Internet Protocols for Multimedia

By Scott Midkiff

ECpE/CS 5516, VPI

Spring 1997

(modified by Marc Abrams for Spring 1998)

Multimedia Systems (1)

- Multimedia systems incorporate *continuous* media
  - Voice
  - Video
  - Animations
- Distributed multimedia systems
  - Continuous data transfer over long periods of time
  - Media synchronization
  - Possibly massive storage
  - Possibly special indexing and retrieval

Multimedia Systems (2)

- **Interactive versus non-interactive**
  - **Interactive -- active user**
    - **Live**
    - **Stored**
  - **Non-interactive -- passive user**

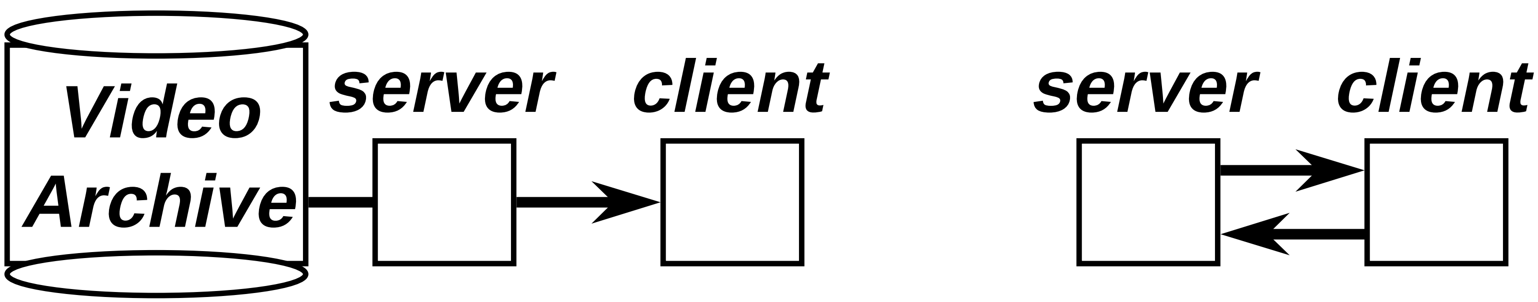

Stored **Live**

Example Multimedia Applications

- Multimedia-on-demand
  - Video-on-demand
  - Audio-on-demand
- Live video
  - News
- Collaboration and video conferencing
- Remote sensing and imaging

System Demands by Multimedia (1)

- **Processing**
  - **Multimedia file systems and formats**
  - **Software codecs (coder/decoders)**
- **Server architecture**
  - **High bus bandwidth**
  - **Efficient input/output**
- **Operating system**
  - **New data types**
  - **Real-time scheduling**
  - **Fast interrupts**

System Demands by Multimedia (2)

- **Storage**
  - **Large capacity**
  - **Fast access**
  - **High transfer times**
- **Network**
  - **High data rate**
  - **Low latency (delay)**
  - **Low jitter (variation in delay)**

Opportunity and Need for Video Compression

- Video characteristics
  - Demanding with respect to storage and/or data rate

$$\left(640 \times 480 \text{ pixels / f}\right)\left(24 \text{ b / pixel}\right)\left(30 \text{ f / s}\right) = 221\,\text{Mbps}$$

  - Highly redundant -- duplicated information
    - Compression ratios of 200:1 or even 2000:1 are possible
- Compression is needed to enable
  - Storage
  - Transmission

Video Compression Techniques (1)

- **Information may be lost (but not missed)**
  - **Lossy compression -- information is lost**
  - **Lossless compression -- no loss**
- **Lossy techniques may rely on**
  - **Prediction -- predict future information based on previous values**
  - **Importance -- consider the importance of characteristics and information to viewer**
  - **Frequency -- process spatial frequency information**

    ***Example*: Discrete Cosine Transform (DCT). Encodes 8x8 pixel blocks as coefficients in frequency domain. Uses FFT-like algorithm. Coefficients can be reordered to facilitate compression. Used in JPEG and MPEG.**

Video Compression Techniques (2)

- Scope of compression
  - Intraframe -- eliminate or reduce redundancy within a single frame
  - Interframe -- eliminate or reduce redundancy between consecutive frames
  - Sample to take advantage of human perception

- Most practical compression schemes are hybrids -- they utilize multiple approaches

Example Video Compression Standards

- **MPEG-1 and MPEG-2**
- **ITU-T (CCITT) standards**
  - **H.320 (H.261) — ISDN (64 kbps increments)**
  - **H.323 — LAN**
  - **H.324 — POTS**
- **MJPEG (Motion JPEG)**

MPEG:	Motion Pictures Expert Group
ITU-T:	International Telecommunication Union -- Telecommunication
JPEG:	Joint Photographic Experts Group

MPEG Overview (1)

- Features
  - Can achieve compression ratios of 200:1 (would reduce data rate to around 1.2 Mbps for a 640x480 image)
  - MPEG-1 compresses 320x240 images and requires at least 1.5 Mbps
  - MPEG-2 compresses 720x480 images and requires 4 to 10 Mbps
  - Also includes audio compression with compression ratios of 5:1 to 10:1

