Quality of service

Overview

- Congestion control and avoidance
  - Prevent collapse of network
  - Approach full utilization of network
- Quality of Service (QoS)
  - Providing reliable service on a best-effort network
  - Types of applications
  - Integrated services (IntServ)
  - Differentiated services (DiffServ)

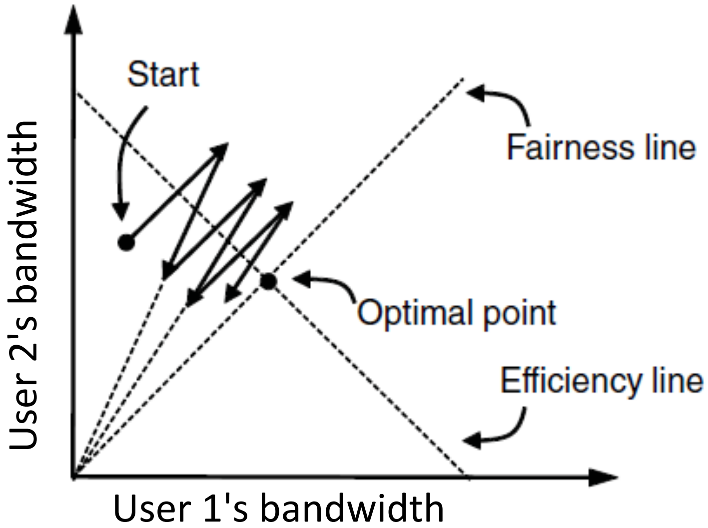
Start
Fairness line
Optimal point
Efficiency line
User 2's bandwidth
User 1's bandwidth

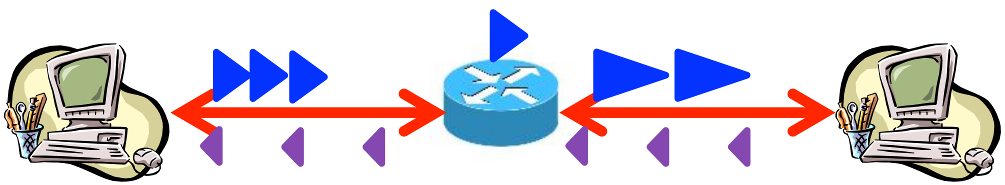

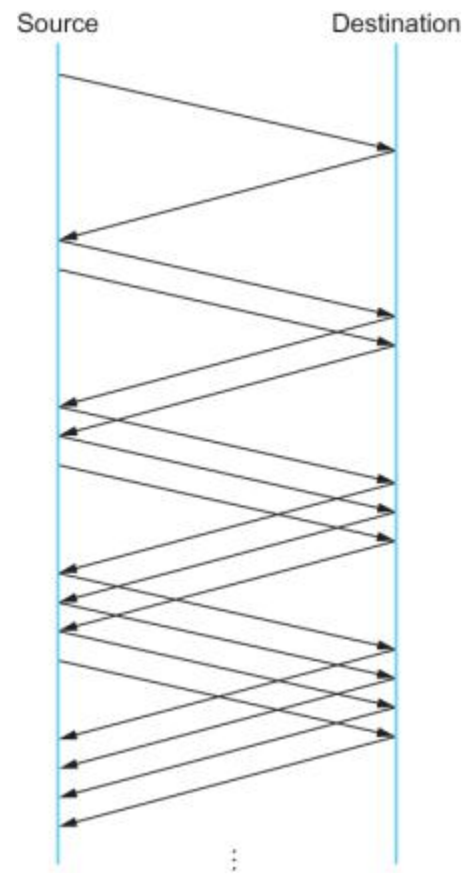
Source
Destination

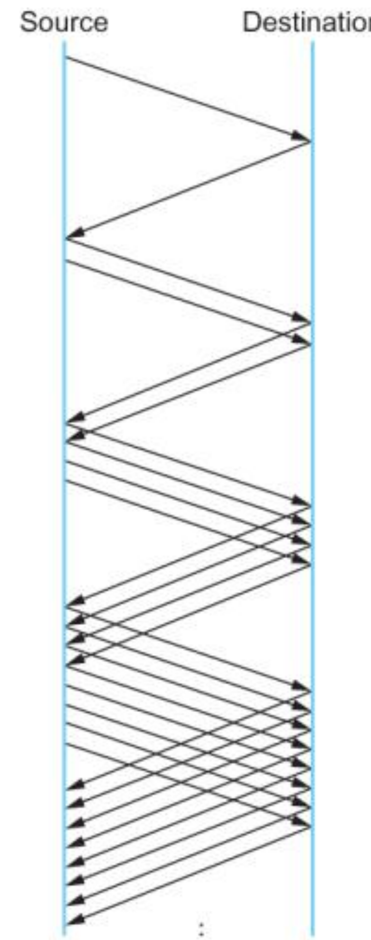
Source
Destination

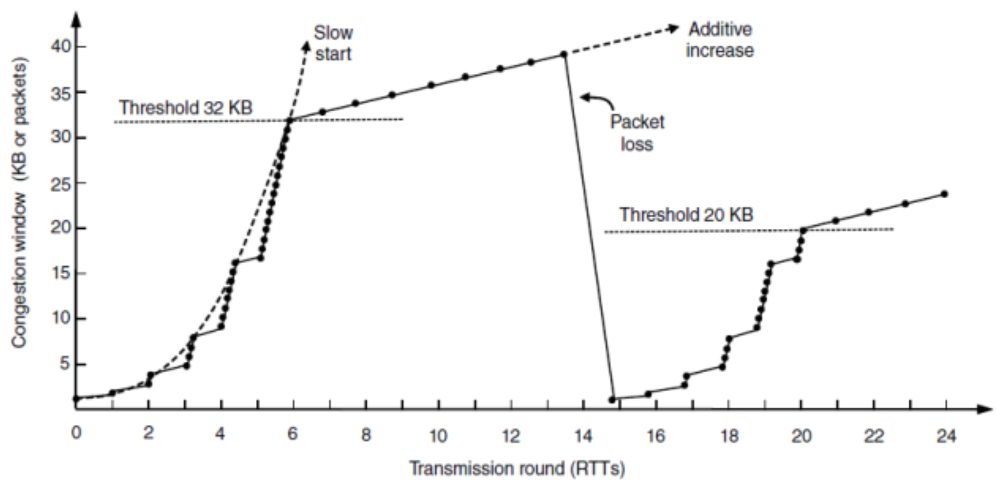
Slow start
Additive increase
Threshold 32 KB
Packet loss
Threshold 20 KB
Congestion window (KB or packets)
Transmission round (RTTs)
0
2
4
6
8
10
12
14
16
18
20
22
24
5
10
15
20
25
30
35
40

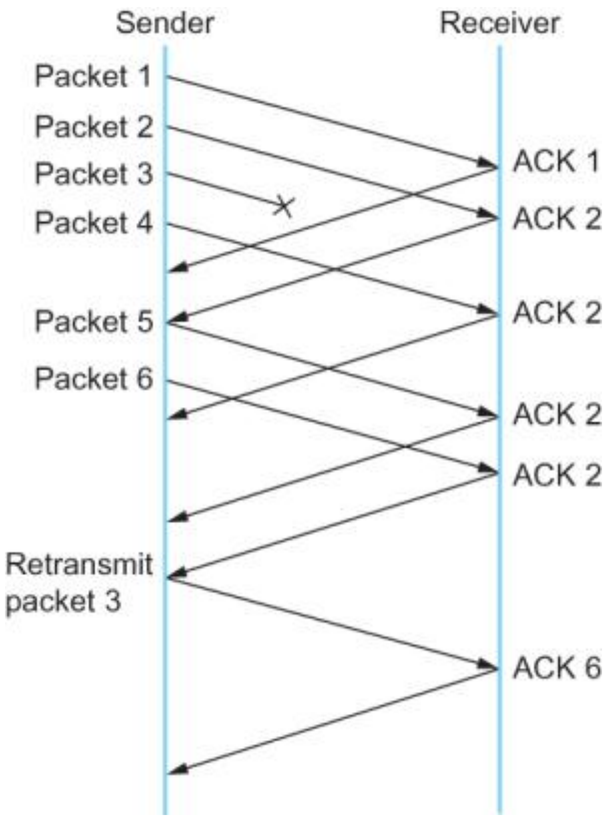
Sender
Receiver
Packet 1
Packet 2
Packet 3
Packet 4
Packet 5
Packet 6
ACK 1
ACK 2
ACK 2
ACK 2
ACK 2
Retransmit packet 3
ACK 6

