Master Course
Computer Networks
IN2097

Prof. Dr.-Ing. Georg Carle
Christian Grothoff, Ph.D.
Stephan Günther

Chair for Network Architectures and Services
Department of Computer Science
Technische Universität München
http://www.net.in.tum.de

Chapter:
Quality of Service Support

Chapter outline – Quality-of-Service Support

- Providing multiple classes of service
- Providing QoS guarantees
- Signalling for QoS

Multiple Classes of Service: Scenario

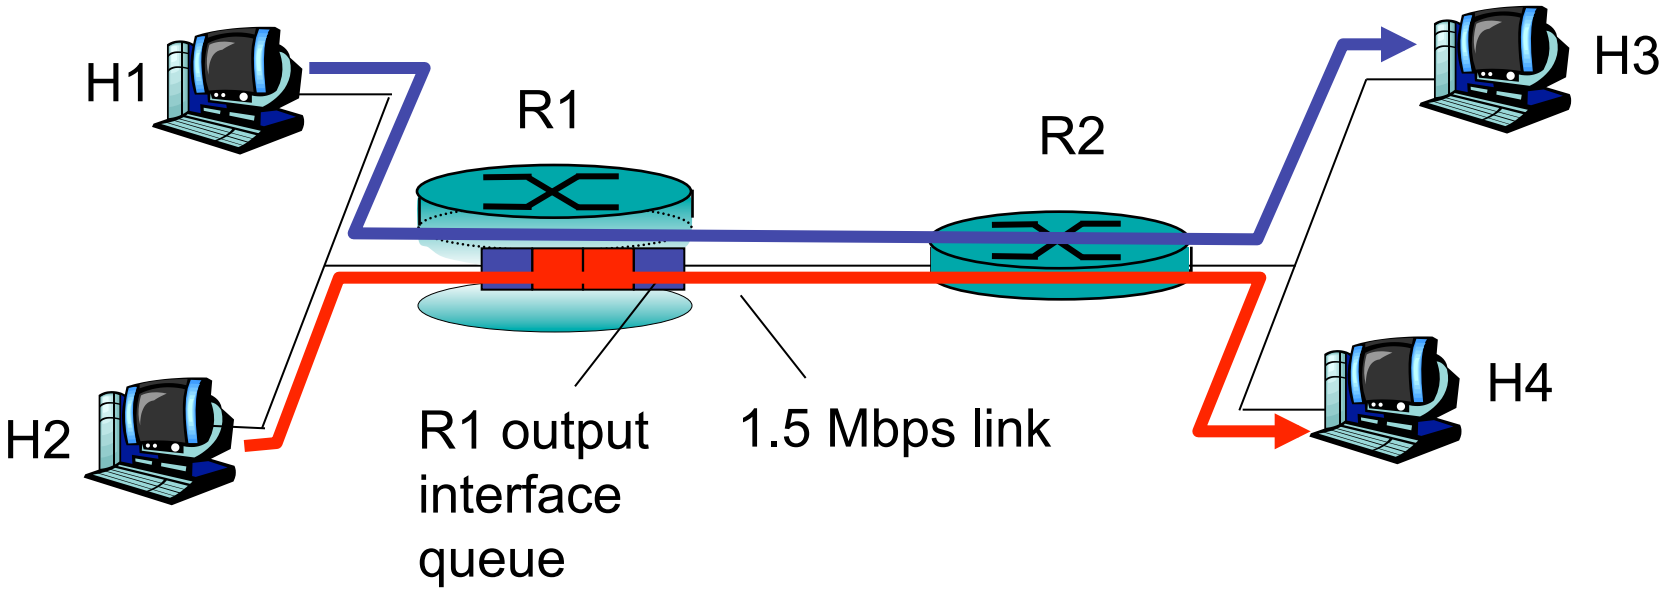

Scenario 1: mixed FTP and audio

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

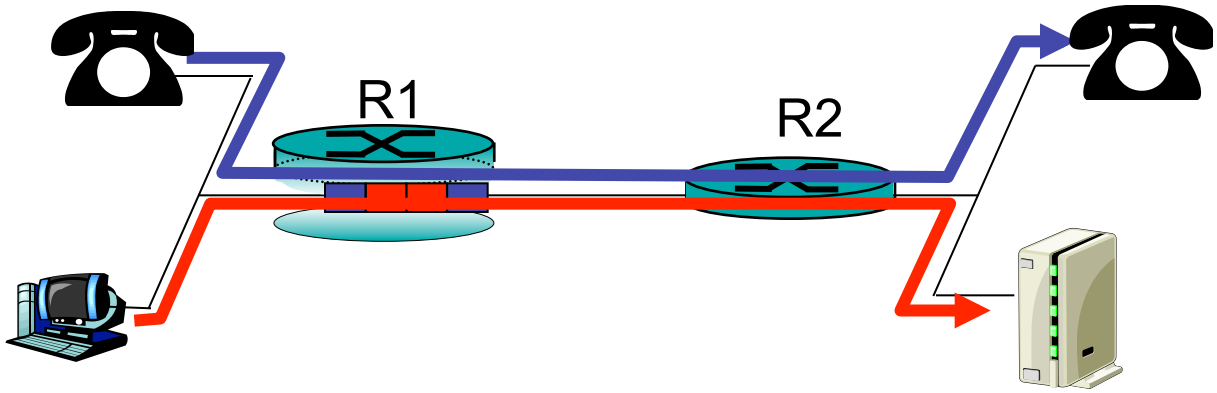

Principle 1

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)

