A Better-Than-Best-Effort Service For Responsive UDP Flows

Kevin Jeffay
F. Donelson Smith

Mark Parris
Jan Borgersen

Department of Computer Science
University of North Carolina at Chapel Hill

http://www.cs.unc.edu/Research/dirt

A Better-Than-Best-Effort Service For Responsive UDP Flows

Outline

- The case for "better-than-best-effort" services
- The INTSERV & DIFFSERV models
- Principles of active queue management
- Extensions to RED for a better-than-best-effort UDP service

A Better-Than-Best-Effort Service For Responsive UDP Flows

Outline

- The case for "better-than-best-effort" services
  - » Application domain(s) of interest
  - » The performance of interactive applications on the Internet today

UNC Multimedia Networking Research

System support for low latency, continuous media transmission

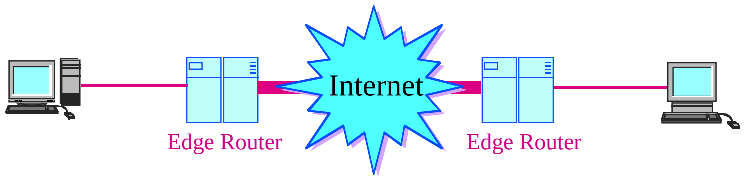

- Focus on real-time media transmission
  - » Periodic media generation (30 Hz or better)
  - » 250 ms (or better) one-way end-to-end latency
  - » Variable levels of loss tolerance
- Applications
  - » Interactive entertainment
  - » Distributed virtual environments
  - » Collaboration support

UNC Multimedia Networking Research

Driving problem

- **The *nanoManipulator* system**
  - A virtual environment interface to a scanning-probe microscope
  - Provides *telepresence* on sample surfaces scaled 1,000,000:1

