Transporting Voice by Using IP

Chapter 2

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 Providers
    - Using Internet Protocol
  - To connect to the Internet
    - Using IP

Interconnecting Networks

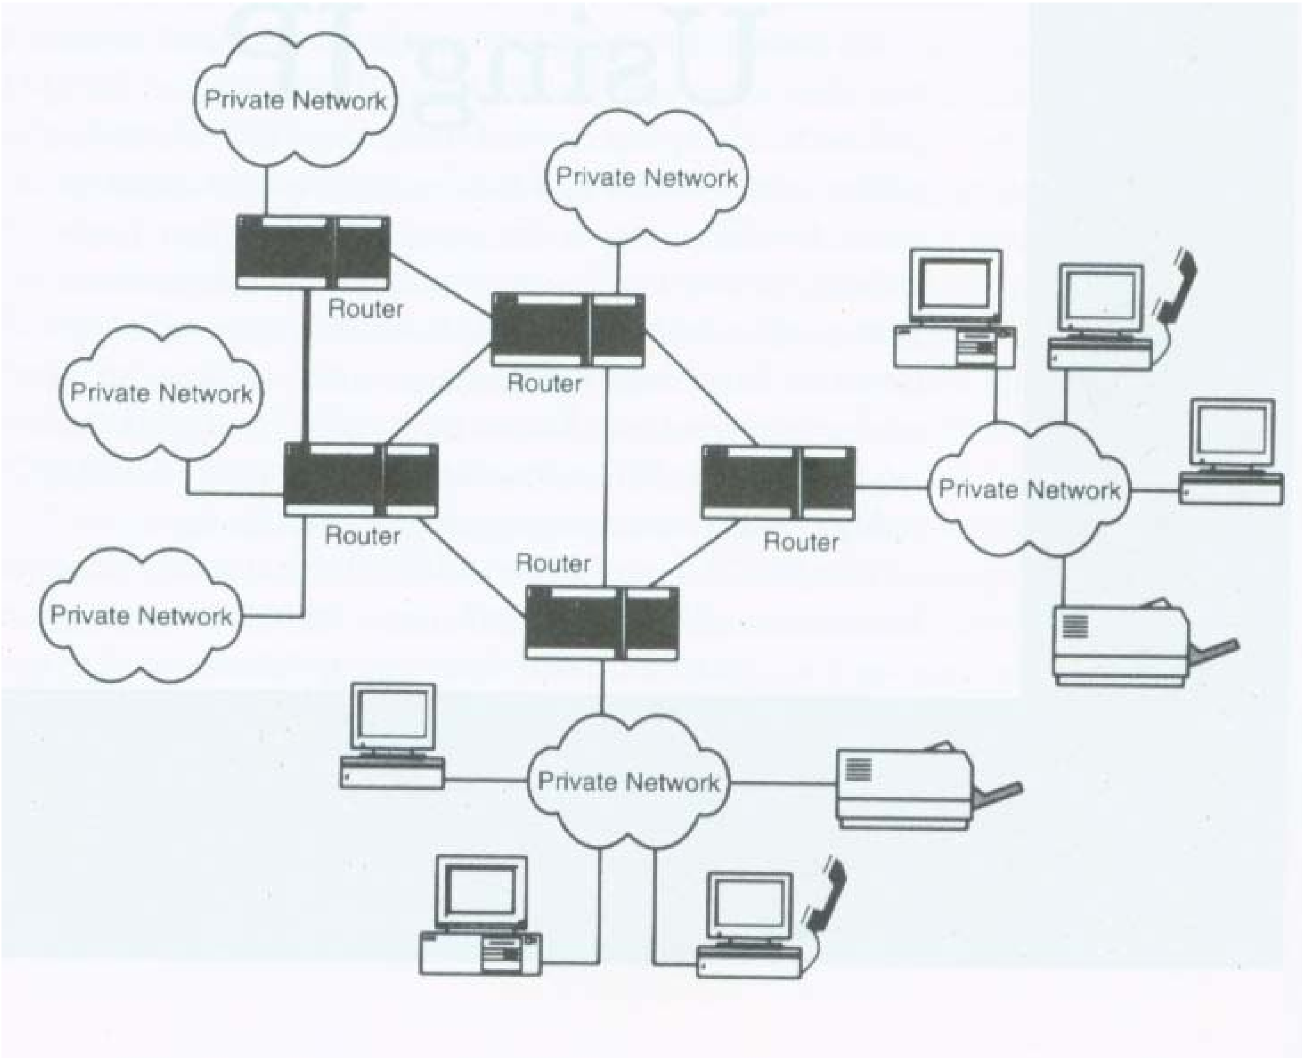

Overview of the IP Protocol Suite

- IP
  - A routing protocol for the passing of data packets
  - Must work in cooperation with higher layer protocols and lower-layer transmission systems
- The OSI seven-layer model
  - The top layer: useable information to be passed to the other side
  - The information must be
    - Packaged appropriately
    - Routed correctly
    - And it must 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 through a network
  - 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 destination for applications
- 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, the encoding and so on
- Application layer
  - Provide an interface to the user
  - File transfer programs and web browsers

The IP suite and the OSI stack

- TCP
  - Reliable, error-free, in-sequence delivery
- UDP
  - No sequencing, no retransmission

