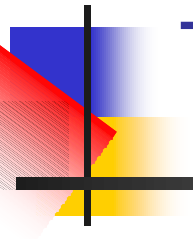

Transporting Voice by Using IP

Internet Overview

- A collection of networks
  - The private networks
    - LANs, WANs
    - Institutions, corporations, business and government
    - May use various communication protocols
  - The public networks
    - ISP: Internet Service Provider
    - Using Internet Protocol (IP)
- To connect to the Internet
  - Using IP
  - Routers provide the connectivity between various networks and form the backbone of the Internet.

Interconnecting Networks

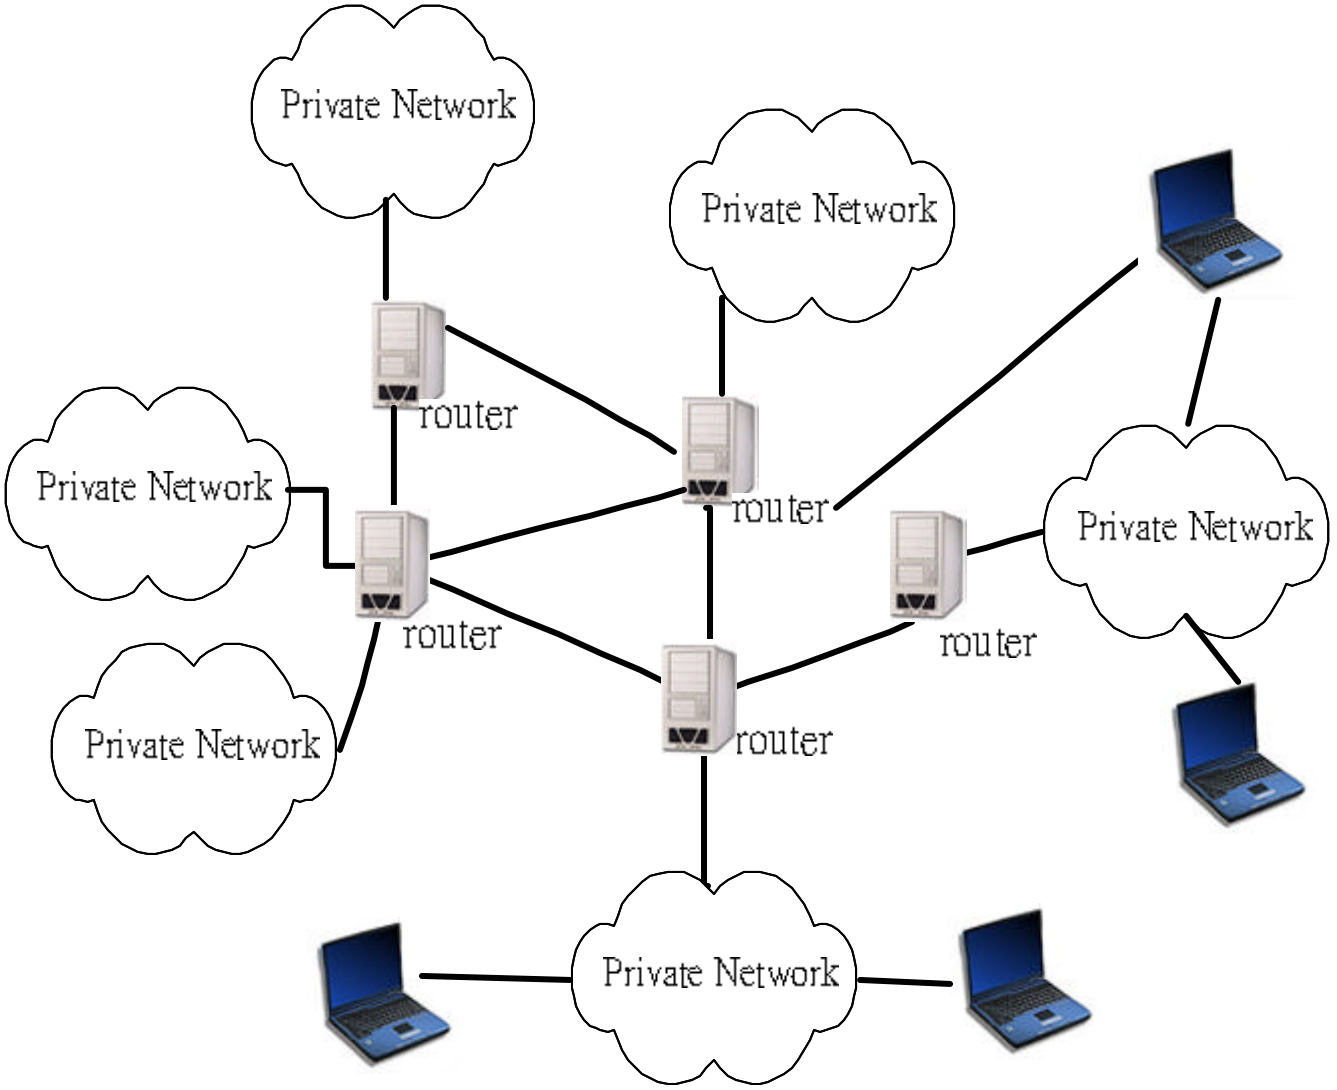

Overview of the IP Protocol Suite

- IP
  - A routing protocol for the passing of data packets
  - Must work in cooperation with higher layer protocols (application or service) and lower-layer transmission systems
- The OSI (Open System Interconnection) seven-layer model
  - The top layer: useable information to be passed to the other side
  - The information must be
    - Packaged appropriately
    - Routed correctly
    - Traverse some physical medium

OSI Model [1/3]

- Physical layer
  - The physical media
  - Coding and modulation schemes for 1's and 0's
- Data link layer
  - Transport the information over a single link
  - Frame packaging, error detection/correction and retransmission
- Network layer
  - Routing traffic from the source to the destination
  - Passing through intermediate points

OSI Model [2/3]

- Transport layer
  - Ensure error-free, omission-free and in-sequence delivery