UNC Multimedia Networking Research

nanoManipulator

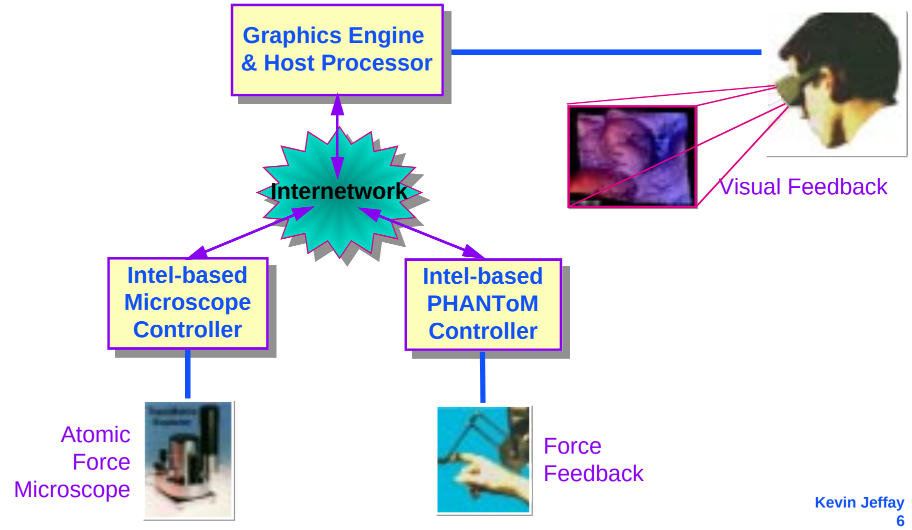

UNC Multimedia Networking Research

nanoManipulator

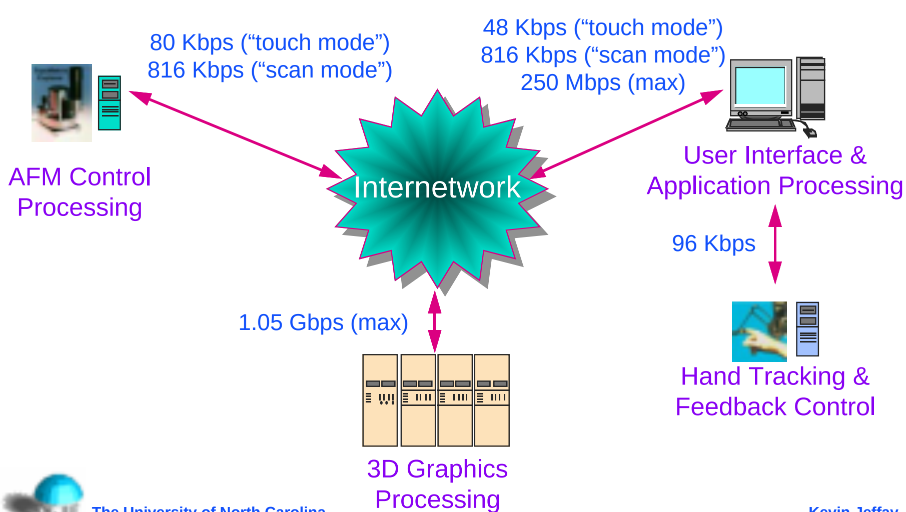

UNC Multimedia Networking Research

OS & network support for the "last mile problem"

- **Operating principle:**
  - Network elements that cannot reserve, or support real-time allocation of resources, will persist for the foreseeable future.
- **Focus on adaptive, best-effort transmission...**
  - Treat the network as a black box — *Assume only that sufficient bandwidth exists for some useful execution of the system*
- **... with real-time media control at the endpoints**

UNC Adaptive Congestion Control
2-Dimensional media scaling for videoconferencing

- **Canonical approach to congestion**
  - » Reduce (video) bit-rate
- **Our approach**
  - » View congestion as a search of a 2-dimensional *bit-rate* x *packet-rate* space
  - » Scale stream bit- and packet-rates simultaneously to find a sustainable *operating point*

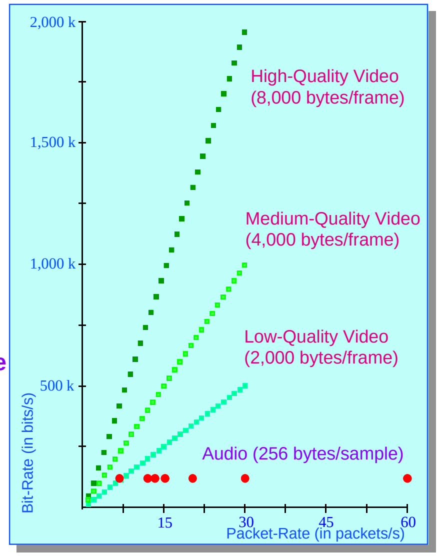

Bit- and Packet-Rate Scaling
An analytic model of media scaling

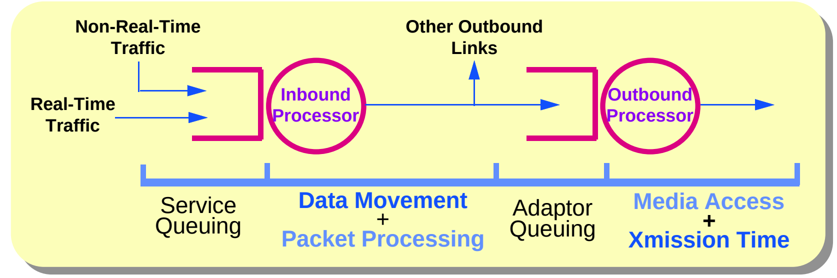

- **Capacity constraints**
  - » The network is incapable of supporting the desired bit rate in any form
- **Access constraints**
  - » The network can not support the desired bit rate with the current packaging scheme

