# On the impact of variability on buffer dynamics in IP networks*

* Y. Joo, V. Ribeiro, A. Feldmann, A. Gilbert, and W. Willinger in Proc. of the 3th Annual Allerton Conference on Communication, Control, and Computing, 1999

Extended version: TCP/IP traffic dynamics and network performance: A lesson in workload modeling , flow control, and trace-driven simulations, ACM CCR 2001

# Objectives

- Main:
  - Highlight by using a toy example the extent to which assumptions underlying the nature of network traffic can influence practical engineering decisions

- Secondary:
  - Introduce some concepts and tools that one usually comes up to when analyzing network traffic

# In more detail

- Toy example: a simple TCP/IP network
- Assumptions underlying the nature of network traffic: wrt. variability of network traffic
- Buffer dynamics of IP routers:
  - A substantial amount of the space, power and cost of a high-end router line card today is used by the buffer memory that stores packets during times of congestion.*
  - If these buffers fill up, they cause queueing delay and delay-variance; when they overflow they cause packet loss, and when they underflow they can degrade throughput.*

* G. Appenzeller, N. McKeown, J. Sommers, and P. Barford, Recent Results on Sizing Router Buffers, in Proc. of the NSDC 2004

# Mechanisms that create variability*

- User behavior (large scales)
  - Application level variability, e.g., Web
- Network characteristics
  - Different delays, round-trip time, cross traffic
- Feedback control (small scales)
  - Reliability and adaptivity

* More information: A. Feldmann, A. Gilbert, P. Huang, and W. Willinger, Dynamics of IP traffic: a study of the role of variability and the impact of control, in Proc. of the SIGCOMM 1999

# Mechanisms that create variability

- Scaling
  - How does traffic behave at different agg. levels
- Large time scales: User dynamics
  - Users act mostly independent of each other
  - Users are unpredictable:
    - Variability in doc size, # of docs, time between docs
- Small time scales: Network dynamics
  - Network protocols effects: TCP flow control
  - Queue at network elements: delay

# Approach

- Study traces from simulations (e.g., ns-2)
  - Complete control over all aspects of network
    - workload model, network model, protocol
- Real network traces used as benchmark
  - Simulation setup and trace analysis
- Simulation traces used for trace-driven simulation
  - Study queuing dynamics
- Correlation of analysis results to simulation setup
  - Identify network features that cause divergence

# Outline

- Simulation setup
  - Workload model
  - Network model
  - Protocol
- Impact of variability at the application layer
  - Mice vs. elephants [V. Jacobson]
- Impact of feedback
  - Open loop vs. closed loop

# Workloads

- No variability:
  - Infinite sources
    - 50 clients requesting big files
- High variability:
  - Web sources
    - 350 clients down loading Web pages
- Simulation:
  - Client startup: random time 0-600 seconds
  - Duration: 4200 seconds
  - Analysis: 900-4200 seconds

