


Quality of Service

Computer Networks, Winter 2010/2011




Chapter 6 outline


- 6.1 Providing multiple classes of service
- 6.2 Differentiated Services
- 6.3 Providing QoS guarantees
- 6.4 Integrated Services
- 6.5 Resource Reservation Protocol
- 6.6 Next Steps in Signaling
- 6.7 Multiprotocol Label Switching

Providing Multiple Classes of Service

- thus far: making the best of best effort service
 - one-size fits all service model
- alternative: multiple classes of service
 - partition traffic into classes
 - network treats different classes of traffic differently (analogy: VIP service vs regular service)
- granularity: differential service among multiple classes, not among individual connections
- history: ToS bits




Multiple classes of service: scenario



Scenario 1: mixed FTP and audio

- Example: 1Mbps IP phone, FTP share 1.5 Mbps link.
 - bursts of FTP can congest router, cause audio loss
 - want to give priority to audio over FTP




Principle 1

packet marking needed for router to distinguish between different classes; and new router policy to treat packets accordingly

Principles for QOS Guarantees (more)

- what if applications misbehave (audio sends higher than declared rate)
 - policing: force source adherence to bandwidth allocations
- marking and policing at network edge:
 - similar to ATM UNI (User Network Interface)




Principle 2

provide protection (*isolation*) for one class from others

Principles for QOS Guarantees (more)

- Allocating fixed (non-sharable) bandwidth to flow: inefficient use of bandwidth if flows doesn't use its allocation




Principle 3

While providing isolation, it is desirable to use resources as efficiently as possible



Scheduling And Policing Mechanisms

- scheduling: choose next packet to send on link
- FIFO (first in first out) scheduling: send in order of arrival to queue
 - real-world example?
 - discard policy: if packet arrives to full queue: who to discard?
 - Tail drop: drop arriving packet
 - priority: drop/remove on priority basis
 - random: drop/remove randomly




Scheduling Policies: more

- Priority scheduling: transmit highest priority queued packet
- multiple classes, with different priorities
 - class may depend on marking or other header info, e.g. IP source/dest, port numbers, etc..
 - Real world example?




Scheduling Policies: still more

- round robin scheduling:
- multiple classes
- cyclically scan class queues, serving one from each class (if available)
- real world example?



Scheduling Policies: still more

- Weighted Fair Queuing:
 - generalized Round Robin
 - each class gets weighted amount of service in each cycle
 - real-world example?




Policing Mechanisms

- **Goal:** limit traffic to not exceed declared parameters
- Three common-used criteria:
- **(Long term) Average Rate:** how many pkts can be sent per unit time (in the long run)
 - crucial question: what is the interval length: 100 packets per sec or 6000 packets per min have same average!
- **Peak Rate:** e.g., 6000 pkts per min. (ppm) avg.; 1500 ppm peak rate
- **(Max.) Burst Size:** max. number of pkts sent consecutively (with no intervening idle)

Policing Mechanisms



- **Token Bucket:** limit input to specified Burst Size and Average Rate.



- bucket can hold b tokens
- tokens generated at rate r token/sec unless bucket full
- over interval of length t : number of packets admitted less than or equal to $(r t + b)$.

Policing Mechanisms (more)

- token bucket, WFQ combine to provide guaranteed upper bound on delay, i.e., **QoS guarantee!**



Chapter 6 outline

- 6.1 Providing multiple classes of service
- 6.2 Differentiated Services
- 6.3 Providing QoS guarantees
- 6.4 Integrated Services
- 6.5 Resource Reservation Protocol
- 6.6 Next Steps in Signaling
- 6.7 Multiprotocol Label Switching

IETF Differentiated Services

- want “qualitative” service classes
 - “behaves like a wire”
 - relative service distinction: Platinum, Gold, Silver
- **scalability**: simple functions in network core, relatively complex functions at edge routers (or hosts)
 - signaling, maintaining per-flow router state difficult with large number of flows
- don't define service classes, provide functional components to build service classes


Diffserv Architecture

Edge router:

- per-flow traffic management
- marks packets as **in-profile** and **out-profile**


