The Internet

ITS323: Introduction to Data Communications

Sirindhorn International Institute of Technology
Thammasat University

Prepared by Steven Gordon on 14 November 2014
ITS323Y14S1L14, Steve/Courses/2014/s1/its323/lectures/internet.tex, r3417

Contents

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 using routers
- Routers connect two or more LANs or WANs together
- Routers are packet switches that operate at network layer

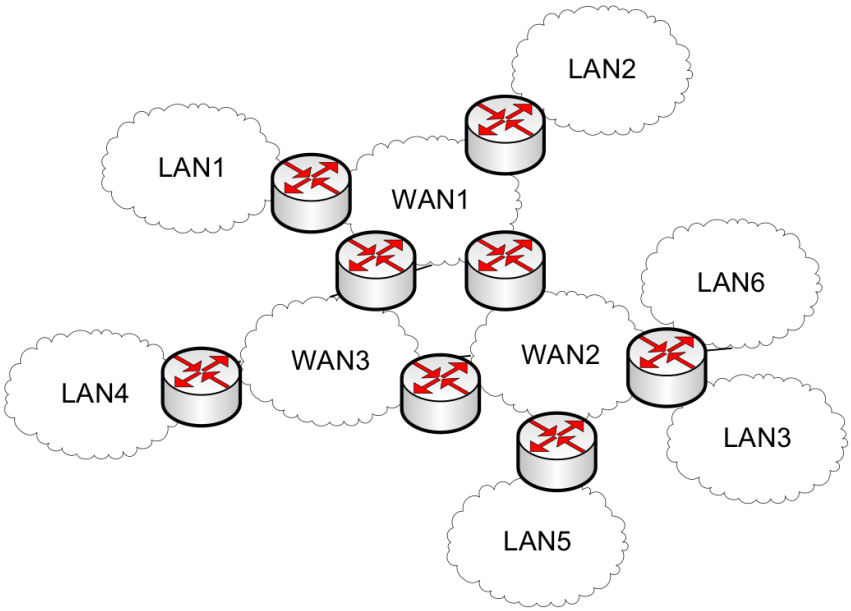

The Internet Protocol

- IP is the internetworking protocol used in the Internet
- Implemented in hosts and routers
- Features:
  - Datagram packet switching
  - Network layer
  - Connnection-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 a Circuit Switching network
- Ethernet switch: an IEEE 802.3 switch (either Ethernet, Fast Ethernet or Gigabit Ethernet). Operates at data link layer

