Providing Quality of Service in the Internet

Based on Slides from Ross and Kurose

Providing Quality of Service

- The Internet is based on best effort
- But, we want to provide guarantees for QoS
  - QoS sensitive applications
  - Paying customers ($$)
- QoS guarantees conflicts with
  - Scalability
  - Efficient resource use
- Hwo can we work around it?

The QoS Problem Illustrated

- Simple model for sharing and congestion studies:
  - Competing for a scarce resource
- The aspects of the problem
  - **Who** is allowed? Admission Control
  - **Which** packets are which? Packet classification
  - **How** do I share? Scheduling
  - Use resources efficiently

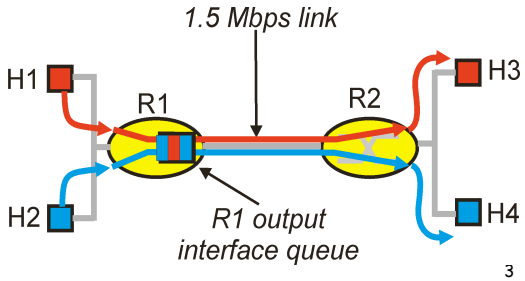

Principles for QOS Guarantees

- Consider a phone application at 1Mbps and an FTP application sharing a 1.5 Mbps link.
  - bursts of FTP can congest the router and cause audio packets to be dropped.
  - want to give priority to audio over FTP
- **PRINCIPLE 1: Marking of packets is needed for router to distinguish between different classes; and new router policy to treat packets accordingly**

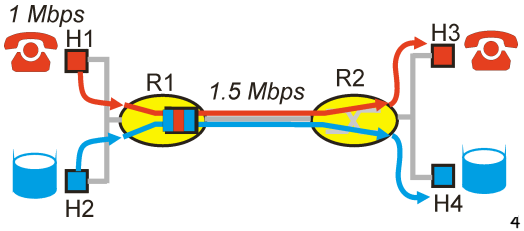

Principles for QOS Guarantees (more)

- Applications misbehave (audio sends packets at a rate higher than 1Mbps assumed above);
- **PRINCIPLE 2: provide protection (isolation) for one class from other classes**
- Require Policing Mechanisms to ensure sources adhere to bandwidth requirements; Marking and Policing need to be done at the edges:

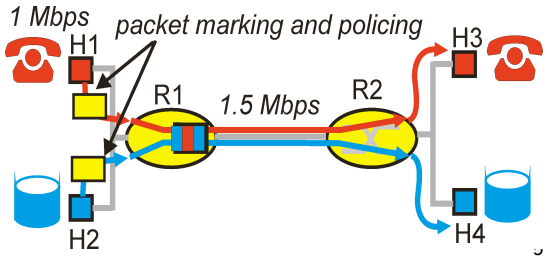

Principles for QOS Guarantees (more)

- Alternative to Marking and Policing: allocate a set portion of bandwidth to each application flow; can lead to inefficient use of bandwidth if one of the flows does not use its allocation
- **PRINCIPLE 3: While providing isolation, it is desirable to use resources as efficiently as possible**

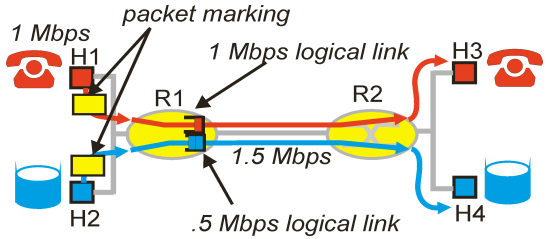

Principles for QOS Guarantees (more)

- Cannot support traffic beyond link capacity
- PRINCIPLE 4: Need a Call Admission Process; application flow declares its needs, network may block call if it cannot satisfy the needs

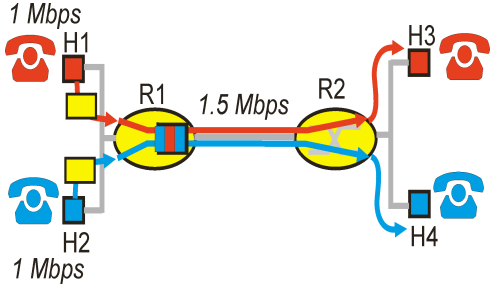

The QoS "Architecture"

