Review: Network link technologies

1. What happens when client send packets back-to-back?
   - What happens when the network is saturated?
2. What about bandwidth/latency/jitter?
3. Nature of packet loss?

- Ethernet: CSMA/CD - collisions
- Wireless LAN: CSMA/CA - wireless channel
- Wireless Cellular - handovers
- Point-to-point - dialup -
- Token-Ring - 4/16 Mbps - fair access
- ATM - cell based - 155 Mbps - pre-determined

TCP components

- Sliding window protocol to achieve reliable transmission
  - Ideally, bandwidth x delay product worth window
    - RTT estimation algorithms
  - Advertised window for flow control - receiver restriction
  - Congestion window for congestion - network restriction
    - Basically, advertised window and congestion window reduces the amount of data that we can send
- TCP uses end-point probing to learn about network
  - Causing congestion creates implicit feedback - cannot expect explicit feedback in the Internet
  - Nice = slow to detect network limits, fast = congestion
  - When congestion happens, be nice and backoff

TCP congestion response

- AIMD - increment slowly to probe network, backoff multiplicatively when congestion happens
- Slow Start - increment multiplicatively -> causing congestion sooner
- Using hybrid, slow start to quickly catchup (like in the beginning) and then use AIMD when fine tuning

Other TCP mechanisms

- Sequence number space is important to avoid packets from previous (independent) TCP connections
  - Initial sequence no. chosen to be less likely to overlap
  - Sequence number space (+ newer timestamp)
  - SYN+SYN/ACK
  - FIN can be initiated by either side
- SACK and D-SACK are newer mechanisms to allow the receiver to let the sender know what packets have been received (to avoid unnecessary retransmissions in the case of loss)

TCP and Networks

1. Sequence number management
   - Initial sequence number
   - total space
2. Handshakes
3. AIMD or Slow-Start
4. SACK/D-SACK

- Ethernet
- Wireless LAN
- Cellular
- Token Ring
- ATM
- dialup

TCP scenarios

- High speed web server clients:
  - Dialup user - bandwidth is low
  - Wireless lan user
    - Network errors are bursty
    - Networks can be lossy
  - ADSL user
    - TCP paces itself with ACKs
    - Downlink != uplink
  - Cellular user
    - Handoff delays can throw off TCP
  - Satellite user
    - Asymmetric
    - High bandwidth x delay networks

Quality of Service

- Some applications require some sort of guarantee of what kind of service that they can get