Two Types of Congestion Constraints
Two dimensions of adaptation

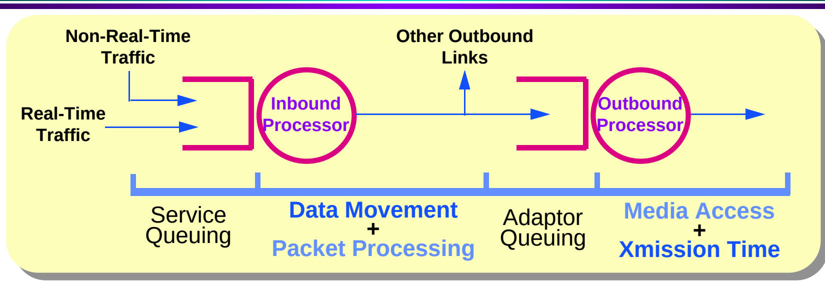

- **Reduce the packet-rate to adapt to an access constraint**
  - » Change the packaging or send fewer video frames
  - » Primary Trade-off: higher latency (potentially)
- **Reduce the bit-rate to adapt to a capacity constraint**
  - » Send fewer video frames or fewer bits per video frame
  - » Primary Trade-off: lower fidelity

Adaptive, 2-Dimensional Media Scaling
Does it work?

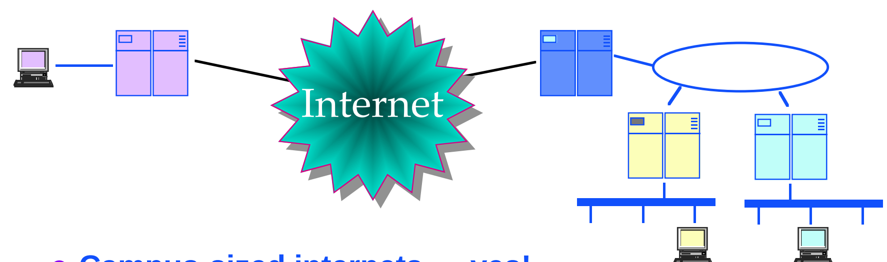

- **Campus-sized internets — yes!**
  - » It “solves” the first-mile/last-mile problem
- **The Internet? — *well...***
  - » Does our necessary condition for success hold?
  - » Does it hold often enough to be useful?

Media Scaling Evaluation on the UNC Campus

Performance with no media scaling

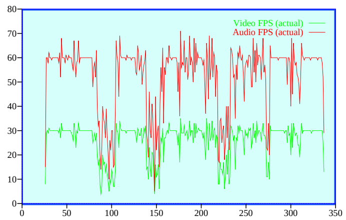

Throughput (frames/sec)

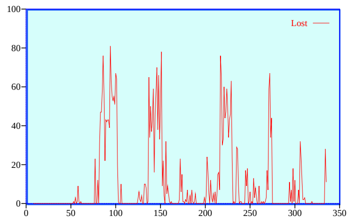

Packet Loss

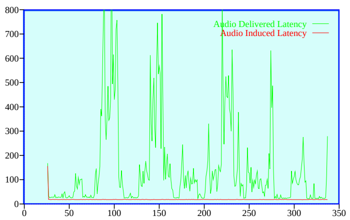

Audio Latency (ms)

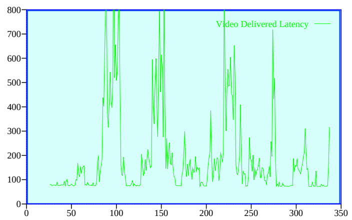

Video Latency (ms)

Media Scaling Evaluation on the UNC Campus

Performance with video scaling only

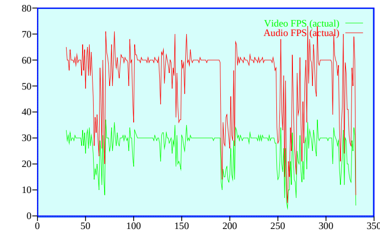

Throughput (frames/sec)