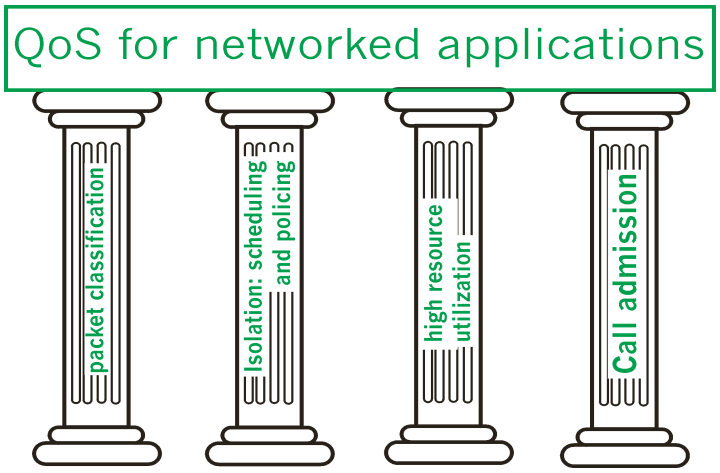

Overview of Current QoS Trends

- RSVP: Resource reservation Protocol
  - Receivers propagate their QoS needs along the path
  - Not scalable, more meaningful in multicasting
- DiffServ: Differentiated Services
  - "Colour" packets on entrance, treat different colours
  - Per Hop Behavior: packets carry with the "routing state" namely how the expect to be treated.
- IntServ: Integrated Services
  - Routers maintain state per flow (!!!)

Scheduling And Policing Mechanisms

- Scheduling: choosing the next packet for transmission on a link can be done following a number of policies;
- FIFO: in order of arrival to the queue; packets that arrive to a full buffer are either discarded, or a discard policy is used to determine which packet to discard among the arrival and those already queued

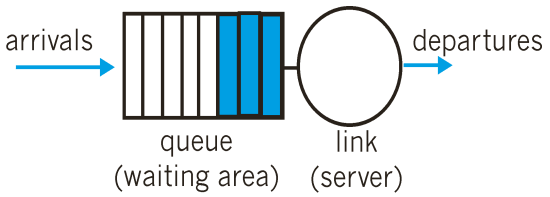

Scheduling Policies

- Priority Queuing: classes have different priorities; class may depend on explicit marking or other header info, eg IP source or destination, TCP Port numbers, etc.
- Transmit a packet from the highest priority class with a non-empty queue
- Preemptive and non-preemptive versions
- Issue: low classes can starve

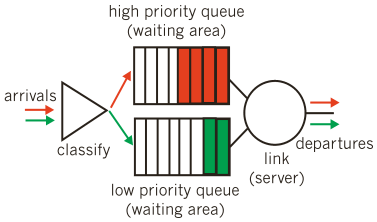

2 is forwarded after 3!

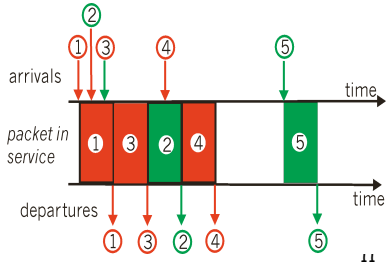

Scheduling Policies (more)

- Round Robin: scan class queues serving one from each class that has a non-empty queue
- Provides better QoS to higher class only if higher class has fewer packets

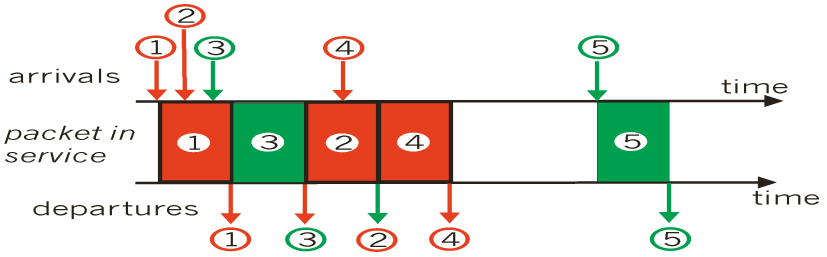

Scheduling Policies (more)

- Weighted Fair Queuing: is a generalized Round Robin in which an attempt is made to provide a class with a differentiated amount of service over a given period of time

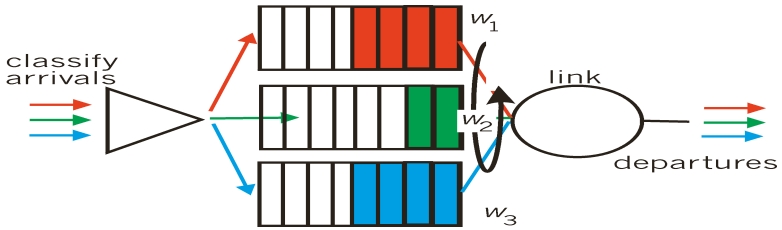

Policing Mechanisms

- Three criteria:
  - (Long term) **Average Rate** (100 packets per sec or 6000 packets per min??), crucial aspect is the interval length
  - **Peak Rate**: e.g., 6000 p p minute Avg and 1500 p p sec Peak
  - (Max.) **Burst Size**: Max. number of packets sent consecutively, ie over a short period of time

Policing Mechanisms

- Token Bucket mechanism, provides a means for limiting input to specified Burst Size and Average Rate.

