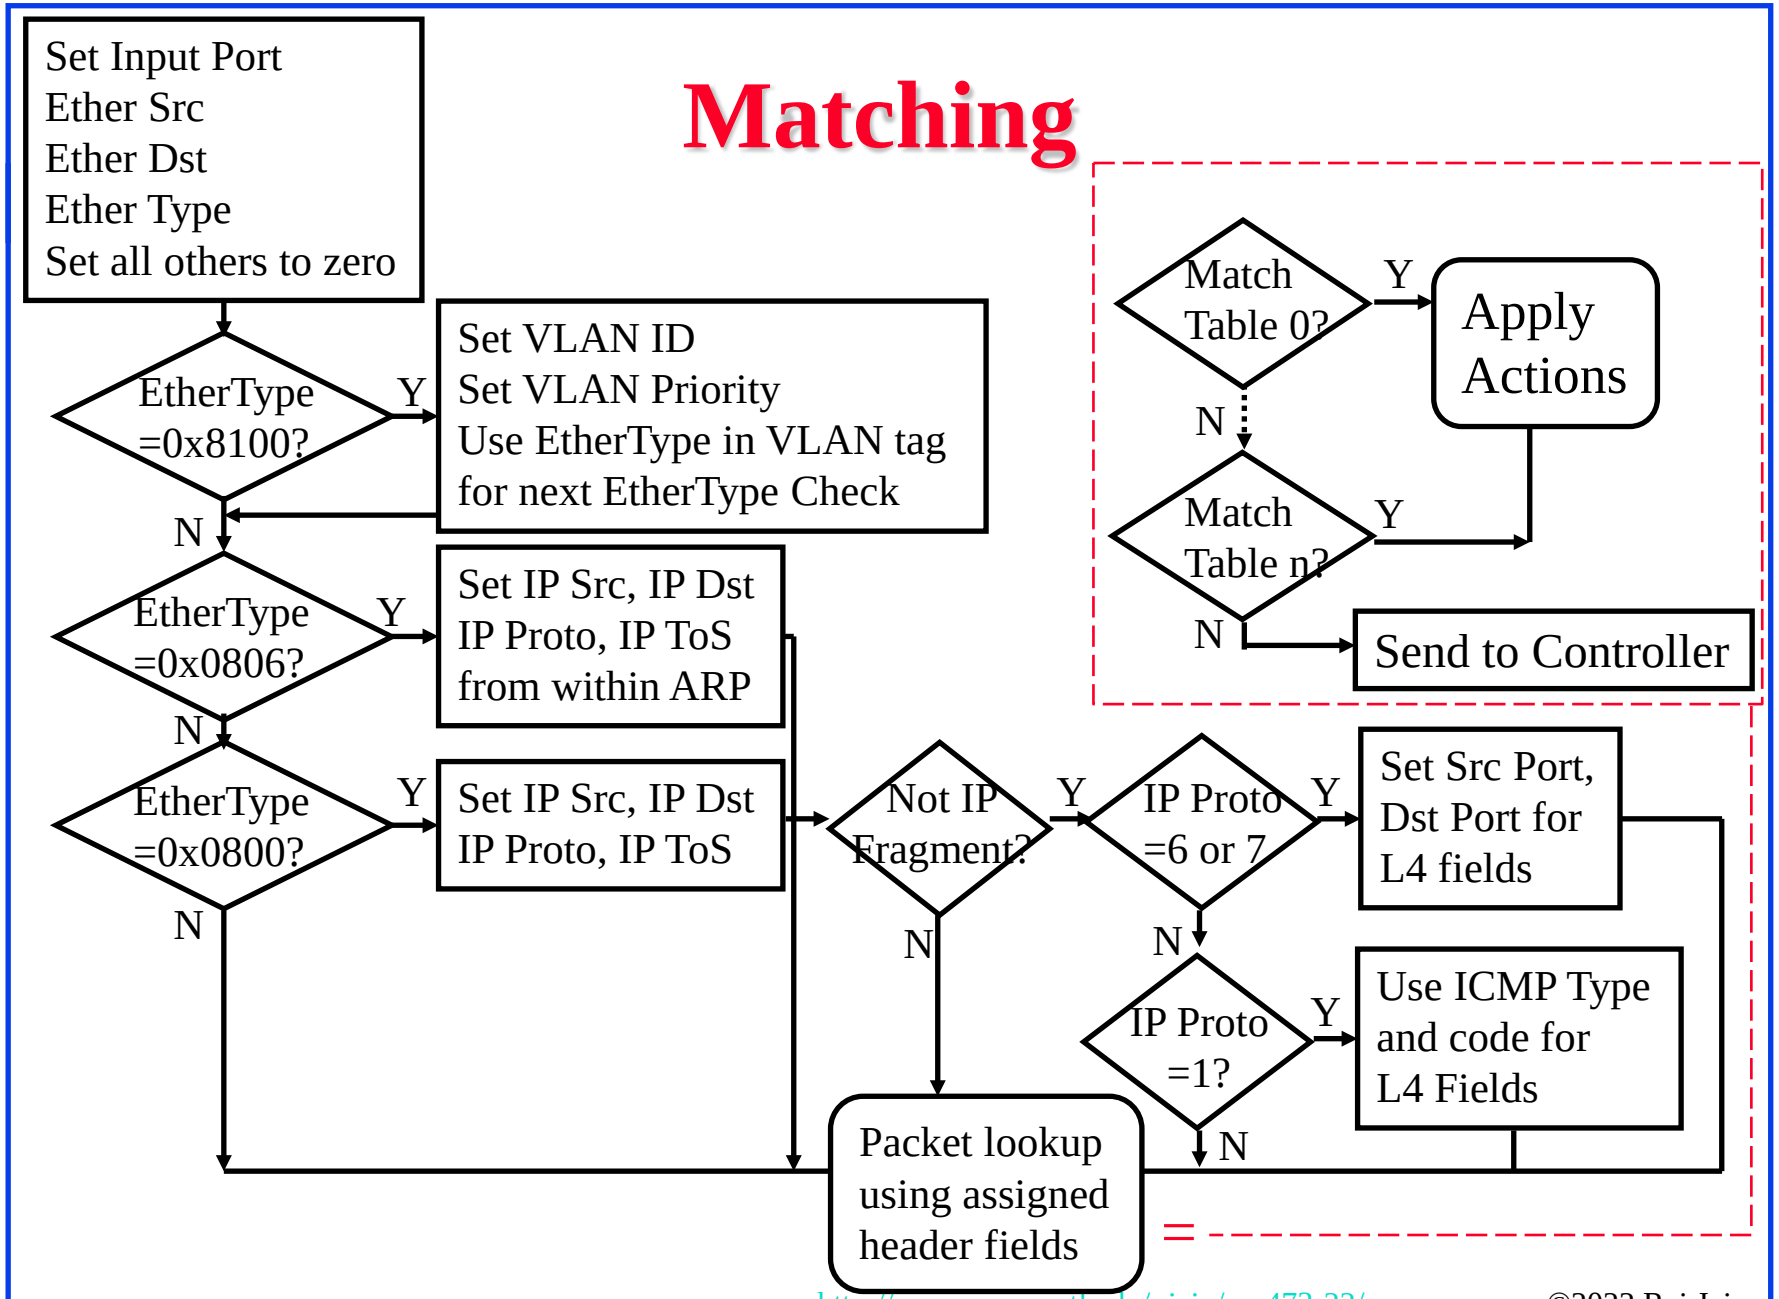

Student Questions

- To clarify, are only the fields necessary for the EtherType command set and others left blank?

No. The top box indicates fields that are used in the left 3 decision boxes.

- Do we need this chart for exam 2?

YES

- The upper right corner doesn't seem to be connected to the rest, how can we get there?