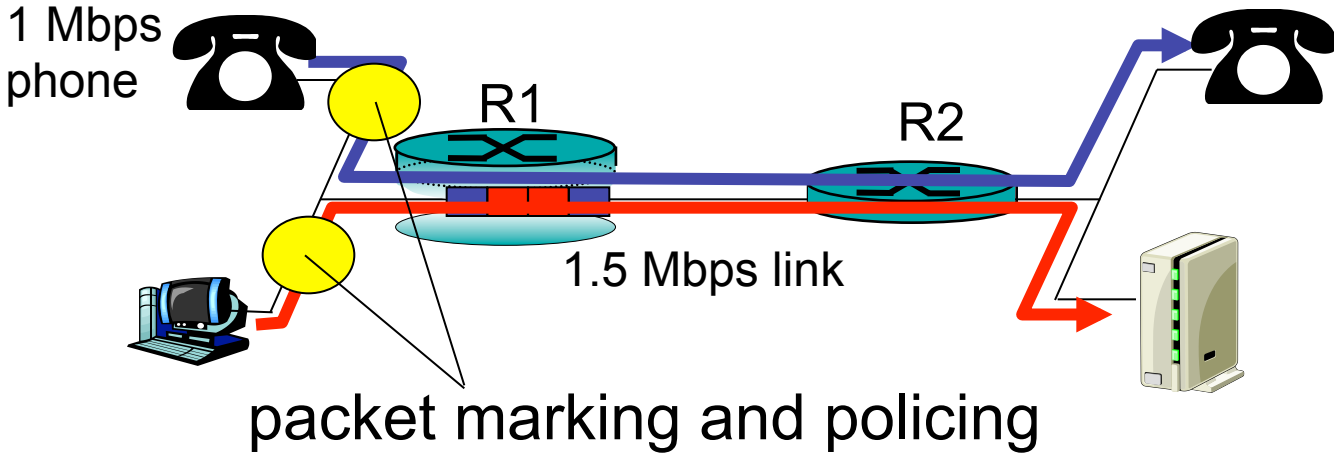

Principle 2

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

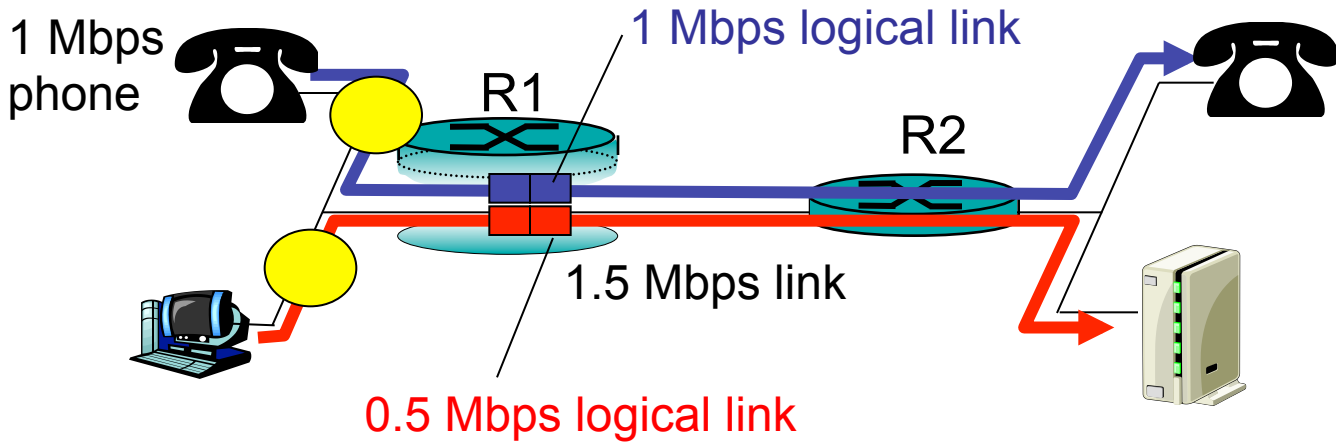

Principle 3

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

⇨ allocate sharable bandwidth to logical link
issue: sharing policy (scheduling, discarding) to be defined

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
- 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

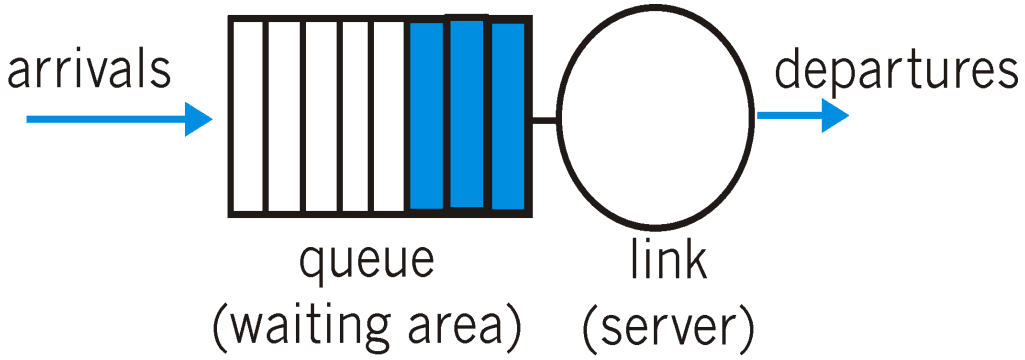

Scheduling Policies: more

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..

