EEE 545

Computer Networks 1

7. IP Version 4

LANs and WANs

LANs

- Different types: different topologies, different technologies, different purposes
- Many LANs operate at layers 1 and 2 (Physical and Data Link Layer) using switches and hubs
- Bridges can connect LANs of similar technologies together

WANs

- Can interconnect LANs over a larger distance
- Point-to-point link (e.g. ADSL, PDH) or a network (e.g. ATM, SDH, telephone) using packet or circuit switching
- Device that interconnects the WAN to LAN must support both technologies
- WANs typically operate at Layers 1 and 2

Connect Multiple LANs and WANs

- Organisations have different requirements of their network, and therefore may choose different technologies for their LANs/WANs
- Aim: allow any computer to communicate with any other computer, independent of what LAN/WAN they are connected to
- **Internetworking** involves connecting the many different types of LANs/WANs together to achieve this aim
- An internetworking protocol supports data delivery across different types of LANs/WANs
- E.g. the **Internet Protocol (IP)**

Internetworking with Routers

- Internetworking is performed by using **routers.**
- Routers connect two or more LANs or WANs together. The idea of the router is to allow us to communicate between any combination of LAN and WAN irrespective of the technologies used internally but still be able communicate with anyone else
- Routers are packet switches that operate at **network layer**

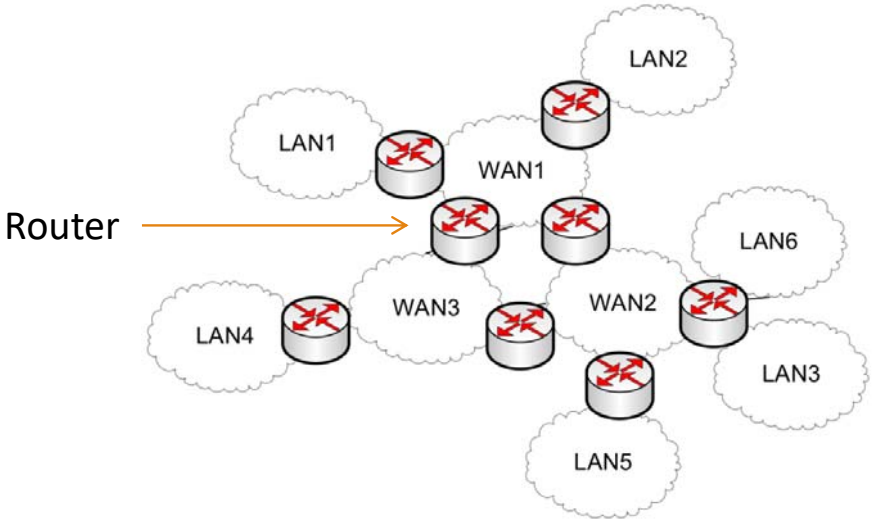

In this figure, we have 9 sub-networks to form one large network, which make 10 networks in total. The larger network is called Internet. We have interconnected many networks together. Small networks are called **subnetworks,** which can be abbreviated as **subnets.**
As a result, this internet is formed from joining nine different subnets together.
Routers are packet switchers

The Internet Protocol (IP)

- IP is the internetworking protocol used in the Internet Implemented in hosts and routers
- Features:
  - Datagram packet switching
  - Network layer
  - Connection - less
  - Addressing
  - Fragmentation-and-reassembly
- IP version 4 most widely used; IPv6 is available
- Features IP does NOT provide:
  - Connection control, error control, flow control (TCP)
  - Status reporting (ICMP)
  - Priority, quality of service (DiffServ, IntServ)
  - Security (IPsec)

Terminology

- **Routers:** nodes that connect networks (LANs/WANs) together; operate at network layer
- **Subnetworks:** individual networks (LANs and WANs)
- **Internetworking:** connect two or more subnets together using routers
- An internetwork or an internet: the resulting network from internetworking
- **The Internet:** an internet that uses the Internet Protocol (IP) and used today to connect networks across the globe
- **Routing:** process of discovering a path from source to destination through a network
- **Forwarding:** process of sending data along a path through a network
- **Packet Switch:** a generic device that performs switching in a Packet Switching network. May operate at data link or network layer. A packet switch at network layer is called a router
- **Circuit Switch:** a generic device that performs circuit switching in Circuit Switching network
- **Ethernet switch:** an IEEE 802.3 switch (either Ethernet, Fast Ethernet or Gigabit Ethernet). Operates at data link layer

