Differentiated Services (DiffServ) in the Internet

N. C. State University

CSC557 ♦ Multimedia Computing and Networking

Fall 2001

Lecture # 22

"Roadmap" for Multimedia Networking

1. Introduction
   - why QoS?
   - what are the problems?

Lecture 17

2. Basic operations
   - jitter buffers (at hosts)
   - task scheduling (at hosts)
   - packet shaping (at hosts)
   - packet dropping (at routers)
   - packet scheduling (at routers)

Lecture 16

Lecture 18

Lecture 19

3. Types of service
   - Integrated Services (IntServ) and Resource Reservation Protocol (RSVP)
   - **Differentiated Services (DiffServ)**

Lectures 20 and 21

Today's Lecture

"Roadmap" (cont'd)

4. Application-level feedback and control
   - Real-time Protocol (RTP), Real-time Control Protocol (RTCP)
   - Real-time Streaming Protocol (RTSP)
5. Application signaling and device control
   - Session Announcement Protocol (SAP)
   - Session Description Protocol (SDP)
   - Session Initiation Protocol (SIP)
   - Media Gateway Control Protocol (MGCP)
6. Routing
   - Multi-protocol Label Switching (MPLS)
   - multicasting

(Reminder of Problems, + Solutions)

- ✓ Less-than-ideal average delays and loss rates
- Variations in traffic loads in the network
- ✓ TCP's congestion control
- ✓ Retransmission-based error recovery
- 5. Simplistic routing algorithms
- ✓ "Burstiness" or variability of a single traffic source
  - ✓ peak rate, average rate, maximum burst size

- Schedule packet transmission carefully to control delays/losses
- Manage traffic loads through reservations and admission control (IntServ/RSVP)+DiffServ
- Improved packet dropping policies
- Use with jitter buffers, and to restore anchor frames of video
- …
- "Shape" traffic to reduce variability

Introduction

- QoS: some levels of network service are better than others
- Intserv: QoS managed on a per-flow basis
  - per-flow state stored in all routers in the path
  - per-flow scheduling, policing, shaping
  - hop-by-hop reservations → signaling overhead, complexity

Another Approach: Airline Seating!

- First-class, business-class, and coach-class
  - Coach class (best-effort) carries bulk of traffic
  - business/first-class: small amount of traffic, but quite profitable
- Differentiated services
  - not expected to comprise all traffic in the Internet
  - goal: healthy service offerings and profit opportunities

Another Approach: HOV Highway Lanes!

- One lane reserved for exclusive use of High-Occupancy Vehicles (HOVs) during rush hour
  - outside rush hour, other vehicles may also use the HOV lane
- HOVs experience little congestion, less delay
- Work Conservation law: total queueing delay remains constant over all cars
  - improved service for HOVs means worse service for everyone else

DiffServ Goals

- Ease of use and generality
  - limited flexibility
- Simple processing in core routers
  - push complexity to network edge
- Accommodate mostly best-effort traffic