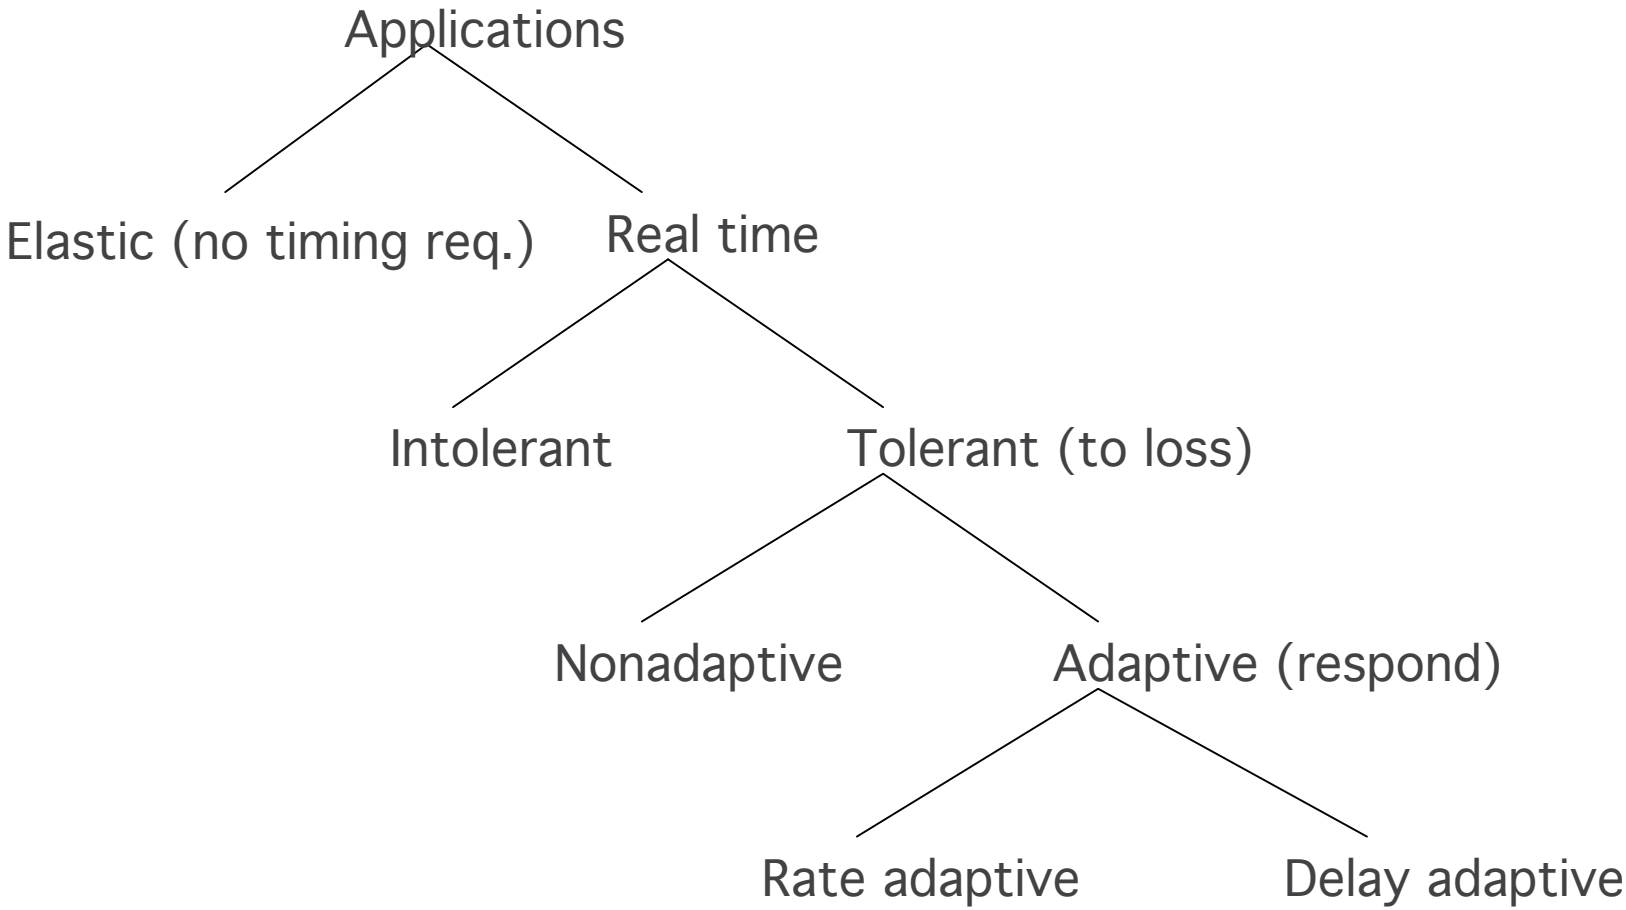

Approaches to QoS

- Fine-grained approaches: QoS on per flow basis
  - E.g. RSVP
- Coarse-grained approaches: provides QoS for large class of aggregated traffic
  - Differentiated services

Quality of Service

- Outline
  - Realtime Applications
    - Networking with specified delay components
  - Integrated Services
    - Per flow QoS
  - Differentiated Services
    - QoS for aggregated traffic

Streaming Audio

The media player buffers input from the media server and plays from the buffer rather than directly from the network.

