EPL606

Quality of Service and Traffic Classification

Multimedia, Quality of Service: What is it?

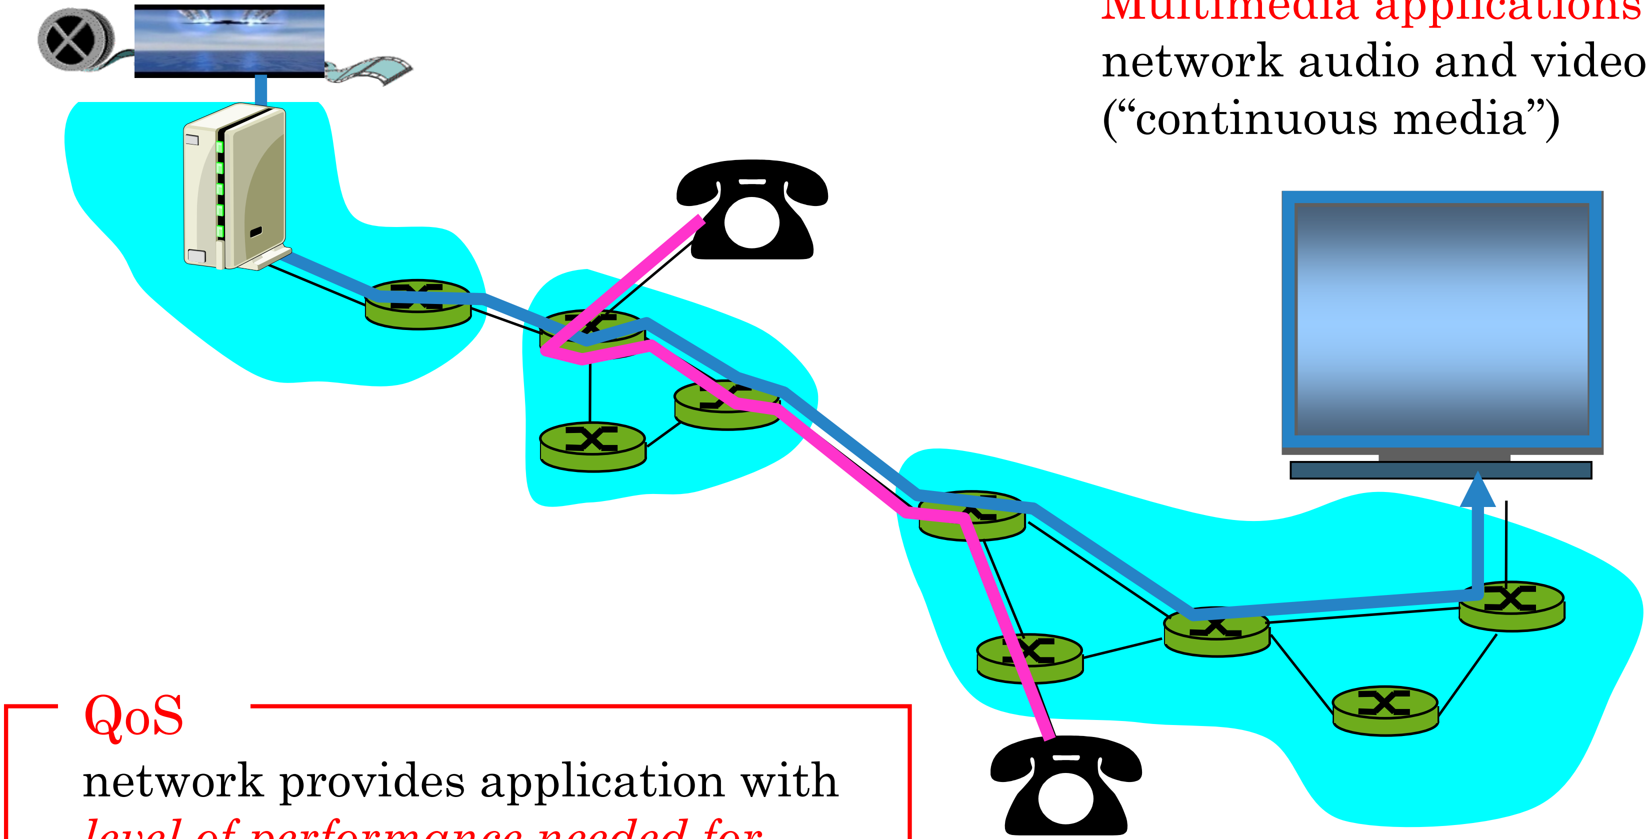

Multimedia applications: network audio and video ("continuous media")

QoS

network provides application with *level of performance needed for application to function.*

Goals

Principles

- Classify multimedia applications
- Identify the network services the apps need
- Making the best of best effort service
- Mechanisms for providing QoS

Protocols and Architectures

- Specific protocols for best-effort
- Architectures for QoS

Multimedia Applications and Requirements

- Multimedia applications and services:
  - combine (simultaneously) information (data) in different forms (e.g. voice, video, images, text, animation)
  - distributed in nature, and involve networking
- Multimedia stimulates the senses:
  - sight (pictures, graphs, text, motion)
  - sound (embedded narration)
  - motion (animation)

Multimedia Applications and Requirements

- Multimedia Applications involve many types of information:
  - Voice
  - Data
    - file transfers
    - distributed computing
  - Audio and Video
    - Stored, e.g. Entertainment, Training
    - Live Interactive, e.g. Conferencing
    - Non-interactive, e.g. TV Broadcasting
  - Still Images
    - Interactive, e.g. Browsing
    - Non-interactive, e.g. Archiving

Multimedia Networking Applications: Quality of Service

- Due to the differing requirements of multiservice and multimedia traffic concept of
  - Quality of Service (QoS) for individual users has become necessary (in order to guarantee the quality of service required by a user and also increase network utilisation)
    - user required to identify and define
      - bandwidth demand (full characterisation of the traffic source behaviour)
      - loss tolerance
      - delay tolerance
      - delay variation tolerance

Multimedia Networking Applications: Requirements

- Either **Bursty, Variable Bit Rate (VBR)** or **Constant Bit Rate (CBR)**
- Typically **delay sensitive**
  - end-to-end delay
  - delay jitter
- But **loss tolerant**: infrequent losses cause minor glitches
- Antithesis of data, which are loss intolerant but delay tolerant.

Multimedia Networking Applications: Data Rate Classification

Rate Type	Descriptions
Stream	Predictable delivery at a relatively constant bit rate (CBR). Quantifiable upper bound.
Burst	Unpredictable delivery, variable bit rate (VBR). FTP applications. (use all available bandwidth)

Multimedia Networking Applications: Delay Sensitivity Classification

Delay Tolerance	Delivery Type	Description
High	Asynchronous	No constraints on delivery time (elastic)
	Synchronous	Time sensitive data, but flexible.
	Interactive	Delays may be noticeable to users or applications, but don not adversely affect usability or functionality.
	Isochronous	Time sensitive data to an extend that adversely affects usability.
Low	Mission-Critical	Data delivery delays disable functionality.

Multimedia Networking Applications: Traffic characteristrics

	Data Traffic	Voice	Multimedia Traffic
Data rate	Low	Very low	High
Traffic pattern	Bursty	Stream-oriented	Stream-oriented and/or Highly Bursty
Correctness required	No Loss	Loss can be tolerated	Some loss can be tolerated
Latency required	None	Small (e.g.~ 30 msec)	May be small (e.g. 20msec)
Mode of connection	Point-to-point	Point-to-point	Point-to-point or Multipoint
Temporal relationships	None	Synchronised transmissions	Synchronised transmissions
Type of service	Single traffic	Single traffic	Multiple traffic

Broadband Services: Source characterisation

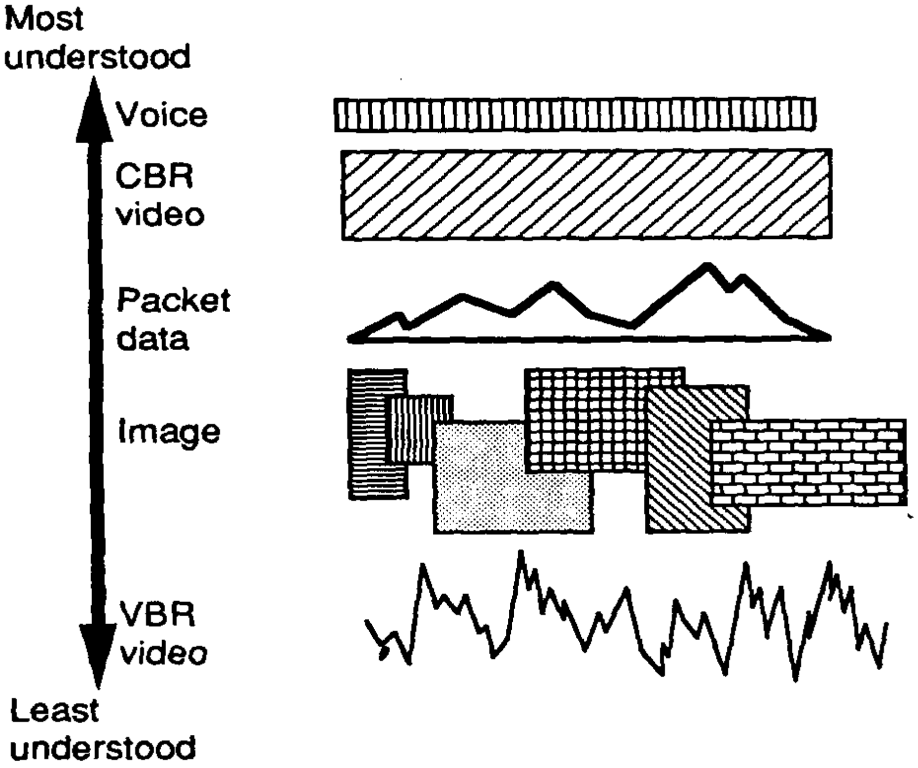

Figure 3.1 Current state of information on source characterization in ATM networks.

Observe diverse nature of broadband sources and difficulty in characterization/ classification

Q. Why do we need to characterize / classify sources?

Examples of source bandwidth requirements

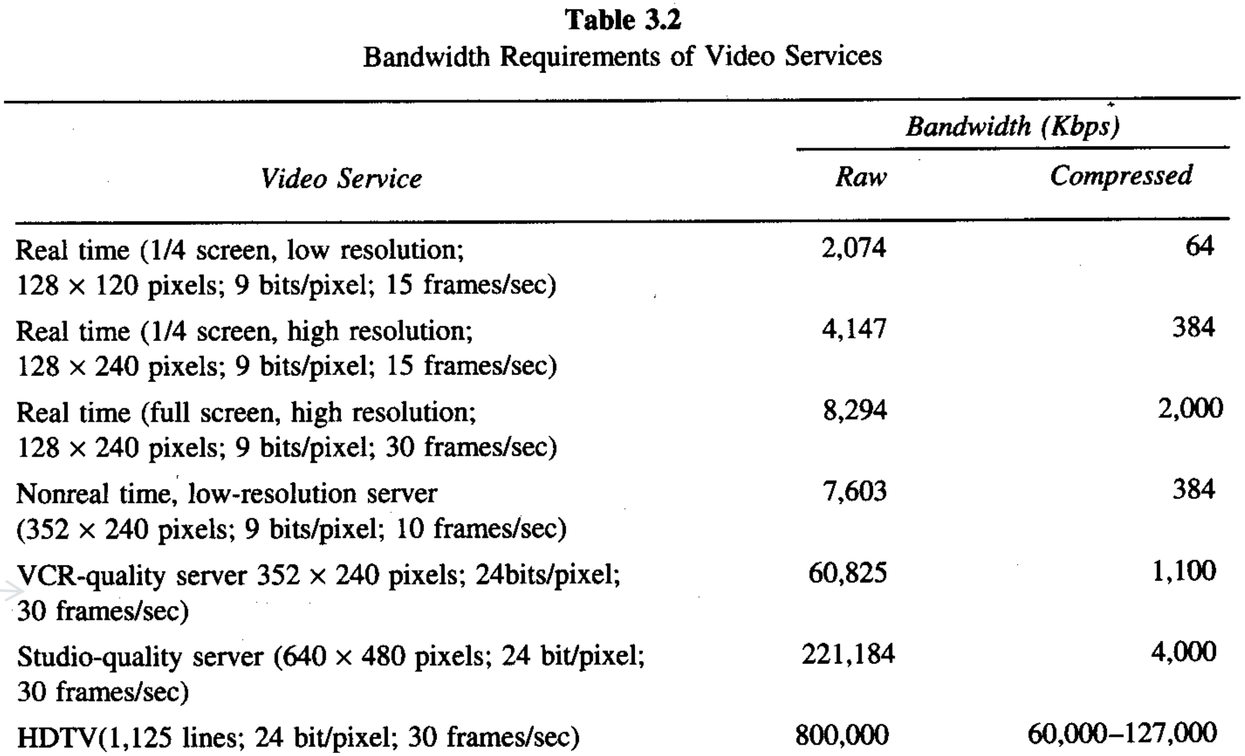

Table 3.2
Bandwidth Requirements of Video Services

Video Service	*Bandwidth (Kbps)*	
	Raw	*Compressed*
Real time (1/4 screen, low resolution; 128 × 120 pixels; 9 bits/pixel; 15 frames/sec)	2,074	64
Real time (1/4 screen, high resolution; 128 × 240 pixels; 9 bits/pixel; 15 frames/sec)	4,147	384
Real time (full screen, high resolution; 128 × 240 pixels; 9 bits/pixel; 30 frames/sec)	8,294	2,000
Nonreal time, low-resolution server (352 × 240 pixels; 9 bits/pixel; 10 frames/sec)	7,603	384
VCR-quality server 352 × 240 pixels; 24bits/pixel; 30 frames/sec)	60,825	1,100
Studio-quality server (640 × 480 pixels; 24 bit/pixel; 30 frames/sec)	221,184	4,000
HDTV(1,125 lines; 24 bit/pixel; 30 frames/sec)	800,000	60,000–127,000

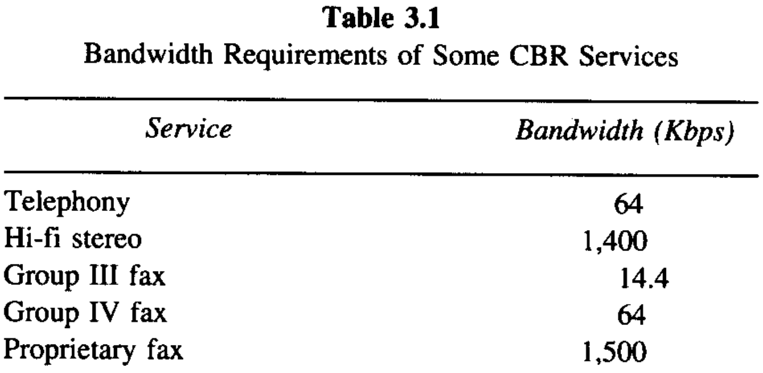

Table 3.1
Bandwidth Requirements of Some CBR Services

Service	*Bandwidth (Kbps)*
Telephony	64
Hi-fi stereo	1,400
Group III fax	14.4
Group IV fax	64
Proprietary fax	1,500

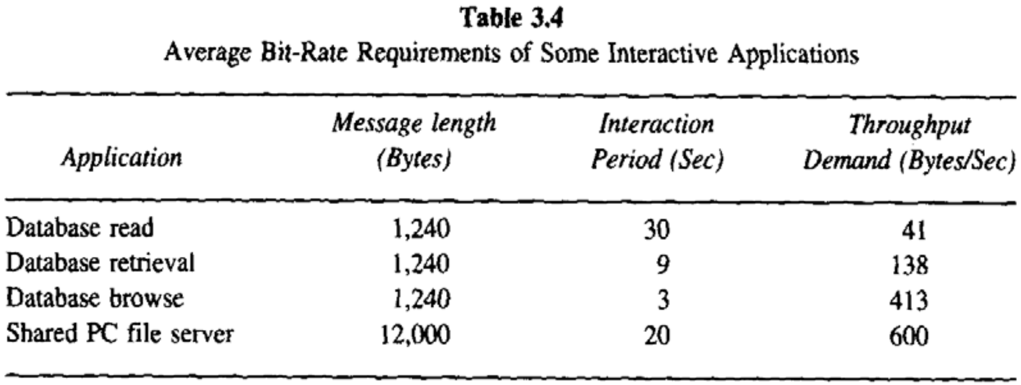

Table 3.4
Average Bit-Rate Requirements of Some Interactive Applications

Application	*Message length (Bytes)*	*Interaction Period (Sec)*	*Throughput Demand (Bytes/Sec)*
Database read	1,240	30	41
Database retrieval	1,240	9	138
Database browse	1,240	3	413
Shared PC file server	12,000	20	600