Core vs. Access Networks

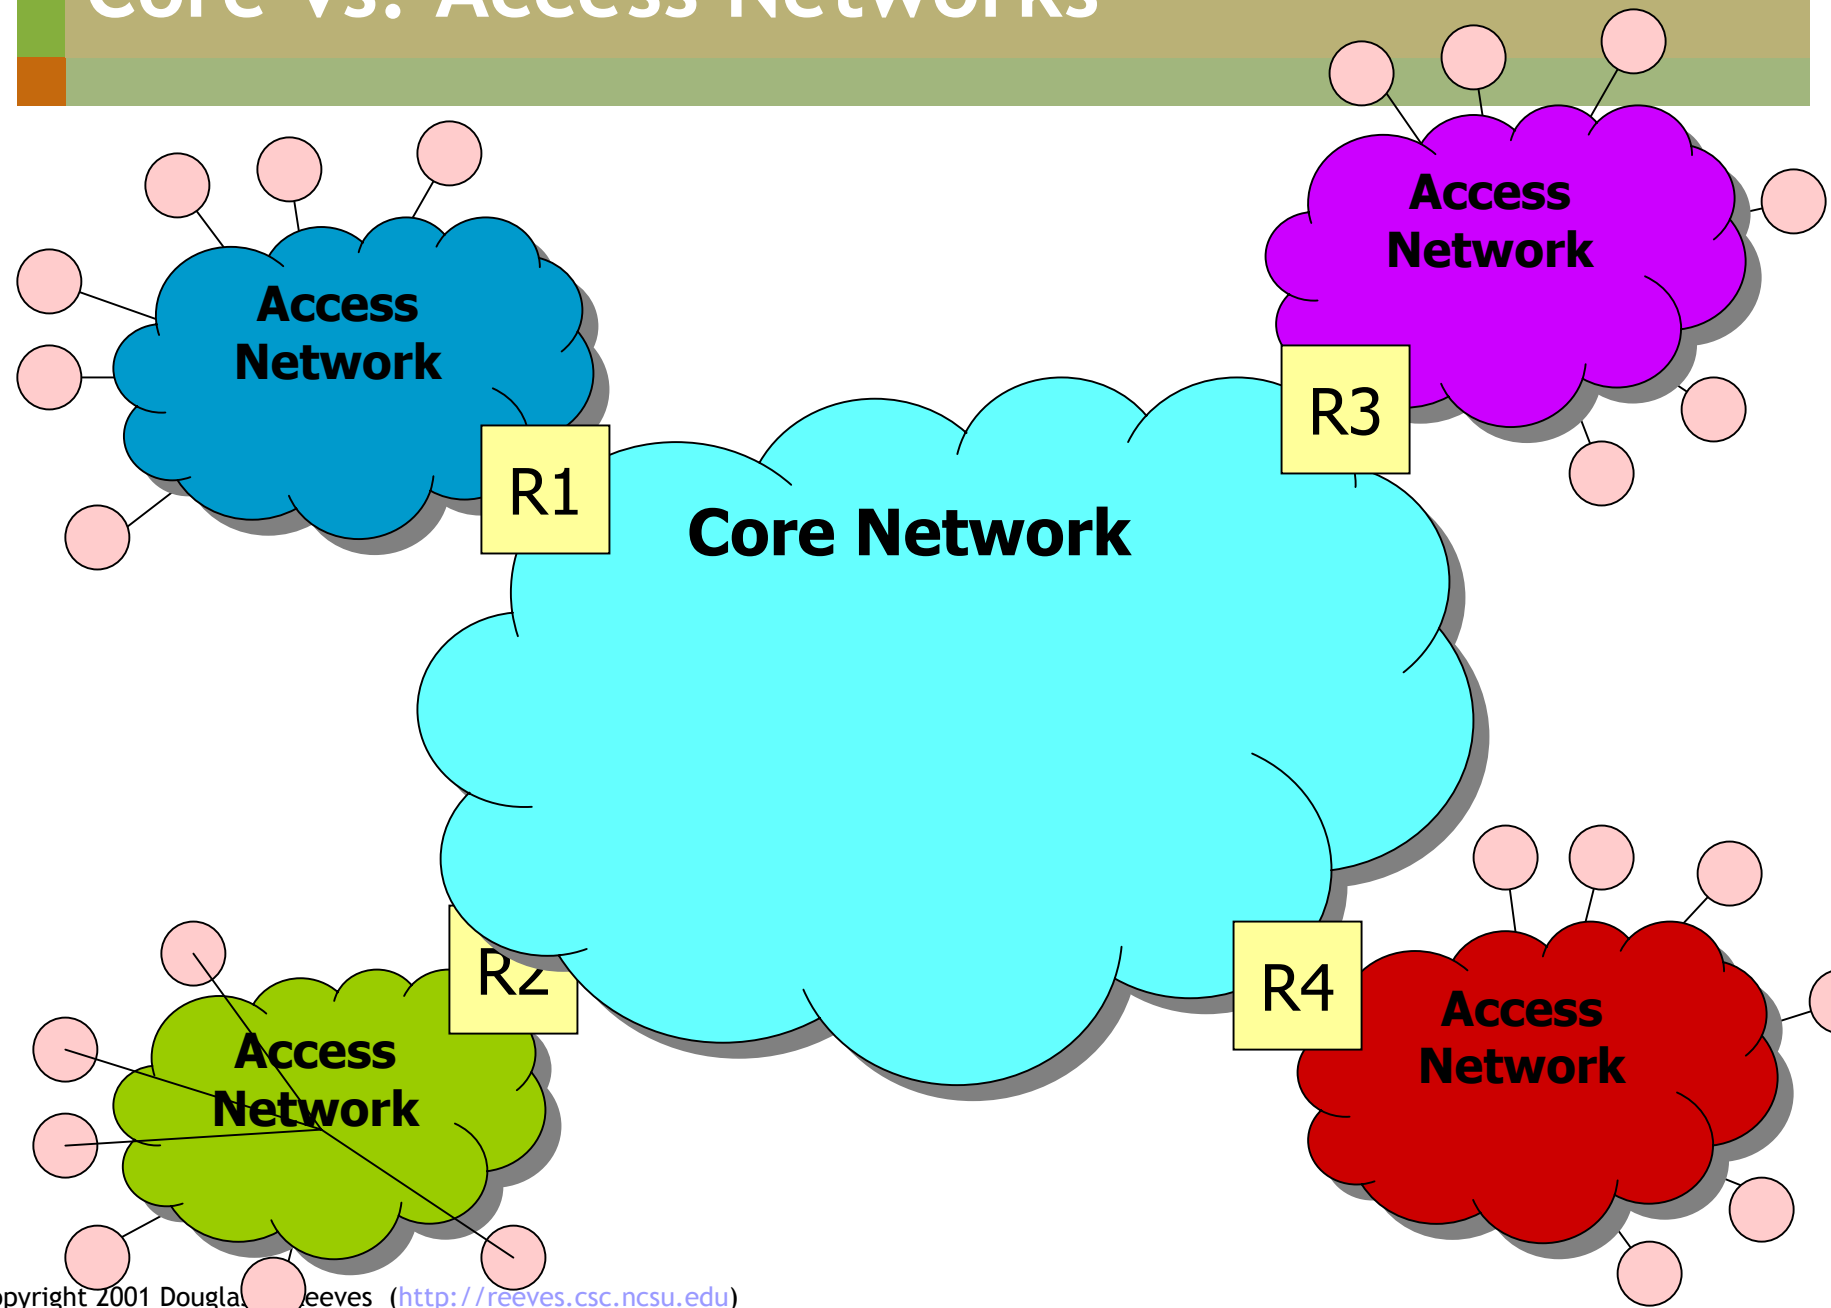

Architecture

- Neither…
  - best-effort (connectionless) model
  - guaranteed service (connection-oriented) model
- In-between: service guarantees for aggregations of flows
  - implemented in the core network only

	IntServ	DiffServ
Focus is on…	Users, applications	Network owners / administrators
Standardizes…	End-to-end service	Per-hop service (behavior)

Diffserv Codepoint (DSCP)

- Field in the IP header specifying the class of service the packet is to receive
  - replaces the (8-bit) TOS field

