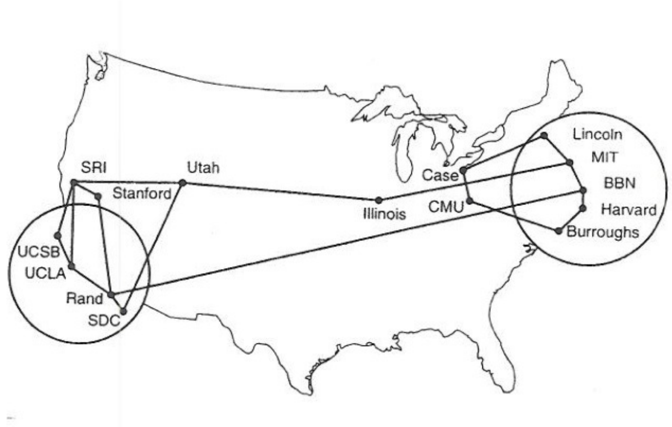

QoS - Quality of Service

Outline

Real-time Applications -- timely arrival

Integrated Services

Differentiated Services

New Applications

- Require Delivery Assurances
  - Voice, Video, Industrial Control, etc.
  - Need more than Best Effort
  - Late data worthless
- IETF Efforts - Requires underlying Internet infrastructure
  - Building Protocol Extensions
    - admission control
    - Controlled load
    - Guaranteed service
  - Still need to support Best Effort

Example: Real-time Application

- Require "deliver on time" assurances
  - must come from *inside* the network

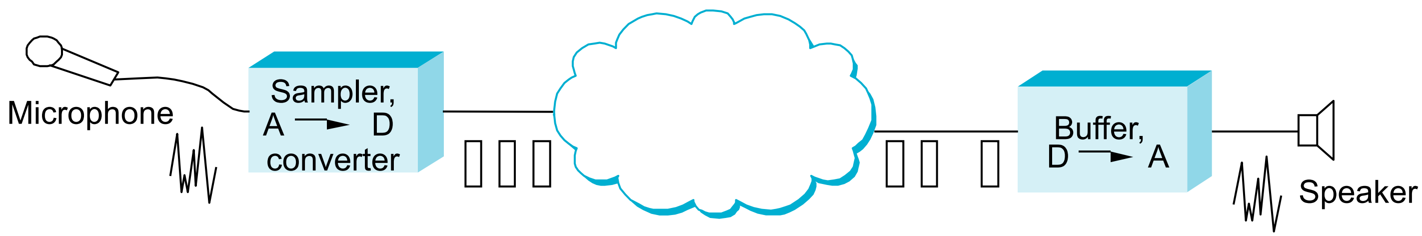

- Example application (audio)
  - sample voice once every 125us
  - each sample has a *playback time* (when needed at Receiver)
  - packets experience variable delay in network
  - add constant factor to playback time: *playback point*