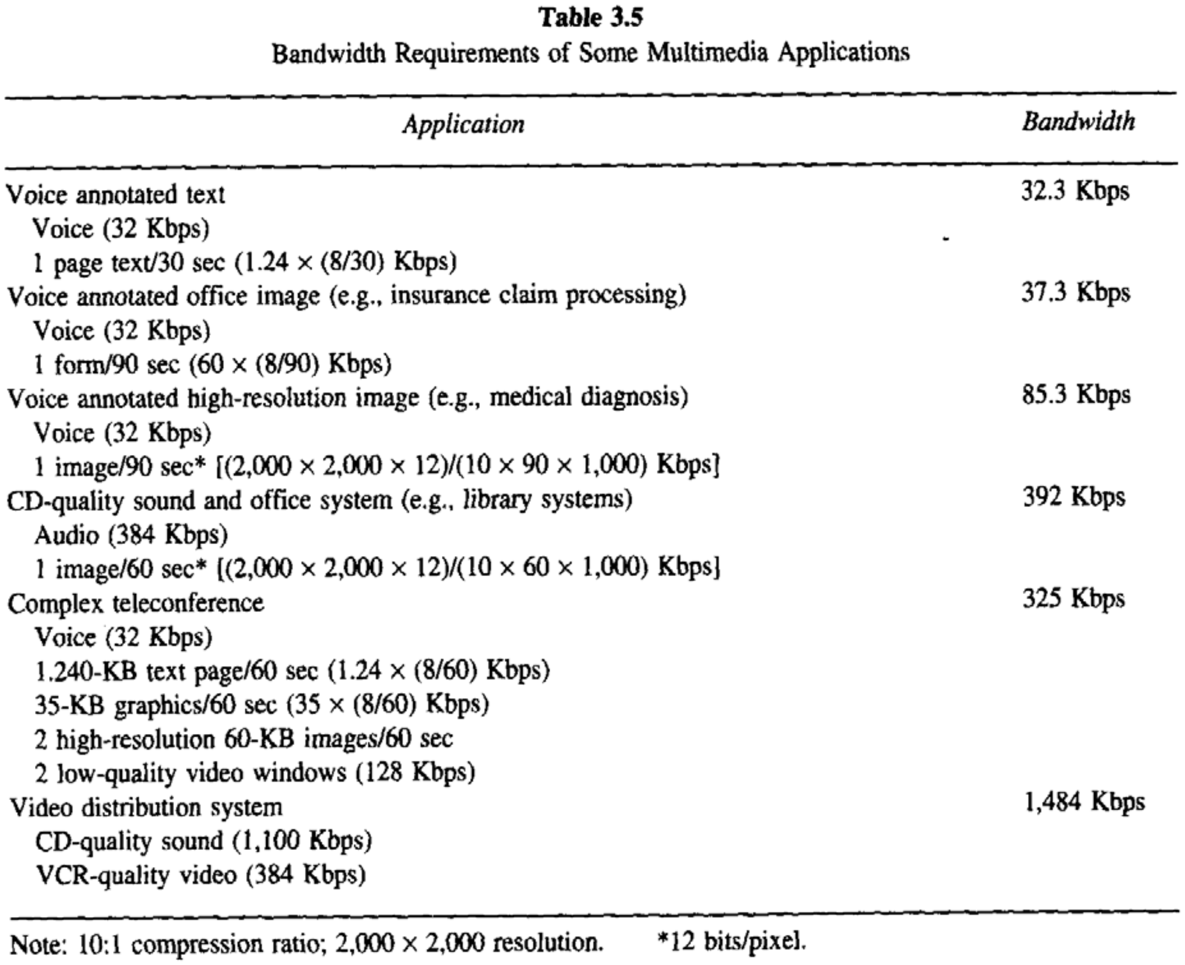

Table 3.5

Bandwidth Requirements of Some Multimedia Applications

Application	*Bandwidth*
Voice annotated text	32.3 Kbps
Voice (32 Kbps)	
1 page text/30 sec (1.24 × (8/30) Kbps)	
Voice annotated office image (e.g., insurance claim processing)	37.3 Kbps
Voice (32 Kbps)	
1 form/90 sec (60 × (8/90) Kbps)	
Voice annotated high-resolution image (e.g., medical diagnosis)	85.3 Kbps
Voice (32 Kbps)	
1 image/90 sec* [(2,000 × 2,000 × 12)/(10 × 90 × 1,000) Kbps]	
CD-quality sound and office system (e.g., library systems)	392 Kbps
Audio (384 Kbps)	
1 image/60 sec* [(2,000 × 2,000 × 12)/(10 × 60 × 1,000) Kbps]	
Complex teleconference	325 Kbps
Voice (32 Kbps)	
1.240-KB text page/60 sec (1.24 × (8/60) Kbps)	
35-KB graphics/60 sec (35 × (8/60) Kbps)	
2 high-resolution 60-KB images/60 sec	
2 low-quality video windows (128 Kbps)	
Video distribution system	1,484 Kbps
CD-quality sound (1,100 Kbps)	
VCR-quality video (384 Kbps)	

Note: 10:1 compression ratio; 2,000 × 2,000 resolution. *12 bits/pixel.

observe diverse nature of source behaviour (in terms of required bandwidth) for even a similar activity-- that of a conference between two offices

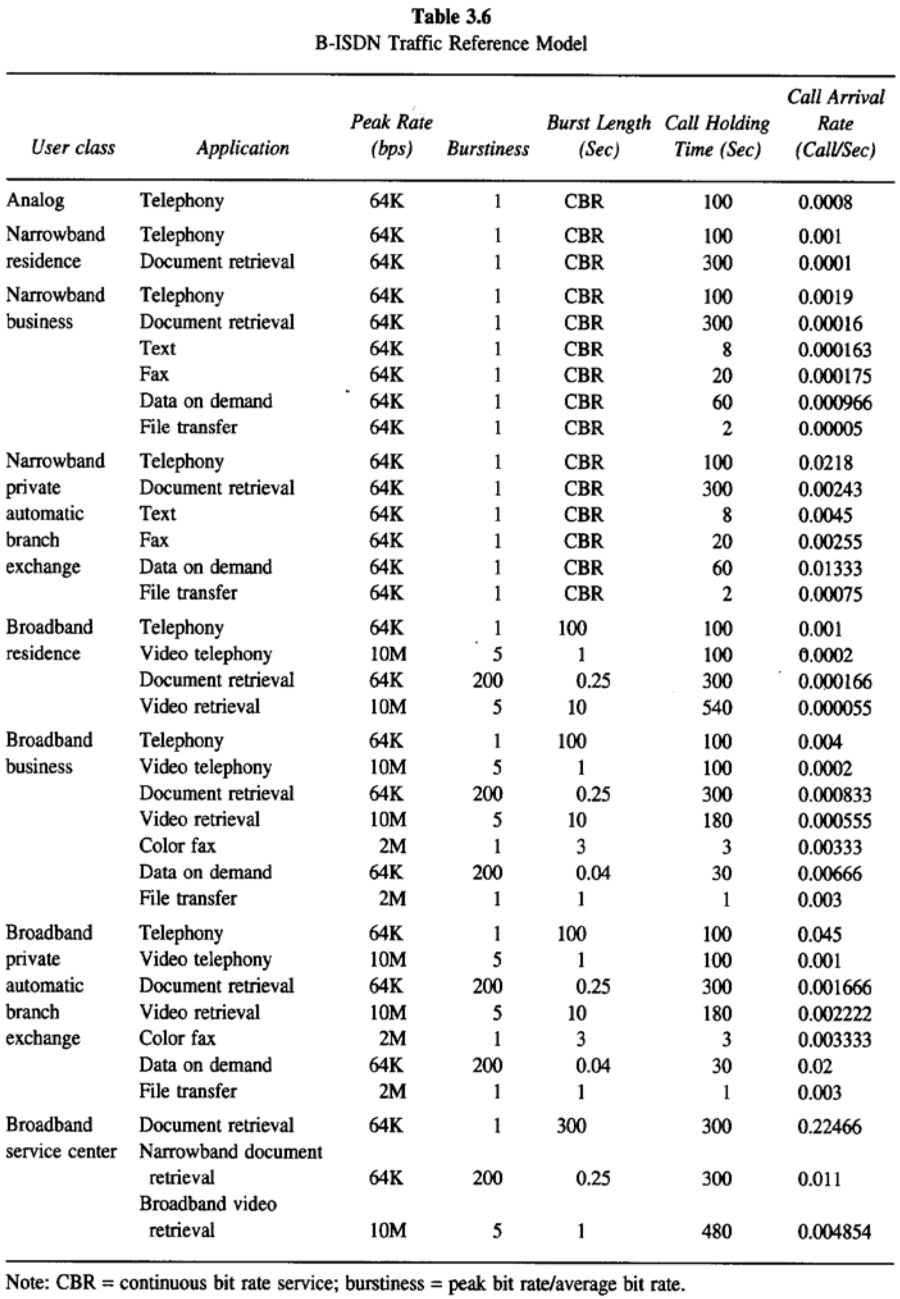

Table 3.6
B-ISDN Traffic Reference Model

User class	*Application*	*Peak Rate (bps)*	*Burstiness*	*Burst Length (Sec)*	*Call Holding Time (Sec)*	*Call Arrival Rate (Call/Sec)*
Analog	Telephony	64K	1	CBR	100	0.0008
Narrowband residence	Telephony	64K	1	CBR	100	0.001
	Document retrieval	64K	1	CBR	300	0.0001
Narrowband business	Telephony	64K	1	CBR	100	0.0019
	Document retrieval	64K	1	CBR	300	0.00016
	Text	64K	1	CBR	8	0.000163
	Fax	64K	1	CBR	20	0.000175
	Data on demand	64K	1	CBR	60	0.000966
	File transfer	64K	1	CBR	2	0.00005
Narrowband private automatic branch exchange	Telephony	64K	1	CBR	100	0.0218
	Document retrieval	64K	1	CBR	300	0.00243
	Text	64K	1	CBR	8	0.0045
	Fax	64K	1	CBR	20	0.00255
	Data on demand	64K	1	CBR	60	0.01333
	File transfer	64K	1	CBR	2	0.00075
Broadband residence	Telephony	64K	1	100	100	0.001
	Video telephony	10M	5	1	100	0.0002
	Document retrieval	64K	200	0.25	300	0.000166
	Video retrieval	10M	5	10	540	0.000055
Broadband business	Telephony	64K	1	100	100	0.004
	Video telephony	10M	5	1	100	0.0002
	Document retrieval	64K	200	0.25	300	0.000833
	Video retrieval	10M	5	10	180	0.000555
	Color fax	2M	1	3	3	0.00333
	Data on demand	64K	200	0.04	30	0.00666
	File transfer	2M	1	1	1	0.003
Broadband private automatic branch exchange	Telephony	64K	1	100	100	0.045
	Video telephony	10M	5	1	100	0.001
	Document retrieval	64K	200	0.25	300	0.001666
	Video retrieval	10M	5	10	180	0.002222
	Color fax	2M	1	3	3	0.003333
	Data on demand	64K	200	0.04	30	0.02
	File transfer	2M	1	1	1	0.003
Broadband service center	Document retrieval	64K	1	300	300	0.22466
	Narrowband document retrieval	64K	200	0.25	300	0.011
	Broadband video retrieval	10M	5	1	480	0.004854

Note: CBR = continuous bit rate service; burstiness = peak bit rate/average bit rate.

← These are some (technical) parameters that aid in the characterisation / classification of traffic source bahaviour for demanded bandwidth

Can you list some of the cases that you think these could be useful? network design, dimensioning, management and control, etc...

Quality of Service (QoS)

- What is QoS?
  - Specifies a set of performance characteristics
  - It is used to manage the network resources more efficiently.
  - QoS doesn’t create bandwidth
- Two types of QoS:
  - Resource reservation (integrated services)
  - Prioritisation (differentiated services)
- These QoS protocols complement each other.

QoS Characteristics

- Throughput (bandwidth)
- Delay (Latency)
- Delay variation
- Packet loss rate
- Service availability

QoS Characteristics

- Bandwidth
  - Bandwidth is the ideal capacity that the network can operate. The networks never work on ideal maximum capacity since there are negative factors that cause deterioration of the quality of the network. Such as factors can be transmission delay, noise and etc.
- Packet Loss
  - Packet loss takes in place when we are experiencing congestion on our network. In the event of the congestion the network can discard this packet, which are defined by this parameter.
- Service Availability
  - Availability is the reliability of the user's connection to the Internet service. To be able to do this we use Service Level Agreement (SLA).

QoS Characteristics

- Latency
  - Latency or a propagation time is referred to the time it takes to send a message from the sender until to the time the receiver receives.
- Router Latency
  - It’s the time it takes to the router to retransmit the packet that it had received from the time it had arrived to the router.
- Jitter (Delay variation)
  - Refers to the variation in time duration in all packets in stream taking the same route. For instance, when sending a video or audio stream over the network and the packets don’t arrive in the order that was sent on a timely basis. This creates a distortion of the signal, which is very harmful to multimedia.

Examples of QoS measures

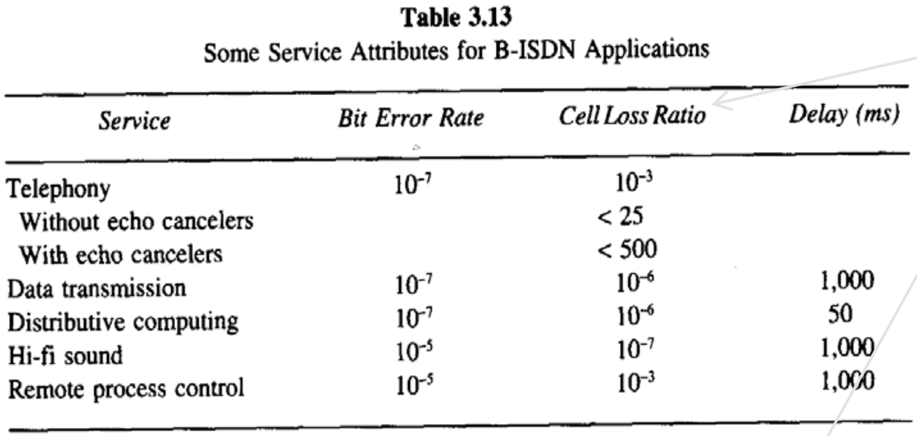

Table 3.13
Some Service Attributes for B-ISDN Applications

Service	*Bit Error Rate*	*Cell Loss Ratio*	*Delay (ms)*
Telephony	10^{-7}	10^{-3}	
Without echo cancelers		< 25	
With echo cancelers		< 500	
Data transmission	10^{-7}	10^{-6}	1,000
Distributive computing	10^{-7}	10^{-6}	50
Hi-fi sound	10^{-5}	10^{-7}	1,000
Remote process control	10^{-5}	10^{-3}	1,000

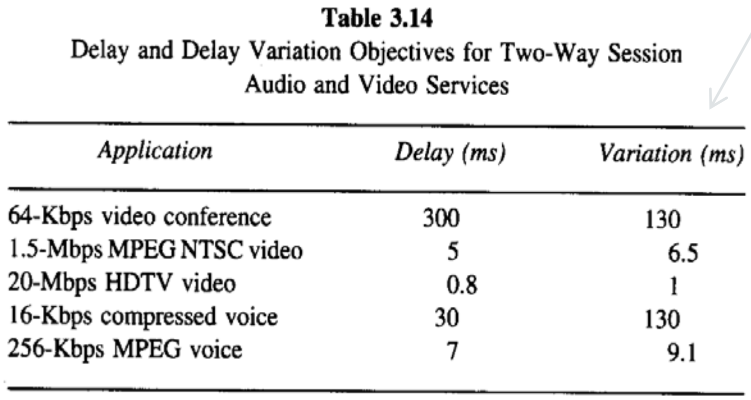

Table 3.14
Delay and Delay Variation Objectives for Two-Way Session Audio and Video Services

Application	*Delay (ms)*	*Variation (ms)*
64-Kbps video conference	300	130
1.5-Mbps MPEG NTSC video	5	6.5
20-Mbps HDTV video	0.8	1
16-Kbps compressed voice	30	130
256-Kbps MPEG voice	7	9.1

Any loss, delay, or delay variation better than these figures is acceptable to the user

Improving QoS in IP networks

Quality of Service

- Approaches to QoS Support
  - *fine-grained approaches, which provide QoS to individual applications or flows*
  - *coarse-grained approaches, which provide QoS to large classes of data or aggregated* traffic

  - In the first category we find “Integrated Services,” a QoS architecture developed in the IETF and often associated with RSVP (Resource Reservation Protocol).
  - In the second category lies “Differentiated Services,” which is probably the most widely deployed QoS mechanism.

QoS Policies

- QoS Policies
  - To be able to enable QoS on the Internet we need policies to include preferential queuing or dropping, admitting or denying access, or encrypting the packet's payload.
- Policy is comprised of the following:
  - Decision-making
    - Using an application-specific policy check the current state of the network with the desired state of the network and decides how to achieve the desired state of the network.
  - Enforcement
    - By using different mechanisms configures and modifies devices so can achieve the desired policy of the network.
  - Policing
    - Policing is an active or passive examination of the network, checking the state of the network if it's healthy. This policy is being continuously work around the clock.

