

Distributed Systems

2. Quality of Service

Paul Krzyzanowski
pxk@cs.rutgers.edu

What factors matter for network performance?

- **Bandwidth (bit rate)**
 - Average number of bits per second through the network
- **Delay (latency)**
 - Average time for data to get from one endpoint to its destination
- **Jitter**
 - Variation in end-to-end delay
- **Errors (packet loss)**
 - Percentage of packets that don't reach their destination

Why do we care?

- We want to provide users with
 - Good interactive performance
 - Mostly glitch-free streaming video
 - Good quality IP telephony

Quality of Service Goal

- Provide an applications with the data flow performance they need
- Examples
 - VoIP: low bandwidth, low latency (<150 ms), low jitter (<50 ms)
 - Video streaming: high bandwidth, longer latency OK
 - Remote login: low bandwidth, relatively low latency, high reliability
 - File transfer: high bandwidth, high latency OK, high jitter OK , high reliability

QoS Aware Networks

- A QoS-aware network will allow an application to specify its needs and will reserve sufficient capacity in the network to meet those goals
- **Admission Control**
 - Applications request a particular quality of service from the network
 - If the request is granted, the application has the contract
- **Traffic Control**
 - Classifies, queues, and schedules traffic within the network

ATM: Asynchronous Transfer Mode

- Late 1980's
- Goal: Merge voice & data networking



low but constant
bandwidth



high but bursty
bandwidth

ATM

- Traditional voice networking
 - Circuit switching
 - Too costly
 - Poor use of resource
 - Does not lend to multicasting
- ATM
 - Based on **fixed-size packets over virtual circuits**
 - Fixed-size cells provide for **predictive scheduling**
 - Large cells will not hold up smaller ones
 - Rapid switching

ATM

- Current standard:
 - 53-byte cell: 48-byte data, 5-byte header
- Sender specifies traffic type upon connecting:

CBR	Constant bit rate	<i>Bandwidth</i>	Uncompressed video, voice
VBR	Variable bit rate	<i>Avg, peak bandwidth</i>	Compressed video, voice
ABR	Available bit rate	<i>Minimum bandwidth</i>	web access
UBR	Unspecified bit rate	–	ftp

ATM

- Small cells → lots of interrupts
 - >100,000/second
- ATM hardware supports an **ATM Adaptation Layer (AAL)**
 - Converts cells to variable-sized (larger) packets:
 - AAL 1: for CBR
 - AAL 2: for VBR
 - AAL 3/4: ABR data
 - AAL 5: ABR data, simplified
 - AAL 6: MPEG-2 video
- IP traffic: transmitted as data using AAL5

Quality of Service in IP

- IP was not designed with QoS controls in mind
- IP cannot take advantage of QoS controls offered by an underlying network
 - And you can't count on a specific underlying network for end-to-end service
- Some QoS mechanisms created as add-ons

Quality of Service Problems in IP

- Too much traffic
 - Leads to **congestion**
- Inefficient packet transmission (not really a QoS issue)
 - 59 bytes to send 1 byte in TCP/IP!
 - 20 bytes TCP + 20 bytes IP + 18 bytes ethernet
- Unreliable delivery
 - Software to the rescue – TCP/IP
- Unpredictable packet delivery
 - No controls on bit rate, delay, jitter
 - Packets may take different routes, resulting in different performance

Inefficient Packets

- Some software generates lots of tiny packets
 - Head-of-line blocking
 - Nagle's algorithm:
 - buffer new data if unacknowledged data outstanding
- Header/packet compression
 - Link-to-link
 - Header compression (RFC 3843)
 - Payload compression (RFC 2393)
 - Trade-off: \$ delivery vs. \$ compression

Approaches to QoS

- Best effort
 - Avoid congestion by overprovisioning the network
- Differentiated service
 - Identify one type of service (data flow) vs. another
 - Use this information to affect packet queuing and scheduling
- Guaranteed service
 - Have every device in a route commit to a level of service for a specific data flow

IP Flow Detection


- Flow detection in routers:
 - Flow: set of packets from one *address:port* to another *address:port* with same protocol
 - Network controls flow rate by dropping or delaying packets
 - With flow detection
 - drop TCP packets over UDP
 - Discard specific UDP flows to ensure QoS for other flows