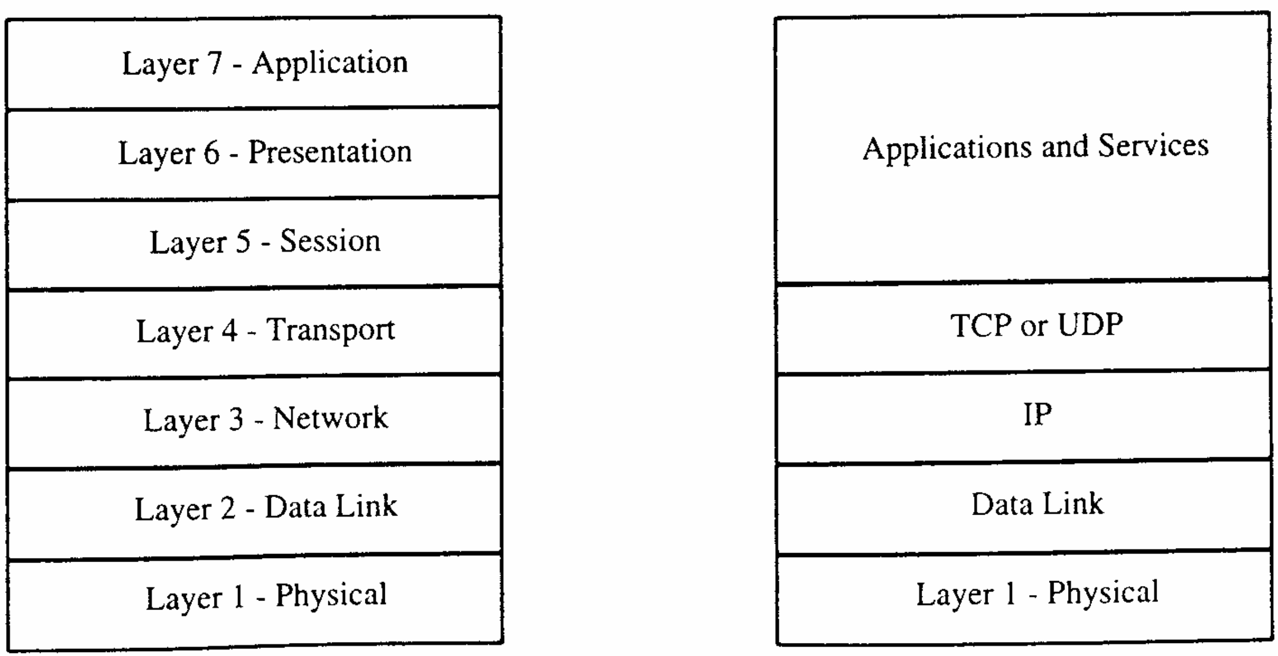

Internet Standards and the Process

- The Internet Society
  - 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

Internet Standards and the Process

- IAB
  - The Internet Architecture Board
  - The technical advisory group
  - Providing technical guidance to Internet Society
  - Overseeing the Internet standards process
- IETF
  - The Internet Engineering Task Force
  - Comprising a huge number of volunteers
    - Equipment vendors, network operators, research institutions etc.
  - Developing Internet standards
  - Detailed technical work
  - Working groups
    - megaco, iptel, sip, sigtran

Internet Standards and the Process

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

The Internet Standards Process

- The process
  - 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

The Internet Standards Process

- RFC
  - 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 be demonstrated

The Internet Standards Process

- A standard
  - The IESG is satisfied
  - The spec. is stable and mature
  - Significant operational experience
  - A standard (STD) number
- Not all RFCs are standards
  - Some document Best Current Practices (BCPs)
    - Processes, policies, or operational considerations
  - Others applicability statements
    - How a spec be used, or different specs work together

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
  - Best-effort protocol
    - No guarantee that a given packet will be delivered

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
  - Identify data fragments
  - Flags
    - a datagram can be fragmented or not
    - Indicate the last fragment
- TTL
  - A number of hops (not a number of seconds)

IP Header [2/2]

- Protocol
  - The higher-layer protocol
  - TCP (6); UDP (17)
- Source and Destination IP Addresses