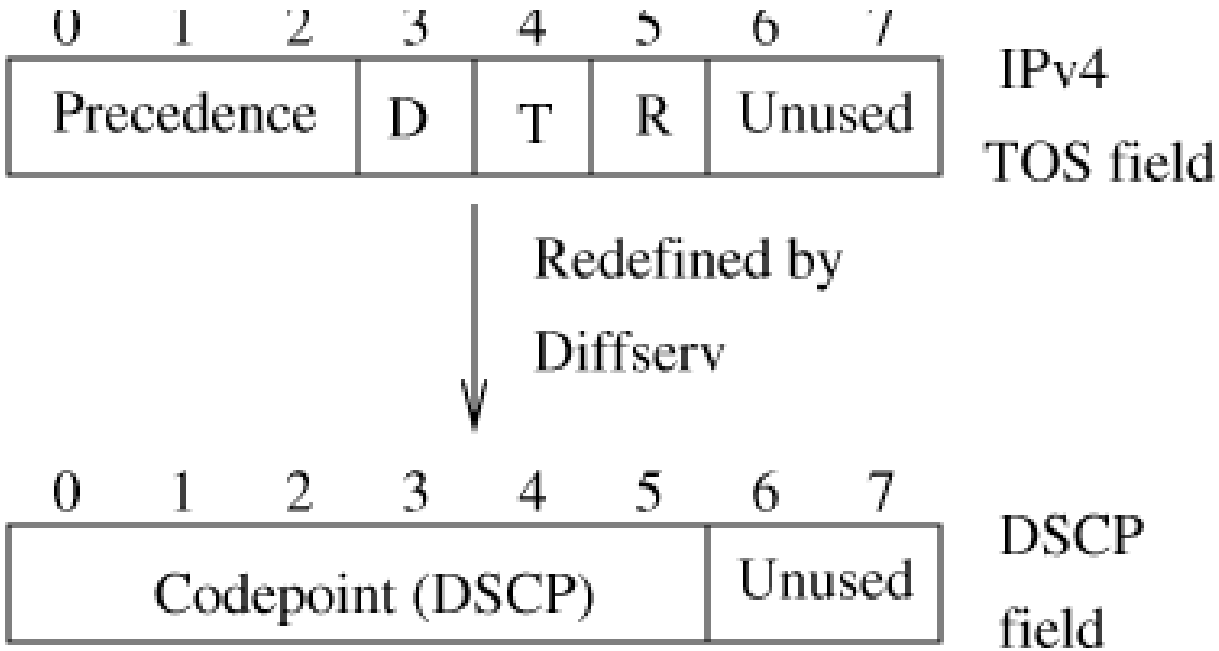

Per-Hop Behavior (PHB)

- *Behavior aggregate (BA)* = a collection of flows with the same Diffserv codepoint (DSCP) , and sharing a link
- *Per-hop behavior (PHB)* = the observable forwarding behavior of a BA, in terms of…
  - absolute QoS metric (on delay and/or loss), or
  - relative QoS metric (e.g., lower packet loss than another BA)
- DSCP maps to underlying service definition

Diffserv Domains

- A *DiffServ Domain* is …
  - a contiguous set of routers implementing DiffServ capabilities, and
  - supporting a common set of provisioning policies and PHBs
- Boundary routers (ingress + egress) and interior routers

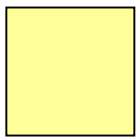

= border router, meters, polices, shapes, marks, and implements PHB

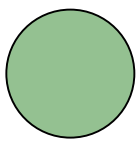

= interior router, implements PHB

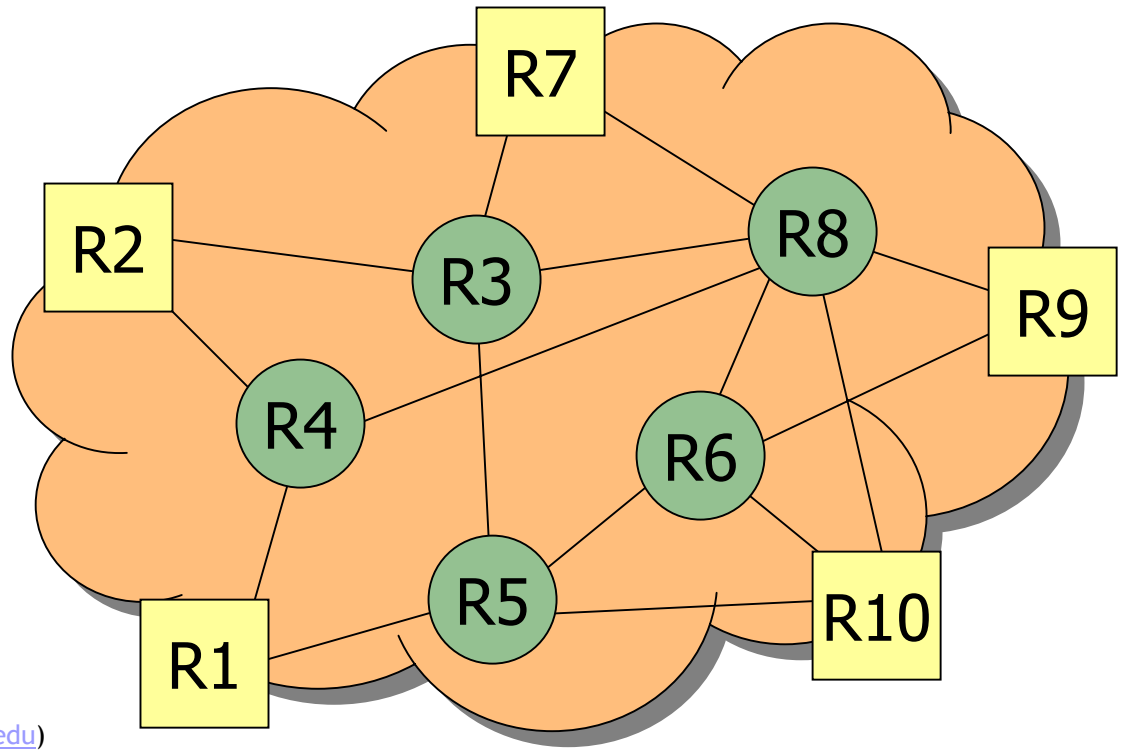

Best Effort and Class Selector PHBs

- RFC 2474: Best Effort (BE) PHB
  - default PHB when no other agreement is in place
  - corresponds to today's best-effort service
  - routers must reserve some minimal resources to avoid starvation
- Class Selector (CS) PHB: for backwards compatibility with routers using IPv4's precedence field occupied by the DSCP

Edge Router Operations

- Classifies/remarks traffic in a BA (sets the DSCP)
- Meters traffic in a BA
  - measures performance and arrival statistics
- Polices (shapes, drops) traffic in a BA

Expedited Forwarding (EF) PHB -- RFC 2598

- Rate guarantee: minimum bandwidth is guaranteed
- Qualitatively: low loss, delay, jitter
- The BA is policed by token bucket, excess traffic is dropped
- The EF PHB can be used to provide the Virtual Wire end-to-end (*Per-Domain Behavior*, or *PDB*) service
- EF PHB can be implemented using various scheduling disciplines
  - including static priority and weighted round-robin