QoS Protocols

QoS	Net	App	Description
most	X		Provisioned Resources end-to-end
	X	X	RSVP [IntServ Guarantee Services]
	X	X	RSVP [IntServ Controlled Services]
	X		Multi-Protocol Label Switching [MPLS]
	X	X	DiffServ.
	X	X	DiffServ or SBM
	X		Diffserv applied at network core ingress.
	X		Fair queuing applied by network elements (e.g. CFQ,WFQ,RED)
least			Best effort service

Improving QOS in IP Networks

Thus far: “making the best of best effort”

Future: next generation Internet with QoS guarantees

- RSVP: signaling for resource reservations
- Differentiated Services: differential guarantees
- Integrated Services: firm guarantees

- simple model for sharing and congestion studies:

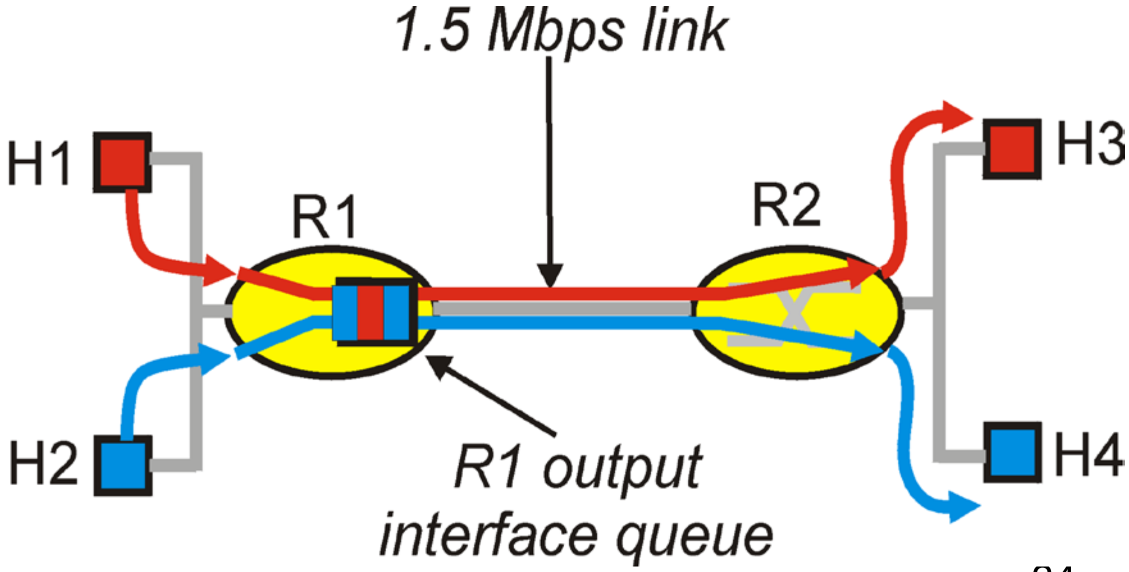

Principles for QOS Guarantees

- 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

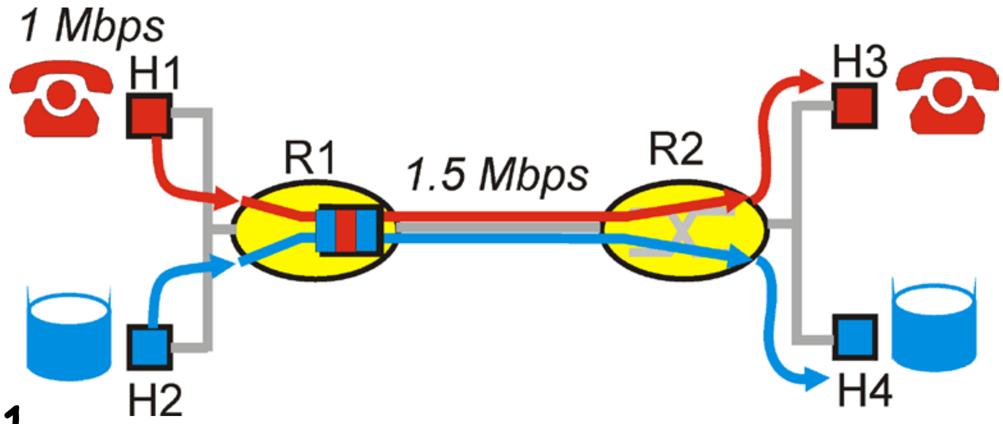

Principle 1

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)

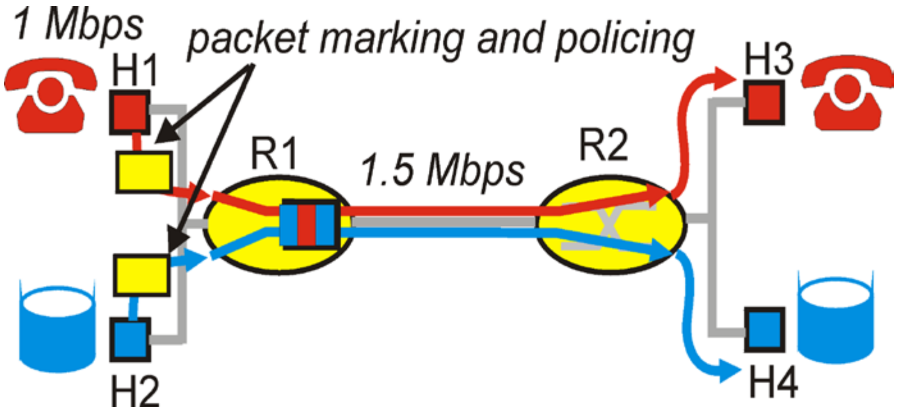

Principle 2

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

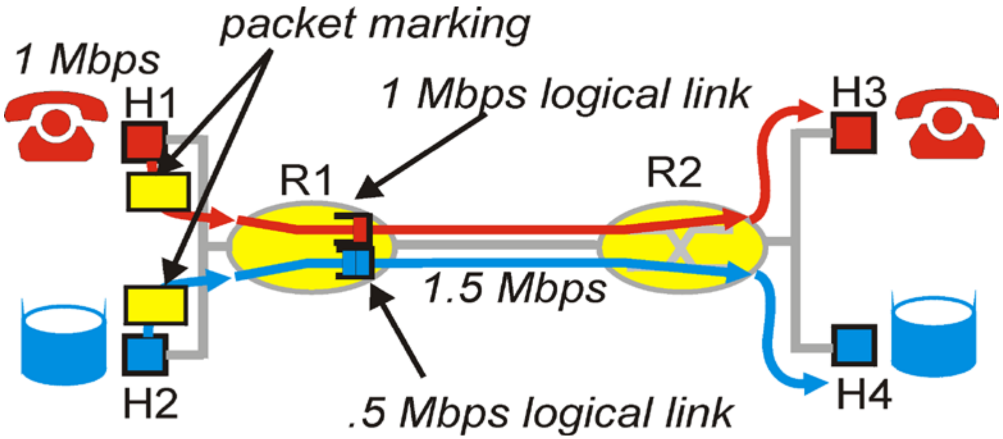

Principle 3

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

Principles for QOS Guarantees (more)

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

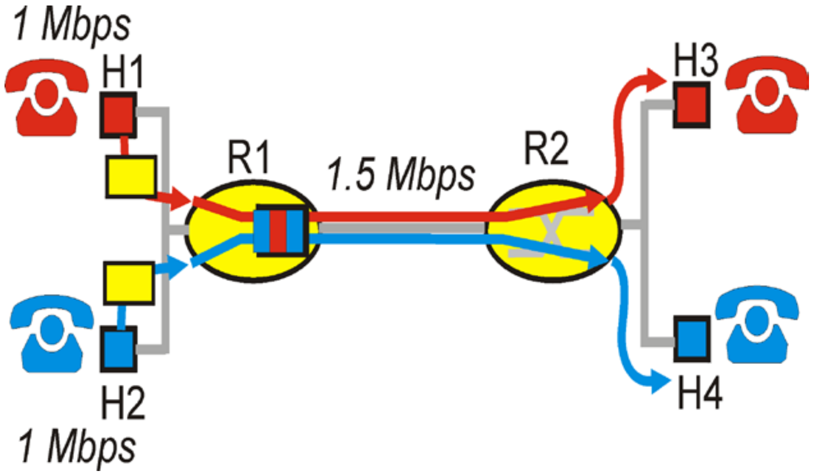

Principle 4

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

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

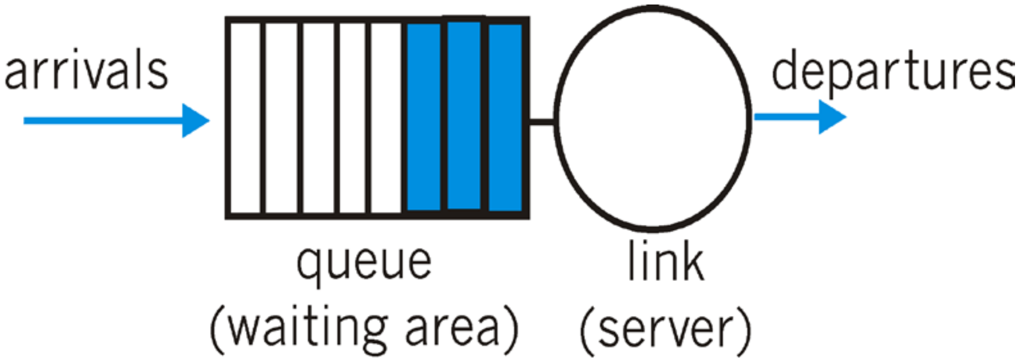

Scheduling Policies: more

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?

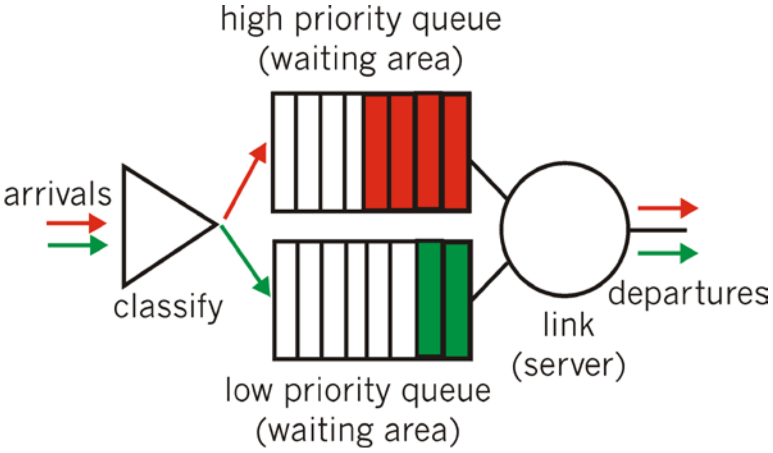

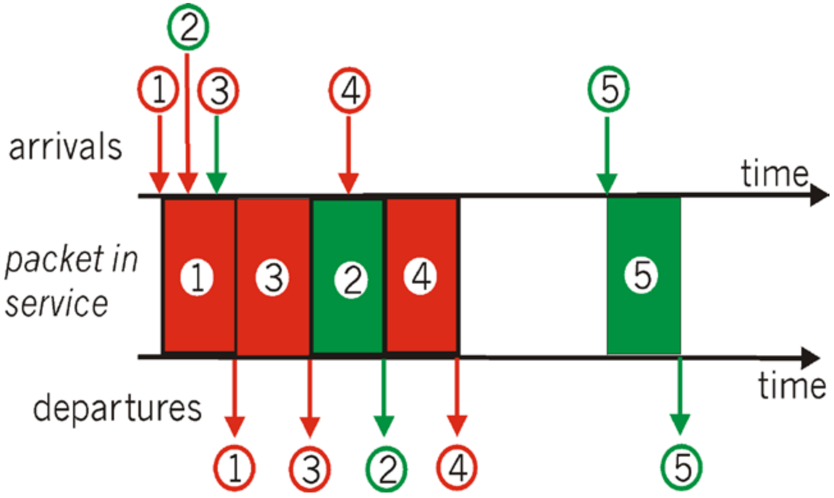

Scheduling Policies: still more

round robin scheduling:

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

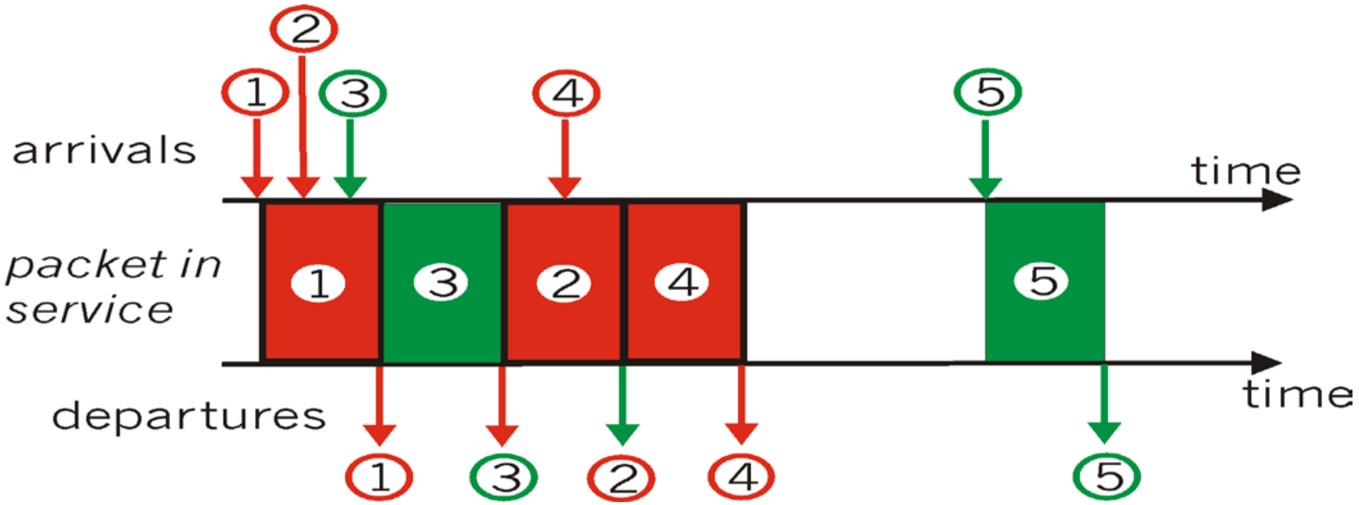

Scheduling Policies: still more

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

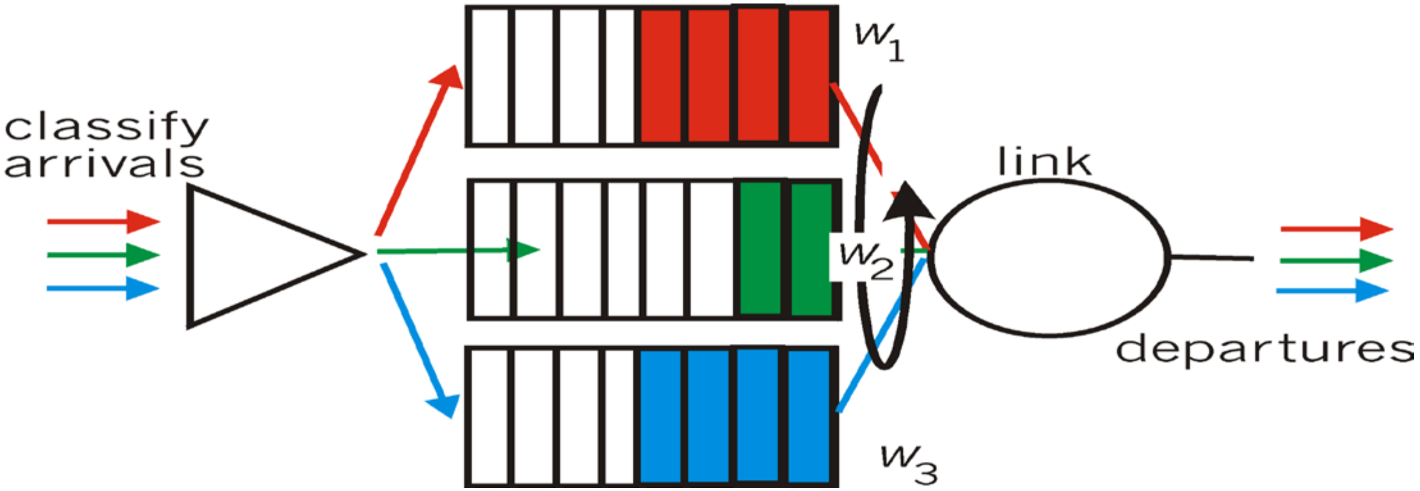

Policing Mechanisms

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.