MPEG Overview (2)

- **Compression techniques**
  - **Int*ra*-frame compression**
  - **Int*er*-frame compression by storing differences between successive frames**
- **There are three frame types**
  - ***Intraframes*** **(I frames) are encoded using JPEG's DCT w/ Huffman coding for intraframe compression**
  - ***Predicted frames*** **(P frames) are predicted from previous I frames using a vector to represent displacement of a 16x16 pixel macroblock**
  - ***Bidirectional frames*** **(B frames) are interpolated from previous and future frames (e.g., for a scene that is moving)**

MPEG Overview (3)

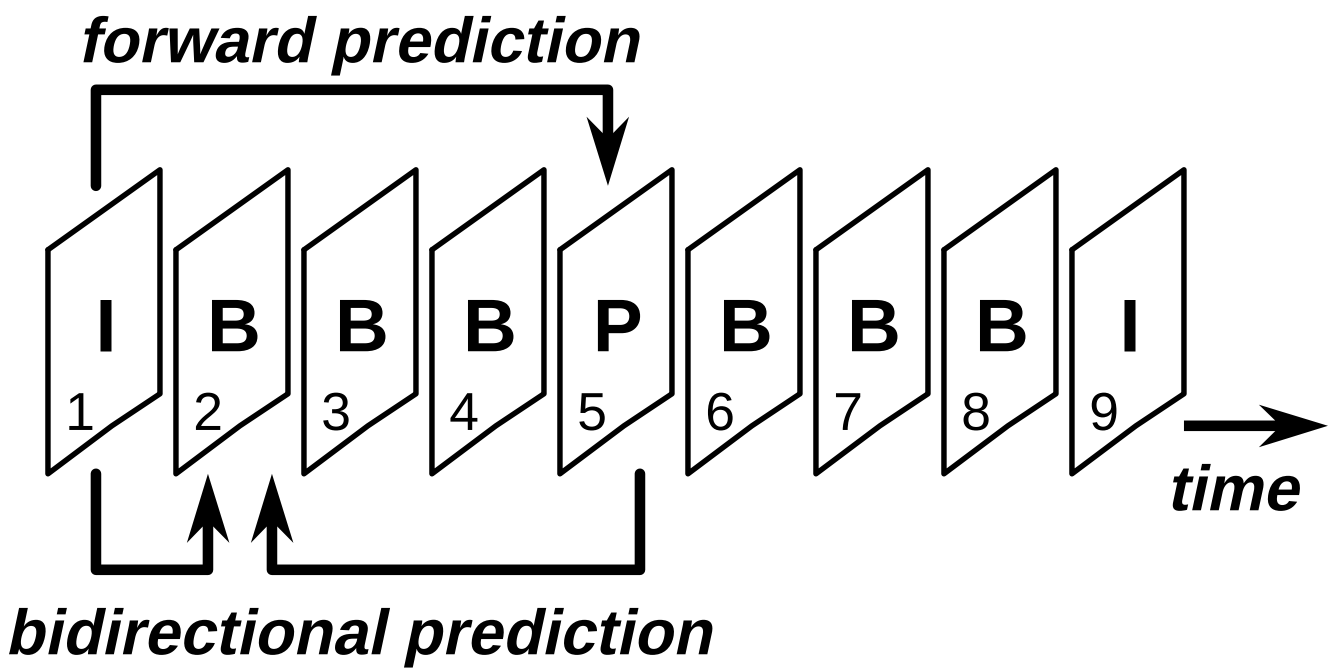

Quality of Service -- QoS (1)

- **Quality of service (QoS) is critical for multimedia traffic.**
- **QOS may specify acceptable values of ...**
  - **Data rate**
  - **Average latency (delay)**
  - **Variance in delay or jitter**
    - **Jitter is the instantaneous difference in two synchronized events or streams**
  - **Error rate**
    - **Bit error rate (BER)**
    - **Packet error rate (PER)**

Quality of Service -- QoS (2)

- ***Average delay***
  - **Delay must be minimized for interactive applications, especially live**
  - **Delay alone does not affect non-interactive applications**