Core router:

- per class traffic management
- buffering and scheduling based on **marking** at edge
- preference given to **in-profile** packets



Edge-router Packet Marking

- **profile**: pre-negotiated rate A, bucket size B
- packet marking at edge based on **per-flow** profile



Possible usage of marking:

- class-based marking: packets of different classes marked differently
- intra-class marking: conforming portion of flow marked differently than non-conforming one


Classification and Conditioning

- Packet is marked in the Type of Service (TOS) in IPv4, and Traffic Class in IPv6
- 6 bits used for Differentiated Service Code Point (DSCP) and determine PHB that the packet will receive
- 2 bits are currently unused



Classification and Conditioning

- may be desirable to limit traffic injection rate of some class:
 - user declares traffic profile (e.g., rate, burst size)
 - traffic metered, shaped if non-conforming



Forwarding (PHB)

- PHB result in a different observable (measurable) forwarding performance behavior
- PHB does not specify what mechanisms to use to ensure required PHB performance behavior
- Examples:
 - Class A gets $x\%$ of outgoing link bandwidth over time intervals of a specified length
 - Class A packets leave first before packets from class B

Forwarding (PHB)


- PHBs being developed:
- **Expedited Forwarding:** pkt departure rate of a class equals or exceeds specified rate
 - logical link with a minimum guaranteed rate
- **Assured Forwarding:** 4 classes of traffic
 - each guaranteed minimum amount of bandwidth
 - each with three drop preference partitions

Chapter 6 outline

- 6.1 Providing multiple classes of service
- 6.2 Differentiated Services
- 6.3 Providing QoS guarantees
- 6.4 Integrated Services
- 6.5 Resource Reservation Protocol
- 6.6 Next Steps in Signaling
- 6.7 Multiprotocol Label Switching

Principles for QOS Guarantees (more)

