University of North Carolina at Chapel Hill

Beyond Audio & Video

Multimedia Networking Support for Distributed, Immersive Virtual Environments

Kevin Jeffay
Thomas Hudson
Mark Parris

EUROMICRO, September 2001

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

Multimedia Networking

Beyond Audio and Video

- Support for real-time delivery of audio and video on the Internet was the "killer app" of the 1990's
- What did we learn?
  - Per flow bandwidth/delay guarantees are too hard to support and are likely unnecessary
  - If the Internet could be made to act like a "lightly loaded LAN" then end-system media adaptation was sufficient
- So what's next? What are other interesting continuous-media (CM) flows to study?
  - How well do network mechanisms and end-system adaptations scale to meet the requirements of "next generation" CM applications?

Beyond Audio and Video

Support for distributed virtual environments

- Goal: Use resources distributed across the Internet to provide users with a sense of immersion in a virtual world

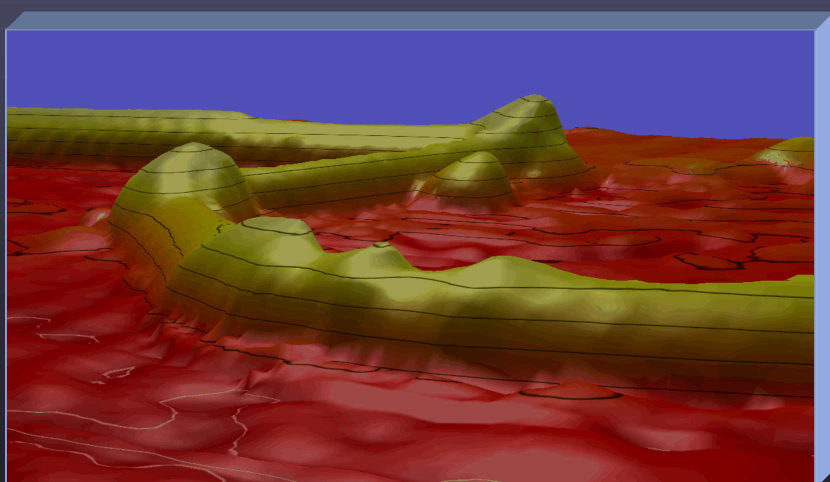

Distributed Virtual Environments

Distributed virtual laboratories

- Computers and computer interfaces are fundamental to modern scientific instruments

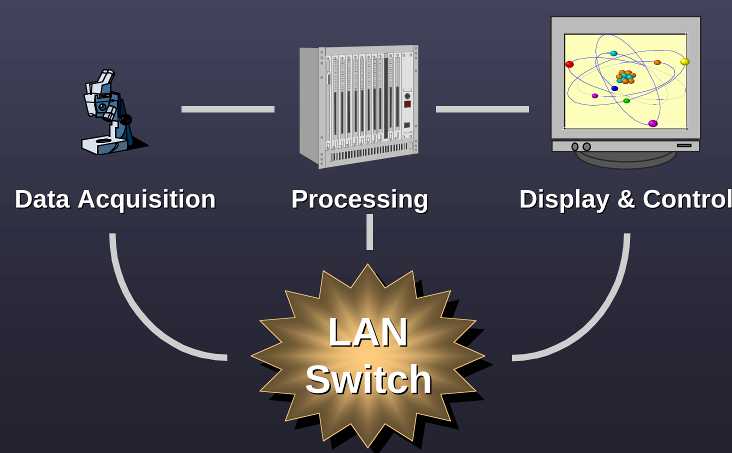

Distributed Virtual Environments
The UNC nanoManipulator system

- A virtual environment interface to a scan-probe microscope
- Provides telepresence on sample surfaces scaled 1,000,000:1

