Quality of Service

An example of why we need QoS: Jitter

Jitter is variation in network delay

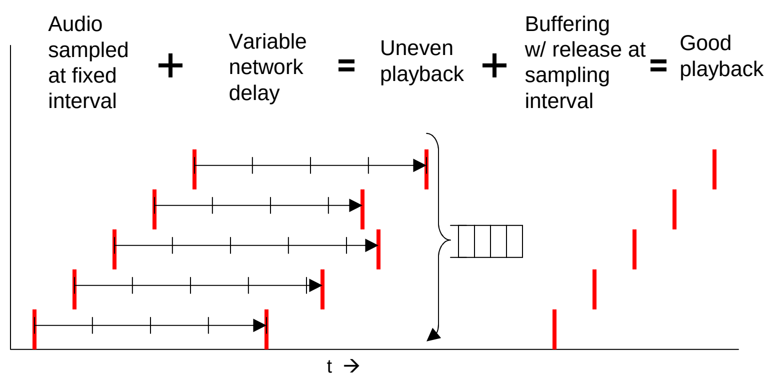

For streaming media, it is usually acceptable to delay the start of the stream (for buffering) in order to get a jitter-free playback. But we can't do this for interactive real-time traffic (e.g, telephony).

It would be good if…

cs570

The University Of Alabama in Huntsville Computer Science

- For real-time interactive traffic, it would be good if we could ask the network to guarantee a bound on jitter
  - Example: “99% of packets will be delayed 24 to 25 msec”
- There are other types of performance guarantees we might like from a network.
  - Priority levels
  - Max delay
  - Constant bandwidth available end-to-end

cs570

The University Of Alabama in Huntsville Computer Science

Approaches to QoS

Overprovisioning

- Provide so much network bandwidth that packets are almost never delayed
- Problems:
  - Cost
  - Load tends to grow faster than capacity

Traffic shaping

The University Of Alabama in Huntsville Computer Science

cs570

- The idea:
  - Regulate the average and burst rates of traffic

The University Of Alabama in Huntsville Computer Science

Bursty traffic patterns are a problem

cs570

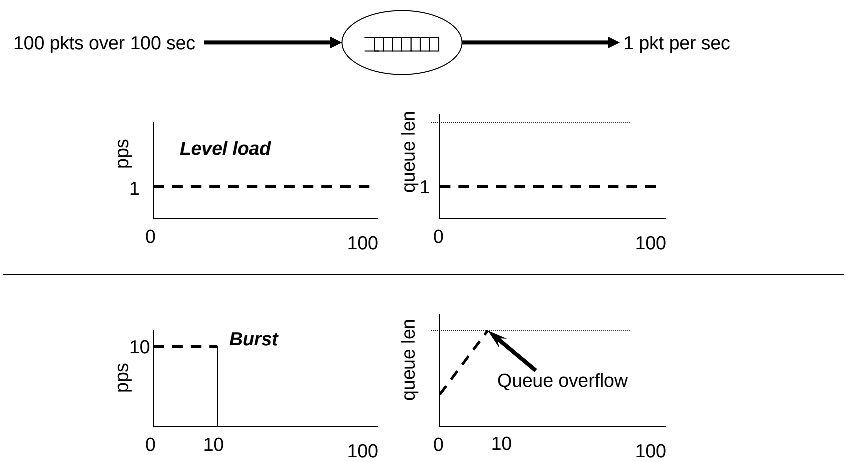

The University Of Alabama in Huntsville Computer Science

Leaky bucket algorithm

cs570

Applied at the sender's network interface to regulate traffic it sends (rate and burstiness)

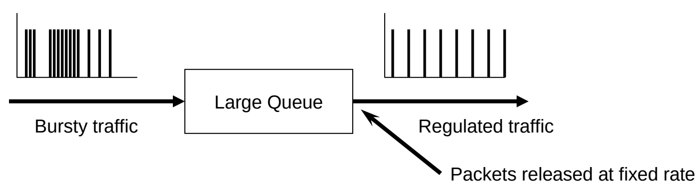

Algorithm ("Byte-counting" Leaky Bucket) – limits traffic bandwidth to n bytes / t sec

Each *t* units of time,
    counter = *n*   (note: *n* must be >= max packet size)
    while *counter* > 0
        If packet at top of queue has fewer than *counter* bytes
            send it
            counter ← *counter* - # bytes in the packet
    end
end