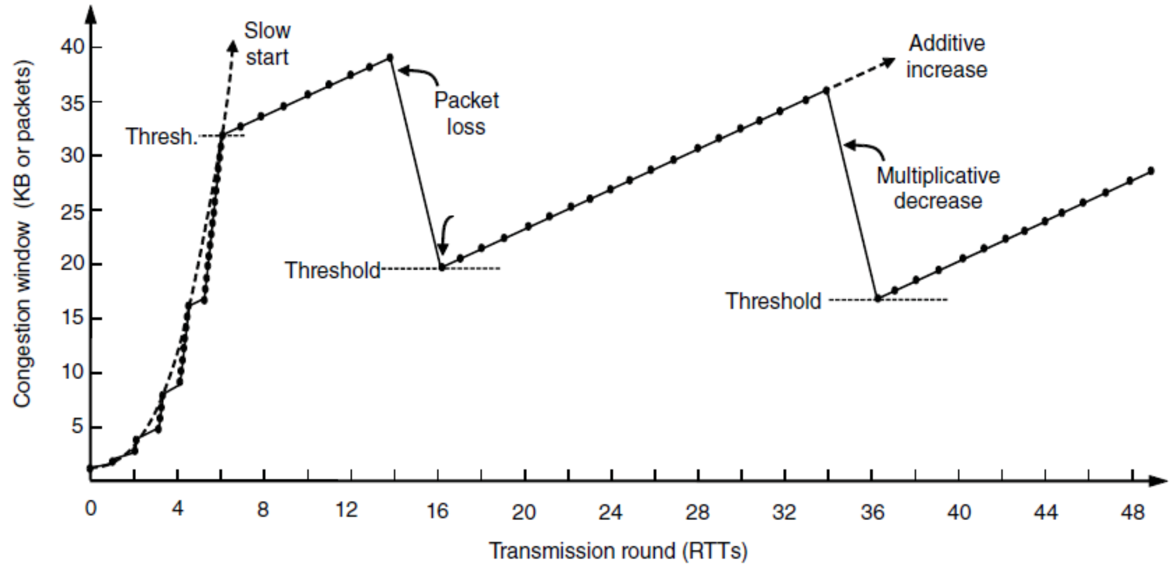
Slow start
Additive increase
Packet loss
Multiplicative decrease
Thresh.
Threshold
Threshold
Congestion window (KB or packets)
Transmission round (RTTs)
0
4
8
12
16
20
24
28
32
36
40
44
48
5
10
15
20
25
30
35
40

Type of applications

- Different applications have differing network needs
  - Stream of packets = flow
  - Determine Quality of Service (QoS) a flow requires

Application	Bandwidth	Delay	Jitter	Loss
Email	Low	Low	Low	Medium
File sharing	High	Low	Low	Medium
Web access	Medium	Medium	Low	Medium
Remote login	Low	Medium	Medium	Medium
Audio on demand	Low	Low	High	Low
Video on demand	High	Low	High	Low
Telephony	Low	High	High	Low
Videoconferencing	High	High	High	Low

What are people doing?

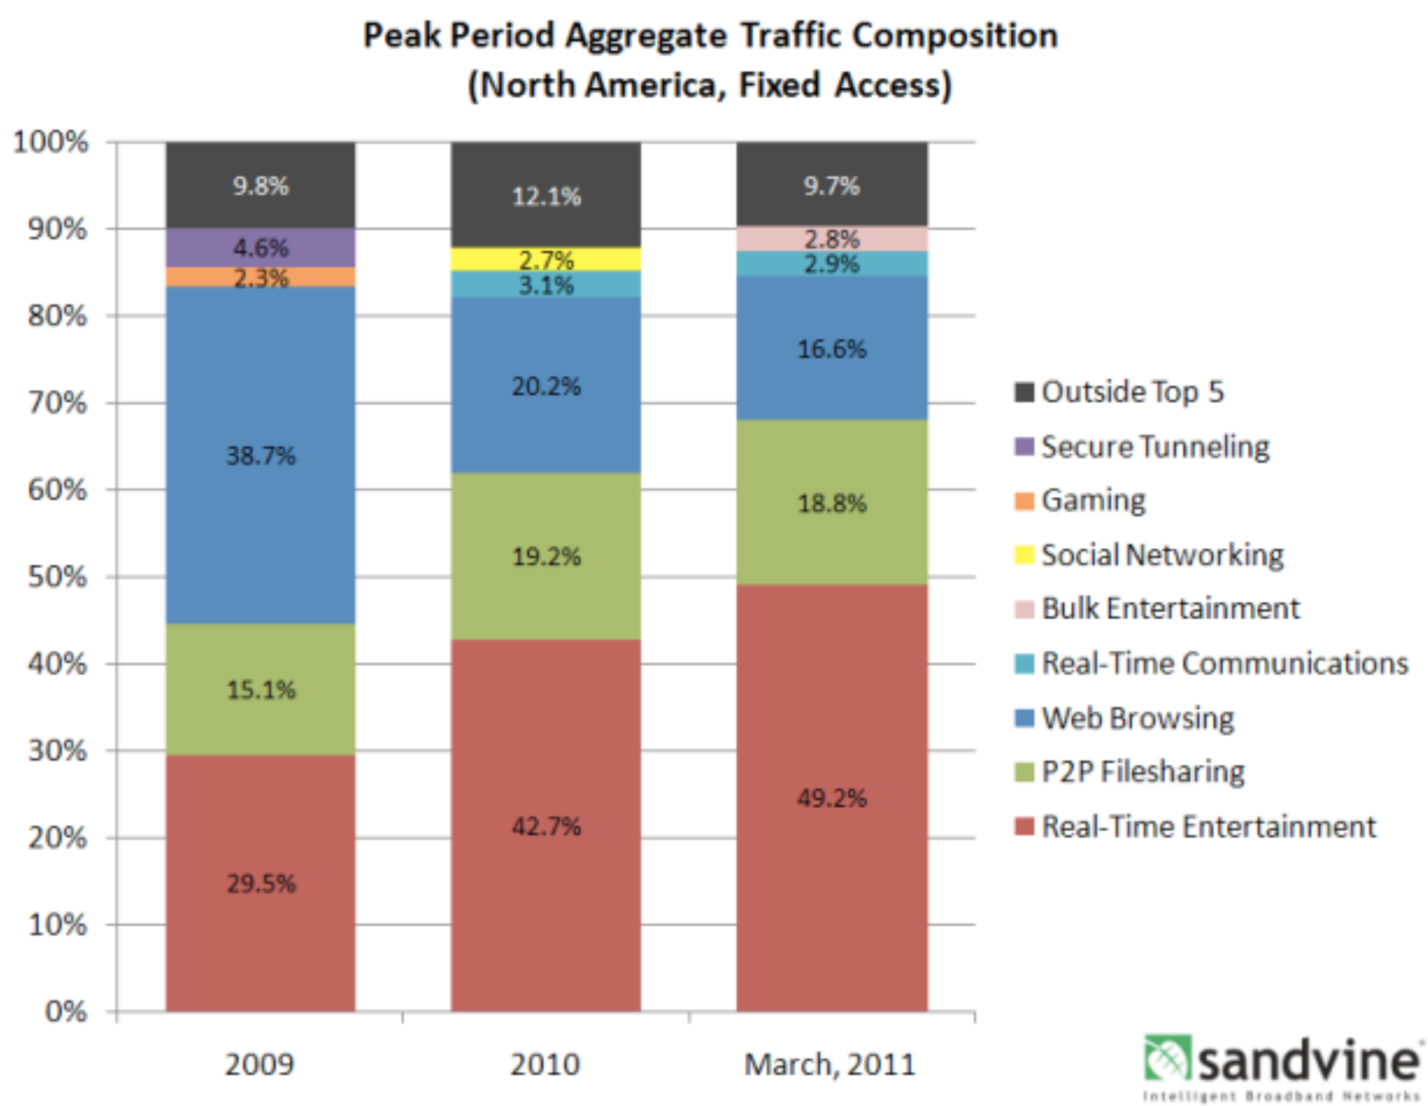

	Upstream		Downstream		Aggregate	
Rank	Application	Share	Application	Share	Application	Share
1	BitTorrent	47.55%	Netflix	32.69%	Netflix	29.03%
2	HTTP	11.45%	HTTP	17.48%	HTTP	16.59%
3	Netflix	7.69%	YouTube	11.32%	BitTorrent	13.47%
4	Skype	4.27%	BitTorrent	7.62%	YouTube	9.90%
5	SSL	3.57%	Flash Video	3.41%	Flash Video	3.04%
6	Facebook	2.19%	RTMP	3.12%	RTMP	2.81%
7	PPStream	1.73%	iTunes	3.05%	iTunes	2.69%
8	YouTube	1.64%	Facebook	1.78%	SSL	1.96%
9	Xbox Live	1.31%	MPEG	1.72%	Facebook	1.84%
10	Teredo	1.25%	SSL	1.69%	MPEG	1.49%
	Top 10	82.63%	Top 10	83.88%	Top 10	82.83%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS

sandvine

Table 1 - Top Peak Period Applications by Bytes (North America, Fixed Access)