This is the detail of the bottom box.

Counters

Per Table	Per Flow	Per Port	Per Queue
Active Entries	Received Packets	Received Packets	Transmit Packets
Packet Lookups	Received Bytes	Transmitted Packets	Transmit Bytes
Packet Matches	Duration (Secs)	Received Bytes	Transmit overrun errors
	Duration (nanosecs)	Transmitted Bytes	
		Receive Drops	
		Transmit Drops	
		Receive Errors	
		Transmit Errors	
		Receive Frame Alignment Errors	
		Receive Overrun erorrs	
		Receive CRC Errors	
		Collisions	

Student Questions

- What is a switch here?

A switch is a layer 2 device, e.g., Ethernet Switch or Wi-Fi Access Point. It speaks Layer 2 protocol only.

Actions

- Forward to Physical/**Virtual Port** *i*
- Enqueue: To a particular **queue** in the port ⇒ QoS
- Drop
- Modify Field: E.g., add/remove VLAN tags, ToS bits, Change TTL
- Masking allows matching only selected fields, e.g., Dest. IP, Dest. MAC, etc.
- If header matches an entry, corresponding actions are performed and counters are updated
- If no header match, the packet is queued and the **header is sent to the controller**, which sends a new rule. Subsequent packets of the flow are handled by this rule.
- Secure Channel: Between controller and the switch using TLS

Student Questions

- Were there ever attacks on OpenFlow networks by generating and sending lots of distinct packets with distinct headers to force queries of the controller?

No. Even if these were to happen, these can easily be overcome by rate control.

- Would you elaborate on the TLS mechanism?

Transport layer security (TLS) will be discussed in Chapter 8.

- What is the benefit of using virtual ports?

Virtualization allows one device to act as many devices.

- What is QoS?

Quality of Service.

- How does the controller decide which port the packet is sent to?

It prepares a routing table for the whole network.

Actions (Cont)

- Modern switches already implement flow tables, typically using Ternary Content Addressable Memories (TCAMs)
- Controller can change the forwarding rules if a client moves
  $\Rightarrow$ Packets for mobile clients are forwarded correctly
- Controller can send flow table entries beforehand (**Proactive**) or Send on demand (**Reactive**). OpenFlow allows both models.

Student Questions

- What does "client moves" mean?

The user goes from one room to the next and so is now connected to another switch or Wi-Fi access point.

SDN Data Plane: Summary

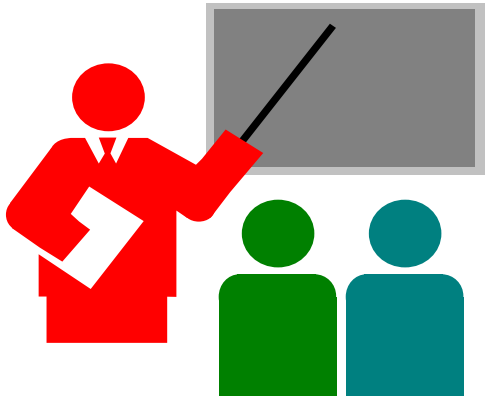

1. **Data plane** consists of packets sent by the users
2. OpenFLow separates data plane from the **control plane** and centralizes the control plane
3. The **controller** makes rules for forwarding and sends to switches
4. Switches match the rules and take specified actions

Student Questions

- Unsure but the second to last question may have had the wrong answer selected.

In OpenFlow, forwarding decisions are made by matching flow table entries with packet headers.

- Are duties of the control plane ever carried out by end systems?

End systems also have control plane and data plane division applies to all systems.

- Will new grade rankings from the exam be released?

Sure.

- How is OpenFlow implemented? Since it is a part of SDN, is it a downloadable program?

Open-source implementations are available.

