CSCI1680
Network Layer: IP & Forwarding

Nick DeMarinis

Based partly on lecture notes by Rodrigo Fonseca, David Mazières, Phil Levis, John Jannotti

Administivia

- Snowcast: was due last night
- IP Project: Out later today/tomorrow (Sep 30)
  - **Fill out group preference form by 11:59pm tomorrow (Sep 30)**
- **HW2: Announcement soon**
  - **Stuff we've covered + warmup for IP!**

Today

Start of network layer

- Network layer: Internet Protocol (IP) (v4)
- Mechanics of IP forwarding
- Intro to IP project

LAST LECTURE'S PLQ

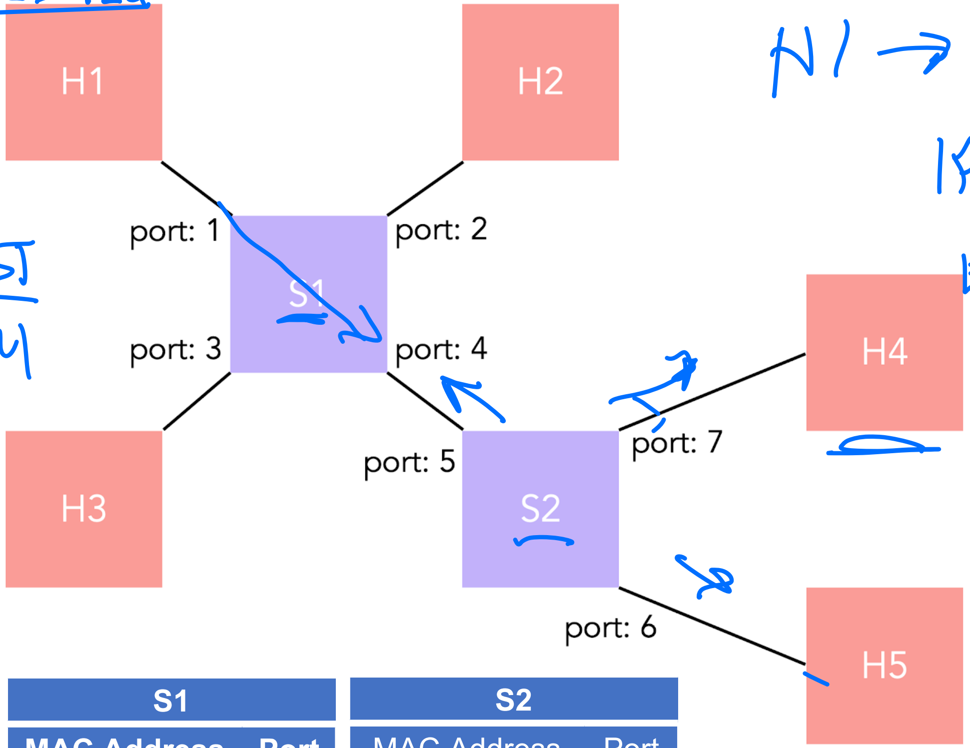

SRC
H1

DST
H4

H1 → H4

IF IN TABLE
↳ DIRECT
ELSE
SEND TO ALL

S1	
MAC Address	Port
H1	1
H4	4

S2	
MAC Address	Port
H1	5

Bridges and Extended LANs

- Single Ethernet collision domain has limitations
  - Limits performance, distance, …
- Next step: separate collision domains with bridges
  - Operates on Ethernet addresses
  - Forwards packets from one collision domain to others
- Modern ethernet uses switches: all hosts directly connected to a bridge

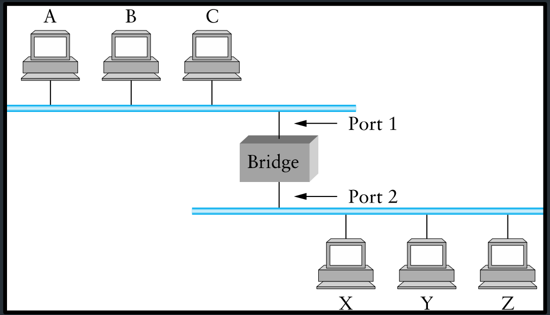

MORE SWITCHING CHALLENGES

- DEALING W/ LOOPS
- LINE-RATE FORWARDING...
- MORE LATER!

Layers, Services, Protocols

Application — Service: user-facing application.
Application-defined messages

Transport — Service: multiplexing applications
Reliable byte stream to other node (TCP),
Unreliable datagram (UDP)

Network — Service: move packets to any other node in the network
Internet Protocol (IP)

L2 **Link** — Service: move frames to other node across link.
May add reliability, medium access control

Physical — Service: move bits to other node across link

Internet Protocol (IP) Goals

How to connect everyone?

- Glue lower-level networks together
- A network of networks!

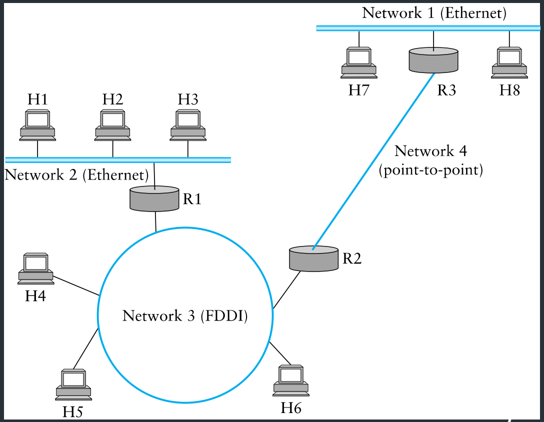

Internet Protocol (IP) Goals

How to connect everyone?

- Glue lower-level networks together
- A network of networks!
- Router: device that forwards packets between networks

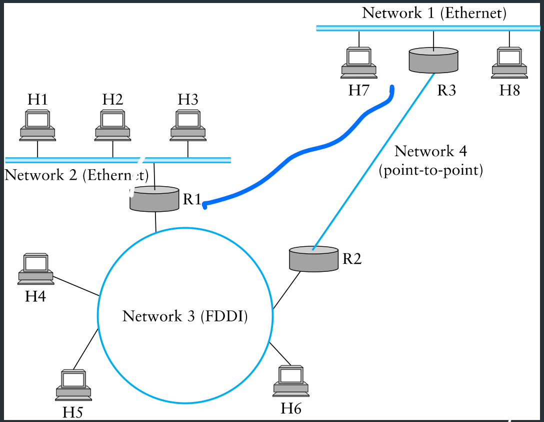

Internet Protocol (IP) Goals

How to connect everyone?

- Glue lower-level networks together
- A network of networks!
- Router: device that forwards packets between networks

Doesn't this sound like switching?

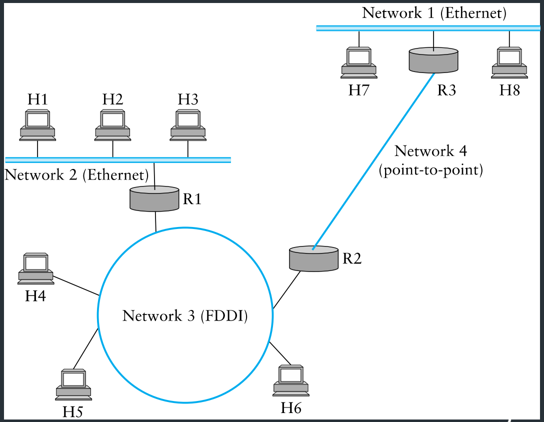

Inter-networking Challenges

- Networks are heterogeneous (eg. Wifi vs. Ethernet)
  - Different frame formats
  - Different service models
  - Different packet sizes/bandwidths

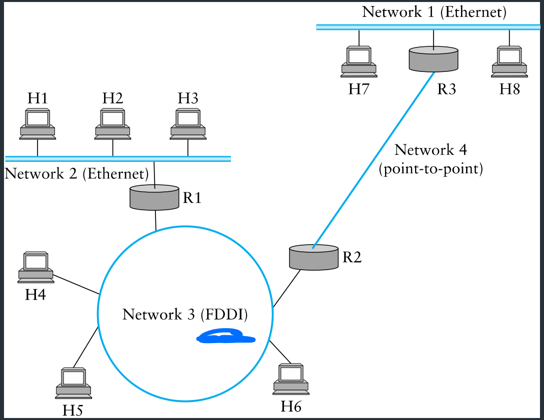

Inter-networking Challenges

- Networks are heterogeneous (eg. Wifi vs. Ethernet)
  - Different frame formats
  - Different service models
  - Different packet sizes/bandwidths
- Scaling
  - Link-layer forwarding strategies don't scale to Internet!

