

Networks with Deterministic Quality-of-Service Guarantees

Jörg Liebeherr


Department of Computer Science

University of Virginia

jorg@cs.virginia.edu

Motivation

- Transmission of video and audio over packet-switched networks.



- Requires *new networks and protocols*.


Overview

- Background
- *Traffic Characterization*
- *Packet Scheduling*
- Conclusions

Quality-of-Service

- Video and audio need *Quality-of-Service (QoS)* guarantees:
 - *delay*
 - jitter
 - throughput
 - loss rate
- A *deterministic service* gives worst-case guarantees.


Multimedia Networks



- Multimedia connections have *QoS* and *traffic* parameters.
- Multimedia networks need *resource reservation*.


Why is Resource Reservation Difficult?

- Compressed digital video has a *variable bit rate*.





- *Problem:* How do we provide deterministic QoS without peak-rate reservation?

Design Space of a Multimedia Network






What is Traffic Characterization?



- A *traffic characterization* is a bound for the traffic over any interval.
 - *Time-invariant*: $A^*(t) \geq A[\tau, \tau + t]$, $\forall t, \tau$
 - *Subadditive*: $A^*(t_1 + t_2) \leq A^*(t_1) + A^*(t_2)$, $\forall t_1, t_2$
- Traffic characterization must map to traffic policer.

The "Leaky Bucket" Traffic Characterization




- Used in: ATM, Integrated-services Internet

Traffic Characterization Problem


- Given a video sequence, how do I select leaky bucket parameters?
- Previous approaches:
 - *Candidate Sets* (Low and Varaiya 1991).
 - Choose B according to buffer space availability (Pancha and El Zarki 1995).
 - Relative importance of buffer space and bandwidth (Guillemin et. al. 1995).
 - *Empirical envelope* (Wrege, Knightly, Zhang, and Liebeherr 1996).

Empirical Envelope



- The best possible characterization for a video source is its *empirical envelope* E^* .
- $E^*(t) := \sup_{\tau \geq 0} A[\tau, \tau + t]$, for all $t \geq 0$.
- *Difficult to police, expensive to compute.*

Our Approach



- Approach: approximate the empirical envelope E^* .
- Use only a subset of E^* .
- Select leaky bucket parameters.


Evaluation

- How much information do we need from the envelope?
- How good is our approximation?

Experimental Setup


- Single 155 Mbps switch.
- Characterize a “typical” MPEG-compressed traffic source.
- Frame pattern: IBBPBBPBBPBB
- Video frames partitioned into 53-byte cells.

How much of E^* do we need?



- *200 points* of the envelope are sufficient.
- Empirical envelope has 40,000 points.

How Good is Our Method?



- We plot a normalized utilization $U(d) = \#A^* / \#E^*$.

Design Space of a Multimedia Network


**Packet
Scheduling**



Traffic
Characterization


**Admission
Control**

Packet Scheduling



- A connection j has a *delay bound* d_j .
- *Packet scheduling discipline* determines *delay*.

What is a good scheduler?



Approximate EDF with FIFO queues

Approximations that require *no sorting*:


- *HOL-PJ* (Lim/Kobza 1990)
- *Relabeling Architecture* (Peha/Tobagi 1991)
- *Rotating-Priority-Queues* (RPQ) (Liebeherr/Wrege 1994)

Rotating-Priority-Queues⁺ (RPQ⁺)


Design Principles:

- P priority sets.
- $2P$ FIFO queues with labels.
- Relabel queues every Δ time units.
- One delay bound for each priority set: $d_p = p \cdot \Delta$.



RPQ⁺ Scheduler



RPQ⁺ Queue Rotation



Implementating RPQ⁺ in Shared Memory



- No movement of packets.
- Operations independent of queued packets.

Admission Control Test for RPQ⁺

For all priorities p and all $t \geq d_p$,


$$t \geq \sum_{q=1}^{p-1} \sum_{j \in C_q} A_j^*(t - d_q + \Delta) + \sum_{q=p}^P \sum_{j \in C_q} A_j^*(t - d_q) + \max_{r, d_r > t} s_r^{max}$$

Experimental Setup


- Single 155 Mbps switch.
- Three connection groups *Low*, *Medium*, *High Delay*.

	Index	Delay Bound	Burst Size	Rate
	j	d_j	B_j	r_j
<i>Low</i>	1	12 ms	4,000 cells	10-155 Mbps
<i>Medium</i>	2	24 ms	2,000 cells	10-155 Mbps
<i>High</i>	3	36 ms	4,000 cells	10-155 Mbps

Evaluation




EDF





SP

Evaluation




RPQ+ ($\Delta = 12ms$; 6 FIFOs)

RPQ+ ($\Delta = 6ms$; 12 FIFOs)



RPQ+ ($\Delta = 4ms$; 18 FIFOs)

RPQ+ ($\Delta = 3ms$; 24 FIFOs)



Conclusions

- Relax deterministic service.
- Implement RPQ^+ for IP forwarding.
- Combine advantages of delay schedulers (EDF, RPQ) and rate schedulers (WFQ).
- Reading:
 - IEEE/ACM Transactions on Networking, December 1996.
 - Proc. IEEE Infocom '96, San Francisco, March 1996.
 - Proc. IEEE Infocom '97, Kove, April 1997.