The UNC nanoManipulator
Atomic force microscopy simplified

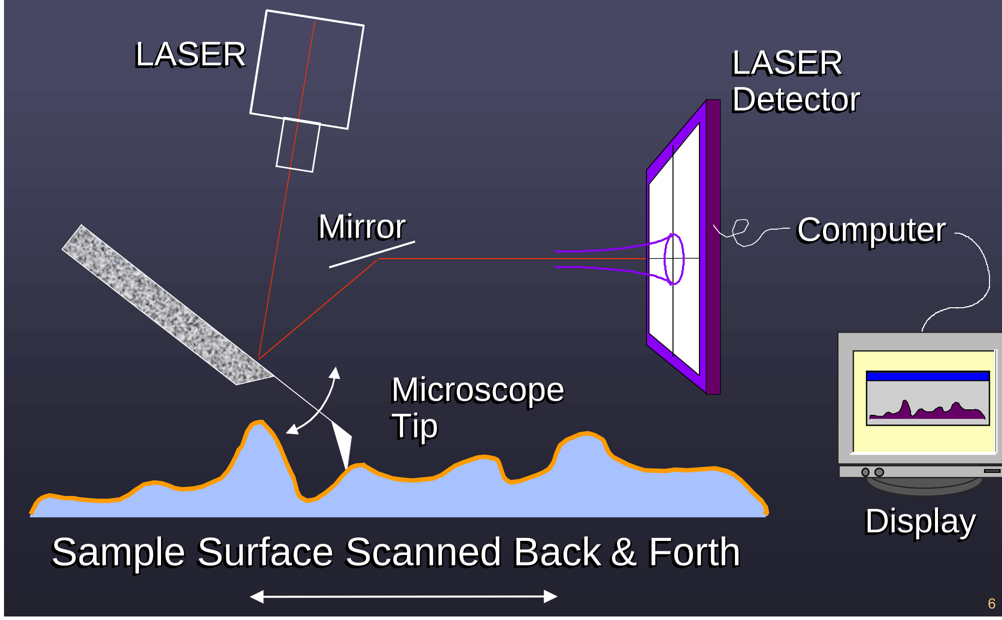

Distributed Virtual Environments
The nanoManipulator as a distributed system

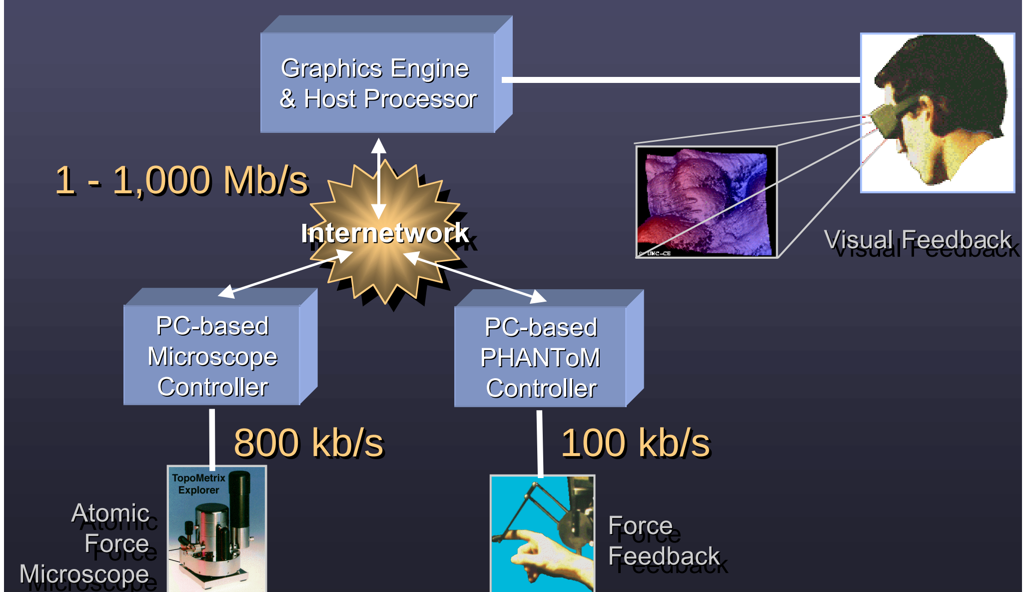

Distributed Virtual Environments
Haptic displays: Feeling a surface

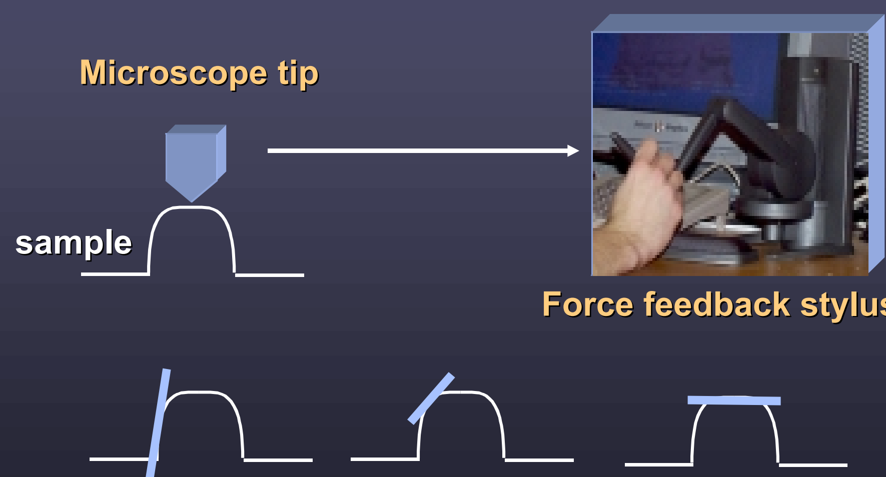

User feels a series of approximating planes
(At a 20 Hz update rate, surface appears smooth)

Supporting DVEs
The impact of Internet pathologies on nM flows

- Packet delay, delay-jitter, and loss lead to “pops” and “gaps” in audio playout
- For haptics display, Internet pathologies lead to incorrect surface shape

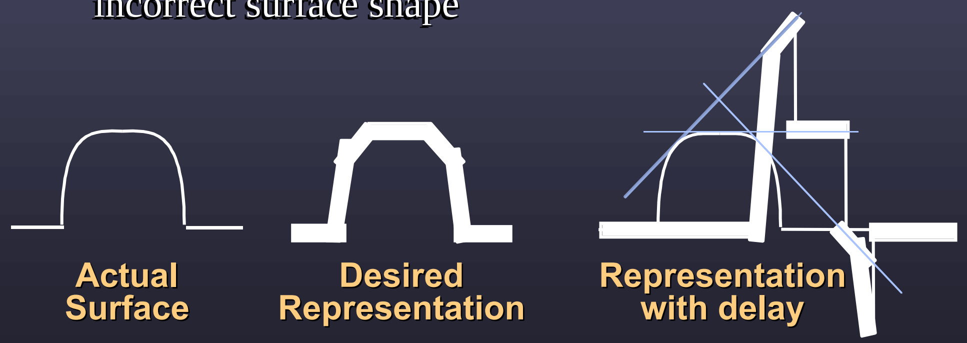

Supporting DVEs
The impact of Internet pathologies on nM flows

- Essential problems: Manage delay, delay-jitter, and loss
  - The abstract requirements are the same as for interactive audio/video applications, just the constants change
- Approach:
  - Investigate the use end-system audio/video media adaptations for ameliorating the effects of delay-jitter and loss
  - Investigate the use of novel router active queue management mechanisms for realizing better-than-best-effort services (differentiated services) for DVE flows

Supporting DVEs
End-system media adaptation

- To ameliorate the effects of loss we use a simple FEC scheme [MMCN 01]
- To ameliorate the effects of delay-jitter we use an elastic display queue management scheme called *queue monitoring* (QM) [MMSJ ‘95]