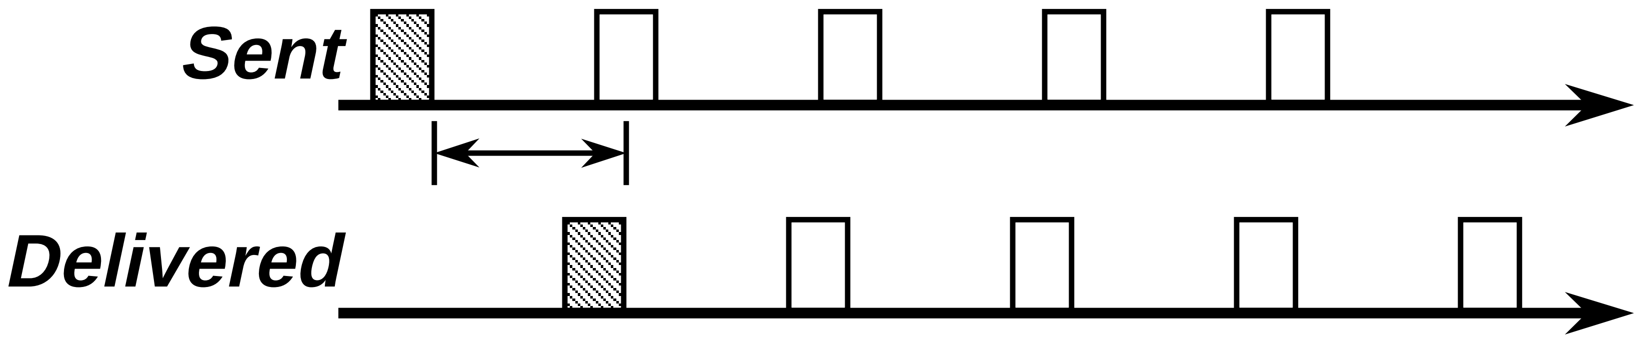

Quality of Service -- QoS (3)

- ***Variance in delay or jitter***
  - **Violates real-time constraints**
  - **Data is not available when it is needed, even though average delay may be satisfactory**

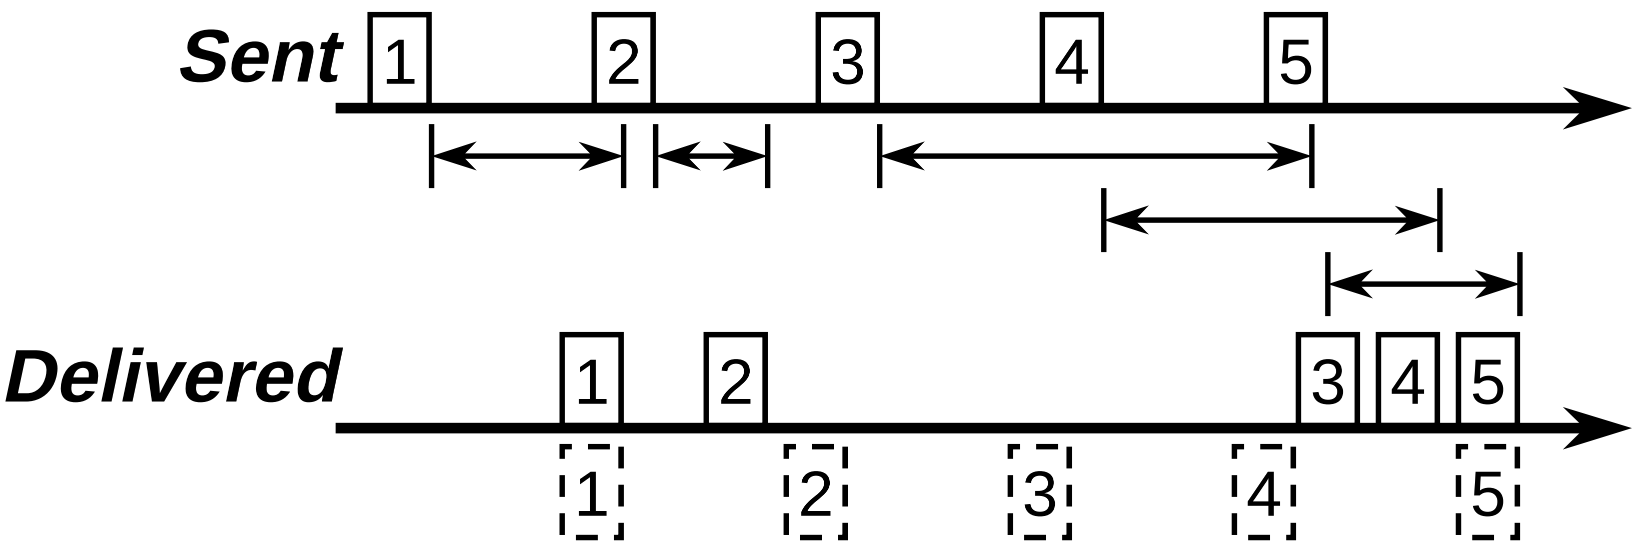

Quality of Service -- QoS (4)

- ***Error rate***
  - **Types of errors**
    - **Bit error -- isolated bit errors**
    - **Packet error -- entire packet (video frame) is lost**
  - **Compression makes video *less robust* in the face of errors**
  - **It may be possible to ignore isolated bit errors, for example a single pixel may be momentarily incorrect**
  - **Packet loss may be significant, especially if it is a critical frame (e.g. I or P frame in MPEG)**

Internet Support for Multimedia

- IP provides a “best effort” service that is optimized for point-to-point delivery
  - IP cannot provide QoS guarantees
  - Point-to-point broadcasts may be wasteful for one-to-many delivery and multipoint conferences
- IPv6 offers some improvements
  - Provides a foundation for a solution through flow identifiers, but does not solve QoS problems
  - Better multicast support

Coping with Insufficient QoS Guarantees

- Reduce bandwidth demands
  - Compression
  - Multicast
- Mask errors
- Error recovery
- Streaming
- New protocols
  - ST-II
  - RSVP
  - RTP and RTCP
- More fundamental changes may be needed!