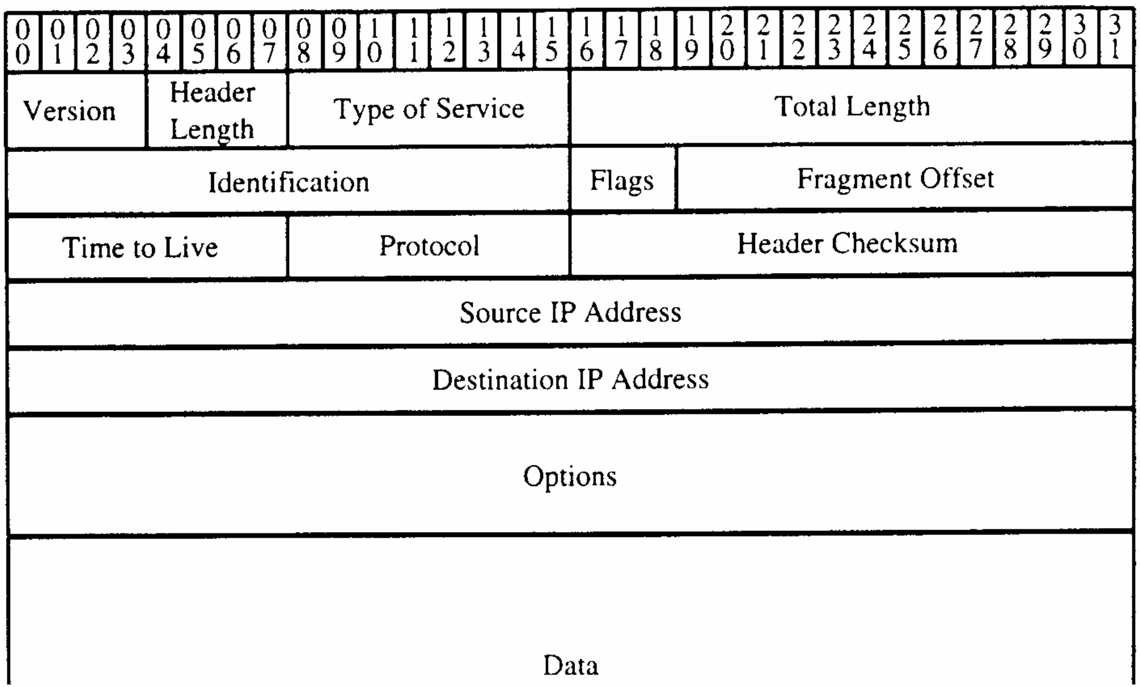

IP Routing

- 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)

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.
    - Area 0: backbone area
    - Border router
  - BGP (Border Gateway Protocol)

OSPF Areas

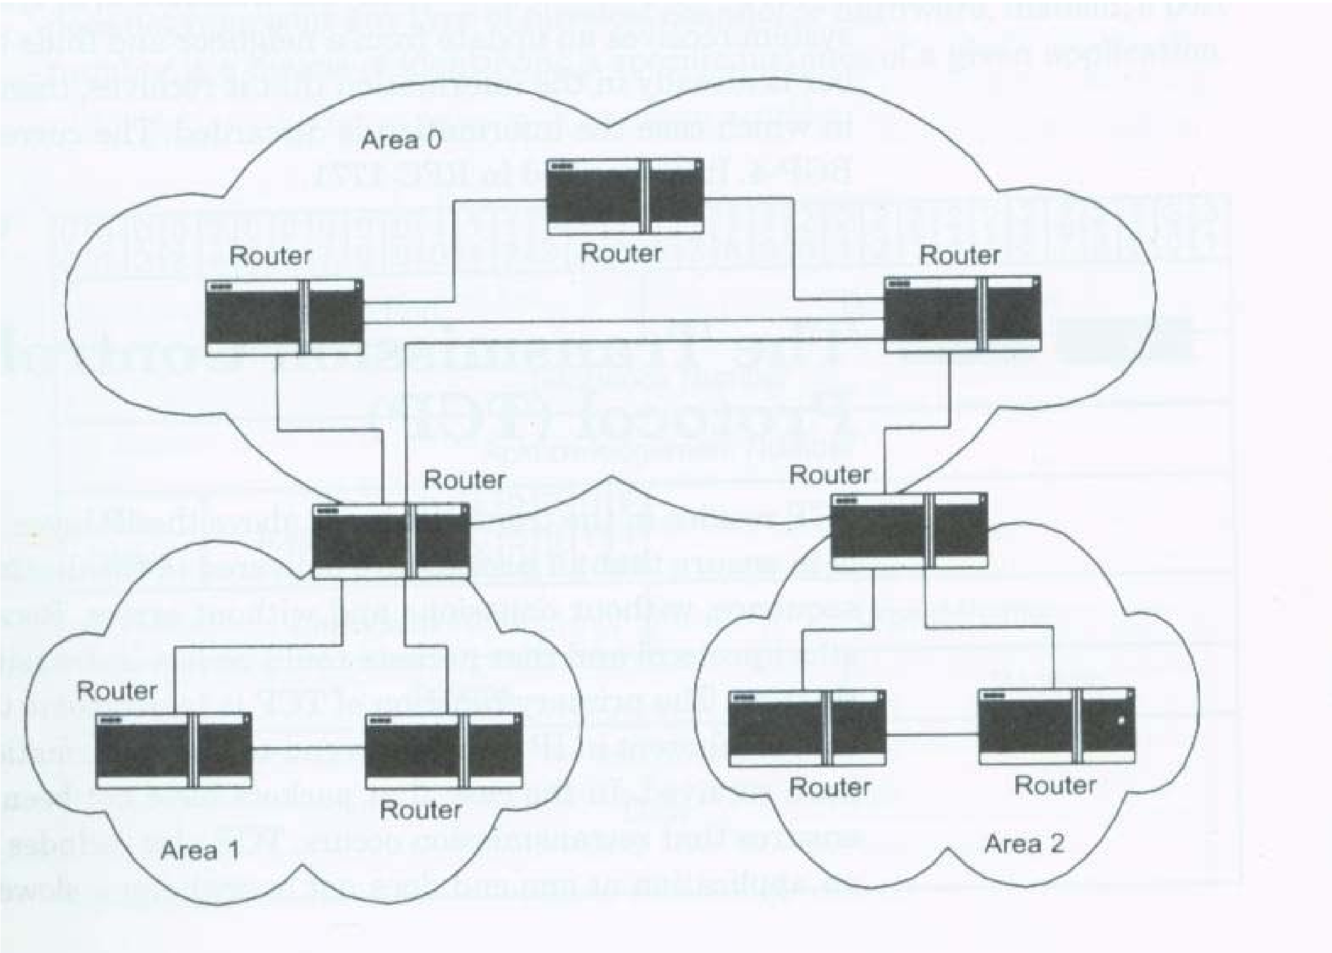

Area 0
Router
Router
Router
Router
Router
Router
Router
Router
Router
Area 1
Area 2