  - Support multiple streams from the source to the destination for applications
    - To identify each individual stream and ensure that each stream is passed upward to the correct application
- Session layer
  - The commencement (e.g., login) and completion (e.g., logout) of a session between applications
  - Establish the dialogue
    - One way at a time or both ways at the same time

OSI Model [3/3]

- Presentation layer
  - Specify the language to be used between applications
    - Encoding
- Application layer
  - Provide an interface to the user
  - File transfer programs and web browsers

The IP suite and the OSI stack

- Transmission Control Protocol (TCP)
  - Reliable, error-free, in-sequence delivery
- User Datagram Protocol (UDP)
  - No sequencing, no retransmission

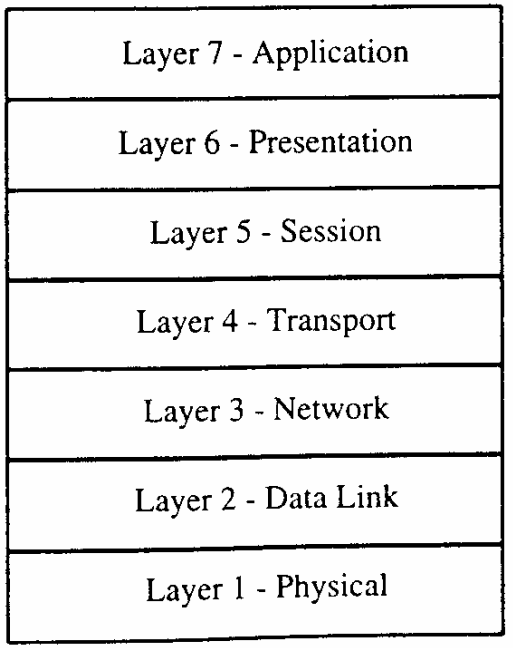

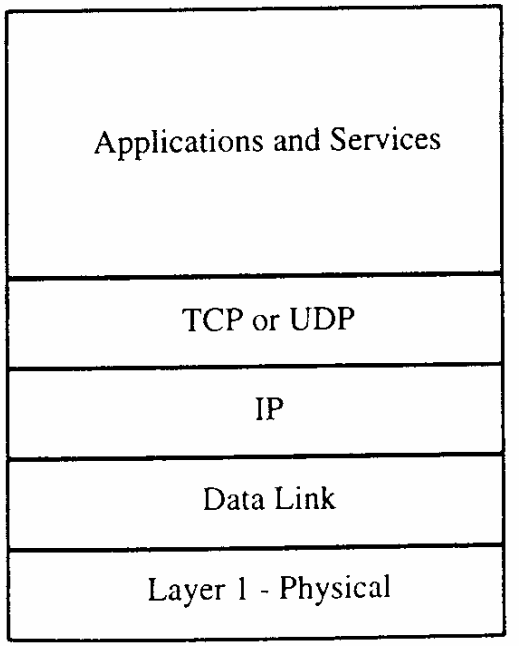

Internet Society [1/3]

- A non-profit organization
- Keep the Internet alive and growing
- “To assure the open development, evolution, and the use of Internet for the benefit of all people throughout the world”
- The tasks include
  - Supporting the development and dissemination of Internet standards
  - Supporting the RD related to the Internet and internetworking
  - Assisting developing countries in implementing Internet Infrastructure
  - Forming liaisons with other organizations for education, collaboration, and coordination related to Internet development

Internet Society [2/3]

