

Optimizing Internet Data Transport

Venkata N. Padmanabhan

Microsoft Research

<http://www.research.microsoft.com/~padmanab>

BU/NSF IMIC Workshop

August 30, 1999

Outline

Two issues

- **Concurrent data streams**
 - discovering presence of shared bottleneck
 - coordinating distributed senders
- **All packets are not equal**
 - shielding vulnerable packets


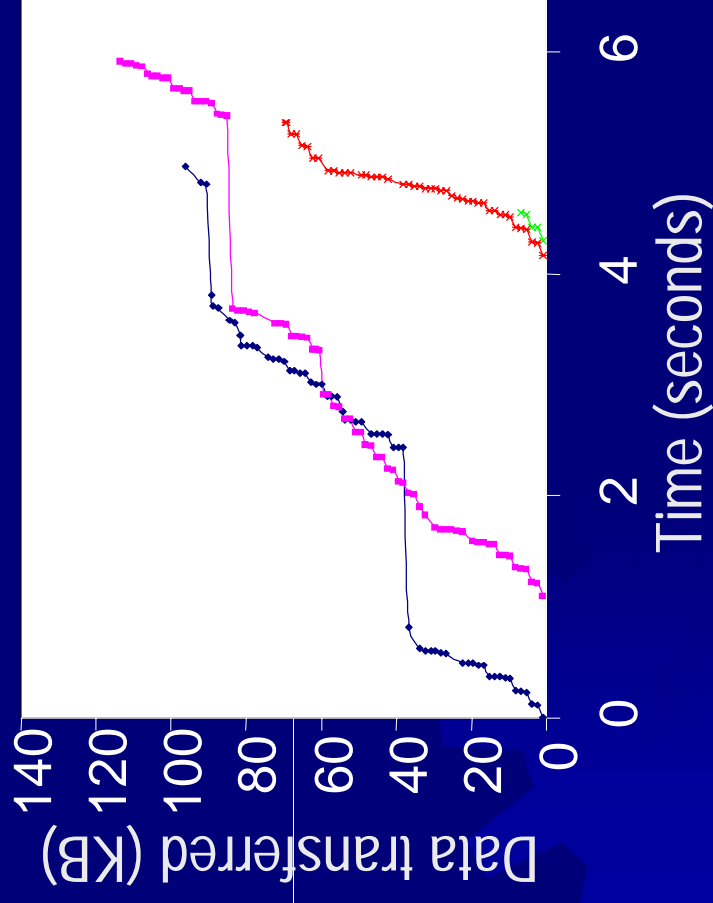
Concurrent Data Streams

- ★ Coordination better than competition
 - shared learning of network conditions
 - reflect user utility rather than network dynamics
 - e.g.: optimal progressive Web page delivery [GB99]

★ TCP Session [Pad98]

