Multimedia Networking

- Application classes
  - streamed stored audio/video
  - one-to-many (multicast) streaming of real-time a/v
  - real-time interactive audio/video
- Typical application issues
  - packet jitter
  - packet loss / recovery
- Internet protocols for multimedia
  - RTSP
  - RTP/RTCP
  - H.323
- Text: Kurose-Ross, Chapter 6

Example Multimedia Apps

- Streamed stored audio/video
  - movies, CS-653 taped lectures (available on MANIC)
- One-to-many streaming
  - News broadcasts, popular movies
- Real-Time Interactive
  - IP telephony, teleconference, distributed gaming

Multimedia terminology

- **Multimedia session**: a session that contains several media types
  - e.g., a movie containing both audio & video
- **Continuous-media session**: a session whose information must be transmitted "continually"
  - e.g., audio, video, but not text (unless ticker-tape)
- **Streaming**: application usage of data during its transmission

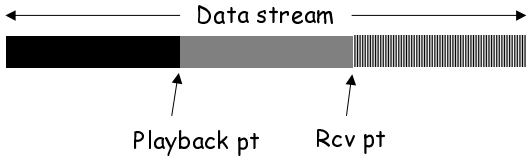

IntServ / DiffServ

- Integrated Services (IntServ)
  - Resource Reservation Protocol: RSVP
- Differentiated Services (DiffServ)
  - Assured Forwarding
  - Expedited Forwarding
  - Comparison: AF vs. EF

- Reading:Kurose-Ross Chapter 6.7-6.9

QoS (Quality of Service) in the Internet

- The Internet Protocol (IP) does not guarantee QoS to applications
- Idea: Re-engineer IP to provide quality of service
  - Let routers distinguish classes of flows
- Q: What is the model for a class?
- Solution must consider:
  - the needs of (some/most/all) applications
  - the add'l state that routers must maintain
  - the add'l communication overhead (add'l packets or bits)
  - # of flows or classes a router must handle

IntServ vs. DiffServ

IntServ (1992)

- per-flow reservations
  - i.e., needs RSVP
- flows provide traffic characterization
- "heavy" state: per-flow
- "strong" guarantees, e.g.,
  - conformance to leaky-bucket characterization [RFC 2215]
  - bound on max e2e delay [RFC 2212]

DiffServ (1997)

- packet classification
- edge & core routers
  - edge: "heavy" state
  - core: "light" state
- "weak" guarantees, e.g.,
  - Flow A gets better service than Flow B

Integrated Services (1992)

As described in [RFC 1633] from 1994:

- Philosophy: "guarantees cannot be achieved without reservations"
- Four components to IntServ architecture:
  - packet scheduler
  - classifier
  - admission control routine
  - reservation setup protocol

IntServ Components

All components implemented at all routers!

- Packet Scheduler
  - Manages forwarding of different streams
  - Required resources: sets of queues, timers
  - Example: Implementation of Weighted Fair Queuing (WFQ)
- Classifier
  - Maps packets to a class
  - Packets in same class treated similarly
  - Examples:
    - per-flow class
    - video-packet class

IntServ Components (cont'd)

- Admission Controller
  - Determines whether or not to admit a new flow
  - Q: why would a flow be rejected?
  - Requirements:
    - Knowledge of available resources at router
    - (conservative) model of flow's resource consumption
      - e.g., leaky bucket
  - The hard part: getting apps to characterize their flows
- Reservation Setup Protocol
  - Sets up and maintains (distributed) flows' network resource usage
    - "negotiates" between admission controllers at routers
    - establishes active classifiers at routers
  - e.g., RSVP protocol

RSVP protocol

- The commonly suggested reservation setup protocol
- Designed for multicast sessions (unicast is a special case)
- Receiver-oriented: receivers initiate requests
  - allows for receiver heterogeneity
- Reservation styles allow merging of reservations (i.e., use the style that's appropriate for the app)
- Uses soft-state: reservations need to be refreshed or they expire. Why soft-state?
- Dynamic: able to reconfigure reservation rather than perform complete teardown / setup

RSVP messaging

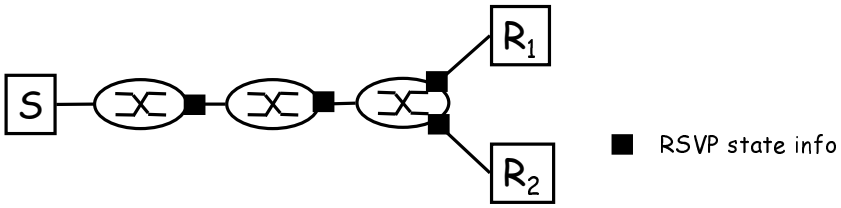

- Rcvrs make requests for reservations
- Sufficient resources: Router reserves per outgoing interface (i.e., link) and forwards request upstream
- Insufficient: send ResvError message downstream
- Path messages: from sender toward rcvr so that routers know where to forward receiver requests.
  - Why not just head toward sender using Internet routing tables?