LAN < 1000 HOSTS

INTERNET ≈ 50B

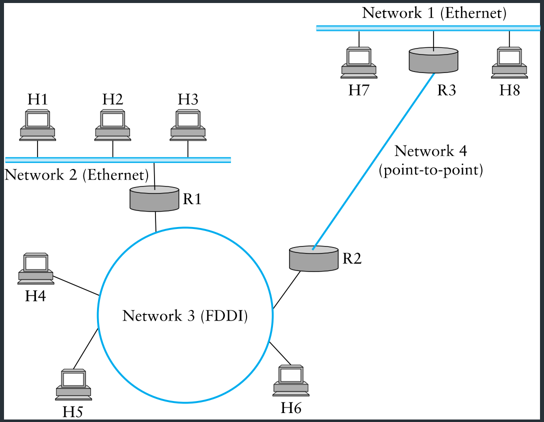

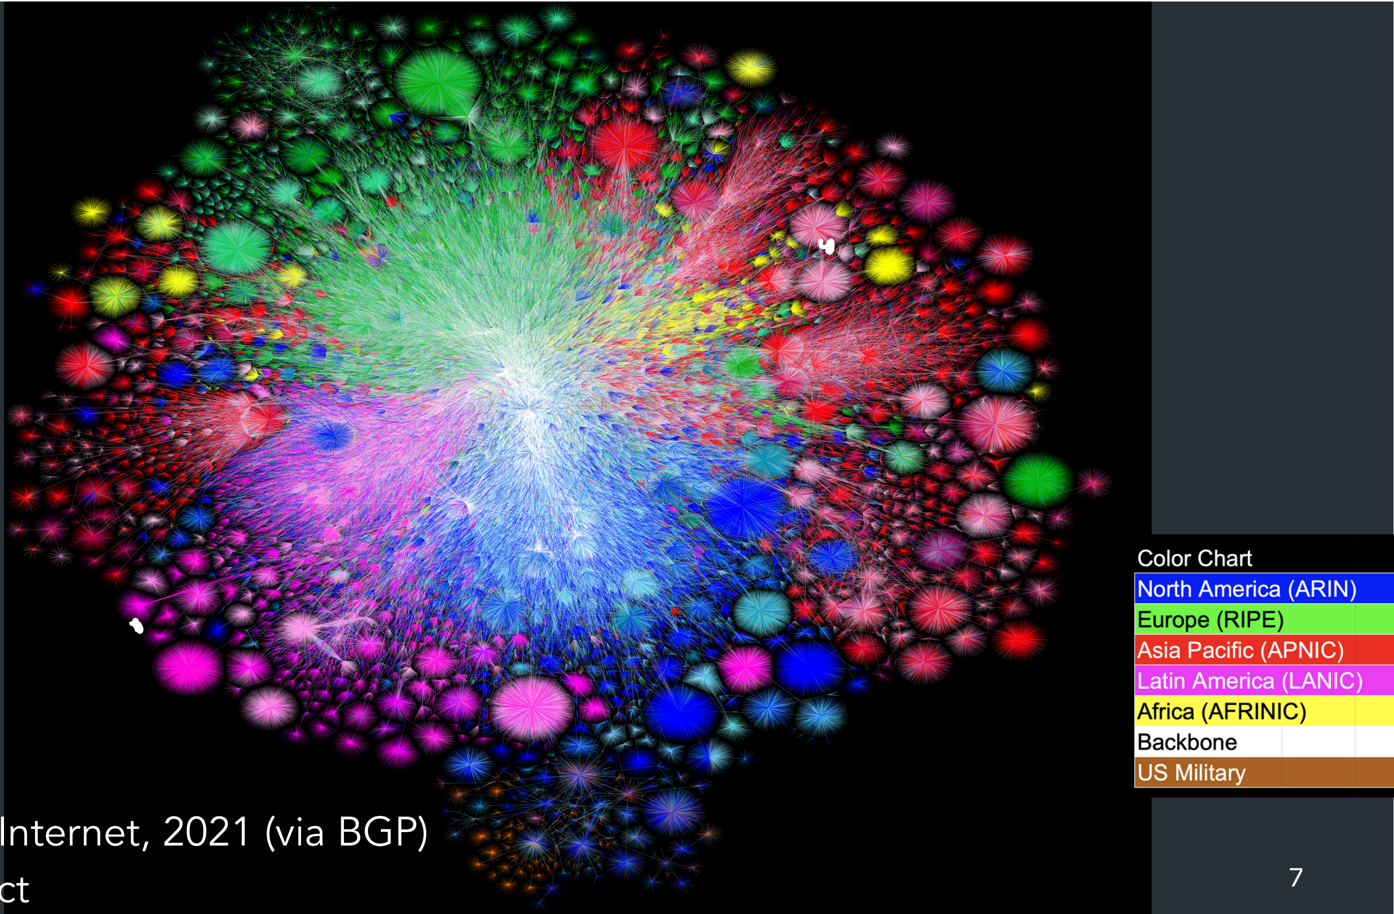

Map of the Internet, 2021 (via BGP)
OPTE project

A Bit of History

- Packet switched networks: Arpanet's IMPs
  - Late 1960's
  - RFC 1, 1969!
  - Segmentation, framing, routing, reliability, reassembly, primitive flow control
- Network Control Program (NCP)
  - Provided connections, flow control
  - Assumed reliable network: IMPs
  - Used by programs like telnet, mail, file transfer

A Bit of History

- Packet switched networks: Arpanet's IMPs
  - Late 1960's
  - RFC 1, 1969!
  - Segmentation, framing, routing, reliability, reassembly, primitive flow control
- Network Control Program (NCP)
  - Provided connections, flow control
  - Assumed reliable network: IMPs
  - Used by programs like telnet, mail, file transfer
- Wanted to connect multiple networks
  - Not all reliable, different formats, etc…

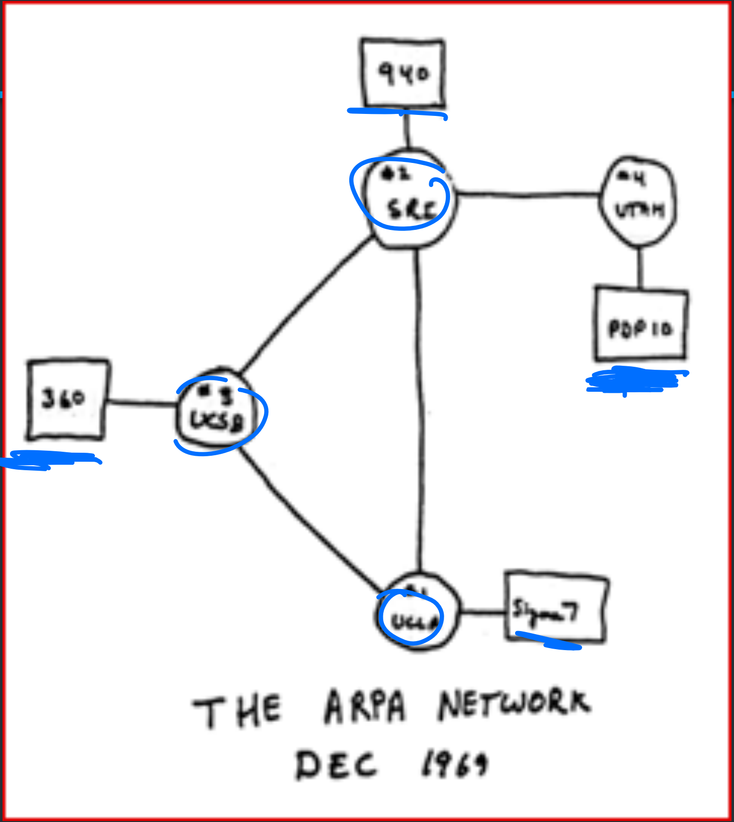
940
#2
SRI
#4
UTAH
PDP 10
360
#3
UCSB
#1
UCLA
Sigma 7
THE ARPA NETWORK
DEC 1969