Internet Multicast Backbone -- MBONE

- MBone allows real-time audio/video/data to be sent over the Internet in a bandwidth efficient manner
- The MBone is a virtual network on top of the Internet
  - Routers that support IP multicast
  - IP tunnels between such routers and/or subnets
- Utilizes IP multicast and the Internet Group Management Protocol (IGMP)

Multicast Routing

- Class D IP addresses are used to identify multicast groups
  - "Well-known" addresses
  - Dynamic groups
- Multicast router maintains table of multicast groups that are active on its networks

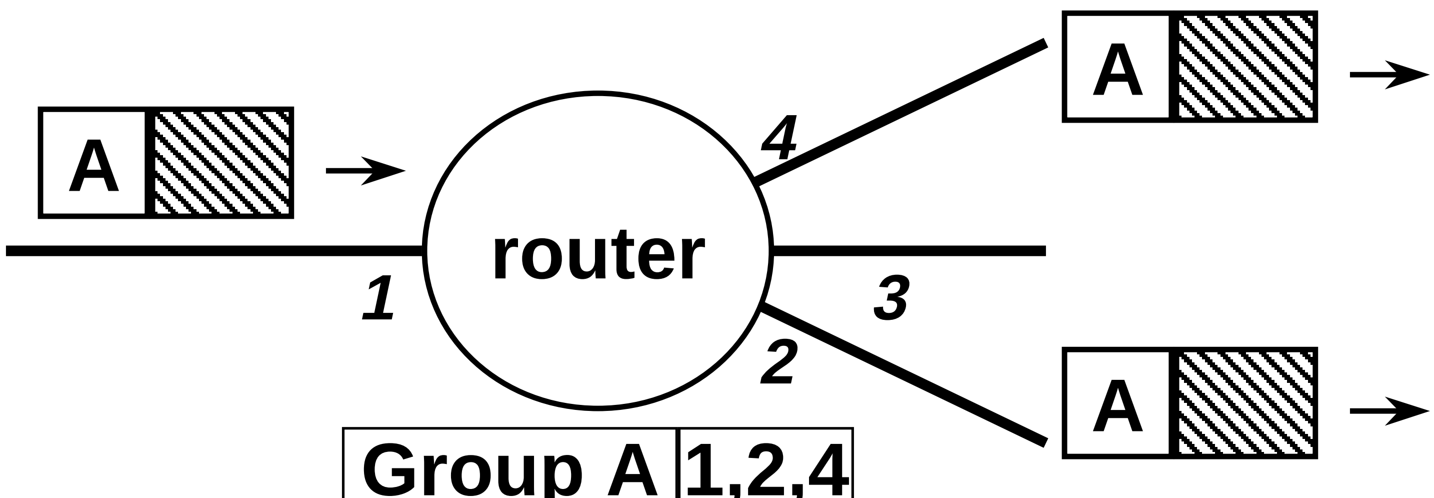

Internet Group Management Protocol -- IGMP (1)

- **IGMP provides information to routers so that it can build its multicast routing table**
  - **Hosts send *reports* of all groups with at least one joined process**
  - **Routers send *queries* for reports**
- **IGMP message is carried by IP**

IP datagram

IP header	IGMP message
20 bytes	*8 bytes*

Internet Group Management Protocol -- IGMP (2)

0 3	*4 7*	*8 15*	*16 31*
vers	type	not used	checksum
group address (same as IP destination)			

- **IGMP message format**
  - **4-bit IGMP version (=1)**
  - **4-bit IGMP type**
    - **1: Query sent by a router**
    - **2: Report sent by a host**
  - **32-bit group address (Class D IP address)**
  - **16-bit checksum**
- **IP header specifies address of sending host**

Internet Group Management Protocol -- IGMP (3)

- **Joining a group**
  - **Host sends group report when the first process joins a given group**
- **Maintaining table at the router**
  - **Multicast router periodically queries for group information; sent to all-hosts address (1's for host id)**
  - **Any host wanting to remain in group must reply with an IGMP report**
  - **So host does not notify router when the last process leaves a group -- this is discovered through the lack of a report for a query**

Internet Group Management Protocol -- IGMP (4)