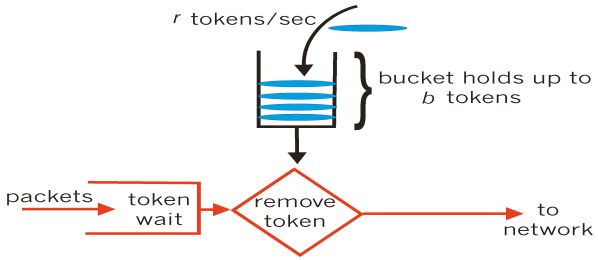

Policing Mechanisms (more)

- Bucket can hold b tokens; token are generated at a rate of *r token/sec* unless bucket is full of tokens.
- Over an interval of length *t*, the number of packets that are admitted is less than or equal to $(r\ t + b)$.
- Token bucket and WFQ can be combined to provide upper bound on delay.

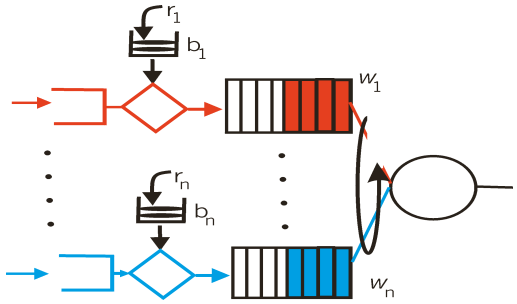

Handling Bursty Sources

- Token bucket is good for "well behaved" sources
  - Approximately near the average sending rate
  - Few big bursts (that may get clipped)
- What about an application that has one big burst?
  - No credit for idle period
  - Slaughtered during its peak

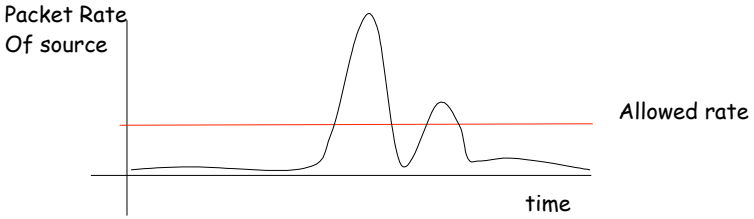

Two Buckets Instead Of One

- Published in Global Internet (Globecom) 2002
- Key: if you are idle, you get credit for big burst
- Provide a second buffer to collect credit during idle times (call it burst bucket)
- During peak rate burst bucket provides extra tokens to token bucket at some rate
- Problem: what if all sources are quite and then burst altogether?
- Parameter finetuning:
  - How big the burst bucket should be?

Integrated Services

- An architecture for providing QOS guarantees in IP networks for individual application sessions
- relies on resource reservation, and routers need to maintain state info (Virtual Circuit??), maintaining records of allocated resources and responding to new Call setup requests on that basis

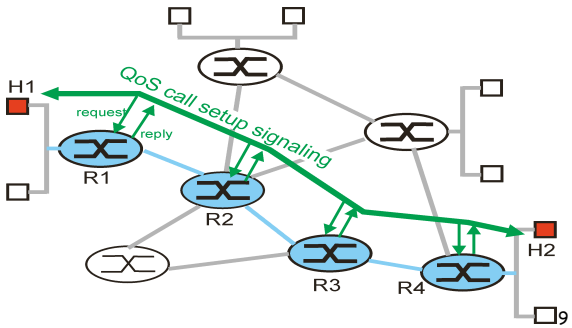

Call Admission

- Session must first declare its QOS requirement and characterize the traffic it will send through the network
- **R-spec**: defines the QOS being requested
- **T-spec**: defines the traffic characteristics
- A signaling protocol is needed to carry the R-spec and T-spec to the routers where reservation is required; RSVP is a leading candidate for such signaling protocol

Call Admission

- Call Admission: routers will admit calls based on their R-spec and T-spec and base on the current resource allocated at the routers to other calls.

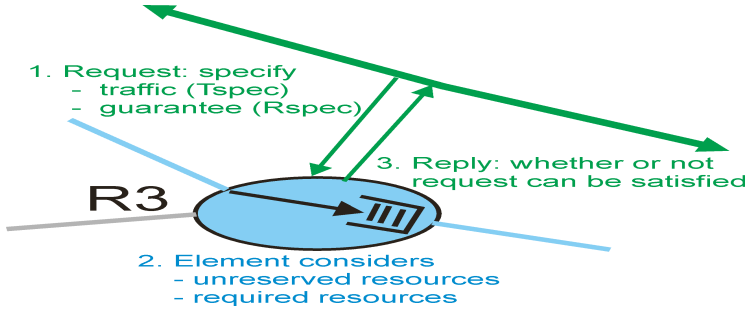

Integrated Services: Classes

- **Guaranteed QOS**: this class is provided with firm bounds on queuing delay at a router; envisioned for hard real-time applications that are highly sensitive to end-to-end delay expectation and variance
- **Controlled Load**: this class is provided a QOS closely approximating that provided by an unloaded router; envisioned for today's IP network real-time applications which perform well in an unloaded network