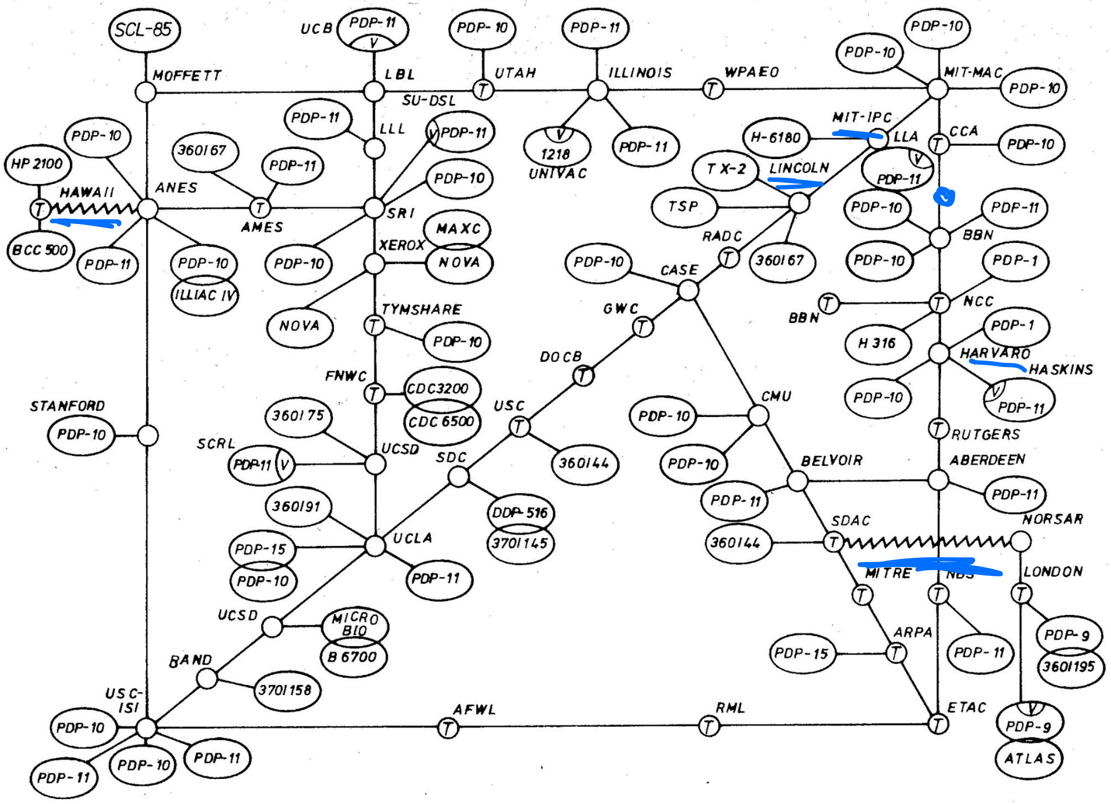

Abb. 4 ARPA NETwork, topologische Karte. Stand Juni 1974.

How would you design such a protocol?

- Circuits or packets?
  - Predictability

CIRCUIT SWITCHING

- NODES DECIDE PATH BEFOREHAND

PACKET SWITCHING - LESS STATE.

- BREAK MESSAGES INTO PACKETS, EACH NODE DECIDES WHAT TO DO

How would you design such a protocol?

- Circuits or packets?
  - Predictability
- What service model?
  - Reliability, timing, bandwidth guarantees

How would you design such a protocol?

- Circuits or packets?
  - Predictability
- What service model?
  - Reliability, timing, bandwidth guarantees
- How to enable connectivity?
  - How do you find a particular host?
  - How do you get a message there?
  - What happens when a host joins/leaves?

1974: TCP/IP Introduced

- Vint Cerf, Robert Kahn build protocol to replace NCP
- Initial design: single protocol providing a unified reliable pipe

1974: TCP/IP Introduced

- Vint Cerf, Robert Kahn build protocol to replace NCP
- Initial design: single protocol providing a unified reliable pipe
- Different requirements soon emerged, and the two were separated
  - IP: basic datagram service among hosts
  - TCP: reliable transport
  - UDP: unreliable multiplexed datagram service

TERM FOR "MESSAGE" AT NETWORK LAYER

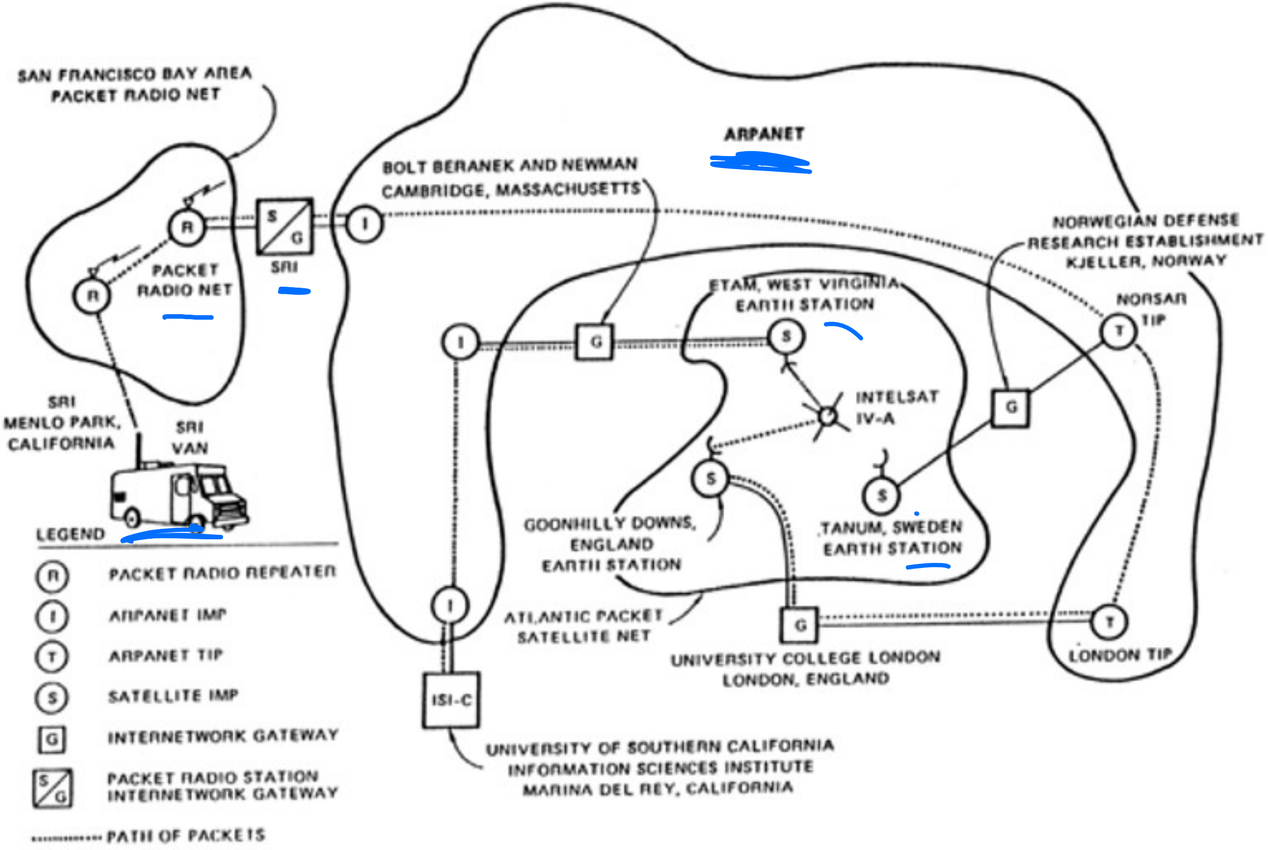

SAN FRANCISCO BAY AREA
PACKET RADIO NET
PACKET
RADIO NET
SRI
BOLT BERANEK AND NEWMAN
CAMBRIDGE, MASSACHUSETTS
ARPANET
NORWEGIAN DEFENSE
RESEARCH ESTABLISHMENT
KJELLER, NORWAY
NORSAR
TIP
ETAM, WEST VIRGINIA
EARTH STATION
INTELSAT
IV-A
SRI
MENLO PARK,
CALIFORNIA
SRI
VAN
GOONHILLY DOWNS,
ENGLAND
EARTH STATION
TANUM, SWEDEN
EARTH STATION
ATLANTIC PACKET
SATELLITE NET
UNIVERSITY COLLEGE LONDON
LONDON, ENGLAND
LONDON TIP
ISI-C
UNIVERSITY OF SOUTHERN CALIFORNIA
INFORMATION SCIENCES INSTITUTE
MARINA DEL REY, CALIFORNIA
LEGEND
PACKET RADIO REPEATER
ARPANET IMP
ARPANET TIP
SATELLITE IMP
INTERNETWORK GATEWAY
PACKET RADIO STATION
INTERNETWORK GATEWAY
PATH OF PACKETS

IP’s Decisions

IP's Decisions

- Circuits or packets? Packets.

  ↳ ROUTERS DON'T NEED TO REMEMBER PER CONNECTION STATE.