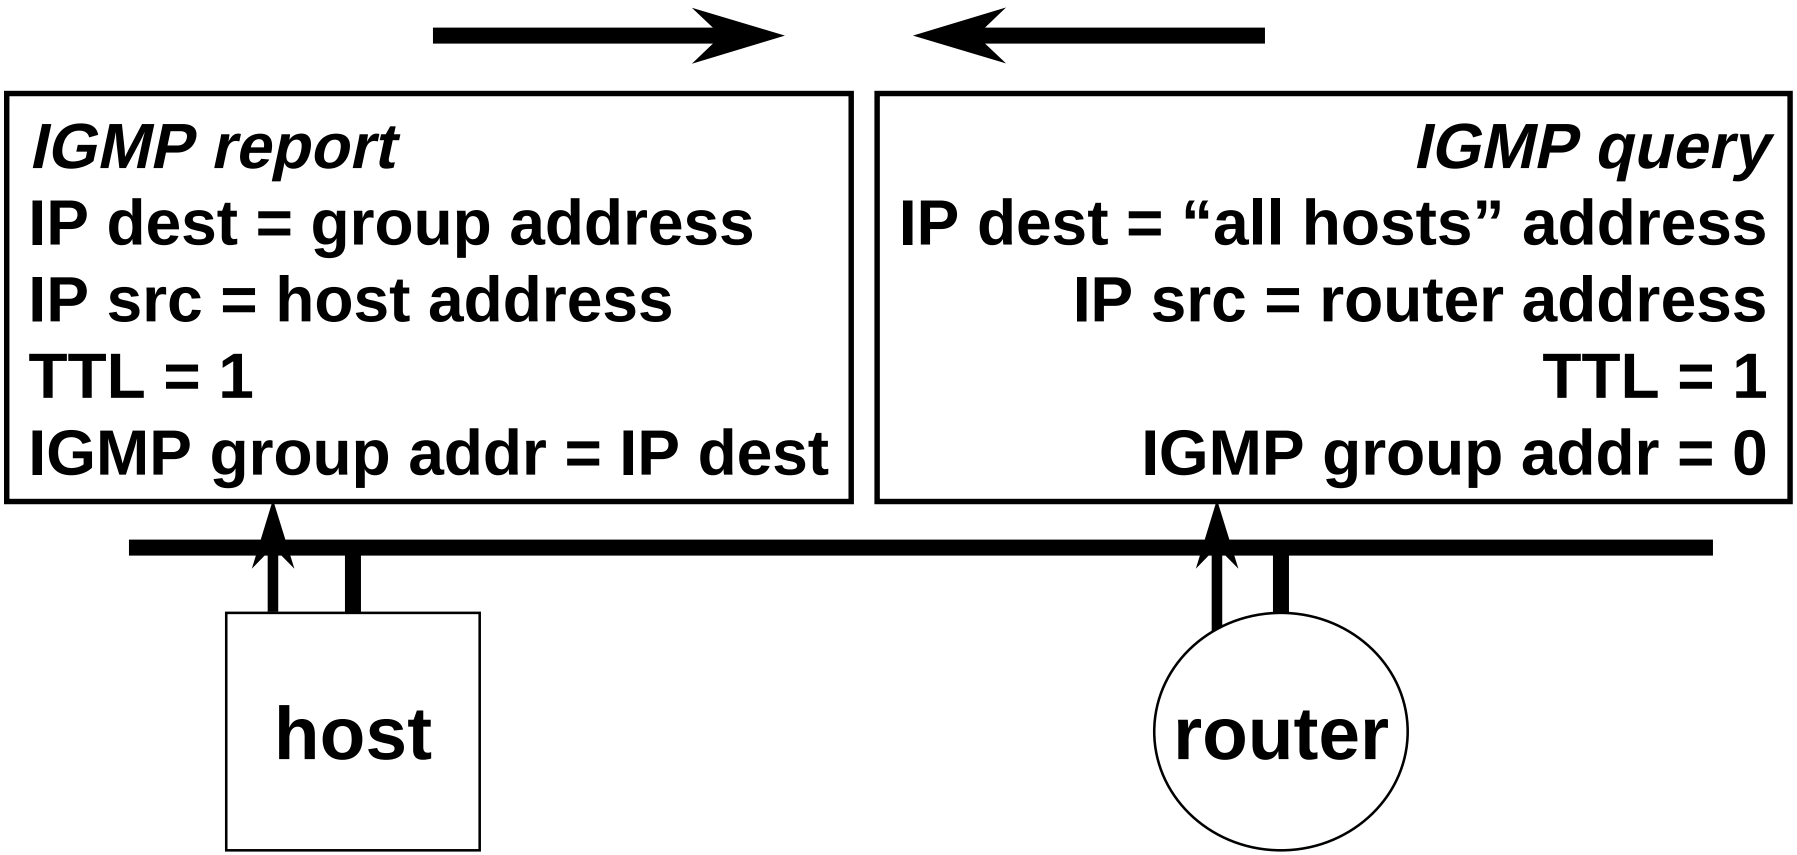

Streaming Applications (1)

- Streaming applications attempt to overcome variation in latency by "smoothing" input using a buffer
- Recall the problem with delay variance

