Describing network traffic

- Traffic patterns
- Application requirements
- QoS parameters and descriptions

Mixed traffic in the network [1]

- Different applications:
  - traffic (generation) profiles
  - traffic timing constraints
- Routers use FCFS queues:
  - no knowledge of application
  - no knowledge of traffic patterns
- Different traffic types share same network path
- Consider three different applications …

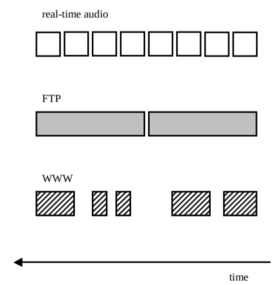

Mixed traffic in the network [2]

- Router:
  - 3 input lines: serviced FCFS at a router
  - 1 output line (1 output buffer)

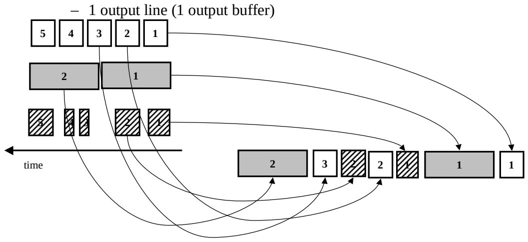

Mixed traffic in the network [3]

- Different traffic patterns:
  - different applications
  - many uses of an application
  - different requirements
- Traffic aggregation:
  - core: higher aggregation
  - many different sources
  - hard to model
- Routing/forwarding:
  - destination-based
  - single metric for all traffic
  - queuing effects
- Large packet size:
  - good for general data
  - “router friendly”
  - “slows” real-time traffic
- Small packet size:
  - good for real-time data
  - less end-to-end delay
  - better tolerance to loss
  - (less jitter?)
  - less efficient (overhead)
  - “not router-friendly”

Elastic applications

Elastic
- Interactive — e.g. Telnet, X-windows
- Interactive bulk — e.g. FTP, HTTP
- Asynchronous — e.g. E-mail, voice-mail

Examples of elastic applications

- E-mail:
  - asynchronous
  - message is not real-time
  - delivery in several minutes is acceptable
- File transfer:
  - interactive service
  - require “quick” transfer
  - “slow” transfer acceptable
- Network file service:
  - interactive service
  - similar to file transfer
  - fast response required
  - (usually over LAN)
- WWW:
  - interactive
  - file access mechanism(!)
  - fast response required
  - QoS sensitive content on WWW pages

Inelastic applications

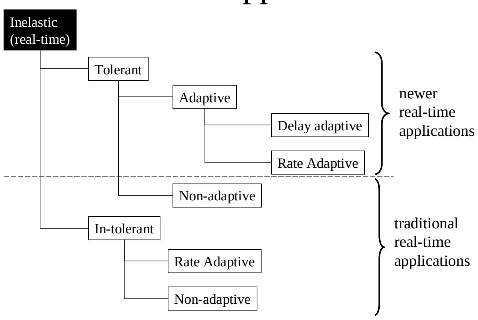

Examples of inelastic applications

- Streaming voice:
  - not interactive
  - end-to-end delay not important
  - end-to-end jitter not important
  - data rate and loss very important
- Real-time voice:
  - person-to-person
  - interactive
- Important to control:
  - end-to-end delay
  - end-to-end jitter
  - end-to-end loss
  - end-to-end data rate

QoS parameters for the Internet [1]

Delay

- Not possible to request maximum delay value
- No control over end-to-end network path
- Possible to find actual values for:
  - maximum end-to-end delay, D_{MAX}
  - minimum end-to-end delay, D_{MIN}

Jitter

- Not possible to request end-to-end jitter value
- Approximate maximum jitter:
  - $D_{MAX} – D_{MIN}$
  - evaluate D_{MIN} dynamically
  - D_{MAX}? 99th percentile?
- Jitter value:
  - transport-level info
  - application-level info

QoS parameters for the Internet [2]

Loss

- Not really a QoS parameter for IP networks
- How does router honour request?
- Linked to data rate:
  - hard guarantee?
  - probabilistic?
  - best effort?
- (Traffic management and congestion control)

Packet size

- Restriction: path MTU
- May be used by routers:
  - buffer allocation
  - delay evaluation

QoS parameters for the Internet [3]

- Data rate:
  - how to specify?
- Data applications are bursty:
  $\frac{\text{peak data rate}}{\text{mean data rate}} >> 1$
- Specify mean data rate?
  - peak traffic?
- Specify peak data rate?
  - waste resources?
- Real-time flows:
  - may be constant bit rate
  - can be variable bit rate
- Application-level flow:
  - **application data unit (ADU)**
- Data rate specification:
  - application-friendly
  - technology neutral

Delay

End-to-end delay

- Propagation:
  - speed-of-light
- Transmission:
  - data rate
- Network elements:
  - buffering (queuing)
  - processing
- End-system processing:
  - application specific

Delay bounds?

