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

Anja Feldmann[1]

Youngmi Joo[2]
Vinay Ribeiro[3]
Anna C. Gilbert[1]
Walter Willinger[1]

[2] Stanford University, [3] Rice University
[1] AT&T Labs-Research,

Mechanism that create variability

User behavior
- Application level variability, e.g., Web

Network characteristics
- Different delays, round-trip time, cross traffic

Feedback control
- Reliability and adaptivity

Our approach

Study traces from simulations

- 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 dynamic

Correlation of analysis results to simulation setup

- Identify network features that cause divergence

Outline

Simulation setup

- Workload model
- Network
- Protocol

Impact of variability at the application layer

- Mice vs. elephants [V. Jacoson]

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

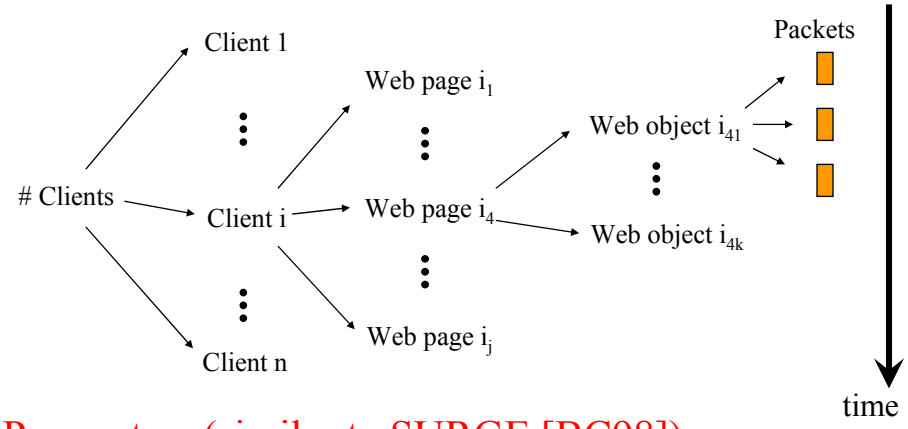

Parameters (similar to SURGE [BC98]):

- Number of clients, pages, objects, packets per object
- Time between Web pages, Web objects

A simple network topology

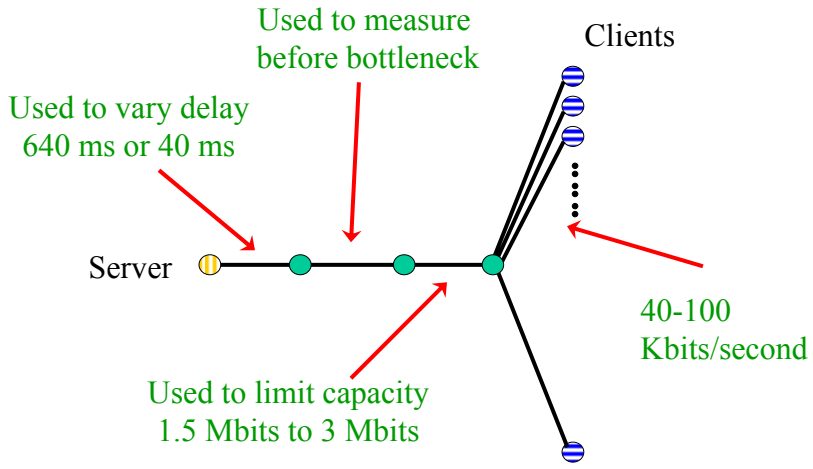

TCP provides a reliable byte stream

Data segmented into segments

Segments are acknowledged by receiver (cumulative)

Timer for every segment

Segments retransmitted

- Timer goes off
- Four duplicate Acks received

Flow control

- Sliding window protocol avoids losses at receiver
- Bandwidth limits impose congestion window
  - Slow start increases cwnd exponentially
  - Congestion avoidance increases cwnd linearly
  - Packet losses triggers congestion window changes

Outline

Simulation setup

Impact of variability at the application layer

- Mice vs. elephants

Impact of feedback

- Open loop vs. closed loop

Elephants: Infinite sources

Packet rate process and buffer occupancy process

- Network round trip time 1.3 seconds

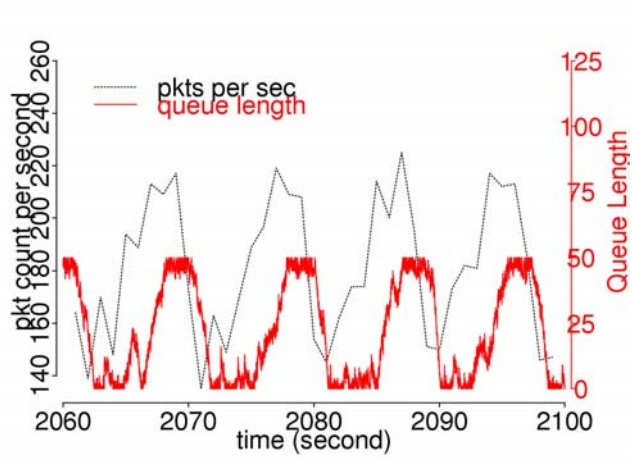

Elephants (cont.)