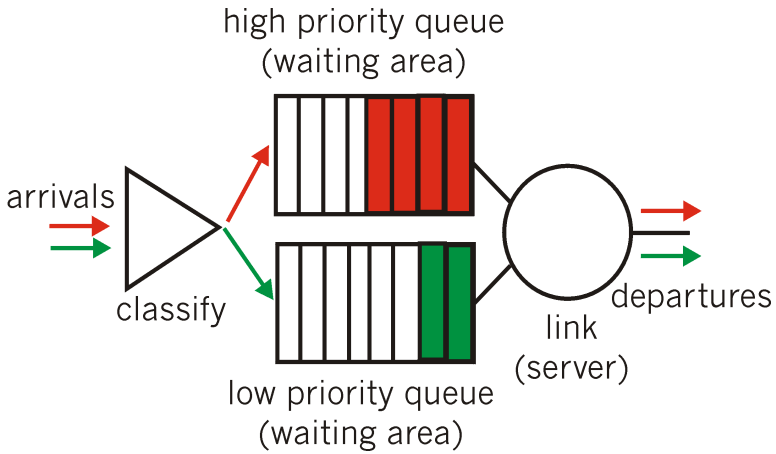

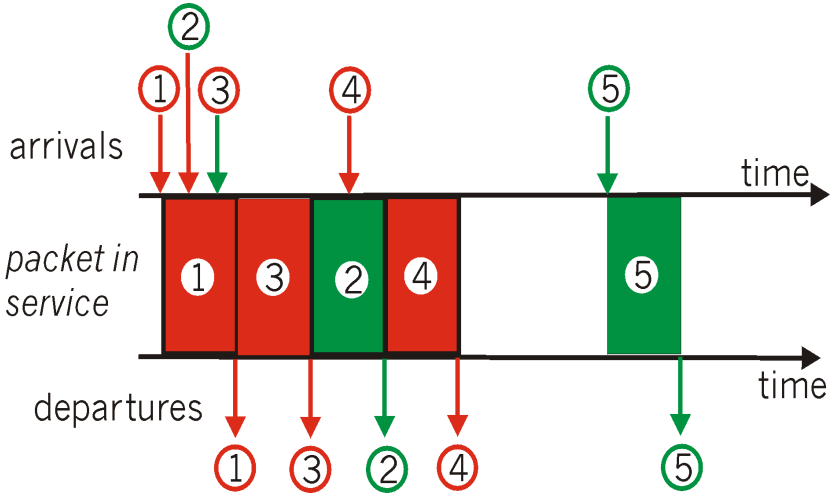

Scheduling Policies: more

Round robin scheduling:

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

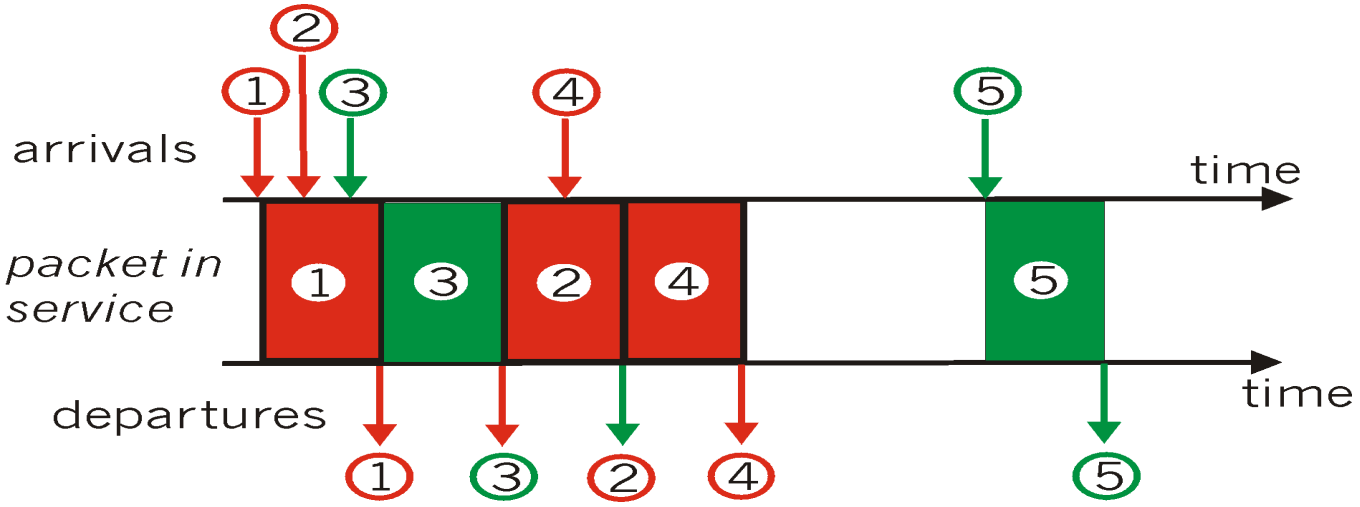

Scheduling Policies: more

Weighted Fair Queuing:

- each class gets weighted amount of service in each cycle
- when all classes have queued packets, class i will receive a bandwidth ratio of $w_i/\Sigma w_j$
  (for all j classes that have packets in queue)
- ill-behaved traffic classes only punish themselves
- Parekh and Galagher showed that combination with
  → *leaky bucket policing* allows end-to-end delay bounds