Contents

IP in the TCP/IP Stack

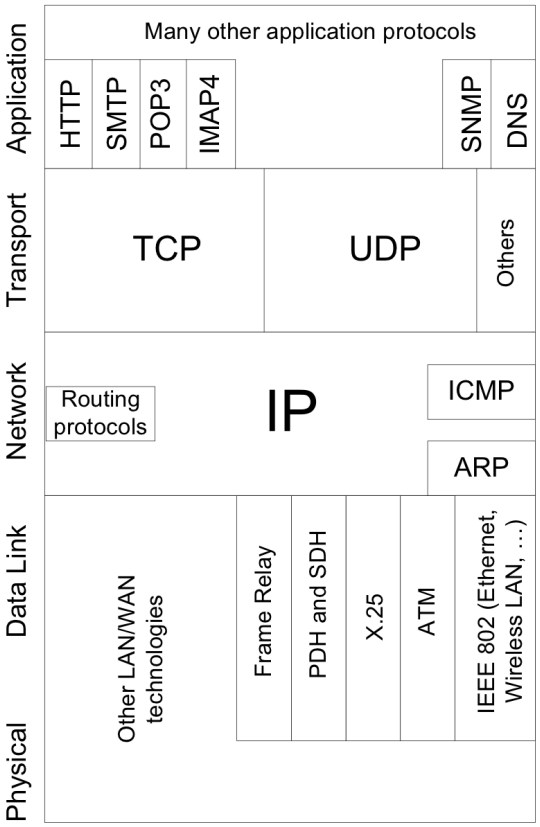

IP Hosts and Routers

- 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 Hosts and Routers

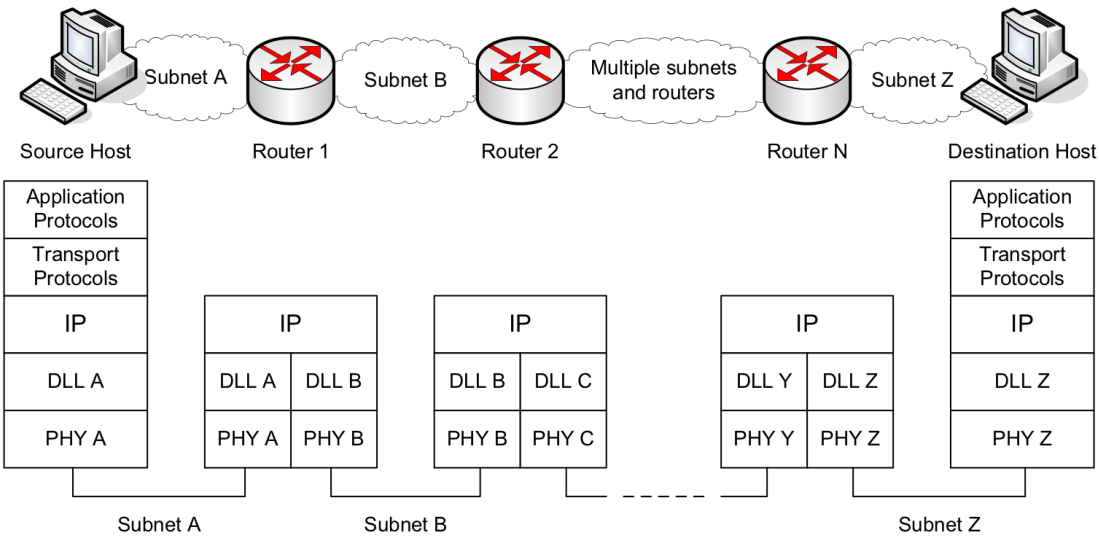

IP Datagram

- 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

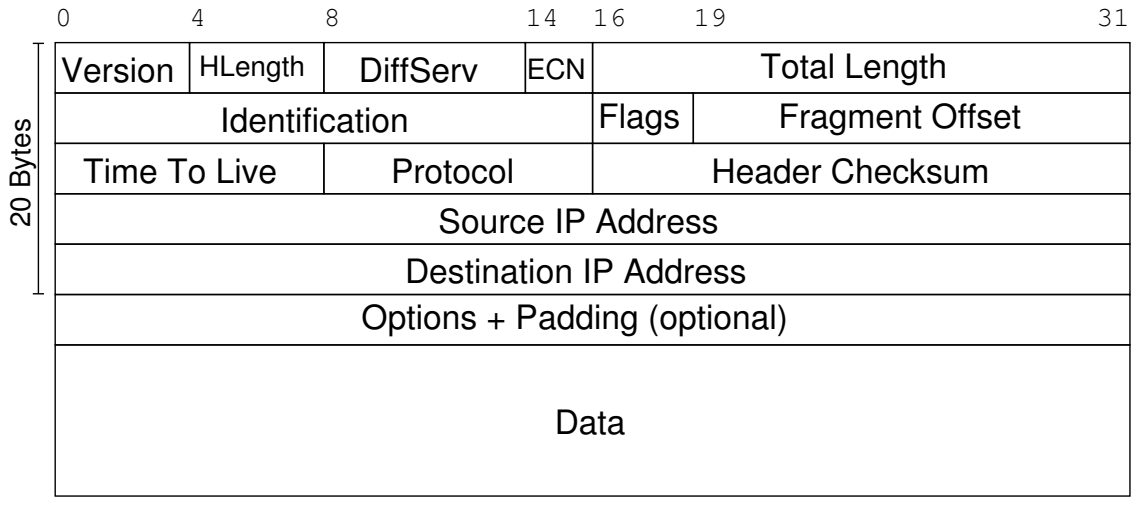

IP Datagram Fields

- 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 sents "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)

Contents

IPv4 Addresses

- IPv4 addressess are 32 bits in length
- Split into network portion and host portion: first 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
```

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