IP in the TCP/IP Stack

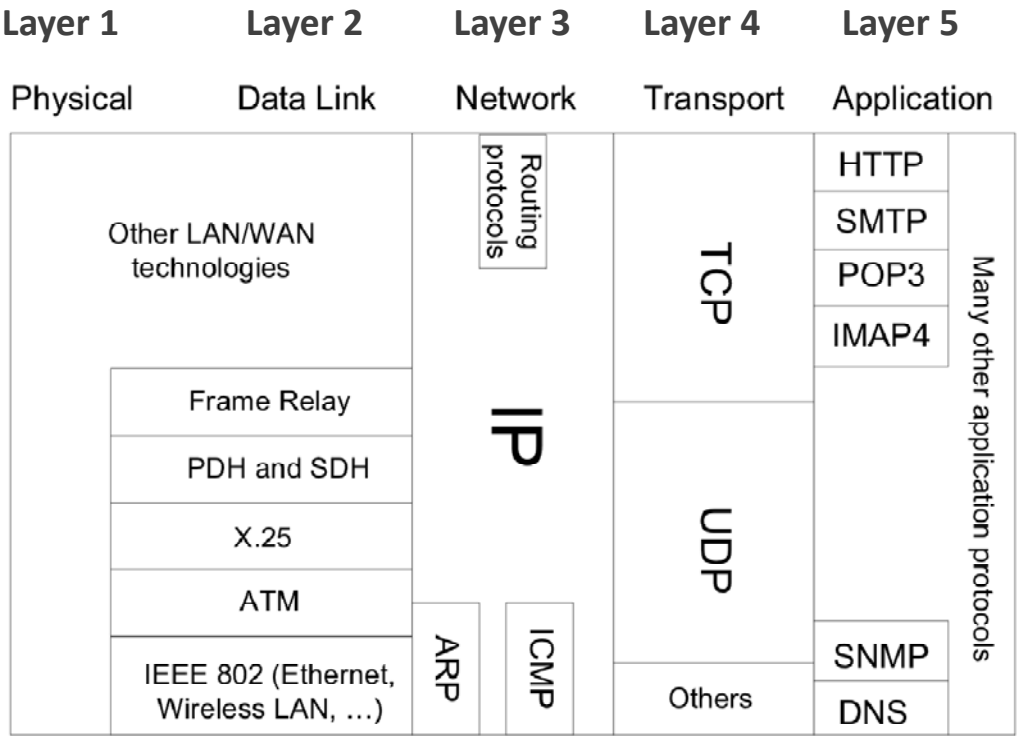

This figure is the stack of 5 layers, physical up to application. This figure contains examples of different technologies and protocols at some of the layers.

IP is located in the Network Layer and it really forms of the core of the internet. It doesn't matter what lower technology we are using and what application we are using at the top. It all centers around IP.

There are routing protocols which are separate from IP. IP is not a routing protocol but a forwarding protocol. A routing protocol finds the best paths and puts the information in routing tables. IP just uses the routing tables.

IP Hosts and Routers

- **Hosts** are the end-devices (stations)
  - Usually only use single network interface at a time
  - Hosts do not forward IP datagrams
  - Either source or destination
- **Routers** are the datagram packet switches
  - Routers have two or more interfaces (since they connect LANs/WANs together)
  - Routers forward datagrams
  - Routers can act as a source or destination of datagrams (however this is mainly for management purposes)
- **IP routing** is the process of discovering the best path between source and destination; store destination and next router in routing table
  - E.g. RIP, EIGRP, OSPF, BGP
- **IP forwarding** is the process of delivering an IP datagram from source to destination; read next router from routing table

IP Datagram (IP Packet)

- Variable length header and variable length data
- **Header:** 20 Bytes of required fields; optional fields may bring header size to 60 Bytes
- **Data:** length must be integer multiple of 8 bits; maximum size of header + data is 65,656 Bytes