The *Virtual Wire* PDB

- Appears to endpoints as dedicated, point-to-point connection of a minimum bandwidth
  - guaranteed bandwidth service with negligible queueing delay
- User promises not to exceed peak rate (must be policed)
- Similar to telephone service
  - customer expects capacity to be there when receiver lifted
  - may be unused a good deal of the time
- Use: video broadcasts, voice-over-IP, virtual leased lines
  - Virtual Wire traffic will comprise a small fraction of total traffic, but will be priced much higher
  - unused capacity will be used by best-effort traffic

Virtual Wire PDB (cont'd)

- Access router...
  - filters packets entering the network
  - sets the DSCP on packets matching service specification
  - performs traffic shaping on the flow to smooth traffic bursts
- No changes in hosts required
  - (DiffServ-aware hosts: do own shaping, entry routers merely police)
- Forwarding treatment
  - two levels of priority queueing
  - send EF packets first
  - all Virtual Wire flows aggregated in a single high priority queue

EF Policing and Shaping

- Virtual Wire bandwidth allocations must not be oversubscribed
- Policing and shaping needed only at...
  - access routers, on a per-flow basis
  - DS border routers, on behavior aggregate
- Policing based on simple token bucket
  - packets exceeding the rate are dropped

Assured Forwarding (AF) PHB -- RFC 2597

- For some applications, Virtual Wire service too restrictive (hard limits) and too conservative
- *Assured service*: traffic will be forwarded with "high probability" as long as within expected profile
  - assured service traffic may exceed its profile, but excess traffic not given same assurance level
- Firmness of guarantee depends on how well individual links are provisioned for bursts of AF packets

AF PHB (cont'd)

- Four AF classes; in each class, packets…
  - are allocated a certain amount of resources (buffers, bandwidth)
- May be marked with one of three drop precedence values
  - specifies relative packet importance during periods of congestion
- Level of forwarding assurance of an IP packet depends on...
  1. amount of resources reserved for this AF class
  2. current load of this AF class
  3. drop precedence of packet
- Resource provisioning, admission control: difficult?!
- Network must not reorder packets in the same "microflow" whether they are in or out of profile

"Olympic" Service

- Three service classes: bronze, silver, gold
- Gold load < silver load << bronze load
- Each class can be mapped to one AF PHB class
- Packets within each class may be given low, medium, or high drop precedence
  - map to corresponding AF drop precedence level

Markers to Implement the Two Services

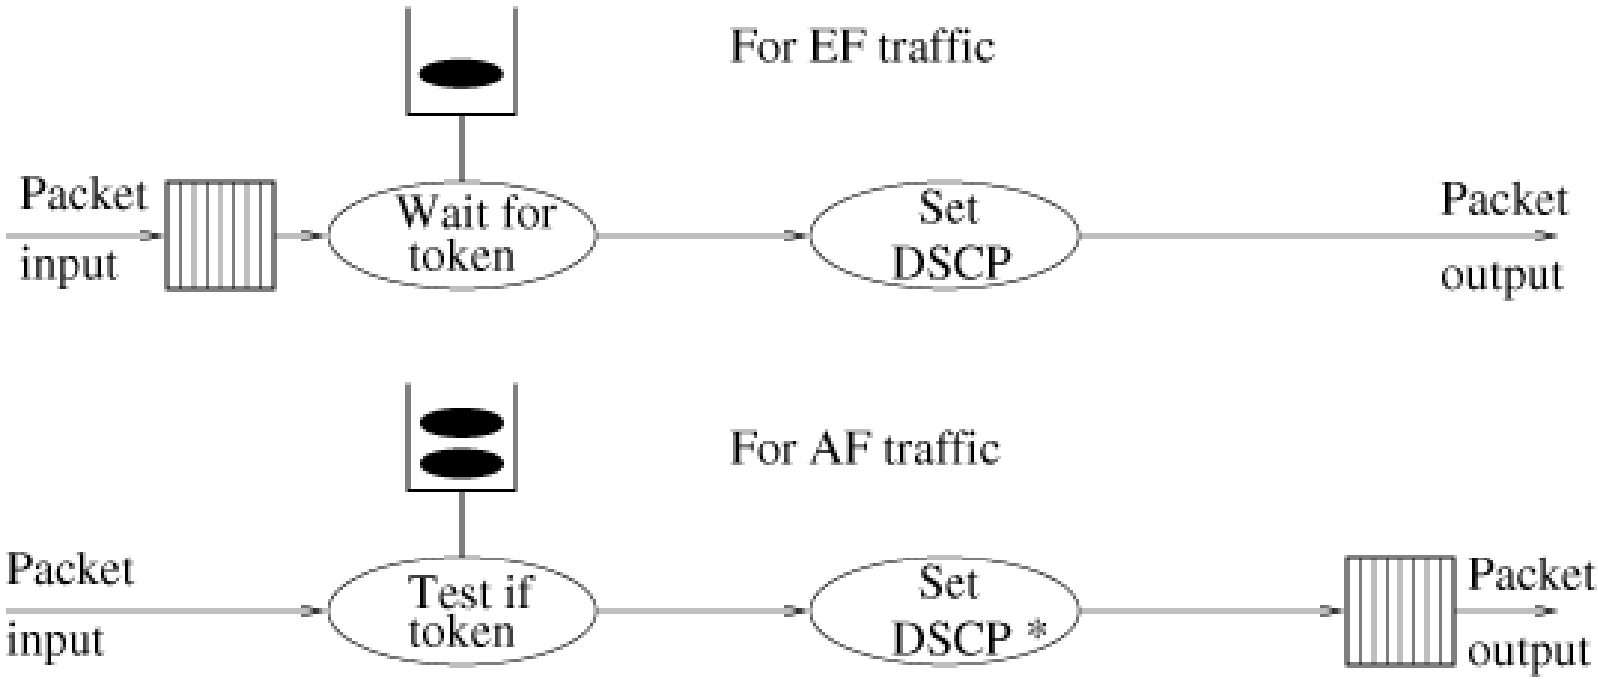

* appropriate AF class and precedence level

Token Bucket

- Virtual Wire service
  - bucket only allowed to fill to maximum packet size
  - only real queue the flow sees in the network
- Assured rate service: bucket may fill to configured burst parameter