```
IP:   10000010000100010010100110000001
Mask: 11111111111111111111110000000000
```


Special Case IP Addresses

Selected IP addressess are used for special purposes; they cannot be used to identify a host

Network Address identifies a subnet in the internet; all bits in host portion are 0

Directed Broadcast Address identifies all hosts on a specific subnet; all bits in host portion are 1

Local Broadcast Address identifies all hosts on the current subnet; all bits are 1

Loopback Address identifies current host; first 8 bits are 01111111; also called localhost

Startup Source Address identifies host if currently it has no address; all bits are 0

Selected addresses reserved for private networks (e.g. not connected to Internet; behind NAT)

- 10.0.0.0—10.255.255.255
- 172.16.0.0—172.31.255.255
- 192.168.0.0—192.168.255.255

Example of IP Addressing

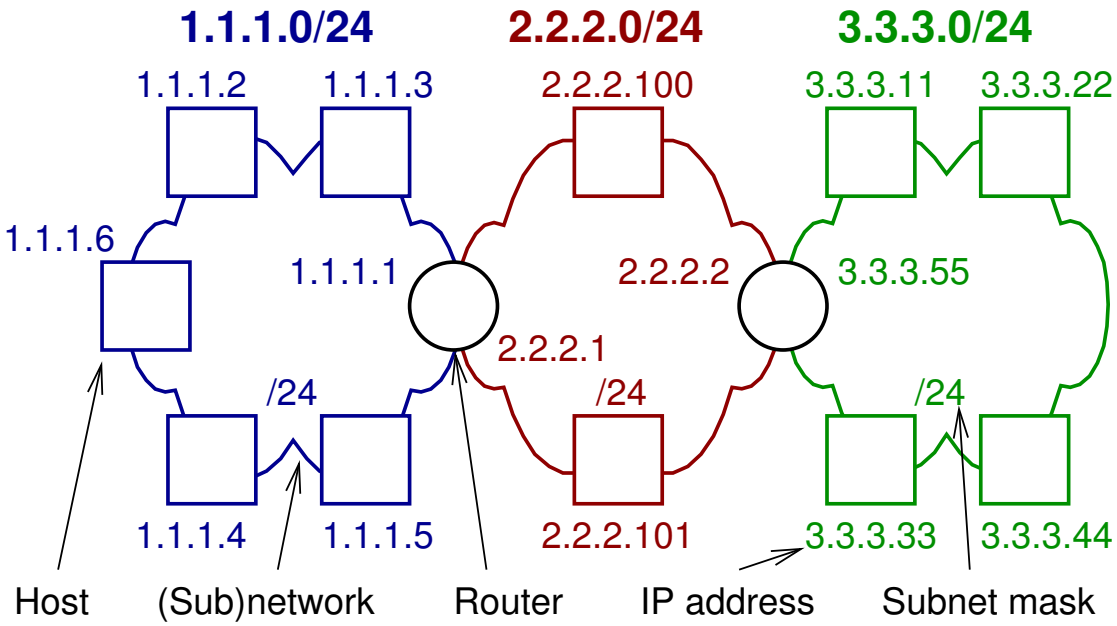

Example of Unicast

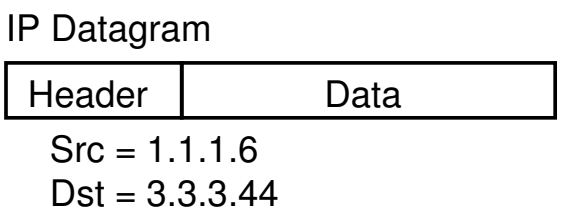

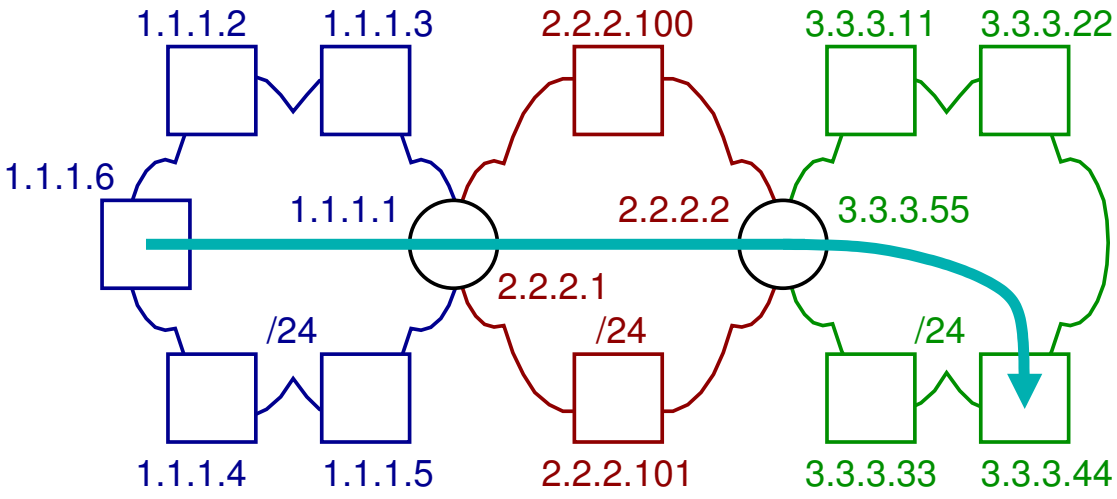

Example of Directed Broadcast

IP Datagram

Header	Data

Src = 1.1.1.6
Dst = 2.2.2.255

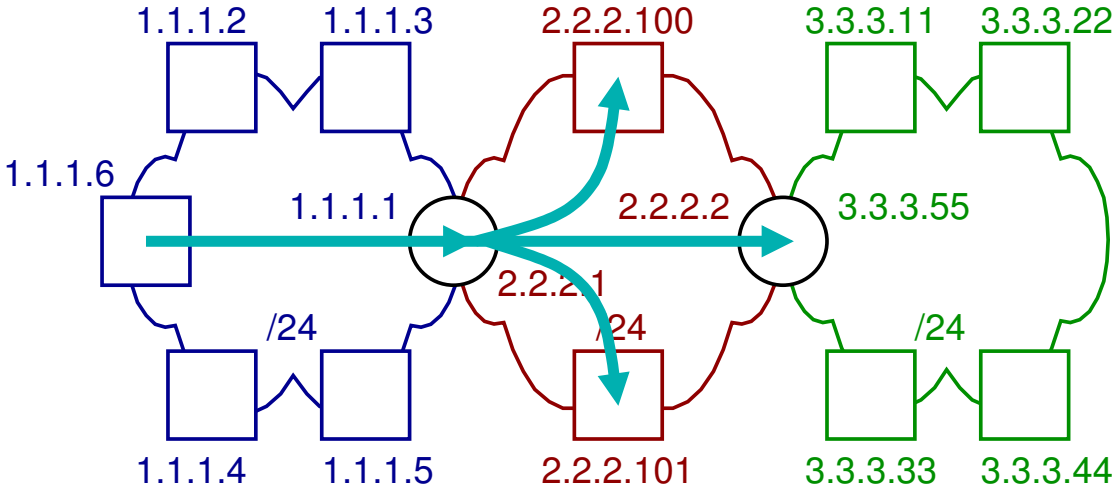

Example of Startup Source and Local Broadcast

IP Datagram

Header	Data

Src = 0.0.0.0
Dst = 255.255.255.255

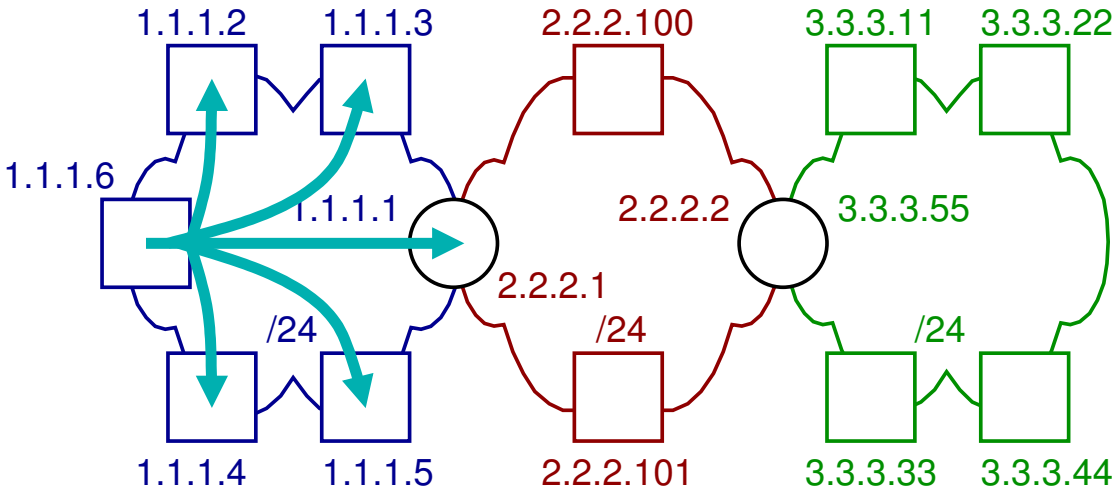

Example of Loopback Address

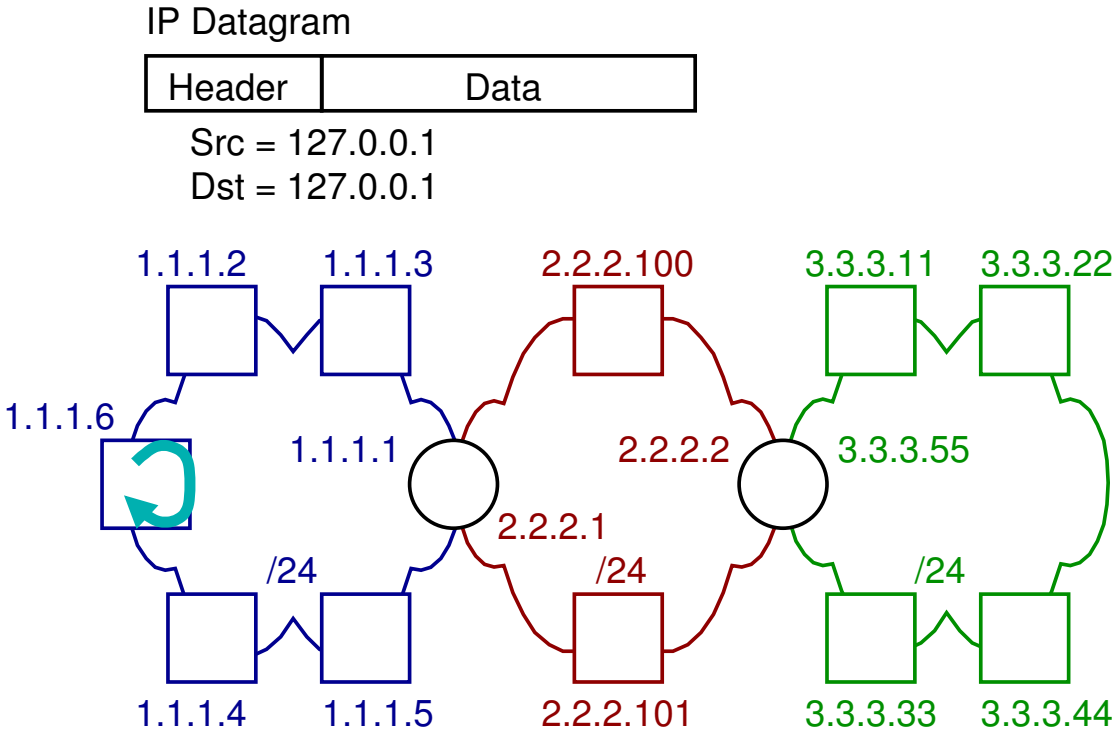

IP Addressing Example

My office computer has address `104.209.61.169/18`. What is the network address and directed broadcast address for my network? How many IP devices can be attached to my network?