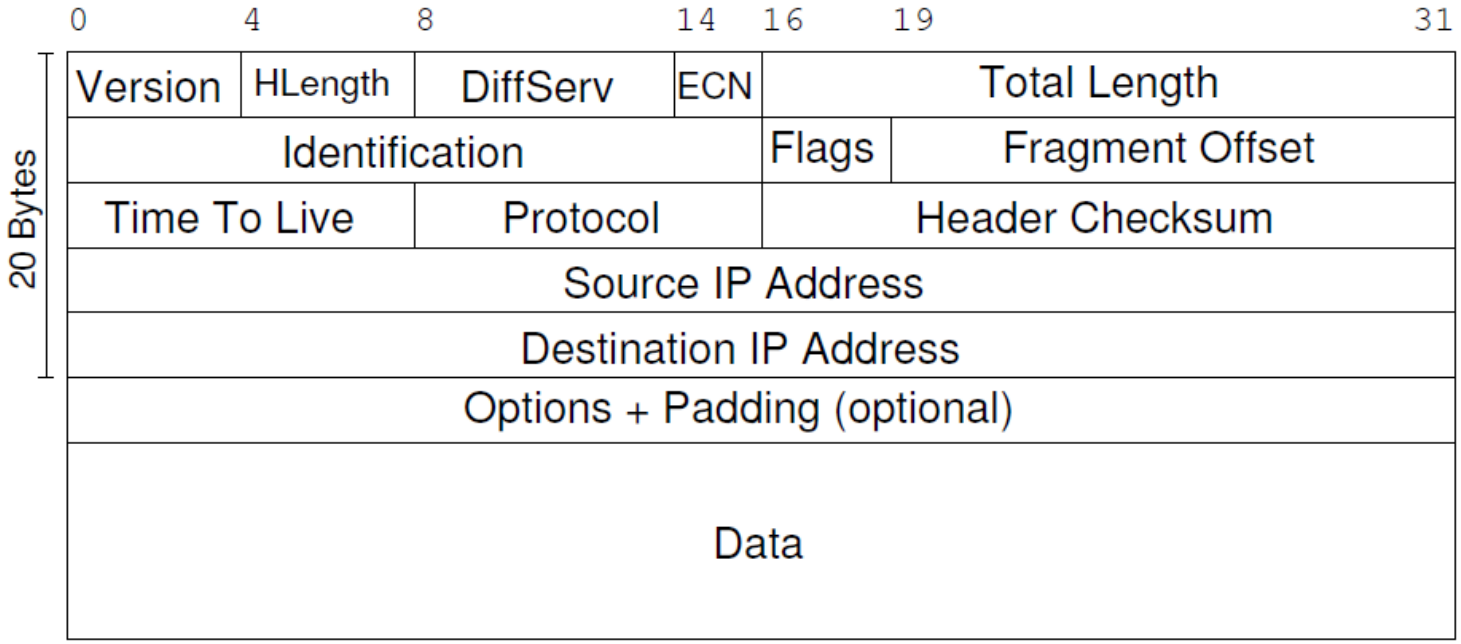

IP Datagram Fields

- **Version [4 bits]:** version number of IP; current value is 4 (IPv4)
- **Header Length [4 bits]:** length of header, measured in 4 byte words
- **DiffServ [6 bits]:** Used for quality of service control
- **ECN [2 bits]:** Used for notifying nodes about congestion
- **Total Length [16 bits]:** total length of the datagram, including header, measured in bytes
- **Identification:** sequence number for datagram
- **Flags:** 2 bits are used for Fragmentation and Re-assembly, the third bit is not used
- **Fragment Offset [13 bits]:** See Fragmentation and Re-assembly
- **Time To Live [8 bits]:** datagram lifetime
- **Protocol [8 bits]:** indicates the next higher layer protocol
- **Header Checksum [16 bits]:** error-detecting code applied to header only; recomputed at each router
- **Source Address [32 bits]:** IP address of source host
- **Destination Address [32 bits]:** IP address of destination host
- **Options:** variable length fields to include options
- **Padding:** used to ensure datagram is multiple of 4 bytes in length
- **Data:** variable length of the data

IP Routing and Forwarding

Routing Tables

- Store address of destination and next node
- Created manually or by routing protocols

Routing Protocols in the Internet

- Collect network status information, calculate least cost paths and update routing tables
- Adaptive routing protocols: OSPF, RIP, EIGRP, BGP

Forwarding

- Routers forward IP datagrams from source host to destination host
- Destination host address in IP datagram header
- Lookup destination address in routing table

Other Features

- IP includes:
  - **Fragmentation and reassembly**: source host and routers may divide datagrams into smaller fragments; destination host reassembles fragments into full datagram
  - **Time To Live (TTL):** source sends “lifetime” of datagram in header; decremented by each router; if 0, datagram is discarded
- Other network layer features:
  - ICMP: error reporting, ping
  - ARP: map IP addresses to Ethernet addresses
  - IPv6
  - Multicasting
  - Quality of Service (DiffServ)
  - Mobility (Mobile IP)
  - Security (IPsec)

IPv4 Addresses

- IPv4 addresses are 32 bits in length
- Split into **network** portion and **host** portion: rst *N* bits identify a subnet in the Internet; last *H* bits identify an IP device (host/router) in that subnet
- All subnets in the Internet have unique network portion
- All IP devices in a subnet have same network portion, but unique host portions
- Where/how to split has changed over time: Classful, Subnet addressing, Classless addressing
- Focus on classless addressing
- Why split? Allows hierarchical addressing, makes routing in Internet scalable

Representing IPv4 Addresses

- Writing and remembering 32 bits is difficult for humans
- IP addresses usually written in **dotted decimal notation**
- Decimal number represents the bytes of the 32-bit address
- Decimal numbers are separated by dots

```
IP: 11000000111001000001000100111001= 192.228.17.57
```