# High variability workload: Web

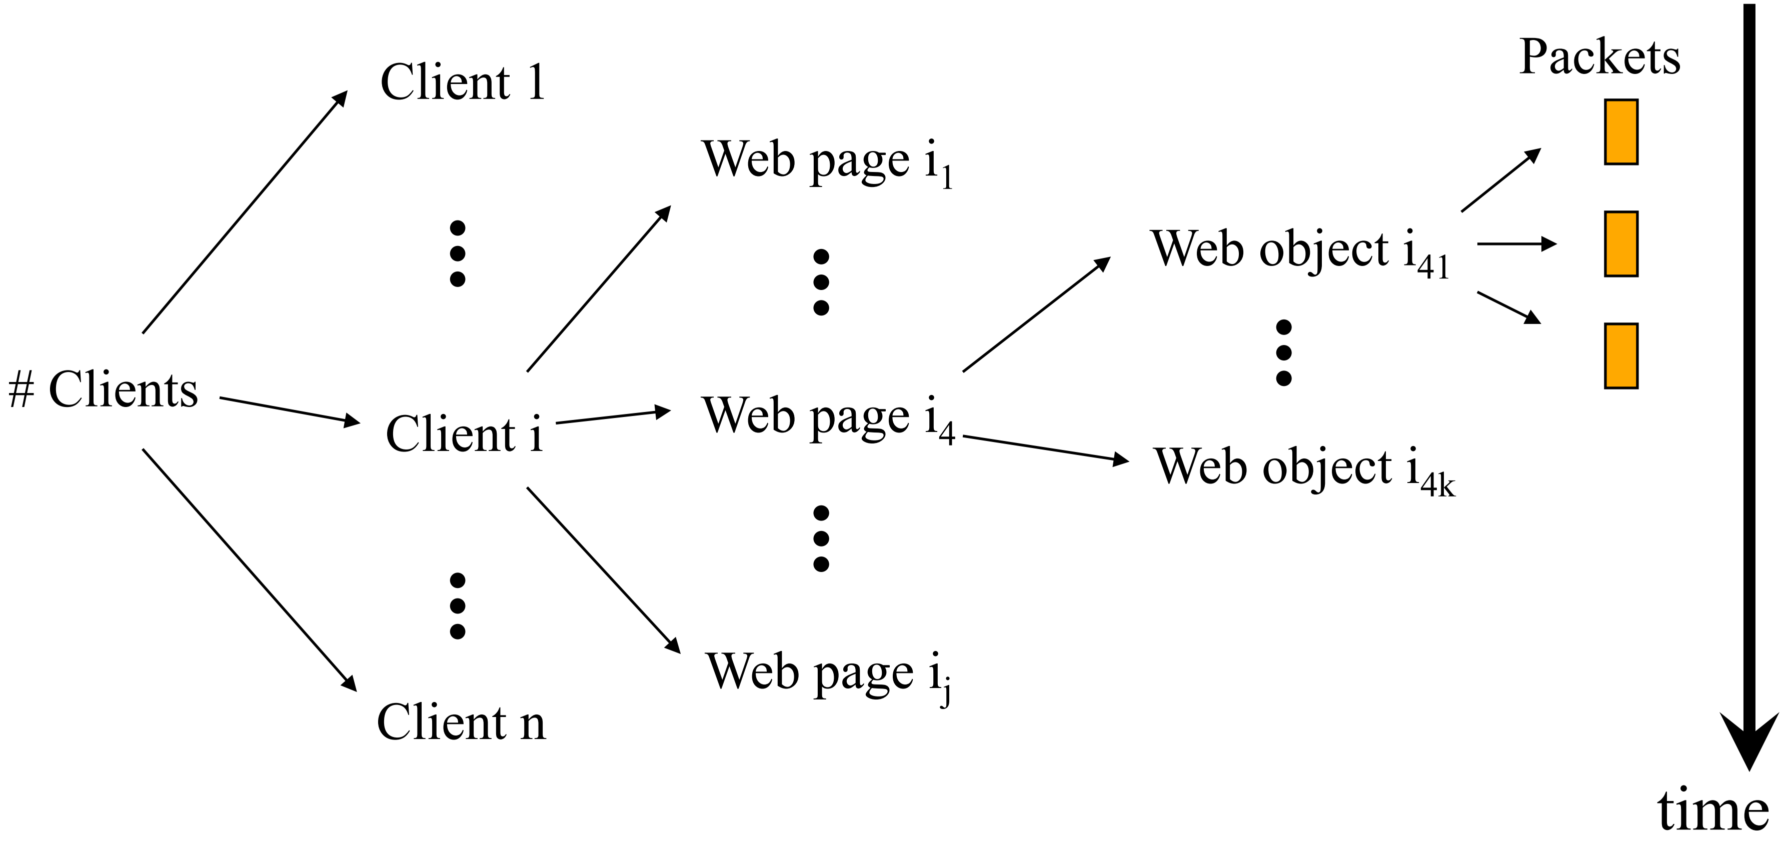


- Parameters (similar to SURGE [BC98]):
  - Number of clients, pages, objects, packets per object
  - Time between Web pages, Web objects

