QoS over Data Networks

Raj Jain, Jennifer Hou, Jason Han

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
http://www.cse.wustl.edu/~jain/

Integrated Services

- Best Effort Service: Like UBR.
- Controlled-Load Service: Performance as good as in an unloaded datagram network. No quantitative assurances. Like nrt-VBR or UBR w MCR
- Guaranteed Service: rt-VBR
  - Firm bound on data throughput and delay.
  - Delay jitter or average delay not guaranteed or minimized.
  - Every element along the path must provide delay bound.
  - Is not always implementable, e.g., Shared Ethernet.
  - Like CBR or rt-VBR

Problems with RSVP and Integrated Services

- Complexity in routers: packet classification, scheduling
- Scalable in number of receivers per flow but
  Per-Flow State: O(n) $\Rightarrow$ Not scalable with # of flows.
  Number of flows in the backbone may be large.
  $\Rightarrow$ Suitable for small private networks
- Need a concept of “Virtual Paths” or aggregated flow groups for the backbone
- Need policy controls: Who can make reservations?
  Support for accounting and security.
  $\Rightarrow$ RSVP admission policy (rap) working group.

Problems (Cont)

- Receiver Based: Need sender control/notifications in some cases. Which receiver pays for shared part of the tree?
- Soft State: Need route/path pinning (stability).
  Limit number of changes during a session.
- RSVP does not have negotiation and backtracking
- Throughput and delay guarantees require support of lower layers. Shared Ethernet $\Rightarrow$ IP can't do GS or CLS. Need switched full-duplex LANs.
- Can't easily do RSVP on ATM either
- Most of these arguments also apply to integrated services.

Differentiated Services

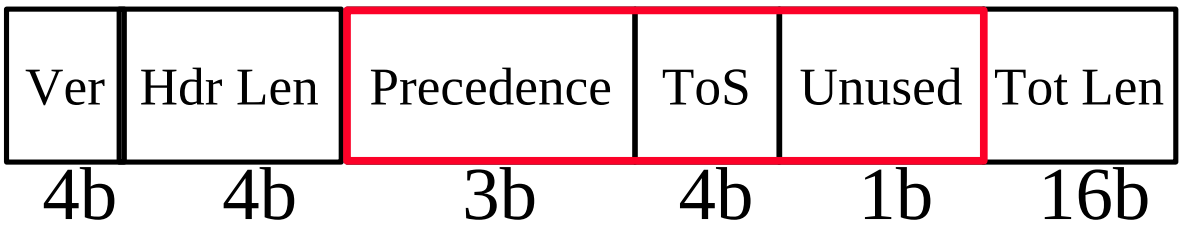

- IPv4: 3-bit precedence + 4-bit ToS
- Many vendors use IP precedence bits but the service varies ⇒ Need a standard ⇒ Differentiated Services
- DS working group formed February 1998
- Charter: Define ds byte (IPv4 ToS field)
- Per-Hop Behavior: Externally Observable Forwarding Behavior, e.g., x% of link bandwidth, or priority

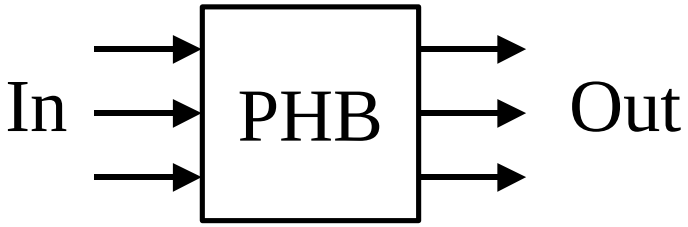

Expedited Forwarding

- Also known as “Premium Service”
- Virtual leased line
- Similar to CBR
- Guaranteed minimum service rate
- Policed: Arrival rate < Minimum Service Rate
- Not affected by other data PHBs
  ⇒ Highest data priority (if priority queueing)