Classless IP Addressing

- **Subnet mask** or address mask identifies where the IP address is split between network and host portion
- Mask is 32 bits: a bit 1 indicates the corresponding bit in the IP address is the network portion; a bit 0 indicates the corresponding bit in the IP address is the host portion
- The mask can be given in dotted decimal form or a shortened form, which counts the number of bit 1's from left
- Each IP address comes with a subnet mask

```
IP:   11000000101010000000000100000001= 192.168.1.1/24
Mask: 11111111111111111111111000000000= 255.255.255/24
```

Network portion, subnet in the internet

Host portion, the device in that subnet

Example of IP Addressing

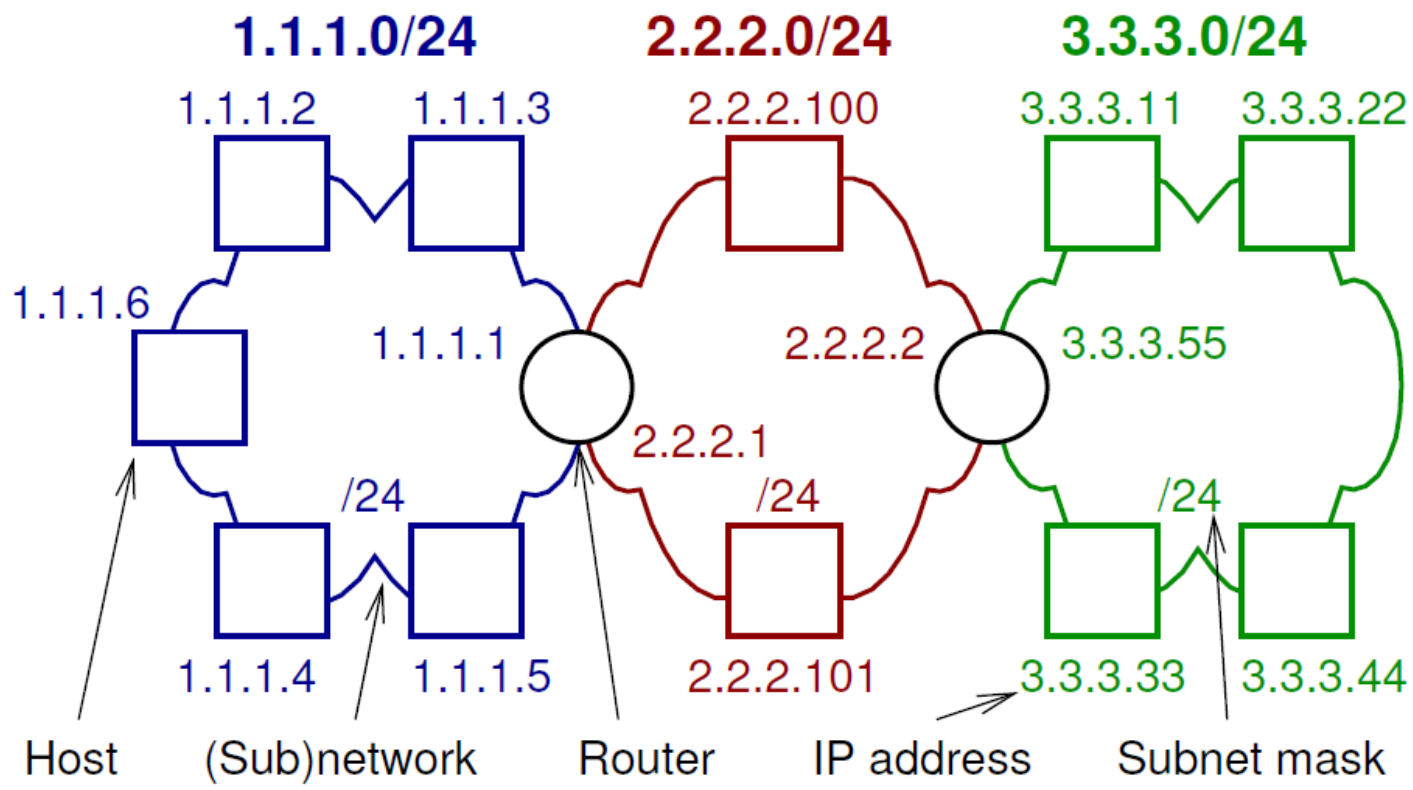

In internet, we don't need 48-bit Ethernet addresses, we need an addressing scheme such that any device on any subnet can talk to any other device. That is what IP addresses do for us. They allow us to have a common addressing scheme so that each device in theory has a unique address across the internet.