- Internet Architecture Board (IAB)
  - The technical advisory group
  - Providing technical guidance to Internet Society
  - Overseeing the Internet standards process
- Internet Engineering Task Force (IETF)
  - Comprising a huge number of volunteers
    - Equipment vendors, network operators, research institutions etc.
  - Developing Internet standards
  - Detailed technical work
  - Working groups (addressing specific areas or topics)
    - Most of the work is done by individuals, and shares with others through the use of mailing lists.
    - The IETF as a whole meets three times a year.
    - megaco, iptel, sip, sigtran

Internet Society [3/3]

- Internet Engineering Steering Group (IESG)
  - Managing the IETF's activities
  - Approving an official standard
- Internet Assigned Numbers Authority (IANA)
  - Administration of unique numbers and parameters used in Internet standards
  - Be registered with the IANA

Internet Standards Process

- The process is documented in RFC 2026.
- First, Internet Draft
  - The early version of spec.
  - Can be updated, replaced, or made obsolete by another document at any time
  - IETF's Internet Drafts directory
  - Six-month life-time
  - Any reference made to an Internet draft must be done with great care and with emphasis that the draft is a work in progress.

Internet Standards Process [1/2]

- RFC
  - An Internet draft is considered sufficiently complete
  - Request for Comments
  - An RFC number
- Proposed standard
  - A stable, complete, and well-understood spec.
  - Has garnered significant interest
- Draft standard
  - Two independently successful implementations
  - Interoperability must have been be demonstrated.

Internet Standards Process [2/2]

- A standard
  - The IESG is satisfied that the spec. is stable and mature and can be readily and successfully implemented on a large scale.
  - Significant operational experience
  - A standard (STD) number
- Not all RFCs are standards, nor do all RFCs document technical specifications
  - Some document best current practices (BCPs)
    - To outline processes, policies, or operational considerations related to the Internet
  - Others are known as applicability statements
    - How a spec. be used to achieve a particular goal, or different specs work together

Internet Protocol (IP)

- RFC 791
  - Amendments: RFCs 950, 919, and 920
  - Requirements for Internet hosts: RFCs 1122, 1123
  - Requirements for IP routers: RFC 1812
- IP datagram
  - Data packet with an IP header that contains information about the originator and destination addresses.
  - The header information is used by routers to route the packet to its destination.
- Best-effort protocol
  - No guarantee that a given packet will be delivered
  - A packet may be lost due to transmission errors, congestion in buffers or transmission facilities, link failures, and so on.

IP Header [1/2]

- Version 4
- Header Length
- Type of Service
- Total Length
- Identification, Flags, and Fragment Offset
  - A datagram can be split into fragments in the case where the size of the datagrams is greater than the maximum that can be handled by a given link
  - Identify data fragments that belong together
  - Flags
    - a datagram can be fragmented or not
    - Indicate the last fragment
  - The fragment offset is a number describing where the fragment belongs in the overall datagram.
    - Enabling the destination to put the different pieces together correctly

IP Header [2/2]

- TTL
  - A number of hops (not a number of seconds)
- Protocol
  - The higher-layer protocol
  - TCP (6); UDP (17)
- Source and Destination IP Addresses

