Quality of Service in Integrated Networks

Henning Schulzrinne
GMD Fokus, Berlin
hgs@fokus.gmd.de

Tubkom'94/95 Berlin

February 1, 1995

Overview

- definitions, requirements
- ATM and QOS
- QOS in heterogeneous networks
- protocols for measuring QOS: RTP
- assuring QOS: jitter compensation
- adaptive applications

Definitions of Quality of Service

The revised *Devil's Dictionary* says:

Quality of Service, n., liberal use of term assures that your RACE/ACTS/. . . proposal gets funded. Use together with ATM and multimedia for best effect.

Applications

Integrated networks should handle at least:

- file transfer, remote backup, . . . (ftp, FTAM)
- electronic mail (smtp, X.400)
- remote login (telnet, rlogin)
- directory access (DNS, X.500)
- remote control, meter reading, . . .
- interactive data (irc, telnet)
- non-real time multimedia (hypermedia, e.g., WWW; video on demand, teleshopping, . . .)
- real-time multimedia = interactive audio, video
- . . .

QOS/Service Model

- applications have a network worldview
- technology changes, but application view remains
- few, well-described, broadly applicable classes
- useful for application
- implementable in switches, routers
- measurable
- source characteristics ↔ service requirements

QOS

- application view:
  - frame rate
  - video quality (SNR, quantization level)
  - file transfer throughput
  - query/keyboard response time
  - . . .
- network view – exhaustive:
  - throughput (bytes/sec)
  - delay (seconds)
  - cell/frame/packet loss (fraction) → delay by retransmission
  - delay variation (jitter) → delay by jitter compensation

QOS Description

bounds: "never $>$" – hard to assure with failures

probabilities: hard to contest 10^{-9}, interval?

Application Requirements

Holy grail: characterize application requirements on networks

Unfortunately, . . .

- most data, many real-time applications depend on network performance (closed loop)
- usage depends on price, responsiveness, applications, . . .
- aggregate can change quickly with new applications (WWW)

QOS Requirements

application	throughput	resp. time	errors
audio	fixed	40 ms	1%
video	(fixed)	500 ms	10^{-9} – 10%
telnet, rlogin	x	100 ms	none
ftp	xx	seconds	none
smtp (mail)	–	minutes	none
http (WWW)	xx	ms	none

In case of difficulty . . .

Choices if network is “full”:

1. reserve ahead (large-scale teleconferences)
2. displace low-fare customers and pay more
3. wait until resource available
4. “don’t call us, we’ll call you”

Admission Control

Need:

- number(s) representing application
- work for homogeneous and mixes
- additive
- efficient: full means full
- computable from user-level characteristics
- *equivalent bandwidth*

→ doesn't exist without approximations

The Fundamental Laws of Networks

1. you can't store bandwidth for later
2. sending an additional bit costs zero (incremental cost)
3. network costs are peak capacity costs (even more extreme than trains)
4. in a free market, prices reflect incremental cost

→ send bits as soon as possible, as fast as possible → ATM ABR service

ATM: The Short Version

We'll see ATM often, so recap. . .

- cell switching: fixed 48 byte payload, 5 byte header
- cells combined into frame by *ATM adaptation layer* (AAL)
- speed independent: 25 Mb/s to 2.4 Gb/s used
- no hop-by-hop flow control, loss recovery, error detection
- connection oriented: virtual circuit (VC) $\in$ virtual path (VP)
- VCI/VPI are local: translated by each switch
- switches operate on VCs or VPs
- cell-loss priority (CLP) bit: in case of emergency, discard first
- ISDN-like signaling to set up connections

The ITU Model

Initial view driving AALs:

	class A	class B	class C	class D
timing	required		not required	
bit rate	constant	variable		
mode	connection oriented			CL
AAL	1	2	3	4
category	CBR	VBR	UBR/ABR	

- mixes source characteristics and requirements
- only AAL1, AAL3/4, AAL5 used; AAL 3/4 future iffy
- no quantitative measures

ATM Forum

Traffic type and service categories:

	description	application
CBR	constant bit rate	T1 emulation
VBR	variable bit rate	video
UBR	unspecified bit rate	data, loss-tolerant
ABR	available bit rate	data, loss-sensitive

Internet Service Model

currently: best-effort, FIFO

Scott Shenker, Xerox (IETF Integrated Services Working Group):

	uncontrolled delay	controlled delay
uncontrolled bandwidth	email	video
controlled bandwidth	bulk data	audio

→ no loss as parameter

Uncontrolled Delay, Bandwidth

elastic applications:

- can gracefully adapt to variable per-packet delays
- can gracefully adapt to variable bandwidth
- e.g., file transfer, WWW, . . .

proposed service: ASAP (= best effort)