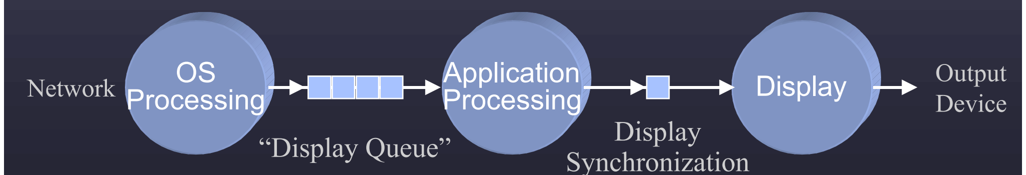

Media receiver’s processing pipeline

Delay-Jitter Adaptation
Principles of display queue management

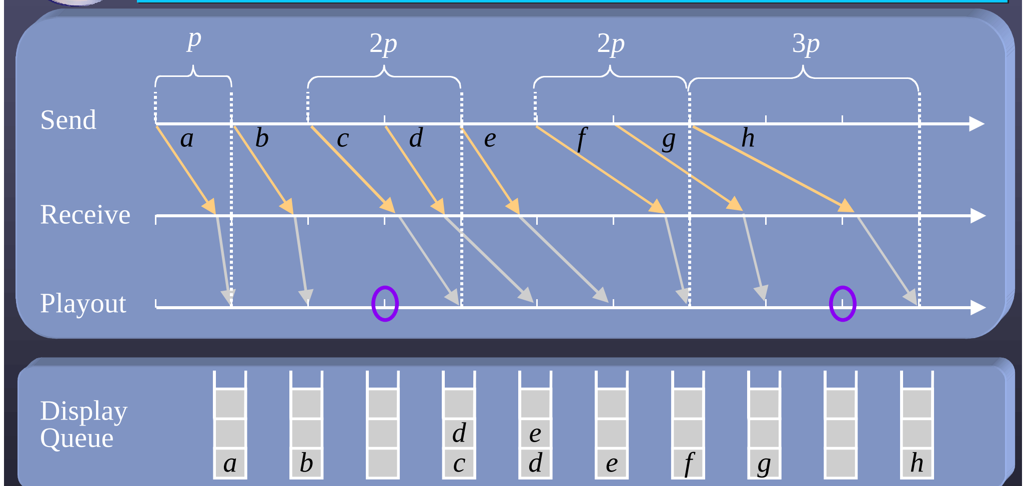

- Enqueue all arriving samples — Variable playout delay but minimal gap-rate
- Discard late samples — Constant playout delay but high gap-rate

Delay-Jitter Adaptation
Principles of display queue management

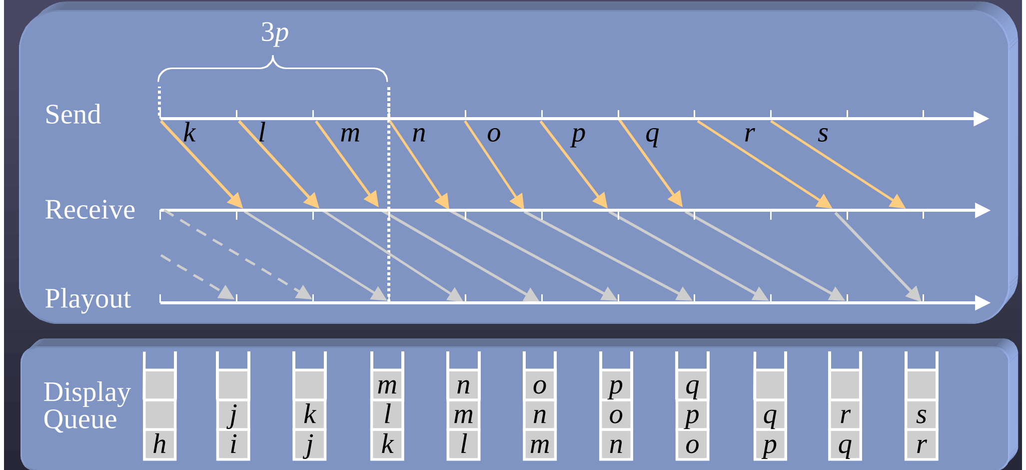

- If display queue length grows, network delay is decreasing
- If queue length shrinks, network delay is increasing
- If queue remains constant, network delay is stable