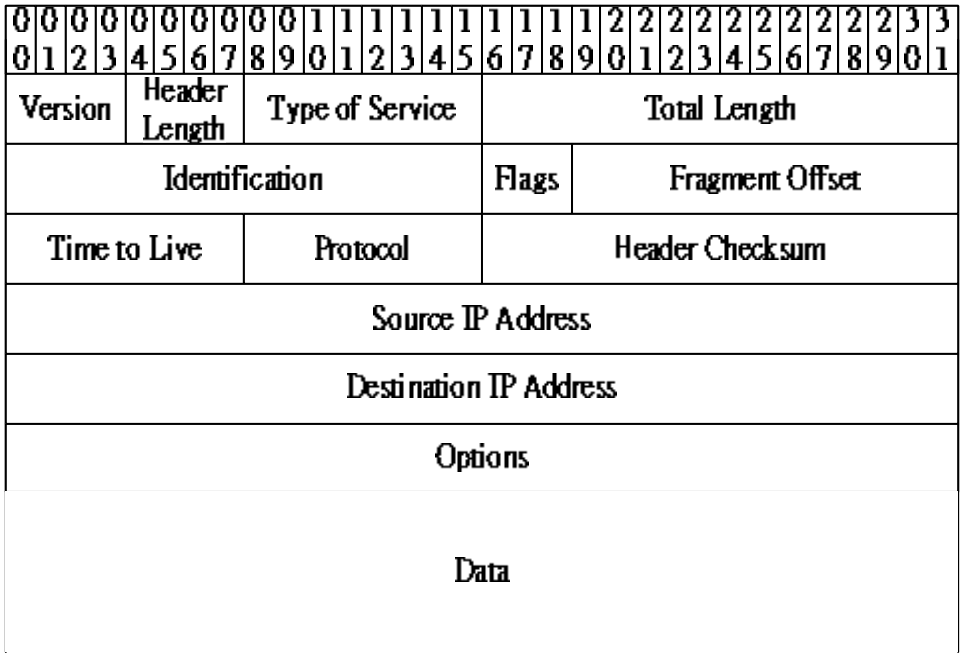

IP Routing

- Based on the destination address in the IP header
- Routers
  - Can contain a range of different interfaces
  - Determine the best outgoing interface for a given IP datagram
  - Routing table
    - Destination
    - IP route mask
    - For example, any address starting with 182.16.16 should be routed on interface A. (IP route mask 255.255.255.0)
    - Longest match

Populating Routing Tables

- Issues
  - The correct information in the first place
  - Keep the information current in a dynamic environment
  - The best path?
- Protocols
  - OSPF (Open Short Path First)
    - An AS (Autonomous System) is a group of routers that share routing information between them.
      - The AS is further divided into areas.
    - Area 0: backbone area
    - Border router
  - BGP (Border Gateway Protocol)

OSPF Areas

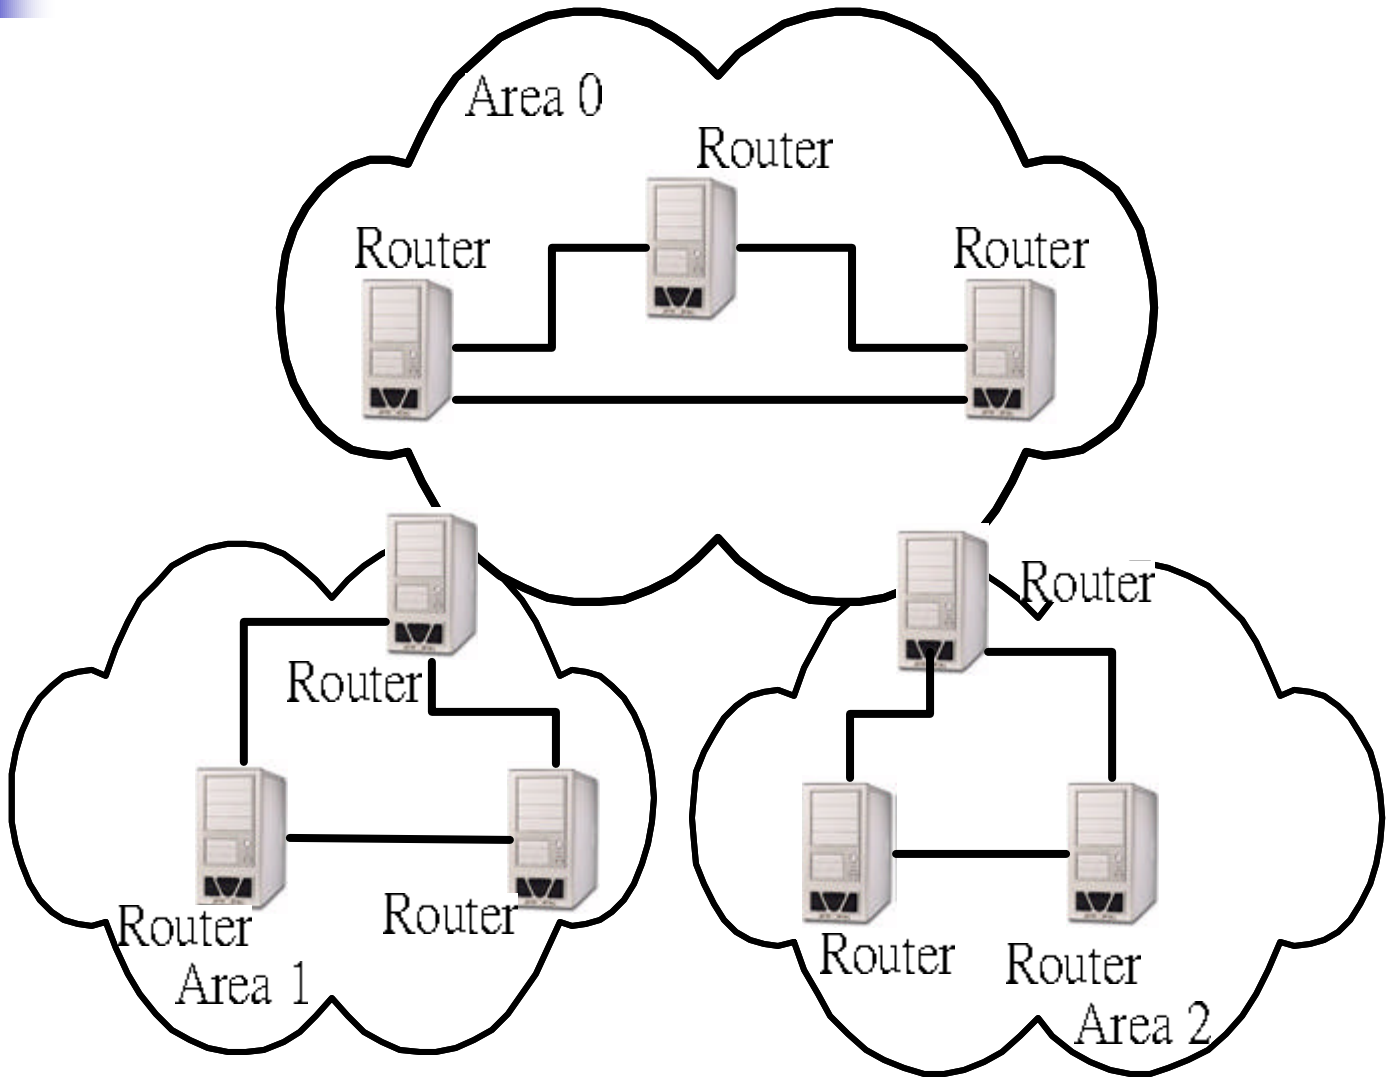
Area 0
Router
Router
Router
Router
Router
Router
Router
Router
Area 1
Router
Router
Area 2