Network Layer Data Plane: Summary

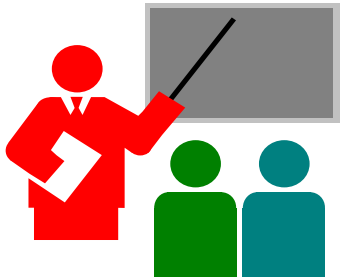

1. **Forwarding** consists of matching the destination address to a list of entries in a table. **Routing** consists of making that table.
2. IP is a forwarding protocol. IPv4 uses 32 bit addresses in **dot-decimal notation**. IPv6 uses 128 bit addresses in **Hex-Colon notation**.
3. **DHCP** is used to assign addresses dynamically.
4. **Private addresses** are used inside an enterprise network. **NAT** allows a single public address to be used by many internal hosts with private addresses.
5. **OpenFlow** separates data plane from control plane and centralizes the control plane

Student Questions

Acronyms

- ACK Acknowledgement
- ACM Automatic Computing Machinery
- AQM Active Queue Management
- ARP Address Resolution Protocol
- ATM Asynchronous Transfer Mode
- BGP Border Gateway Protocol
- BUM Broadcast, Unknown, and Multicast
- CAMs Content Addressable Memories
- CBR Constant bit rate
- CCR Computer Communications Review
- CIDR Classless Inter-Domain Routing
- CPU Central Processing Unit
- DHCP Dynamic Host Control Protocol
- DNS Domain Name Service
- FCAPS Fault, Configuration, Accounting, Performance and Security
- FCFS First Come First Served

Student Questions

Acronyms (Cont)

FTP	File Transfer Protocol
GFR	Guaranteed Frame Rate
HTTP	Hyper-Text Transfer Protocol
ICMP	IP Control Message Protocol
ID	Identifier
IP	Inter-Network Protocol
IPv4	IP Version 4
IPv6	IP Version 6
ISP	Internet Service Provider
KISS	Keep it simple stupid
LAN	Local Area Network
MAC	Media Access Control
MS	Microsoft
MTU	Maximum Transmission Unit
NAT	Network Address Translation
PBX	Private Branch Exchange

Student Questions

Acronyms (Cont)

PHY	Physical Layer
QoS	Quality of Service
RED	Random Early Drop
RFC	Request for Comment
RIP	Routing Information Protocol
RTT	Round Trip Time
SDN	Software Defined Networking
SMTP	Simple Mail Transfer Protocol
SSL	Secure Socket Layer
TCAM	Ternary Content Addressable Memory
TCP	Transmission Control Protocol
TLS	Transport Level Security
ToS	Type of Service
TTL	Time to live
UBR	Unspecified bit rate
UPnP	Universal Plug and Play

Student Questions

Acronyms (Cont)

- VBR Variable bit rate
- VCI Virtual Circuit Identifiers
- VLAN Virtual Local Area Network
- VPN Virtual Private Network
- WAN Wide Area Network
- WiFi Wireless Fidelity

Student Questions

Scan This to Download These Slides

Raj Jain
http://rajjain.com

http://www.cse.wustl.edu/~jain/cse473-22/i_4nld.htm

Student Questions

Related Modules

CSE 567: The Art of Computer Systems Performance Analysis
https://www.youtube.com/playlist?list=PLjGG94etKypJEKjNAa1n_1X0bWWNyZcof

CSE473S: Introduction to Computer Networks (Fall 2011),
https://www.youtube.com/playlist?list=PLjGG94etKypJWOSPMh8Azcgy5e_10TiDw

CSE 570: Recent Advances in Networking (Spring 2013)
https://www.youtube.com/playlist?list=PLjGG94etKypLHyBN8mOgwJLHD2FFIMGq5

CSE571S: Network Security (Spring 2011),
https://www.youtube.com/playlist?list=PLjGG94etKypKvzfVtutHcPFJXumyyg93u

Video Podcasts of Prof. Raj Jain's Lectures,
https://www.youtube.com/channel/UCN4-5wzNP9-ruOzQMs-8NUw

Student Questions