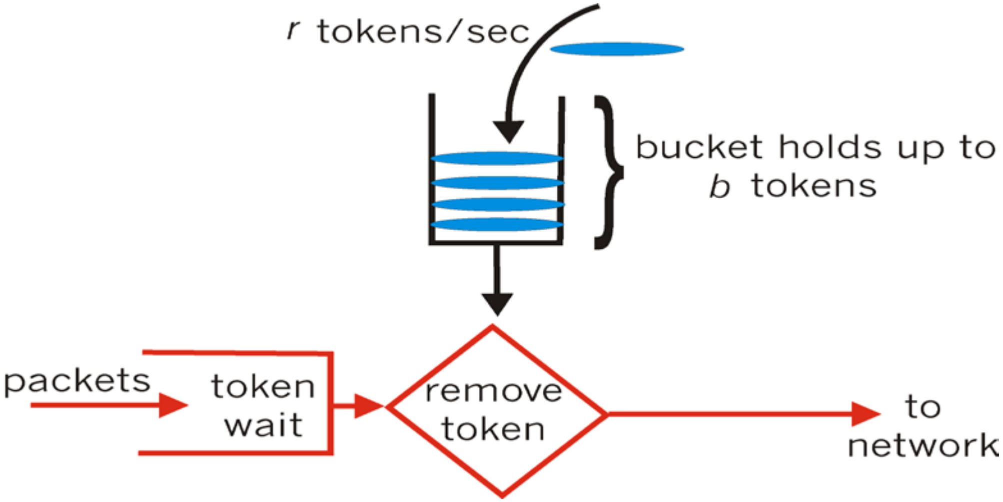

- bucket can hold b tokens
- tokens generated at rate *r token/sec* unless bucket full
- *over interval of length t: number of packets admitted less than or equal to (r t + b).*

Policing Mechanisms (more)

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

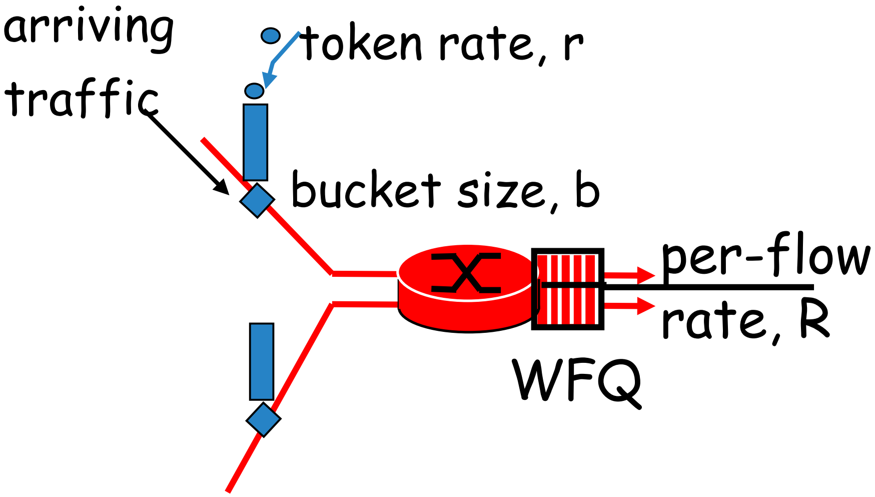

arriving traffic

token rate, r

bucket size, b

per-flow rate, R

WFQ

$D_{max} = b/R$

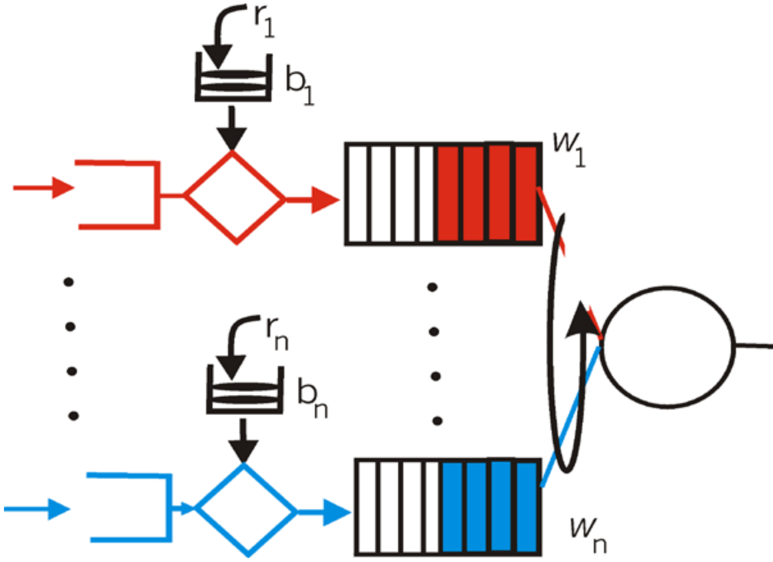

IETF Integrated Services

- 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 violating QoS guarantees made to already admitted flows?

Intserv: QoS guarantee scenario

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

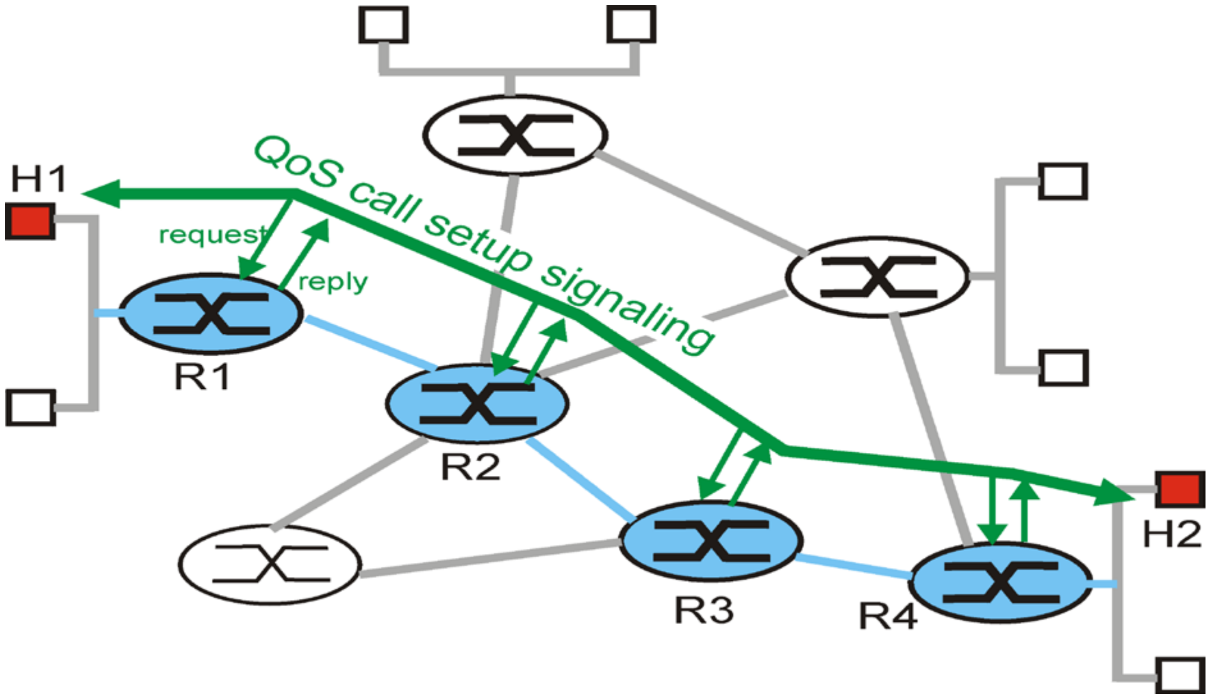

Call Admission

Arriving session must :

- declare its QOS requirement
  - R-spec: defines the QOS being requested
    - controlled-load: none
    - guaranteed: delay target
- characterize traffic it will send into network
  - T-spec: defines 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

IntServ – Signaling Protocol

- Signaling Protocol
  - needed to carry R-spec and T-spec to routers (where reservation is required)
  - RSVP – Resource Reservation Protocol

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."

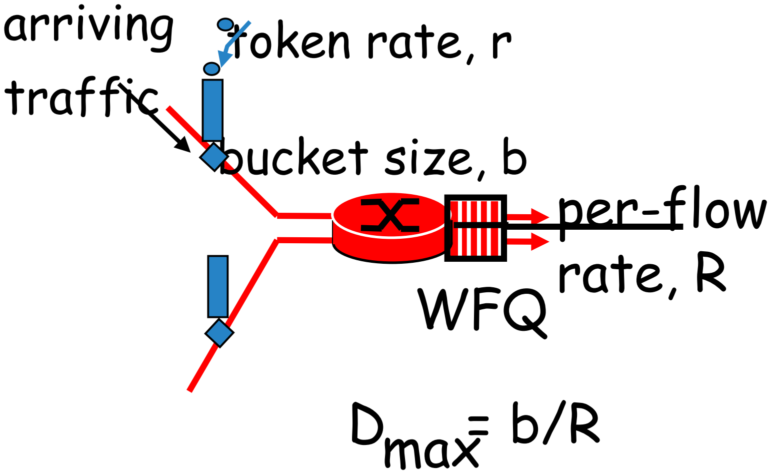

$D_{max} = b/R$

Quality of Service

- Integrated Services (RSVP)
  - Overview of Mechanisms
    - Flowspec
      - With a best-effort service we can just tell the network where we want our packets to go and leave it at that, a real-time service involves telling the network something more about the type of service we require
      - The set of information that we provide to the network is referred to as a *flowspec*.
    - Admission Control
      - When we ask the network to provide us with a particular service, the network needs to decide if it can in fact provide that service. The process of deciding when to say no is called *admission control*.
    - Resource Reservation
      - We need a mechanism by which the users of the network and the components of the network itself exchange information such as requests for service, flowspecs, and admission control decisions. We refer to this process as *resource reservation*

Quality of Service

- Integrated Services (RSVP)
  - Overview of Mechanisms
    - Packet Scheduling
      - Finally, when flows and their requirements have been described, and admission control decisions have been made, the network switches and routers need to meet the requirements of the flows.
      - A key part of meeting these requirements is managing the way packets are queued and scheduled for transmission in the switches and routers.
      - This last mechanism is *packet scheduling*.

Quality of Service

- Integrated Services (RSVP)
  - Flowspec
    - There are two separable parts to the flowspec:
      - The part that describes the flow's traffic characteristics (called the *TSpec)* and
      - The part that describes the service requested from the network (the *RSpec).*

      - The RSpec is very service specific and relatively easy to describe.
      - For example, with a controlled load service, the RSpec is trivial: The application just requests controlled load service with no additional parameters.
      - With a guaranteed service, you could specify a delay target or bound.

Quality of Service

- Integrated Services (RSVP)
  - Flowspec
    - Tspec
      - We need to give the network enough information about the bandwidth used by the flow to allow intelligent admission control decisions to be made
      - For most applications, the bandwidth is not a single number
        - It varies constantly
      - A video application will generate more bits per second when the scene is changing rapidly than when it is still
        - Just knowing the long term average bandwidth is not enough

Quality of Service

- Integrated Services (RSVP)
  - Flowspec
    - Suppose 10 flows arrive at a switch on separate ports and they all leave on the same 10 Mbps link
    - If each flow is expected to send no more than 1 Mbps
      - No problem
    - If these are variable bit applications such as compressed video
      - They will occasionally send more than the average rate
    - If enough sources send more than average rates, then the total rate at which data arrives at the switch will be more than 10 Mbps
    - This excess data will be queued
    - The longer the condition persists, the longer the queue will get

Quality of Service

- Integrated Services (RSVP)
  - Flowspec
    - One way to describe the Bandwidth characteristics of sources is called a Token Bucket Filter
    - The filter is described by two parameters
      - A token rate r
      - A bucket depth B
    - To be able to send a byte, a token is needed
    - To send a packet of length n, n tokens are needed
    - Initially there are no tokens
    - Tokens are accumulated at a rate of r per second
    - No more than B tokens can be accumulated

Quality of Service

- Integrated Services (RSVP)
  - Flowspec
    - We can send a burst of as many as *B* bytes into the network as fast as we want, but over significant long interval we cannot send more than *r* bytes per second
    - This information is important for admission control algorithm when it tries to find out whether it can accommodate new request for service

Quality of Service

- Flowspec

  - The figure illustrates how a token bucket can be used to characterize a flow’s Bandwidth requirement
  - For simplicity, we assume each flow can send data as individual bytes rather than as packets
  - Flow A generates data at a steady rate of 1 MBps
    - So it can be described by a token bucket filter with a rate *r* = 1 MBps and a bucket depth of 1 byte
    - This means that it receives tokens at a rate of 1 MBps but it cannot store more than 1 token, it spends them immediately

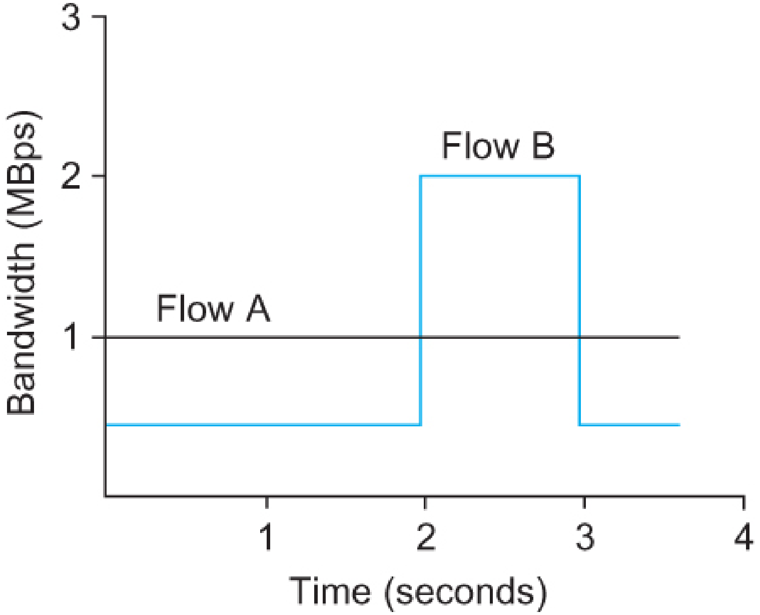

Quality of Service

- Flowspec
  - Flow B sends at a rate that averages out to 1 MBps over the long term, but does so by sending at 0.5 MBps for 2 seconds and then at 2 MBps for 1 second
  - Since the token bucket rate *r* is a long term average rate, flow B can be described by a token bucket with a rate of 1 MBps
  - Unlike flow A, however flow B needs a bucket depth *B* of at least 1 MB, so that it can store up tokens while it sends at less than 1 MBps to be used when it sends at 2 MBps

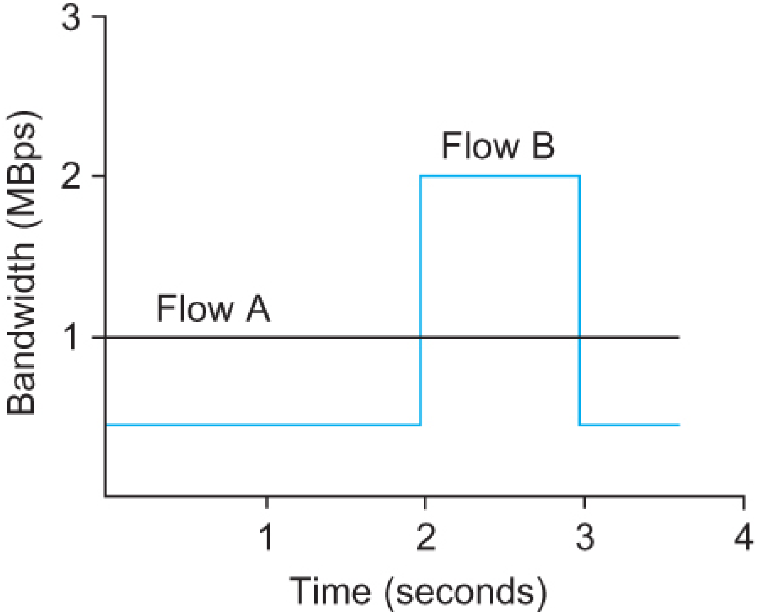

Quality of Service

- Flowspec

- For the first 2 seconds, it receives tokens at a rate of 1 MBps but spends them at only 0.5 MBps,
  - So it can save up 2 × 0.5 = 1 MB of tokens which it spends at the 3rd second

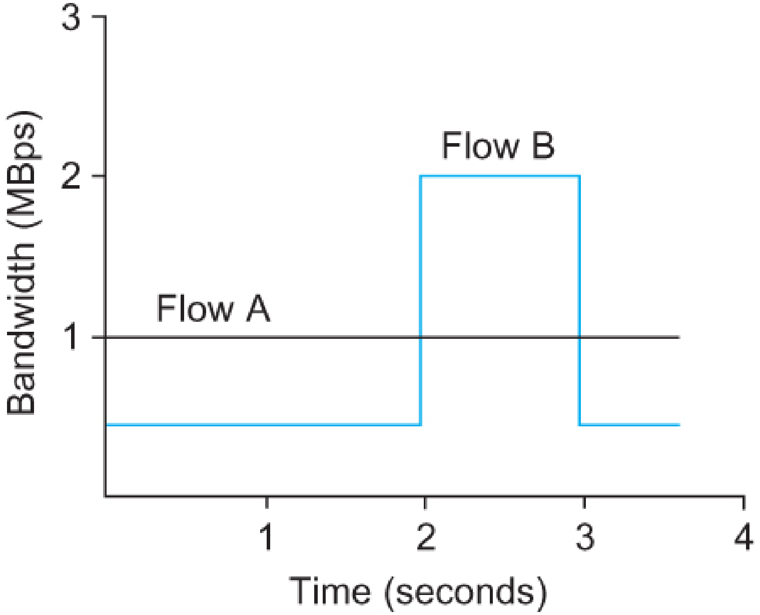

Quality of Service

- Integrated Services (RSVP)
  - Admission Control
    - The idea behind admission control is simple: When some new flow wants to receive a particular level of service, admission control looks at the TSpec and RSpec of the flow and tries to decide if the desired service can be provided to that amount of traffic, given the currently available resources, without causing any previously admitted flow to receive worse service than it had requested. If it can provide the service, the flow is admitted; if not, then it is denied.