Transmission Control Protocol (TCP)

- In sequence, without omissions and errors
- End-to-end confirmation, packet retransmission flow control, congestion control
- RFC 793
- Break up a data stream in segments
- Attach a TCP header
- Sent down the stack to IP
- At the destination, checks the header for errors
  - Send back an ack
- The source retransmits if no ack is received within a given period.

The TCP Header [1/5]

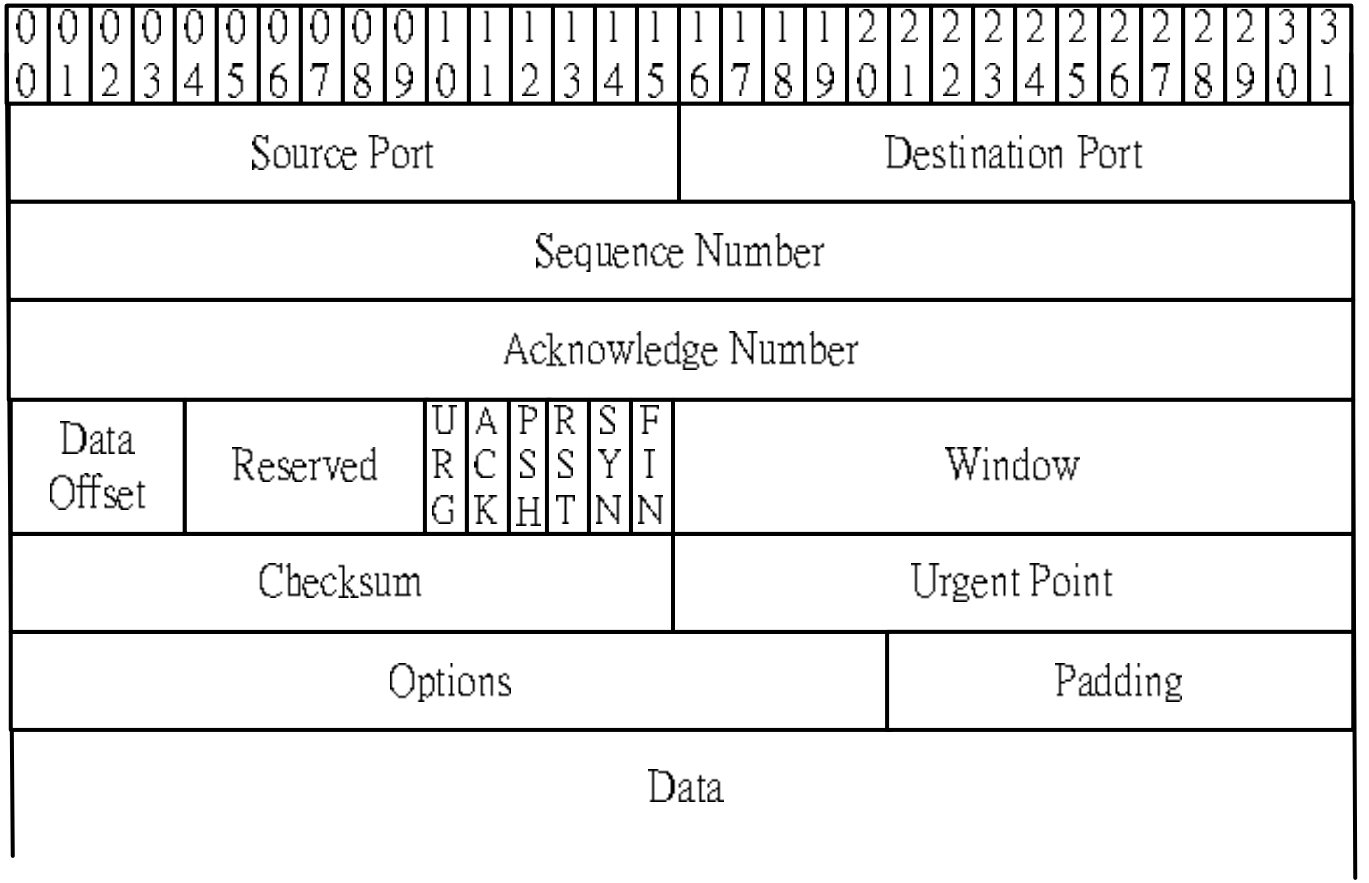

The TCP Header [2/5]

- TCP Port Numbers
  - Identifying a specific instance of a given application
  - A unique port number for a particular session
  - Well-known port numbers
    - IANA, 0-1023
    - 23, telnet; 25, SMTP
  - Many clients and a server
    - TCP/IP
    - Source address and port number + Destination address and port number
    - A socket address (or a transport address)

The TCP Header [3/5]

- Sequence and acknowledge numbers
  - Identify individual segments
  - Actually count data octets transmitted
  - A given segment with a SN of 100 and contains 150 octets of data
    - The ack number will be 250
    - The SN of the next segment is 250