TCP

- Transmission Control Protocol
  - In sequence, without omissions and errors
  - End-to-end confirmation, packet retransmission, Flow 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 within a given period

The TCP Header [1/5]

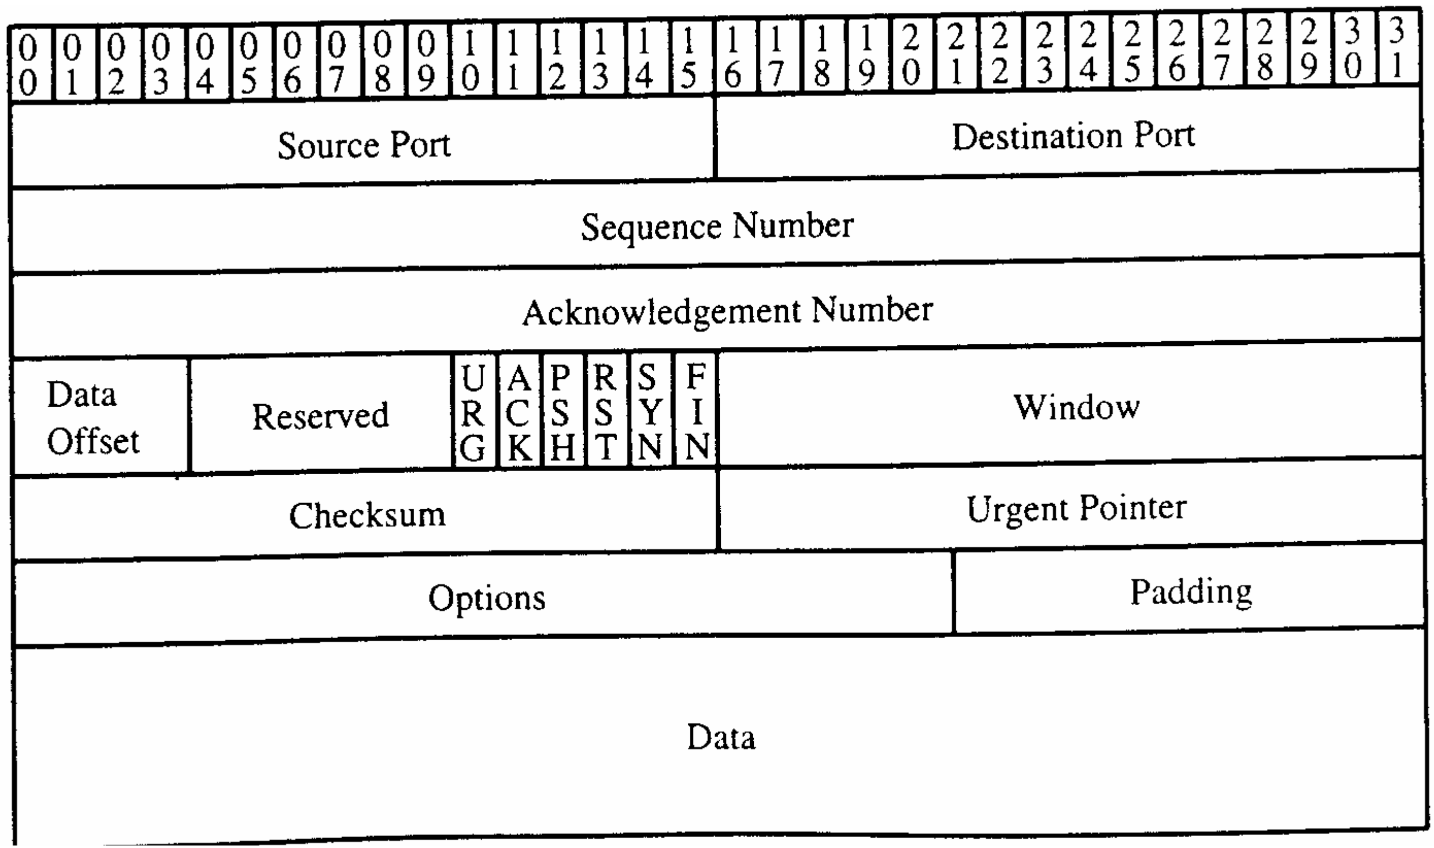

The TCP Header [2/5]

- 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
  - ACK: 1, an ACK
  - PSH: a push function, be delivered promptly

The TCP Header [4/5]

- 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

<table>
<tr><td>00</td><td>01</td><td>02</td><td>03</td><td>04</td><td>05</td><td>06</td><td>07</td><td>08</td><td>09</td><td>10</td><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>17</td><td>18</td><td>19</td><td>20</td><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>27</td><td>28</td><td>29</td><td>30</td><td>31</td></tr>
<tr><td colspan="32">Source IP Address</td></tr>
<tr><td colspan="32">Destination IP Address</td></tr>
<tr><td colspan="8">0</td><td colspan="8">Protocol (6 for TCP)</td><td colspan="16">TCP Length</td></tr>
</table>