- congestion control
- loss recovery
- bandwidth sharing

TCP Session Performance

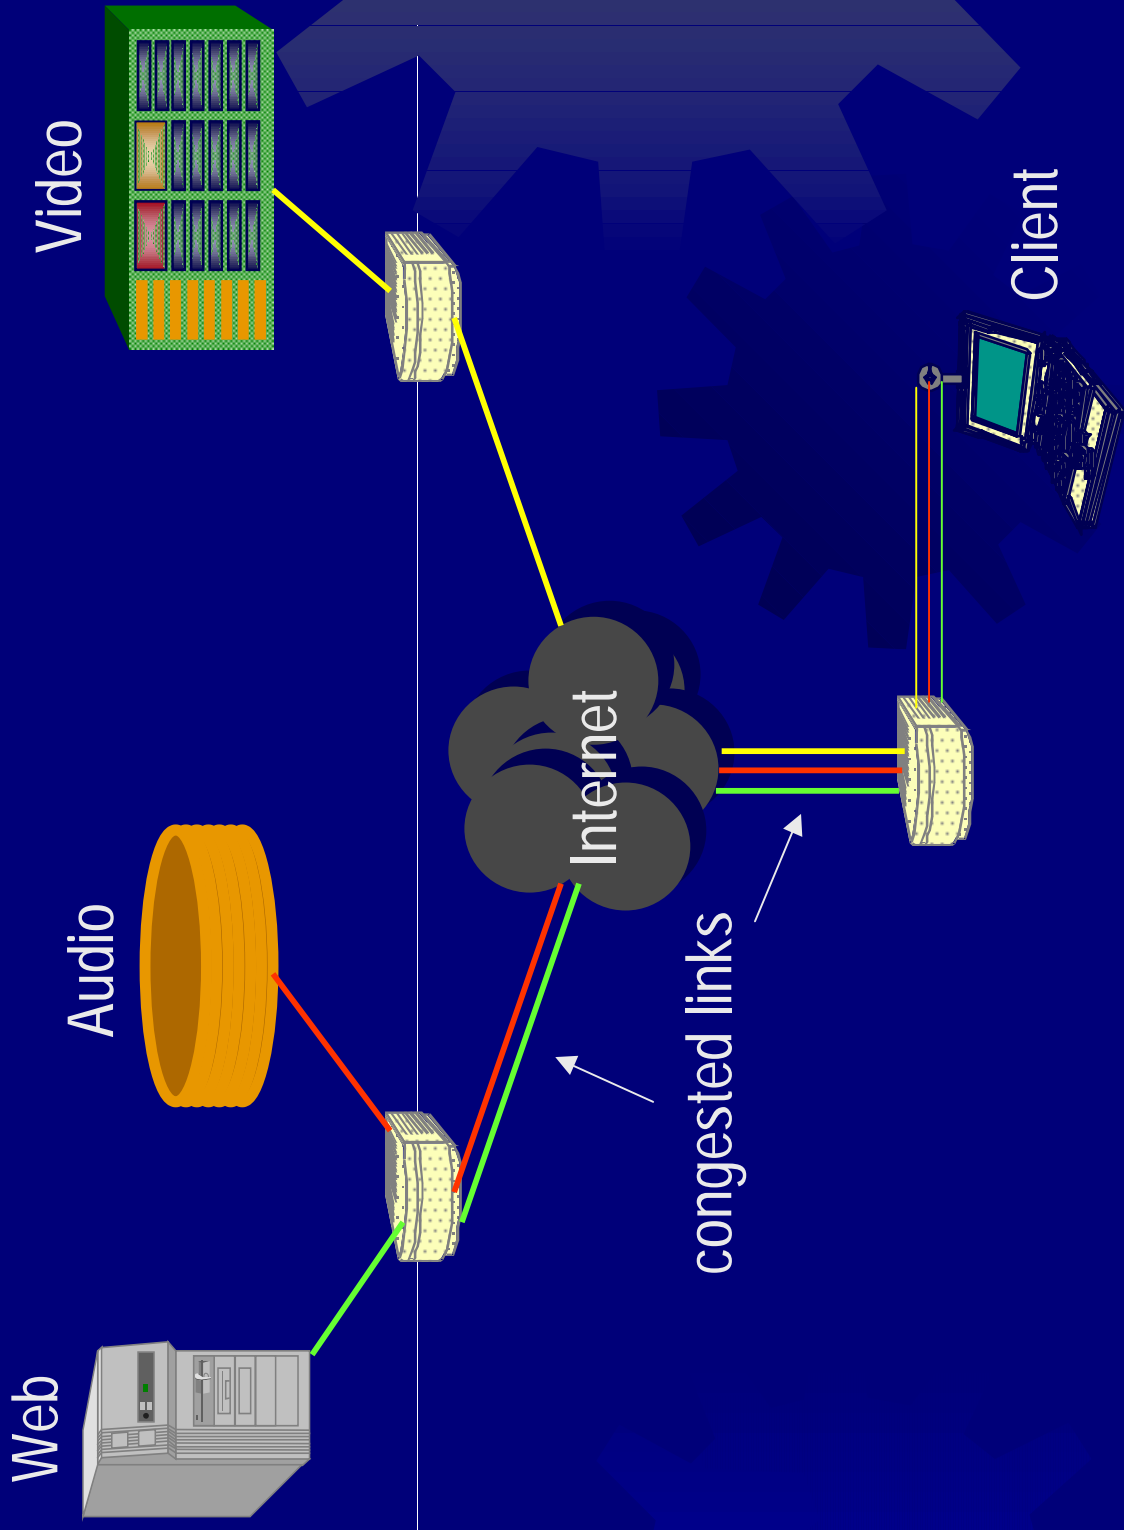


Coordination \Rightarrow more predictable performance

Concurrent Heterogeneous Streams

- ★ Heterogeneous often implies distributed
 - ★ specialized servers
 - ★ new usage scenarios
 - e.g., listening Internet radio while surfing
- ★ Looking beyond TCP \Rightarrow looking beyond host-pair coordination