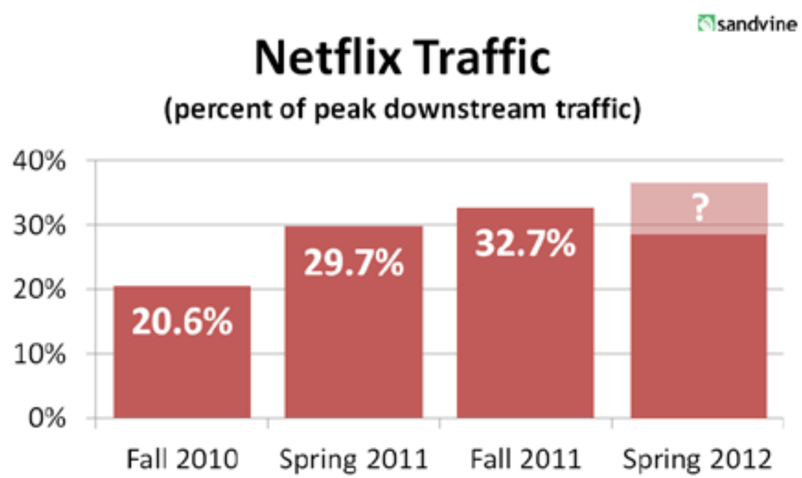

What are they doing mobile?

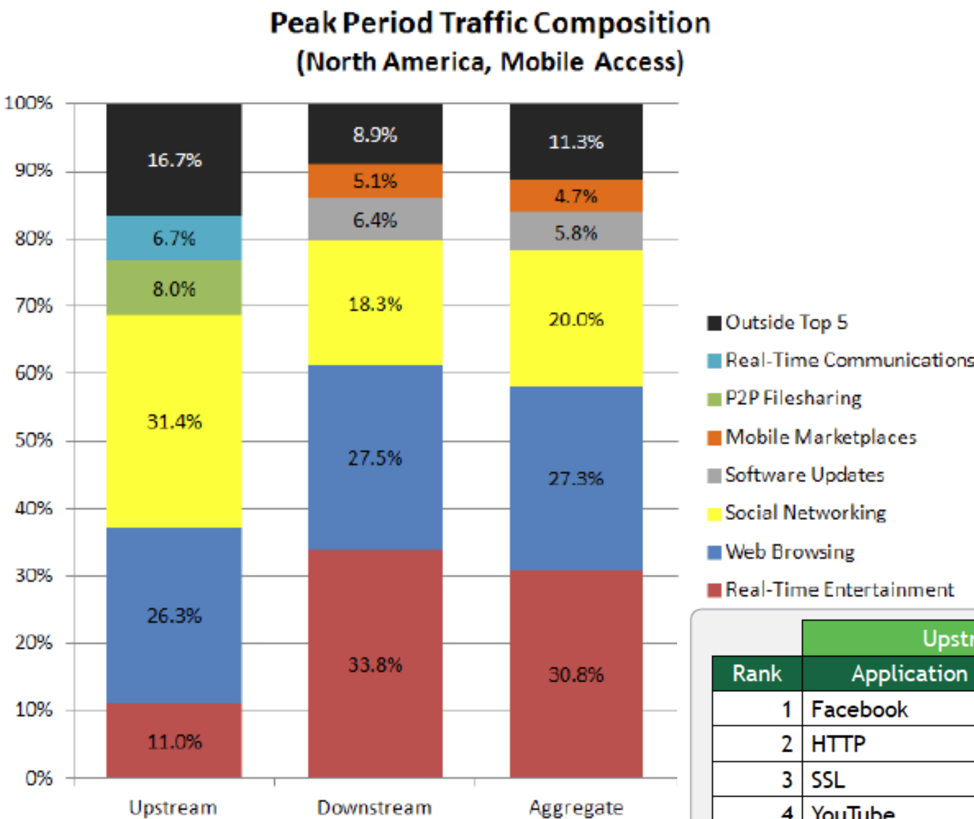

	Upstream		Downstream		Aggregate	
Rank	Application	Share	Application	Share	Application	Share
1	Facebook	30.85%	HTTP	27.46%	HTTP	27.31%
2	HTTP	26.24%	YouTube	19.99%	Facebook	19.29%
3	SSL	6.05%	Facebook	17.62%	YouTube	18.23%
4	YouTube	6.01%	Windows Update	5.17%	Windows Update	4.70%
5	BitTorrent	3.83%	Android Market	4.09%	Android Market	3.75%
6	Ares	3.45%	Flash Video	2.96%	Flash Video	2.66%
7	Oovoo	2.57%	SSL	1.97%	SSL	2.48%
8	Skype	1.81%	RTSP	1.89%	RTSP	1.67%
9	Gmail	1.49%	Shockwave Flash	1.75%	Shockwave Flash	1.63%
10	Windows Update	1.48%	MPEG	1.67%	MPEG	1.53%
	Top 10	83.77%	Top 10	84.57%	Top 10	83.26%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS

sandvine

Table 2 - Top Peak Period Applications by Bytes - North America, Mobile Access

When are they doing it?