- Basic fact of life: can not support traffic demands beyond link capacity

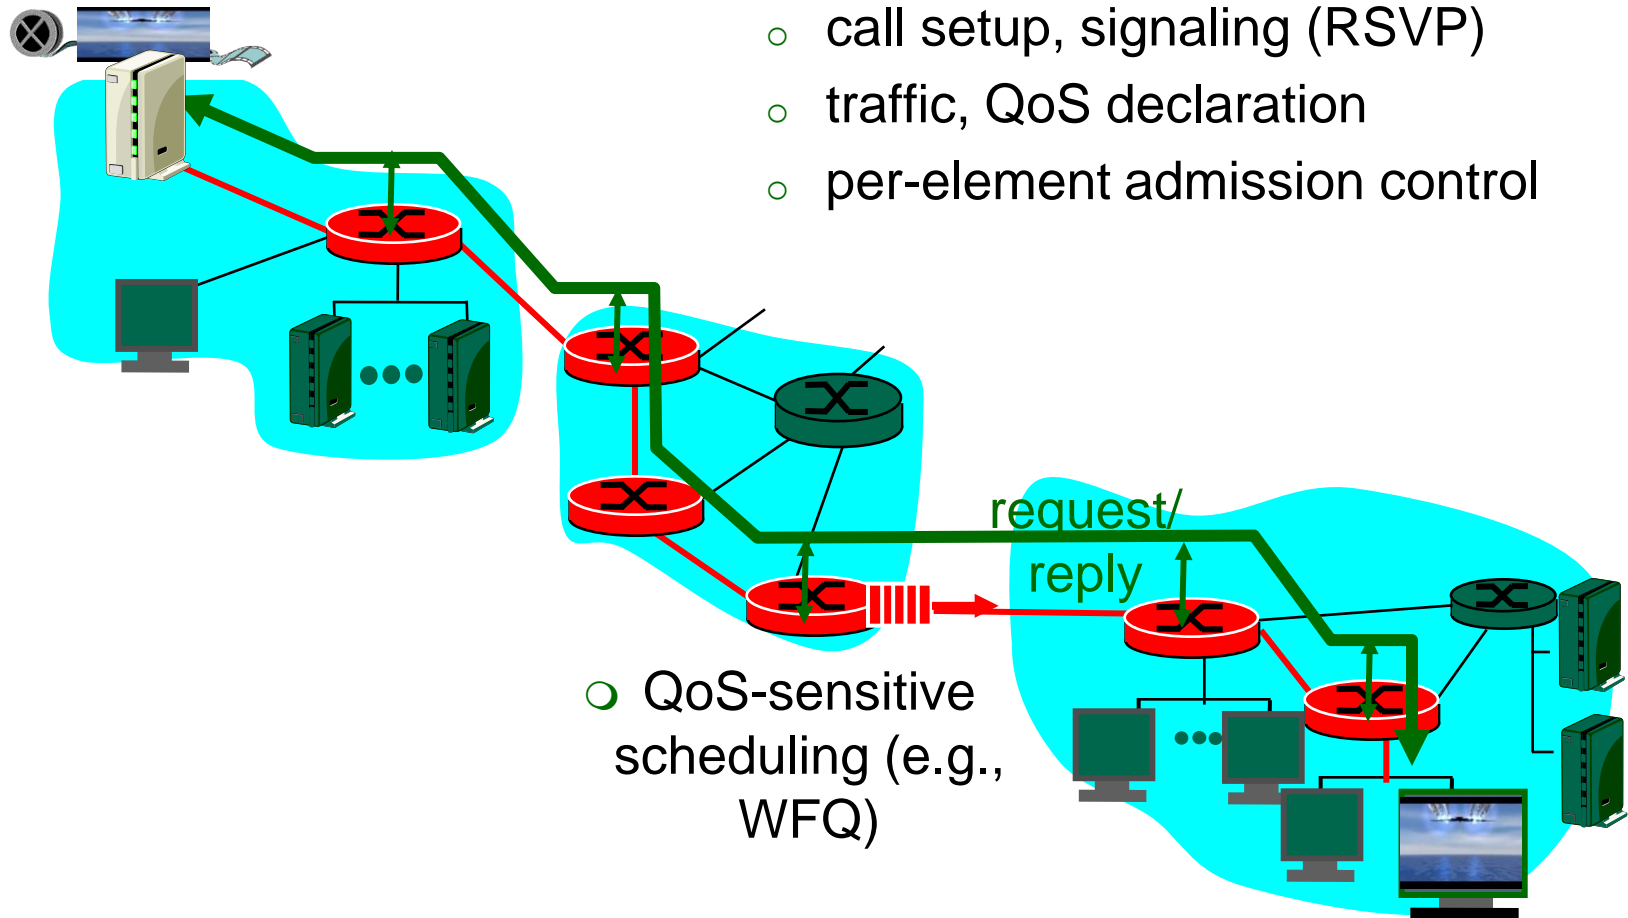


Principle 4

Call Admission: flow declares its needs, network may block call (e.g., busy signal) if it cannot meet needs

QoS guarantee scenario

- Resource reservation
 - call setup, signaling (RSVP)
 - traffic, QoS declaration
 - per-element admission control



- QoS-sensitive scheduling (e.g., WFQ)

Chapter 6 outline

- 6.1 Providing multiple classes of service
- 6.2 Differentiated Services
- 6.3 Providing QoS guarantees
- **6.4 Integrated Services**
- 6.5 Resource Reservation Protocol
- 6.6 Next Steps in Signaling
- 6.7 Multiprotocol Label Switching

IETF Integrated Services

- architecture for providing QoS guarantees in IP networks for individual application sessions
- resource reservation: routers maintain state info (a la VC) of allocated resources, QoS req's
- admit/deny new call setup requests:

Question: can newly arriving flow be admitted with performance guarantees while not violated QoS guarantees made to already admitted flows?