IP's Decisions

- Circuits or packets? Packets.
- Service model?
  - Lowest common denominator: best effort, connectionless datagram

IP's Decisions

- Circuits or packets? Packets.
- Service model?
  - Lowest common denominator: best effort, connectionless datagram
- Enabling connectivity?
  - IP header: common message format
  - IP address: each host has an address, based on hierarchical structure of network

An excellent read

David D. Clark, "The design Philosophy of the DARPA Internet Protocols", 1988

- Primary goal: multiplexed utilization of existing interconnected networks
- Other goals:
  - Communication continues despite loss of networks or gateways
  - Support a variety of communication services
  - Accommodate a variety of networks
  - Permit distributed management of its resources
  - Be cost effective
  - Low effort for host attachment
  - Resources must be accountable

Internet Protocol

IP runs on all hosts and routers

- Provides addressing: how we name nodes in an IP network
- Provides forwarding: how routers move packets based on the destination address
- Later: routing: how routers build forwarding rules

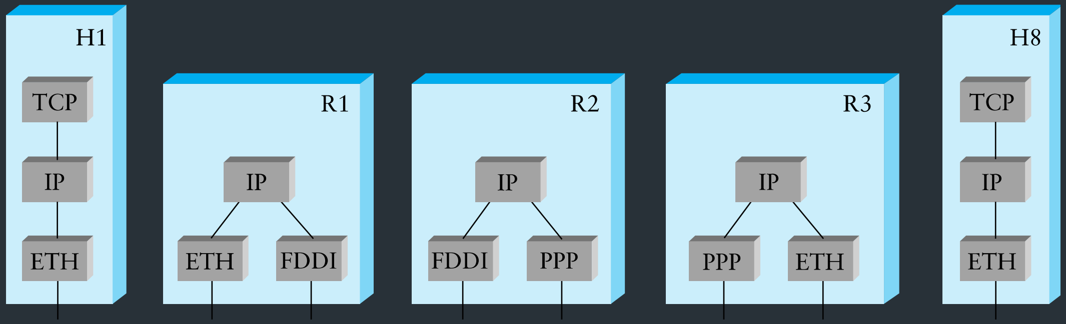

IP's Service Model

- Connectionless (datagram-based)
- Best-effort delivery (unreliable service)
  - packets may be lost
  - packets may be delivered out of order
  - duplicate copies of packets may be delivered
  - packets may be delayed for a long time
- It's the lowest common denominator
  - A network that delivers no packets fits the bill!
  - All these can be dealt with above IP (if probability of delivery is non-zero…)

- CAN HANDLE NETWORKS OF DIFF TYPES.
- ROUTERS CAN BE SIMPLE ⇒ SCALE

IP Addressing

128.148.16.7

IP Version 4: Each address is a 32-bit number:

128.148.16.7

10000000 10010100 00010000 00000111

IP Addressing

IPv6: 128 BITS

128.148.16.7

IP Version 4: Each address is a 32-bit number:

128.148.16.7

OCTET

THINK ABOUT THIS LIKE IT REFERS TO ONE HOST.

10000000 10010100 00010000 00000111

2^{32} POSSIBLE ADDRESSES

$\approx$ 4 BILLION POSSIBLE IPs

<u>Notation</u>

- Write each byte ("octet") as a decimal number
- This is called "dotted decimal" or "dotted quad" notation

IP Addressing

An IP address identifies…

HOST

- Who a host is: A unique number #1242
- Where it is on the Internet

→ BROWN NETWORK

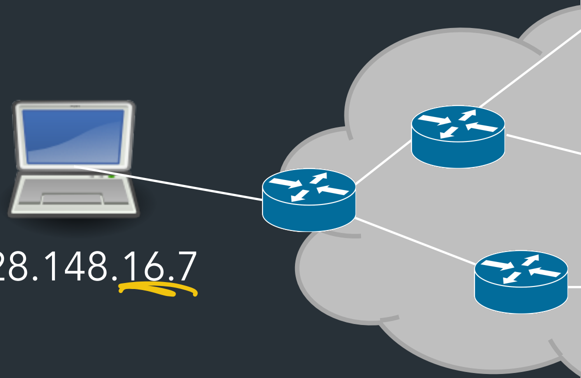

128.148.16.7

IP Addressing

An IP address identifies…

- Who a host is: A unique number
- Where it is on the Internet

- Networks are allocated ranges of IPs by global authority (ICANN)

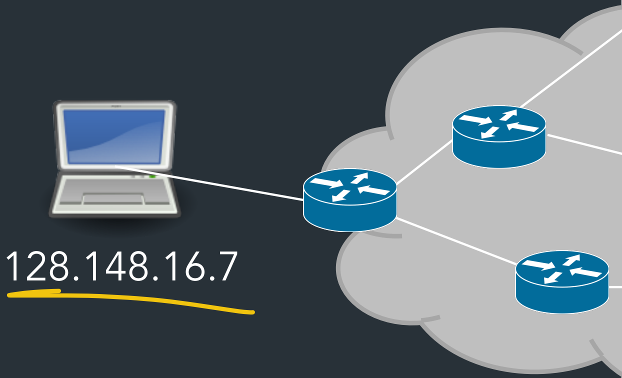

*ICANN (Internet Corporation for Assigned Names and Numbers)

IP Addressing

An IP address identifies…

- Who a host is: A unique number
- Where it is on the Internet

- Networks are allocated ranges of IPs by global authority (ICANN)
  - Further subdivided by regions, ISPs, …
- Some IPs have special uses (eg. 127.0.0.1)

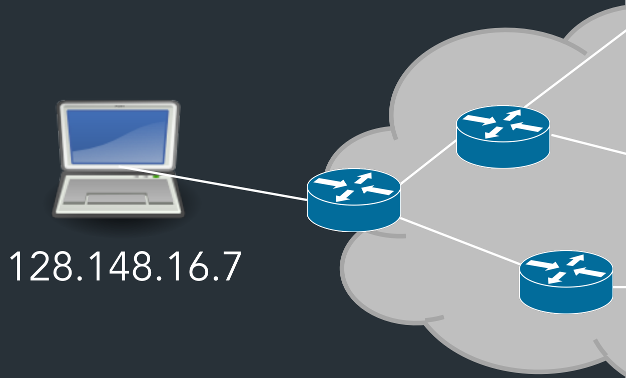

*ICANN (Internet Corporation for Assigned Names and Numbers)

IP Addressing

An IP address identifies…

- Who a host is: A unique number
- Where it is on the Internet

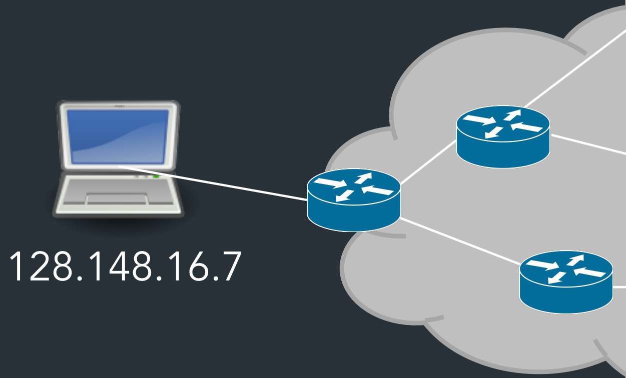

- Networks are allocated ranges of IPs by global authority (ICANN)
  - Further subdivided by regions, ISPs, …
- Some IPs have special uses (eg. 127.0.0.1)

eg. Brown owns 128.148.xxx.xxx, 138.16.xxx.xxx

*ICANN (Internet Corporation for Assigned Names and Numbers)

IP Addressing

Brown owns the range:

128.148.xxx.xxx

2^{16} HOSTS

// ≈ 65K HOSTS

10000000 10010100 xxxxxxxx xxxxxxxx

NETWORK

16 BITS

16 BITS

Network part

Identifies Brown (to the Internet)

Host part

Denotes individual hosts within the Brown Network

128.148.0.0/16

128.148/16

IP Addressing

A network can designate IP addresses for its own hosts within its address range

For 128.148.xxx.xxx:

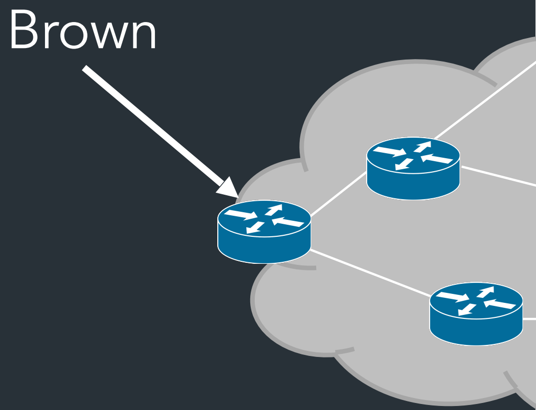

IP Addressing

A network can designate IP addresses for its own hosts within its address range

0 - 255

For 128.148.xxx.xxx:

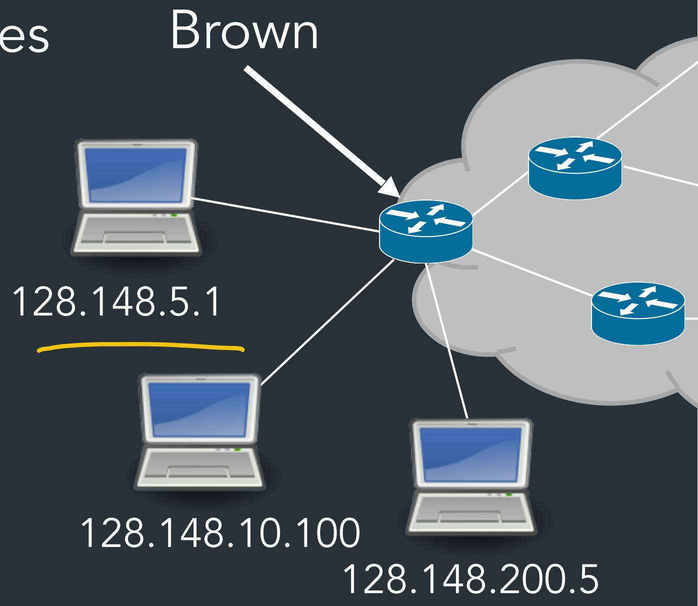

IP Addressing

A network can designate IP addresses for its own hosts within its address range

For 128.148.xxx.xxx:

`10000000 10010100 xxxxxxxx xxxxxxxx`

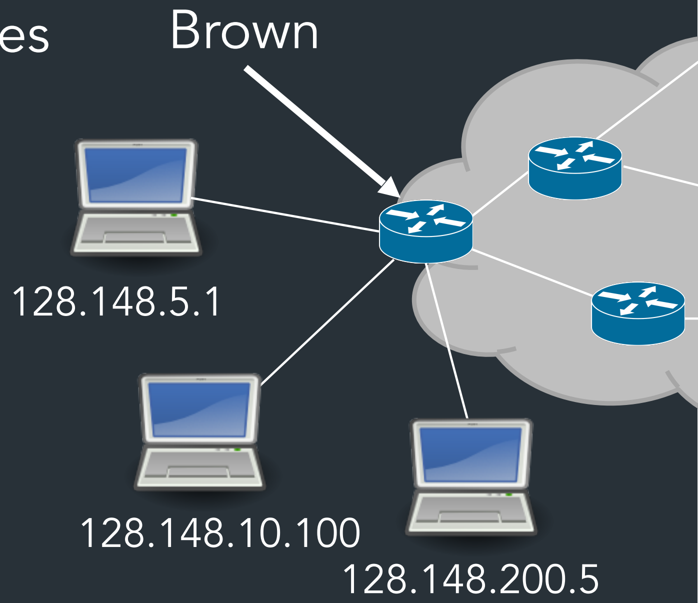

Common prefixes

1.2.0.0/16	`00000001 00000010 xxxxxxxx xxxxxxxx`
8.0.0.0/8	`00001000 xxxxxxxx xxxxxxxx xxxxxxxx`
	$2^{24} \approx$ 16M HOSTS
123.10.1.0/24	`01111011 00001010 00000001 xxxxxxxx`
	$2^8 \approx$ 256 HOSTS.
201.112.10.200/30	`11001001 01110000 00001010 110010xx`
	$2^2 = 4$ HOSTS

Example

How many addresses are in the network 1.1.0.0/20?

1 . 1 . 0 . 0 / 20 — FIRST 20 BITS ARE NETWORK PART

0 7 15 23 31

0000 0001 | 0000 0001 | 0000 0000 | 0000 0000

20 BITS => NETWORK PART

12 BITS = HOST PART

2^{12} = $\boxed{\text{4096 POSSIBLE HOSTS}}$

How do we move
packets between networks?

$\Rightarrow$ FORWARDING.

How IP forwarding works

Assume:

- Communicating on same network is easy—this is the link-layer's job!

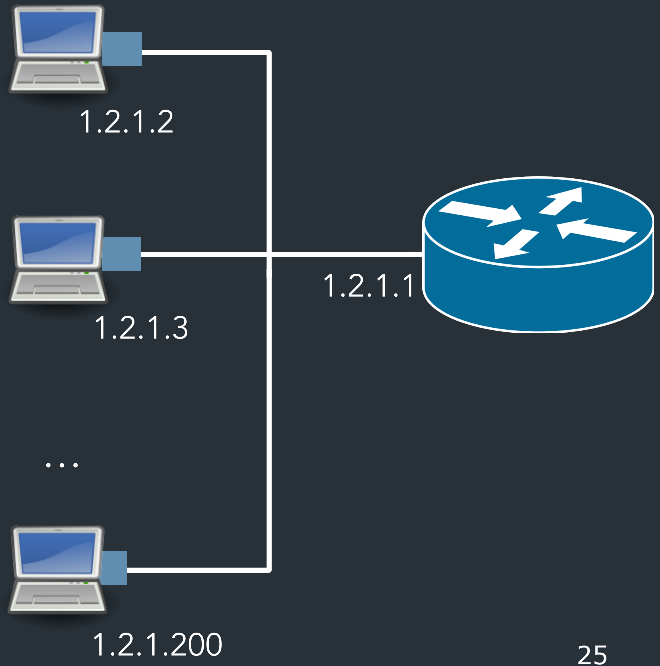

How IP forwarding works

Assume:

- Communicating on same network / DIRECTLY CONNECTED VIA ETHERNET WIFI, ETC.) is easy—this is the link-layer's job!
- Can map IP addresses to MAC addresses (more on this later)

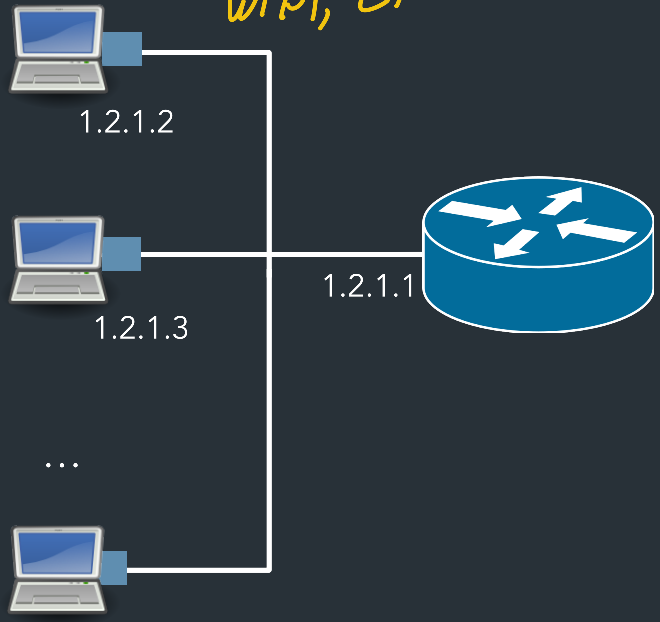

How IP forwarding works

Assume:

- Communicating on same network is easy—this is the link-layer's job!
- Can map IP addresses to MAC addresses (more on this later)

How to reach an address outside this network?