RealTime needs definition

Playback Buffer

Playback Point – limited by application, e.g., voice or video

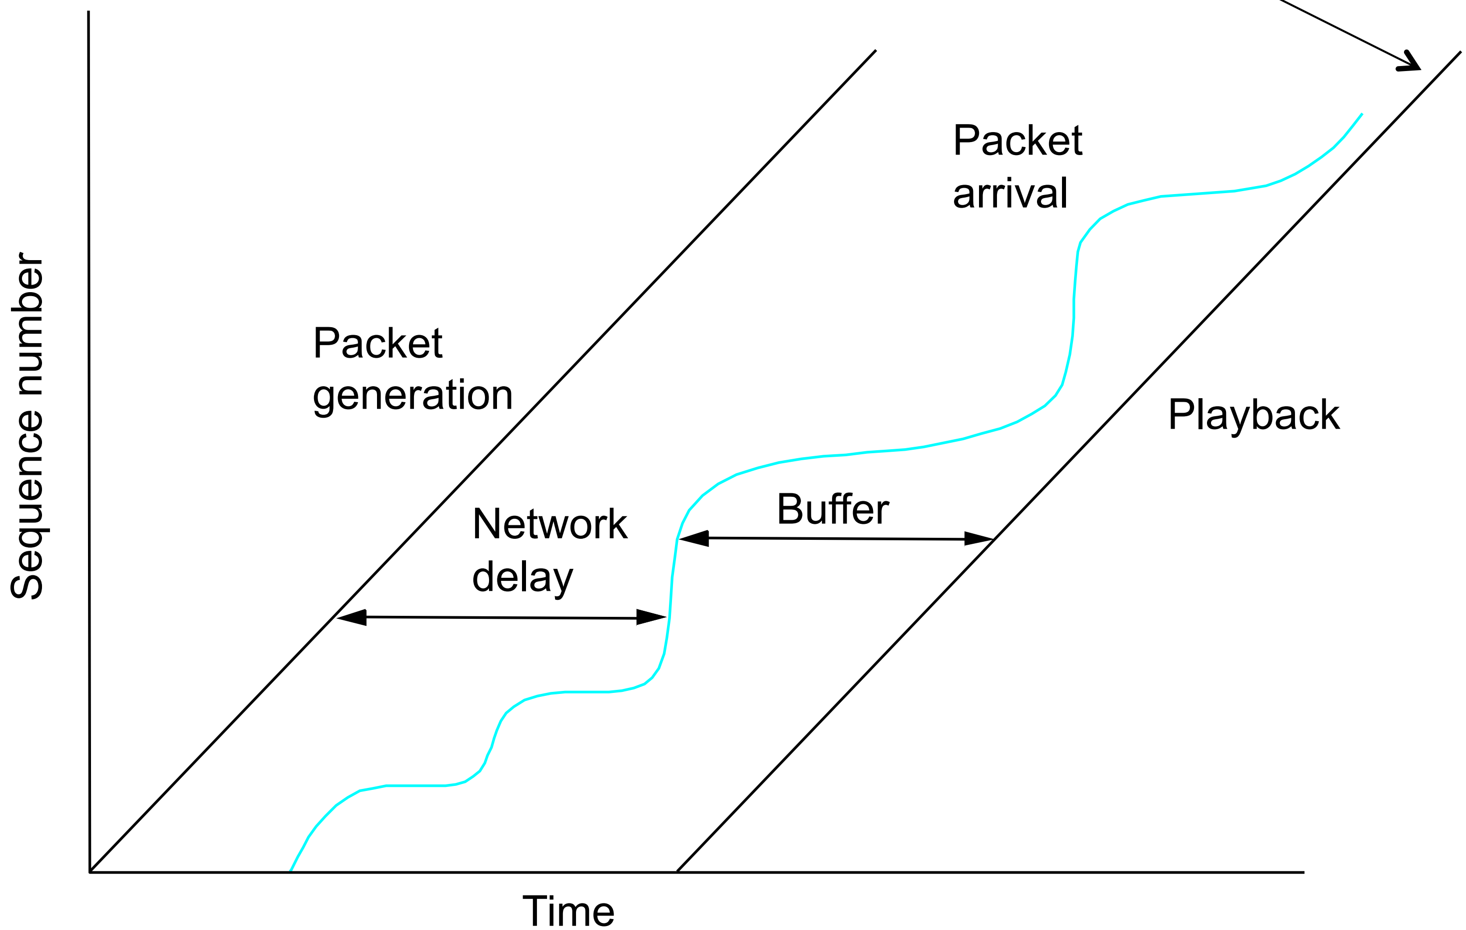