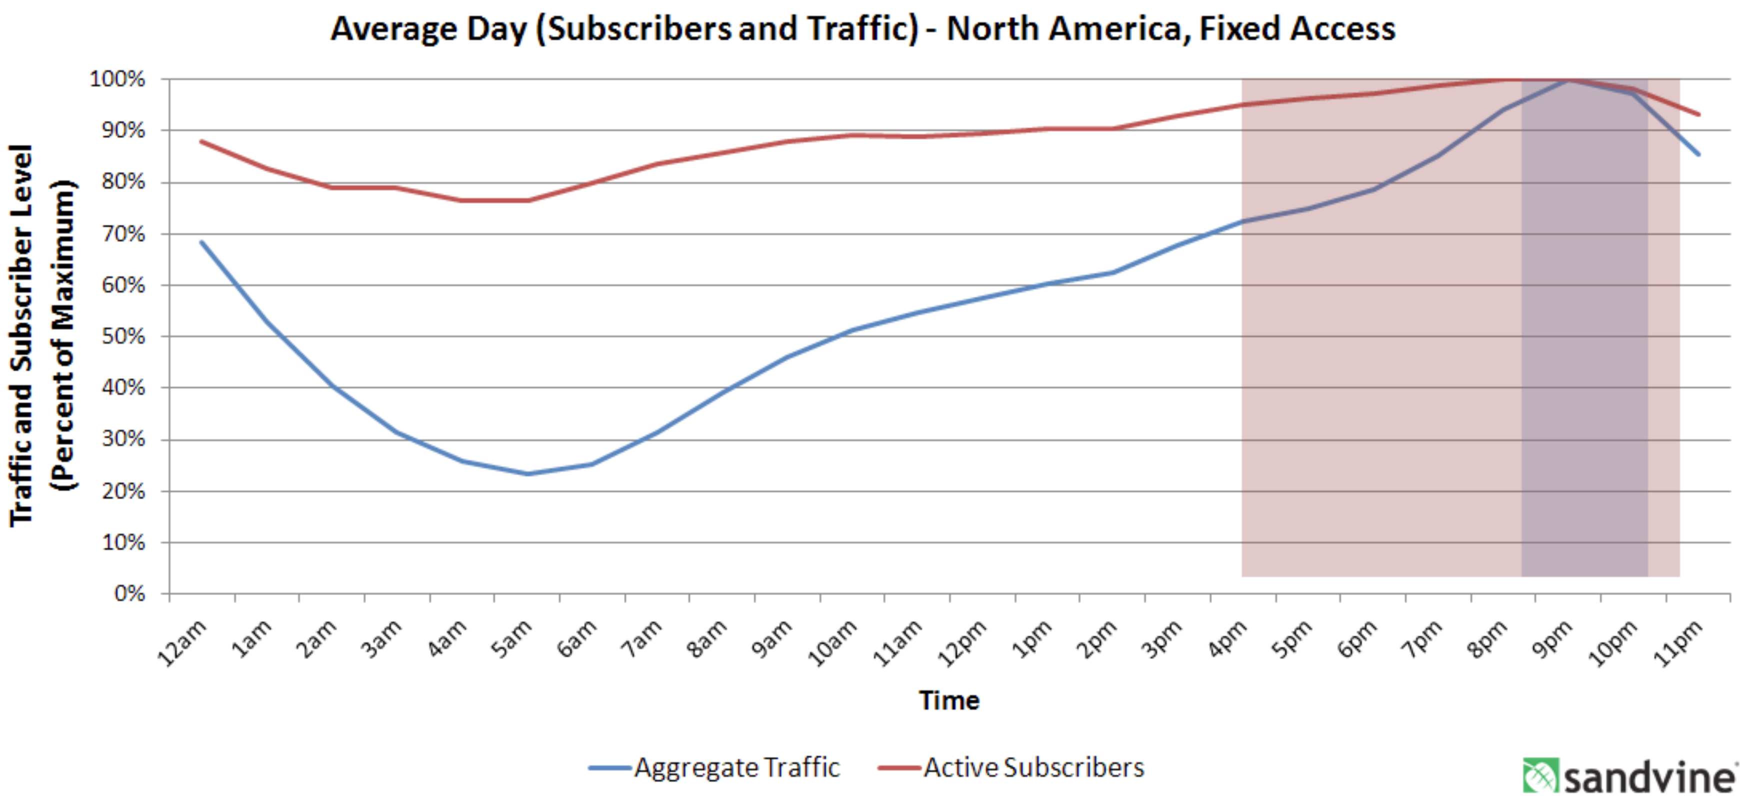

Taxonomy of applications

Just need end-to-end retransmission, e.g. file transfers.

Apps sensitive to timeliness of data, e.g. Skype.

Loss or delay may degrade app, but still works, e.g. dropping some frames of video.

Loss or delay is a big problem, e.g. controlling a surgical robot remotely.

Applications adjust to operate with network realities.

Applications static regardless of what happens.

Applications change buffer based on delay.

Applications change their bandwidth consumption, e.g. use lower resolution

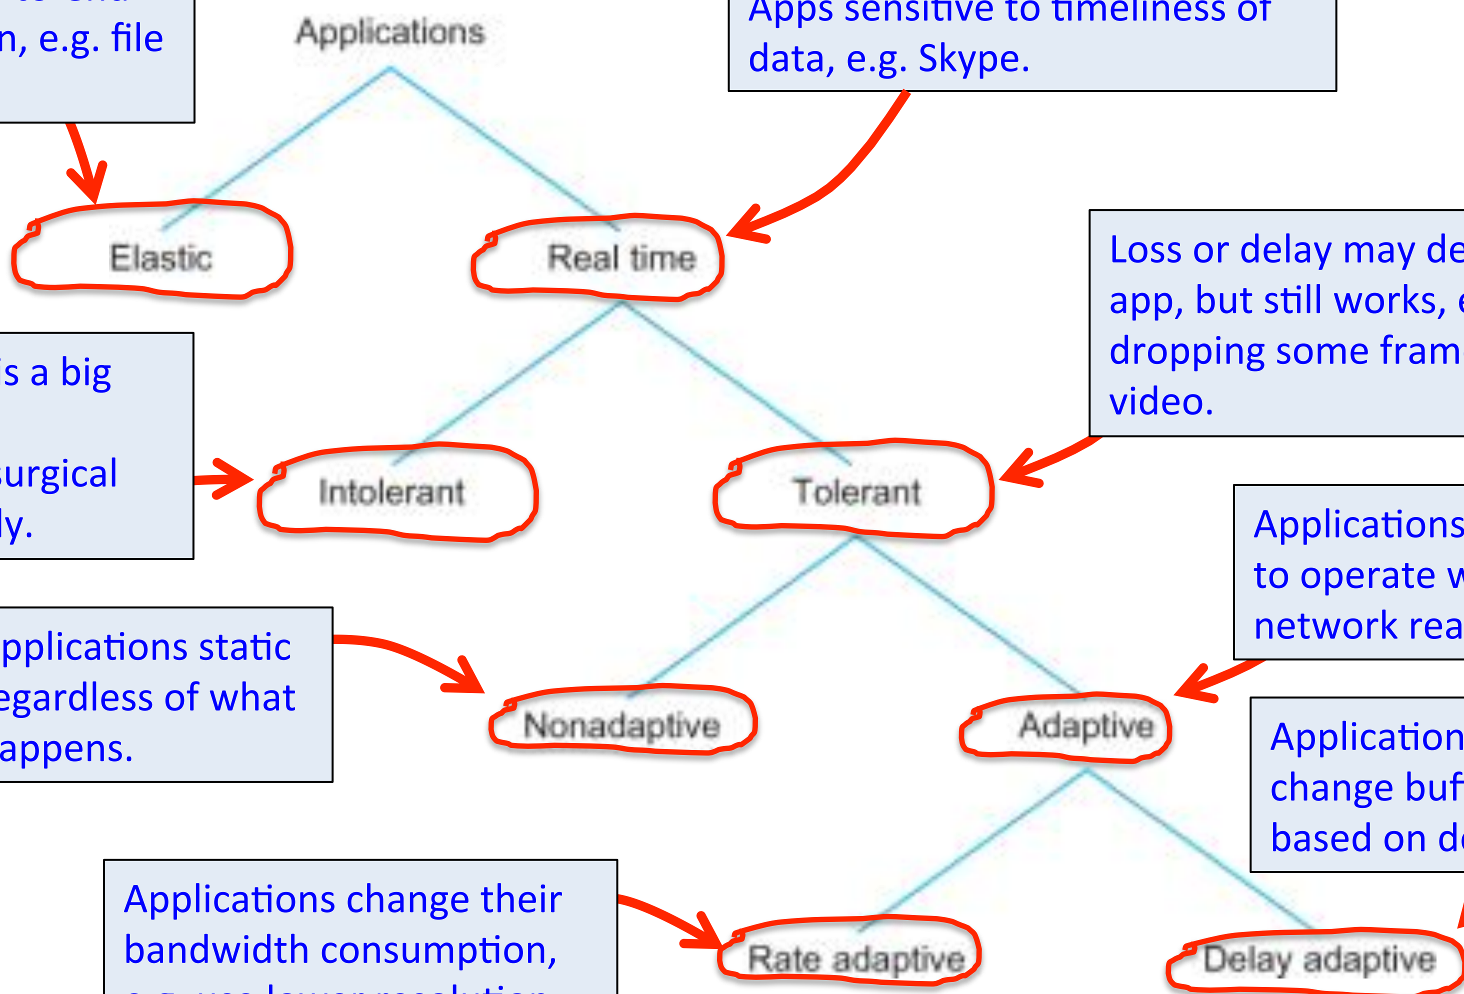

Real-time audio

- Delay adaptive
  - Change playback point
  - Good voice quality ~150ms one-way latency
  - Requires temporary increase/decrease in playback rate
- Rate adaptive
  - Use different sampling rate
  - Use a different vocoder