Router Output Interface

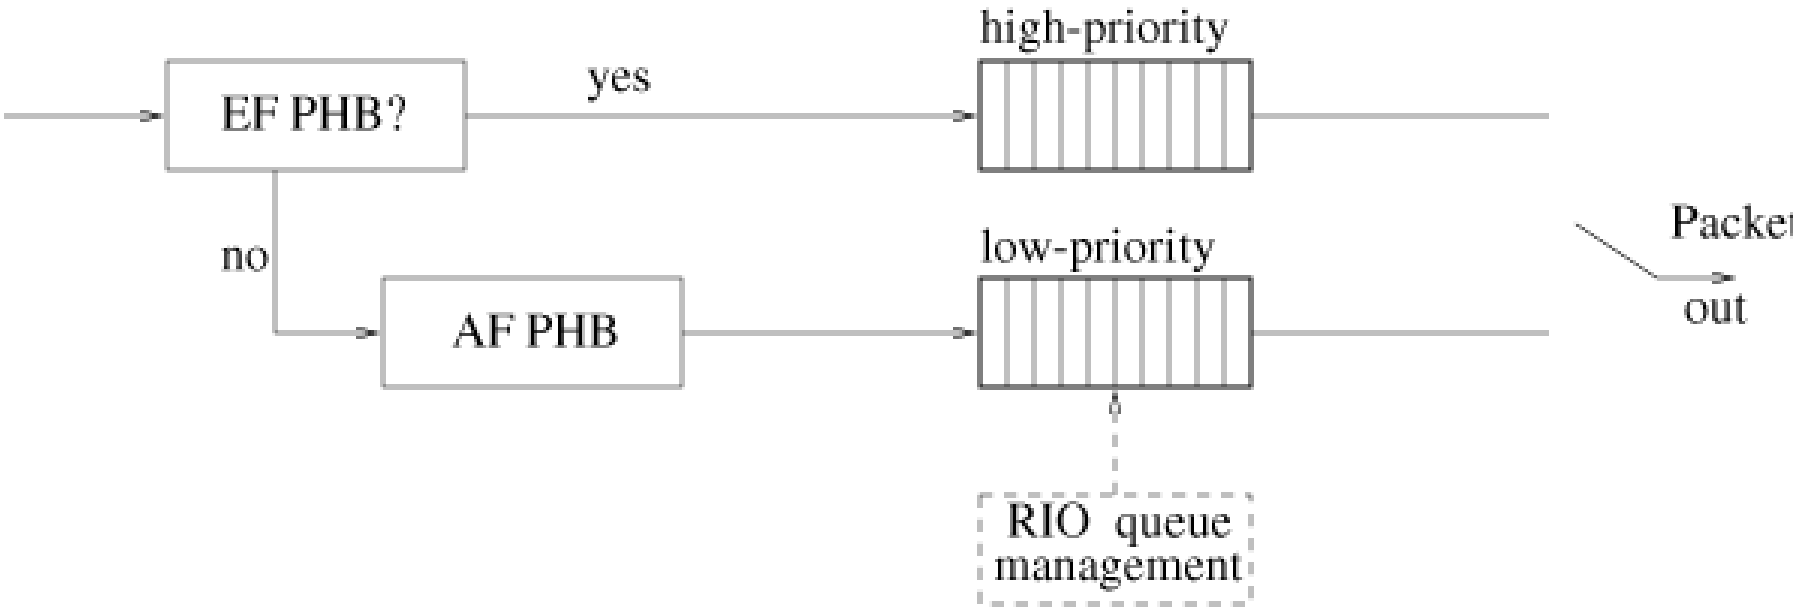

- RED-based RIO mechanism on lower priority queue
  - out-of-profile traffic: low dropping threshold
  - in-profile traffic: higher dropping threshold
- Re-marking of packets to higher drop precedence levels also possible for out-of-profile traffic