- application view: 1 bit/hour better than rejection
- multiple classes of ASAP service (email, remote backup vs. WWW)
- not subject to admission control
- issues: fairness, congestion notification

Uncontrolled Delay, Controlled Bandwidth

Time-constrained bulk transfers:

- want as much as is available
- can gracefully adapt to variable per-packet delays
- can gracefully adapt to variable bandwidth, but need an assured minimal amount of bandwidth

Proposed service: assured ASAP

- assured minimal available bandwidth
- must be subject to admission control
- issues: fairness, congestion notification

Controlled Delay, Controlled Bandwidth

Standard real-time applications:

- need bound on per-packet delays
  - nonadaptive: setting playback point
  - adaptive: a priori estimation of delays
- have an intrinsic bandwidth requirement
- either tolerant or intolerant of occasional late packets (→ loss)

Proposed service: predictive and guaranteed

- guaranteed: worst-case upper bound
- predictive: fairly reliable upper bound
- both give propagation delay lower bound
- must be subject to admission control

Controlled Delay, Uncontrolled Bandwidth

Variable quality real-time applications:

- need bound on per-packet delay
- can function at different rates
- assumed tolerant of rare violations of delay bounds
- example: video

Proposed service: variable bandwidth predictive

- fairly reliable upper bound on delay
- variable bandwidth with assured minimum
- subject to admission control
- related to ATM ABR

ATM Traffic Control

Idea: Assure QOS by bounding traffic

- policing or shaping
- policing: mark (CLP bit) or drop
- non-compliant

ATM and QOS

Myth: ATM supports QOS.

Reality → needed for predictable QOS:

- ATM is very loss sensitive: one cell gone → frame gone
- network adapters: traffic shaping/policing → limit rate, burstiness
- switches: isolation of traffic flows by
  - priorities: e.g., CBR, ABR, VBR?
    - loss priority: who loses if buffer is full
    - delay priority: who goes first (important?)
  - round robin
  - weighted fair queueing

ATM and QOS

- fair queueing: bit-by-bit round robin → order transmission when the last bit would have been sent
- weighted fair queuing → unequal bandwidth allocation
- WFQ + leaky bucket → delay bound
- Need large buffers (currently very small: 300 – 1000)

Reality: 2 priorities, if lucky.

ATM Forum QOS Parameters

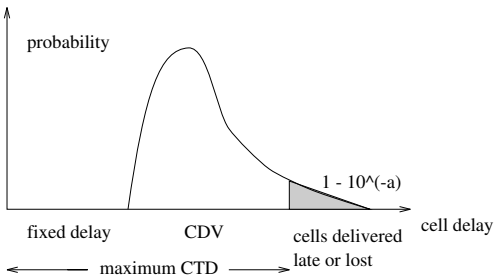

fixed delay: speed of light (5 μs/km), switching (10 μs/switch)

CDV (cell delay variation: queueing delays in switches, NICs

CDT (cell delay tolerance: fixed delay + peak-to-peak CDV

ATM Forum QOS Parameters

parameter	CBR	VBR		ABR	UBR
		real-time	non RT		
p-p CDV	X	X	–	–	–
max. CTD	X	X	?	–	–

Simple QOS Network Model

Two classes may be sufficient:

1. high priority, peak rate allocation → circuit emulation, audio, CBR/VBR video, guaranteed bandwidth data
2. lower priority, fair sharing → "data" applications

Reasonable if 1st class leaves enough room ($\sum peak \ll$ capacity) and 2nd class can fill the gaps.

Old debate: ATM =

- adjustable-rate bit pipes with fine granularity
- packet switching network

ATM Forum Traffic Parameters

PCR	$R_p = 1/T$	peak cell rate
SCR	$R_s = 1/T_s$	sustainable cell rate (average)
MCR		minimum cell rate

PCR: separate for CLP=0 and CLP=0+1.

GCRA: Generic Cell Rate Algorithm

PCR, MCR, SCR defined by GCRA (jitter!):