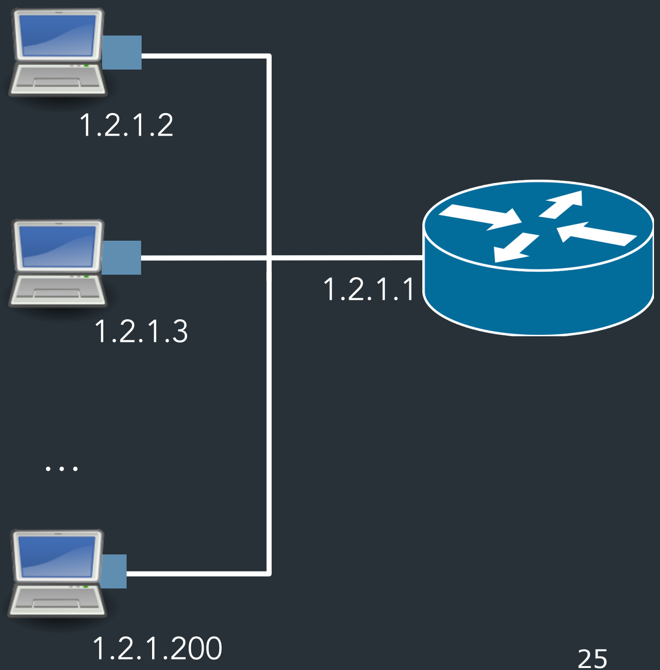

How IP forwarding works

Assume:

- Communicating on same network is easy—this is the link-layer's job!
- Can map IP addresses to MAC addresses (more on this later)

How to reach an address outside this

Send packets to a <u>router</u>, which <u>forwards</u> IP packets to other networks