Quality of Service

- Integrated Services (RSVP)
  - Reservation Protocol
    - While connection-oriented networks have always needed some sort of setup protocol to establish the necessary virtual circuit state in the switches, connectionless networks like the Internet have had no such protocols.
    - However we need to provide a lot more information to our network when we want a real-time service from it.
    - While there have been a number of setup protocols proposed for the Internet, the one on which most current attention is focused is called Resource Reservation Protocol (RSVP).

Quality of Service

- Integrated Services (RSVP)
  - Reservation Protocol
    - One of the key assumptions underlying RSVP is that it should not detract from the robustness that we find in today's connectionless networks.
    - Because connectionless networks rely on little or no state being stored in the network itself, it is possible for routers to crash and reboot and for links to go up and down while end-to-end connectivity is still maintained.
    - RSVP tries to maintain this robustness by using the idea of *soft state in the routers.*

Quality of Service

- Integrated Services (RSVP)
  - Reservation Protocol
    - Another important characteristic of RSVP is that it aims to support multicast flows just as effectively as unicast flows
    - Initially, consider the case of one sender and one receiver trying to get a reservation for traffic flowing between them.
    - There are two things that need to happen before a receiver can make the reservation.

Quality of Service

- Integrated Services (RSVP)
  - Reservation Protocol
    - First, the receiver needs to know what traffic the sender is likely to send so that it can make an appropriate reservation. That is, it needs to know the sender’s TSpec.
    - Second, it needs to know what path the packets will follow from sender to receiver, so that it can establish a resource reservation at each router on the path. Both of these requirements can be met by sending a message from the sender to the receiver that contains the TSpec.
    - Obviously, this gets the TSpec to the receiver. The other thing that happens is that each router looks at this message (called a PATH message) as it goes past, and it figures out the *reverse path that will be used to send* reservations from the receiver back to the sender in an effort to get the reservation to each router on the path.

Quality of Service

- Integrated Services (RSVP)
  - Reservation Protocol
    - Having received a PATH message, the receiver sends a reservation back "up" the multicast tree in a RESV message.
    - This message contains the sender's TSpec and an RSpec describing the requirements of this receiver.
    - Each router on the path looks at the reservation request and tries to allocate the necessary resources to satisfy it. If the reservation can be made, the RESV request is passed on to the next router.
    - If not, an error message is returned to the receiver who made the request. If all goes well, the correct reservation is installed at every router between the sender and the receiver.
    - As long as the receiver wants to retain the reservation, it sends the same RESV message about once every 30 seconds.

Quality of Service

- Integrated Services (RSVP)
  - Reservation Protocol

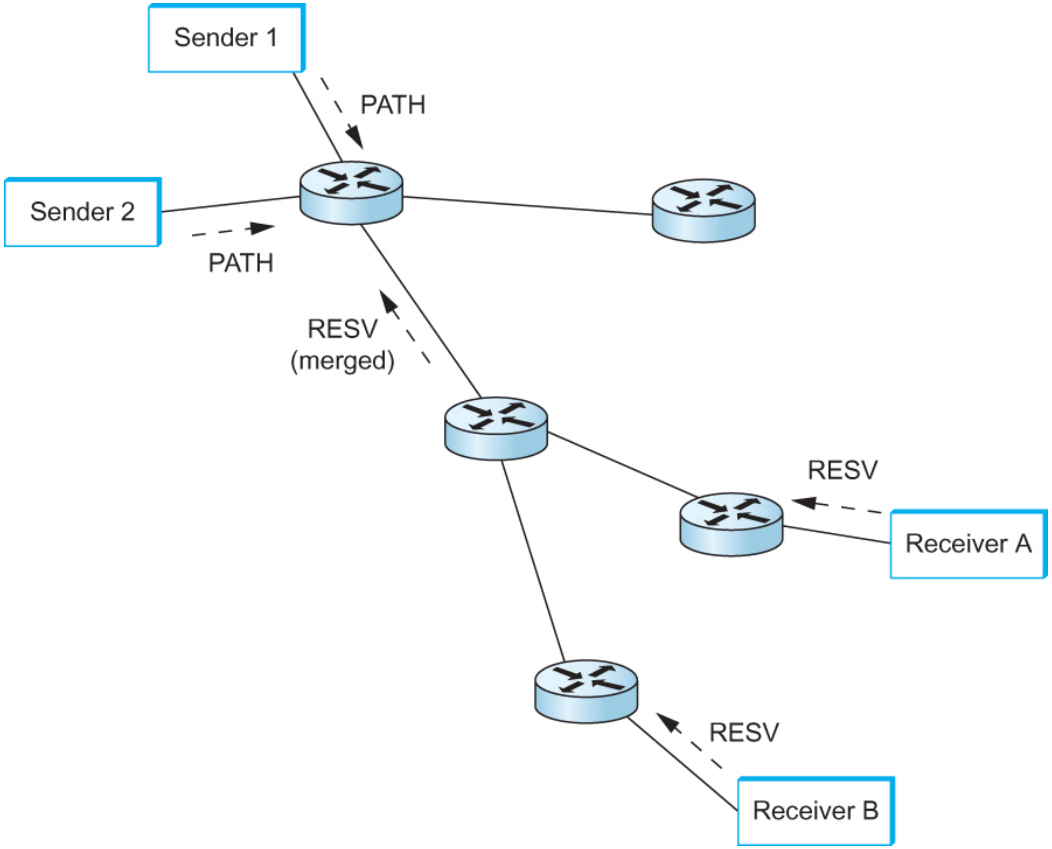

Making reservations on a multicast tree

Quality of Service

- Integrated Services (RSVP)
  - Packet Classifying and Scheduling
    - Once we have described our traffic and our desired network service and have installed a suitable reservation at all the routers on the path, the only thing that remains is for the routers to actually deliver the requested service to the data packets. There are two things that need to be done:
      - Associate each packet with the appropriate reservation so that it can be handled correctly, a process known as *classifying packets*.
      - Manage the packets in the queues so that they receive the service that has been requested, a process known as packet *scheduling*.

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

Path msgs: RSVP *sender-to-network* signaling

- path message contents:
  - *address:* unicast destination, or multicast group
  - *flowspec:* bandwidth requirements spec.
  - *filter flag:* if yes, record identities of upstream senders (to allow packets filtering by source)
  - *previous hop:* upstream router/host ID
  - *refresh time:* time until this info times out
- path message: communicates sender info, and reverse-path-to-sender routing info
  - later upstream forwarding of receiver reservations

RSVP: simple audio conference

- H1, H2, H3, H4, H5 both senders and receivers
- multicast group m1
- no filtering: packets from any sender forwarded
- audio rate: *b*
- only one multicast routing tree possible

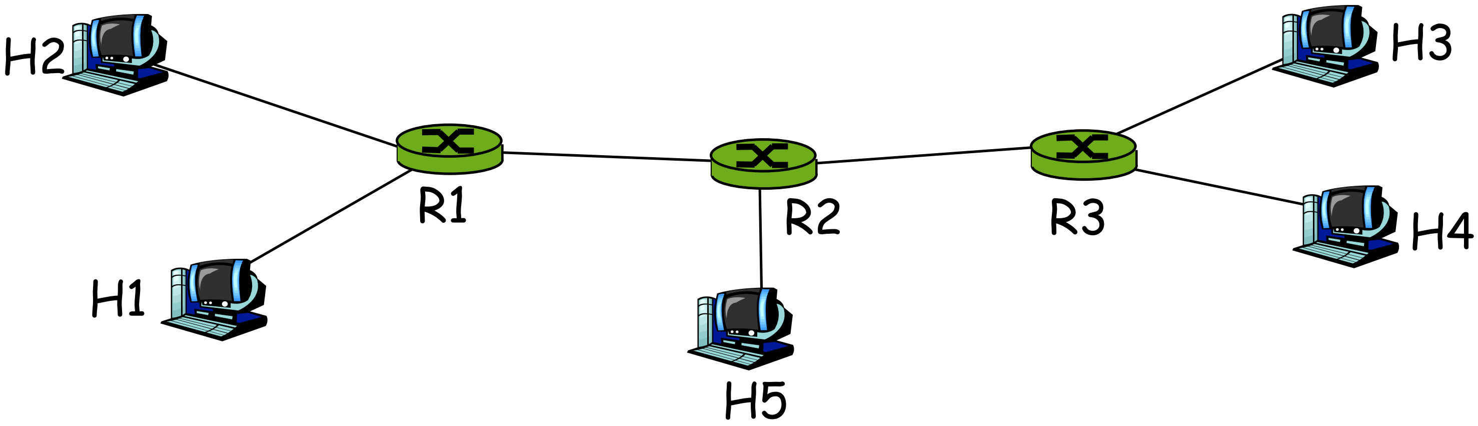

RSVP: building up path state

- H1, …, H5 all send path messages on *m1:*

  (address=*m1*, Tspec=*b*, filter-spec=no-filter,refresh=100)

- Suppose H1 sends first path message

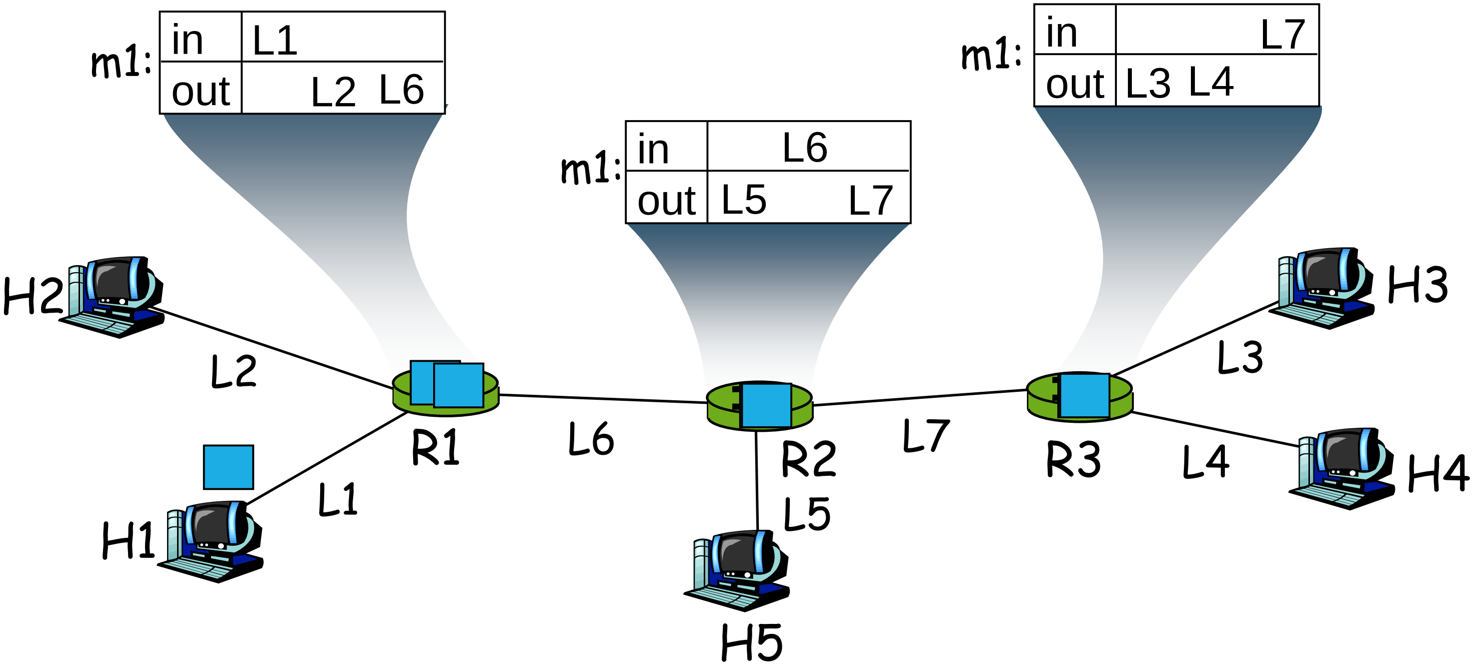

m1:

in	L1
out	L2 L6

m1:

in	L6
out	L5 L7

m1:

in	L7
out	L3 L4

RSVP: building up path state

- next, H5 sends path message, creating more state in routers

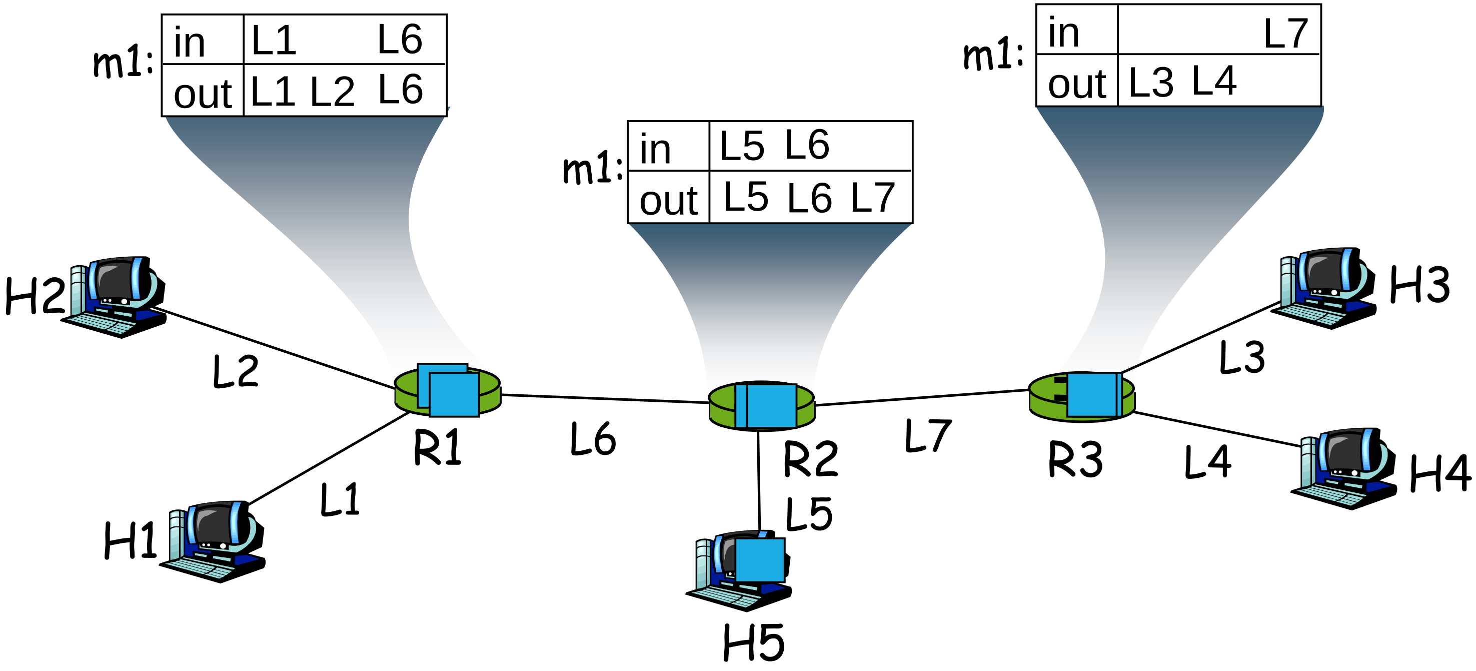

m1:

in	L1		L6
out	L1	L2	L6

m1:

in	L5	L6	
out	L5	L6	L7

m1:

in		L7
out	L3	L4

RSVP: building up path state

- H2, H3, H5 send path msgs, completing path state tables

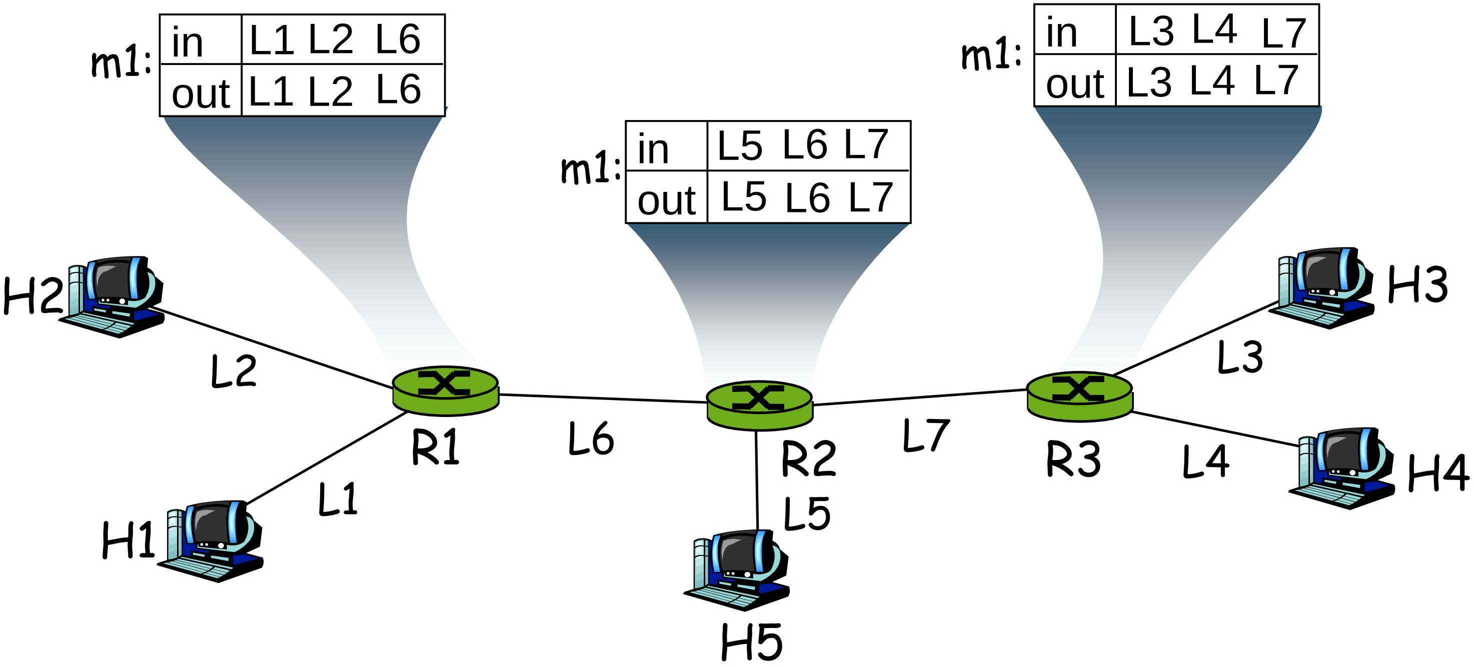

m1:

in	L1 L2 L6
out	L1 L2 L6

m1:

in	L5 L6 L7
out	L5 L6 L7

m1:

in	L3 L4 L7
out	L3 L4 L7

reservation msgs: *receiver-to-network* signaling

- reservation message contents:
  - *desired bandwidth:*
  - *filter type:*
    - no filter: any packets address to multicast group can use reservation
    - fixed filter: only packets from specific set of senders can use reservation
    - dynamic filter: senders who's p[ackets can be forwarded across link will change (by receiver choce) over time.
  - *filter spec*
- reservations flow upstream from receiver-to-senders, reserving resources, creating additional, *receiver-related* state at routers

RSVP: *receiver* reservation example 1

H1 wants to receive audio from all other senders

- H1 reservation msg flows uptree to sources
- H1 only reserves enough bandwidth for 1 audio stream
- reservation is of type "no filter" – any sender can use reserved bandwidth

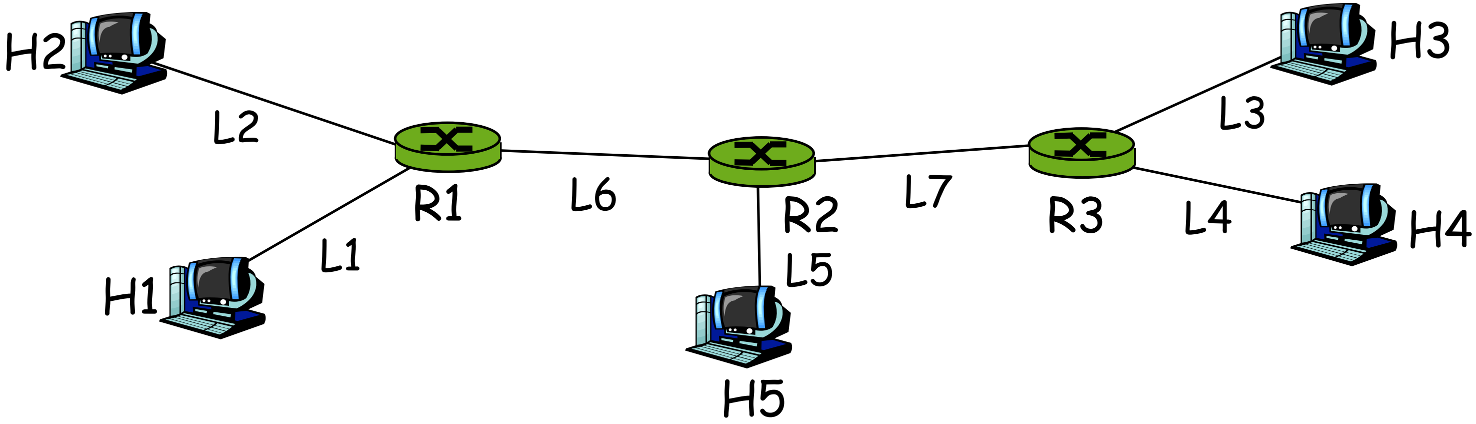

RSVP: *receiver* reservation example 1

- H1 reservation msgs flows uptree to sources
- routers, hosts reserve bandwidth b needed on downstream links towards H1

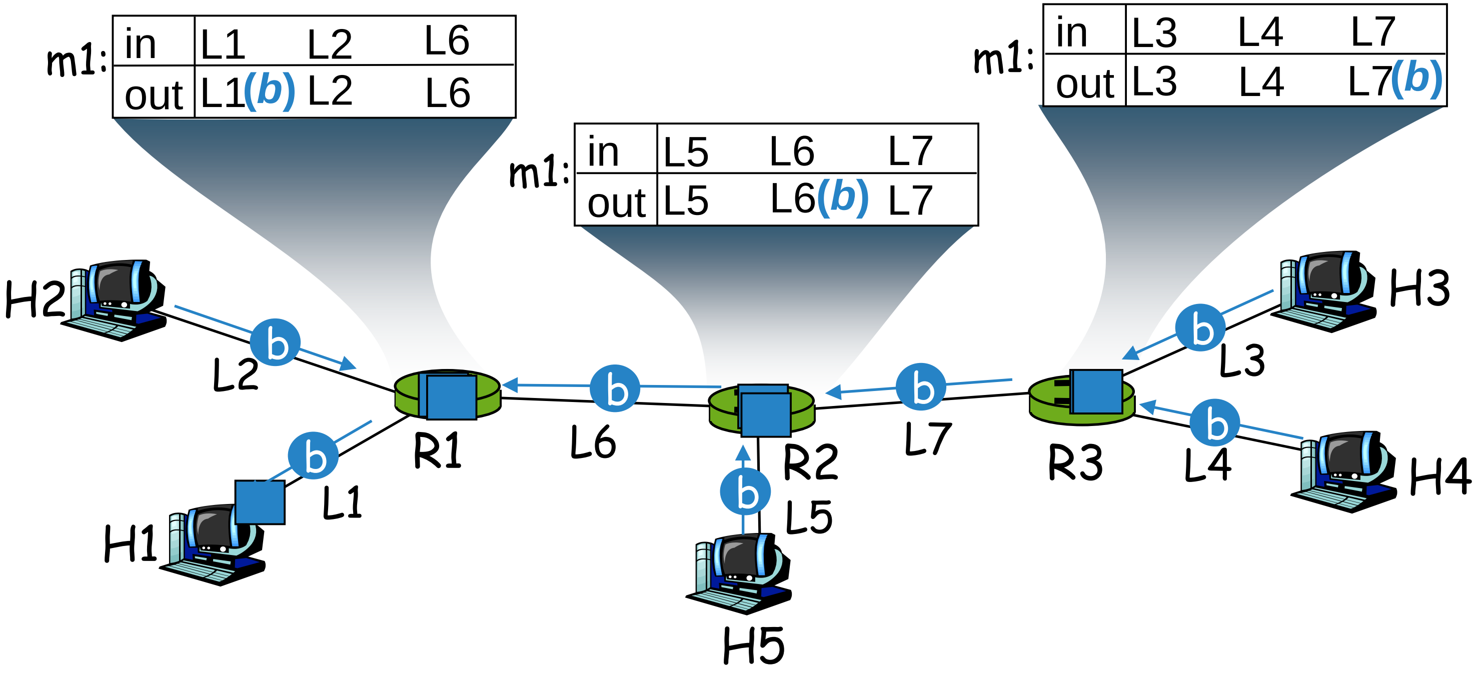

m1:

in	L1	L2	L6
out	L1(*b*)	L2	L6

m1:

in	L5	L6	L7
out	L5	L6(*b*)	L7

m1:

in	L3	L4	L7
out	L3	L4	L7(*b*)

RSVP: *receiver* reservation example 1 (more)

- next, H2 makes no-filter reservation for bandwidth *b*
- H2 forwards to R1, R1 forwards to H1 and R2 (?)
- R2 takes no action, since *b* already reserved on L6

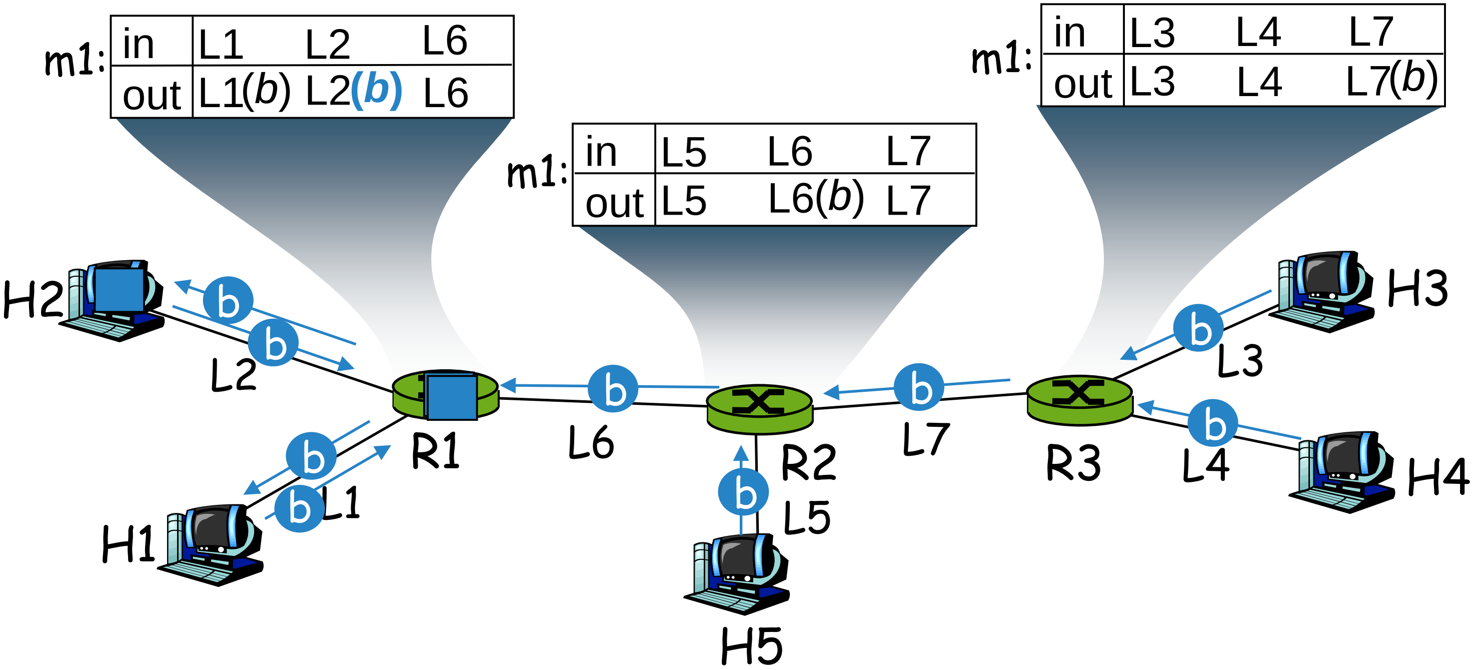

m1:

in	L1	L2	L6
out	L1(*b*)	L2(***b***)	L6

m1:

in	L5	L6	L7
out	L5	L6(*b*)	L7

m1:

in	L3	L4	L7
out	L3	L4	L7(*b*)

RSVP: *receiver* reservation: issues

What if multiple senders (e.g., H3, H4, H5) over link (e.g., L6)?

- arbitrary interleaving of packets
- L6 flow policed by leaky bucket: if H3+H4+H5 sending rate exceeds b, packet loss will occur

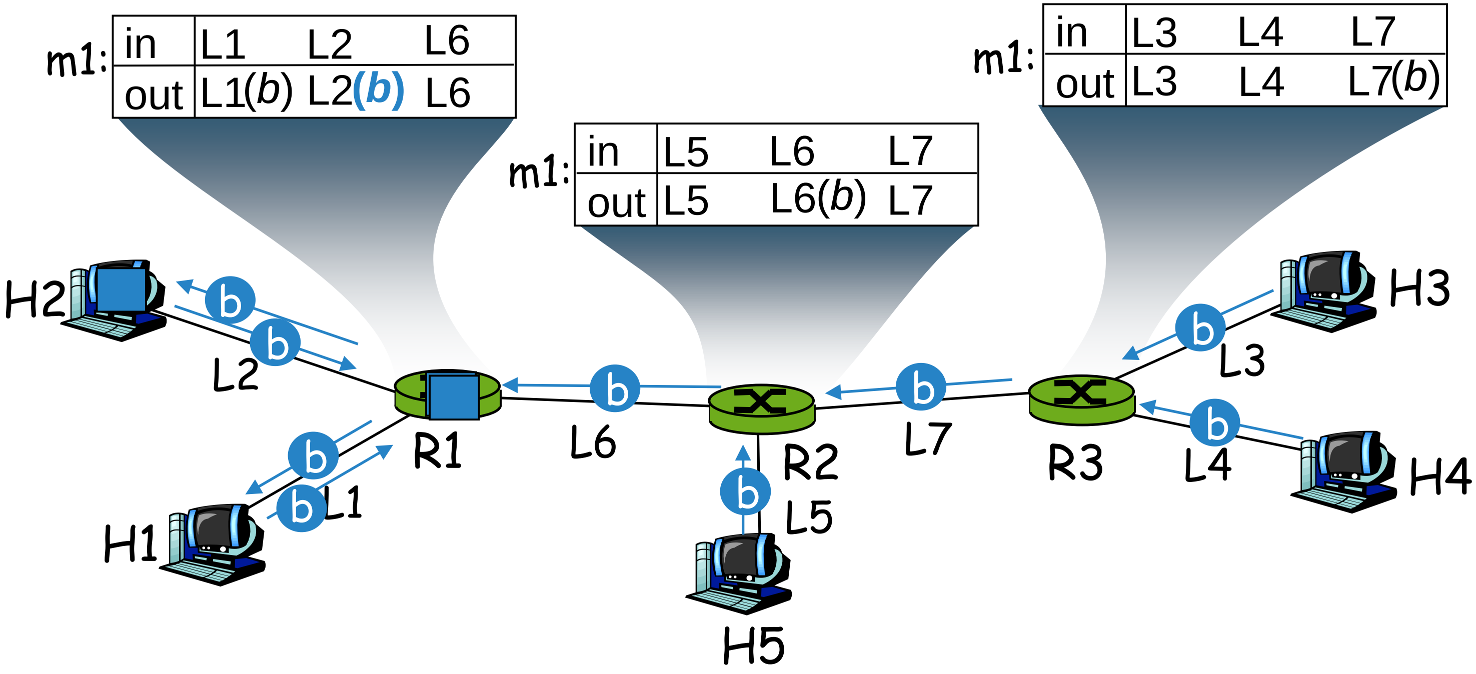

m1:

in	L1	L2	L6
out	L1(*b*)	L2(***b***)	L6

m1:

in	L5	L6	L7
out	L5	L6(*b*)	L7

m1:

in	L3	L4	L7
out	L3	L4	L7(*b*)