# A simple network topology

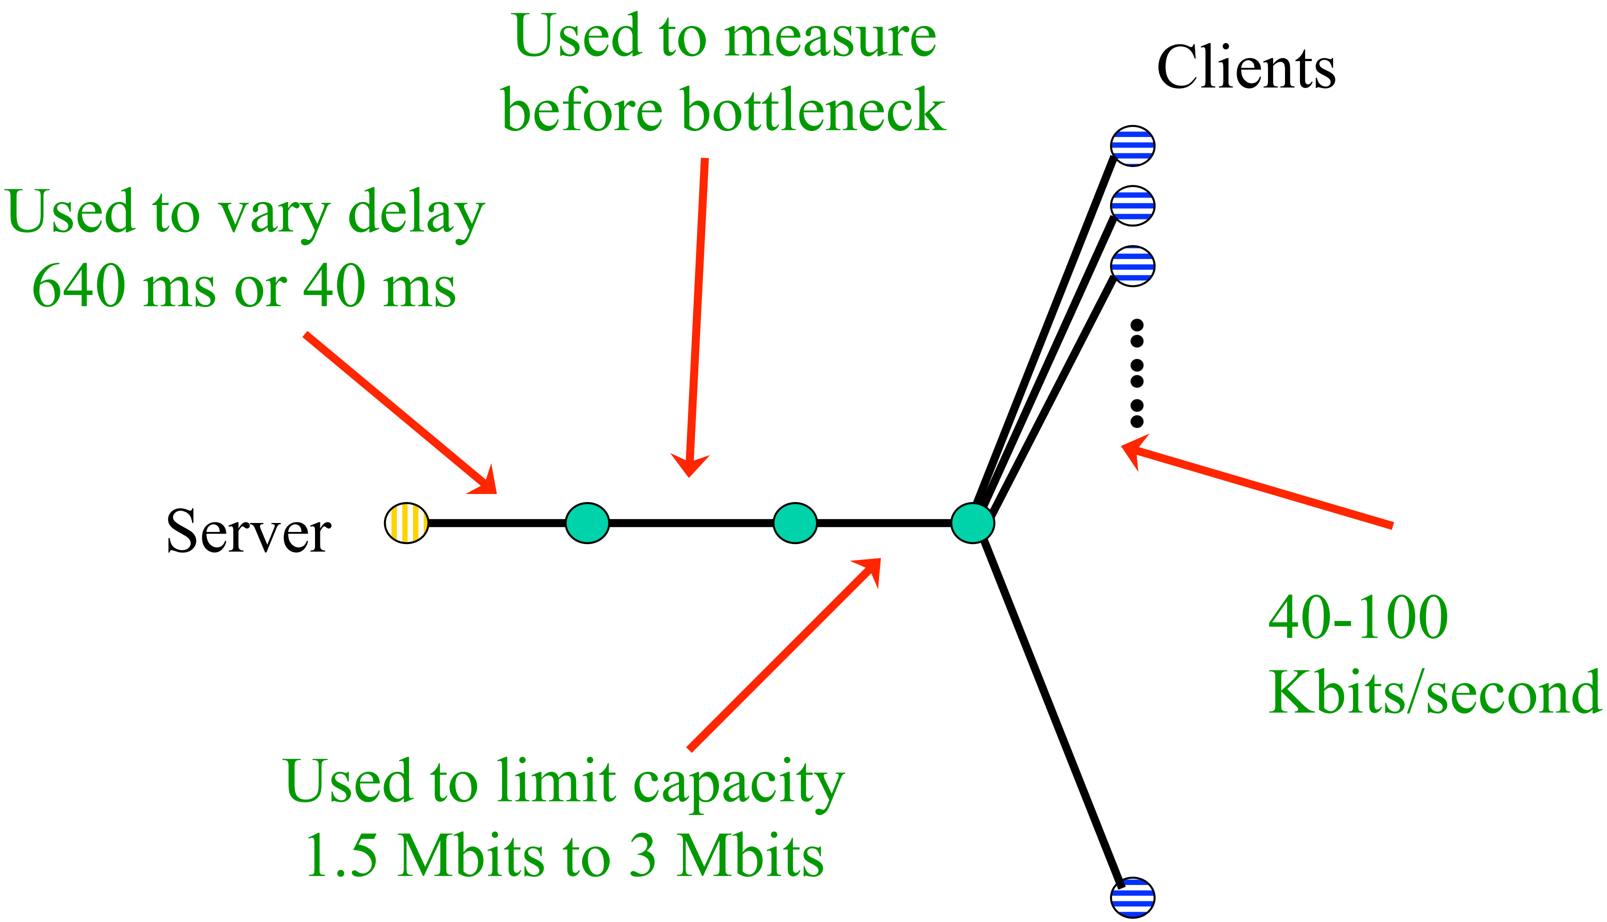

# TCP: Reliable byte stream

- Data divided into segments
- Segments are ACKED by receiver (cumulative)
- Timer for every segment
- Segments retransmitted
  - Timer goes off
  - Four duplicate ACKs received
- Flow control
  - Sliding window protocol avoids losses at the receiver
  - Bandwidth limits impose congestion window (cwnd)
    - Slow start increases cwnd exponentially
    - Congestion avoidance increases cwnd linearly
    - Packet losses triggers cwnd changes

# Elephants: Infinite sources

- Packet rate process and buffer occupancy process
  - Network round trip time 1.3 seconds