Example Distribution of Delays

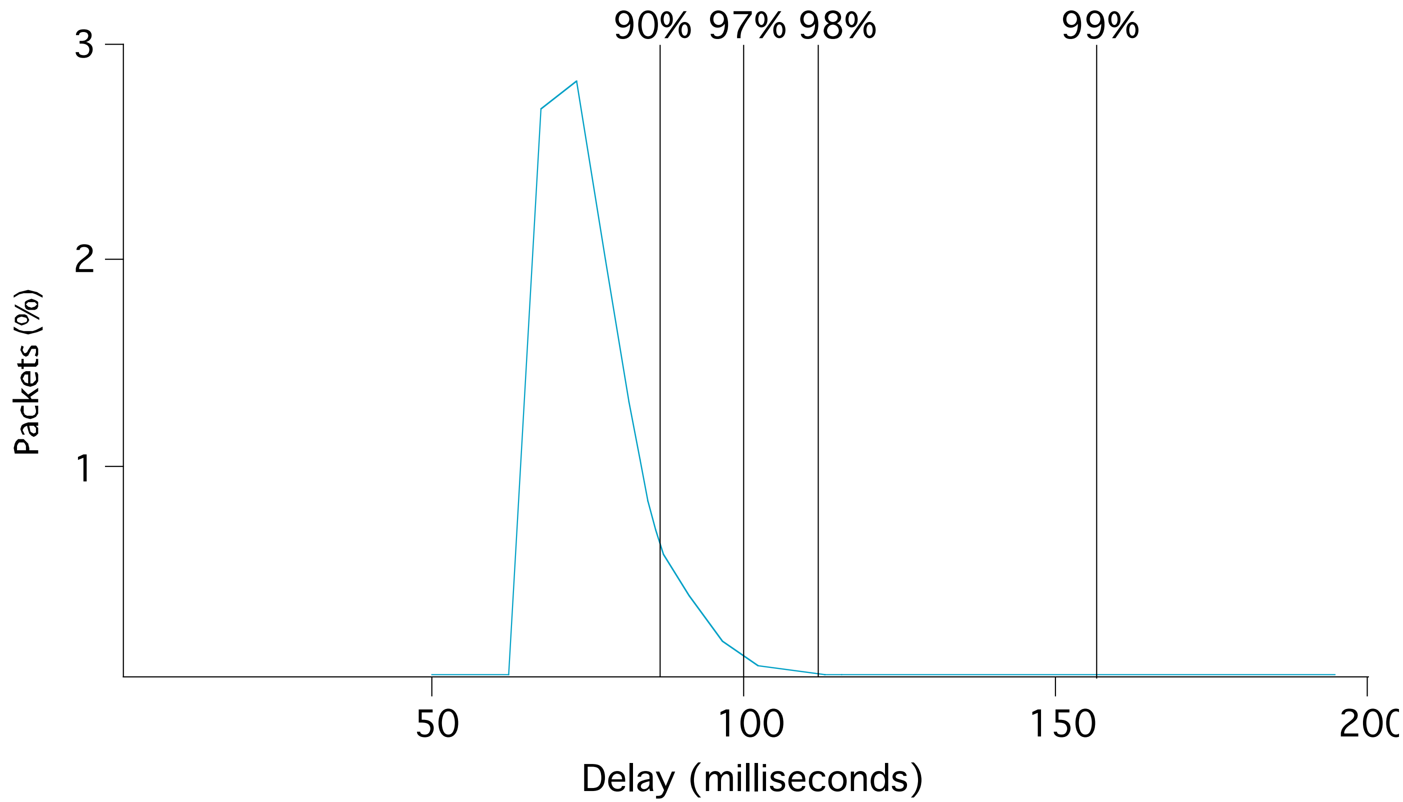

Delay is variable, need to know 'variability of delay' Can you?