Delay-Jitter Adaptation
Queue monitoring

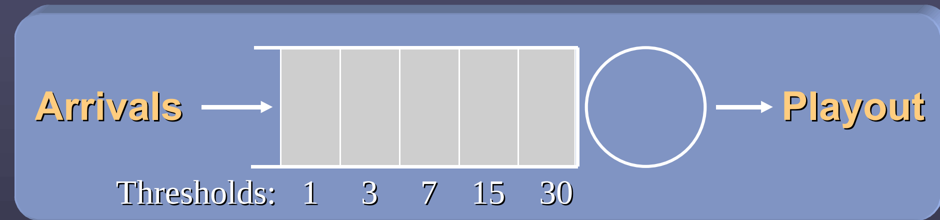

- Define a *threshold value* and a *decay rate* dropping queue elements in an effort to reduce latency
  - If queue element *k* always contains a media sample for
    threshold x (*decay* x (*n*–*k*))
    enqueues, then it is “safe” to drop the element in location *k*
- Drop quickly from long queues, drop slowly from short queues

Delay-Jitter Adaptation
Queue monitoring performance

- In a laboratory testbed simulating a small-scale internetwork, queue monitoring is effective
- For haptics data flow, queue monitoring results in a gap-rate equal to the packet-loss rate
  - At the cost of slightly increased latency
- Remote operation of the nanoManipulator on the East coast of the USA confirms these results
- (See paper for details)

Delay-Jitter Adaptation
Queue monitoring performance

Protocol	Loss	Drop rate	Gap rate	Latency
UDP	9.7%	11.7%	21.5%	89 *ms*
QM (30, 2)	10%	0.6%	10.6%	94 *ms*
QM (150, 2)	9.7%	0.02%	9.7%	96 *ms*
QM (3600, 2)	9.5%	0.001%	9.5%	91 *ms*

- For haptics data flow, queue monitoring results in a gap-rate equal to the packet-loss rate
  - At the cost of slightly increased latency

End System Media Adaptation
What more is needed?

- Queue monitoring (and other adaptations) can work well but…
  - They don't scale well
  - They don't address issues of fair bandwidth allocation between TCP and non-TCP flows
- We are also investigating the use of router-based mechanisms for realizing better-than-best-effort services
  - (Within the context of the *differentiated services* architecture for the Internet)

Towards QoS Networking
The *differentiated services* architecture

- ISPs allocate and sell capacity for a "premium" service

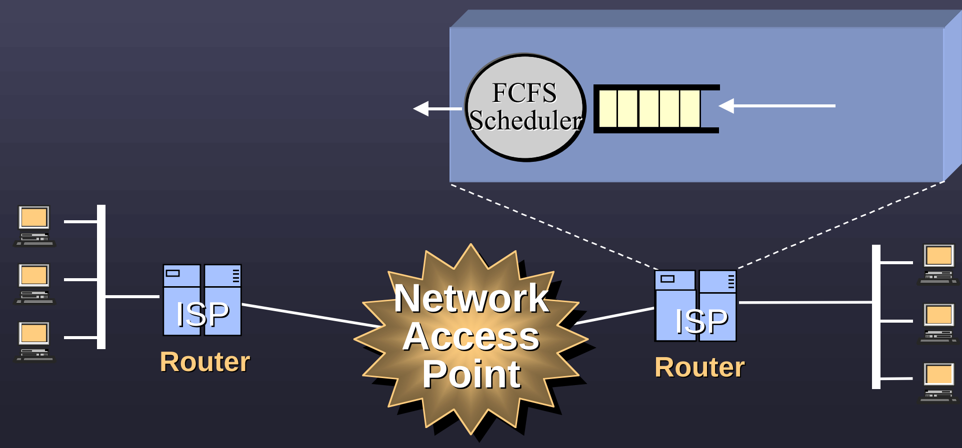

Towards QoS Networking
The *differentiated services* architecture

- ISPs allocate and sell capacity for a "premium" service
- Packets are marked according to "service profiles"