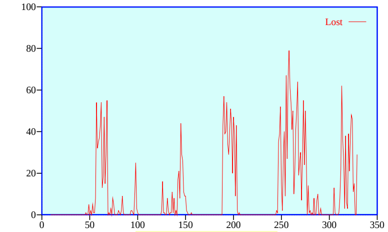

Packet Loss

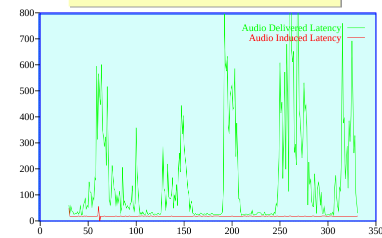

Audio Latency (ms)

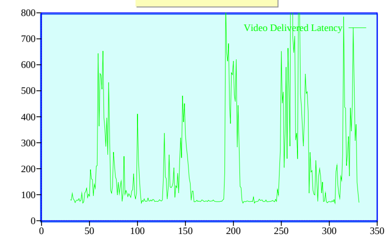

Video Latency (ms)

Media Scaling Evaluation on the UNC Campus

Performance with 2-dimensional scaling

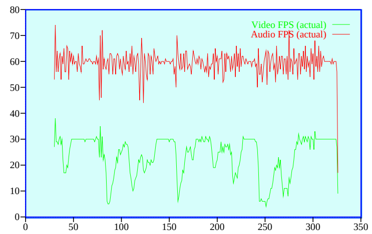

Throughput (frames/sec)

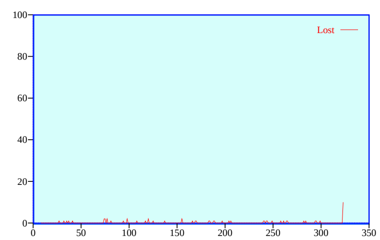

Packet Loss

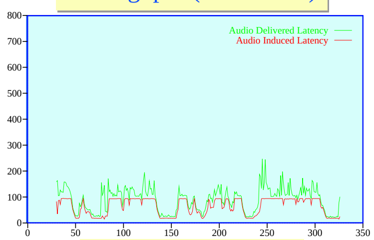

Audio Latency (ms)

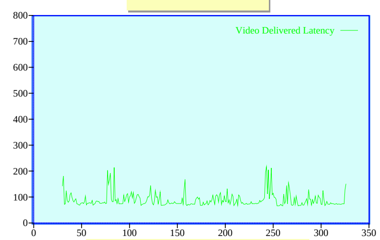

Video Latency (ms)

Media Scaling Evaluation on the Internet

ProShare with no media scaling

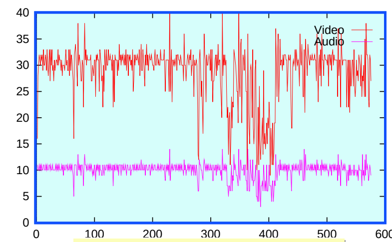

Throughput (frames/sec)

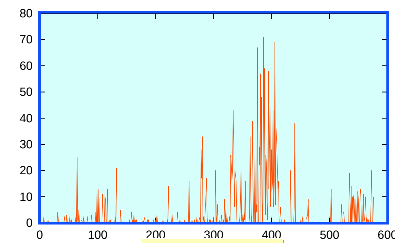

Packet Loss

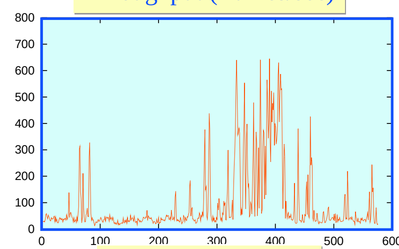

Audio Latency (ms)

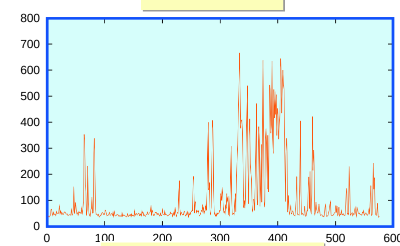

Video Latency (ms)