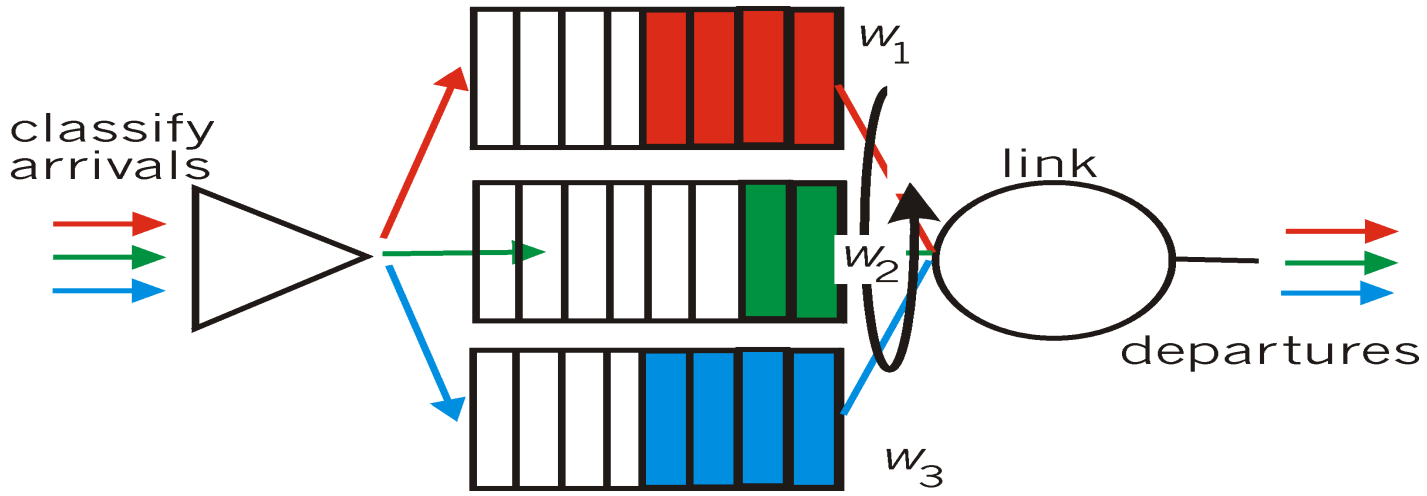

WFQ and Packets

- Fluid Flow System (Processor Sharing)
  - work-conserving scheduling without scheduling overhead
  - fluid flow: conceptually bit-by-bit weighted round robin

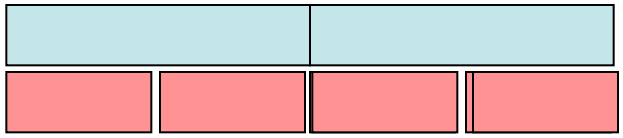

- Packet-by-Packet scheduling
  - approach: use finishing time of packet in fluid system as priority for choosing next packet

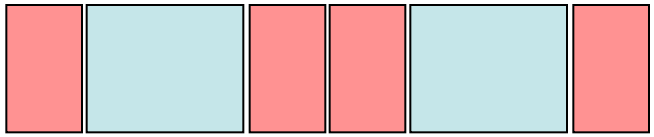

  - issue: arrival of packets of new flow

    ⇨virtual time (round number) finishing time

Policing Mechanisms

Goal: limit traffic to not exceed declared parameters

Three commonly used criteria:

- *(Long term) Average Rate:* how many packets 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 packets per min (ppm) avg.;
  1500 pps (90000 ppm) peak rate
- *(Max.) Burst Size:* max. number of packets sent consecutively

Policing Mechanisms

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

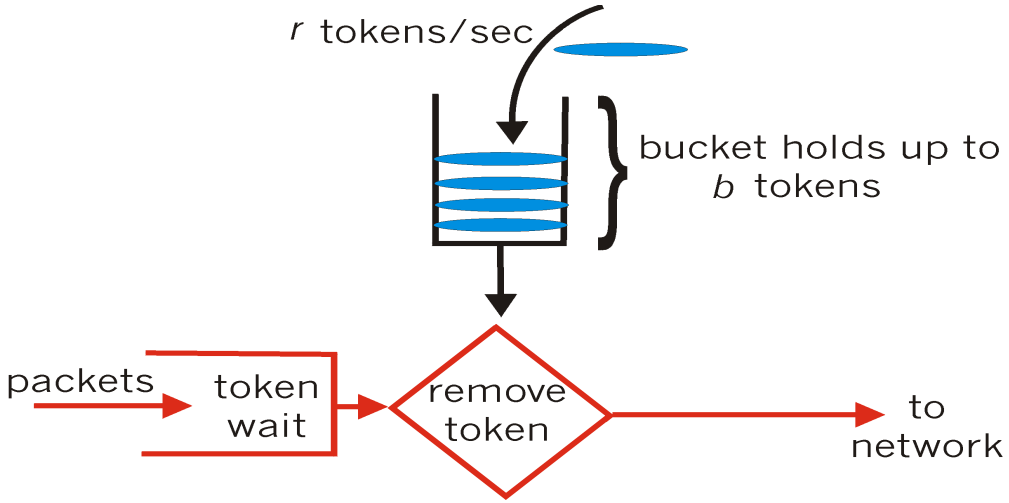

- bucket can hold b tokens ⇒ limits maximum burst size
- tokens generated at rate *r token/sec* unless bucket full
- *over interval of length t: number of packets admitted less than or equal to (r t + b).*