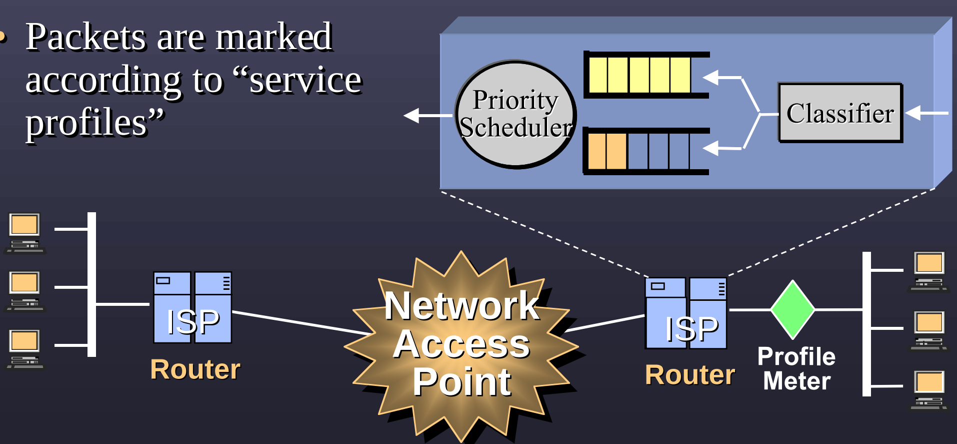

Realizing Differentiated Services
Active queue management

- This is significant utility in realizing differential services with a single router queue
  - In this model, a key technology for realizing differential services is a packet dropping policy

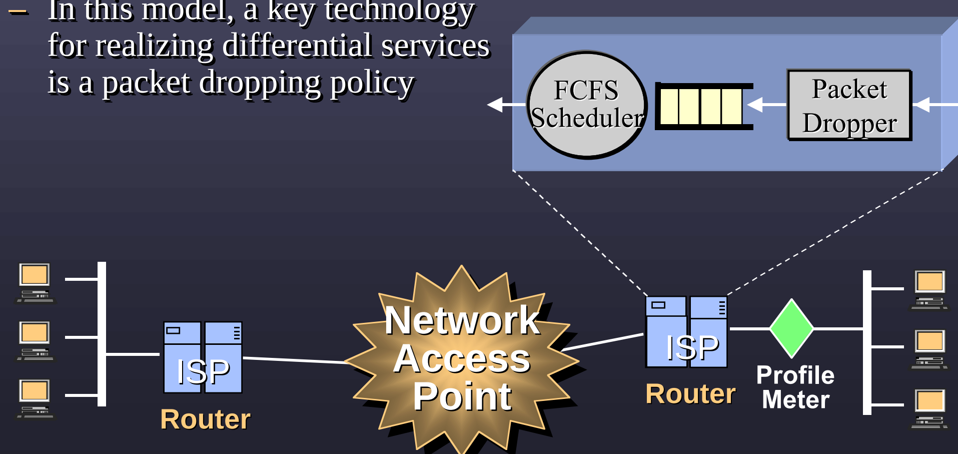

Realizing Differentiated Services
RED active queue management

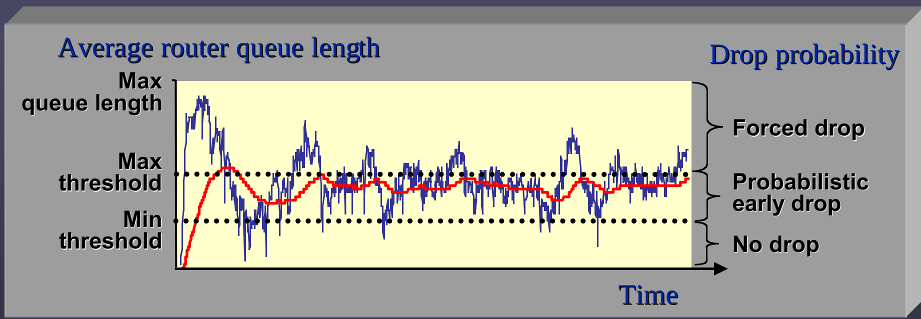

- Basic mechanism for realizing differentiated services is the *random early detection* (RED) congestion avoidance mechanism