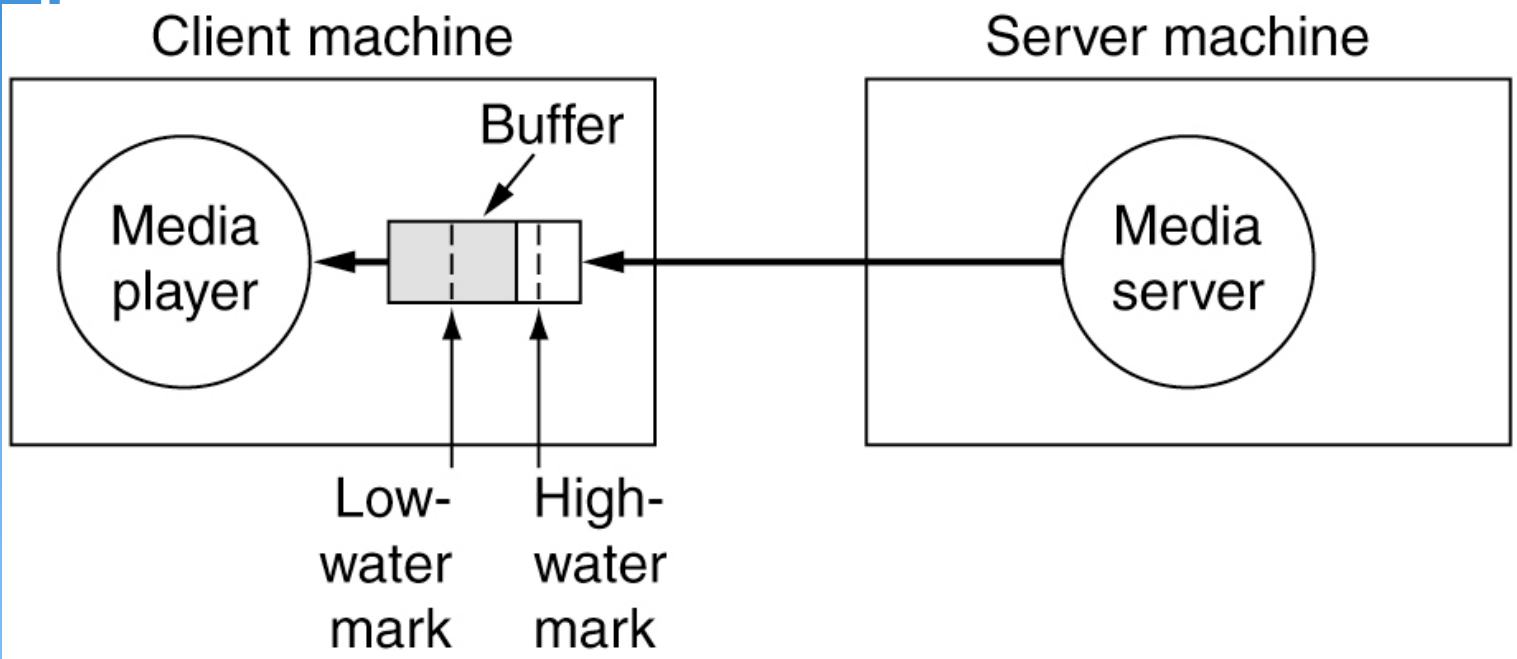

Realtime Applications

- Require “deliver on time” assurances
  - must come from inside the network

- Example application (audio)
  - sample voice once every 125μs
  - each sample has a playback time
  - packets experience variable delay in network
  - add constant factor to playback time: playback point
    - Similar to skip protection in portable CD players

Playback Buffer

- Playback point as insurance against Internet delays
- Multimedia care about delay and jitter (variability within delay)