RSVP Reservation Styles

- Fixed-Filter: Allocation per sender indicated
  - Sample application: multimedia (e.g., send audio (S_1) and video (S_2) at same time)

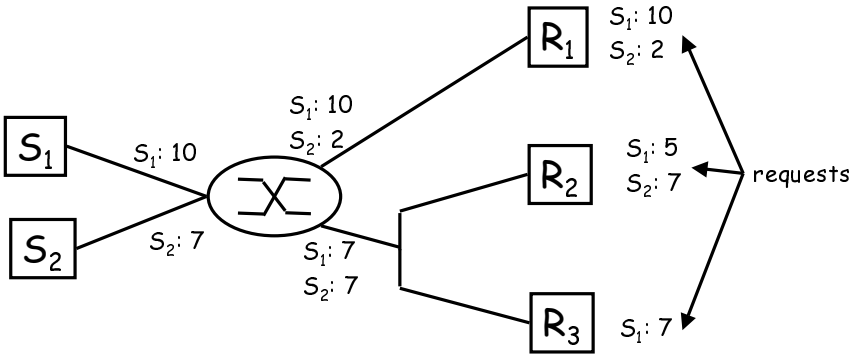

RSVP Reservation Styles

- Shared-Explicit: Allocation shared by list of senders
  - Sample application: multimedia (e.g., debate w/ 2 speakers)

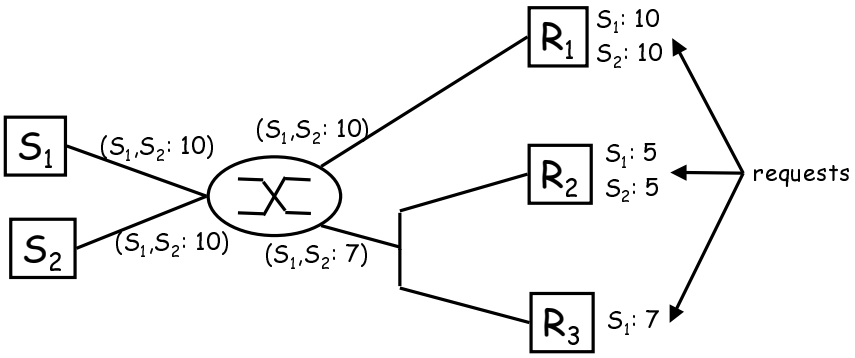

RSVP Reservation Styles

- Wildcard-Filter: Allocation shared by all senders
  - Sample application: town meeting (one sender, but not clear who the speakers might be)

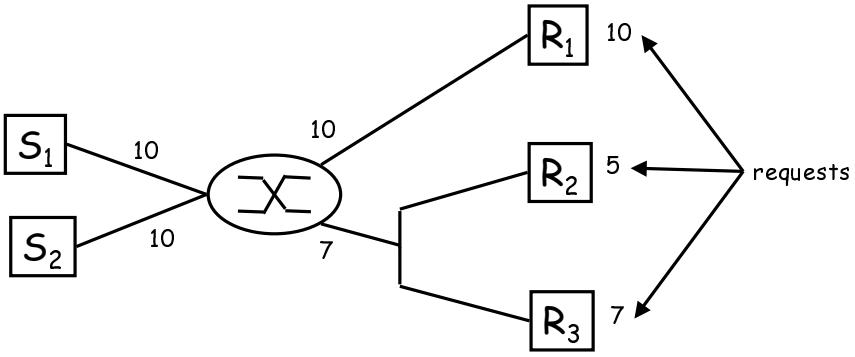

Style Summary

- Fixed-Filter: reservation per sender
  - Senders don't "share" bandwidth
  - Dynamic event: rcvr wants to change a sender allocation
- Shared-Explicit: reservation per list-of-senders
  - Fixed set of senders "share" bandwidth
  - Dynamic event: rcvr wants to add/remove sender or change group allocation
- Wildcard-Filter: no sender specified w/ reservation
  - Any sender can "share" bandwidth
  - Dynamic event: new sender begins transmitting, rcvr wants to increase its receiving allocation