TAT: theoretical (nominal) arrival time

I**:** interval (1/cell rate) = T

L**:** limit, determines burstiness = τ

non-conforming:] tag (CLP=1) or drop

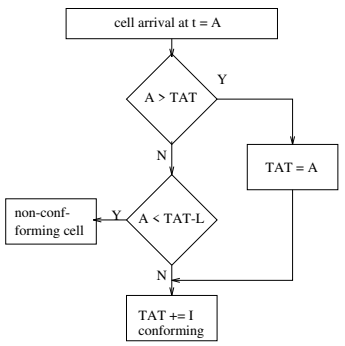

$$\text{burst size} = \left\lfloor 1 + \frac{\tau}{T - \tau} \right\rfloor$$

GCRA for Video

Usage parameter control at network access point: trade L for I

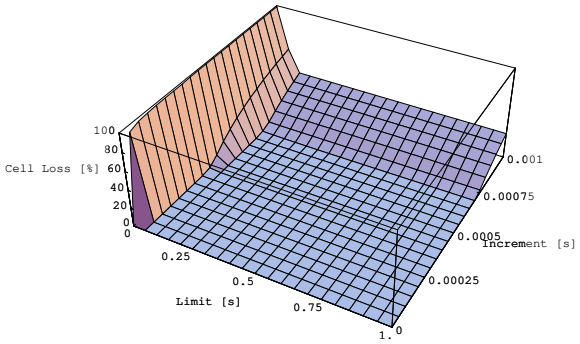

traffic shaping: trade delay for peak rate

Playout Buffer

- Motivation
  - networks, operating systems introduce delay variations between sender and receiver (jitter)
  - audio and video streams have to be resynchronized at the receiver
- Approach
  - playout buffer introduces additional delay to compensate for delay variations
  - playout delay is adjusted dynamically during talkspurts
  - conference (interactive) mode: low delay, some late losses
  - lecture mode: longer delay, few losses
  - use simulation

Estimation of End-End Delay

- clocks are not synchronized
- unknown delays in operating system
- → network delay measurement using cross correlation of audio signals

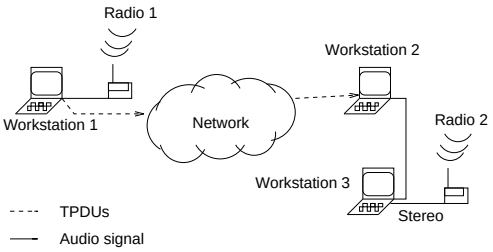

Estimation of End-End Delay

Cross correlation of a recorded stereo signal

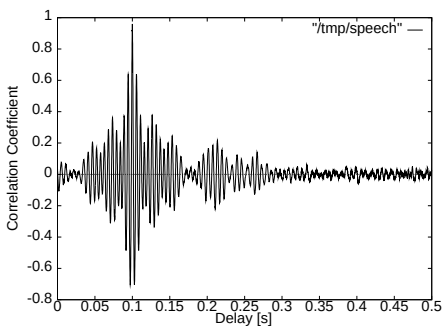

Application Control: Why?

- multimedia traffic is . . .
  - long-lived (hours instead of 3 minutes)
  - high host bandwidth (2 Mb/s instead of 64 kb/s)
  - multiparticipant
  - flexible within bounds
- network costs are peak costs (see Basic Law)
- Hypothesis: most users would rather get lower quality than be turned away
- better to lower frame from 30 to 27 fps than turn away 10% of calls
- → give best possible QOS
- no preference for early risers

Application Controlability

Adjust bandwidth of M-JPEG application using Q factor:

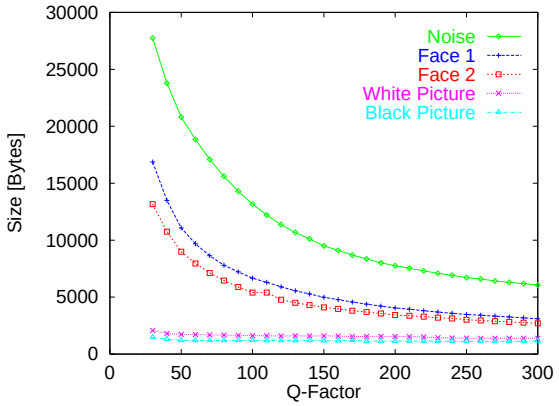

- scales linearly with frame rate, image size
- audio: encoding, number of channels, sampling rate, . . .
- generally: factor 2 to 3 usable range

Application Control: What?

- Idea: provide service class with adjustable QOS
- used to electric tariffs for water heaters
- similar to “brownouts”
- used by electric utilities for blackout prevention
- must work for multicast
- must scale to 10^7 applications

Aside: RTP

RTP: real-time transport protocol has QOS feedback channel

- packet loss → congestion
- delay variation (jitter)
- approximate end-to-end delay

Via ulticast to all session participants

Application Control: How?

Several options:

- application → network manager
  - scales poorly
  - doesn't localize
- RTP QOS options via multicast
  - multicast for all

Application Control: How?

- RTP QOS (or similar) via unicast to sender
  - implemented (roughly) in ivs video tool
  - need multicast anyway to estimate group size
  - others may need network QOS for fault detection
- ATM EFCI, RM (resource management) cells
  - similar to ABR
  - maybe same mechanism in switch?
  - adjustment speed?
  - multicast?

Conclusion

- service models
- fundamental network measurements
- traffic estimation is inherently hard
- want closed-loop control, for multimedia, too
- want simple metrics for applications
- if right traffic mix, details may not matter
- ABR: fairness, utilization, buffer size