Call Admission

Arriving session must :

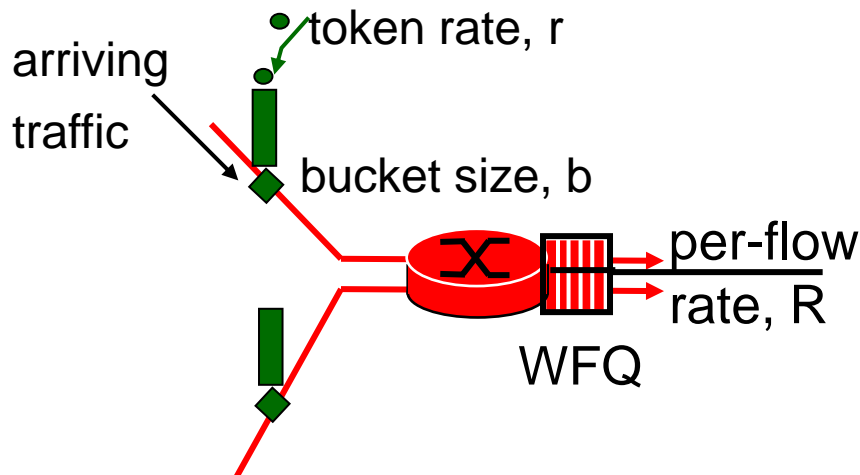
- declare its QOS requirement
 - **R-spec**: defines the QOS being requested
- characterize traffic it will send into network
 - **T-spec**: defines traffic characteristics
- signaling protocol: needed to carry R-spec and T-spec to routers (where reservation is required)
 - **RSVP**

Intserv QoS: Service models

[rfc2211, rfc 2212]

- Guaranteed service:
 - worst case traffic arrival: leaky-bucket-policed source
 - simple (mathematically provable) bound on delay [Parekh 1992, Cruz 1988]

- Controlled load service:
 - "a quality of service closely approximating the QoS that same flow would receive from an unloaded network element."



$$D_{\max} = b/R$$

Chapter 6 outline

- 6.1 Providing multiple classes of service
- 6.2 Differentiated Services
- 6.3 Providing QoS guarantees
- 6.4 Integrated Services
- **6.5 Resource Reservation Protocol**
- 6.6 Next Steps in Signaling
- 6.7 Multiprotocol Label Switching

Signaling in the Internet

connectionless
(stateless) forwarding
by IP routers + best effort
service = no network
signaling protocols
in initial IP design

- **New requirement:** reserve resources along end-to-end path (end system, routers) for QoS for multimedia applications
- **RSVP:** Resource Reservation Protocol [RFC 2205]
 - “ ... allow users to communicate requirements to network in robust and efficient way.” i.e., signaling!
- earlier Internet Signaling protocol: ST-II [RFC 1819]

RSVP Design Goals

1. accommodate **heterogeneous receivers** (different bandwidth along paths)
2. accommodate different applications **with different resource requirements**
3. make **multicast a first class service**, with adaptation to multicast group membership
4. **leverage existing multicast/unicast routing**, with adaptation to changes in underlying unicast, multicast routes
5. **control protocol overhead** to grow (at worst) linear in # receivers
6. **modular design** for heterogeneous underlying technologies

RSVP: does not...

- specify how resources are to be reserved
 - rather: a mechanism for communicating needs
- determine routes packets will take
 - that's the job of routing protocols
 - signaling decoupled from routing
- interact with forwarding of packets
 - separation of control (signaling) and data (forwarding) planes

RSVP: overview of operation

- senders, receiver join a multicast group
 - done outside of RSVP
 - senders need not join group
- sender-to-network signaling
 - *path message*: make sender presence known to routers
 - path teardown: delete sender's path state from routers
- receiver-to-network signaling
 - *reservation message*: reserve resources from sender(s) to receiver
 - reservation teardown: remove receiver reservations
- network-to-end-system signaling
 - path error
 - reservation error

Chapter 6 outline

- 6.1 Providing multiple classes of service
- 6.2 Differentiated Services
- 6.3 Providing QoS guarantees
- 6.4 Integrated Services
- 6.5 Resource Reservation Protocol
- 6.6 Next Steps in Signaling
- 6.7 Multiprotocol Label Switching

Next Steps in Signaling (NSIS)