Border Router Input Interface Profile Meters

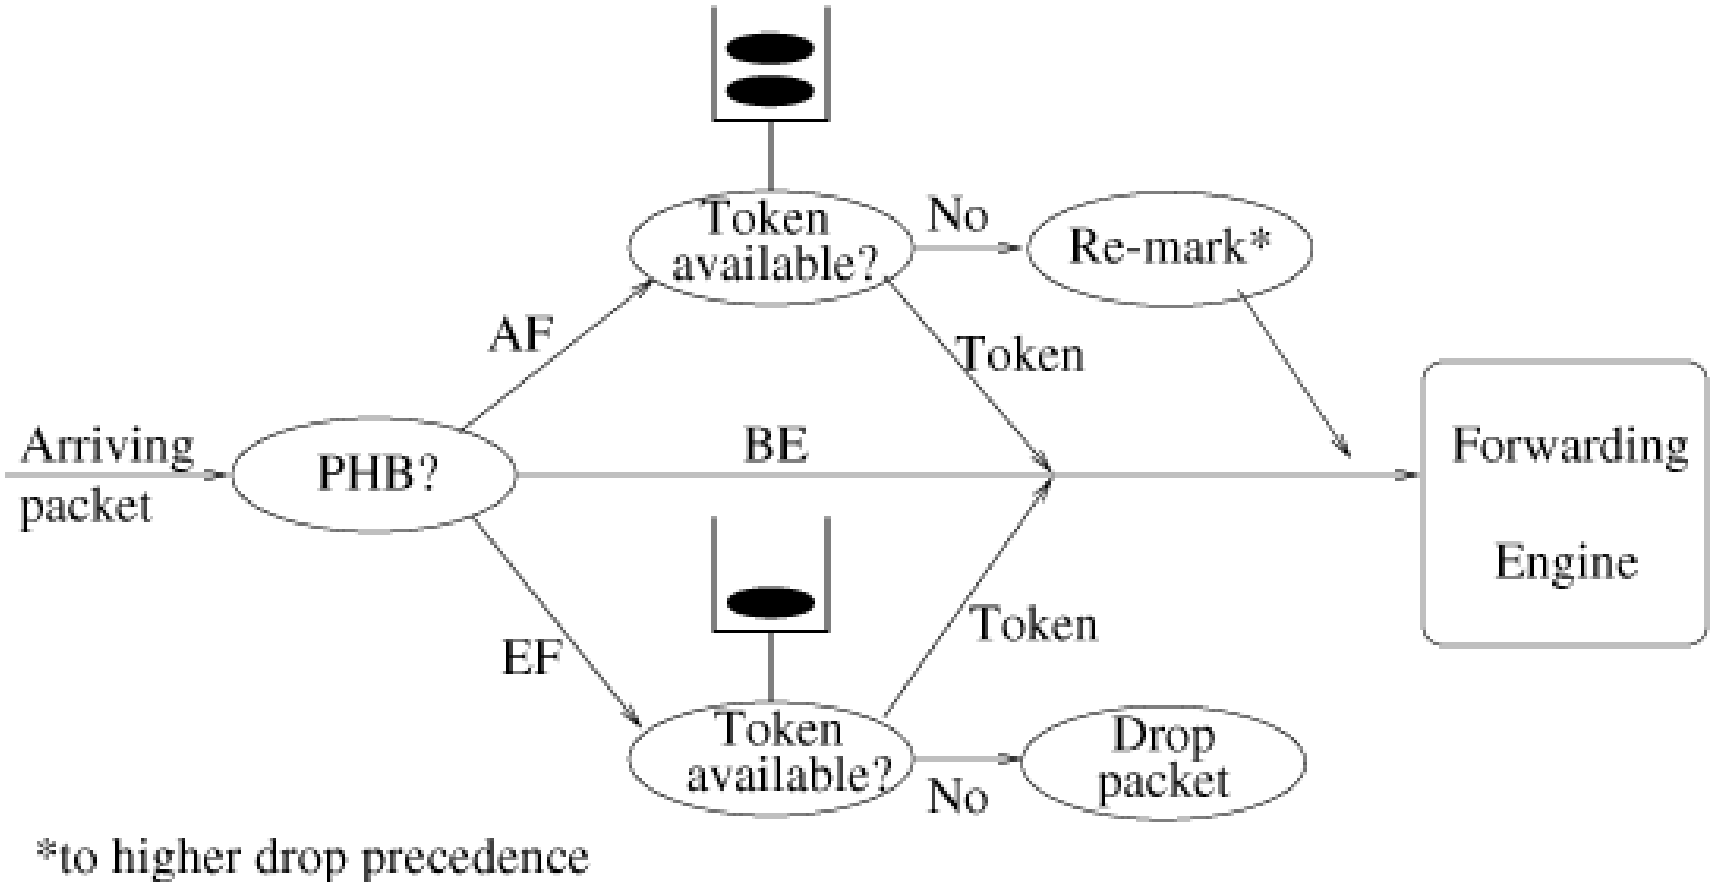

Discussion

- Mechanisms required to initiate requests and configure profiles (signaling function, can use RSVP, SNMP, or other)
- Greatest burden of flow matching and shaping will be at access routers, where the speeds and buffering required should be less that those required deeper in the network
- Shaping would be applied to an aggregation of all the Virtual Wire flows that exit the domain, not to each flow individually

Traffic Allocation

- Provisioning: sets up static levels of Diffserv traffic
  - a "mail-order" house may purchase a fixed amount of bandwidth
  - in/out of web-site (based on hit rates in previous quarter)
- Call set-up: creates an allocation of Diffserv traffic for duration of a single flow
  - makes it possible to respond to particular events (e.g., a broadcast by the company's CEO)
  - required to meet a pre-committed service when source/destination is allowed to be anywhere in the Internet

Traffic Allocation (cont'd)

- Both allocation types important
- Static allocation expected to be useful in initial deployment of differentiated services
- Without a per-call dynamic set up: "paying for bits you don't use" vs. "paying for service one expects to be available all the time"
  - Compare with flat telephone charges

Bandwidth Brokers

- Agents that essentially perform admission control
  - configured with organizational policies
  - keep track of current allocation of Diffserv traffic
  - interpret new requests to mark traffic in light of the policies and current allocation

Bandwidth Brokers (cont'd)

- Needed because neither individuals nor routers have info necessary to decide which packets are more important to the organization
- BBs make it possible for bandwidth allocation to follow organizational hierarchies
- Keep state on an administrative domain basis, confine per-flow state to access routers
- BBs only need to establish relationships of limited trust with peers in adjacent domains
  - Compare with setting flow specifications in routers throughout a path (RSVP)

Bandwidth Brokers (cont'd)

- One BB per administrative domain (initially)
- Responsibilities
  - parcel out the domain's Diffserv traffic allocations and set up access routers → security
  - manage messages sent across boundaries to adjacent regions' BBs (to negotiate and configure a service class and rate)
- Must maintain policy database on "who can do what when"