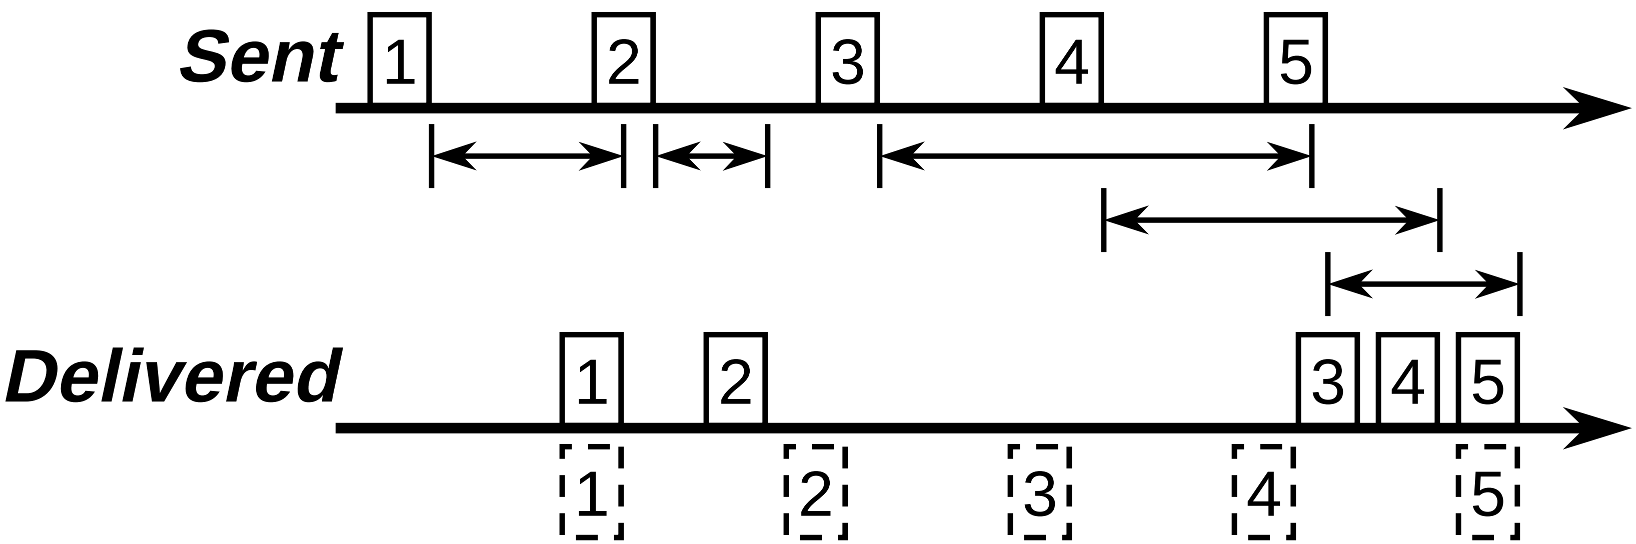

Streaming Applications (2)

- "Elastic" buffer
  - Data is pre-loaded into the buffer
  - Buffer length reduces as delay increases
  - Buffer length increases as delay decreases
  - Protocol may be needed to pace the source
    - "Low water mark" increase source rate
    - "High water mark" decreases source rate

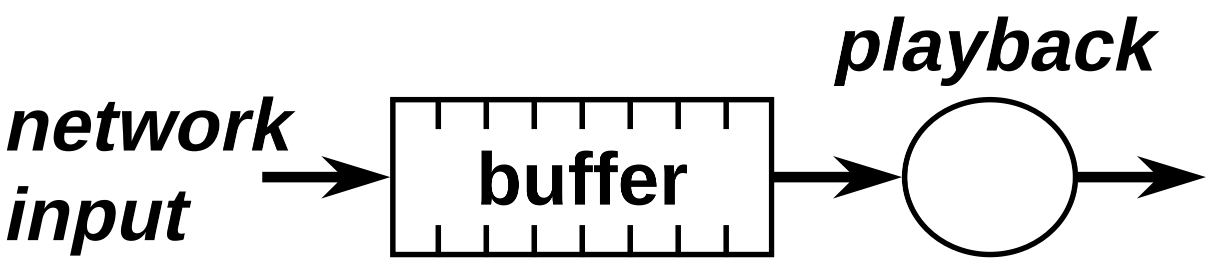

Streaming Applications (3)