Realizing Differentiated Services
RED active queue management

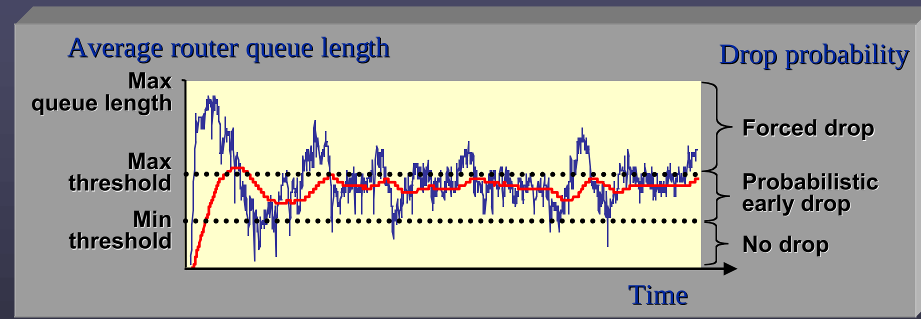

- Random drops avoid lock-out/synchronization effects
  - All flows see the same loss rate
- Early drops avoid full queues
  - Increases effective network utilization (“goodput”)
  - Decreases end-to-end latency by decreasing queuing delay

Realizing QoS Through AQM
“Class-based thresholds”

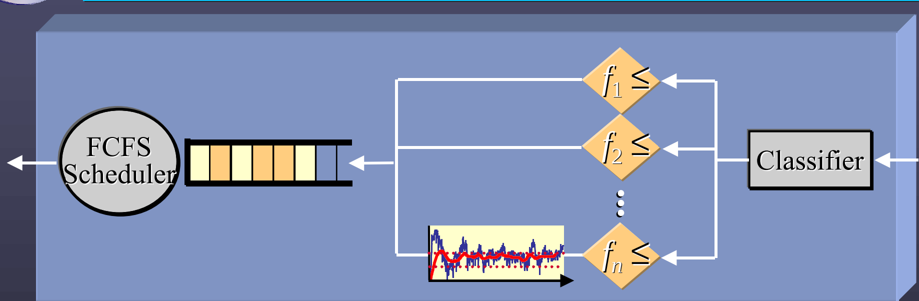

- Designate a set of traffic classes and allocate a fraction of a router’s buffer capacity to each class
- Once a class is occupying its (weighted average) limit of queue elements, discard *all* arriving packets
- Within a traffic class, further active queue management may be performed

Active Queue Management
CBT performance comparison

Queue Management Scheme	Drop Rate for Marked Flows	Latency for Marked Flows	TCP Throughput
FIFO	32.4%	63.2 *ms*	200 *kbps*
RED	30.0%	26.2 *ms*	300 *kbps*
CBT	1.3%	28.4 *ms*	790 *kbps*

- Reserving buffer capacity for multimedia (*nano*) flows improves both TCP and multimedia performance
  - At the cost of small router state complexity

Beyond Audio and Video
Network support for immersive DVEs

Summary

- Distributed virtual environments represent a significant generalization of the traditional multimedia networking problem
- We're attempting to…
  - Allow co-existence of selected non-congestion responsive UDP traffic and responsive TCP traffic through differentiated services
  - Apply audio/video end-system media adaptations to ameliorate remaining congestion effects
- Conclusion: Mechanisms for supporting real-time audio and video flows work well for immersive DVEs

Beyond Audio and Video
The nanoManipulator in detail