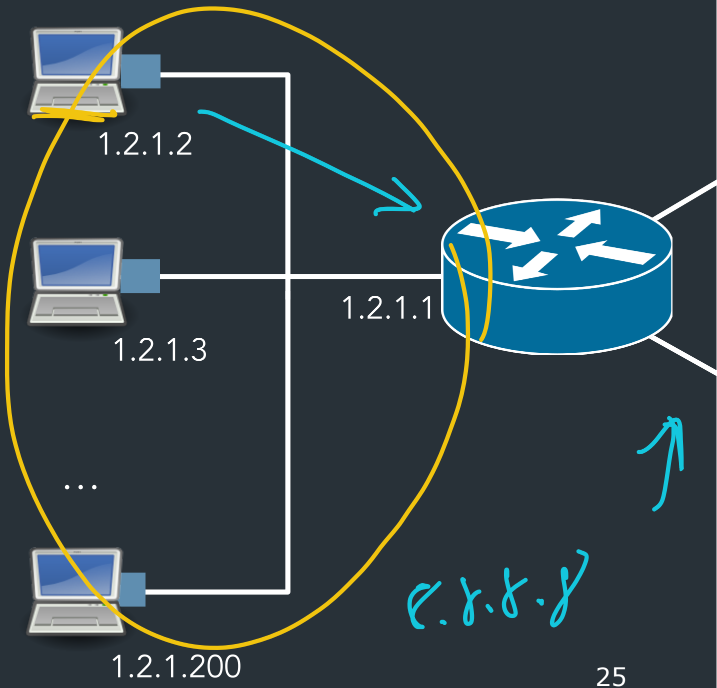

Forwarding IP packets

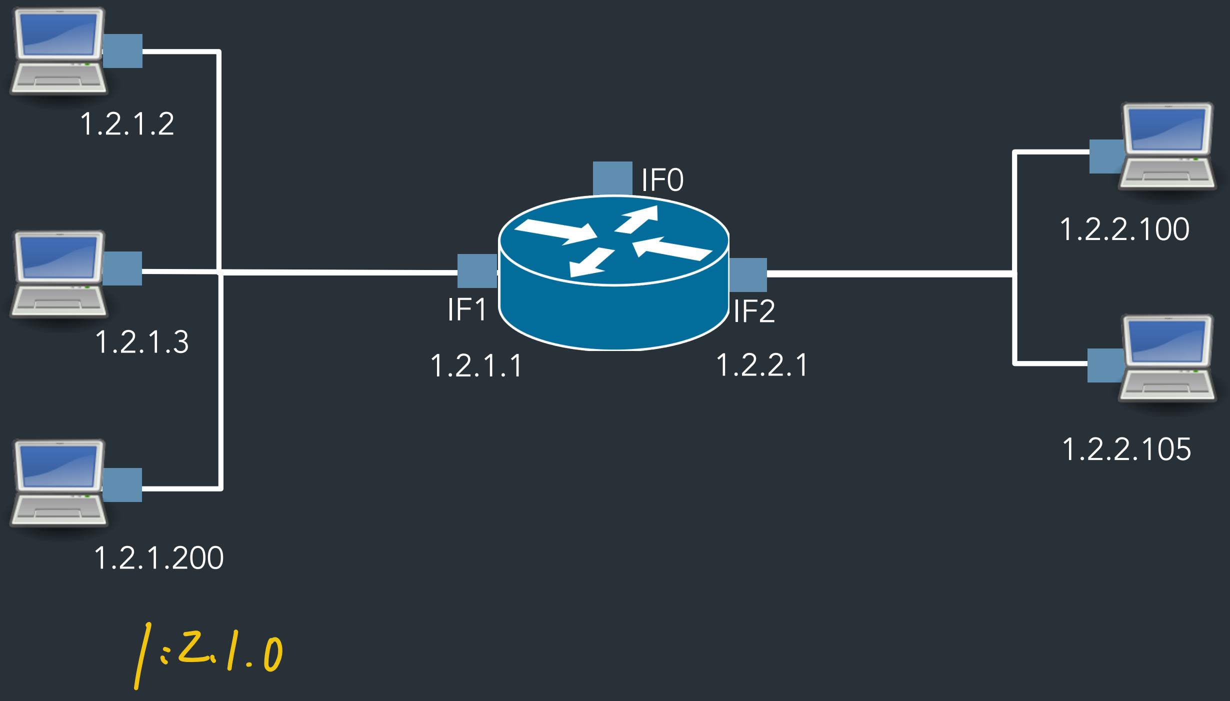

Forwarding IP packets

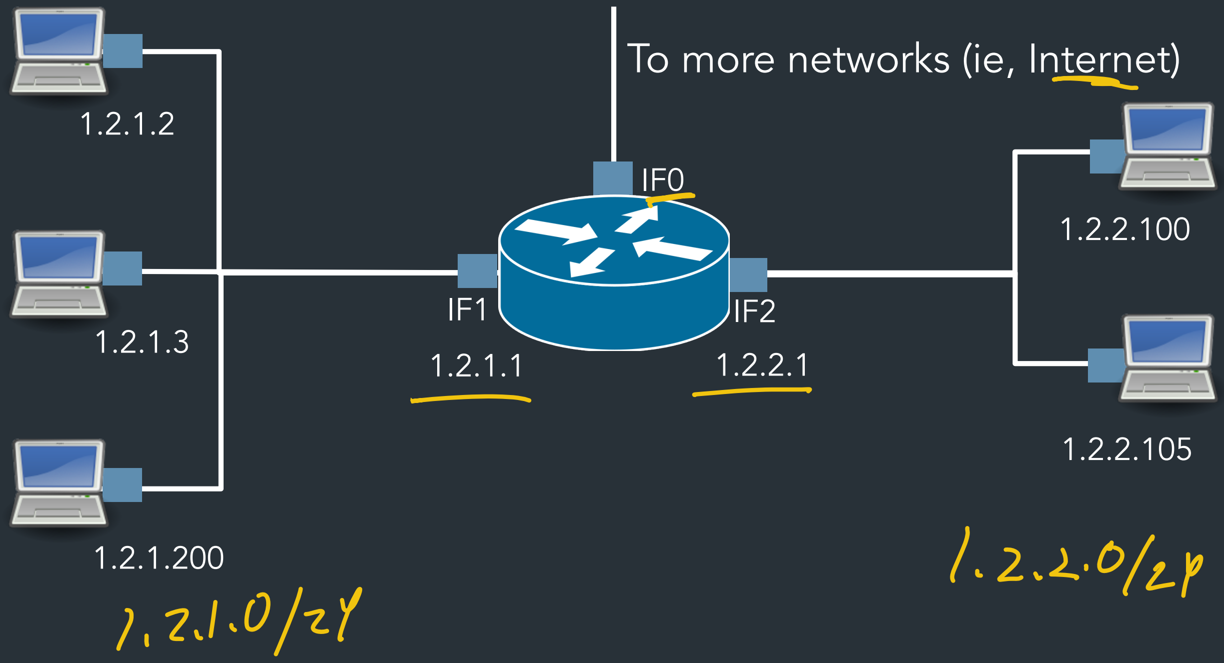

Forwarding IP packets

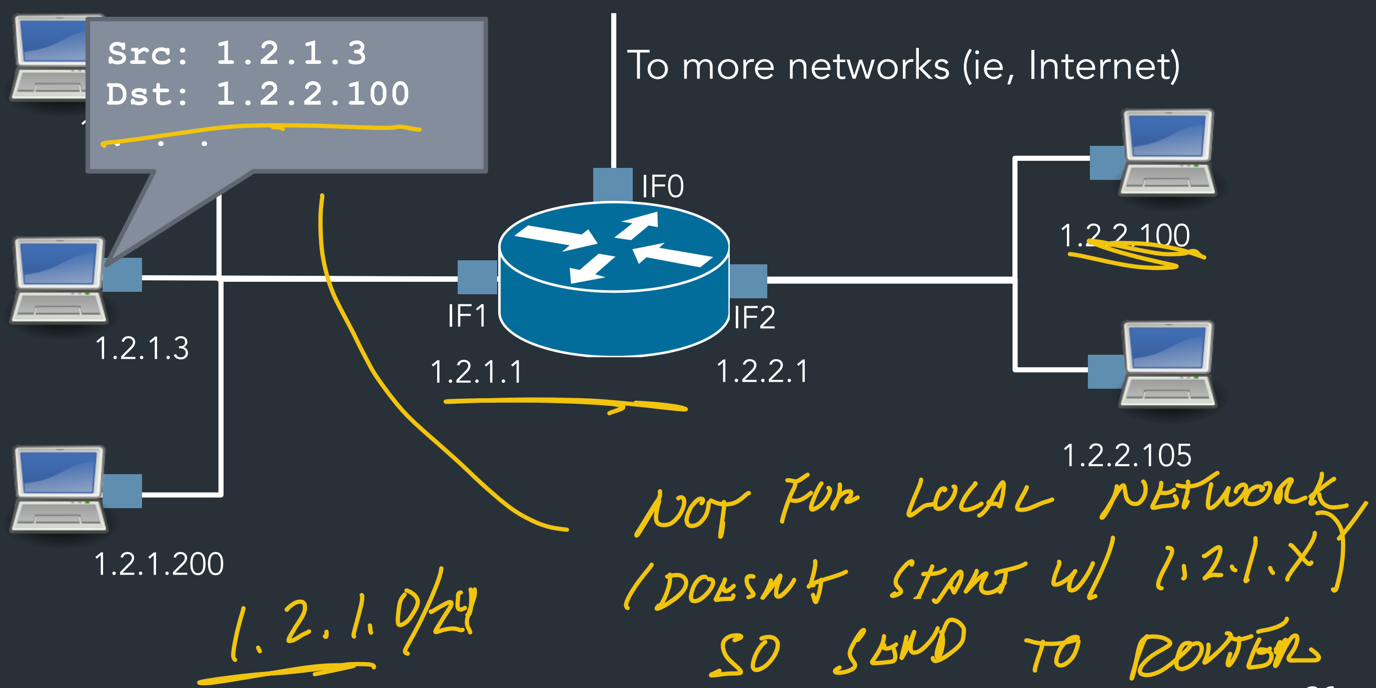

Forwarding IP packets

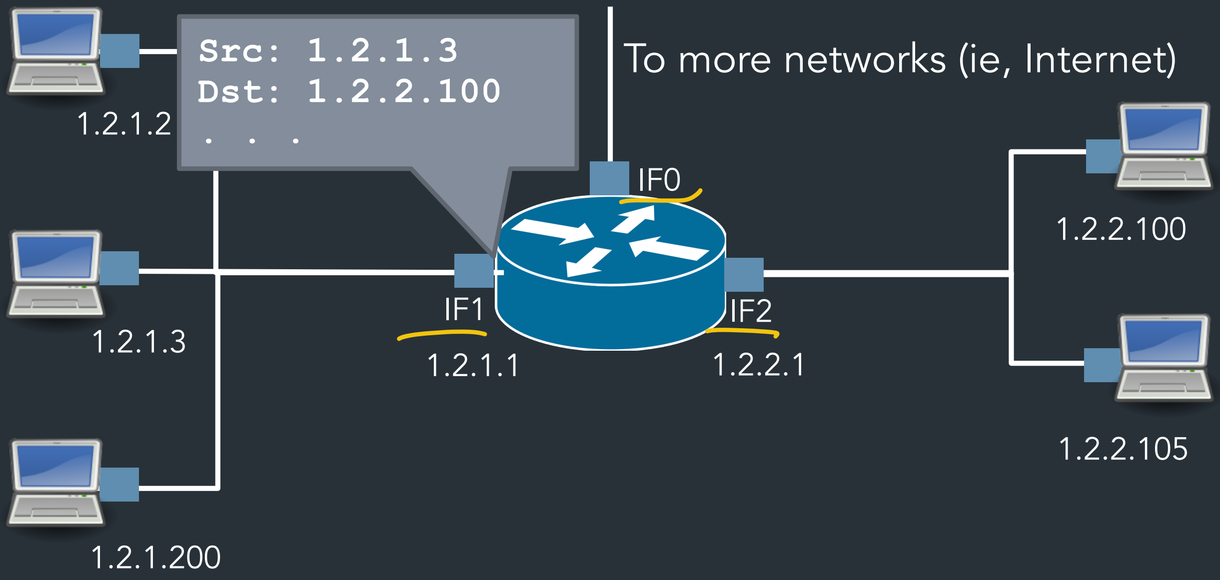

Forwarding IP packets

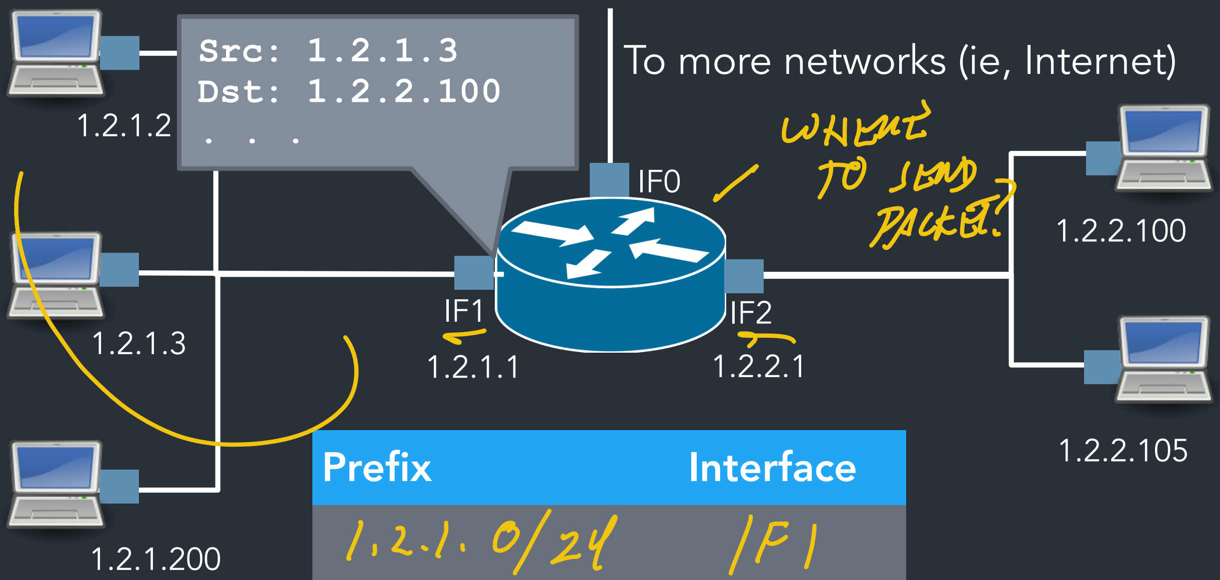

Prefix	Interface
1.2.1.0/24	IF1
1.2.2.0/24	IF2
*	IF0

Forwarding IP packets

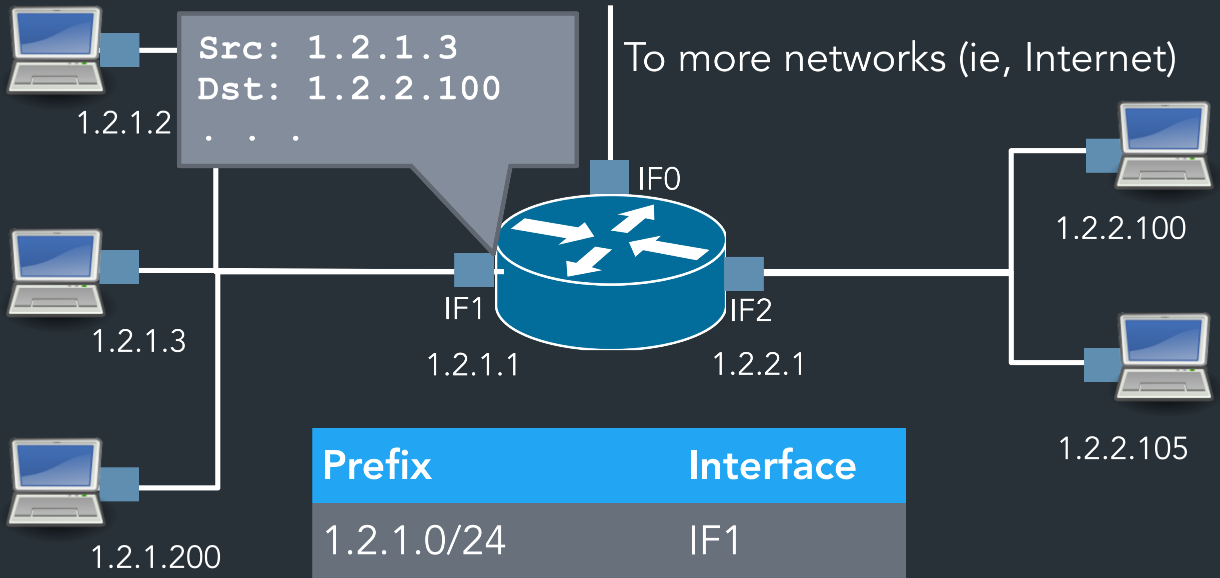

Prefix	Interface
1.2.1.0/24	IF1
1.2.2.0/24	IF2
<everything else>	(IF0)

What about the rest?