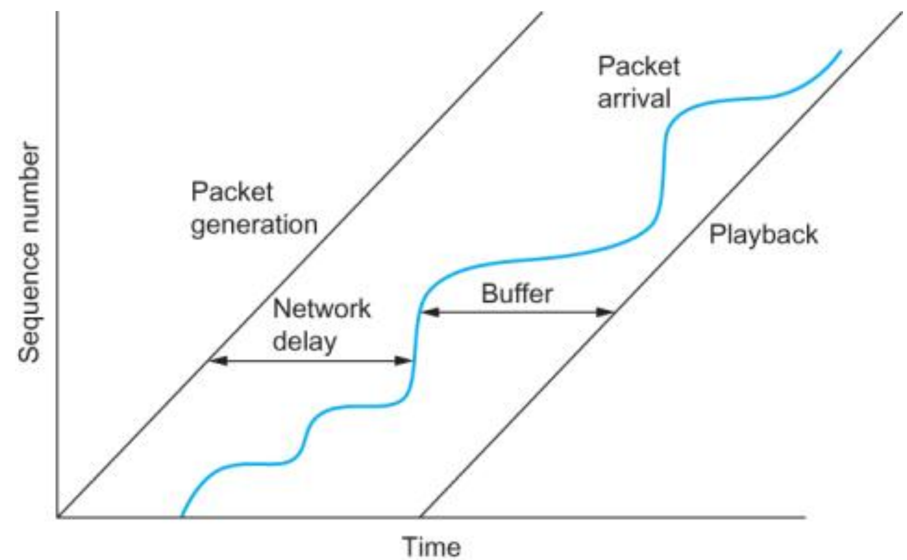

Impact of compression

- Mean opinion score (MOS)
  - Subjective measure of audio quality
  - 5-point Likert scale, 1 (bad) to 5 (excellent)

Method	Bit rate (Kbps)	MOS
G.711 PCM	64	4.1
G.726 ADPCM	32	3.9
G.728 Low Delay Code Excited Linear Predictive (LD-CELP)	15	3.6
G.729 Conjugate Structure Algebraic Code Excited Linear Predictive (CS-ACELP)	8	3.9
G.729a CS-ACELP	8	3.7
G.723.1 MP-MLQ	6.3	3.9
G.723.1 ACELP	5.3	3.7

Network performance

- Bandwidth and delay is variable
  - Long tailed distribution!