Media Scaling Evaluation on the Internet

ProShare with 2-dimensional media scaling

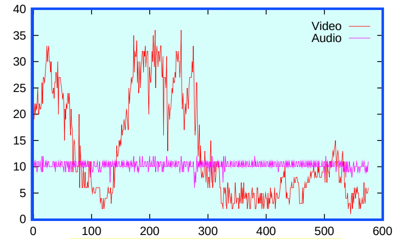

Throughput (frames/sec)

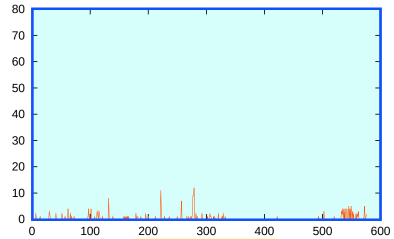

Packet Loss

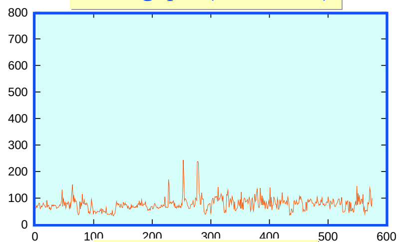

Audio Latency (ms)

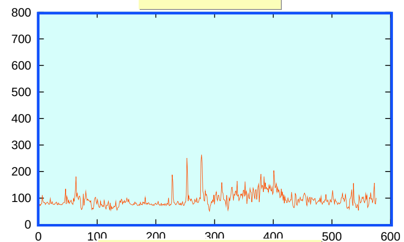

Video Latency (ms)

Adaptive, 2-Dimensional Media Scaling

Sustainability results

- Results of an Internet performance study between UNC and UVa
  - » Repeated trials from 10 am to 7 PM weekdays
  - » Trials separated by at least two hours
  - » Scattered over three months

Time Slot	Sustainable	Not Sustainable
10:00-12:00	67%	33%
12:00-14:00	50%	50%
14:00-16:00	8%	92%
16:00-18:00	25%	75%
18:00-20:00	44%	56%
Percentage	39%	61%

A Better-Than-Best-Effort Service For Responsive UDP Flows

Outline

- The case for "better-than-best-effort" services
- **The INTSERV & DIFFSERV models**
- Principles of active queue management
- Extensions to RED for a better-than-best-effort UDP service

The Integrated Services Architecture for the Internet

Reference implementation components

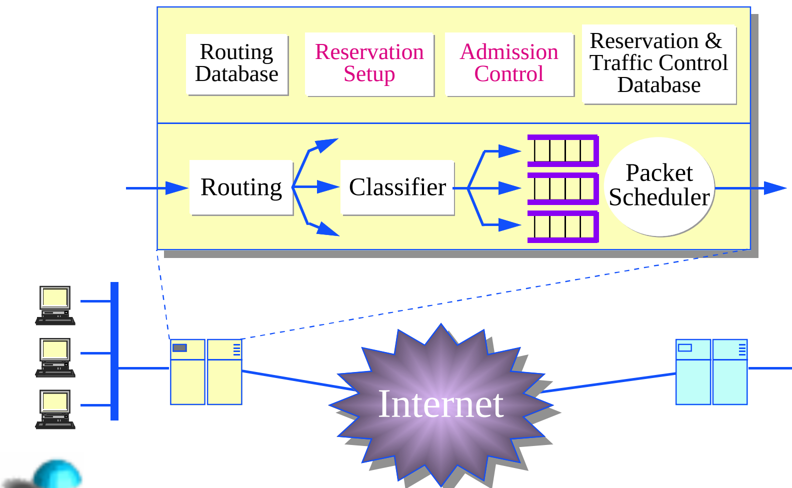

Differentiated Services
Clark *et al.*'s "expected capacity" service

- ISPs allocate and sell capacity for an "assured" service
- Senders/border routers mark packets according to "service profiles"
  - » Simple IN/OUT marking
  - » Markings indicate "drop preference"