IP Flow Detection

- With flow detection:
 - Traffic Shaping
 - Identify traffic flows
 - Queue packets during surges and release later
 - High-bandwidth link to low-bandwidth link
 - Traffic Policing
 - Discard traffic that exceeds allotted bandwidth

Dealing with congestion: Scheduling

- Congestion
 - More incoming traffic than we can route
 - Pick what traffic to route, what to delay (buffer), and what to drop
- FIFO queuing
- Priority queues
- Flow-based weighted fair queuing
 - Group all packets from a flow together
- Class-based weighted fair queuing
 - Based on protocols, access control lists, interfaces, etc.
- Custom queues

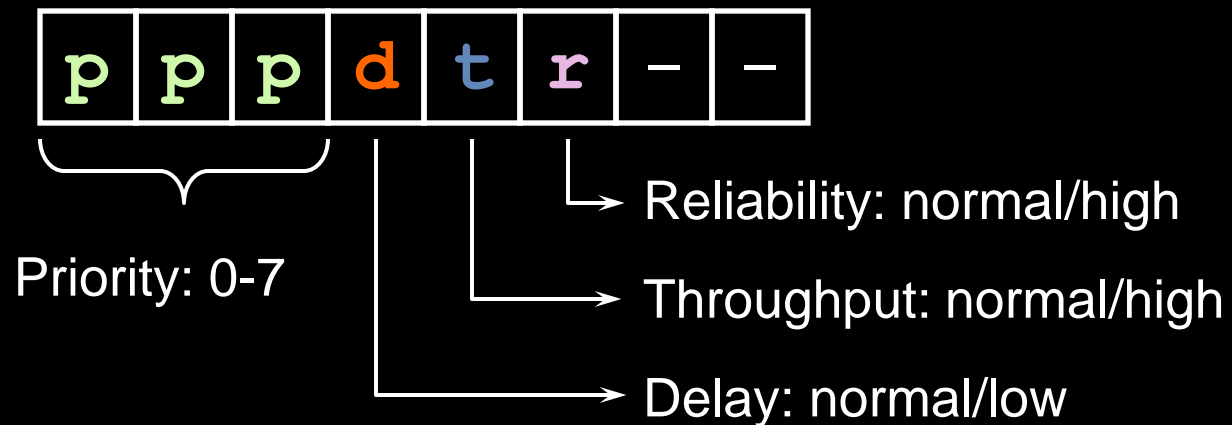


Differentiated Services (soft QoS)

- Treat some traffic as better than other
 - Statistical - no guarantees
- Identify class of service & priority
 - Router can use this data to make scheduling/dropping decisions
- Use on Internet (especially across ISPs) limited due to peering agreement complexities

Differentiated Services (DiffServ)

- ToS field in IP header (bits 9...16)
 - Differentiated Services Control Point (DSCP)



RFC 2475, December 1998

Guaranteed QoS (hard QoS)

- Guarantee via end-to-end reservation

Reservation & Delivery Protocol

- RSVP: ReSerVation Protocol
 - Each unidirectional data stream is a flow
 - Hosts request specific quality of service per flow
 - Routers reserve resources
 - RFC 2205
- Every device through which data flows must support RSVP

Other mechanisms

- Multi-Protocol Layer Switching (MPLS)
 - Make packet routing more efficient ... more like layer 2
 - Each packet is assigned a routing label (based on destination & packet priority)
 - Routers make queuing and forwarding decisions based on this label
 - Similar problems to RSVP

Other mechanism

- 802.1p: control QoS at the MAC layer (layer 2)
 - Eight traffic classes defined in an ethernet frame
 - This is *not* an IP QoS mechanism

Priority	Traffic
0	Background
1	Best effort
2	Excellent effort
3	Critical
4	Video <100 ms latency
5	Voice <10 ms latency
6	Internetwork control
7	Network control

Media Delivery Protocols

- Real-Time Control Protocol (RTCP)
 - Provides feedback on QoS (jitter, loss, delay)
 - RFC 3550
- RTP: Real-Time Transport Protocol
 - Not a routing protocol
 - No service guarantees
 - Provides:
 - Payload identification
 - sequence #
 - time stamp
- RTP/RTCP do not provide QoS controls

The End