They way we do it, each subnet gets a particular Network portion and each device gets a particular portion

For the figure, all three subnets, subnet mask is /24, all subnets use the same subnet mask. Network portion for these subnets are different. Every device in the same subnet must have same first 24 bits in their IP address.

Furthermore, subnet addresses, which are 1.1.1.0, 2.2.2.0 and 3.3.3.0 are special for the subnets and CANNOT be given to the devices in the subnets

Remember IP address is not for a computer or a device, it is for an interface.

A router normally has two or more IP addresses, because it attaches two or more subnets.

In the internet, hosts and routers have routing tables.

Example of IP Addressing

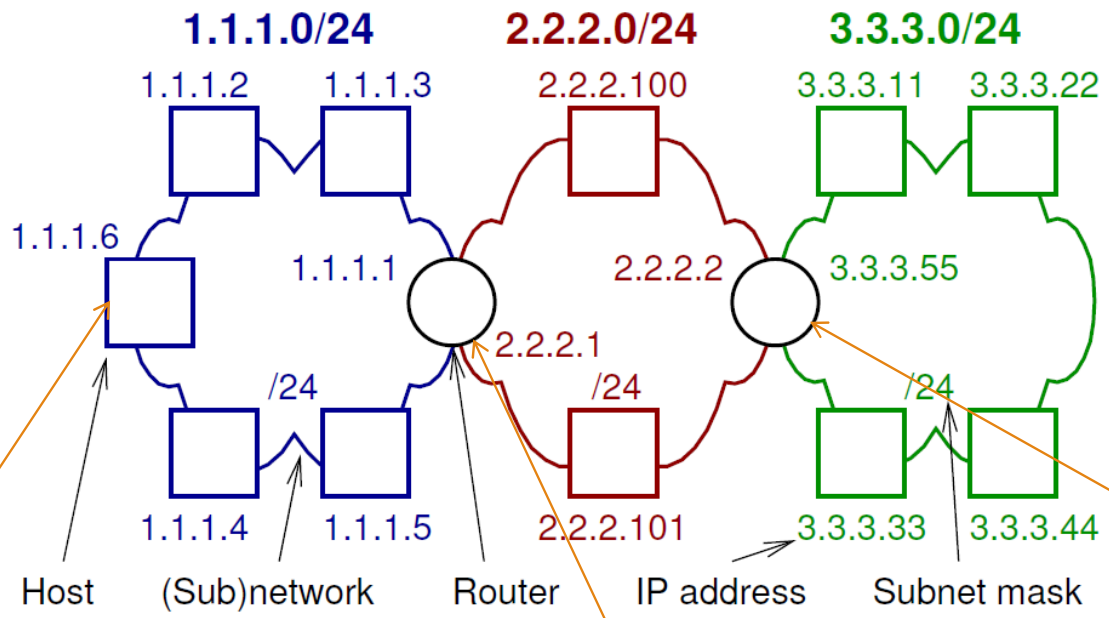

Rounding tables for host and routers

Dest	Next
1.1.1.0	Direct