Differentiated Services

- Intended to address the following difficulties with Intserv and RSVP;
- **Scalability**: maintaining states by routers in high speed networks is difficult sue to the very large number of flows
- **Flexible Service Models**: Intserv has only two classes, want to provide more qualitative service classes; want to provide 'relative' service distinction (Platinum, Gold, Silver, …)
- **Simpler signaling**: (than RSVP) many applications and users may only w ant to specify a more qualitative notion of service

Differentiated Services

- Approach:
  - Only simple functions in the core, and relatively complex functions at edge routers (or hosts)
  - Do not define service classes, instead provides functional components with which service classes can be built

Edge Functions

- At DS-capable host or first DS-capable router
- **Classification**: edge node marks packets according to classification rules to be specified (manually by admin, or by some protocol)
- **Traffic Conditioning**: edge node may delay and forward

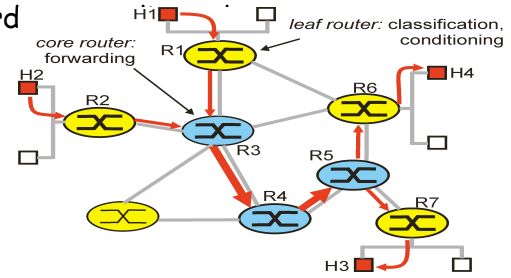

Core Functions

- **Forwarding**: according to "Per-Hop-Behavior" or PHB specified for the particular packet class; such PHB is strictly based on class marking (no other header fields can be used to influence PHB)

- BIG ADVANTAGE:

  No state info to be maintained by routers!

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

- Limit traffic injection rate of each class;
- User declares traffic profile (eg, rate and burst size); packets are dropped if non-conforming
- Traffic shaping takes place at incoming point of a network

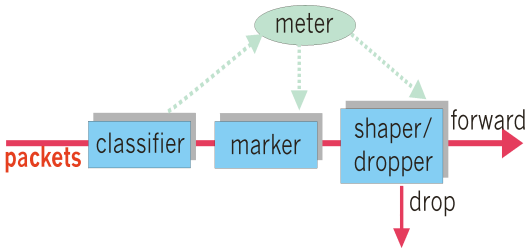

Forwarding (PHB)

- 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 under consideration:
  - **Expedited Forwarding**: departure rate of packets from a class equals or exceeds a specified rate (logical link with a minimum guaranteed rate)
  - **Assured Forwarding**: 4 classes, each guaranteed a minimum amount of bandwidth and buffering; each with three drop preference partitions

Differentiated Services Issues

- AF and EF are not even in a standard track yet… research ongoing
- We need to determine the impact of crossing multiple ASs and routers that are not DS-capable

QoS: Summary

- Internet was not designed with QoS in mind
- Adding QoS over best-effort is not easy
- QoS also requires access limitations
  - Admission control
  - Traffic shaping
- No final solution exists yet

Real-Time Protocol (RTP)

- Provides standard packet format for real-time application
- Typically runs over UDP
- Specifies header fields below
- **Payload Type**: 7 bits, providing 128 possible different types of encoding; eg PCM, MPEG2 video, etc.
- **Sequence Number**: 16 bits; used to detect packet loss

RTP Header

Real-Time Protocol (RTP)

- **Timestamp**: 32 bytes; gives the sampling instant of the first audio/video byte in the packet; used to remove jitter introduced by the network
- **Synchronization Source identifier (SSRC)**: 32 bits; an id for the source of a stream; assigned randomly by the source

RTP Header

RTP Control Protocol (RTCP)

- Protocol specifies report packets exchanged between sources and destinations of multimedia information
- Three reports are defined: Receiver reception, Sender, and Source description
- Reports contain statistics such as the number of packets sent, number of packets lost, inter-arrival jitter
- Used to modify sender transmission rates and for diagnostics purposes

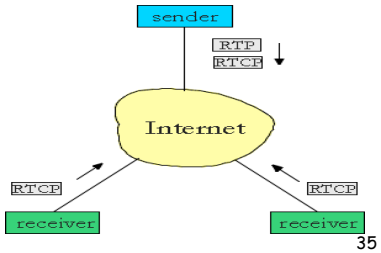

RTCP Bandwidth Scaling

- If each receiver sends RTCP packets to all other receivers, the traffic load resulting can be large
- RTCP adjusts the interval between reports based on the number of participating receivers
- Typically, limit the RTCP bandwidth to 5% of the session bandwidth, divided between the sender reports (25%) and the receivers reports (75%)