Policing Mechanisms (more)

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

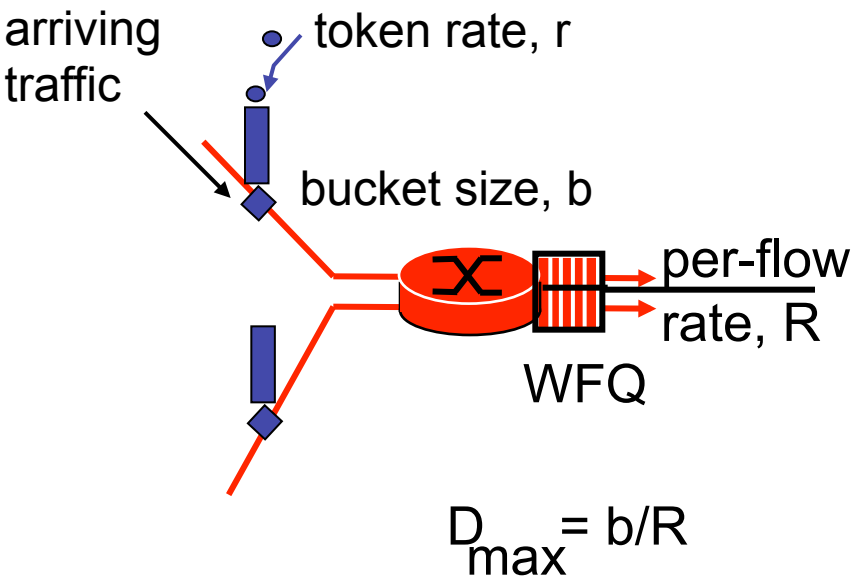

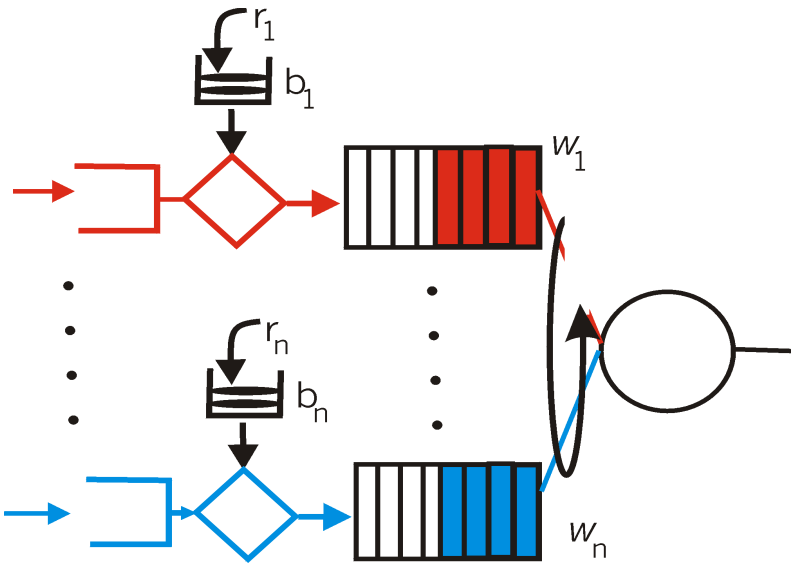

IETF Differentiated Services

- 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)
  - in contrast to IETF Integrated Services: signaling, maintaining per-flow router state difficult with large number of flows
- Don’t define define service classes, provide functional components to build service classes

Diffserv Architecture

Edge router:

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

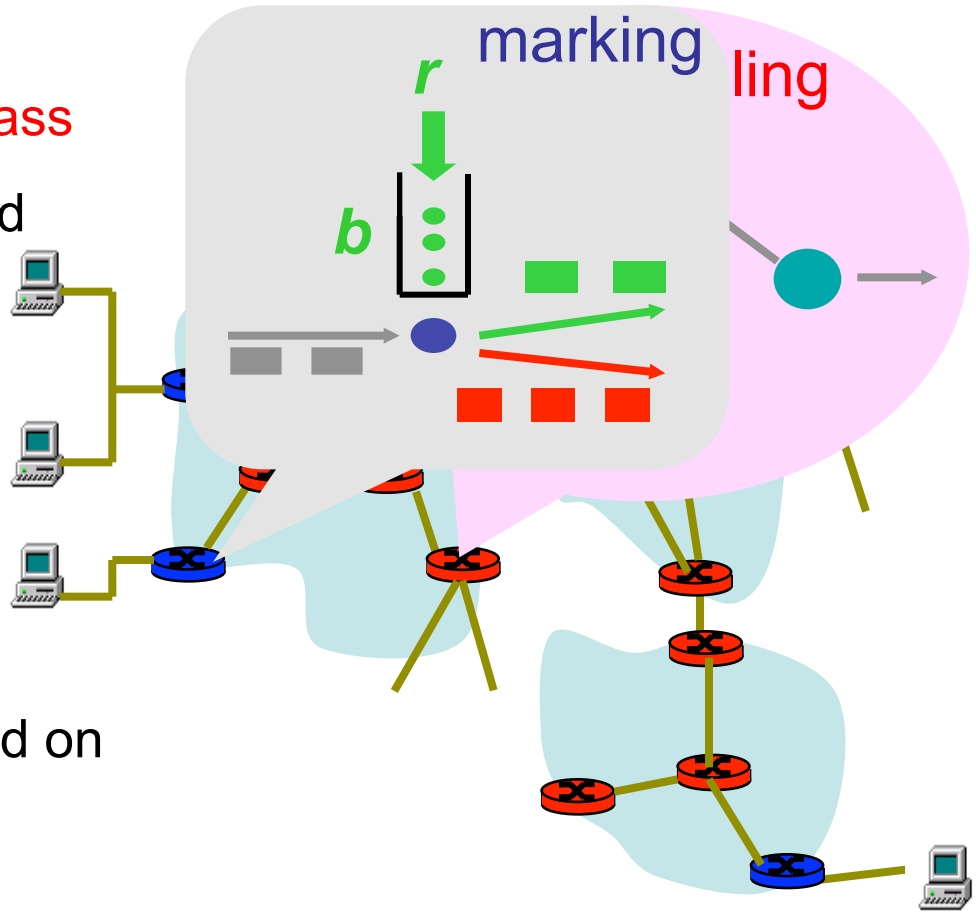

Core router:

- 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

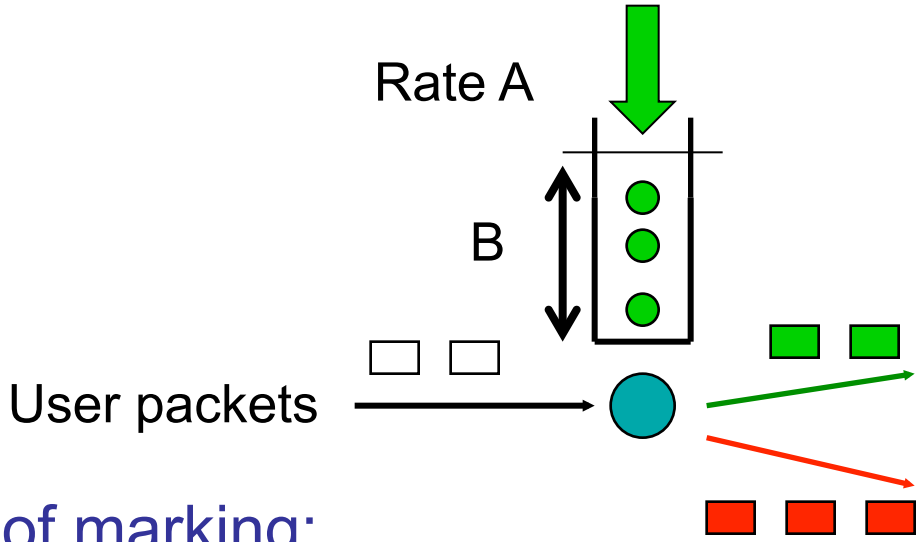

Possible usage of marking:

- 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 can be used for congestion notification: Explicit Congestion Notification (ECN), RFC 3168

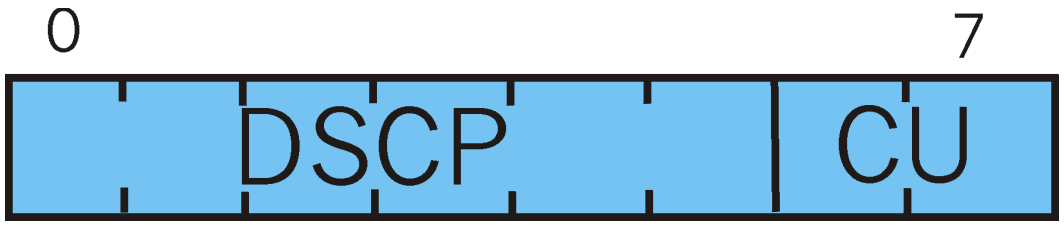

Classification and Conditioning

May be desirable to limit traffic injection rate of some class:

- user declares traffic profile (e.g., rate, burst size)
- traffic metered, shaped or dropped if non-conforming

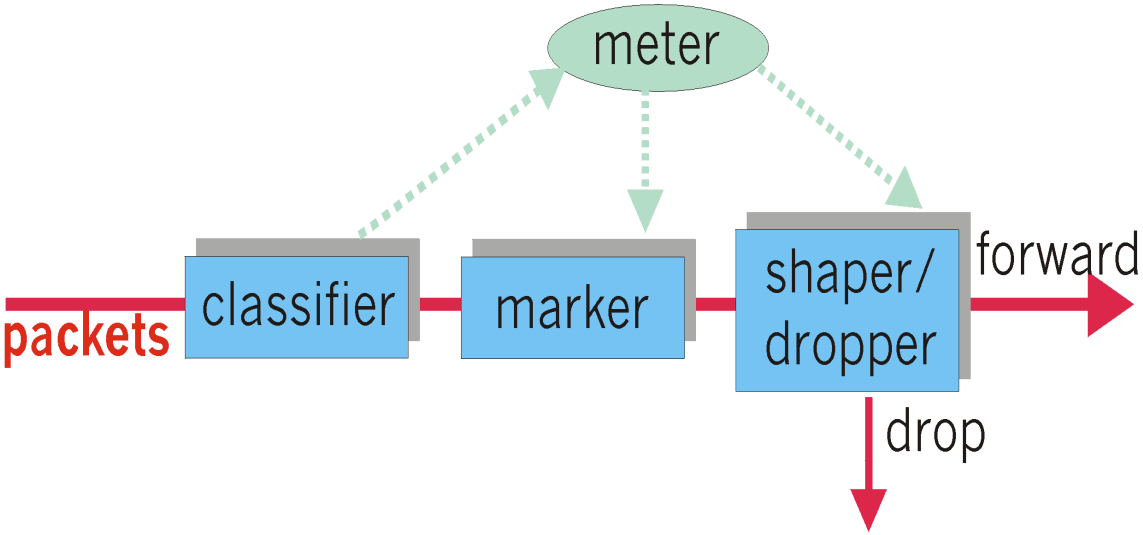

Forwarding (PHB)

- 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: packet departure rate of a class equals or exceeds specified rate
  - logical link with a minimum guaranteed rate
- Assured Forwarding: e.g. 4 classes of traffic
  - each class guaranteed minimum amount of bandwidth and a minimum of buffering
  - packets each class have one of three possible drop preferences; in case of congestion routers discard packets based on drop preference values

Assured Forwarding DiffServ Code Points

- Assured Forwarding behavior definition
  - RFC 2597 - Juha Heinanen, Fred Baker, Walter Weiss, John Wroclawski: Assured Forwarding PHB Group
    - Recommended Codepoints: c.f. table below
  - RFC 3260 - Dan Grossman: New Terminology and Clarifications for Diffserv

RFC 2597 Assured Forwarding (AF) Recommended Code Points

	Class 1 (lowest)	Class 2	Class 3	Class 4 (highest)
Low Drop	001010	010010	011010	100010
Med Drop	001100	010100	011100	100100
High Drop	001110	010110	011110	100110