* (Anyone else)	1.1.1.1

Dest	Next
1.1.1.0	Direct
2.2.2.0	Direct
3.3.3.0	2.2.2.2

Dest	Next
2.2.2.0	Direct
3.3.3.0	Direct
1.1.1.0	2.2.2.1

Example of Unicast

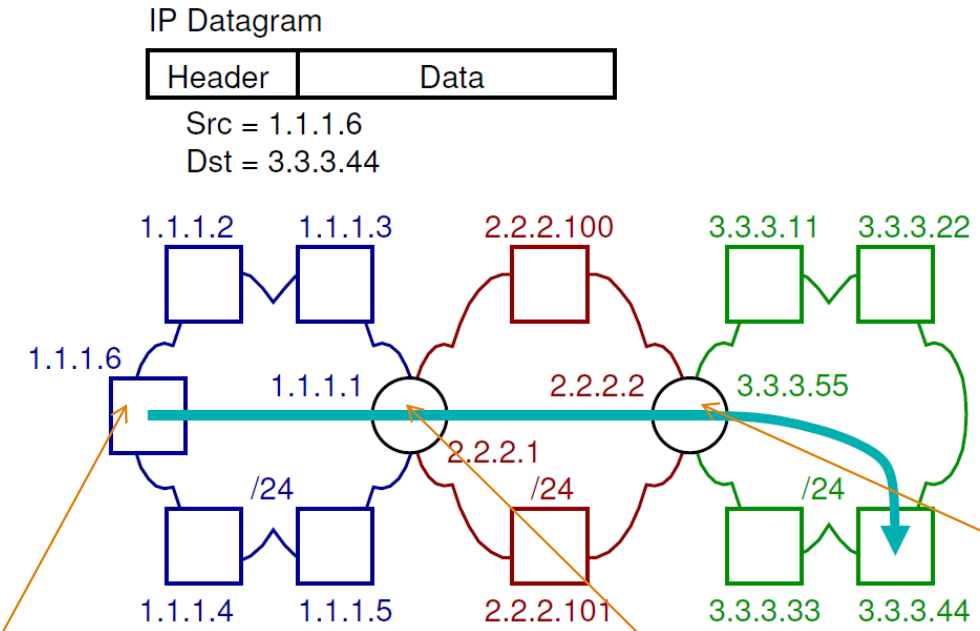

Unicast means “data transfer only to 1 point”. IP datagram at the top contains source and destination addresses.

Dest	Next
1.1.1.0	Direct
* (Anyone else)	1.1.1.1

Dest	Next
1.1.1.0	Direct
2.2.2.0	Direct
3.3.3.0	2.2.2.2

Dest	Next
2.2.2.0	Direct
3.3.3.0	Direct
1.1.1.0	2.2.2.1

Example of Directed Broadcast

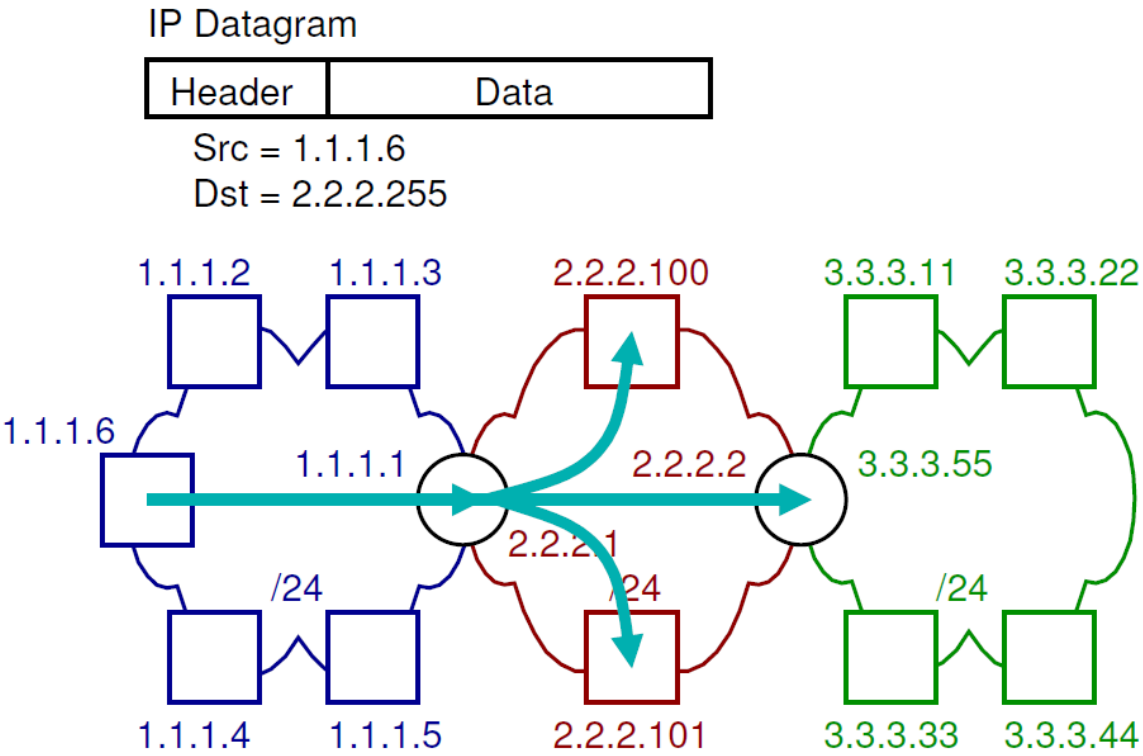

Broadcast means “to all destinations at once”. There are some special cases IP addresses reserved for. If we need to transfer data to everyone in the subnet 2.2.2.0, which is called **directed broadcast.** A directed broadcast address is a special address where the host portion is all 1s (11111111b=255)