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

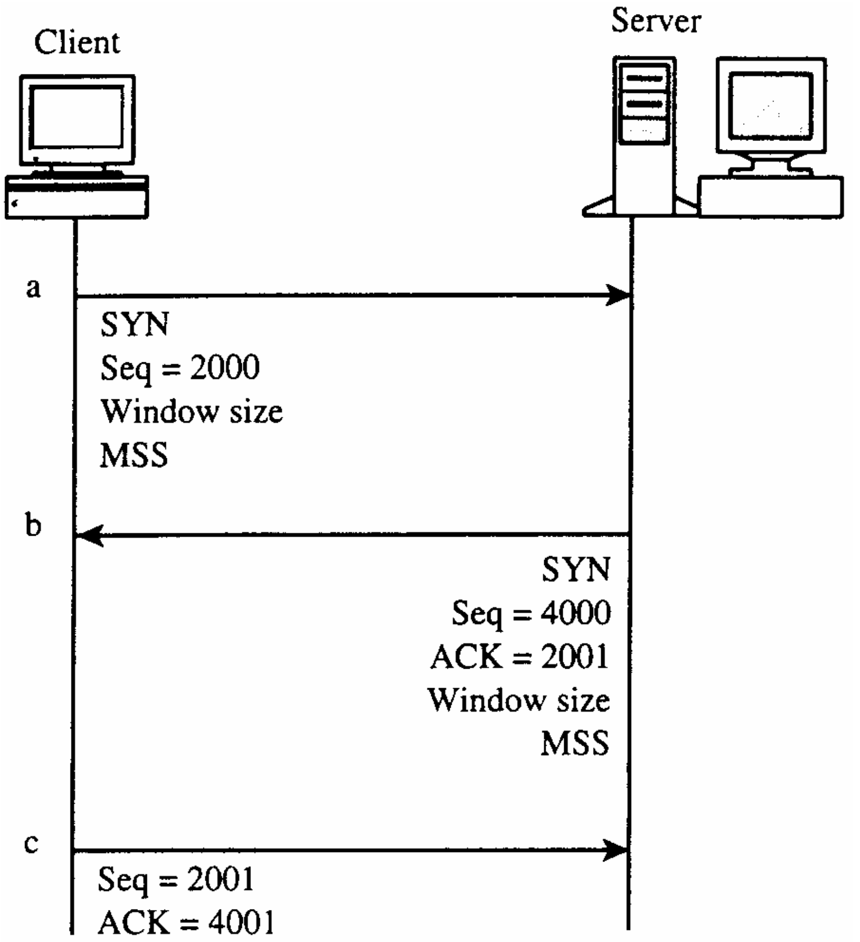

UDP

- 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

<table>
<tr><td>00</td><td>01</td><td>02</td><td>03</td><td>04</td><td>05</td><td>06</td><td>07</td><td>08</td><td>09</td><td>10</td><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>17</td><td>18</td><td>19</td><td>20</td><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>27</td><td>28</td><td>29</td><td>30</td><td>31</td></tr>
<tr><td colspan="16">Source Port</td><td colspan="16">Destination Port</td></tr>
<tr><td colspan="16">Length</td><td colspan="16">Checksum</td></tr>
</table>

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

The 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
  - 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