- Effect of elastic buffer

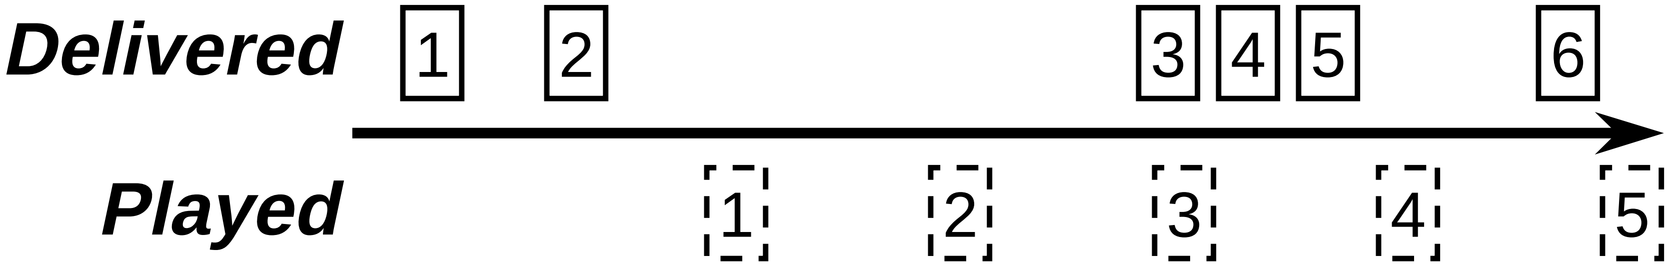

- Application must be able to tolerate delay
  - Works for video- or audio-on-demand or a delivery only broadcast
  - Ineffective for conferencing

Resource Reservation Protocol -- RSVP (1)

- **Network features for guaranteed QoS**
  - ***Routing*** **-- must permit multicast**
  - ***Flow specification***
    - **identifies traffic properties & QoS requirements**
    - **Ex: average, peak BW, level of burstiness**
  - ***Resource class*** **-- e.g., guaranteed delay**
  - ***Resource reservation*** **-- to allocate resources based on QoS needs**
  - ***Admission control*** **-- determines when to accept new reservation requests**
  - ***Packet scheduling*** **-- determines when a packet should be transmitted**
  - ***Payment*** **-- Your ISP charges you extra money for each reservation**

Resource Reservation Protocol -- RSVP (2)

- ***Receivers*** **initiate a reservation**
  - **Reservations can be customized to the needs of each user, not "one size fits all"**
- **Reservation styles may differ**
  - **Multiple video sources from one source**
  - **Single video source from a set of sources**
  - **Single video source sent to each use, but may differ per user**
- **Routers use "soft state," fitting connectionless**
  - **Reservation state is cached in routers**
  - **Periodically refreshed by end stations**

Resource Reservation Protocol -- RSVP (3)