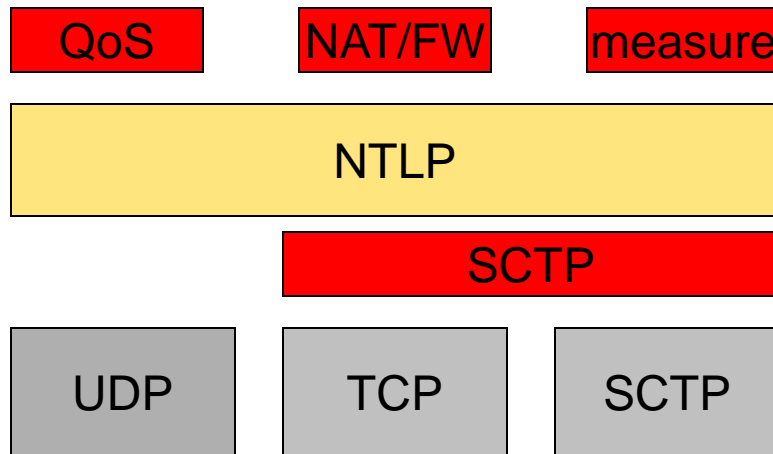
- RSVP not widely used for resource reservation
 - But used for MPLS path setup
 - Design heavily biased by multicast needs
 - Limited ability signaling delivery any size of signaling & over congested situation
 - Marginal and after-the-fact security
 - Limited support for IP mobility

NSIS (cont'd)

- Thus, IETF NSIS working group developing new framework for general state management protocol
 - resource reservation
 - NAT and firewall control
 - traffic and QoS measurement
 - MPLS and lambda path setup
- Split into two components:
 - NSIS Signaling Layer Protocol (NSLP)
 - NSIS Transport Layer Protocol (NTLP)

NSIS (cont'd)

- On-path vs. off-path
 - off-path → bandwidth brokers
- Discovery of next NTLP or NSLP hop
 - use router alert option




Chapter 6 outline

- 6.1 Providing multiple classes of service
- 6.2 Differentiated Services
- 6.3 Providing QoS guarantees
- 6.4 Integrated Services
- 6.5 Resource Reservation Protocol
- 6.6 Next Steps in Signaling
- 6.7 Multiprotocol Label Switching

Multiprotocol Label Switching (MPLS)

- *Original motivation:* improve forwarding speed of IP routers.
- MPLS introduces a **fixed-length label** between layer-2 (i.e. PPP or Ethernet) header and layer-3 (i.e. IP) header.

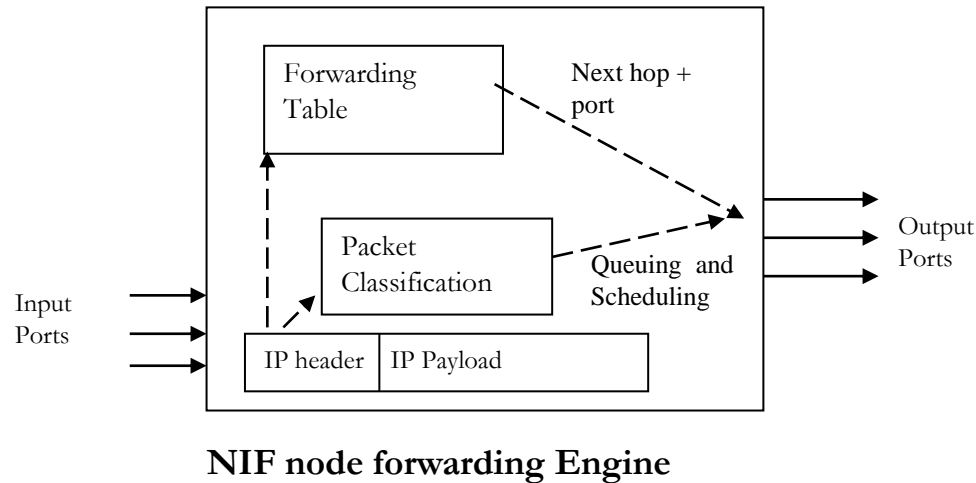


Introduction to MPLS


- **Conventional IP (Layer-3) forwarding:**
Each router analyzes the incoming packet's header and independently chooses a next hop. Routing algorithm and adequate speed are prerequisite.
- **MPLS (Layer-2.5) forwarding:**
All forwarding is driven by the labels, no header analysis needed. Once a packet enters a network, it's assigned a label. Each router forwards packets according to their labels.

