# Analysis of Buffer Size in Core Routers

by Arthur Dick

Supervisor
Anirban Mahanti

# Overview

- Problem Description
- Background Information
  - TCP Reno vs. NewReno
  - The Rule-of-Thumb
  - Synchronization
- Experiment Methodology
- Results

# Problem Description

- Internet routers buffer packets during times of congestion
- Buffers cause queuing delay
- The fast memory required in core routers is expensive
- Recent studies suggest routers in use today contain buffers larger than necessary

# Background Information
## TCP Reno vs. NewReno

- TCP is the protocol in use for many applications such as the web, FTP, and email
- There are a number of variants on the TCP protocol
- The studies cited considered only the Reno variant, whereas NewReno is better at recovering from multiple losses

# Background Information
## TCP Reno vs. NewReno (2)

- TCP is designed to allow multiple connections to share a congested links bandwidth

- Phases
  - Slow Start
  - Congestion Avoidance
  - Fast Retransmit / Fast Recovery

# Background Information
## The Rule-of-Thumb

- B = RTT * C
- Ensures a congested link does not go idle given a single TCP Reno flow
- Core routers today commonly carry in excess of 10 000 TCP flows

# Background Information
## Synchronization

- If the buffer size is large:
  - All flows tend to lose packets at the same time
  - So all flows will reduce their sending rate at the same time

- Therefore the sum of all the flows tends to act as a single amplified flow

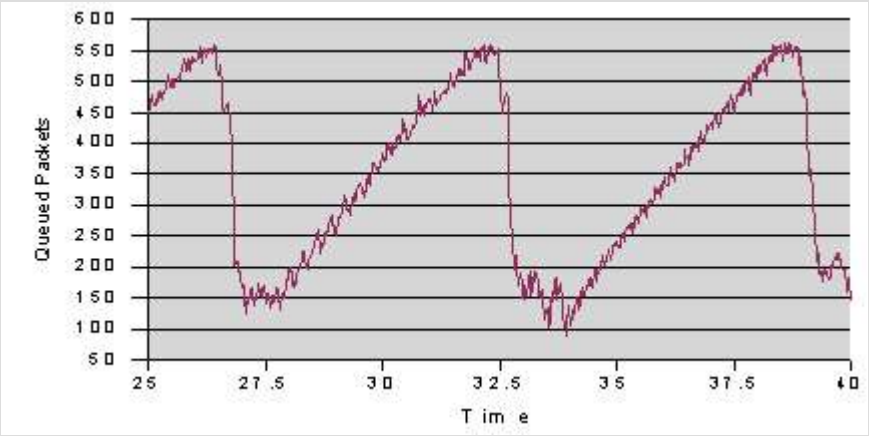


Example: 30 flows through a router with a buffer size of 562 packets

# Background Information
## Synchronization (2)

- If the buffer size is small:
  - Flows do not all lose packets at the same time
  - There is no need to buffer packets to keep the link busy as there is no common sending rate reduction

- The queue length will fluctuate rapidly

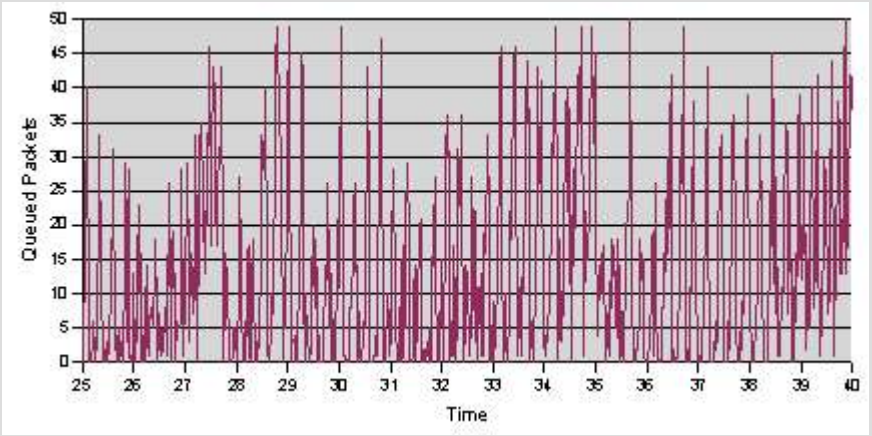


Example: 30 flows through a router with a buffer size of 51 packets

# Experiment Methodology

- Simulation tool: ns-2
- Single congested link
- 45Mb/s data rate on the link
- Constant RTT of 0.1s
- Varied number of flows from 30 to 120

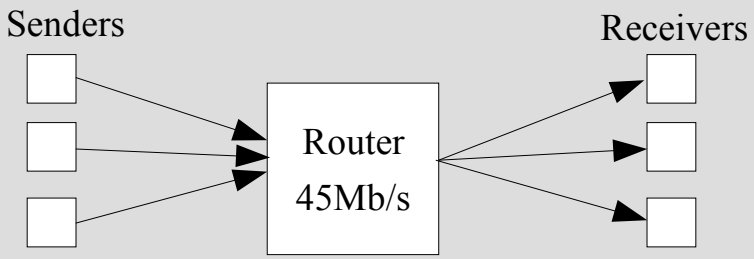

# Experiment Methodology (2)

- The paper “Sizing Router Buffers” suggests a buffer size B=RTT*C/sqrt(n)
- Our buffer size was determined as multiples of the above rule (½x, 1x, 2x, 3x)
- The resulting utilization was measured for both Reno and NewReno flows

# The Results

- The rule B=RTT*C/sqrt(n) resulted in high utilization of the congested link
- Additional buffer reductions are possible assuming TCP NewReno flows

| TCP Flows | RTT*C/sqrt(n) | Buffer Packets | Utilization Reno | Utilization NewReno |
|---|---|---|---|---|
| 30 | 0.5x | 51 | 82.1% | 90.0% |
| 30 | 1x | 103 | 89.6% | 92.3% |
| 30 | 2x | 205 | 94.7% | 96.6% |
| 30 | 3x | 308 | 97.7% | 98.8% |
| 60 | 0.5x | 36 | 80.5% | 86.2% |
| 60 | 1x | 73 | 91.9% | 94.8% |
| 60 | 2x | 145 | 92.1% | 95.7% |
| 60 | 3x | 218 | 96.8% | 98.5% |
| 90 | 0.5x | 30 | 82.0% | 86.8% |
| 90 | 1x | 59 | 95.0% | 96.9% |
| 90 | 2x | 119 | 92.0% | 96.2% |
| 90 | 3x | 178 | 90.5% | 98.0% |
| 120 | 0.5x | 26 | 81.4% | 88.3% |
| 120 | 1x | 51 | 95.5% | 97.3% |
| 120 | 2x | 103 | 91.5% | 98.2% |
| 120 | 3x | 154 | 87.6% | 97.9% |

# Thank You

Questions?