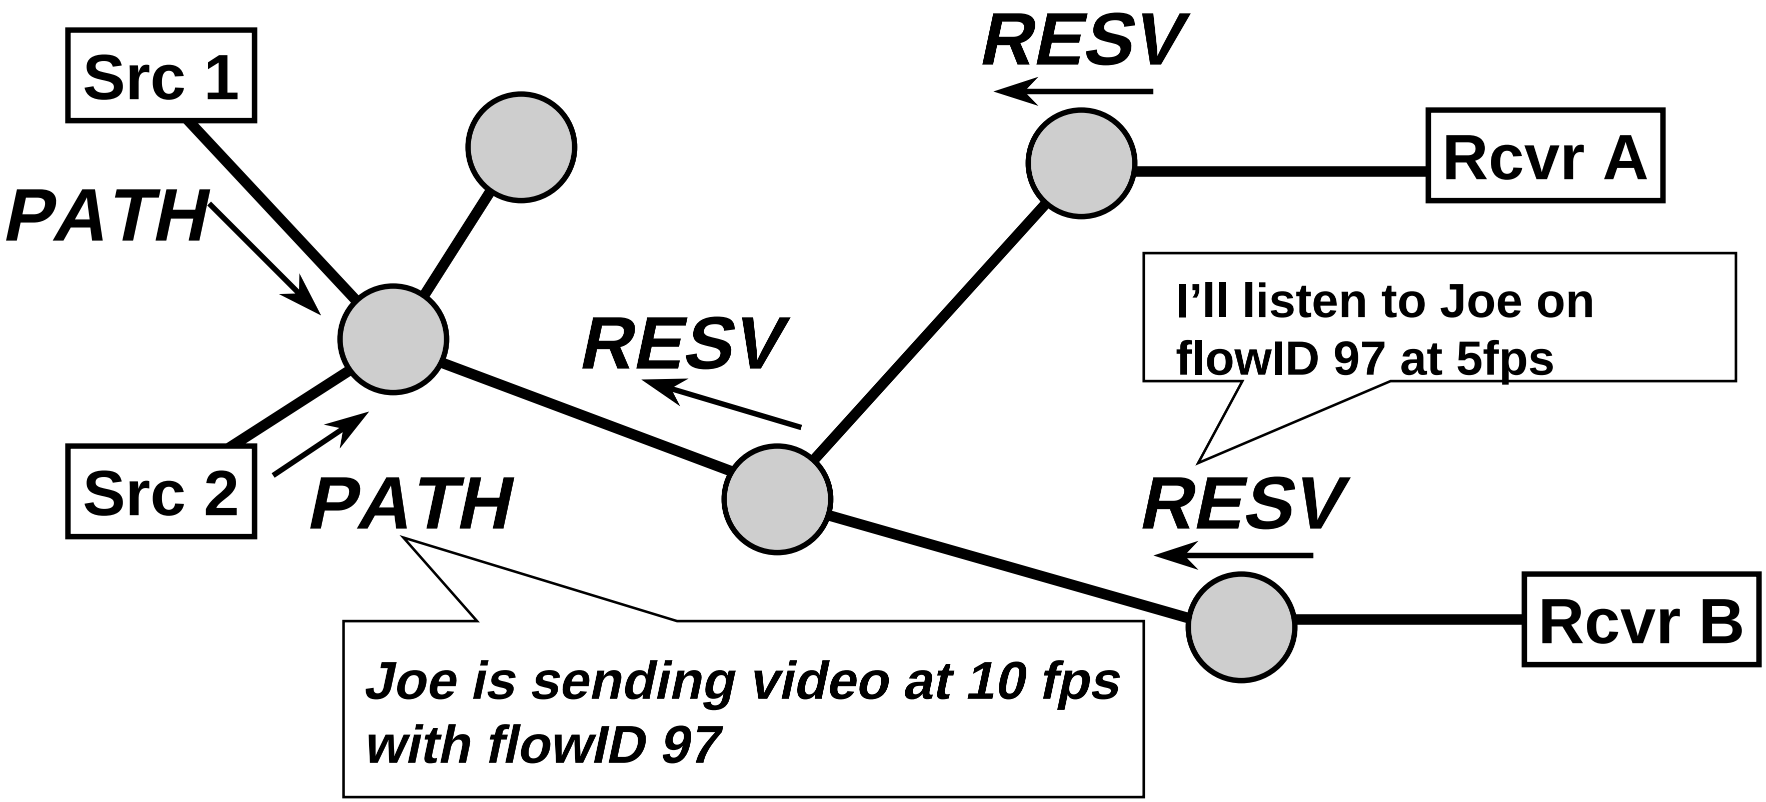

L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: A New Resource ReSerVation Protocol," *IEEE Network*, vol. 7, no. 5, pp. 8-18, Sept. 1993.

Stream Protocol: Version 2 -- ST-II

- ***Alternative to IP*** **-- ST-II is "IPv5"**
- ***Defines streams*** **through a network, providing end-to-end guaranteed service**
- ***Stateful*****, unlike IP**
- **Each ST-II entity in a path from source to destination pre-allocates resources to a stream and maintains state information for the stream**
- **Builds tree from each source to destinations**

CIP Working Group (C. Topolcic, Editor), "Experimental Internet Stream Protocol, Version 2 (ST-II)," RFC 1190, October 1990.

Real-Time Protocol -- RTP (1)

- RTP provides end-to-end network ***transport level*** functions for real-time applications
  - Payload type identification (e.g, encoding)
  - Sequence numbering
  - Time-stamping
  - Delivery monitoring
  - Allows multicast if network allows it
- Usually uses IP or UDP, but design to be independent of lower layers
- Unlike RSVP, does not reserve resources nor guarantee quality-of-service; *does* real-time delivery

Real-Time Protocol -- RTP (2)

- **Real-Time Control Protocol (RTCP) lets applications monitor data delivery**
  - **Scaleable to large multicast networks**
  - **Provides feedback on reception quality, optional identification of multicast receivers**
- **RTP delivers the data, RTCP monitors the quality of the delivery**

H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications," RFC 1889, January 1996.

RTP Audio Multicast Conference Example (1)

- **Addressing**
  - **Multicast group address associated with the conference**
  - **Two ports are used on each host**
    - **One port is used for audio data (carried in RTP over UDP)**
    - **The other port is used for control (RTCP) packets**
  - **Address and port information is distributed to the intended participants**
  - **Allocation and distribution of addresses is *not* a function of RTP**

RTP Audio Multicast Conference Example (2)

- **An audio conferencing application is used by each participant**
  - **Sends audio data in small chunks, 20 ms in duration**
  - **Each chunk of audio data is carried in an RTP packet**
  - **RTP header and data are carried in a UDP packet**
  - **RTP header indicates type of audio encoding (e.g. PCM or ADPCM)**

RTP Audio Multicast Conference Example (3)

- RTP header contains timing information so that the receiver can reconstruct the audio sequence
  - Consecutive audio “chunks” are played every 20 ms
- RTP header contains a sequence number
  - The receiver can estimate lost packets

RTP Audio Multicast Conference Example (4)

- Each instance of the audio application in the conference periodically multicasts a reception report plus the name of its user on the RTCP port
  - Indicates how well the current speaker is being received
  - Can used to control adaptive encodings
  - With multiple sources (e.g., a conference), each could send packets less frequently as number of participants increases
- Site sends the RTCP BYE packet when it leaves the conference

Summary

- Multimedia is a very demanding application area -- it has real-time demands
- Fast growing application area
- Active research area
- Support of QoS in underlying network may be essential
  - Asynchronous Transfer Mode (ATM)
  - 100VG-AnyLAN (voice grade by priority access)
  - FDDI
  - Circuit-switched networks