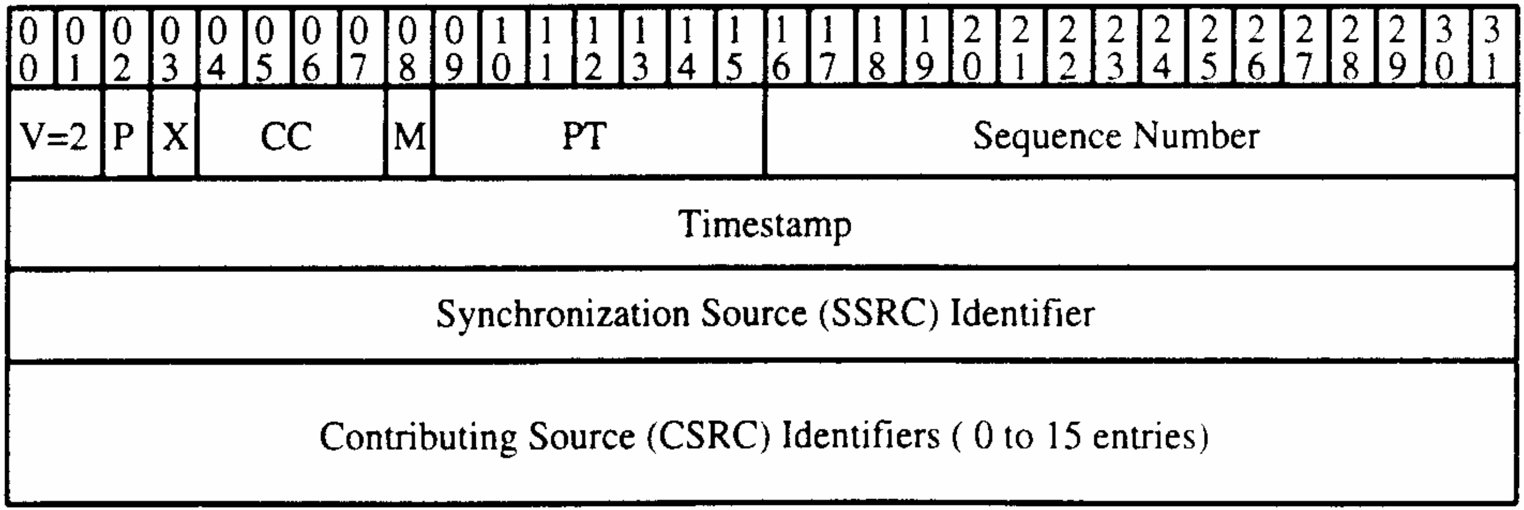

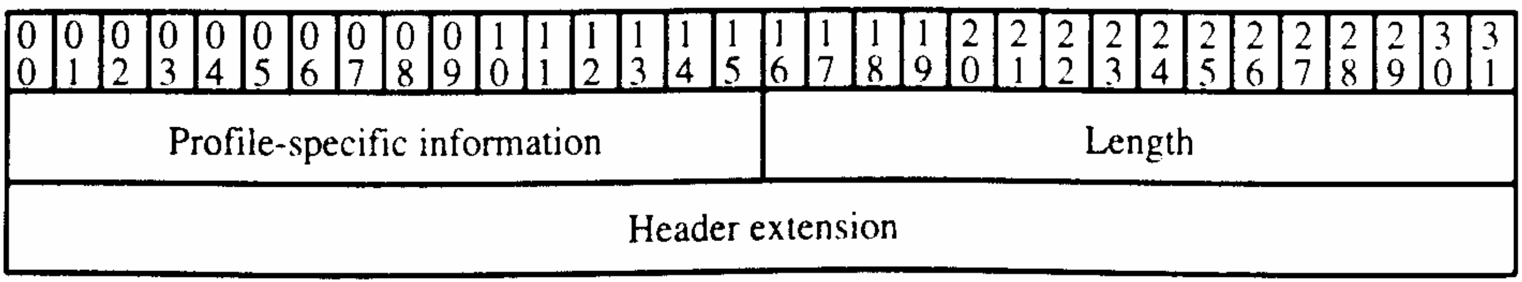

The RTP Header [1/4]

- 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
  - 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
  - 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
  - 0-15 CSRC entries
- RTP Header Extensions

<table>
<tr><td>00</td><td>01</td><td>02</td><td>03</td><td>04</td><td>05</td><td>06</td><td>07</td><td>08</td><td>09</td><td>10</td><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>17</td><td>18</td><td>19</td><td>20</td><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>27</td><td>28</td><td>29</td><td>30</td><td>31</td></tr>
<tr><td colspan="16">Profile-specific information</td><td colspan="16">Length</td></tr>
<tr><td colspan="32">Header extension</td></tr>
</table>

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

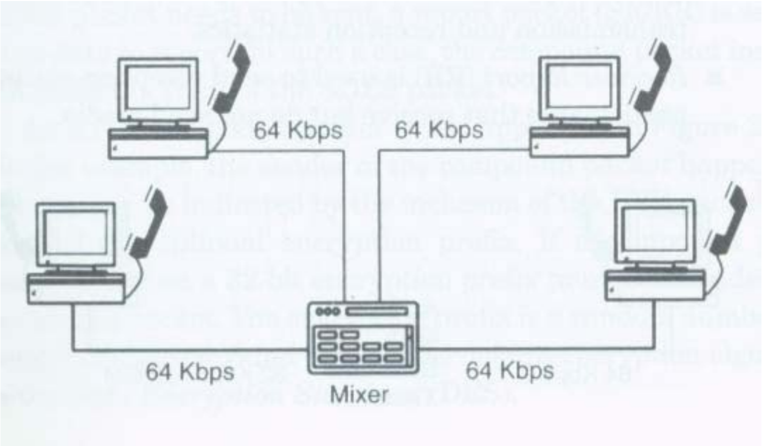

- Translators
  - Manage communications between entities that does not support the same coding scheme
  - The SSRC is the participant, not the translator.

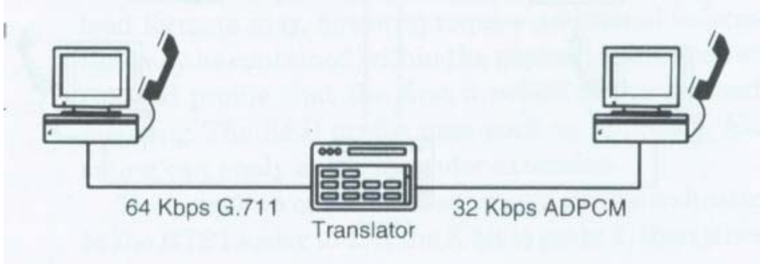

The RTP Control Protocol [1/3]

- RTCP
  - A companion control protocol of RTP
  - Periodic exchange of control information
    - For quality-related feedback
  - A third party can also monitor session quality and detect network problems.
    - Using RTCP and IP multicast
- Five types of RTCP packets
  - Sender Report: transmission and reception statistics
  - Receiver Report: reception statistics