Setting of Profiles in Leaf Routers

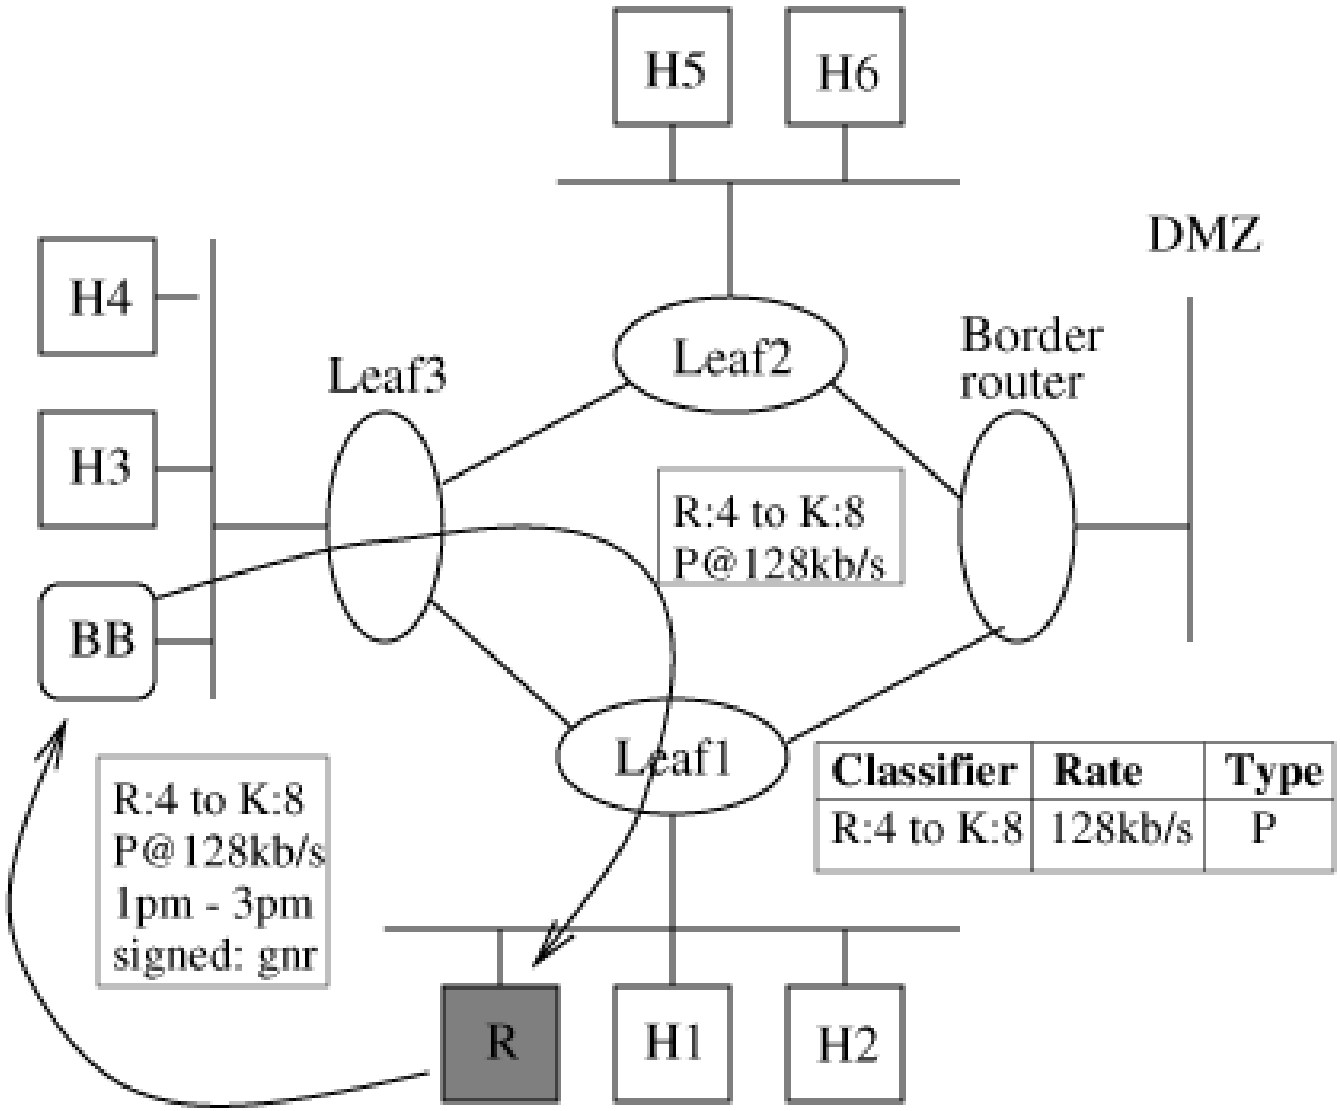

End-to-End Example with Static Allocation

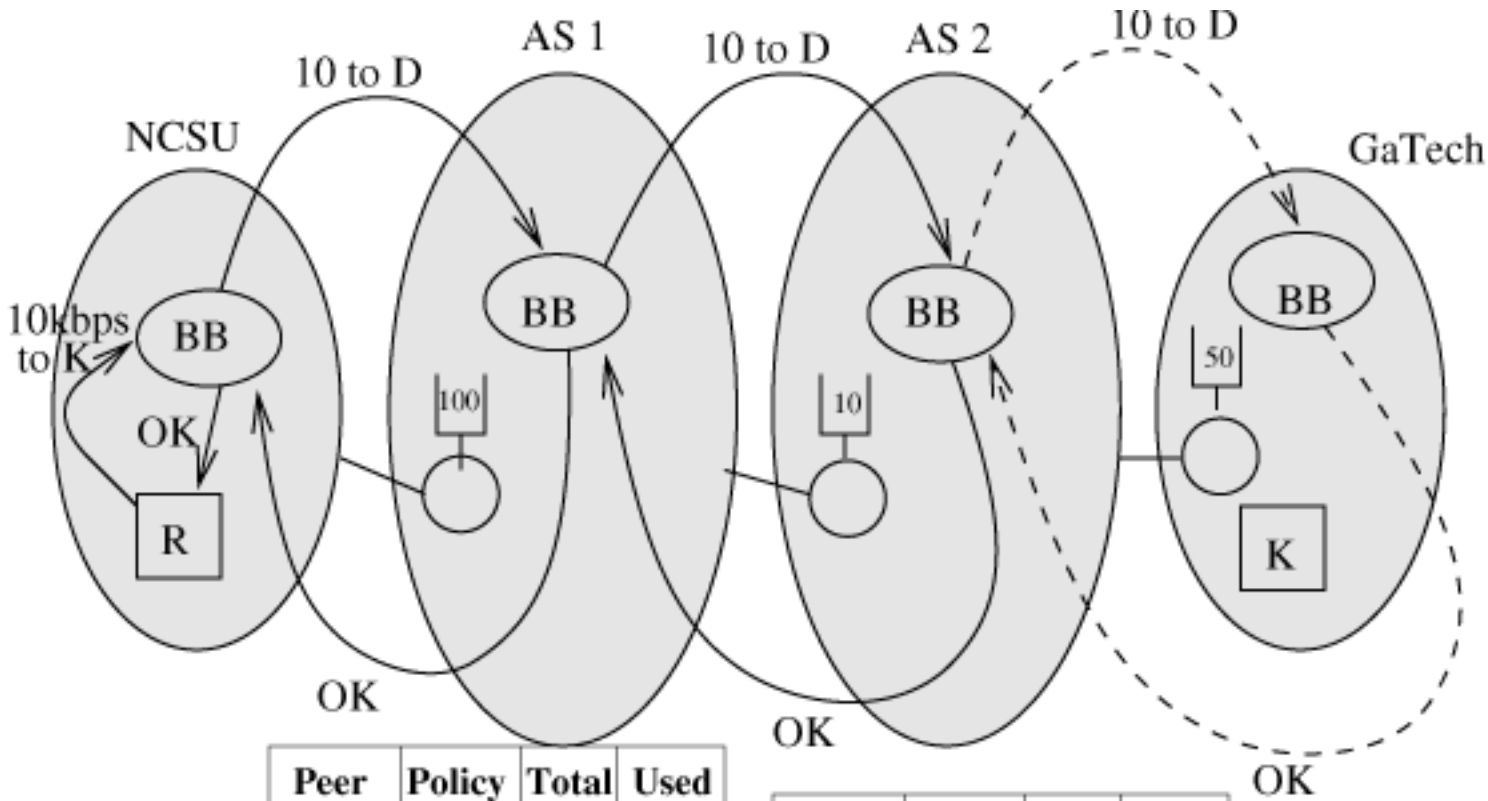

Peer	Policy	Total	Used
AS 2	ask	10	10
NCSU	<50ok	100	30

Peer	Policy	Total	Used
AS 1	ask	0	0
GaTech	<50ok	50	20

End-to-End Example (cont'd)