Obtaining an IP Address

- Internet Assigned Numbers Authority (IANA) manages the assignment of IP addresses
- IANA delegates IP network ranges to regional authorities (e.g. APNIC), delegated further to national registries (e.g. THNIC)
- Organisations obtain network addresses from national/local registries
- Organisations are free to assign addresses as they wish from assigned network address
  - Manually set IP address on each computer
  - Protocol to automatically configure IP addresses in computers on network: Dynamic Host Configuration Protocol

Contents

Internet Applications

- Most Internet applications follow a client/server model of initiating communication:
  1. Server waits for client to initiate communication
  2. Client initiates communication
  3. Once the communication is initiated, data can flow in both directions (client to server and server to client)
- Examples:
  - Web browser (Firefox, Safari) and web server (Apache, IIS)
  - Email client (Thunderbird, Outlook) and email server (MS Exchange, Postfix)
  - Instant messaging client and server (LINE, MSN, TextSecure)
  - Bittorrent (uTorrent, Transmission) and tracker (Opentracker, VUZE)

Issues with Client/Server Applications

- How to make it easy for programmers to create applications without knowing details of communications?
  - Transport protocols implement features common to many apps, e.g. TCP, UDP
- How to allow applications implemented in different languages/OS by different people to communicate?
  - Application layer protocols, e.g. HTTP, SMTP, FTP
  - Use a common API: Sockets