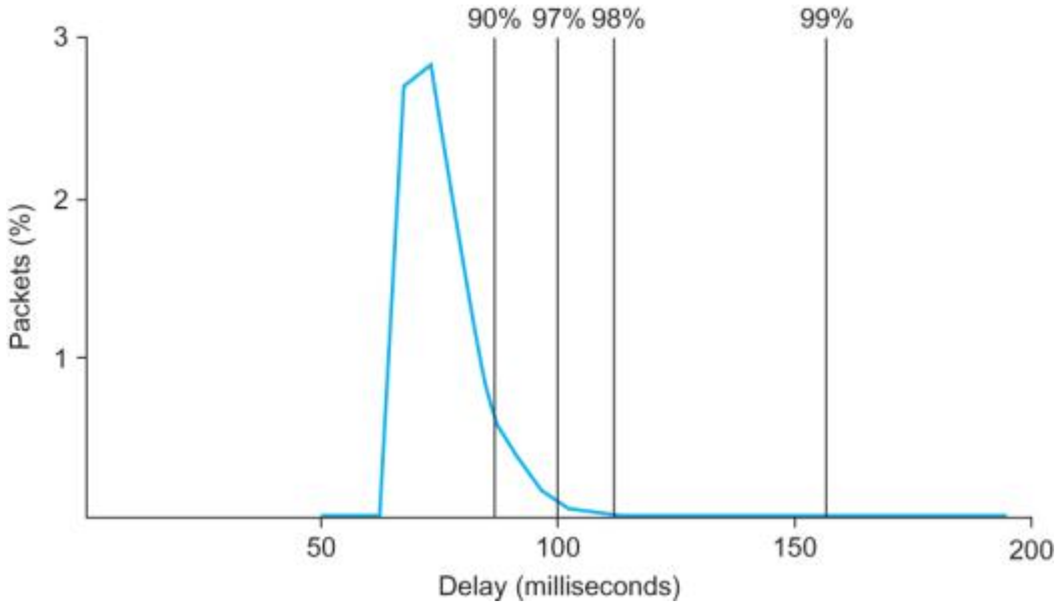

ping times to London

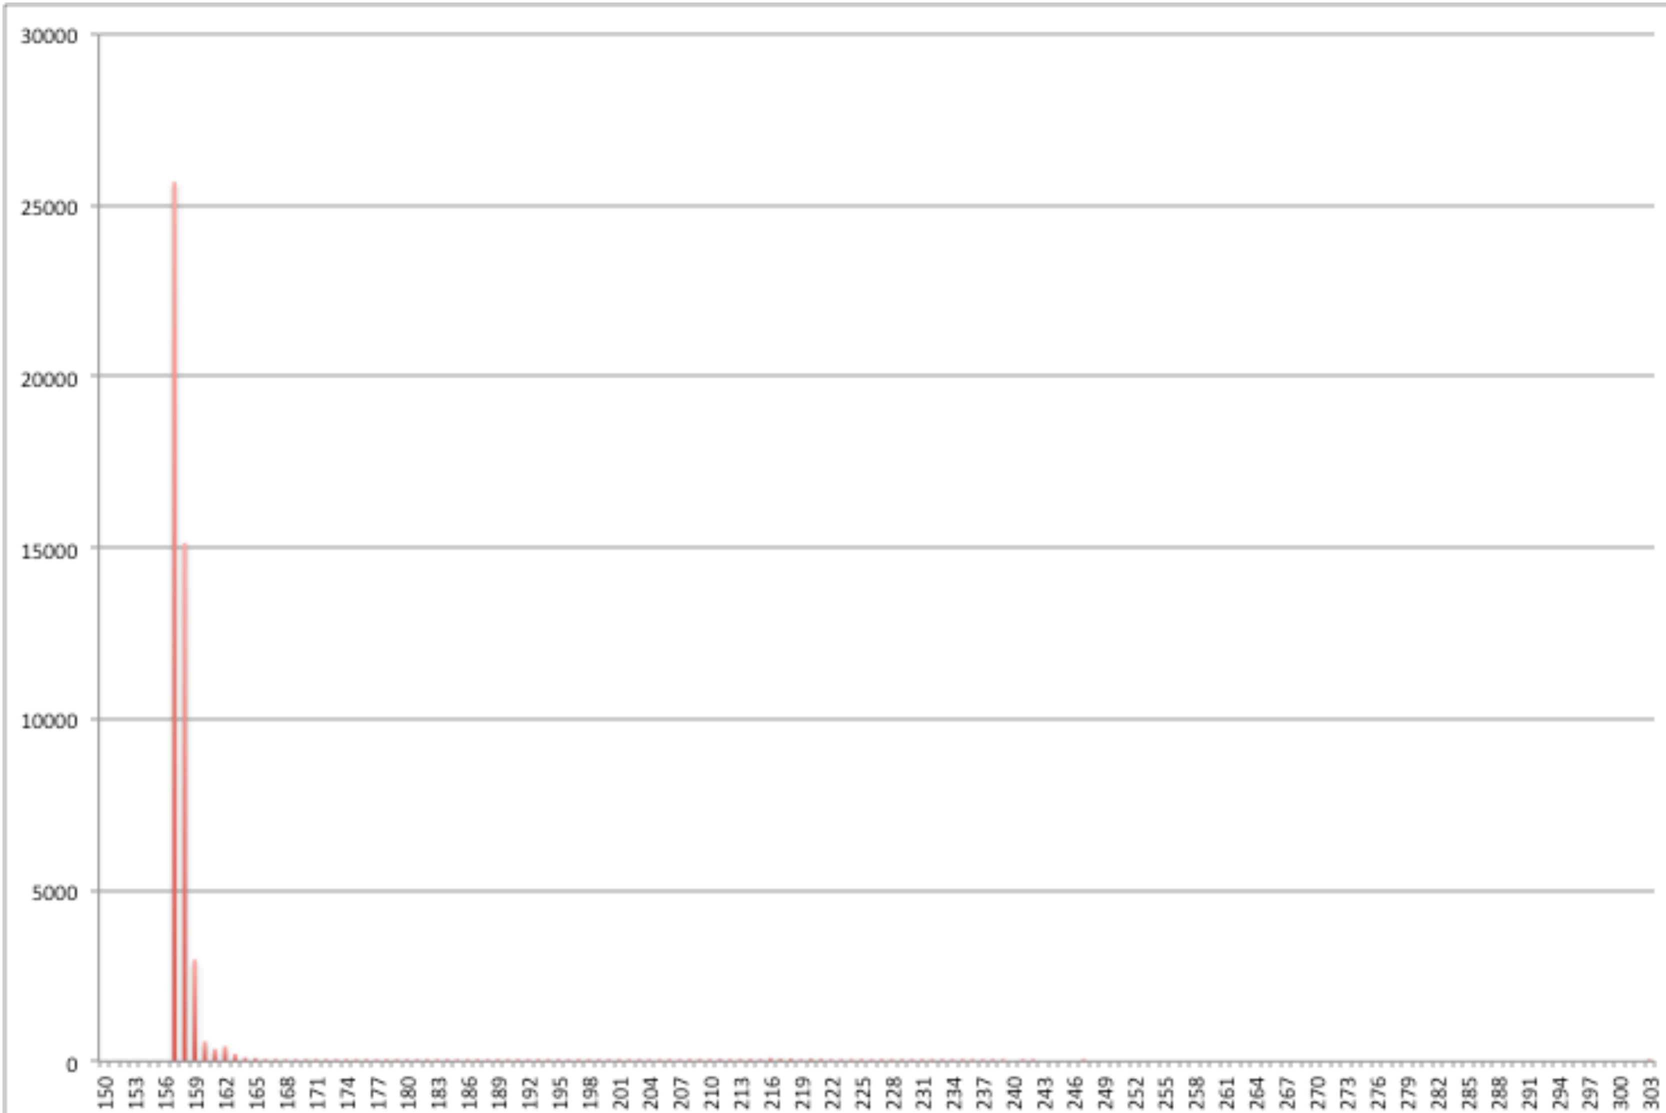

Approaches to QoS

- Overprovisioning
  - Build a network with lots of capacity, e.g. POTS
  - Why might this not be ideal?
- Flow-based
  - Hosts declare what they need
  - Admission control, not all flows get what they want
  - Integrated services (IntServ)
- Class-based
  - Packet classification system
  - Differentiated services (DiffServ)

Flow-based QoS

- Integrated Services (IntServ)
  - Effort of IETF 1995-97, produced two dozen RFFs
  - Largely not adopted
  - Unicast and multicast applications
- Resource reSerVation Protocol (RSVP)
  - Two service classes:
    - Guaranteed service – no packet arrives after playback time
    - Controlled load – emulate lightly loaded network on a heavily loaded one

Make a reservation

- Host sends signal through network
  - Hosts says what they need
  - Reservation spec (RSpec)
    - Guaranteed delivery: delay or bound on latency
    - Controlled load: no parameters
  - Traffic spec (TSpec)
    - Characterization of how the application will use the network
- Admission control
  - Routers decide if they can provide or not

Meeting a reservation's needs

- Routers have limited resources:
  - Bandwidth, link types
  - Buffer space, memory
  - CPU cycles, packets/second
- Packet scheduling:
  - FIFO with tail drop
    - Not suited for providing QoS
  - Priority queuing
    - Burst of high-priority can starve low-priority
  - Weighted fair queuing

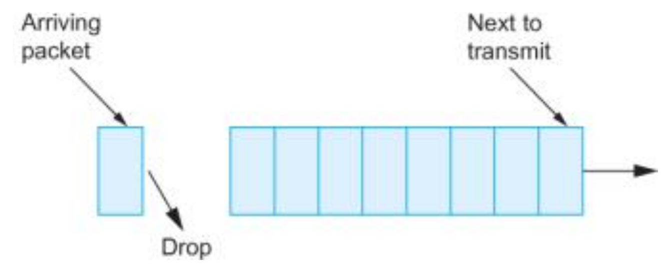

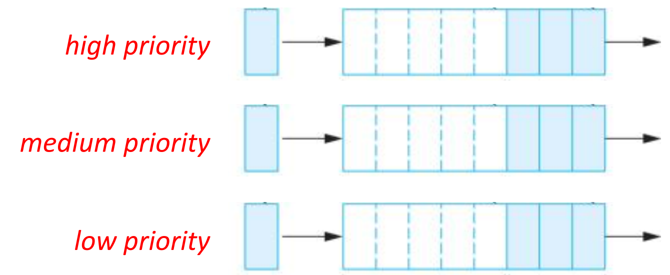

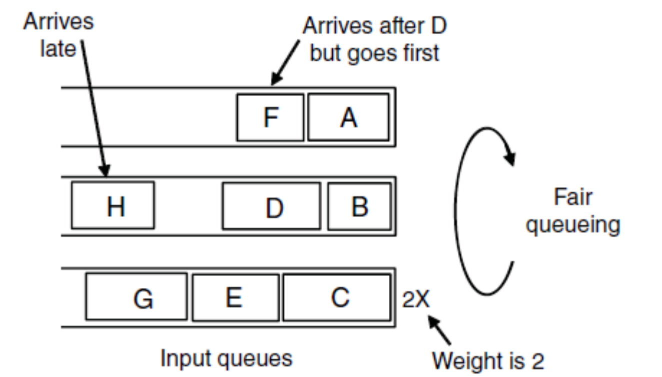

Traffic patterns

- Characterizing a host's network usage
  - Single number not sufficient
    - Constant bandwidth utilization not the same as infrequent bursts
- Leaky / token bucket
  - Outflow constant R bytes per second
  - Bucket can hold B bytes
  - Overflow = dropping packets

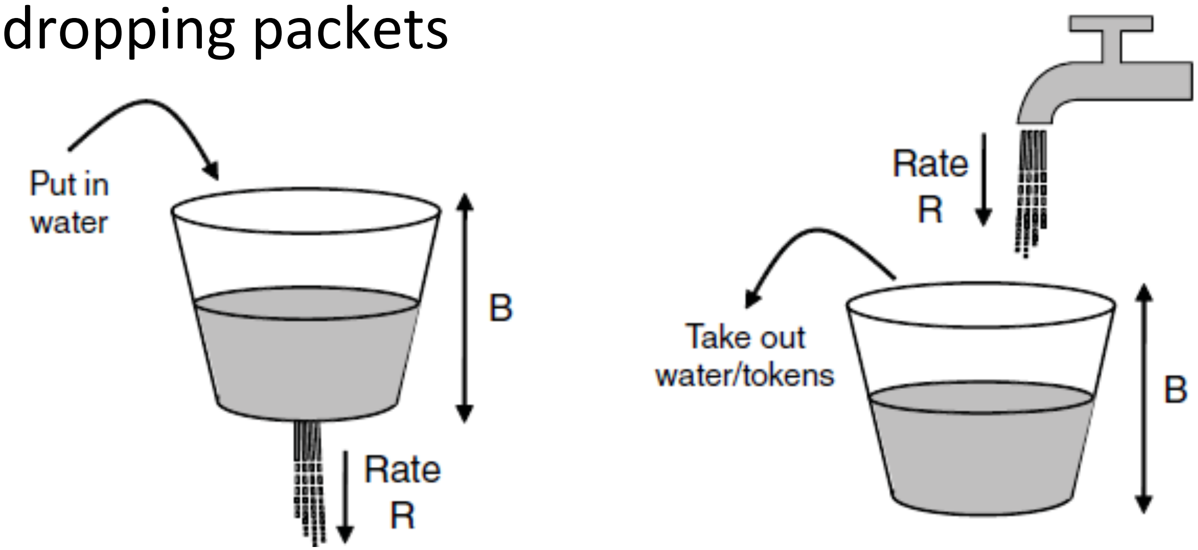

Token bucket

- Token bucket shaping example:
  - Computer produces data at 1000 Mbps (125 MB/sec)
  - First link also 1000 Mbps

Host wants to send a burst at line speed, then a longer sustained transmission.

If network can only handle 200 Mbps, we have to slow down part way through first burst.