The RTP Control Protocol [2/3]

- Source Description (SDES)
  - One or more descriptions related to a particular session participant
  - Must contain a canonical name (CNAME)
    - Separate from SSRC which might change
    - When both audio and video streams were being transmitted, the two streams would have
      - different SSRCs
      - the same CNAME for synchronized play-out
- BYE
  - The end of a participation in a session
- APP
  - For application-specific functions

The RTP Control Protocol [3/3]

- Two or more RTCP packets will be combined
  - SRs and RRs should be sent as often as possible to allow better statistical resolution.
  - New receivers in a session must receive CNAME very quickly to allow a correlation between media sources and the received media.
  - Every RTCP packet must contain a report packet (SR/RR) and an SDES packet
    - Even if no data to report
- An example RTP compound packet

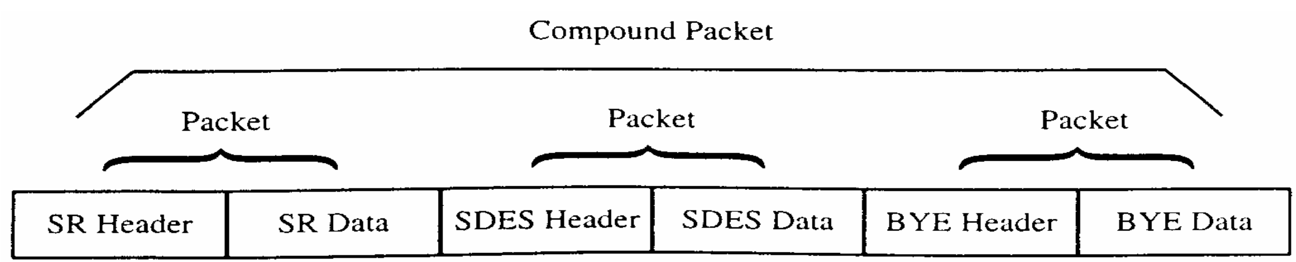

RTCP Sender Report

- SR
  - Header Info
  - Sender Info
  - Receiver Report Blocks
  - Option
    - Profile-specific extension

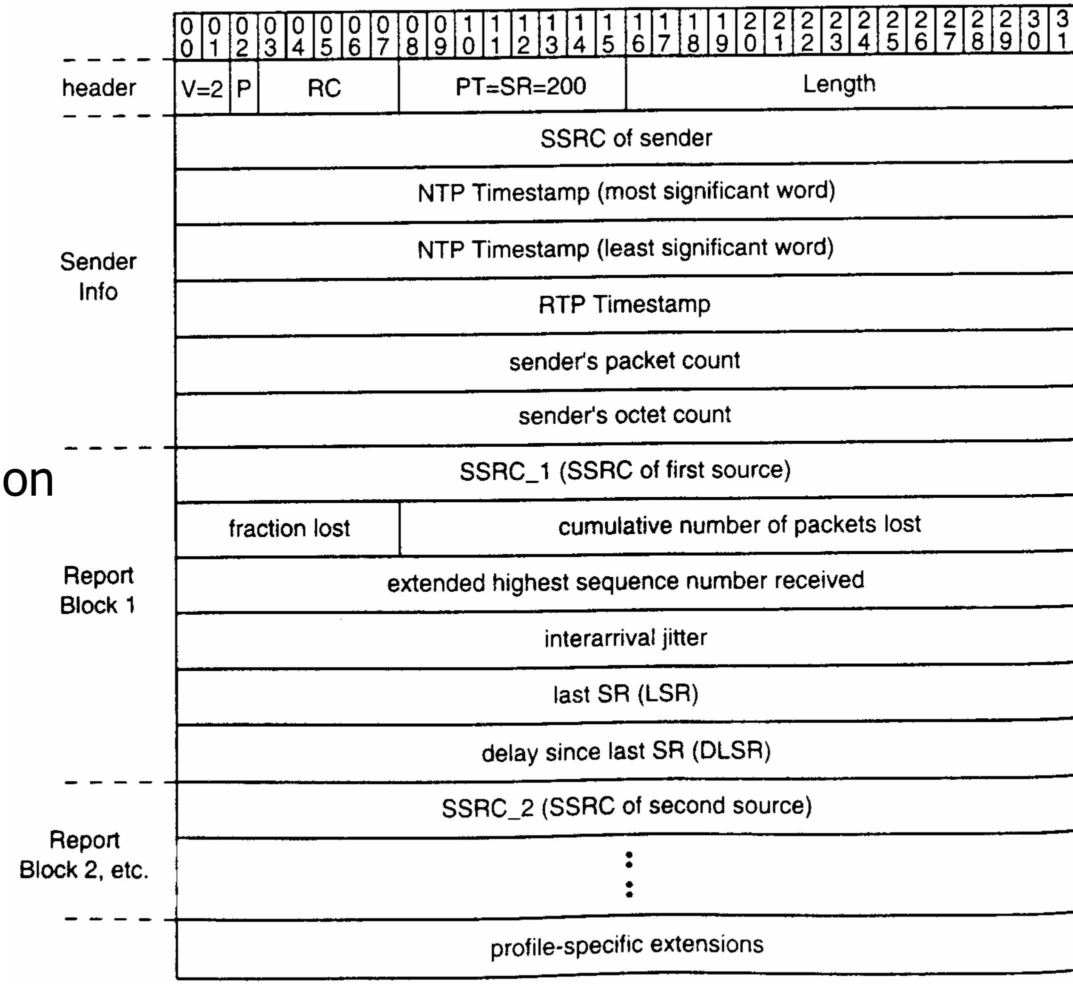

Header Info

- Resemble to an RTP packet
  - Version
    - 2
  - Padding bit
    - Padding octets?
  - RC, report count
    - The number of reception report blocks
    - 5-bit
      - If more than 31 reports, an RR is added
  - PT, payload type