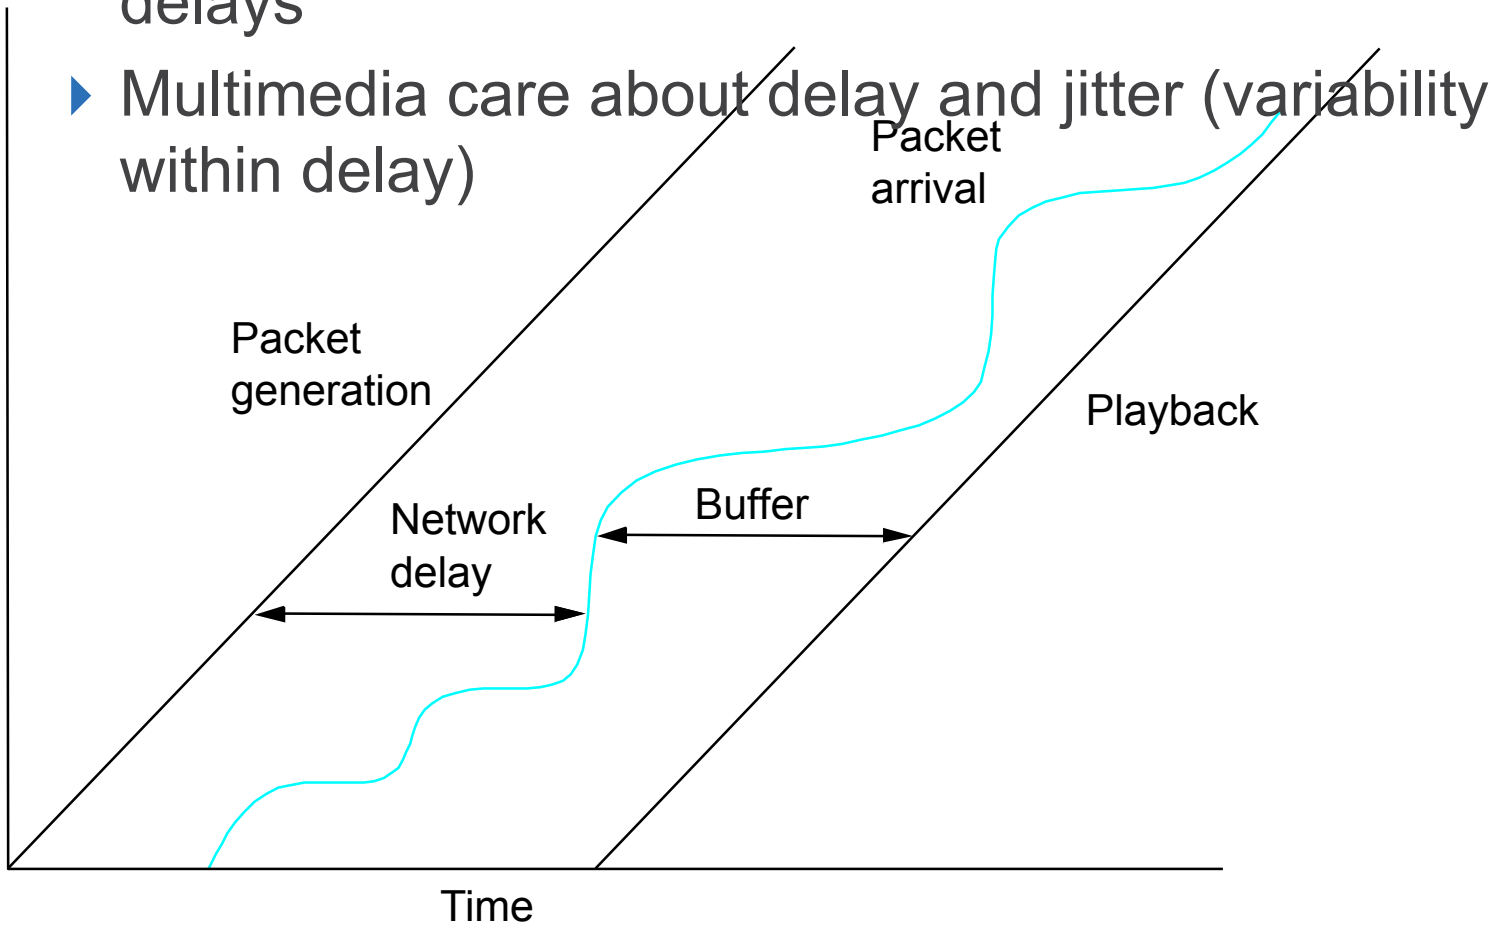

Example Distribution of Delays

- What is a good delay? 200 msec
- Not acceptable for chat application

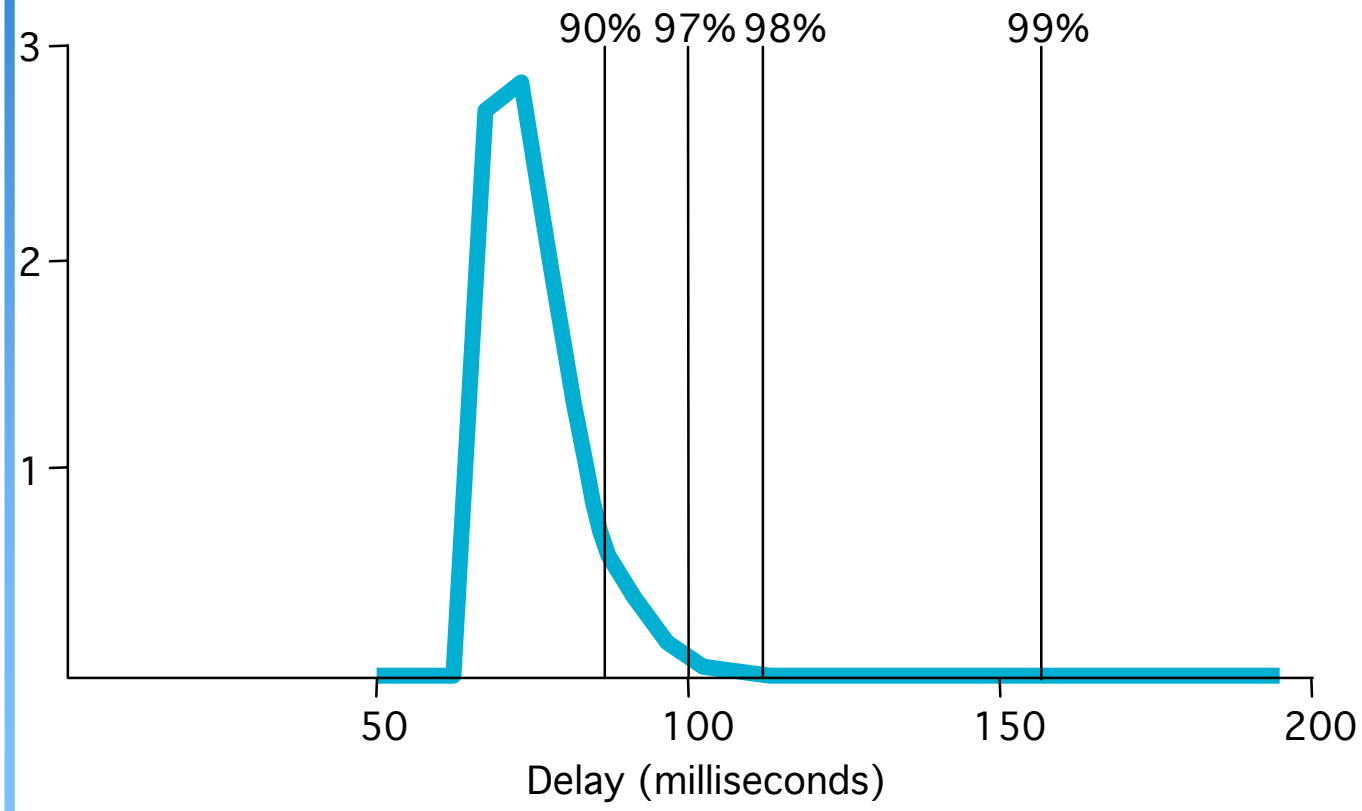

Video transmission

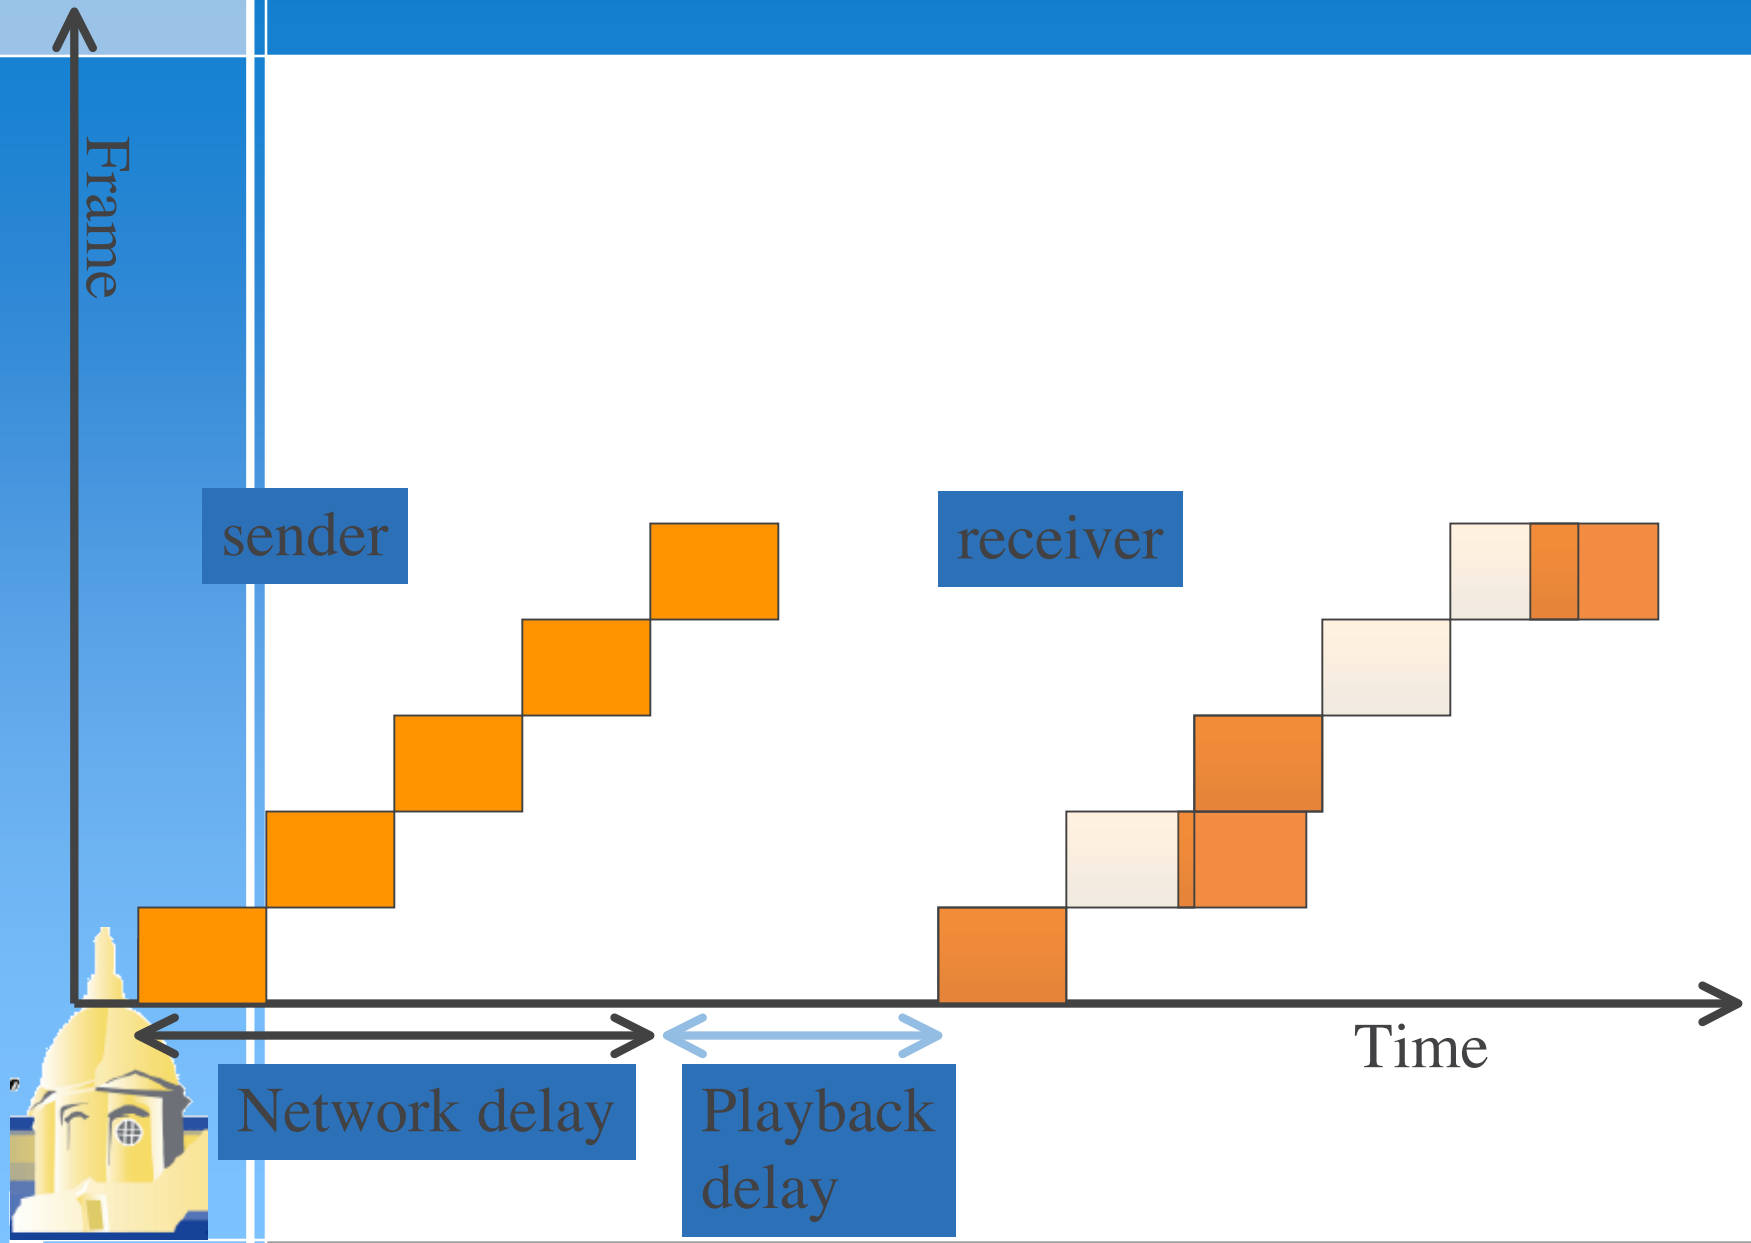

Taxonomy of real time applications

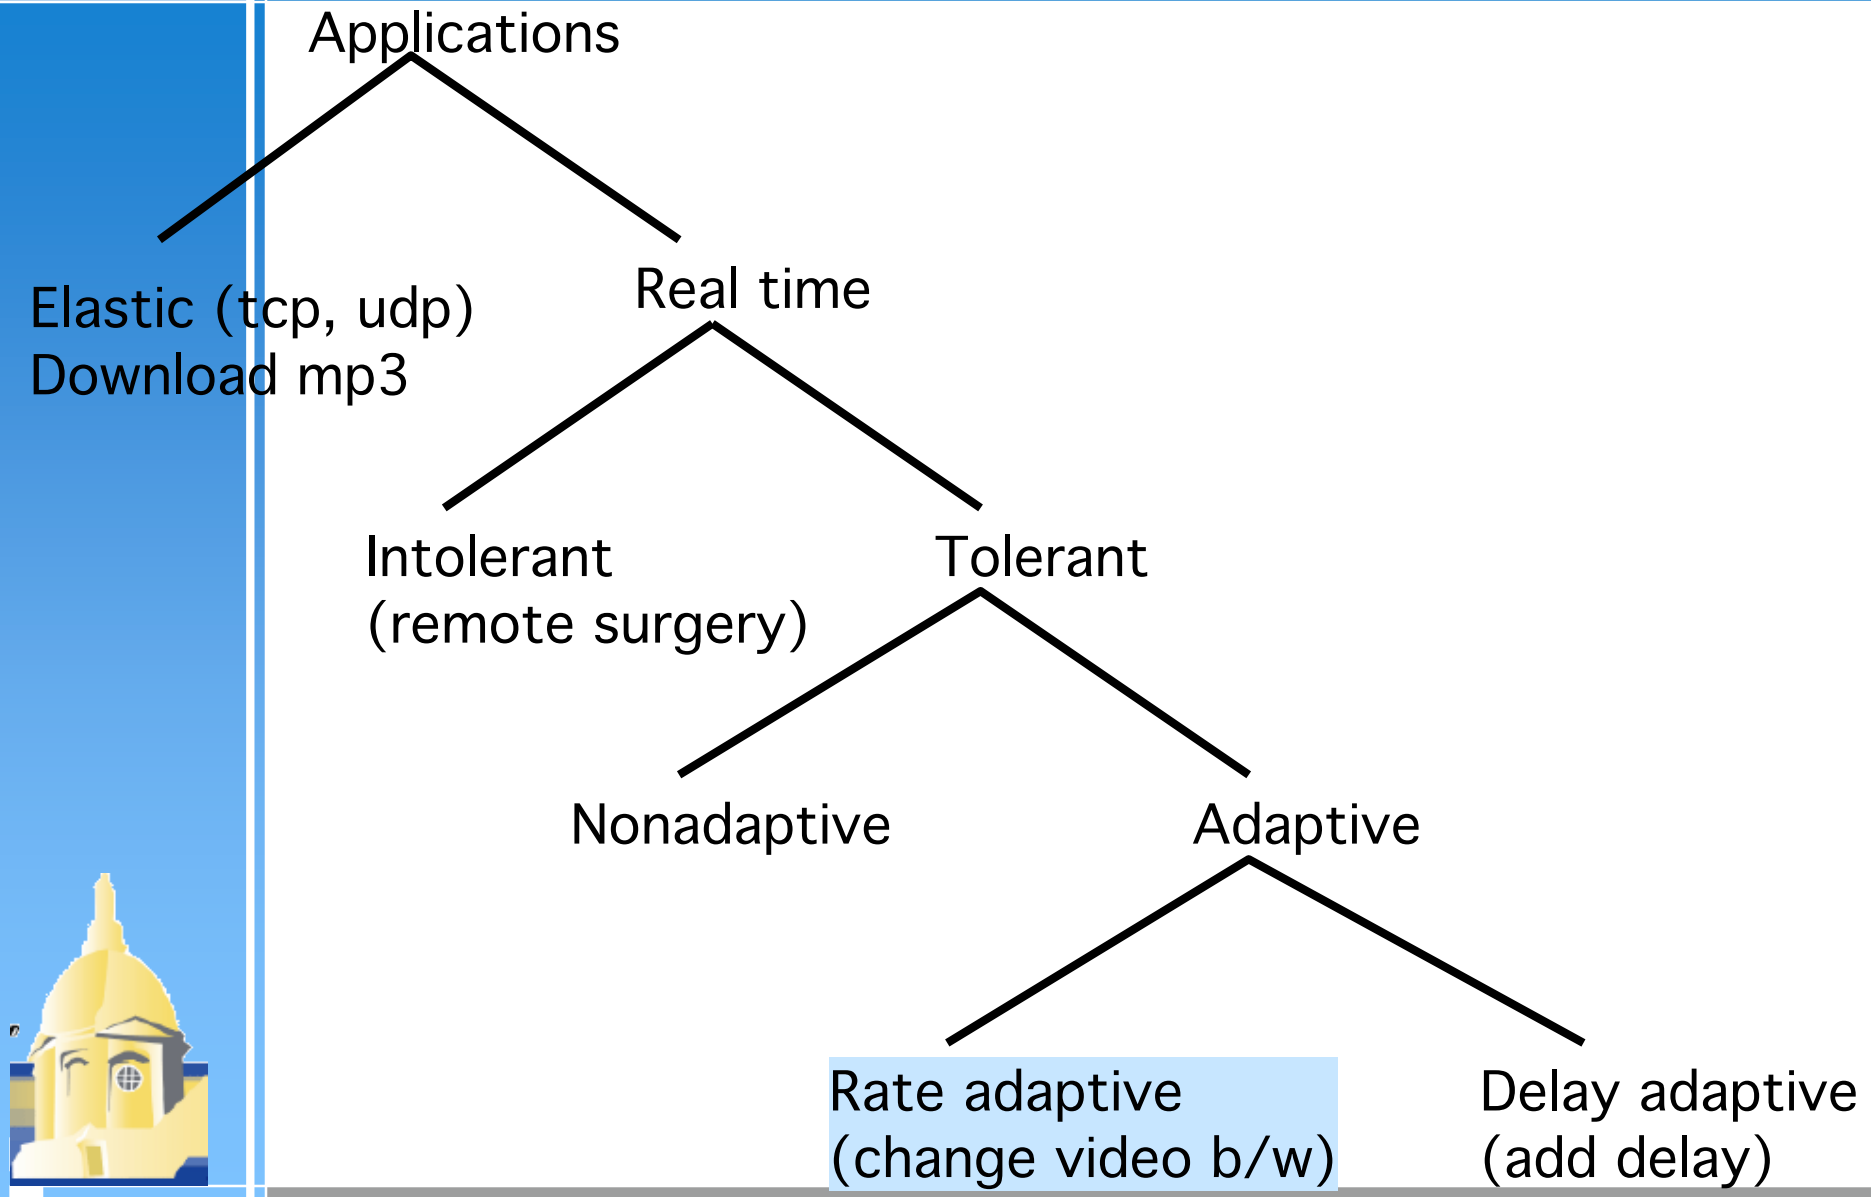

QoS Approaches

- Fine grained - individual application or flows
  - Intserv
  - E.g. for my video chat application
- Coarse grained - aggregated traffic
  - Diffserv
  - E.g. All traffic from CSE (costs $$)