- Other header fields
  - Data offset: header length (in 32-bit words)
  - URG: 1 if urgent data is included, use urgent pointer field

The TCP Header [4/5]

- ACK: 1, an ACK
- PSH: for the push function
- RST: reset; an error and abort a session
- SYN: Synchronize; the initial messages
- FIN: Finish; close a session
- Window
  - The amount of buffer space available for receiving data
- Checksum

The TCP Header [5/5]

- Urgent Pointer
  - An offset to the first segment after the urgent data
  - Indicates the length of the urgent data
  - Critical information to be sent to the user application ASAP

TCP Connections

- An example
- After receiving
  - 100, 200, 300
  - ACK 400
- Closing a connection
  - → FIN
  - ← ACK, FIN
  - → ACK

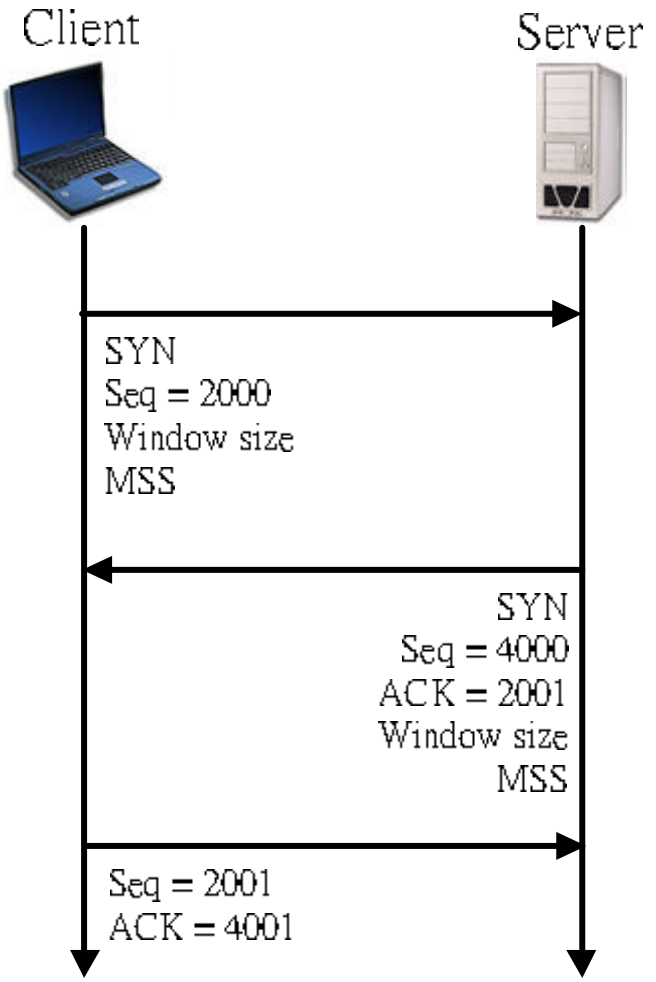

UDP

- User Datagram Protocol
  - Pass individual pieces of data from an application to IP
  - No ACK, inherently unreliable
  - Applications
    - A quick, on-shot transmission of data, request/response
    - DNS
    - If no response, the AP retransmits the request
    - The AP includes a request identifier
  - The source port number is optional
  - Checksum