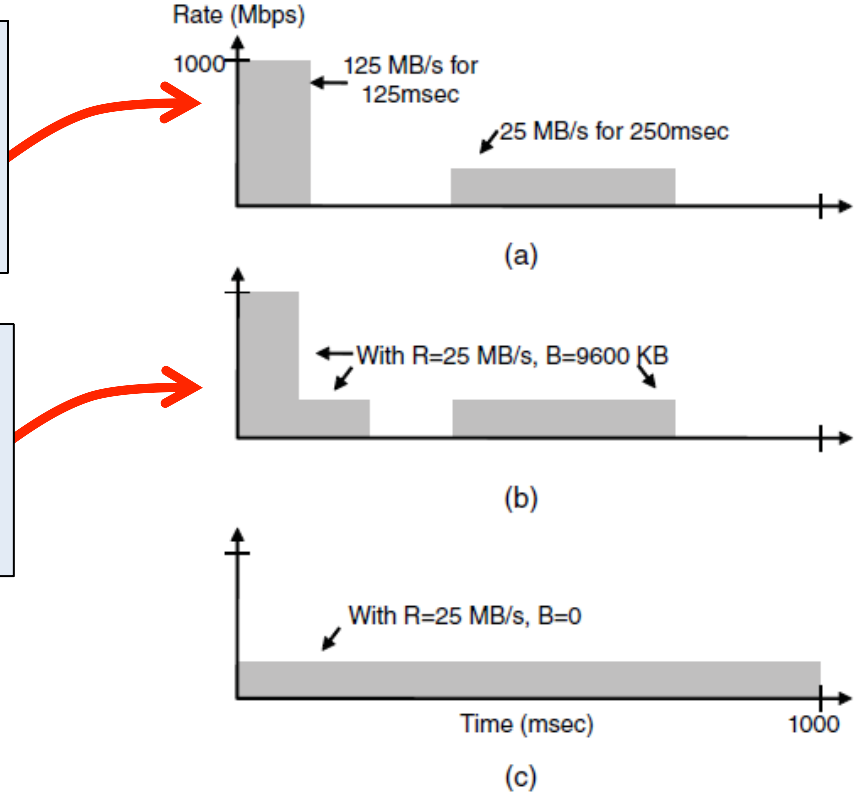

Flow specification

- Integrated services request
  - RFC 2210, 2211
  - Host provides five parameters:

Parameter	Unit
Token bucket rate	Bytes/sec
Token bucket size	Bytes
Peak data rate	Bytes/sec
Minimum packet size	Bytes
Maximum packet size	Bytes

  - Routers use for admission control and packet scheduling

Reservation example

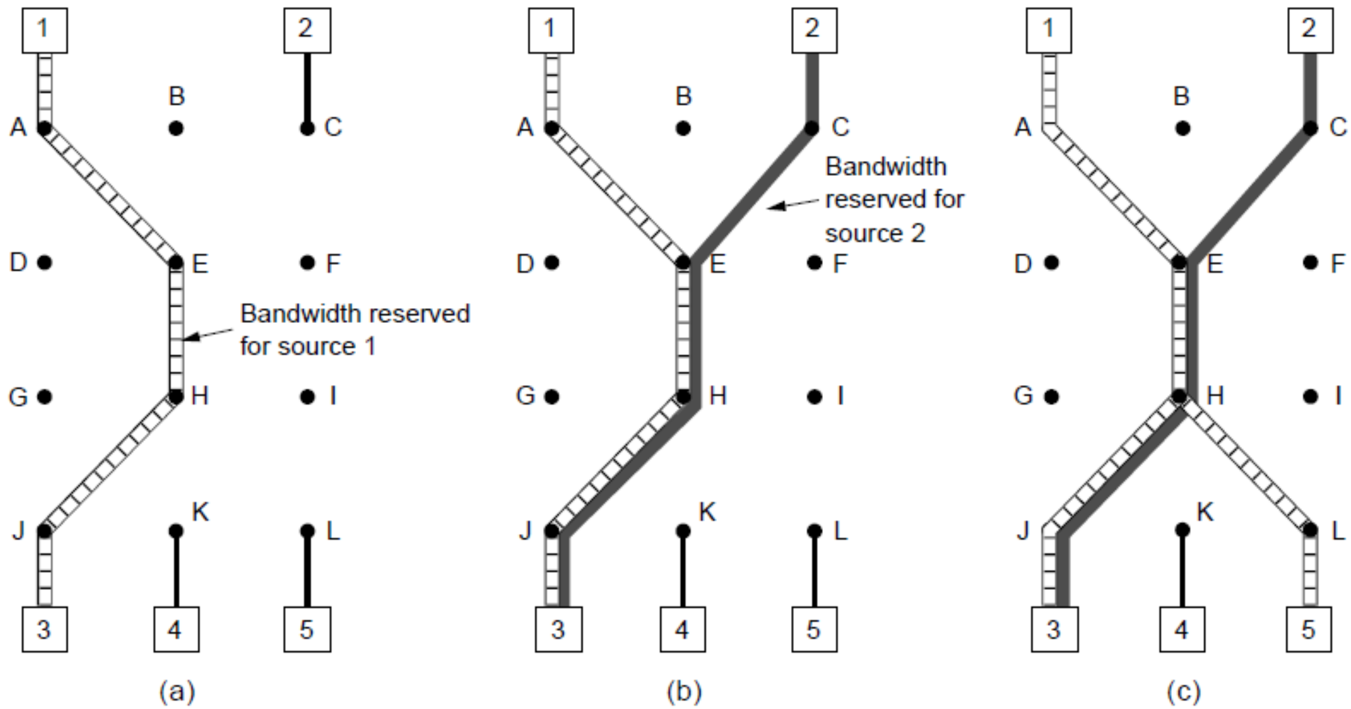

- Host 3 wants to watch channel from host 1, flow admitted
- Packets flow from 1 to 3 without congestion
- Host 3 simultaneous starts watching channel from host 2
- Host 5 starts watching channel from host 1 as well

Integrated Services

- Adoption of IntServ
  - Despite early development, not widely deployed
  - Scalability
    - Every flow passing through a router may need a reservation, requiring router memory
    - Flows need to be established and policed by the router, requiring CPU time and added router complexity

Classed-based QoS

- Differentiated services (DiffServ)
  - Classify packets into a small # of traffic classes
    - Perhaps as simple as normal and high priority
  - Routers apply different per-hop behaviors (PHBs)

PHB determined by looking at 6 bits in the TOS byte of the IP header. Each 6-bit value is a different DiffServ code point (DSCPs)

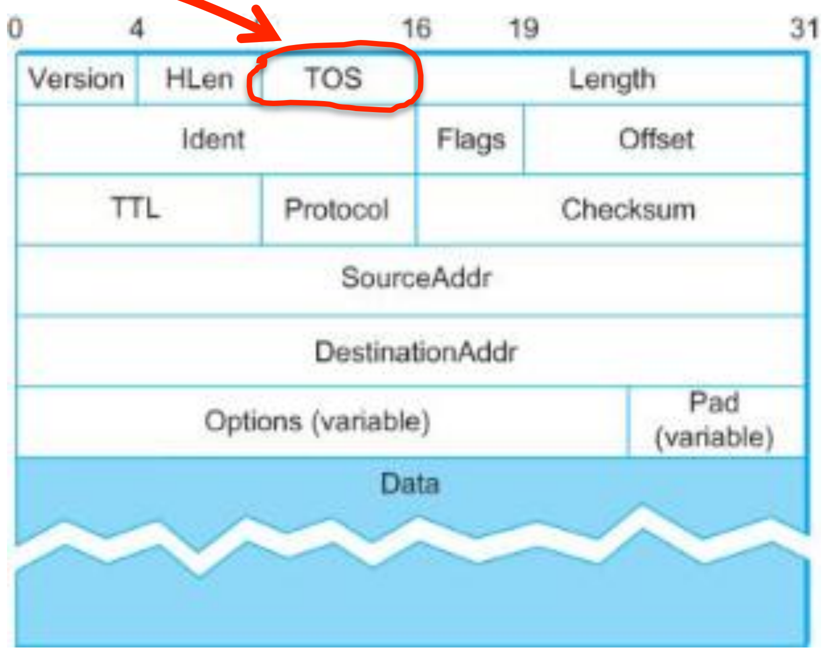

Expedited Forwarding

- Expedited forwarding (EF) PHB
  - Marked packets get priority treatment at routers
    - Strict priority
    - Weighted fair queuing
  - An AS could rate limit at boundary, so never more EF packets than slowest link in network
  - e.g. Voice over IP, residential companies, universities