Example of Startup Source and Local Broadcast

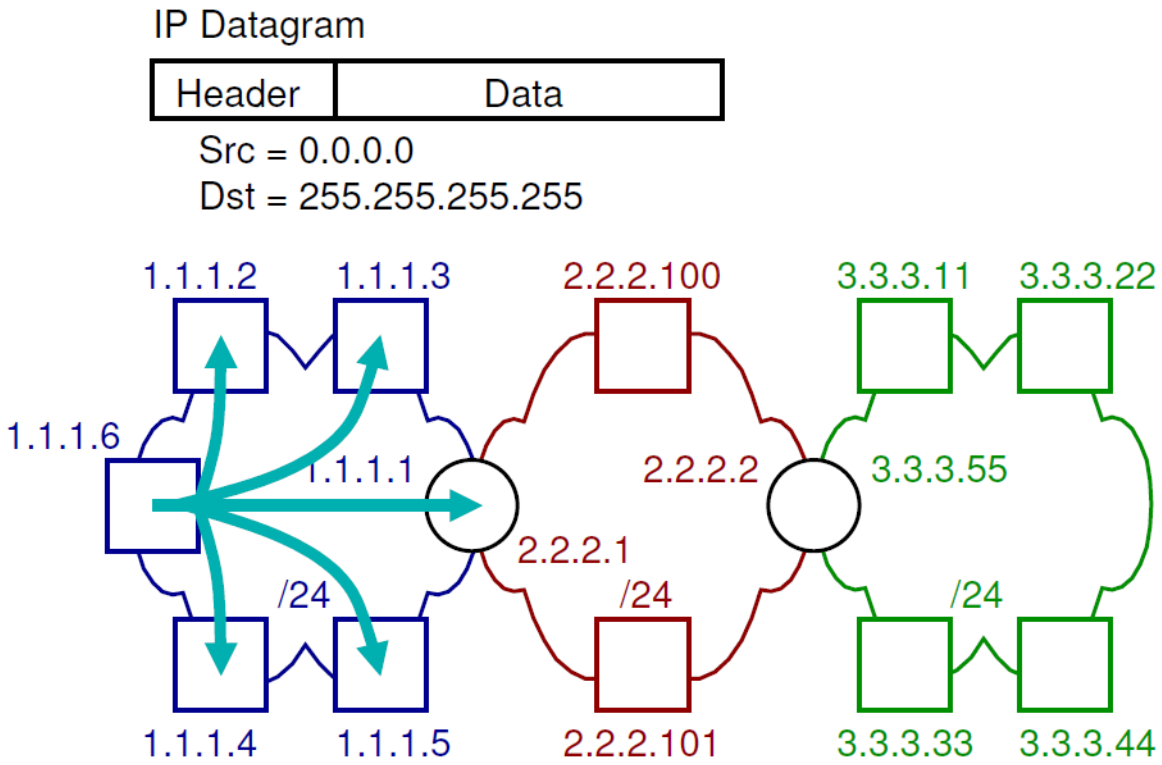

In case you startup computer and it doesn't have an IP address and it needs to discover an IP address. The computer sends a special IP datagram to everyone in the same subnet saying "Can anyone in my subnet give me an IP address?" Computer does that by setting source address as all binary 0s and set the destination address to all binary 1s. This is called **local broadcast.** Our computer hopes one of other devices will reply and give it an IP address. It hopes one of them is a special server