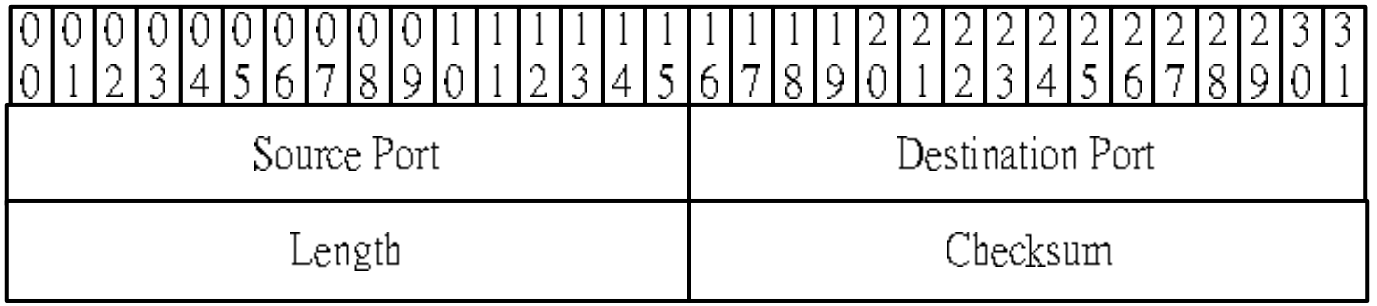

Voice over UDP, not TCP

- Speech
  - Small packets, 10 – 40 ms
  - Occasional packet loss is not a catastrophe
  - Delay-sensitive
    - TCP: connection set-up, ack, retransmit → delays
  - 5 % packet loss is acceptable if evenly spaced
    - Resource management and reservation techniques
    - A managed IP network
  - In-sequence delivery
    - Mostly yes
- UDP was not designed for voice traffic

Real-Time Transport Protocol

- RTP: A Transport Protocol for Real-Time Applications
  - RFC 1889
  - RTP – Real-Time Transport Protocol
  - RTCP – RTP Control Protocol
- UDP
  - Packets may be lost or out-of-sequence
- RTP over UDP
  - A sequence number
  - A time stamp for synchronized play-out
  - Does not solve the problems; simply provides additional information

RTCP

- A companion protocol
- Exchange messages between session users
- # of lost packets, delay and inter-arrival jitter
- Quality feedback
- RTCP is implicitly open when an RTP session is open
- E.g., RTP/RTCP uses UDP port 5004/5005

RTP Payload Formats [1/2]

- RTP carries the actual digitally encoded voice
  - RTP header + a payload of voice/video samples
  - UDP and IP headers are attached
- Many voice- and video-coding standards
  - A payload type identifier in the RTP header
    - Specified in RFC 1890
    - New coding schemes have become available
    - See Table 2-1 and Table 2-2
  - A sender has no idea what coding schemes a receiver could handle.

RTP Payload Formats [2/2]

- Separate signaling systems
  - Capability negotiation during the call setup
  - SIP and SDP
  - A dynamic payload type may be used
    - Support new coding scheme in the future
    - The encoding name is also significant.
      - Unambiguously refer to a particular payload specification
      - Should be registered with the IANA
- RED, Redundant payload type
  - Voice samples + previous samples
  - May use different encoding schemes
  - Cope with packet loss