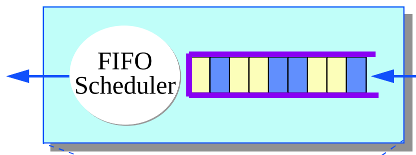

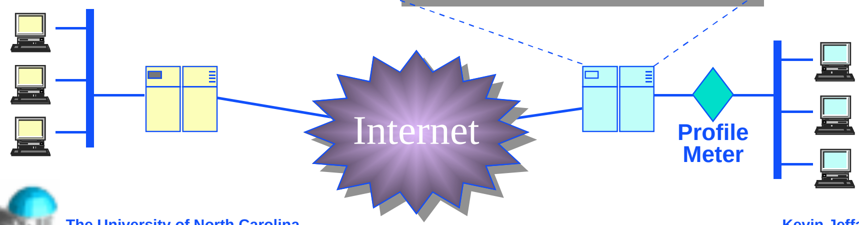

Clark *et al.*'s "Expected Capacity" Service
Realizing differentiated service

- Routers maintain counts of IN and OUT packet populations
  - » OUT packets probablistically dropped when queue population exceeds min threshold
  - » IN packets probablistically dropped when IN packet queue population exceeds (separate) min IN threshold

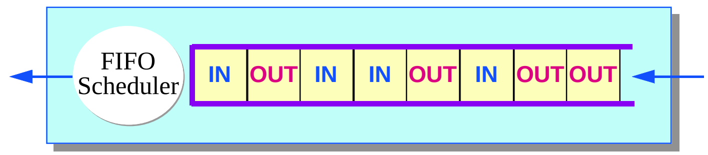

- Key assumption: Flows are "well-behaved"

Differentiated Services
Jacobson *et al.*'s 2-bit differentiated service

- Routers maintain a separate queue for a low-delay, low-jitter "premium" service
- Senders/Edge routers again mark packets
  - » "Premium" service
  - » "Assured" service
  - » Best-effort (default)

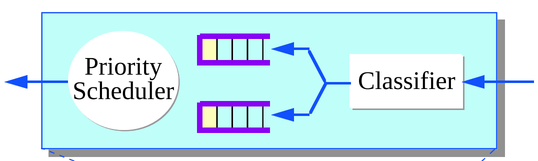

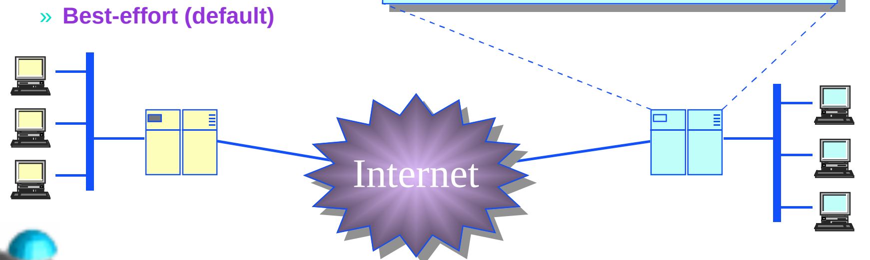

Jacobson *et al.* 2-bit differentiated service
Queue management

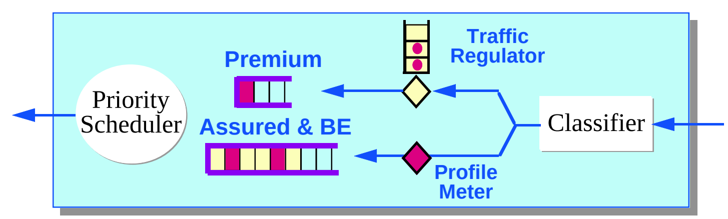

- Capacity explicitly allocated for Premium traffic
  - » Premium flows shaped to eliminate bursts
- Assured and best-effort traffic share a queue
  - » Unmarked packets randomly discarded when queue exceeds a threshold
  - » Assured packets randomly discarded when assured population exceeds a threshold