How to reach networks that aren't directly connected?

Prefix	Interface
1.2.1.0/24	IF1
1.2.2.0/24	IF2
<everything else>	IF0

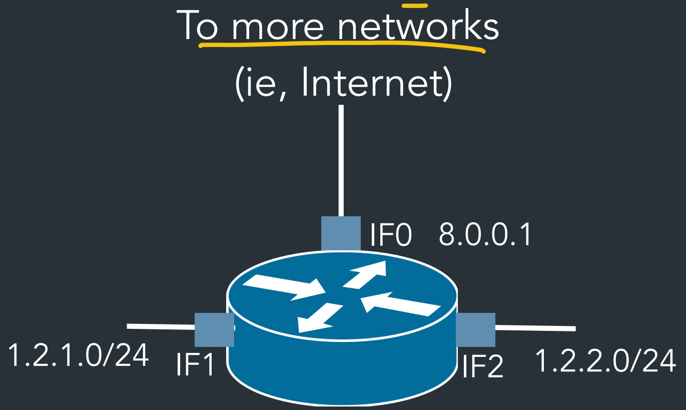

What about the rest?

DEFAULT | 8.0.0.2

- Need "next hop" IP: another router that knows about other networks
  - How to reach it? Check table again!
- "Default gateway": where to send to reach anything not in the table

Prefix	IF/Next hop
1.2.1.0/24	IF1
1.2.2.0/24	IF2
8.0.0.0/30	IF0
128.148.0.0/16	1.2.1.5
Default	8.0.0.2

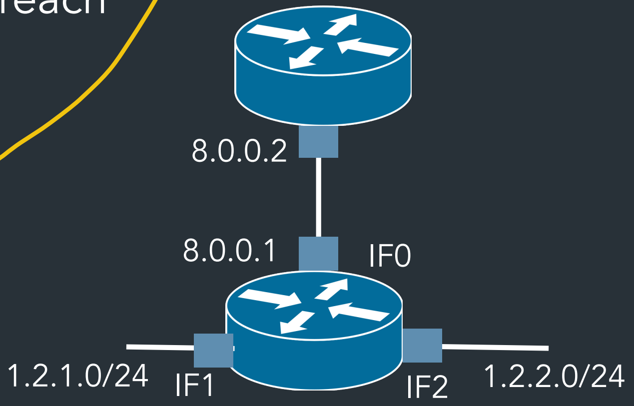

The forwarding table

Exploits hierarchical structure of addresses: know how to reach networks, not individual hosts

Prefix	IF/Next hop
1.2.1.0/24	IF1
1.2.2.0/24	IF2
8.0.0.0/30	IF0
128.148.0.0/16	1.2.1.5
Default	8.0.0.2

- Table is keyed is a network prefix, not a whole address
- Select best prefix with longest prefix matching (more on this later)

A forwarding table

```
# ip route
127.0.0.0/8 via 127.0.0.1 dev lo
172.17.44.0/24 dev enp7s0  proto kernel  scope link  src 172.17.44.22  metric 204
default via 172.17.44.1 dev eth0 src 172.17.44.22  metric 204
```

The IPv4 Header

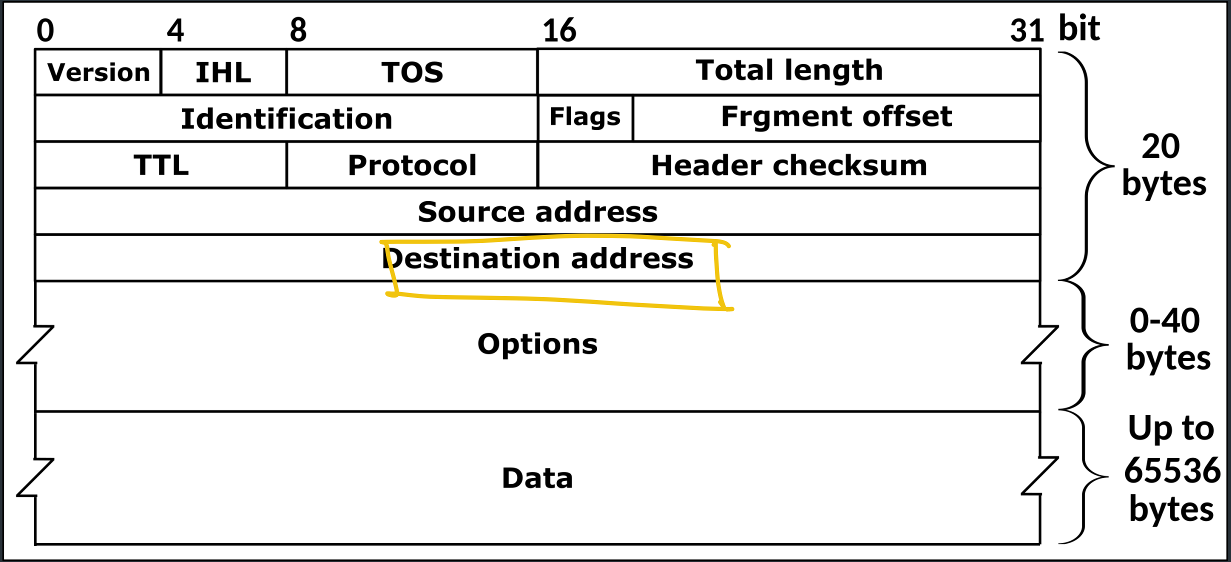

IP: THE ASSIGNMENT

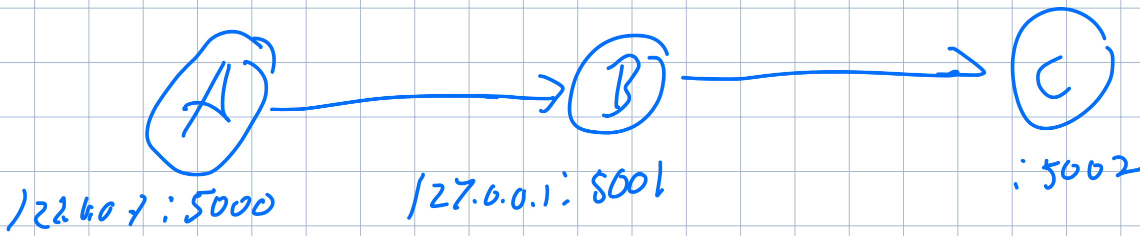

- FOR US, THE "LINK LAYER" IS A UDP SOCKET

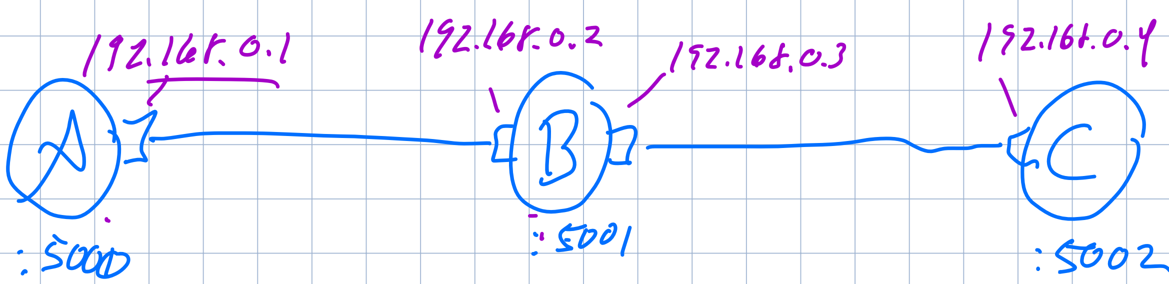

A → B UDP 127.0.0.1:5000 → 127.0.0.1:5001
IP 192.168.0.1 192.168.0.2

A → C 127.0.0.1:5000 → 127.0.0.1:5001
IP 192.168.0.1 192.168.0.4