Concurrent Heterogeneous Streams



Challenges & Potential Solutions

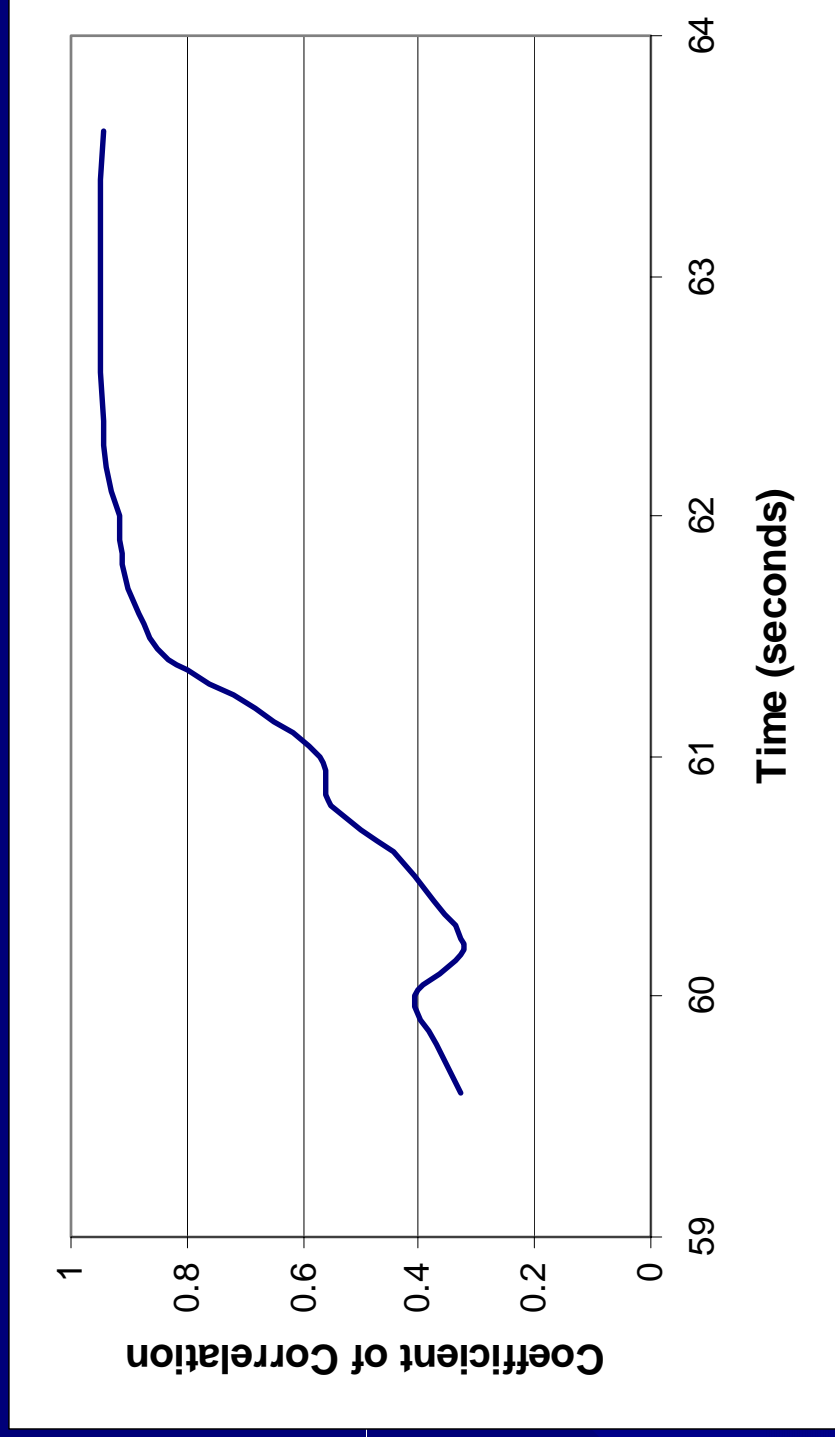
- ★ Discovering presence of shared bottleneck
 - ★ indirect: correlation of delay/loss patterns
 - ★ [Bolot93], MINC [CDH+99]
 - ★ do not need synchronized clocks
 - ★ direct: enhanced ECN
- ★ Coordinating distributed senders
 - ★ explicit: receiver-driven flow control
 - ★ implicit: congestion feedback filtering

Congested Intranet Link




Wide-area inverse-muxed T1 links (3 Mbps)

Congested Intranet Link




Significant correlation in queuing delay

Single Congested Link




Connections 1 & 2 share a congested link

Single Congested Link




Shared bottleneck \Rightarrow high delay correlation

Multiple Congested Links



Connection 1 traverses two congested links
only one of which is shared with connection 2

Multiple Congested Links





Presence of multiple congested links reduces effectiveness of delay correlation technique

Enhanced ECN

- ★ ECN plus unique router tag
 - ★ tag helps discover shared bottleneck
- ★ Non-unique tags can be used for efficiency
 - ★ periodic re-hashing to avoid persistent collisions
- ★ Works better with multiple congested links
- ★ Could complement delay correlation technique


Coordinating Distributed Senders



All Packets Are Not Equal


- Certain packets more important than others
 - TCP
 - SYN packet
 - packets sent when window is small
 - retransmission
 - RM repair request
- Vulnerable to loss of important packets

Retransmission Timeouts



Very little data typically sent between timeouts

Shielding Vulnerable Packets



- Diffserv coupled with protocol-specific knowledge
- High drop priority for vulnerable packets
- FIFO scheduling avoids reordering

Summary

- ★ Need to coordinate concurrent data streams
 - ★ E2E/router-assisted detection of shared bottleneck
 - ★ congestion feedback filtering
- ★ All packets are not equal
 - ★ diffserv to shield vulnerable packets

<http://www.research.microsoft.com/~padmanab>