

Quality of Service and Congestion Management in High Speed Networks

Sonia Fahmy
Purdue University
fahmy@cs.purdue.edu
<http://www.cs.purdue.edu/homes/fahmy/>




- What is Quality of Service (QoS)?
- Four approaches for QoS
- What is congestion management?
- Old and new myths

Quality of Service

- Predictable quality. Metrics:
 - Delay (in time) e.g., round trip delay, one way delay
 - Jitter = delay variation
 - Throughput e.g., in bits per second
 - Loss
 - Error
- Triangle
 - Sender wants to send at any time, with high load, burstiness
 - Receiver expects good service (low delay, high throughput, etc)
 - Carrier wants to minimize infrastructure (e.g., link) cost

Traffic Management

- Traffic management is important when there are multiple services, e.g., for real-time and bulk data, statistically multiplexed
- A dynamic problem. A resource allocation problem.
- Resource = link, router, switch, host, server



Mechanisms

- Traffic management components:
 - Capacity planning
 - Admission control
 - Shaping
 - Policing
 - Scheduling
 - Buffer management
 - Feedback control

QoS can be...

- Deterministic: all packets
 - Or:
- Statistical: no more than x% will see poor performance
- If statistical:
 - Steady state
 - Or:
 - Over specific intervals of time, e.g., no more than x% of the intervals of length I will...

QoS Challenges

- Traffic sources exhibit correlated time-varying behavior
- Granularity of QoS requirements is per-session, not aggregate
- Performance must be evaluated in a network multi-hop setting = intra and inter-session packet interactions due to multiplexing (scheduling)

Tightly Controlled Approaches

- The queuing (scheduling) discipline preserves a session's traffic characteristics
- Example: Stop and go queuing = next output frame
- Performance bounds are easy to compute
- Problems:
 - Per-session non-work conserving scheduling
 - Bandwidth reservation based on peak rate (if peak-to-average ratio large)
 - High delay

Approximate Approaches

- Model traffic sources by simple models, e.g., on/off
- Analyze queuing behavior
- Advantages:
 - Simple
 - Statistical multiplexing
- Disadvantages:
 - Conservative approximations
 - Complex sources modeling
 - Markovian assumptions at nodes do not hold
 - Local versus end-to-end QoS


Bounding Approaches

- Accounts for changes in traffic characteristics as traffic passes through a work conserving multiplexer
- Computes performance bounds for both deterministic and statistical guarantees
- Bounds are computed for each session's traffic after it passes through each multiplexer along its path in the network
- References: Cruz and Parekh
- Assume bound on queue busy period

Observation-based Approaches

- Previously made measurements are used to characterize traffic
- Does not require sources to characterize their traffic
- Source must belong to one of a predefined set of classes
- No firm guarantees = predictive service
- High network utilization (average rather than worst case)
- Ref: Measurement-based admission control

Congestion?



- Congestion = overload on network resources
 - Sigma Demand > Capacity of Resource
- Heterogeneity continues to make congestion control important
- Also configurations where load is not balanced
- Congestion occurs in computer networks even with increase in:
 - buffers, bandwidth and processing power

Traffic Patterns

- Backbones: high speed or low speed?
 - High speed links shared by large numbers of users
 - Mitigates congestion
- Low speed hosts
- Traffic: delay or loss sensitive?
 - Stream: Video conferencing, Telephone
 - Elastic: File transfer, E-mail
 - Interactive graphics/computing
- Telecommunications and data networks merging


Window or Rate?

- Data: TCP/IP = window
- Telecommunication: rate
- Window
 - Designed when memory was bottleneck
 - Back-to-back transmission = bursty traffic
 - Unsuitable for stream-oriented traffic
- Rate
 - Specify burst size and inter-burst arrival
 - Hop-by-hop = need for connections
 - Large queues when input rate close to capacity = feedback required

Open loop or feedback?

- Call, e.g., admission control
- Packet, e.g., scheduling, packet discard
- Performance concerns become on-line
- High speed = propagation delay much higher than packet transmission time
- Number of packets in the "pipe" is high
- Open loop
 - Router-based
 - Reservation
 - Backpressure

End System or Network?



- An important design decision:
 - Division of functionality among hosts and routers
 - Division of functionality among end systems and networks
- Problems with source-based control: large delay, non-cooperative sources, overhead, heterogeneity
- Routers necessary for fairness, but complex and do not avoid congestion
- Source=long time scale, router=short time scale

Backpressure?

- Hop-by-hop
- On/off
- Data-link layer
- Short time scale
 - Or:
- Small networks
- Unfair = everyone affected

Reservation or Walk-in?

- Reservation at setup:
 - Voice/video resources known at setup
 - Data traffic short-lived
 - Gives guarantees
 - Easier to manage resources
- Problems:
 - Low resource utilization?
 - Difficult to predict traffic
 - High overhead and larger time scale

One Scheme or Many?

- Type of scheme depends on duration of overload
- The longer the duration, the higher the layer at which control should be exercised
- No one scheme can solve all congestion problems
- Example: ATM
 - Connection admission
 - Leaky buckets
 - Drop policies
 - Feedback control