

Computer Communications - CSC 551

Copyright © William C. Cheng

1

CS551

Differentiated Services (DiffServ)

Bill Cheng

<http://merlot.usc.edu/csc551-f12>

Computer Communications - CSC 551

Copyright © William C. Cheng

2

Key Ideas

- Traffic classes instead of flows
- Forwarding behaviors instead of end-to-end service guarantees
- Tune applications to network services rather than network services to applications
- Discrete v.s. continuous space
- No resource reservation
- Somewhere between Best Effort and IntServ

Computer Communications - CSC 551

Copyright © William C. Cheng

3

Types of Service

- Premium service: (type P)
 - admitted based on peak rate
 - conservative, virtual wire services
 - unused premium goes to best effort (subsidy!)
- Assured service: (type A)
 - based on expected capacity usage profiles
 - traffic unlikely to be dropped if user maintains profile
 - out-of-profile traffic marked
 - traffic is *in* or *out* (of profile)
 - police traffic to keep *in* within limits
 - use provisioning and/or admission control to limit amount of *in*
 - preferentially drop *out* traffic

Computer Communications - CSC 551

Copyright © William C. Cheng

4

Differences With Integrated Services

- No need for reservations: just mark packets
- Packet marking done at administrative boundaries before injecting packets into network
- Significant savings in signaling, much simpler overall

Computer Communications - CSC 551

Copyright © William C. Cheng

5

Service Differentiation

- Analogy:
 - airline service, first class, coach, various restrictions on coach as a function of payment
 - Best-effort expected to make up bulk of traffic, but revenue from first class important to economic base (will pay for more plentiful bandwidth overall)
 - Not motivated by real-time but by economics and assurances

Computer Communications - CSC 551

Copyright © William C. Cheng

6

A Two-bit Differentiated Services Architecture

[Nichols99a]

Bill Cheng

<http://merlot.usc.edu/csc551-f12>

Copyright © William C. Cheng

7

Premium vs. Assured Forwarding Behaviors

- **Premium** packets receive virtual circuit type of treatment
- **Assured** packets receive "better than best effort" type of treatment
- Appropriate for intolerant and rigid applications
- Appropriate for adaptive applications

Computer Communications - CSC1 551

Copyright © William C. Cheng

8

2-bit Differentiated Service

- Precedence field encodes P & A type packets
- P packets are BW limited, shaped and queued at higher priority than ordinary best effort
- A packets treated preferentially with respect to dropping probability in the normal queue
- Leaf and border routers have input and output tasks - other routers just output

Computer Communications - CSC1 551

Copyright © William C. Cheng

11

First-hop Router Input Functionality

➤ **Markers:** service class, rate, permissible burst size

Computer Communications - CSC1 551

Copyright © William C. Cheng

12

Marker Function in Routers

- First-hop routers have traffic profiles - they classify packets based on packet header
- if no profile present, pass as best effort
- if profile is for A:
- mark packets *in or out* (in-profile packets with A, forward others unmarked)
- if profile is for P:
- delay or drop out-of-profile packets to *shape* into profile
- Routers at administrative boundaries
- must make sure traffic exchange agreements are met

Computer Communications - CSC1 551

Copyright © William C. Cheng

9

Premium Service

- User sends within profile, network commits to delivery with requested profile
- Simple forwarding: classify packet in one of two queues, use priority
- Shaping at trust boundaries only, using token bucket
- Signaling, admission control may get more elaborate, but still not end-to-end

Computer Communications - CSC1 551

Copyright © William C. Cheng

10

Premium Traffic Flow

Computer Communications - CSC1 551

Copyright © William C. Cheng

Markers to Implement Two Different Services

• if profile is for P:

- Wait for token
- Drop on overflow
- Set P bit
- Packet output

• if profile is for A:

- Test if token
- Drop on overflow
- Set A bit
- Packet output

• mark packets *in or out* (in-profile packets with A, forward others unmarked)

- if no profile present, pass as best effort

Computer Communications - CSC1 551

Copyright © William C. Cheng

Output Forwarding

• 2 queues: P packets on higher priority queue

- Lower priority queue implements RED "in or out" scheme (Clark98a)

Computer Communications - CSC1 551

Copyright © William C. Cheng

Per-hop Behaviors (PHBs)

• Define behavior of individual routers rather than end-to-end services

- there may be much more services than behaviors
- Multiple behaviors - need more than one bit in the header
- Six bits from IP tos field are taken for DiffServ code points (DSCP)

Computer Communications - CSC1 551

Copyright © William C. Cheng

Markers to Implement Two Different Services

• delay or drop out-of-profile packets to *shape* into profile

- if profile is for P:
 - Wait for token
 - Drop on overflow
 - Set P bit
 - Packet output
- if profile is for A:
 - Test if token
 - Drop on overflow
 - Set A bit
 - Packet output

Computer Communications - CSC1 551

Copyright © William C. Cheng

Border Router Input Interface Profile Meters

• At border routers profile meters test marked flows:

- drop P packets out of profile
- unmark A packets

Computer Communications - CSC1 551

Copyright © William C. Cheng

Red With In or Out (RIO)

• For Assured Services

- Similar to RED, but with two separate probability curves
- Has two classes, "in" and "out" (or profile)
- "Out" class has lower Minthresh, so packets are dropped from this class first
- As avg queue length increases, "in" packets are dropped
- More drop probability curves (WRED)

Computer Communications - CSC1 551

- Context
 - = a lot of people had been looking at Integrated Services
 - o realizing that Integrated Services were unworkable
- Discussion
 - = nice, implementable schemes
 - = no real driver for this
 - o it's worthwhile to note that when you build a network and put in services, people might find a way to use those
 - o when the network is mature, need an economic driver to introduce new services

Discussions

- Resources to aggregated traffic, not flows
- Traffic policing at the edges, class forwarding in the core
- Define forwarding behaviors, not services
- Guarantees by provisioning and Service Level Agreements, not reservations
- Focus on single domain, not e2e (need Bbs for e2e)


Diffserv vs. Intserv Summary

- Where?
 - = static (long-term): done out-of-band
 - = dynamic: from leaf to Bandwidth Broker and from BB in one domain to another BB
- How?
 - = not clear, but maybe RSVP

Signaling

- Reservation based strategies can provide more varied QoS than feedback-based schemes
- Will this be the end of TCP?
 - = highly unlikely
 - = applications are established, heterogeneous networks, etc.
- Diffserv is middle ground: no intelligence v.s. high intelligence with Intserv
- Will we see a deployment? Jury is still out..

Open Issue: Inside or Outside the Network?



Signaling: Bbs