Taxonomy

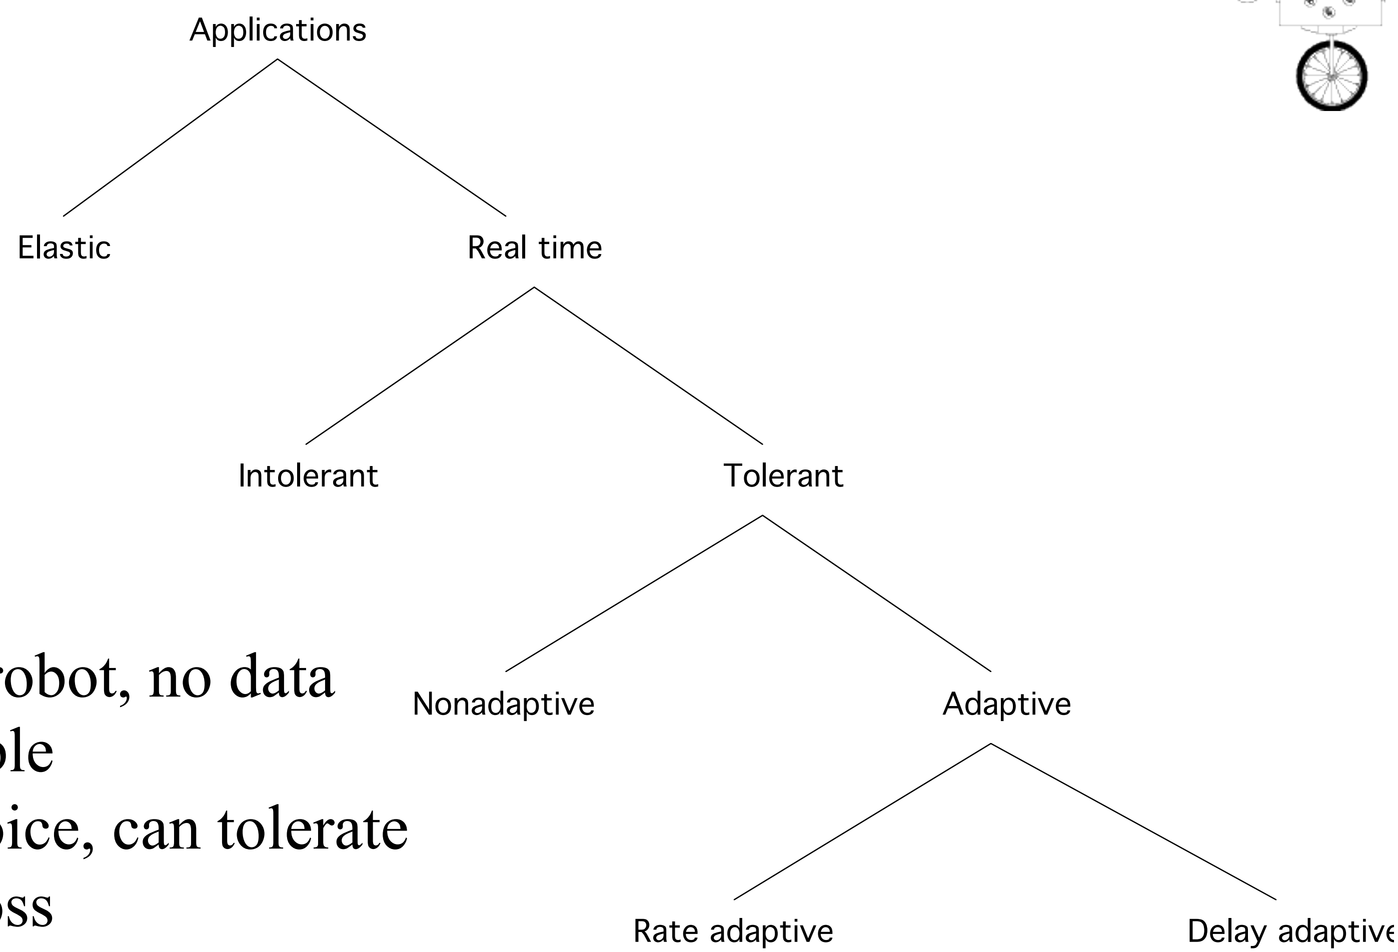

Intolerant – robot, no data loss acceptable
Tolerate – voice, can tolerate occasional loss
Adaptive – can move playback point

QoS Approaches

- End-to-end QoS guarantees- extreme is circuit based network
- Abstraction is unidirectional flow
- Initiated by endpoint
- Fine Grained
  - Setup by individual apps or flows
  - Integrative Services
- Course Grained
  - Apply to large classes of network data
  - Differentiated Services

Integrated Services - RSVP

- Service Classes
  - guaranteed – packet never arrives late
  - controlled-load – tolerant, adaptive apps, VAT Audio uses WFQ
- Mechanisms
  - signaling protocol – what app needs to tell network - Flowspec
  - admission control – router yes or no
  - policing – penalty for badness
  - packet scheduling – what router does to handle Flowspec

Integrated Services - RSVP

An endpoint uses RSVP to request a simplex flow through an IP Internet with specified QoS bounds. If routers along the path agree to honor the request, they approve it, otherwise, they deny it. If an application needs QoS in two directions, each endpoint must use RSVP to request a separate flow.

Issue from start: Overhead in setting up….Router now has to constantly manage its environment so that it can respond to requests….