Assured Forwarding

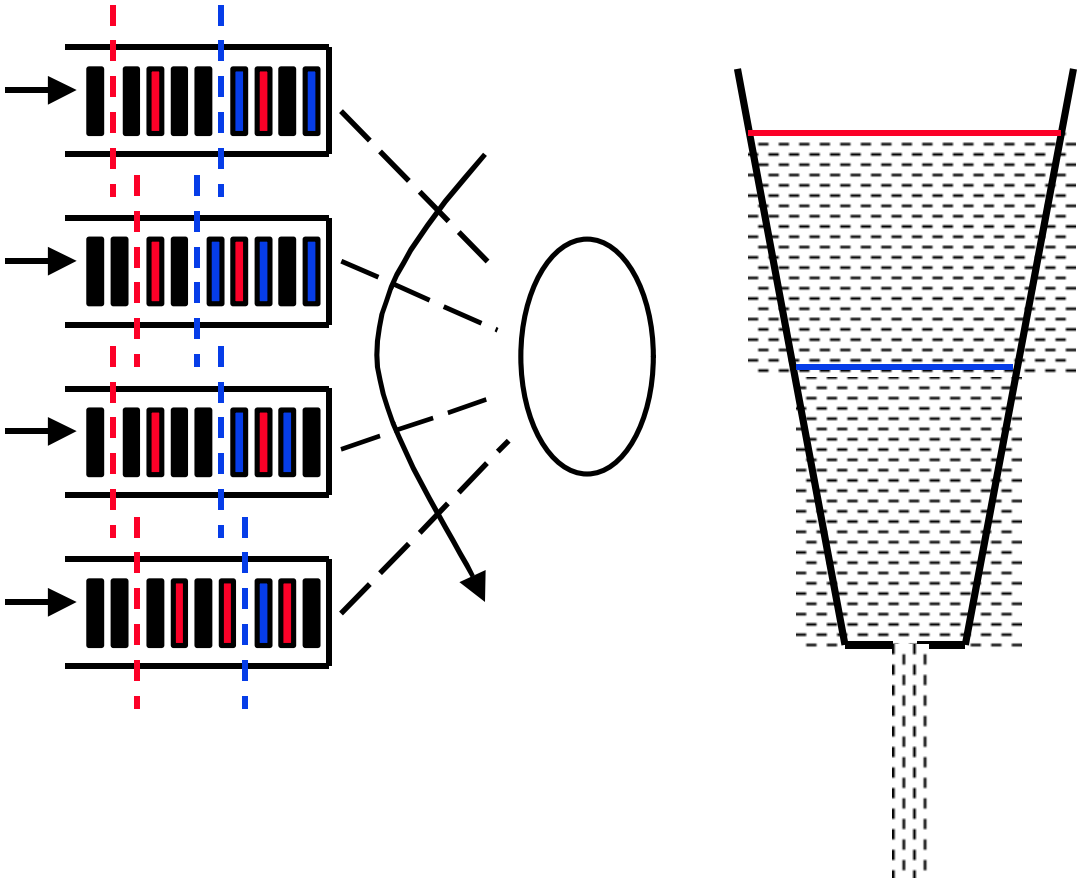

- PHB Group
- Four Classes: Decreasing weights in WFR/WFQ
- Three drop preference per class (one rate and two bucket sizes)

DiffServ: Issues

- per-hop $\Rightarrow$ Need at every hop
  One non-DiffServ hop can spoil all QoS
- End-to-end $\neq \Sigma$ per-Hop
  Designing end-to-end services with weighted guarantees at individual hops is difficult. Only EF will work.
- QoS is for the aggregate not micro-flows. Not intended/useful for end users. Only ISPs.
  - Large number of short flows are better handled by aggregates.

DiffServ Issues (Cont)

- Long flows (voice and video sessions) need per-flow guarantees.
- High-bandwidth flows (1 Mbps video) need per-flow guarantees.

- All IETF approaches are open loop control ⇒ Drop.
  Closed loop control ⇒ Wait at source
  Data prefers waiting ⇒ Feedback
- Guarantees ⇒ Stability of paths
  ⇒ Connections (hard or soft)
  Need route pinning or connections.