IntServ: Problems

- Reservation protocols and structure complicated
  - lots of message passing
  - coordination problems
- All routers maintain state
  - state maintenance requires additional processing / memory resources
  - Lots of flows traverse core (backbone) routers
    - Lots of state: need more memory
    - Lots of RSVP msgs to process: slows transfer speeds
    - Scheduler and Classifier have too much to deal with

DiffServ

- Q: What if IntServ is too complex/costly to deploy?
- A: Build a simpler scheme that takes into account
  - many apps have simple requirements (e.g., need fixed bandwidth, low jitter)
  - App can't/doesn't always conform to/provide "strict" model of resource usage
  - different levels of functionality can be placed at different "types" of routers
    - network edge
    - network core

Differentiated Services

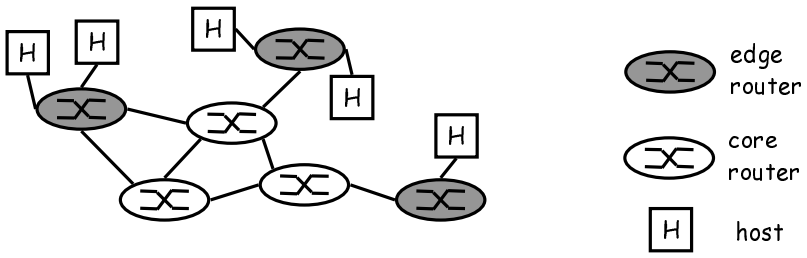

- Idea: keep the architecture simple within the core.
  - higher complexity permitted at edge routers
- Just provide service differences, no explicit guarantees
  - i.e., high and low priority classes (extra $$$ for high)

DiffServ Architecture

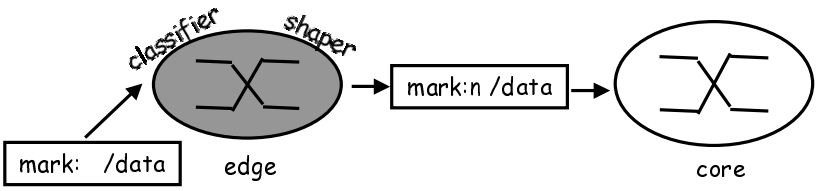

- Edge router
  - classify packet and mark packet
  - shape flow (control entry rate into core, drop pkts, change mark, etc.)
- Core router
  - handle packet based on its mark
  - possibly remark at peering points
- Maintain **Per-Hops Behavior** (PHB): the desired service (e.g. rate) provided to a class at a given hop (router)

2 Competing PHBs

- Expedited Forwarding (EF) [RFC 2598]
  - Router must support classes' configured rates
  - EF class allocated fixed portion of router processing per unit time, e.g.,
    - Class-based queueing (CBQ) w/ priority to EF queue
    - Weighted Fair Queuing
- Assured Forwarding (AF) [RFC 2597]
  - N classes (current standard: N=4)
  - M possible drop preferences w/in class (current standard: M=3)
  - Each classes' traffic handled separately
  - Packet drop "likelihood" increases w/ drop preference

PHB Specs Omit...

- EF and AF PHBs do not specify mechanism, e.g., not specified:
  - edge classification, shaping or marking policy
  - core router queuing mechansim
  - ranges of rates, relative class/preference service ratios, etc.
- Why are these details omitted?
  - Allow flexibility - as long as specified requirements are met.
  - DiffServ is a new idea - still unclear on which mechanism is best - so standardize later
- Which is better, EF or AF?

Comparing PHB Models [Sahu]

- How does isolating traffic (EF) compare with preferential treatment (AF w/ preferences)?

- Measures:
  - expected loss rate
  - expected delays
- Reqm't: overall buffer / bandwidth fixed

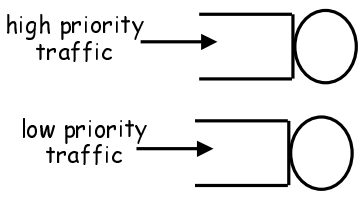

- Queue models:
  - EF: separate queues per class. High priority queue always serviced first (when non-empty)
  - AF: one queue w/ threshold for accepting high drop-preference pkts