Packet rate process and buffer occupancy process

– Network round trip time 0.14 seconds

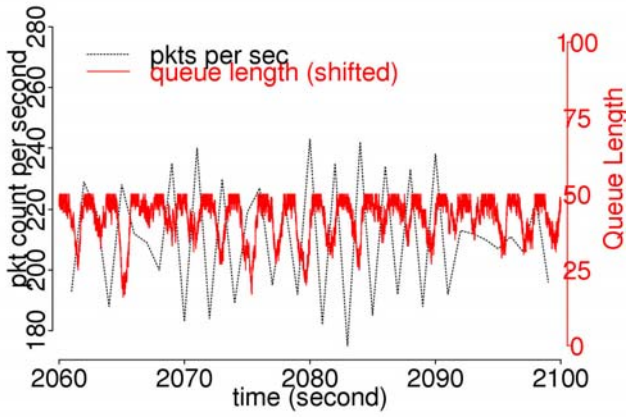

Mice and elephants: Web sources

Significant portion of connections are short

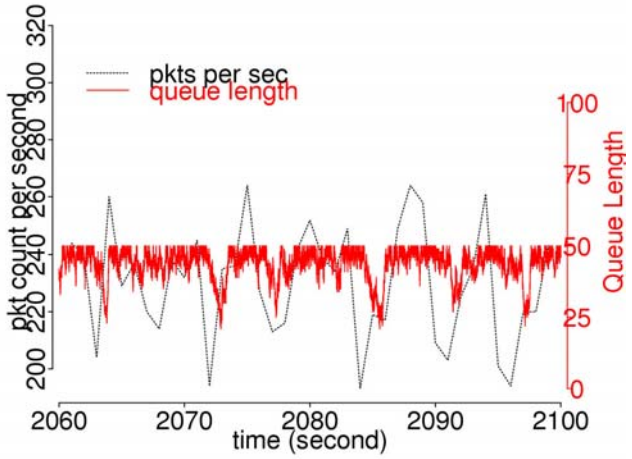

Elephants vs. mice

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 ΔT

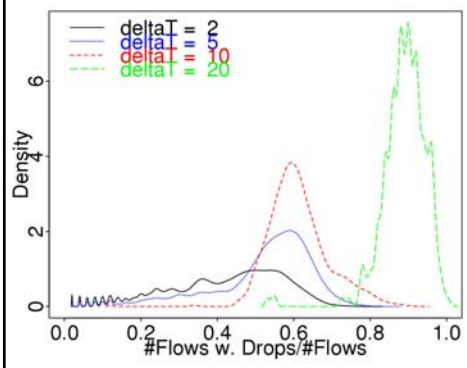

Infinite sources

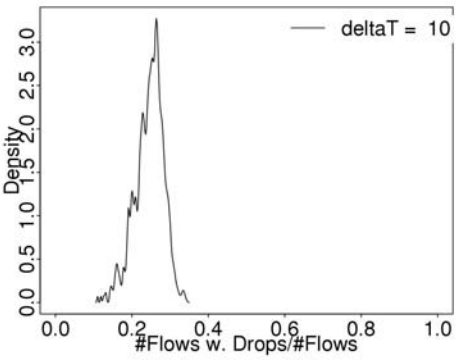

Web sources

Effect of synchronization (cont.)

Fractions of connections with losses in slow-start

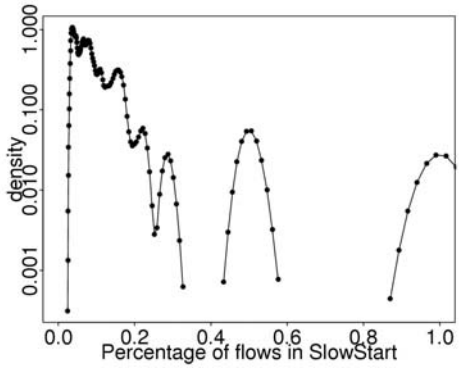

Infinite sources

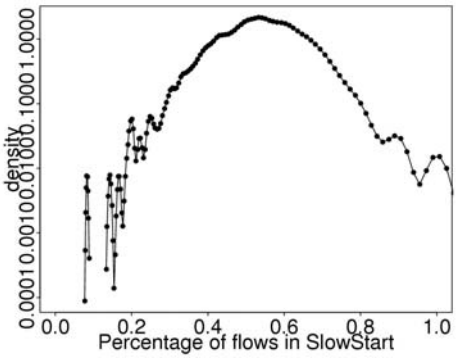

Web sources

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 sources

Outline

Simulation setup

Impact of variability at the application layer

- Mice vs. elephants

Impact of feedback

- Open loop vs. closed loop

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 simulation with approx. finite buffer space (50 pkts)
- 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

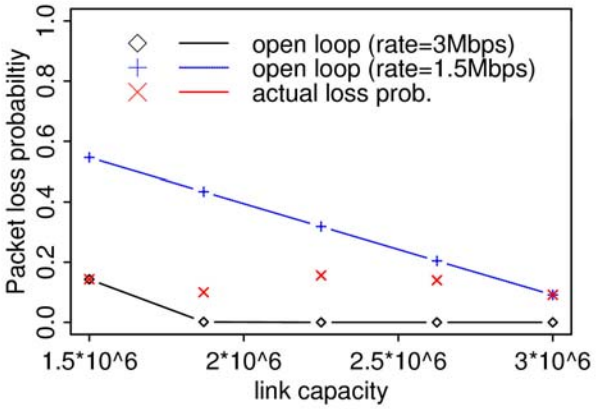

Either extremely conservative or overly aggressive

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

Simulation setup

User behavior: workload model

- Infinite sources — no variability
- Web sources — high variability

Network characteristics

- Unrealistic in its simplicity

Feedback control

- TCP adapts to congestion in the network