Multiprotocol Label Switching

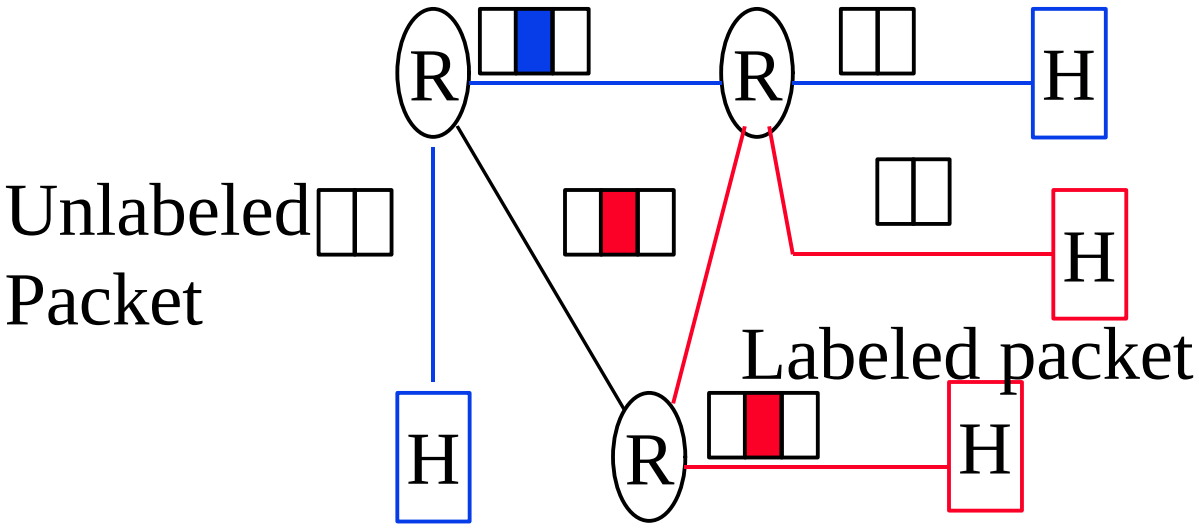

- Entry “label switch router (LSR)” attaches a label to the packet based on the route
- Other LSRs switch packets based on labels. Do not need to look inside
- Labels have local significance
  ⇒ Different label at each hop
  (similar to VC #)
- Exit LSR strips off the label

Traffic Engineering Using MPLS

- Traffic Engineering
  = Performance Optimization
  = Efficient resource allocation,
  Path splitting
  ⇒ Maximum throughput, Min delay, min loss
  ⇒ Quality of service
- In MPLS networks:
  "Traffic Trunks" = SVCs
  Traffic trunks are routable entities like VCs
- Multiple trunks can be used in parallel to the same egress.
- Each traffic trunk can have a set of associated characteristics, e.g., priority, preemption, policing, overbooking

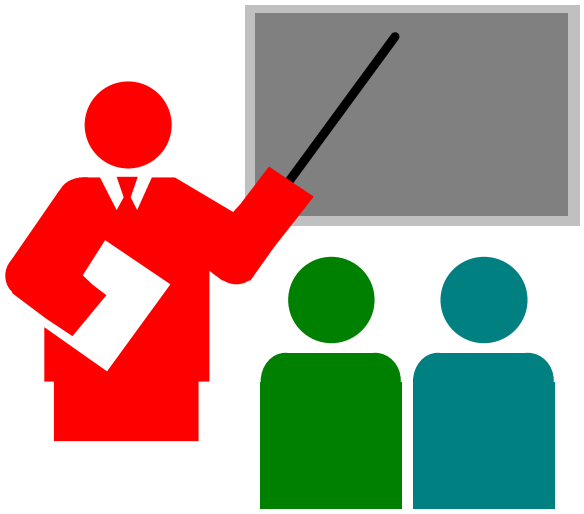

Summary

- Multiple drop preferences does not help data (TCP) or Voice/Video
- Voice/video need multiple leaky bucket rates for layered/scalable coding.
- Need additivity or mathematical aggregatability.
  CBR (EF) should be the first step for IP.
- Excess allocation is useful with closed loop. Network/application dynamics
  $\Rightarrow$ Need closed loop