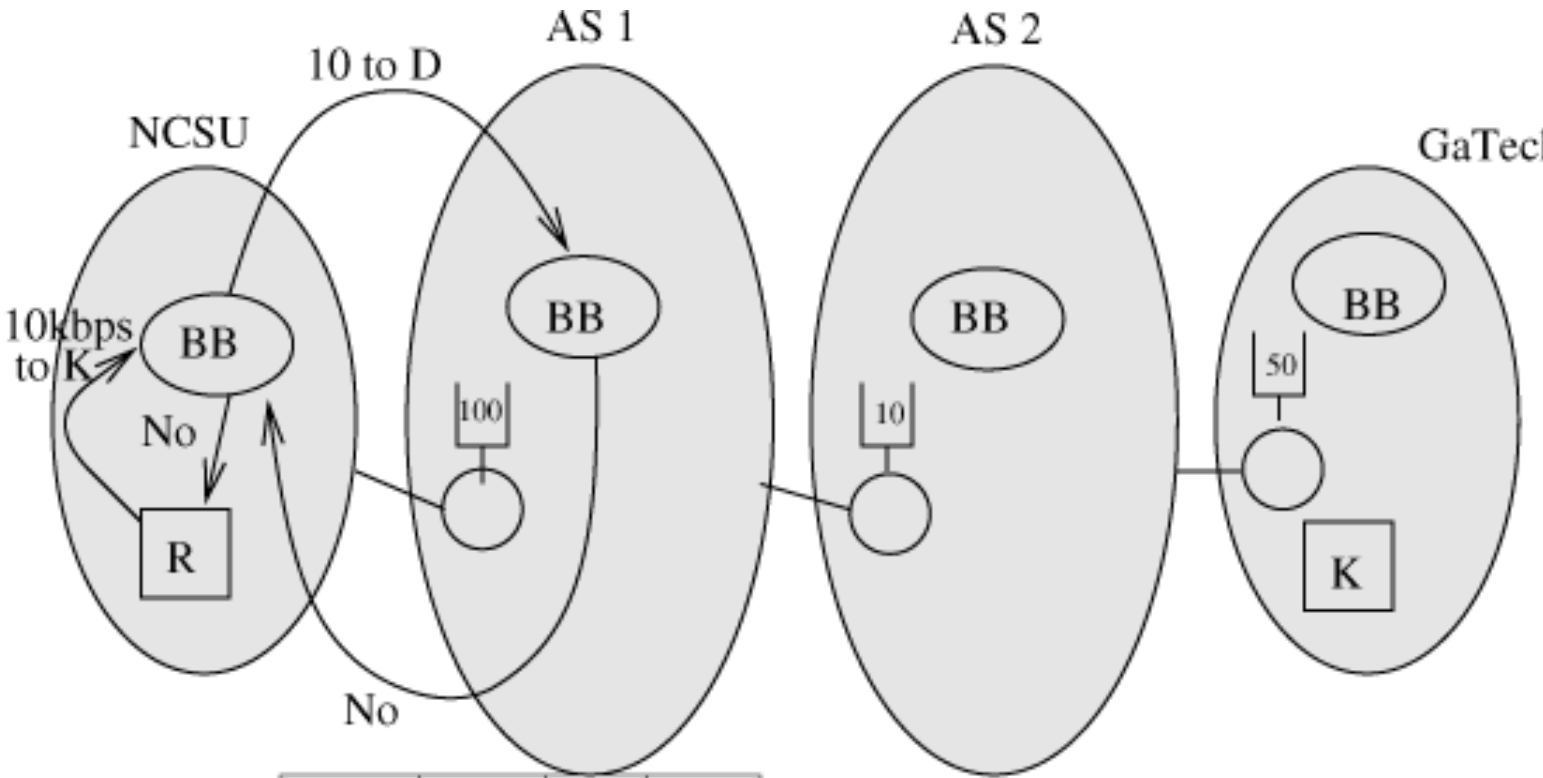

Peer	Policy	Total	Used
AS 2	ask	10	10
NCSU	<50ok	100	30

Peer	Policy	Total	Used
AS 1	ask	0	0
GaTech	<50ok	50	20

Sources of Information

- Web
  - RFC2465: An Architecture for Differentiated Services
  - RFC2597: Assured Forwarding PHB Group
  - RFC2587: An Expedited Forwarding PHB