Integrated Services

- IETF - 1995-97 time frame
- Service Classes
  - guaranteed
  - controlled-load (tolerant, adaptive applications)
    - Simulates lightly loaded link
- Mechanisms
  - signaling protocol: signals required service
  - admission control: rejects traffic that cannot be serviced
  - Policing: make sure that senders stick to agreement
  - packet scheduling: manage how packets are queued

Flowspec

- Rspec: describes service requested from network
  - controlled-load: none
  - guaranteed: delay target
- Tspec: describes flow's traffic characteristics
  - average bandwidth + burstiness: token bucket filter
    - token rate r and bucket depth B
  - must have a token to send a byte
  - must have n tokens to send n bytes
  - start with no tokens
  - accumulate tokens at rate of r per second
  - can accumulate no more than B tokens

Per-Router Mechanisms

- Admission Control
  - decide if a new flow can be supported
  - answer depends on service class
  - not the same as policing
- Packet Processing
  - classification: associate each packet with the appropriate reservation
  - scheduling: manage queues so each packet receives the requested service

Reservation Protocol

- Called signaling in ATM
- Proposed Internet standard: RSVP
- Consistent with robustness of today's connectionless model
- Uses soft state (refresh periodically)
- Designed to support multicast
- Receiver-oriented
- Two messages: PATH and RESV
- Source transmits PATH messages every 30 seconds
- Destination responds with RESV message
- Merge requirements in case of multicast
- Can specify number of speakers

RSVP Example (multicast)