Allowing some burst traffic

cs570

The University Of Alabama in Huntsville Computer Science

- In some cases, it's OK to have short bursts, as long as the long-term behavior is regulated
- Basic Leaky Bucket does not allow this, but a variant does

Token Bucket algorithm

cs570

The University Of Alabama in Huntsville Computer Science

- Allows senders to accumulate permission to send later
- Bursts are allowed, but their length is controlled

Algorithm ("Byte-counting" Token Bucket)

```
counter = 0
Each t units of time,
        counter = counter + n      (n is the # bytes corresponding to one token)
        while counter > 0
                If packet at top of queue has fewer than counter bytes
                        send it
                        counter ← counter - # bytes in the packet
        end
end
```

Terminology: Flow

- In QoS discussions, “flow” is used to describe a stream of data from a given source to a given destination
  - Circuit-switching: flow = data flowing over the circuit – all goes over the same path
  - Packet-switching: the set of packets going from the source to the destination – take different paths

Resource reservation

- The idea: before a flow starts, network resources are reserved to provide it the service it needs
- This is difficult to implement if you are using “pure” packet switching where packets can take any route
- Most connectionless resource-reservation networks (e.g, ATM) use a Virtual Circuit construct where all packets in a flow follow the same route

Elements of resource reservation-based QoS

1. User specifies level of service desired ("Service Spec") and type of traffic it will submit ("Traffic Spec")
2. Network determines if it can support the spec ("Admission Control")
3. Network must manage resources to satisfy the agreement ("Packet Scheduling")
4. Network ensures that the user conforms to the agreement ("Policing")

Service Spec

The University Of Alabama in Huntsville Computer Science

cs570

- Specifies the type of service that the user is requesting
- Examples:
  - Guaranteed max delay ("90% delivered in 100ms")
  - Jitter limitation ("95% delivered between 14 and 15ms")
  - Bandwidth guarantee ("100Mbps end-to-end")
  - Minimization of interference ("controlled load service")
  - Guaranteed space in receive buffer

G. W. Cox -- Fall 2007

QoS

Traffic Spec

- The user's statement of the flow's parameters
- Example:
  - Token bucket rate + size
  - Peak data rate
  - Minimum packet size
  - Maximum packet size

The University Of Alabama in Huntsville Computer Science

Admission control

cs570

- Not standardized
- A general approach:
  - Request (Flow spec + Traffic spec) is forwarded through the network over the (real or virtual) circuit
  - Each router receiving the request inspects it, considering:
    - the router's total capacity
    - commitments already in place
    - operating rules (e.g., "no more than 90% of output link bw committed")
  - The router either:
    - Accepts the request, forwarding it on to the next router, or
    - Declines the request, sending a notice back to the sender node
    - (A variant: instead of declining, the router can reduce the request to what it can support, then forward the reduced request)
  - When the request reaches the destination node, it sends a success notice to the sender (assuming the receiver also accepts the request)

G. W. Cox -- Fall 2007

QoS

Packet scheduling

cs570

- While the flow is active, the network manages the resources to satisfy the agreement made during admission control process
  - Weighted Fair Queuing, etc

The University Of Alabama in Huntsville Computer Science

G. W. Cox -- Fall 2007

QoS

Policing

cs570

- Routers monitor the flow to ensure that original Traffic Spec is not violated
  - Reduce rate using RED
  - Choke packets
  - Etc.

The University Of Alabama in Huntsville Computer Science

G. W. Cox -- Fall 2007

QoS

Differentiated services

cs570

- Flow-based methods have some problems:
  - Complex
  - Advance setup required
  - Not stateless – vulnerable to crashes
- A simpler approach: fixed classes of services defined by net admin ("class-based QoS" or "Differentiated Services")
  - No setup per flow
  - Less complexity at routers
  - Stateless

The University Of Alabama in Huntsville Computer Science

A DS example: Expedited Forwarding

cs570

- The idea: Packets are marked as either "normal" or "expedited". Expedited packets get priority over normal packets (RFC 3246)
- Example: use WFQ to allocate more bandwidth to expedited packets

The University Of Alabama in Huntsville Computer Science

QoS support in IPv6

- 8-bit “Traffic Class” field supports DS
  - A similar field was in IPv4, but it was not widely used
- “Flow Label” field identifies this packet with a particular flow – could be used to support flow-based QoS