RSVP: example 2

- H1, H4 are only senders
  - send *path messages* as before, indicating filtered reservation
  - Routers store upstream senders for each upstream link
- H2 will want to receive from H4 (only)

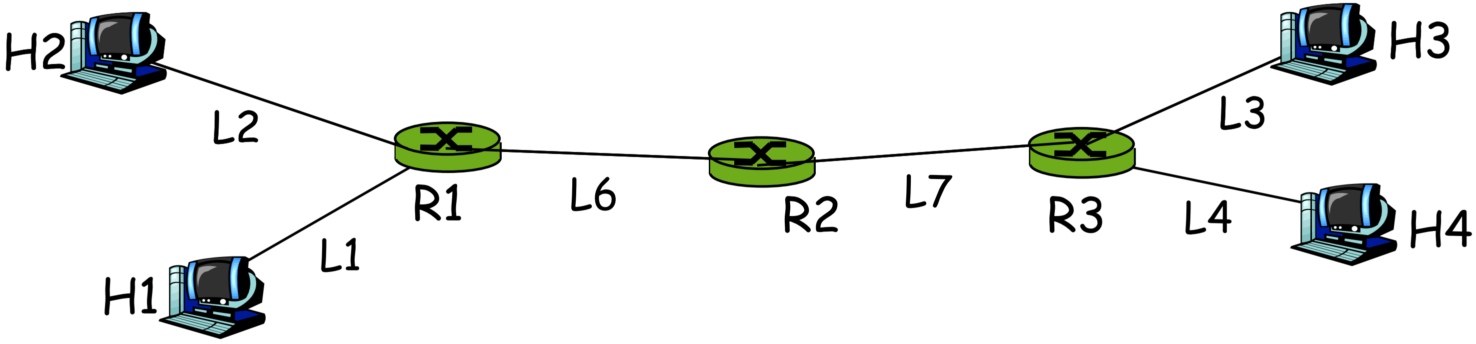

RSVP: example 2

- H1, H4 are only senders
  - send *path messages* as before, indicating filtered reservation

in	L1, L6

out	L2(H1-via-H1 ; H4-via-R2)
	L6(H1-via-H1)
	L1(H4-via-R2)

in	L4, L7

out	L3(H4-via-H4 ; H1-via-R3)
	L4(H1-via-R2)
	L7(H4-via-H4)

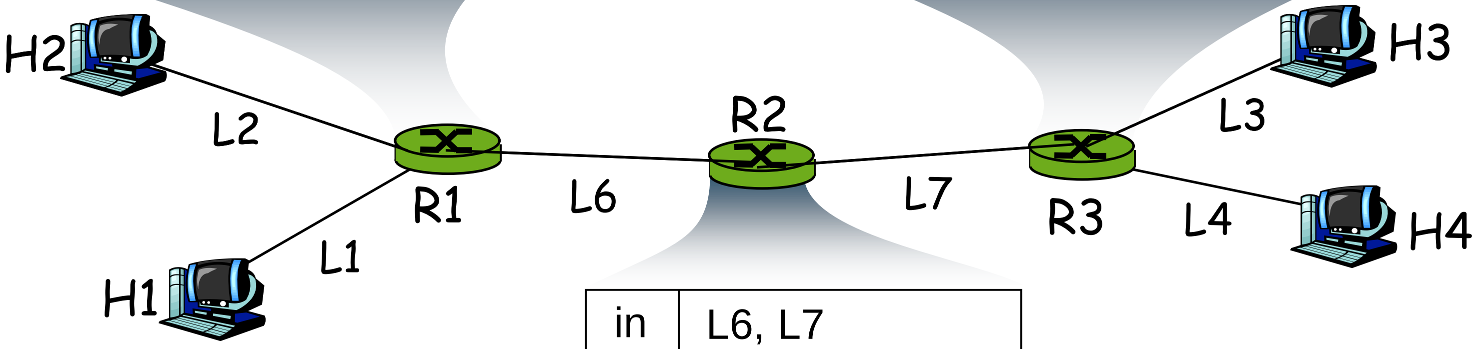

in	L6, L7

out	L6(H4-via-R3)
	L7(H1-via-R1)

RSVP: example 2

- receiver H2 sends reservation message for source H4 at bandwidth *b*
  - propagated upstream towards H4, reserving *b*

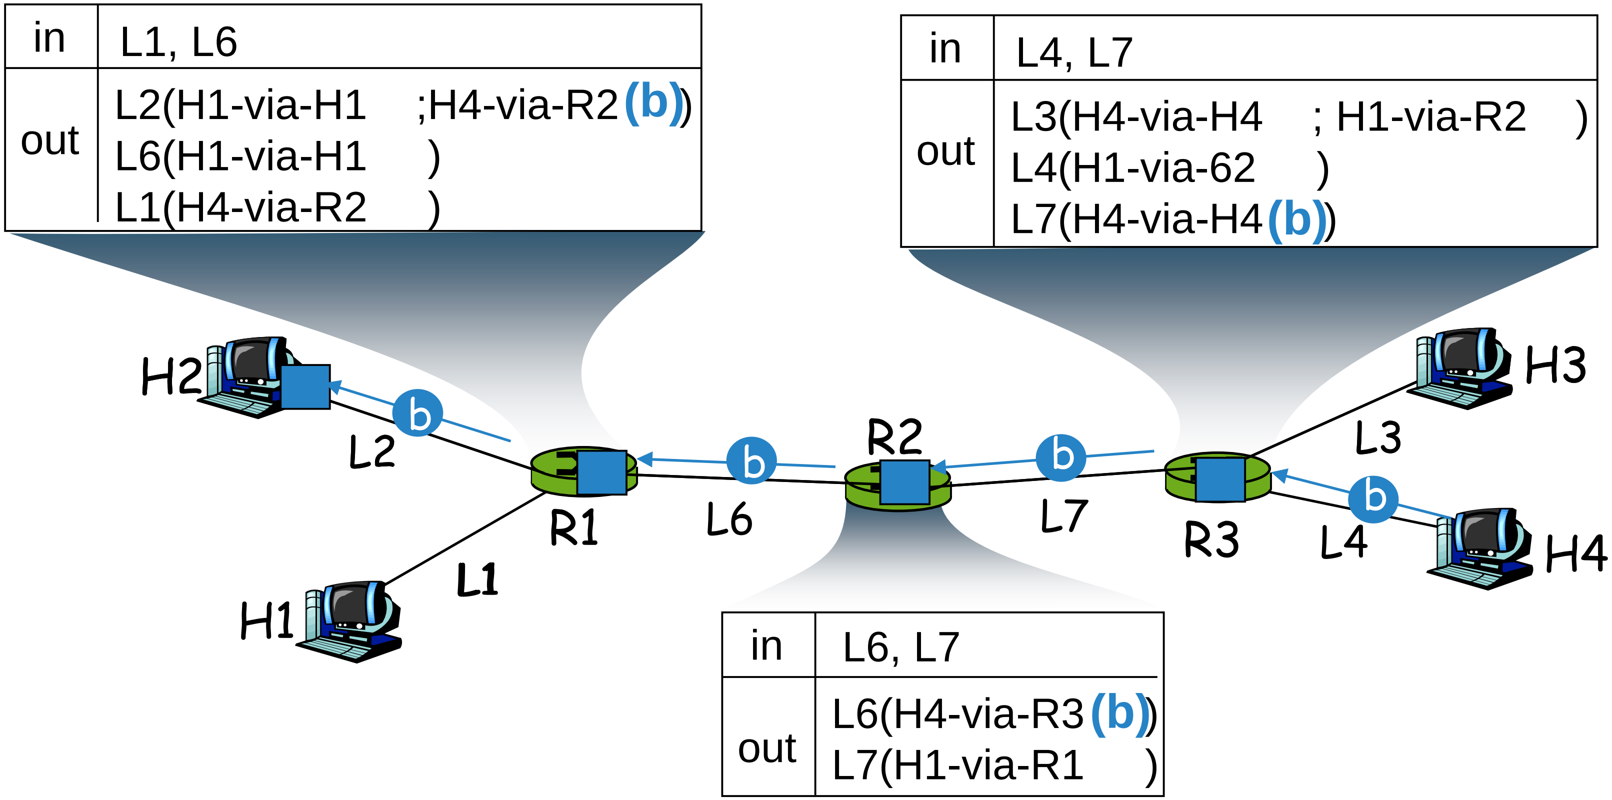

in	L1, L6

out	L2(H1-via-H1 ;H4-via-R2 **(b)**)
	L6(H1-via-H1)
	L1(H4-via-R2)

in	L4, L7

out	L3(H4-via-H4 ; H1-via-R2)
	L4(H1-via-62)
	L7(H4-via-H4 **(b)**)

in	L6, L7

out	L6(H4-via-R3 **(b)**)
	L7(H1-via-R1)

RSVP: *soft-state*

- senders periodically resend path msgs to refresh (maintain) state
- receivers periodically resend resv msgs to refresh (maintain) state
- path and resv msgs have TTL field, specifying refresh interval

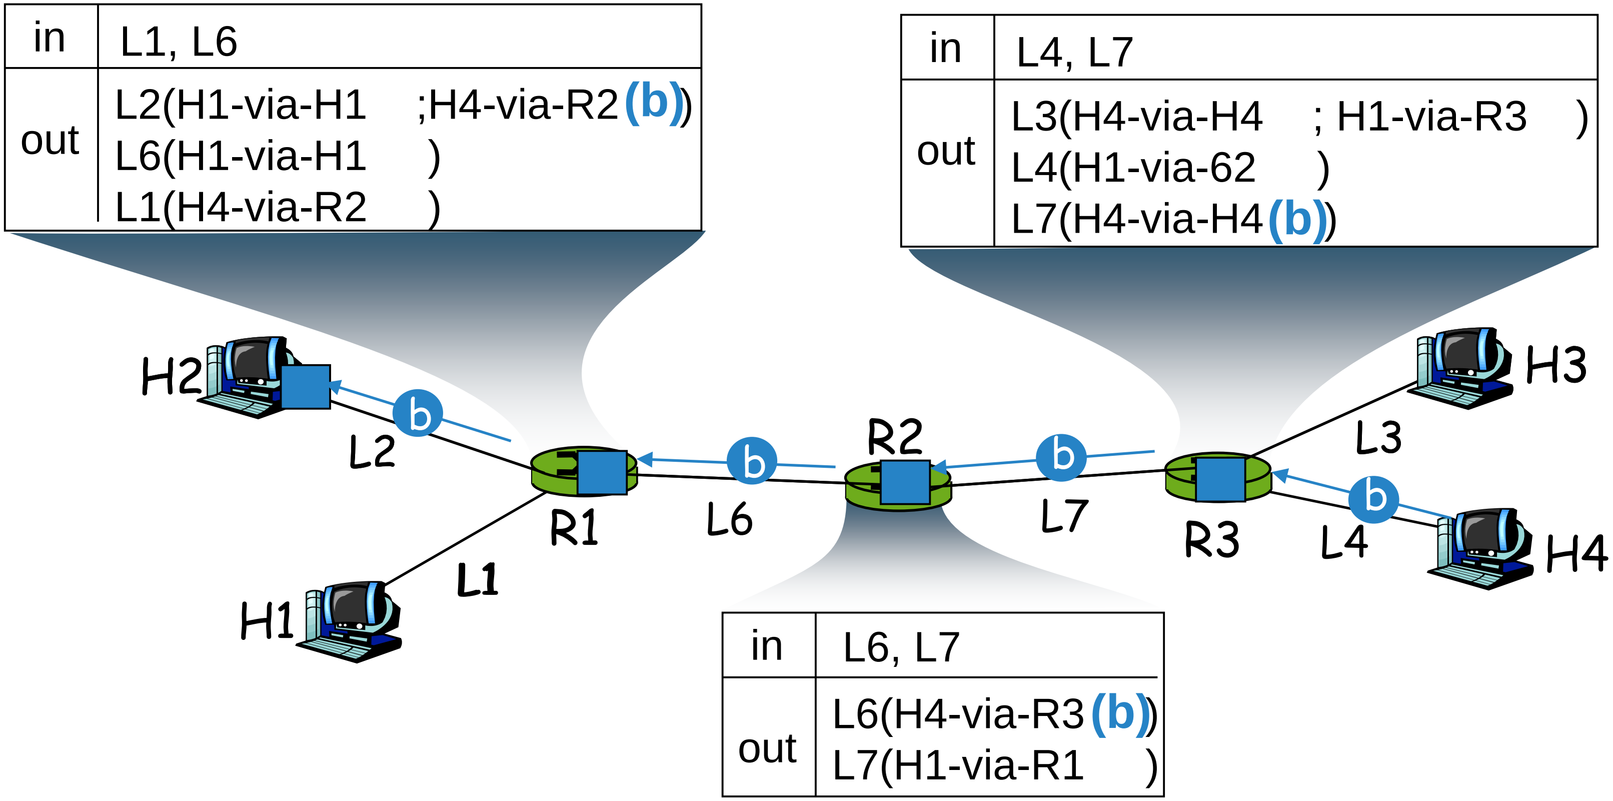

in	L1, L6

out	L2(H1-via-H1 ;H4-via-R2 (b))
	L6(H1-via-H1)
	L1(H4-via-R2)

in	L4, L7

out	L3(H4-via-H4 ; H1-via-R3)
	L4(H1-via-62)
	L7(H4-via-H4 (b))

in	L6, L7

out	L6(H4-via-R3 (b))
	L7(H1-via-R1)

RSVP: *soft-state*

- suppose H4 (sender) leaves without performing teardown
- eventually state in routers will timeout and disappear!

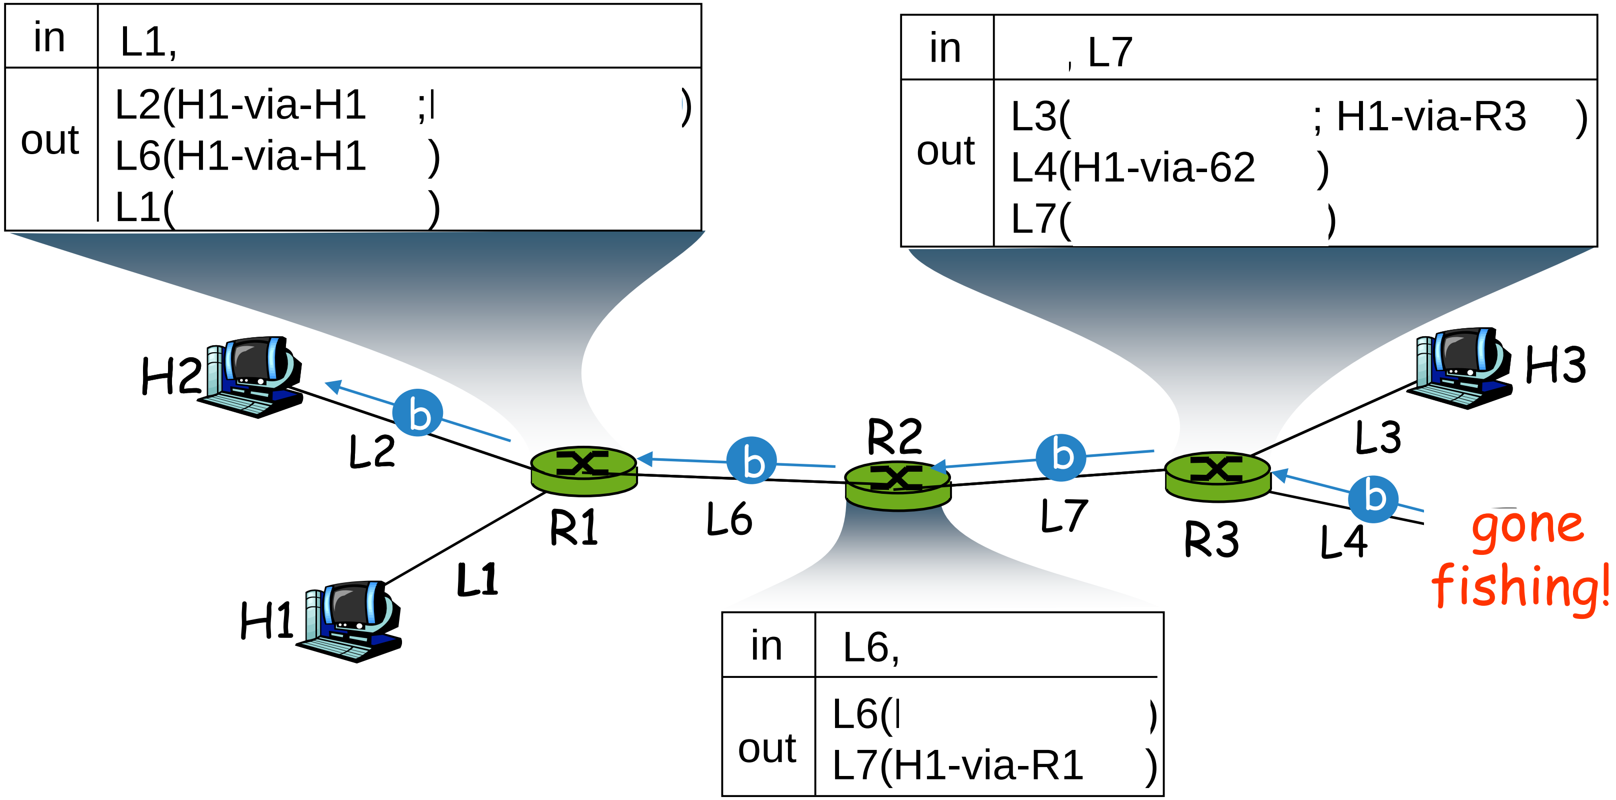

in	L1,

out	L2(H1-via-H1 ;I)
	L6(H1-via-H1)
	L1( )

in	, L7

out	L3(; H1-via-R3)
	L4(H1-via-62)
	L7( )

in	L6,

out	L6(I)
	L7(H1-via-R1)