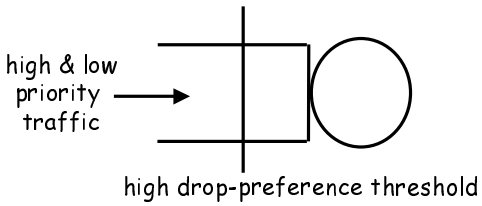

Intuition

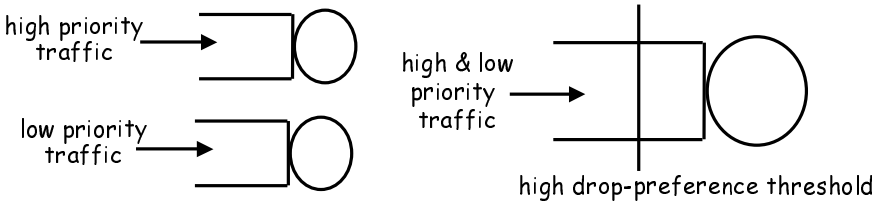

- Which is
  - better for reducing high-priority delay?
  - for high-priority loss?
- How should buffer be allocated in the 2 models to make a fair comparison?
  - EF: low priority gets its own buffer
  - AF: low priority must share its buffer with high

EF vs. AF Comparison

- Choose buffer partitions and threshold such that low-priority traffic sees similar loss rates in two systems
- Examine impact on high priority traffic
- Main Results for high priority traffic:
  - AF router needs to process 30-70% faster than an EF router to maintain same delays (function of partition point and threshold location)
  - EF router needs only 15% add'l buffer to yield same loss rates to low priority traffic as AF

DiffServ Open Issues

- How to decide "how much" to reserve
- How to do DiffServ for multicast
  - Much more complicated
  - Multicast reservation issues significantly complicated IntServ. What about DiffServ?

Summary: Internet Multimedia

- Internet design:
  - flexible
  - easy to extend
  - difficult to support time-bounded applications
- Approach #1: Build on a best-effort network
  - adaptive applications (quality vs. available bandwidth)
  - deal with loss and jitter (e.g., RTP/RTCP)
- Approach #2: Modify (extend) IP design
  - IntServ: guarantee QoS, but takes lots of state
  - DiffServ: create high and low priority customers - give more to high

Multimedia Internet Protocols

- We'll look at 3:
  - RTP/RTCP: transport layer
  - RTSP: session layer for streaming media applications
  - H.323: session layer for conferencing applications

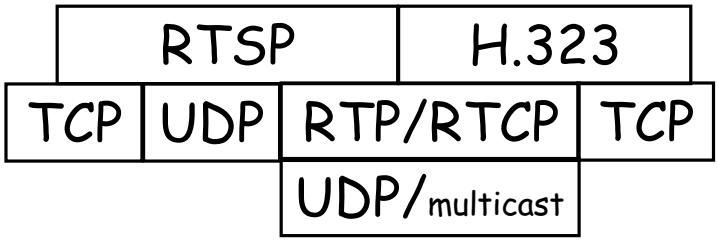

RTP/RTCP [RFC 1889]

- Session data sent via RTP (Real-time Transfer Protocol)
- RTP components / support:
  - sequence # and timestamps
  - unique source/session ID (SSRC or CSRC)
  - encryption
  - payload type info (codec)
- Rcvr/Sender session status transmitted via RTCP (Real-time Transfer Control Protocol)
  - last sequence # rcvd from various senders
  - observed loss rates from various senders
  - observed jitter info from various senders
  - member information (canonical name, e-mail, etc.)
  - control algorithm (limits RTCP transmission rate)

RTP/RTCP details

- All of a session's RTP/RTCP packets are sent to the same multicast group (by all participants)
  - All RTP pkts sent to some even-numbered port, 2p
  - All RTCP pkts sent to port 2p+1
- Only data senders send RTP packets
- All participants (senders/rcvrs) send RTCP packets

RTP header

Payload Type	Sequence #	Timestamp	Synchronization Source Identifier	Misc

- Why do most (all) multimedia apps require
  - sequence #?
  - timestamp?
  - (unique) Sync Source ID?
- Why should every pkt carry the 7-bit payload type?
  - Why not just when sender initiates session?
- Transmission rate: application specific (no congestion control specified in RTP)

RTCP packets

- There are several types of RTCP packets
  - SR: sender report - transmission & reception stats
  - RR: receiver report - reception stats
  - SDES: Source description items
  - BYE: end-of-participation message
  - APP: application-specific functions
- Typically, several RTCP packets of different types are transmitted w/in a single UDP packet

What RTCP provides

- Info to detect colliding Synch source ID's
- Contact info (e-mail, true name) of participants
- Info to count # of session participants
- Reception quality of all participants