Active Queue Management
Random Early Detection (RED)

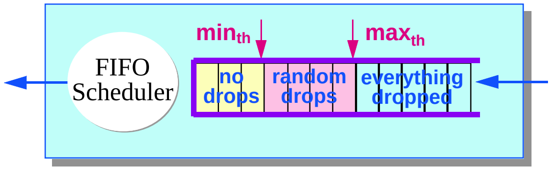

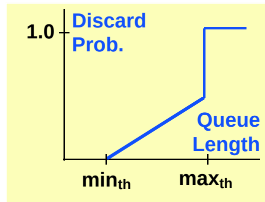

- **Basic mechanism for realizing differentiated services is a RED (random early discard) congestion avoidance mechanism**
- **Powers that be advocate that RED be deployed today**
  - » Protects the network from congestive collapse
  - » Increase effective network utilization
  - » Decrease end-to-end latency

Active Queue Management
Impact of RED on multimedia flows

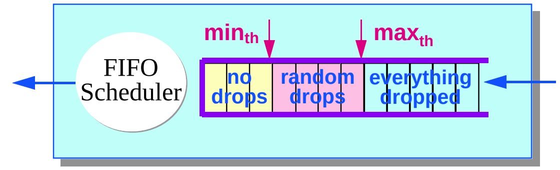

- **Many RED variants...**
  - » Clark et al.'s RED-IN-OUT (RIO)
  - » Floyd & Fall's "RED with Penalty Box"
  - » Lin & Morris "Flow RED" (FRED)
- **... most view UDP as "evil"**
  - » "non-responsive" and/or "non-conformant" flows penalized

Active Queue Management
A "better than best-effort" service for UDP

- **What can be done to improve responsive UDP flow performance with-out sacrificing TCP performance?**
  - » Per flow "threaded queues" with "drop head" discard semantics
  - » CBQ emulation for UDP flows with bounded queues

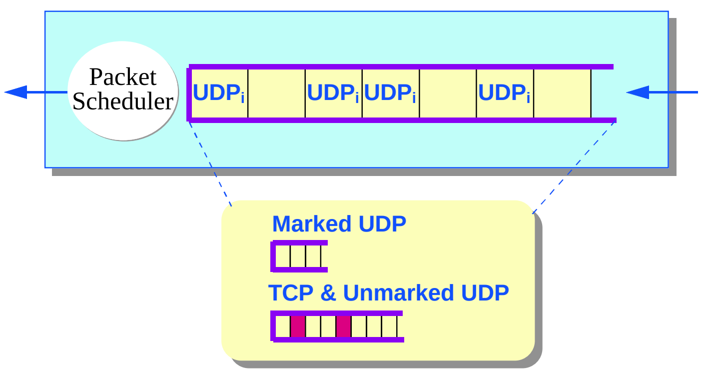

A better-than-best-effort service for UDP
Status

- **Implementation exists in a FreeBSD router**
  - » Using Alt-Q RED implementation
- **Traffic generation engines developed**
- **Early Experimental results promising**

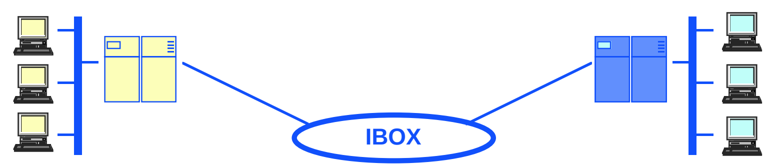

Summary

- Proposals for QoS within the Internet are coming
- In the meantime UDP flows are prime targets for network-based congestion avoidance
  - » This will remain true when INTSERV/DIFFSERV deployed
- We are working to define a simple packet forwarding behavior that will result in a better-than-best effort service for responsive UDP flows