RTP Header Format

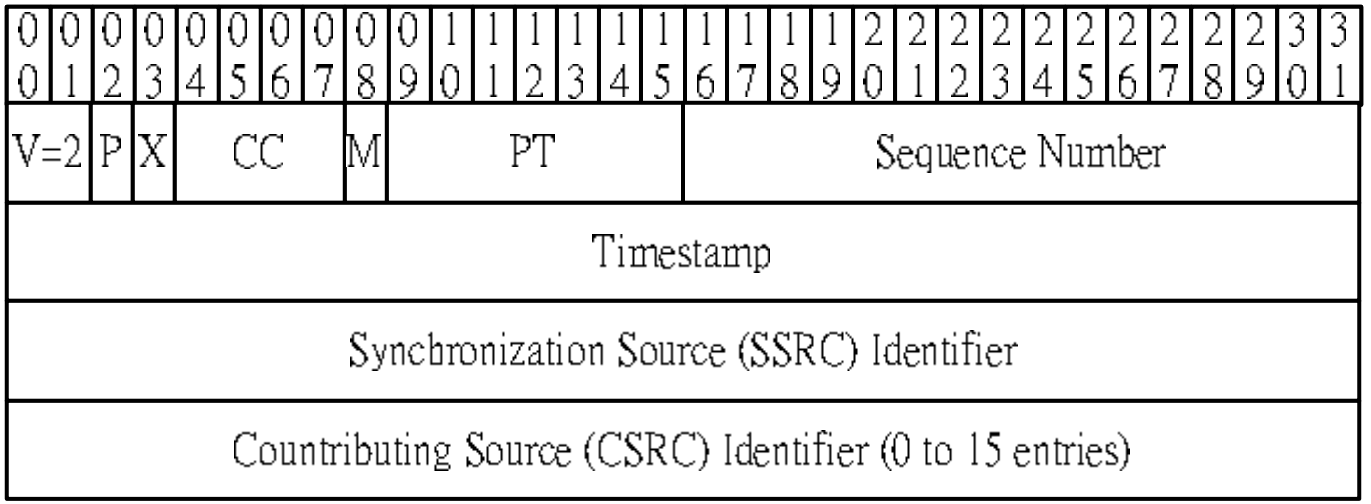

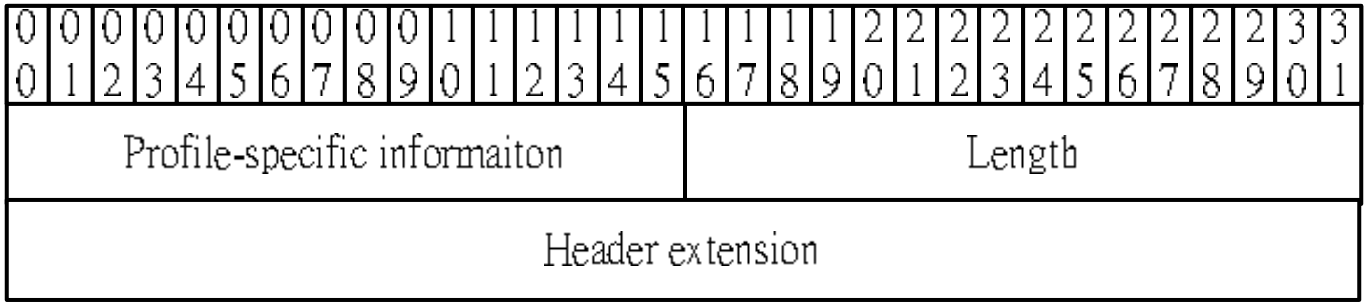

The RTP Header [1/4]

- Version (V)
  - 2
- Padding (P)
  - The padding octets at the end of the payload
  - The payload needs to align with 32-bit boundary
  - The last octet of the payload contains a count of the padding octets.
- Extension (X)
  - 1, contains a header extension

The RTP Header [2/4]

- CSRC Count (CC)
  - The number of contributing source identifiers
- Marker (M)
  - Support silence suppression
  - The first packet of a talkspurt, after a silence period
- Payload Type (PT)
  - In general, a single RTP packet will contain media coded according to only one payload format.
  - RED is an exception.
- Sequence number
  - A random number generated by the sender at the beginning of a session
  - Incremented by one for each RTP packet

The RTP Header [3/4]

- Timestamp
  - 32-bit
  - The instant at which the first sample in the payload was generated
  - The receiver
    - Synchronized play-out
    - Calculate the jitter
    - The clock freq depends on the encoding
      - E.g., 8000Hz
    - Support silence suppression
    - The initial timestamp is a random number chosen by the sending application.

The RTP Header [4/4]

- Synchronization Source (SSRC)
  - 32-bit identifier
  - The entity setting the sequence number and timestamp
    - Normally the sender of the RTP packet
  - Chosen randomly, independent of the network address
    - Meant to be globally unique within a session
  - May be a sender or a mixer
- Contributing Source (CSRC)
  - An SSRC value for a contributor
  - Used to identify the original sources of media behind the mixer
  - 0-15 CSRC entries
- RTP Header Extensions

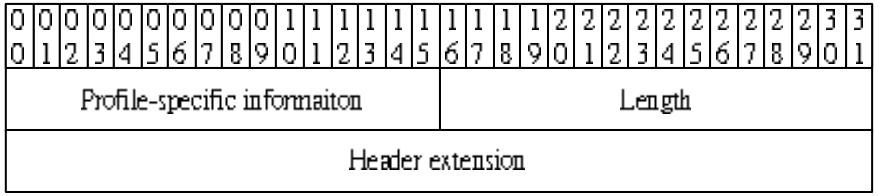

Mixers and Translators

- Mixers
  - Enable multiple media streams from different sources to be combined into a single stream
    - If the capacity or bandwidth of a participant is limited
  - An audio conference
  - The SSRC is the mixer
    - More than one CSRC values
- Translators
  - Manage communications between entities that does not support the same coding scheme
  - The SSRC is the participant, not the translator.