Sender Info

- SSRC of sender
- NTP Timestamp
  - Network Time Protocol Timestamp
    - The time elapsed in seconds since 00:00, 1/1/1900 (GMT)
    - 64-bit
      - 32 MSB: the number of seconds
      - 32 LSB: the fraction of a seconds (200 ps)
- RTP Timestamp
  - Corresponding to the NTP timestamp
  - The same as used for RTP timestamps
  - For better synchronization
- Sender's packet count
  - Cumulative within a session
- Sender's octet count
  - Cumulative within a session

RR blocks [1/2]

- SSRC_n
  - The source identifier of the session participant to which the data in this RR block pertains.
- Fraction lost
  - Fraction of packets lost since the last report issued by this participant
  - By examining the sequence numbers in the RTP header
- Cumulative number of packets lost
  - Since the beginning of the RTP session
- Extended highest sequence number received
  - The sequence number of the last RTP packet received
  - 16 lsb, the last sequence number
  - 16 msb, the number of sequence number cycles

RR blocks [2/2]

- Interarrival jitter
  - An estimate of the variance in RTP packet arrival
- Last SR Timestamp (LSR)
  - The middle 32 bits of the NTP timestamp used in the last SR received from the source in question
  - Used to check if the last SR has been received
- Delay Since Last SR (DLSR)
  - The duration in units of 1/65,536 seconds

RTCP Receiver Report

- RR
  - Issued by a participant who receives RTP packets but does not send, or has not yet sent
  - Is almost identical to an SR
    - PT = 201
    - No sender information

RTCP Source Description Packet

- Provides identification and information regarding session participants
  - Must exist in every RTCP compound packet
- Header
  - V, P, SC, PT=202, Length
- Zero or more chunks of information
  - An SSRC or CSRC value
  - One or more identifiers and pieces of information
    - A unique CNAME
    - Email address, phone number, name

- RTCP BYE Packet
  - Indicate one or more media sources are no longer active
- Application-Defined RTCP Packet
  - For application-specific data
  - For non-standardized application

Calculating Round-Trip Time

- Use SRs and RRs
- E.g.
  - Report A: A, T1 → B, T2
  - Report B: B, T3 → A, T4
  - RTT = T4-T3+T2-T1
  - RTT = T4-(T3-T2)-T1
  - Report B
    - LSR = T1
    - DLSR = T3-T2

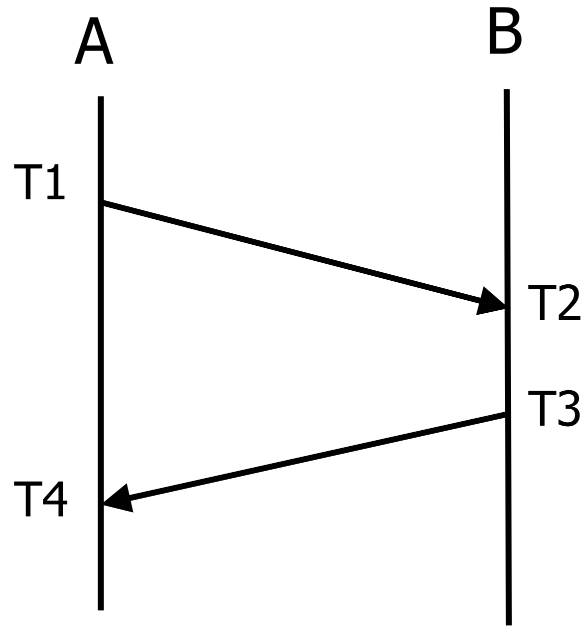

Calculation Jitter

- The mean deviation of the difference in packet spacing at the receiver
  - S_i = the RTP timestamp for packet i
  - R_i = the time of arrival
  - $D(i,j) = (R_j - S_j) - (R_i - S_i)$
- The Jitter is calculated continuously
  - $J(i) = J(i-1) + (|D(i-1,i)| - J(i-1))/16$

Timing of RTCP Packets

- RTCP provides useful feedback
  - Regarding the quality of an RTP session
  - Delay, jitter, packet loss
  - Be sent as often as possible
    - Consume the bandwidth
    - Should be fixed at 5%
- An algorithm, RFC 1889
  - Senders are collectively allowed at least 25% of the control traffic bandwidth.
  - The interval > 5 seconds
  - 0.5 – 1.5 times the calculated interval
  - A dynamic estimate the avg. RTCP packet size

IP Multicast

- An IP diagram sent to multiple hosts
  - Conference
  - To a single address associated with all listeners
- Multicast groups
  - Multicast address
  - Join a multicast group
    - Inform local routers
  - Routing protocols
    - Support propagation of routing information for multicast addresses
    - Minimize the number of datagrams sent