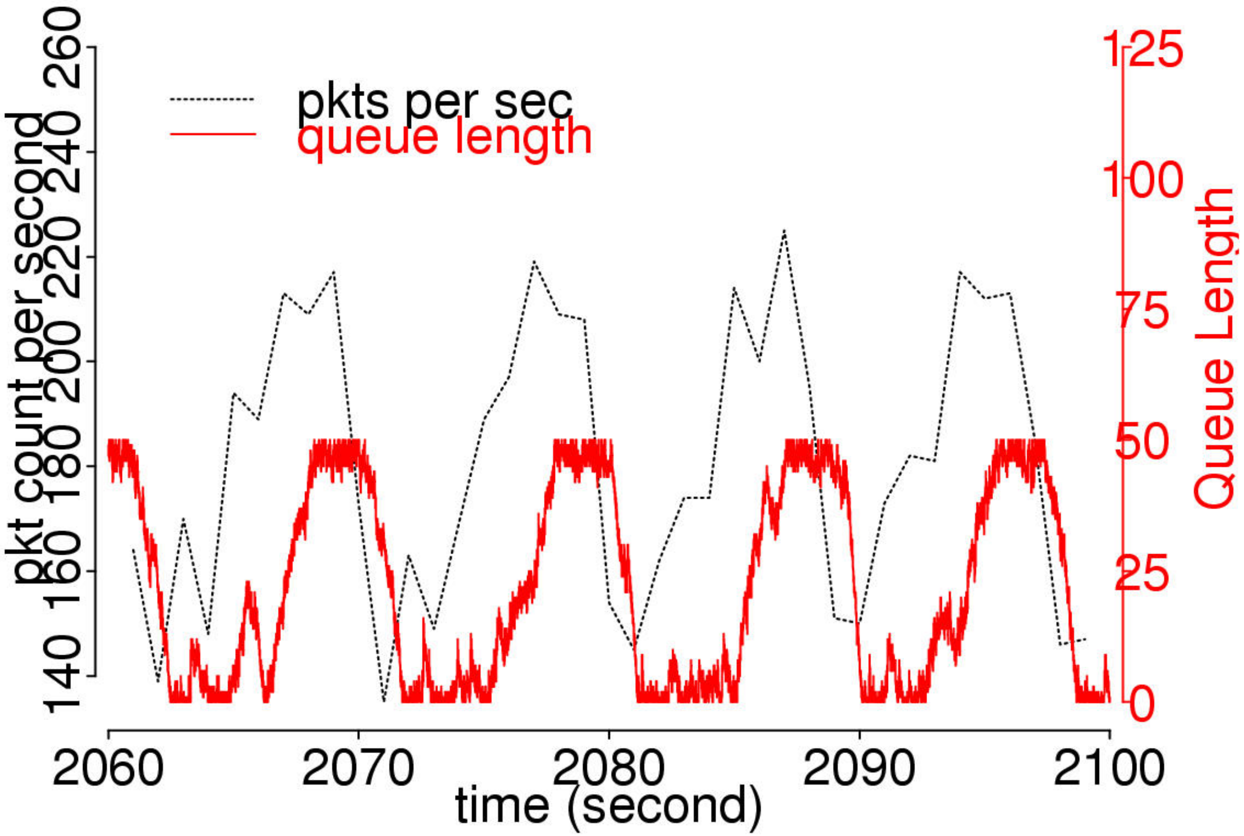

# Elephants (cont.)

- Packet rate process and buffer occupancy process
  - Network round trip time 0.14 seconds