- How to identify different applications on same computer?
  - Addresses to identify applications: Ports

Transport Protocols

- Send data between application processes on source and destination hosts
- End-to-end (or host-to-host) communications
- Transmission Control Protocol
  - Most widely used transport protocol
  - Connection-oriented, error control, flow control, congestion control
- Others: User Datagram Protocol (UDP), SCTP, DCCP, old and domain-specific protocols
- Protocol number: identifies transport protocol used by both hosts
  - 8-bit number; e.g. 6 = TCP, 17 = UDP; 1 = ICMP
  - Included in IP header

`http://www.iana.org/assignments/protocol-numbers/`

How does a client identify a server application?

- Internet contains multiple hosts
  - Host (interface) identified by IP address
- A host may implement multiple transport protocols
  - Transport protocol identified by protocol number
- Multiple applications may use same transport protocol
  - Ports identify application processes on a host
- Five addresses uniquely identify end-to-end communications
  1. Source IP
  2. Destination IP
  3. Protocol number
  4. Source port
  5. Destination port

Multiple Applications, Multiple Transport Protocols

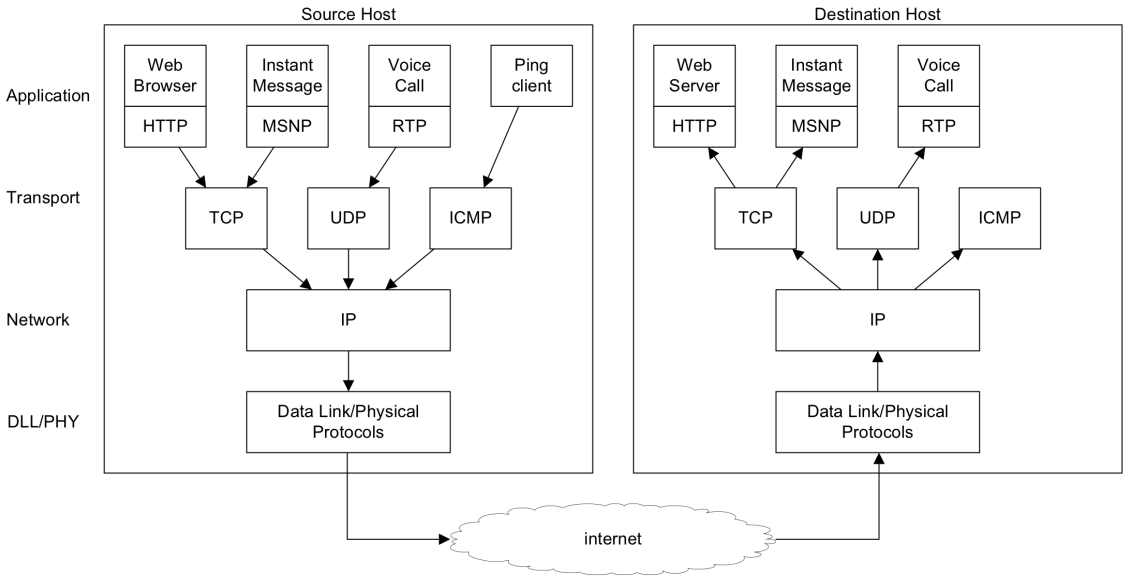

Port Numbers

- Ports are 16-bit numbers
- Source port, destination port in transport protocol header
- On a host, ports are managed by operating system
  - Unique port assigned to processes for Internet communications
  - Ports are local to a host
- Well-known ports: 0–1023