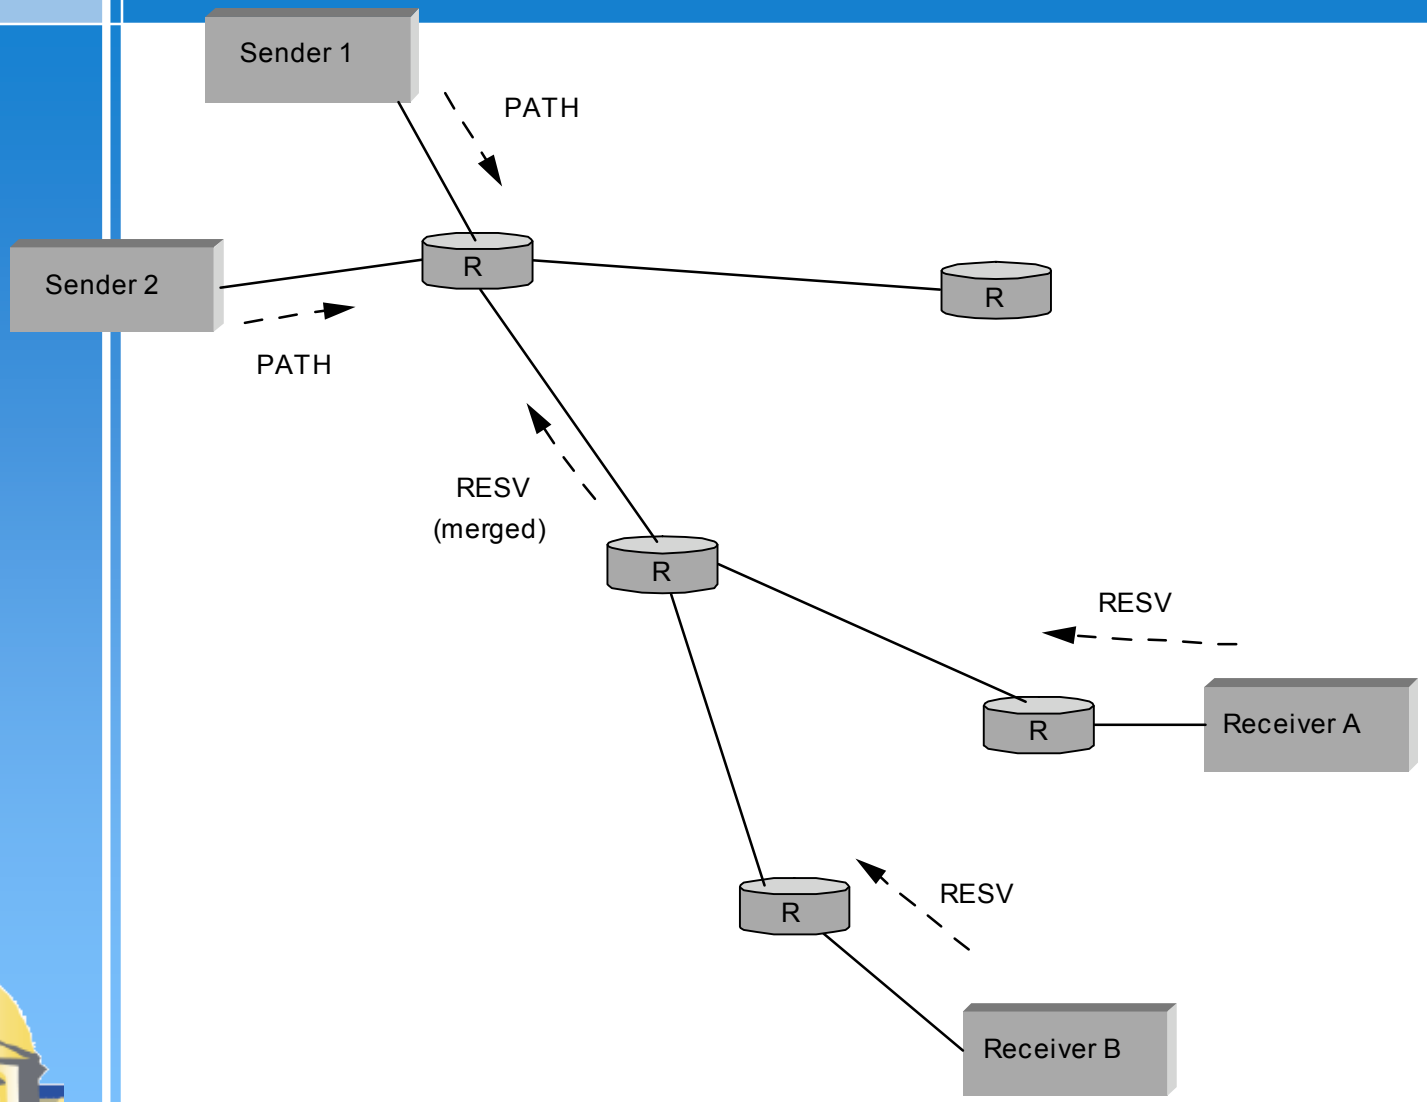

RSVP versus ATM (Q.2931)

- RSVP
  - receiver generates reservation
  - soft state (refresh/timeout)
  - separate from route establishment
  - QoS can change dynamically
  - receiver heterogeneity
- ATM
  - sender generates connection request
  - hard state (explicit delete)
  - concurrent with route establishment
  - QoS is static for life of connection
  - uniform QoS to all receivers

Differentiated Services

- Problem with IntServ: scalability, IntServ operates in a per-flow basis
- Idea: segregate packets into a small number of classes
  - e.g., premium vs best-effort
- Packets marked according to class at edge of network (ND will mark certain packets)
- Core routers implement some per-hop-behavior (PHB)
  - Example: Expedited Forwarding (EF)
    - rate-limit EF packets at the edges
  - PHB implemented with class-based priority queues or Weighted Fair Queue (WFQ)

DiffServ (cont)

- Assured Forwarding (AF)
  - customers sign service agreements with ISPs
  - edge routers mark packets as being “in” or “out” of profile
  - core routers run RIO: RED with in/out