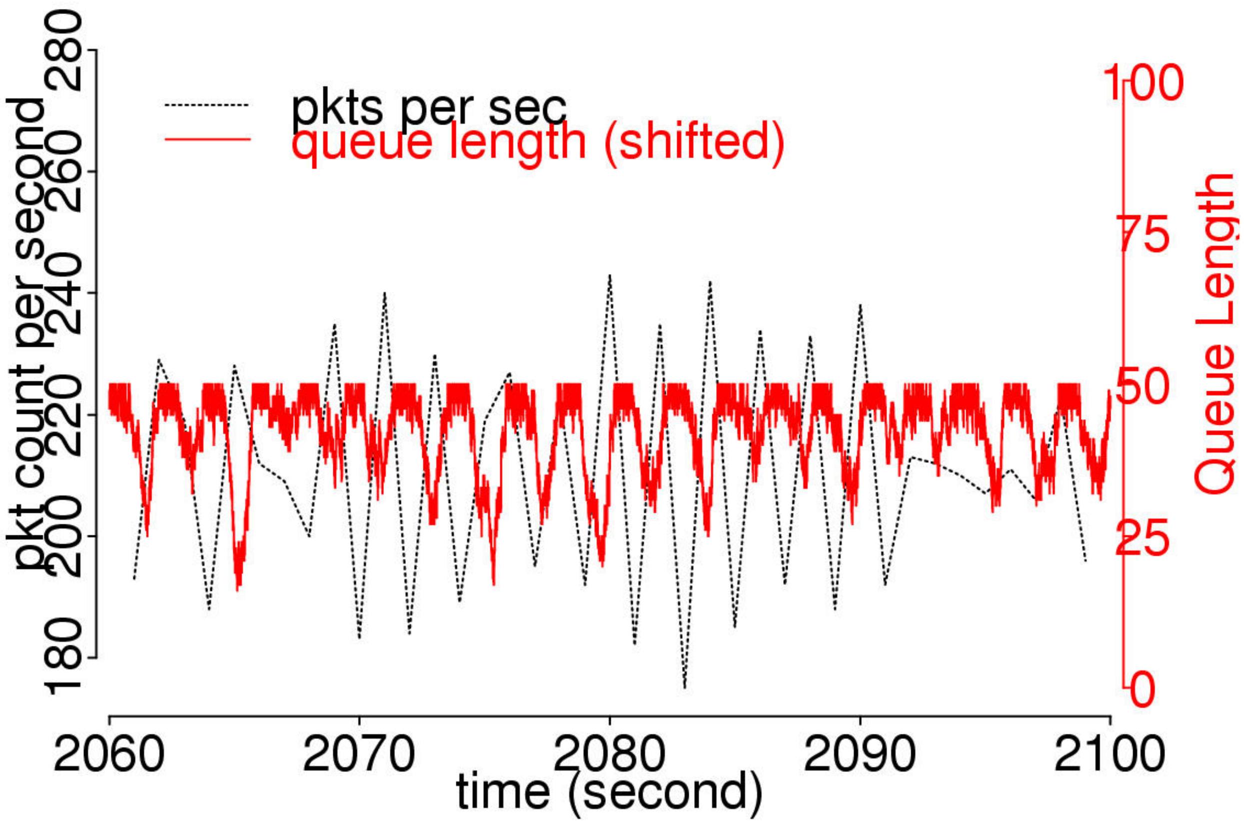

# Mice and elephants: Web sources

- Significant portion of connections are short

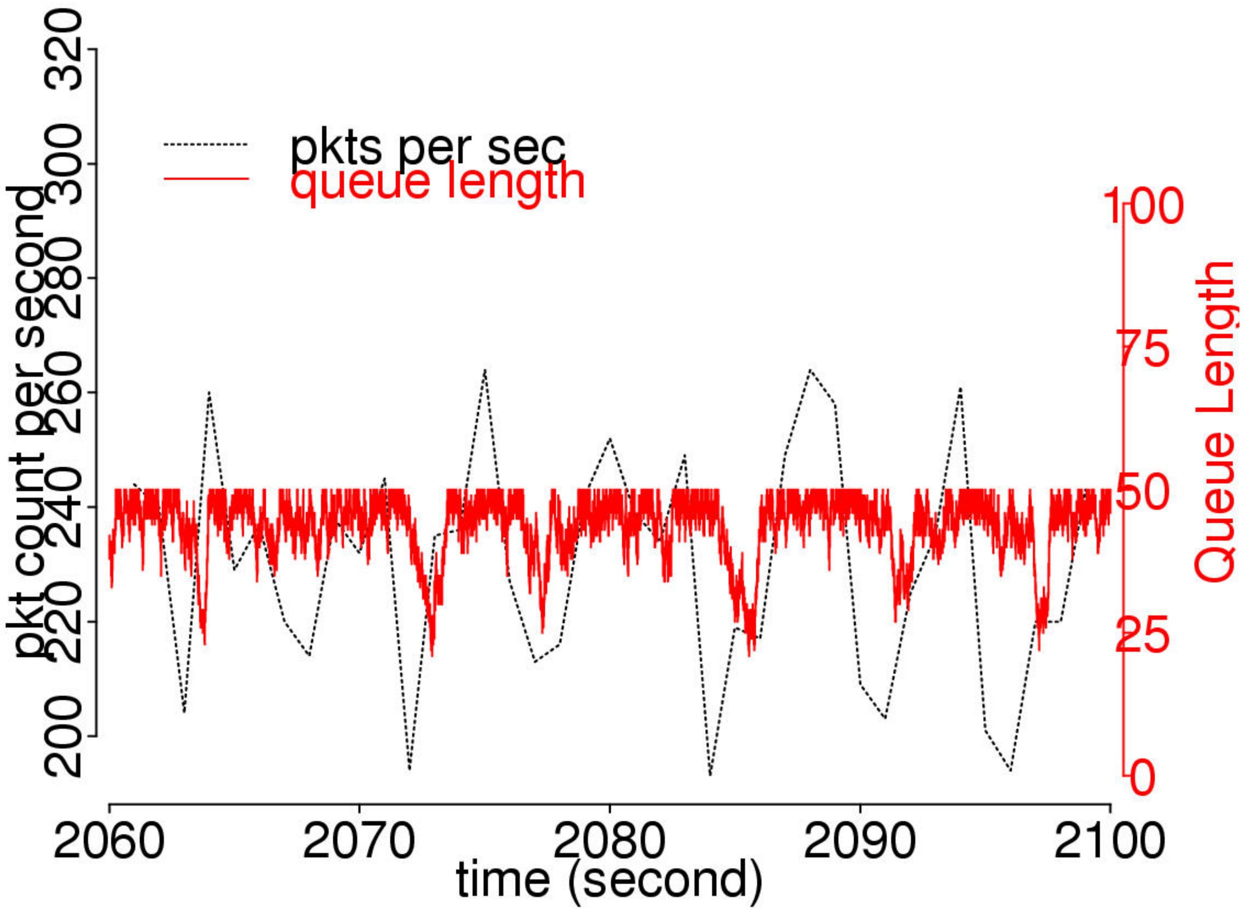

# Elephants vs. Mice

- No variability in workload and network
  - synchronization of packet rate process
  - synchronization of buffer occupancy process
- Mice stop synchronization
  - no apparent synchronization
  - higher packet arrival process
  - higher utilization