  - Common servers use well-known ports by default
  - http = 80, https = 443, ssh = 22, ftp = 20/21, smtp = 25, dns = 53, dhcp = 67, ipp = 631
- Registed ports: 0–49151
  - Servers use registed ports by default
  - openvpn = 1094, mysql = 3306, steam = 27015, ...
- Dynamic ports: 49152–65535
  - Clients use dynamic ports, assigned by OS

`http://www.iana.org/assignments/port-numbers/`

Contents

Transmission Control Protocol

- Most commonly used transport protocol today
  - Web browsing, email, file sharing, instant messaging, file transfer, database access, proprietary business applications, some multimedia applications (at least for control purposes), ...
- Services provided by TCP:
  - Stream-oriented: TCP treats data from application as continuous stream of bytes, sequence numbers count bytes
  - Connection-oriented: setup connection before data transfer
  - Full duplex connection: send data in either direction
  - Flow and error control: Go-Back-N style
  - Congestion control: if network congestion, source slows down

TCP Segment

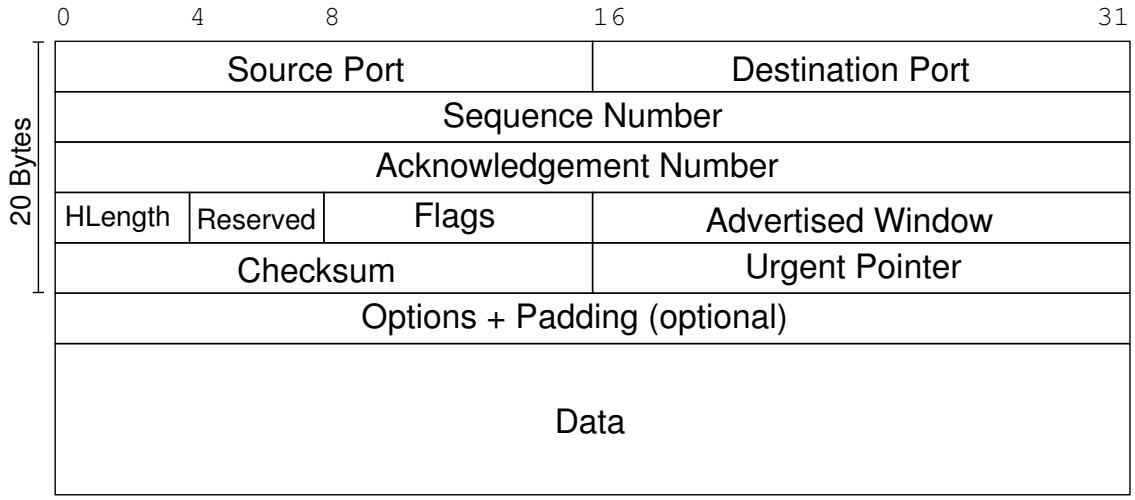

- Header contains 20 bytes, plus optional fields
- Optional fields must be padded out to multiple of 4 bytes

TCP Segment Fields

- Source/Destination port
- Sequence number of the first data byte in this segment (or ISN)
- Acknowledgement number: sequence number of the next data byte TCP expects to receive
- Header Length: Size of header (measured in 4 bytes)
- Window: number of bytes the receiver is willing to accept (for flow control)
- Checksum: error detection on TCP segment