- How does RTCP avoid creating congestion if all participants send RTCP packets?
  - consider a broadcast TV transmission

RTCP congestion control

- A session's aggregate RTCP bandwidth usage should be 5% of the session's RTP bandwidth
  - 75% of the RTCP bandwidth goes to the receivers
  - 25% goes to the senders
- $T_{sender} = \frac{\text{\# senders * avg RTCP pkt size}}{.25 * .05 * \text{RTP bandwidth}}$
- $T_{rcvr} = \frac{\text{\# receivers * avg RTCP pkt size}}{.25 * .05 * \text{RTP bandwidth}}$

Send at (.5 + rand(0,1)) * T : why?

How does each member know # of senders, # rcvrs?

RTCP reconsideration

- Goal: prevent RTCP bandwidth explosion if everybody joins at once
  - Receivers who initially join will count small # of session members
- Solution when first joining:
  1. Compute T, and wait random time interval
  2. At end of interval, reassess # of members
  3. If # of members increased, compute a new T'
  4. If T' < T, send immediately
  5. If T' >= T, wait an additional T', go to step 2
- Other times, use normal wait period

RTSP [RFC 2326]

- RTSP: out-of-band protocol used to control transmission of a media-stream
  - VCR-like functionality (pause/resume, FF, rewind, reposition, etc.)

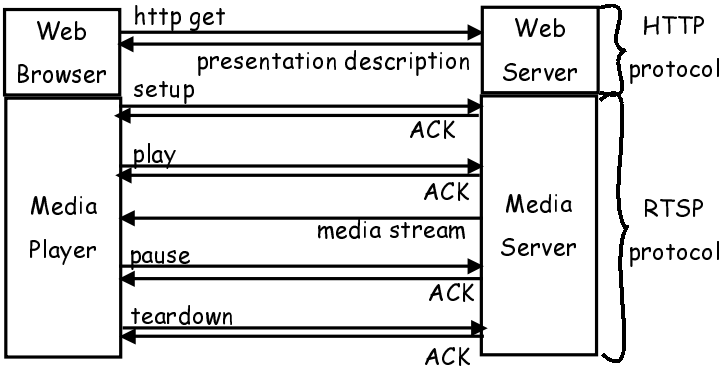

H.323

- A standard for real-time audio / video teleconferncing on the Internet
- Network Components:
  - **end points**: end-host H.323-compliant app
  - **gateways**: interface between H.323-compliant communication and prior technology (e.g. phone network)
  - **gatekeepers**: provide services at gateway (e.g., address translation, billing, authorization, etc.)

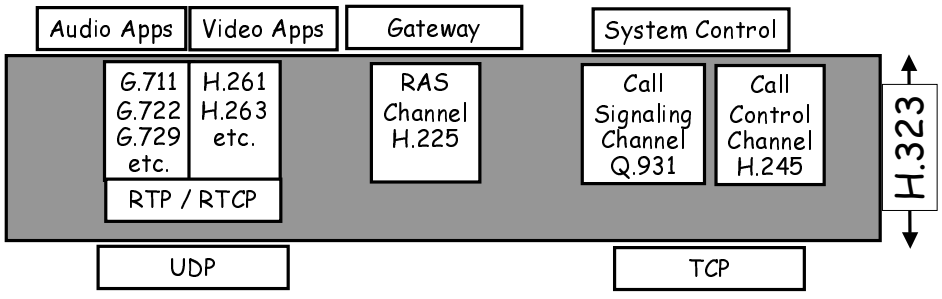

H.323 cont'd
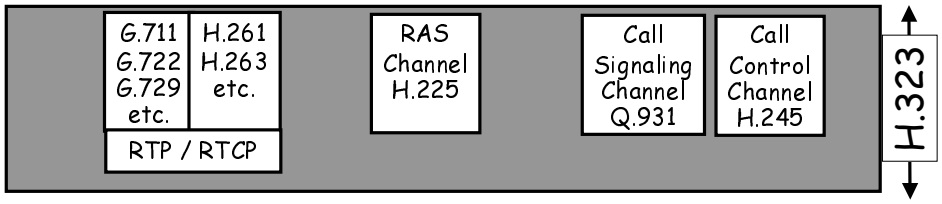
- H.225: notifies gatekeepers of session initiation
- Q.931: signalling protocol for establishing and terminating calls
- H.245: out-of-band protocol negotiates the audio/video codecs used during a session (TCP)