# Effect of synchronization

- Percentage of connections with losses during $\Delta T$

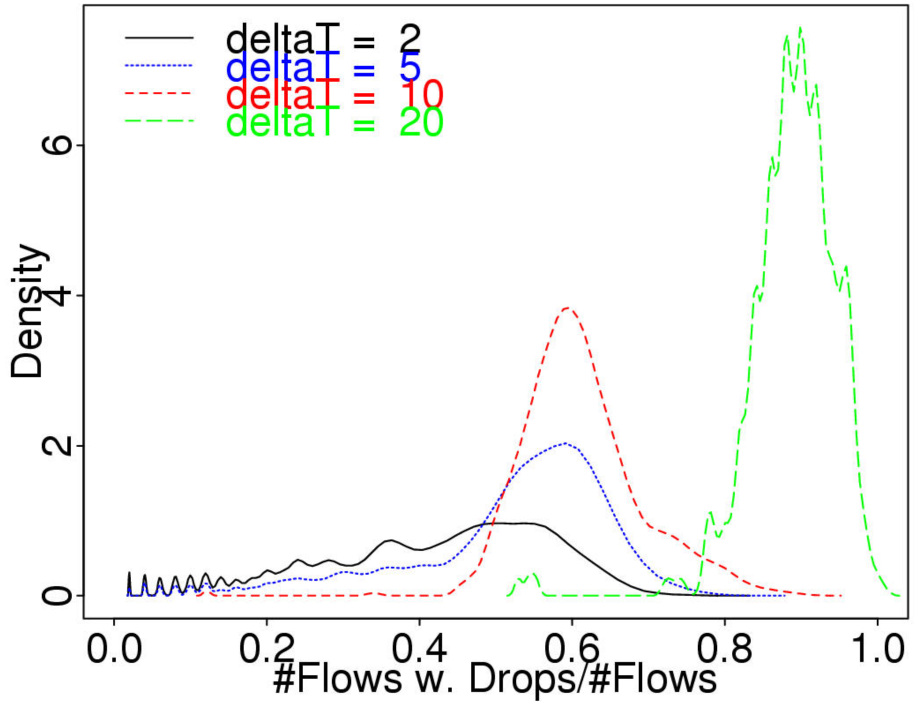


Infinite sources

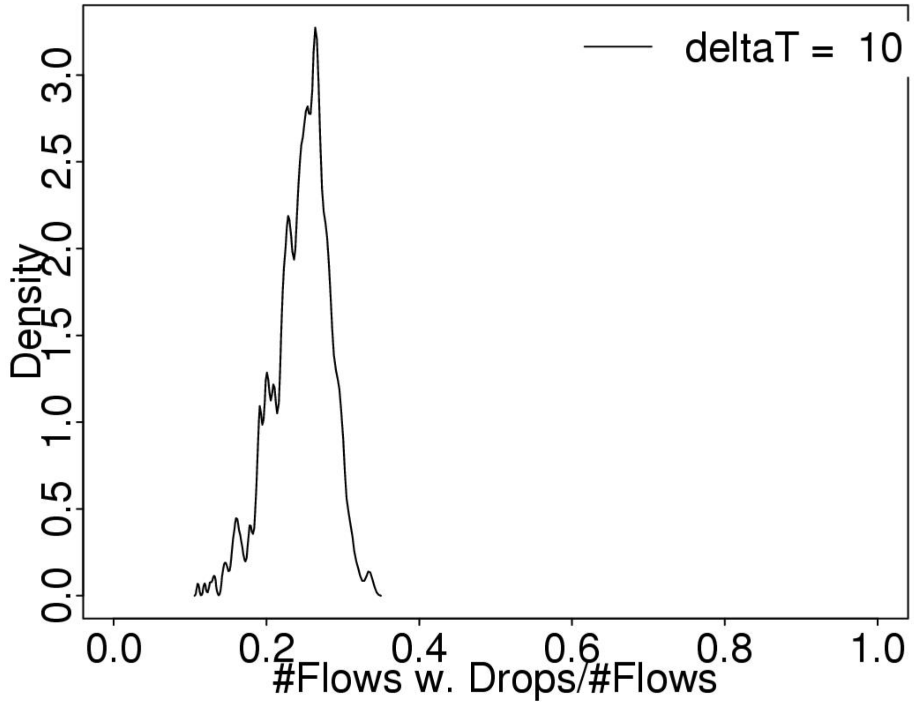


Web sources

# Effect of synchronization (cont.)

- Fractions of connections with losses in slow-start

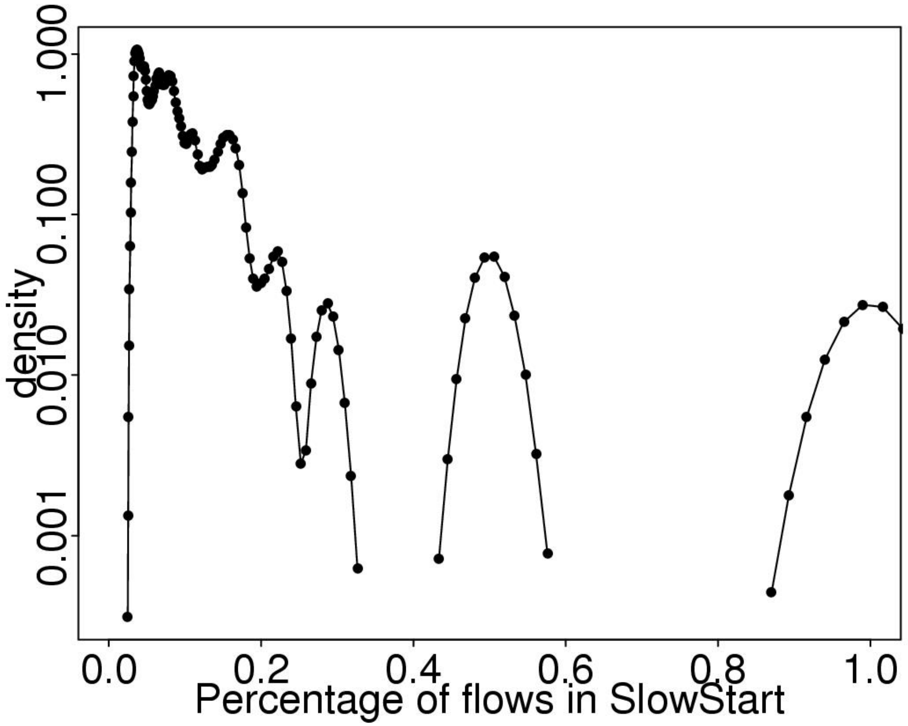


Infinite sources

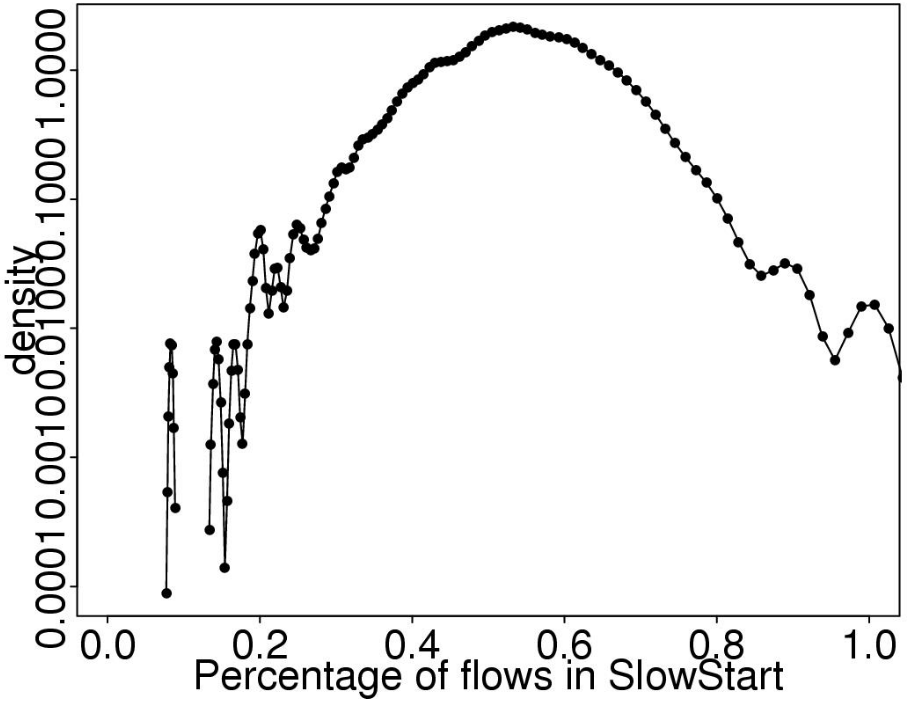


Web sources

# Why mice eliminate synchronization

- Mice
  - Too short for feedback
  - TCP states non-synchronized
  - Arrival highly bursty
  - Large fraction in slow-start
- Elephant
  - Within two cycles losses for almost all connections
  - TCP states synchronized
  - Small percentage in slow start
- Consequence
  - # of burst losses larger for Web than Infinite srcs