Flowspec

- ***Rspec:*** describes service requested from network, i.e., what I want
  - controlled-load: none
  - guaranteed: delay target
- ***Tspec:*** describes flow's traffic characteristics
  - average bandwidth + burstiness: *token bucket* filter
  - token rate *r*
  - 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
  - allow intelligent admission control decisions – rate, bursts, pauses

Token Bucket Filter

- Algorithm:

  I can send a burst of B, but over sufficiently long period, canNOT send more than R bytes/sec. If my flow is 2MBs, can send 1MBs for 2 seconds, then 4 MBs for 1 second, then 2MBs for N secs, etc.

  Allowing bursts, but forcing a constant rate over time

Per-Router Mechanisms

- Admission Control - per flow
  - decide if a new flow can be supported
  - answer depends on service class requested & Router queueing
  - not the same as *policing* - penalty phase
- Packet Processing
  - classification: associate each packet with the appropriate reservation – more than just flow management….
  - scheduling: manage queues so each packet receives the requested service – priority queues in great detail…

RSVP - Resource 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 (TSPEC)
- Two messages: PATH (Tspec) and RESV
- Source transmits PATH messages every 30 seconds
- Destination responds with RESV message (TSPEC + RSPEC)
- Merge requirements in case of multicast
- Can specify number of speakers

RSVP Example

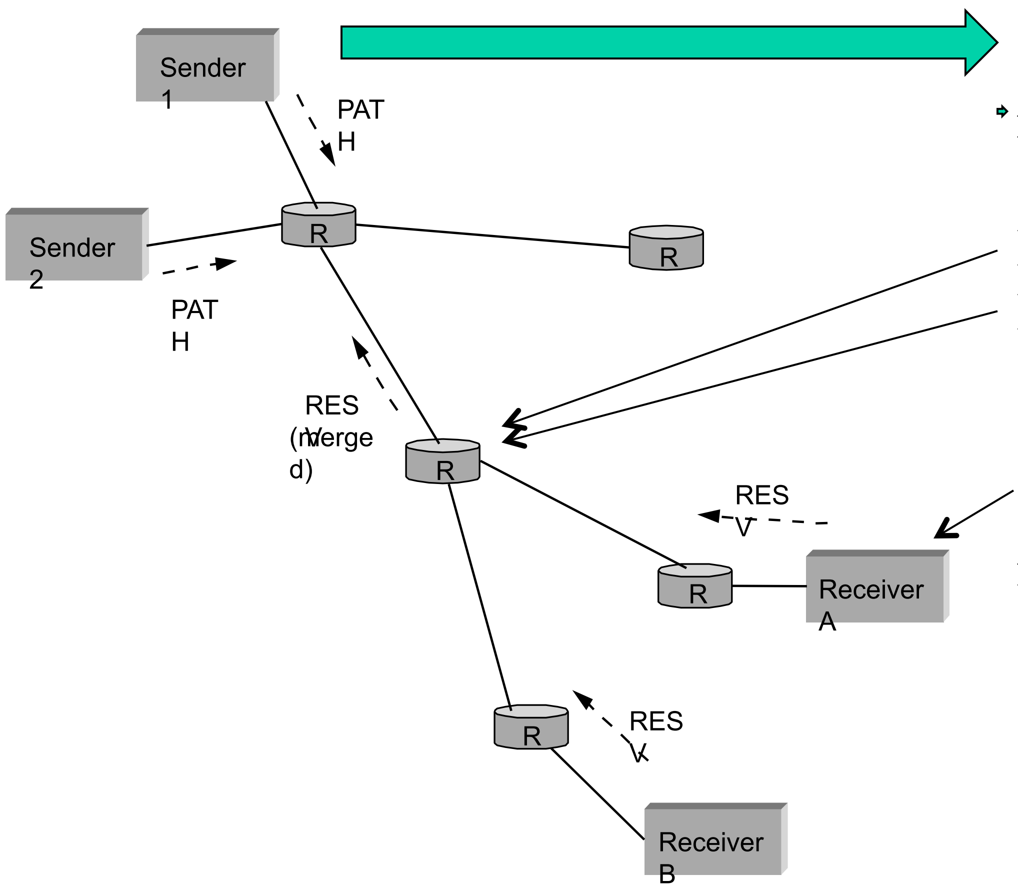

- TSPEC from Sender to anyone receiving info
- Merge requests for multicast
- Has softstate refreshed every 30 secs
- Contains Sender's TSPEC and my RSPEC.. what I need

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

RSVP Notes

- RSVP Defines
  - Message endpoint sends to router to request QoS
  - Messages routers send to other routers
  - Replies
- RSVP does not specify how enforcement done
- Separate protocol needed
- Common Open Policy Services (COPS)
  - Proposed enforcement protocol for RSVP
  - Known as *traffic policing*
  - Checks data sent on flow to ensure that the flow does not exceed pre-established bounds.