Example of Loopback Address

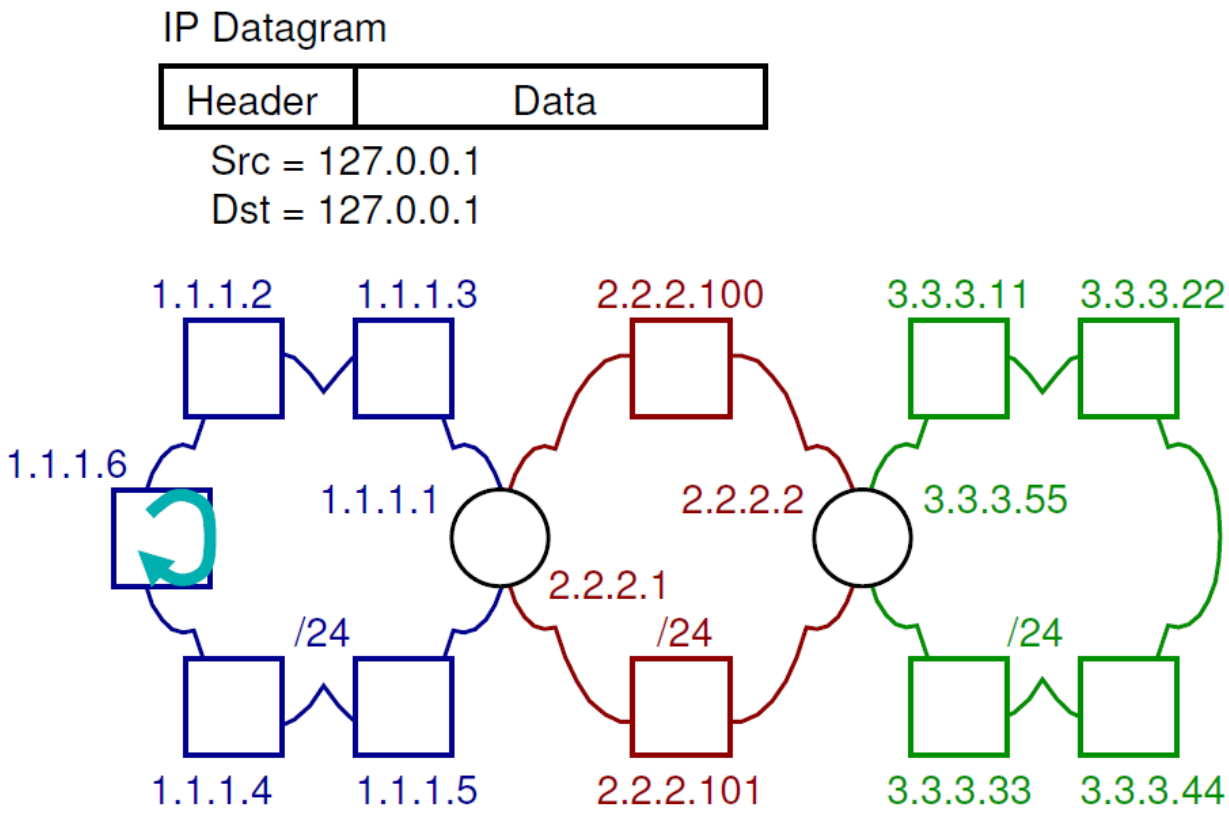

If you need to send data to yourself for testing purposes, this is called **loopback address.** It is generally used to test a software in your computer requiring internet connectivity but not disturbing the other devices.

It is sometimes called local host, which means local computer

Fragmentation

- If the goal is universal connectivity: it must accommodate networks for which the maximum packet size, or Maximum Transfer Unit, MTU, is smaller than the packet that needs forwarding. Otherwise, if we were using IPv4 to join Token Ring (MTU = 4KB, at least originally) to Ethernet (MTU = 1500B), the token-ring packets might be too large to deliver to the Ethernet side, or to traverse an Ethernet backbone en route to another Token Ring.
- So, IPv4 must support fragmentation, and thus also reassembly. There are two potential strategies here: per-link fragmentation and reassembly, where the reassembly is done at the opposite end of the link, and path fragmentation and reassembly, where reassembly is done at the far end of the path. The latter approach is what is taken by IPv4, partly because intermediate routers are too busy to do reassembly, partly because there is no absolute guarantee that all fragments will go to the same next-hop router, and partly because IPv4 fragmentation has always been seen as the strategy of last resort.

Fragmentation

- After fragmentation, the Fragment Offset field marks the start position of the data portion of this fragment within the data portion of the original IPv4 packet. Note that the start position can be a number up to 216, the maximum IPv4 packet length, but the FragOffset field has only 13 bits. This is handled by requiring the data portions of fragments to have sizes a multiple of 8 (three bits), and left-shifting the FragOffset value by 3 bits before using it.
- As an example, consider the following network, where MTUs are excluding the LAN header:

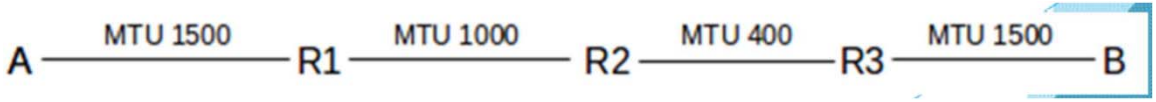

Fragmentation

- Suppose A addresses a packet of 1500 bytes to B, and sends it via the LAN to the first router R1. The packet contains 20 bytes of IPv4 header and 1480 of data.
- R1 fragments the original packet into two packets of sizes 20+976 = 996 and 20+504=544. Having 980 bytes of payload in the first fragment would fit, but violates the rule that the sizes of the data portions be divisible by 8. The first fragment packet has FragOffset = 0; the second has FragOffset = 976. R2 refragments the first fragment into three packets as follows:

- first: size = 20+376=396, FragOffset = 0
- second: size = 20+376=396, FragOffset = 376
- third: size = 20+224 = 244 (note 376+376+224=976), FragOffset = 752.

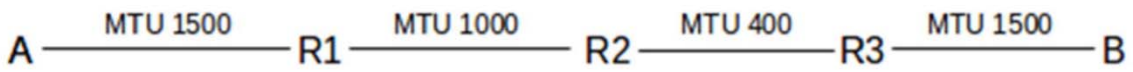

Fragmentation

- R2 refragments the second fragment into two:
  - first: size = 20+376 = 396, FragOffset = 976+0 = 976
  - second: size = 20+128 = 148, FragOffset = 976+376=1352
- R3 then sends the fragments on to B, without reassembly.