The many uses of reservation/path refresh

- recover from an earlier lost refresh message
  - expected time until refresh received must be longer than timeout interval! (short timer interval desired)
- Handle receiver/sender that goes away without teardown
  - Sender/receiver state will timeout and disappear
- Reservation refreshes will cause new reservations to be made to a receiver from a sender who has joined since receivers last reservation refresh
  - E.g., in previous example, H1 is only receiver, H3 only sender. Path/reservation messages complete, data flows
  - H4 joins as sender, nothing happens until H3 refreshes reservation, causing R3 to forward reservation to H4, which allocates bandwidth

RSVP: reflections

- multicast as a “first class” service
- receiver-oriented reservations
- use of soft-state

IETF Differentiated Services

Concerns with Intserv:

- Scalability: signaling, maintaining per-flow router state difficult with large number of flows
- Flexible Service Models: Intserv has only two classes. Also want "qualitative" service classes
  - "behaves like a wire"
  - relative service distinction: Platinum, Gold, Silver

Diffserv approach:

- simple functions in network core, relatively complex functions at edge routers (or hosts)
- Don't define service classes, provide functional components to build service classes

Quality of Service

- Differentiated Services
  - Suppose that we have decided to enhance the best-effort service model by adding just one new class, which we'll call "premium."
  - Clearly we will need some way to figure out which packets are premium and which are regular old best effort.
  - Rather than using a protocol like RSVP to tell all the routers that some flow is sending premium packets, it would be much easier if the packets could just identify themselves to the router when they arrive. This could obviously be done by using a bit in the packet header—if that bit is a 1, the packet is a premium packet; if it's a 0, the packet is best effort

Quality of Service

- Differentiated Services
  - With this in mind, there are two questions we need to address:
    - Who sets the premium bit, and under what circumstances?
    - What does a router do differently when it sees a packet with the bit set?

Quality of Service

- Differentiated Services
  - There are many possible answers to the first question, but a common approach is to set the bit at an administrative boundary.
  - For example, the router at the edge of an Internet service provider's network might set the bit for packets arriving on an interface that connects to a particular company's network.
  - The Internet service provider might do this because that company has paid for a higher level of service than best effort.

Quality of Service

- Differentiated Services
  - Assuming that packets have been marked in some way, what do the routers that encounter marked packets do with them?
  - Here again there are many answers. In fact, the IETF standardized a set of router behaviors to be applied to marked packets. These are called “per-hop behaviors” (PHBs), a term that indicates that they define the behavior of individual routers rather than end-to-end services

Diffserv Architecture

Edge router:

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

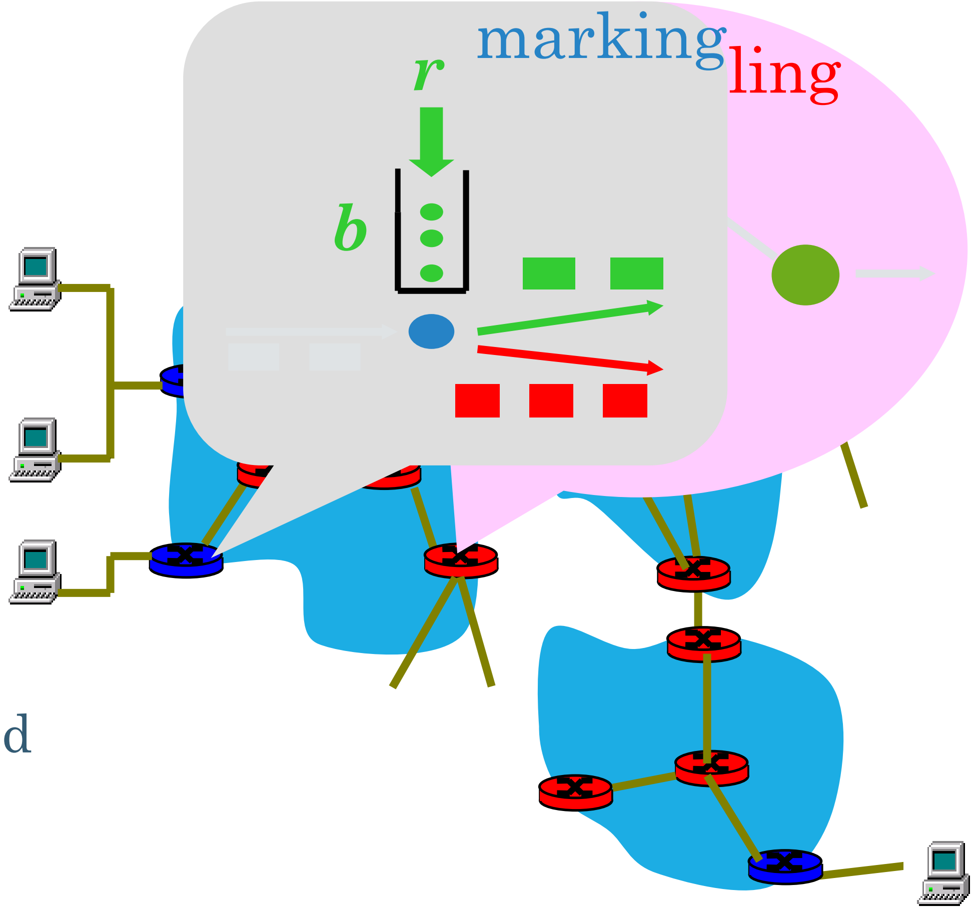

Core router:

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

Diffserv Architecture

Edge router:

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

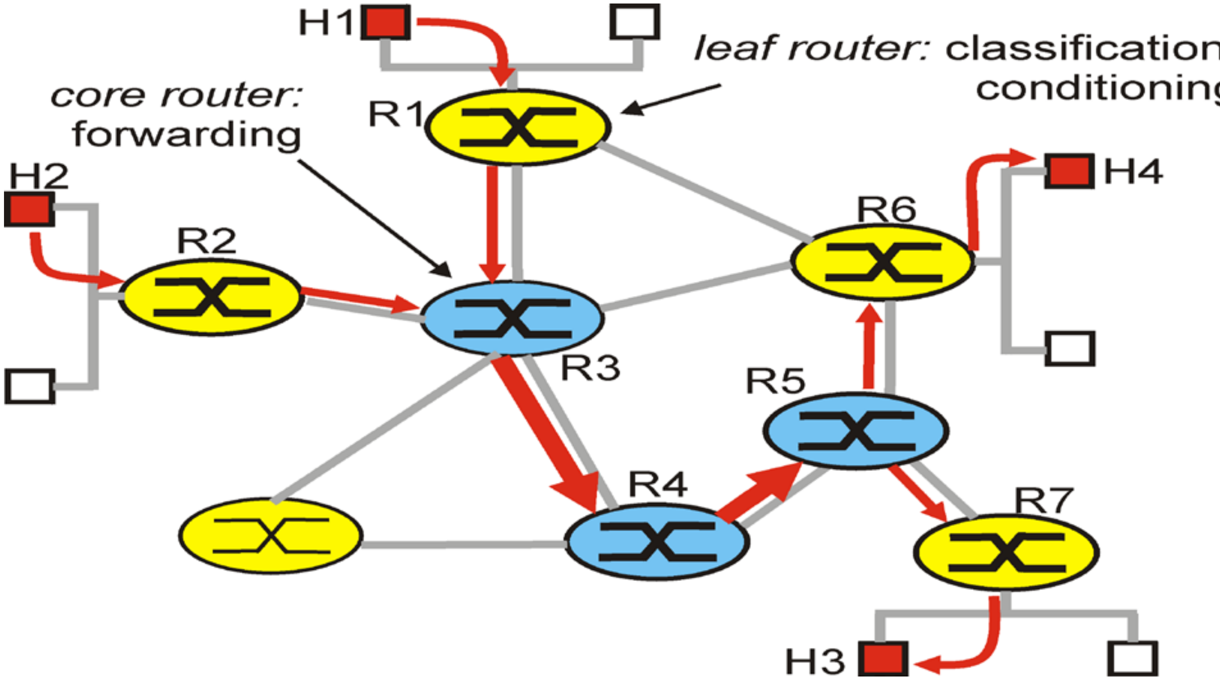

Core router:

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

Edge-router Packet Marking

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

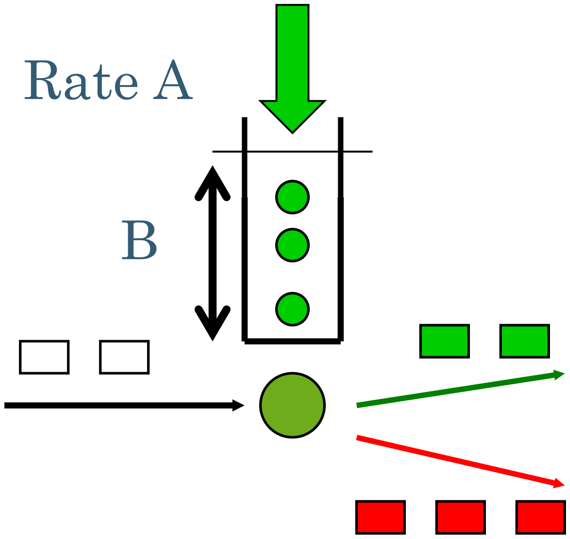

- Possible usage of marking:

- ❒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

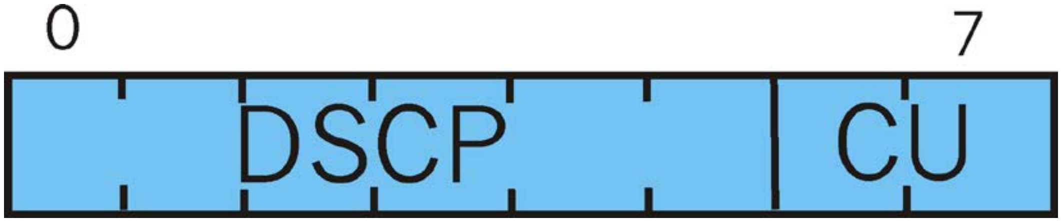

Classification and Conditioning

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

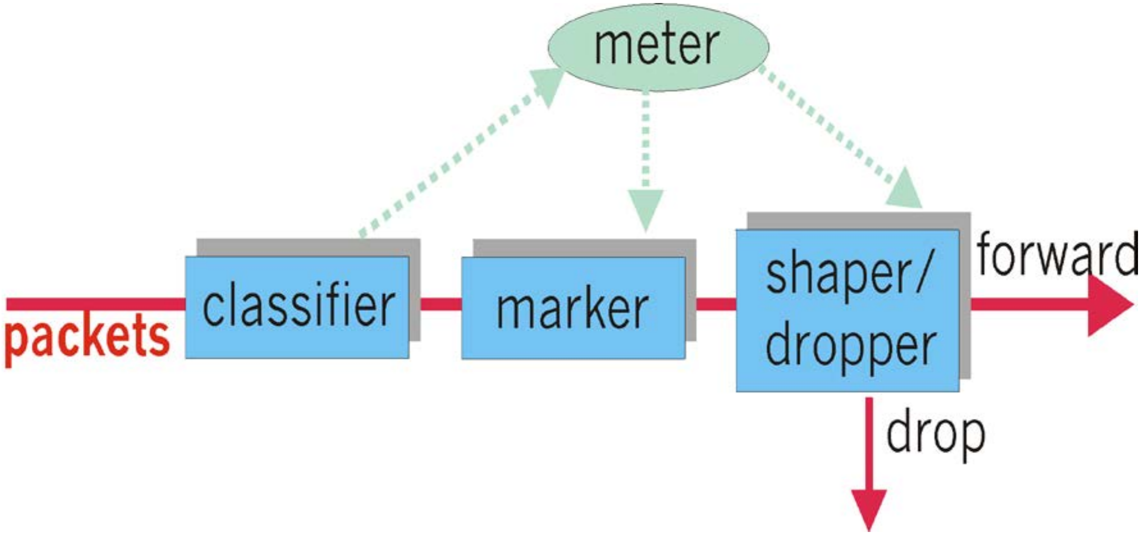

Forwarding (PHB)

- 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

Quality of Service

- Differentiated Services
  - The Expedited Forwarding (EF) PHB
    - One of the simplest PHBs to explain is known as “expedited forwarding” (EF). Packets marked for EF treatment should be forwarded by the router with minimal delay and loss.
    - The only way that a router can guarantee this to all EF packets is if the arrival rate of EF packets at the router is strictly limited to be less than the rate at which the router can forward EF packets.

Quality of Service

- Differentiated Services
  - The Assured Forwarding (AF) PHB
    - The “assured forwarding” (AF) PHB has its roots in an approach known as “RED with In and Out” (RIO) or “Weighted RED,” both of which are enhancements to the basic RED algorithm.
    - For our two classes of traffic, we have two separate drop probability curves. RIO calls the two classes “in” and “out” for reasons that will become clear shortly.
    - Because the “out” curve has a lower MinThreshold than the “in” curve, it is clear that, under low levels of congestion, only packets marked “out” will be discarded by the RED algorithm. If the congestion becomes more serious, a higher percentage of “out” packets are dropped, and then if the average queue length exceeds Min_{in}, RED starts to drop “in” packets as well.

Quality of Service

- Differentiated Services
  - The Assured Forwarding (AF) PHB

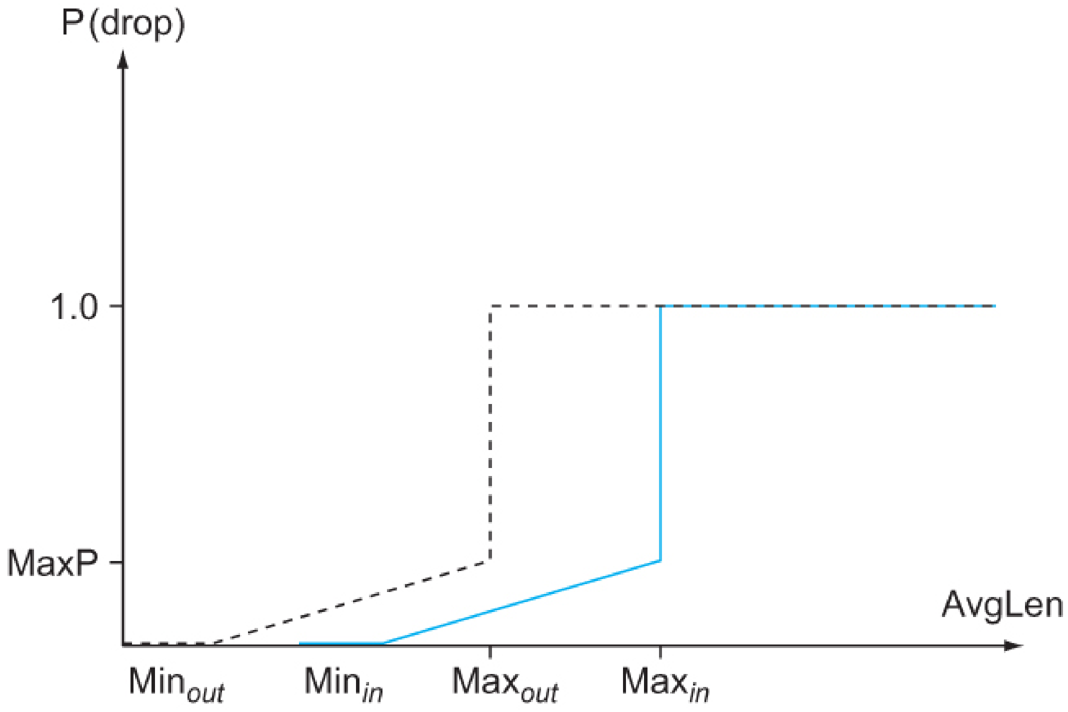

RED with In and Out drop probabilities