# Open loop vs. close loop

- Queuing system with finite buffer
  - Change in link capacity can be modeled by change of service time distribution
- A simple experiment
  - ns-2 simulation with approx. finite buffer space (50 packets)
  - Trace driven queuing analysis of G/D/1 queue with varying service times
- Evaluation
  - Set of ns-2 simulations with different bottleneck speeds
    - Calculate packet loss
  - Compare packet losses from simulations to packet loss predicted from open loop queue system

# Open loop vs. close loop (cont.)

Web sources

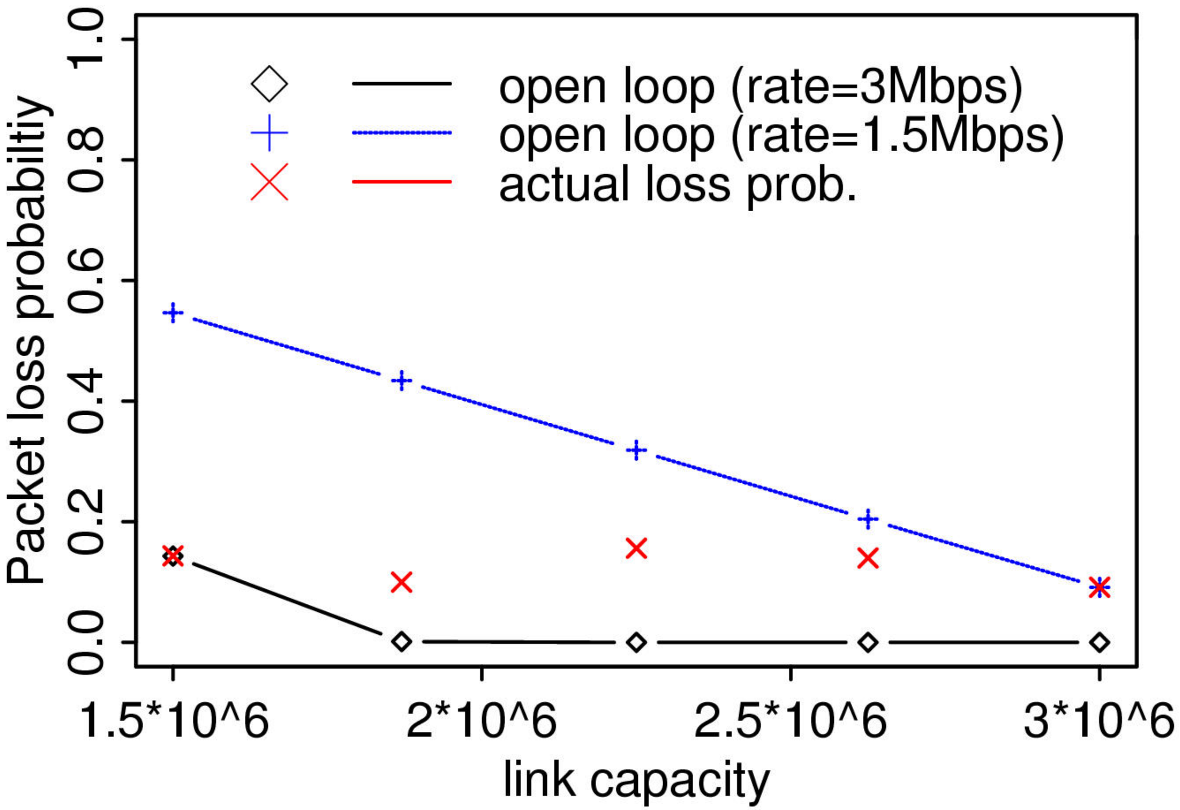


- Either extremely conservative or overly aggressive

# Open loop assumptions

- Queuing system with infinite buffer
  - Buffer occupancy probability P[Q > x] can approximate finite buffer packet loss
- A simple experiment
  - ns-2 simulation (approx. inf. buffer space 1000 pkts)
  - Trace analysis to calculate buffer occupancy
- Evaluation
  - Set of ns-2 simulations with different finite buffers
  - Compare packet losses to buffer occupancy
- Result:
  - P[Q > x] extremely conservative

# Conclusion

- Infinite source models and queue analysis provide necessary simplifications for
  - Analysis
  - Simulations
- Challenge
  - Address variability at
    - User level
    - Network level
  - Judge the impact of feedback