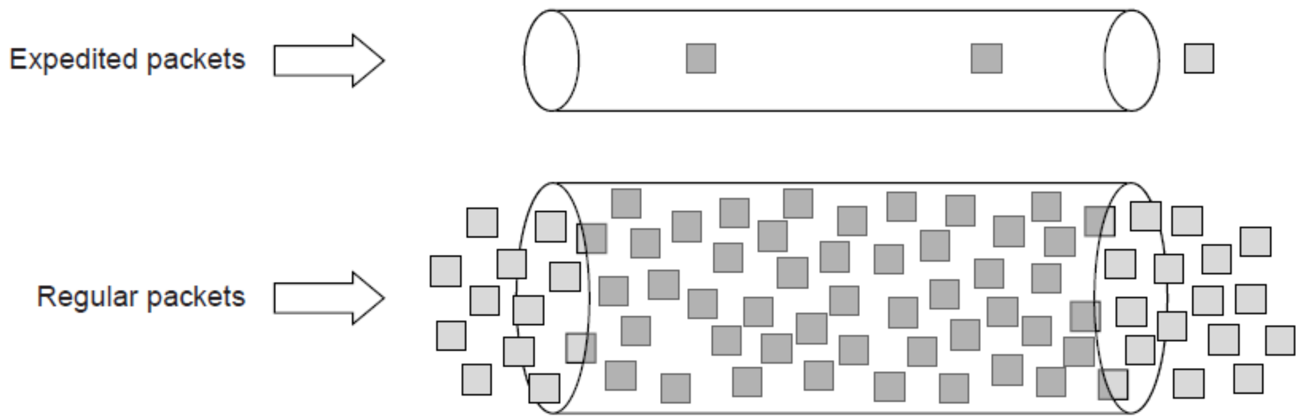

Random early detection

- Random early detection (RED)
  - If router approaching congestion: drop a random packet
  - Source detects packet loss and can adjust send rate
  - Randomness approximates fairness since more likely to signal host sending lots of packets
  - Various parameters controlling drop behavior

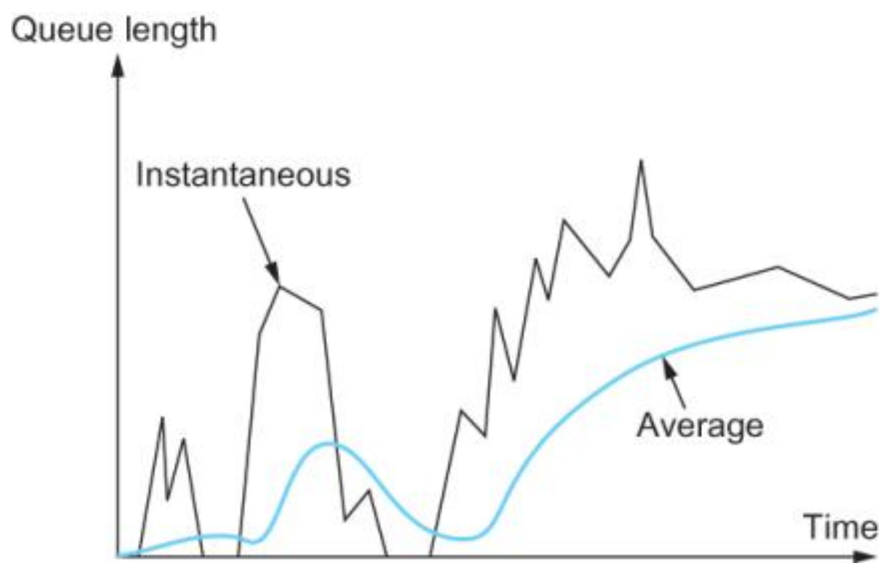

Assured Forwarding

- Assured Forwarding (AF) PHB
  - RED with In and Out (RIO), two classes of traffic:
    - in = important stuff
    - out = other stuff
    - Different drop probability curve for each class
  - Weighted RED (WRED)
    - More than two curves, choose via DSCP value

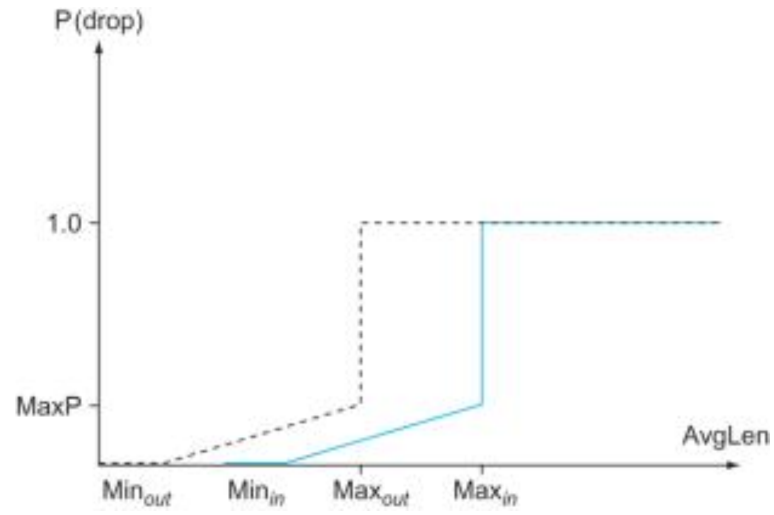

Assured Forwarding

- Assured Forwarding (AF) PHB
  - IETF RFC 2597
    - Four priority classes x 3 discard classes
    - Priority classes go into different WFQ queues

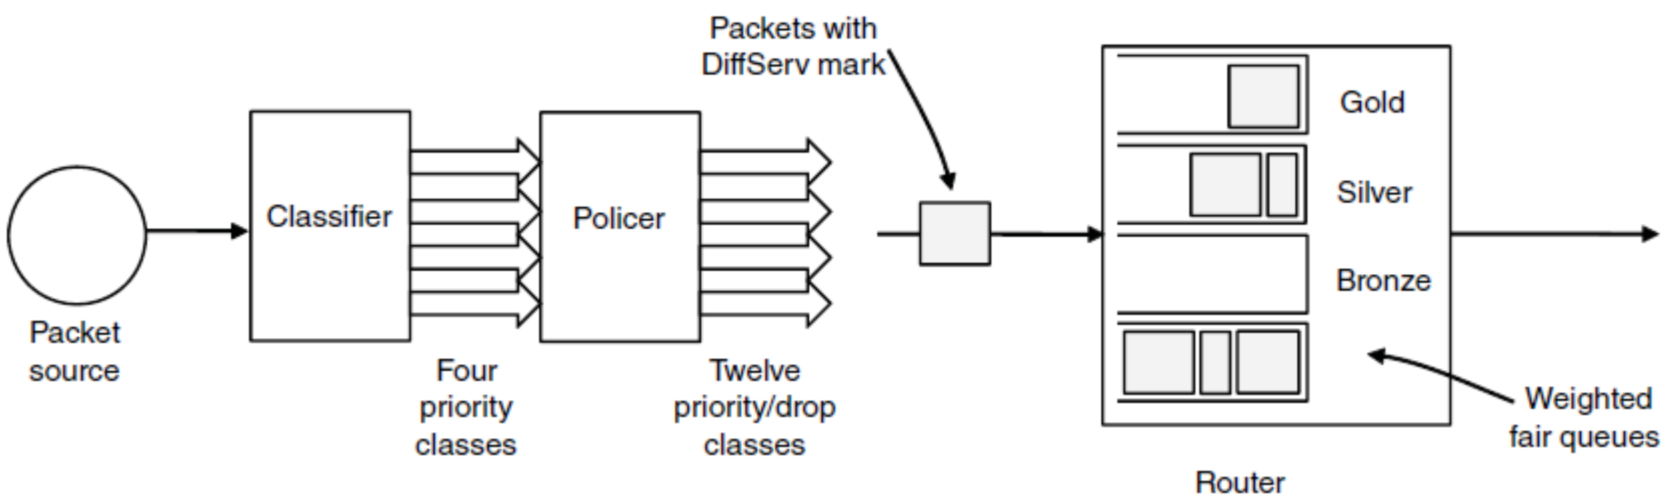

Assured Forwarding

- Host-based QoS
  - TCP congestion control works quite well
    - Requires no cooperation from the network
  - But real-time apps don't want retransmission
    - Data will be too late
  - And real-time apps don't want constantly varying speed
  - Use UDP transport with appropriate rate control so it plays nice with TCP

Summary

- Not all data is created equal
  - Real-time data has special latency needs
  - Increasingly dominates Internet traffic
  - Real-time traffic needs special attention
- Integrated Services
  - Reserving resources in advance
  - Flows must be identified and classified
- Differentiated Services
  - Hosts / providers mark certain packets for preferential treatment by the network