- Internet paths:
  - “unknown” paths
  - dynamic routing
- Other traffic:
  - traffic patterns
  - localised traffic
  - “time-of-day” effects
- Deterministic delay:
  - impractical but not impossible

Jitter (delay jitter)

End-to-end jitter

- Variation in delay:
  - per-packet delay changes
- Effects at receiver:
  - variable packet arrival rate
  - variable data rate for flow
- Non-real-time:
  - no problem
- Real-time:
  - need jitter compensation

Causes of jitter

- Media access (LAN)
- FIFO queuing:
  - no notion of a flow
  - (non-FIFO queuing)
- Traffic aggregation:
  - different applications
- Load on routers:
  - busy routers
  - localised load/congestion
- Routing:
  - dynamic path changes

Loss

End-to-end loss

- Non-real-time:
  - re-transmission, e.g.: TCP
- Real-time:
  - forward error correction and redundant encoding
  - media specific "fill-in" at receiver
- Adaptive applications:
  - adjust flow construction

Causes of loss

- Packet-drop at routers:
  - congestion
- Traffic violations:
  - mis-behaving sources
  - source synchronisation
- Excessive load due to:
  - failure in another part of the network
  - abnormal traffic patterns, e.g. "new download"
- Packet re-ordering may be seen as loss

Data rate

End-to-end data rate

- Short-term changes:
  - during the life-time of a flow, seconds
- Long-term changes:
  - during the course of a day, hours
- Protocol behaviour:
  - e.g. TCP congestion control (and flow control)

Data-rate changes

- Network path:
  - different connectivity
- Routing:
  - dynamic routing
- Congestion:
  - network load – loss
  - correlation with loss and/or delay?
- Traffic aggregation:
  - other users
  - (time of day)

Network probing: a quick note

- Can use probes to detect:
  - delay
  - jitter
  - loss
  - data rate
- Use of network probes:
  - ping
  - traceroute
  - pathchar
- Probes load the network, i.e the affect the system being measured
- Measurement is tricky!
- See:
  - www.caida.org
  - www.nlanr.net

Perceived QoS

- Consider application-specific (media-specific) features
- Real-time – VoIP:
  - packet loss rate of 1/20 for voice
  - what if the first or last phoneme is lost?
  - losing the start of a word leads to lower perceived QoS!
  - other factors (jitter, delay)
- Streaming:
  - can cope with loss by buffering at the receiver
  - what about data rate?
- For example – video:
  - low data rate
  - small picture size
  - low refresh (e.g. 3fps)
  - low colour depth
  - **OK** for adverts, news reels
  - **Not OK** for entertainment

Interactive, real-time media flows

- Audio/video flows:
  - streaming audio/video
  - use buffering at receiver
- Interactive real-time:
  - only limited receiver buffering
  - delay + jitter <150ms
  - (jitter <150ms)
  - keep loss low
- Effects of loss:
  - depend on application, media, and user
- Audio:
  - humans tolerant of "bad" audio for speech
  - humans like "good" audio for entertainment
- Video:
  - humans tolerant of "low" quality video for business
  - humans like "high" quality video for entertainment
- Audio – video sync:
  - separate flows?

Audio

QoS requirements

- Delay < 150ms:
  - including jitter
- Low loss preferable:
  - loss tolerant encodings exist
- Data rates:
  - speech ≤ 64Kb/s
  - "good" music ≥ 128Kb/s

- Time domain sampling
- Example – packet voice:
  - 64Kb/s PCM encoding
  - 8-bit samples
  - 8000 samples per second
  - 40ms time slices of audio
  - 320 bytes audio per packet
  - 48 bytes overhead
    (20 bytes IP header)
    (8 bytes UDP header)
    (20 bytes RTP header)
  - 73.6Kb/s

Video

QoS requirements

- Delay < 150ms:
  - including jitter
  - same as audio
  - inter-flow sync
- Loss must be low
- Data rate – depends on:
  - frame size
  - colour depth
  - frame rate
  - encoding

- Frequency domain:
  - discrete cosine transform (DCT)

Summary

- Different applications have different needs
- Some QoS requirements are application-specific and media-specific:
  - perceived QoS
- Different requirements for real-time multimedia and streamed multimedia