IP Forwarding

- Longest-prefix match based on packet's destination IP address




IP Forwarding (cont'd)



MPLS forwarding

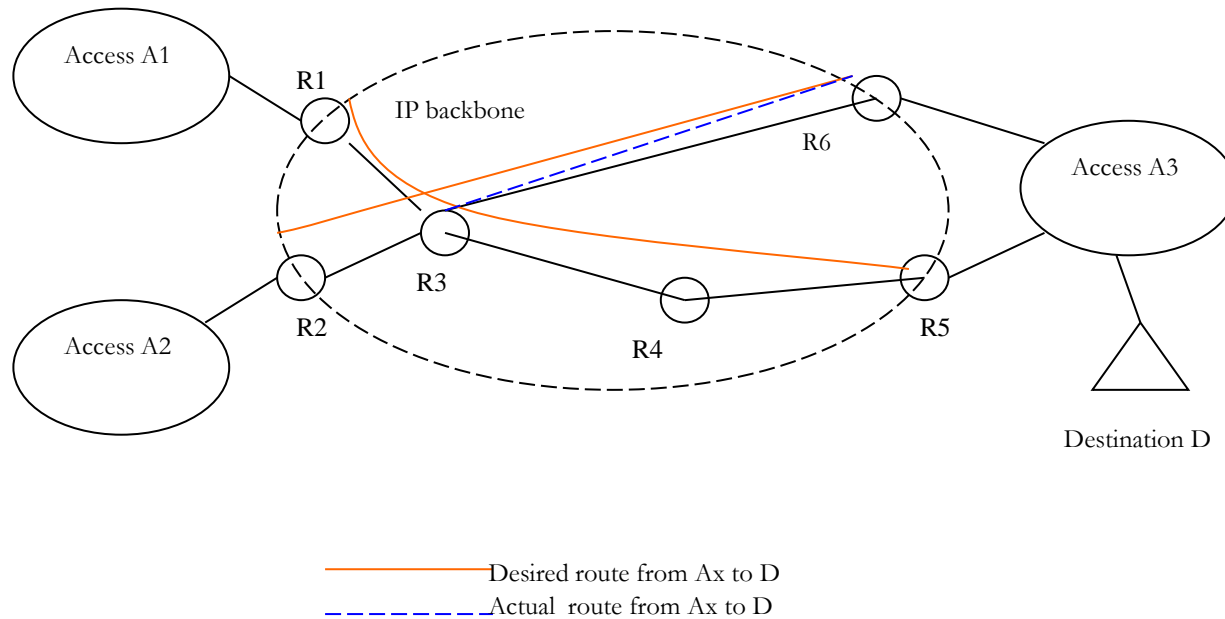
In I/F	In Lab	Address Prefix	Out I/F	Out Lab
4	x	171.68.10	3	5
...

In I/F	In Lab	Address Prefix	Out I/F	Out Lab
0	5	171.68.10	1	3
...



FEC:
Forwarding
Equivalence
Class

MPLS Network



Traffic Engineering to override shortest path route

MPLS Advantages

- Router can use any information in determining label assignment, not limited to packet header.
- How to distribute labels may become more and more complicated, without any impact on the routers that merely forward labeled packets.
- A label can be used to represent a pre-chosen route so that the identity of explicit route need not be carried with the packet.
- Multiprotocol: its techniques are applicable to ANY network layer protocol.

MPLS Applications

- Extensive traffic engineering and management.
- Out-of-band control.
- Quality of Service.
- Traffic separation: Virtual Private Networks.

Summary

- Principles
 - classify multimedia applications
 - identify network services applications need
 - making the best of best effort service
- Protocols and Architectures
 - specific protocols for best-effort
 - mechanisms for providing QoS
 - architectures for QoS
 - multiple classes of service
 - QoS guarantees, admission control

Thank you

Any questions?