- Urgent pointer points to the sequence number of the last byte of urgent data in the segment
- Options: such as maximum segment size, window scaling, selective acknowledgement, ...

TCP Segment Flags

- Flags (1 bit each, if 1 the flag is true or on):
- CWR: Congestion Window Reduced
- ECE: Explicit Congestion Notification Echo
- URG: segment carries urgent data, use the urgent pointer field; receiver should notify application program of urgent data as soon as possible
- ACK: segment carries ACK, use the ACK field
- PSH: push function
- RST: reset the connection
- SYN: synchronise the sequence numbers
- FIN: no more data from sender

TCP Connection Establishment: Three-Way Handshake

Agree upon initial sequence numbers, prepare buffer for data

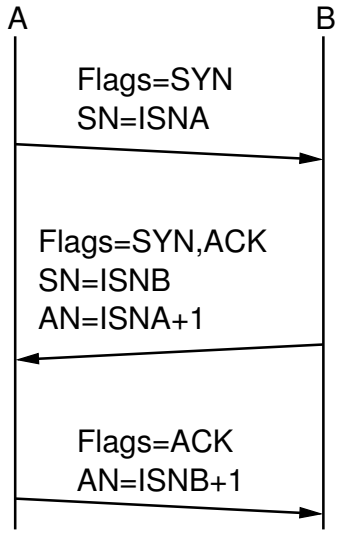

- Initiator A selects an Initial Sequence Number, *ISNA*
- B acknowledges *ISNA* and also chooses its own *ISNAB*
- Data transfer can start after *ISNB* is ACKed
- Optionally, 3rd segment can contain data

TCP Data Transfer

- Segments can contain varying amount of data
- Set ACK flag to indicate an acknowledgement, piggybacking is common
- Speed of data transfer depends on:
  - Flow control: sliding-window
  - Error control: Go-Back-N style
  - Congestion control: loss of segments indicates congestion, sender slows down

Contents

Application Layer Protocols

- Many different protocols to support types of applications
  - HTTP, FTP, SMTP, SSH, Telnet, BitTorrent, SIP, IMAP, RDP, SMB, ...
- Other protocols to support network operation
  - DNS, DHCP/BOOTP, NTP, SNMP, ...