Summary

- Codec translates between analog and digital forms
- RTP used to transfer real-time data
- RTP adds timestamp that sender uses to determine playback time
- Voice over IP uses
  - RTP for digitized voice transfer
- RSVP and COPS provide QoS

Integrated Services

- Routers MUST
  - Classify Packets -- which flow
  - Schedule Packets -- when forwarded
- Classify
  - Source
  - Destination
  - Protocol
  - Port S/D
  - Have TOS bits in IP header, but cannot differentiate flows by TOS
  - e.g., premium vs best-effort
- Why NOT the Internet Standard?
  - Scale
    - Per Flow Info
    - OC48 - 64 kbps for streaming audio
      - 39,000 flows ….
    - Used for MPLS

Another Approach: DiffServ -- Differentiated Services

- Problem with IntServ: scalability
- Idea: segregate packets into a small number of classes
  - e.g., premium vs best-effort
- Packets marked according to class at edge of network
- 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 WFQ
- To Me – Priority based scheduling and Buffering at each Router

Differentiated Services

- Small number of classes of traffic
- Premium bit – set by Admin Boundary
- Router Behaviors - per hop (PHB)
- Use 6 bits to establish -- easy parsing for Router…
  - EF, Expedite Forwarding, minimal delay and loss
    - control input traffic – do not let combined input streams swamp router – edge routers
    - in/out packets
      - in – within service limits
      - out – outside service limits (drop if must, RED)
  - Non EF packets
    - need to provide some service to non-EF packets

DiffServ (cont)

- Assured Forwarding (AF), PHB – Preferred Hop Behavior
  - 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
  - Based on SLA – service level agreement

Probability

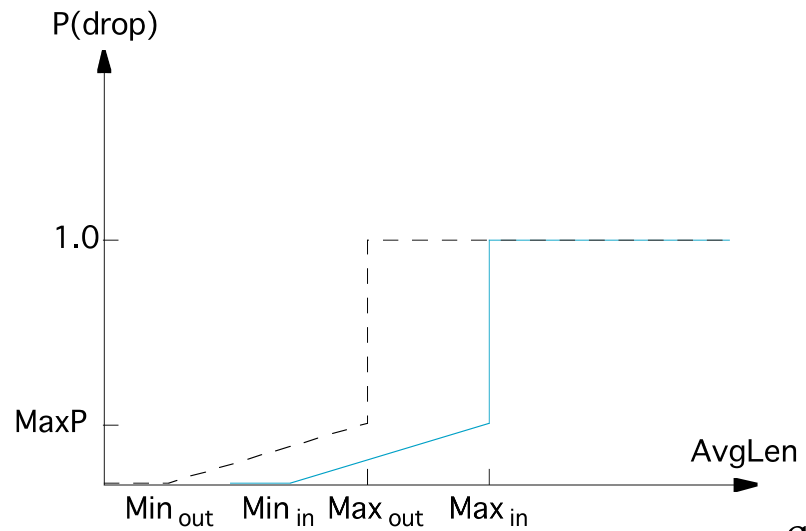

queue

Figure
Under low congestion only "out" packets discarded (policing)
Only when queue > Min(in) do "in" packets get dropped

Future

- TCP & Congestion will continue
- ACKs & timeouts drive end hosts
- Today - QoS
  - Works such as Diffserv Not universal, but VoIP, video driving Diffserv use
  - also helps when Diffserv not enough, i.e., routers like airlines over-subscribe for Hosts to be involved.
- UDP?
  - no connection overhead
  - no reliability overhead
  - real time UDP competes with TCP
  - right now, TCP, only one that backs off
  - open research area