Rate Based Congestion Control for the Internet

(work in progress)

Rui Zhang-Shen
Nandita Dukkipati

Outline

- Problem Statement
- The Model
- Rate Control Protocol (RCP)
- Future Work

Internet Congestion Control

Goals of a congestion control scheme:

- Limit flow rates to avoid “congestion”
- Use the network resources efficiently to minimize flow durations
- Ensure fairness in resource allocation among flows

Problem Statement

- Finding congestion control scheme for high bandwidth Internet
  - Link capacities will continue to increase
  - Conjecture: Flow sizes remain relatively constant
  - An increasing number of flows could finish within a round-trip time (RTT)
- Current congestion schemes:
  - Feedback based
  - Force flows to last multiple RTTs
- Short flows: flow size / link capacity << RTT

Example: A Short Flow

TCP Congestion Control

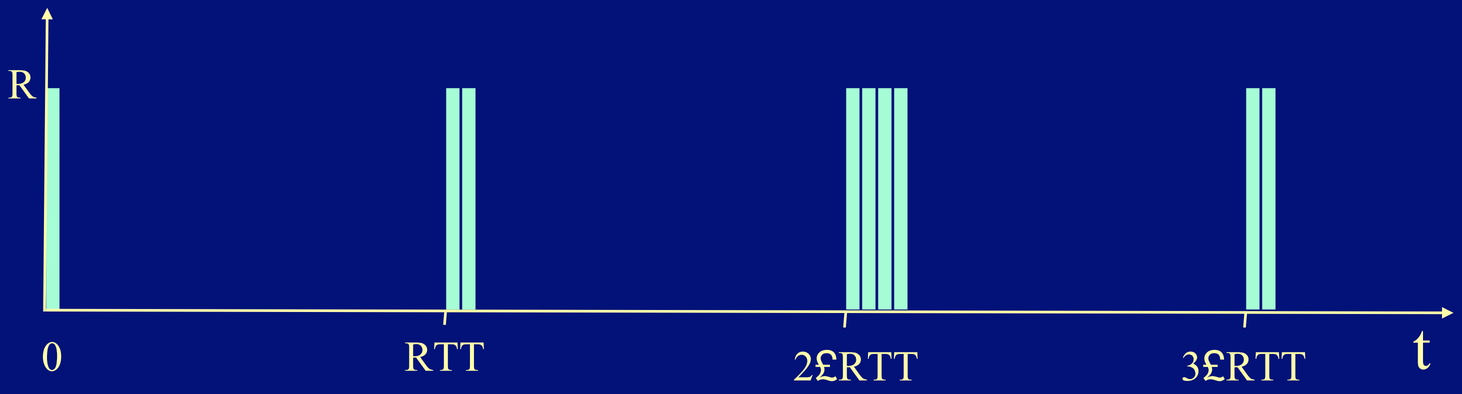

But the flow could finish within one RTT!

We want something like this!

Characteristics of the New Scheme

- Open loop based:
  - No feedback
  - Flow rate determined at start by interaction between routers and end-host
- Flow/user centric instead of packet centric
- Low flow response times
- Fair amongst flows

Assumptions

- Flows arrive according to a Poisson Process
- Flow sizes are independent and identically distributed

- Network consists of short flows only*

*Will be removed later

Optimal Solution

- Single bottleneck link
- Average flow delay minimized when:
  - Flows arrive as single entities
  - Shortest Remaining Processing Time
- Problems
  - Scheduler at output queue
  - Knowledge of flow sizes and per flow state
  - Large buffer

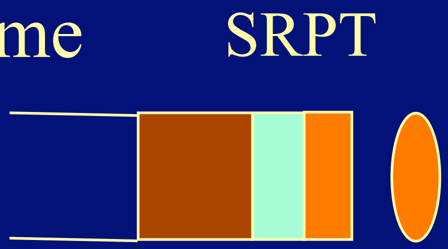

System Model

M/G/1 input process to a single server queue

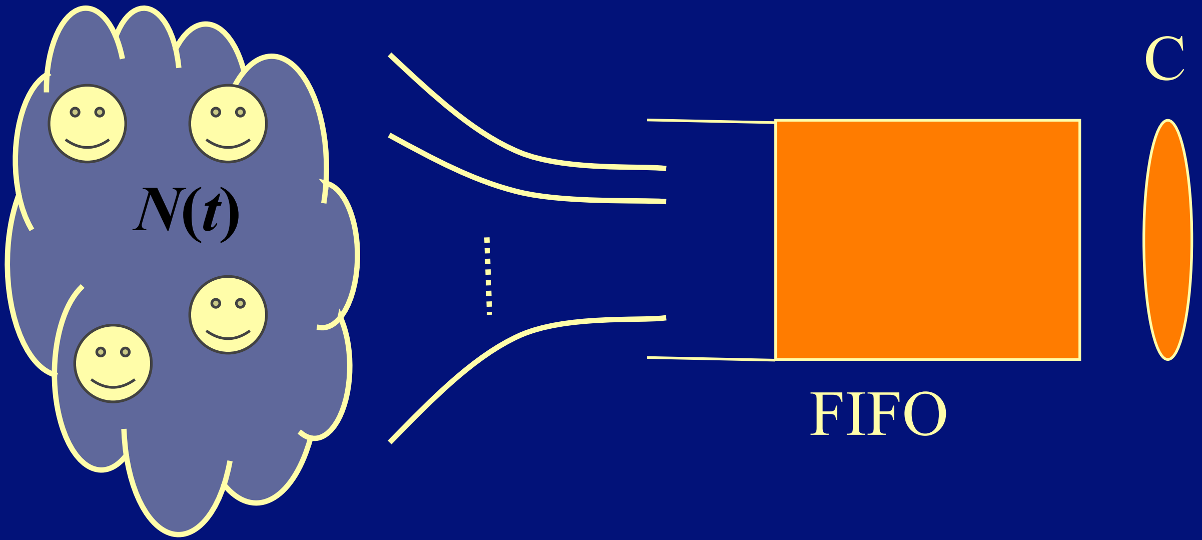

Outline

- Problem Statement
- The Model
- Rate Control Protocol (RCP)
- Future Work

Observations and Intuition

- With an infinite buffer and □<1
  - If flows come at maximum rate, they will be served
  - But maximum buffer occupancy may be large
- If flow rates are small
  - Load is smoother
  - Buffer requirement is lower

An Example

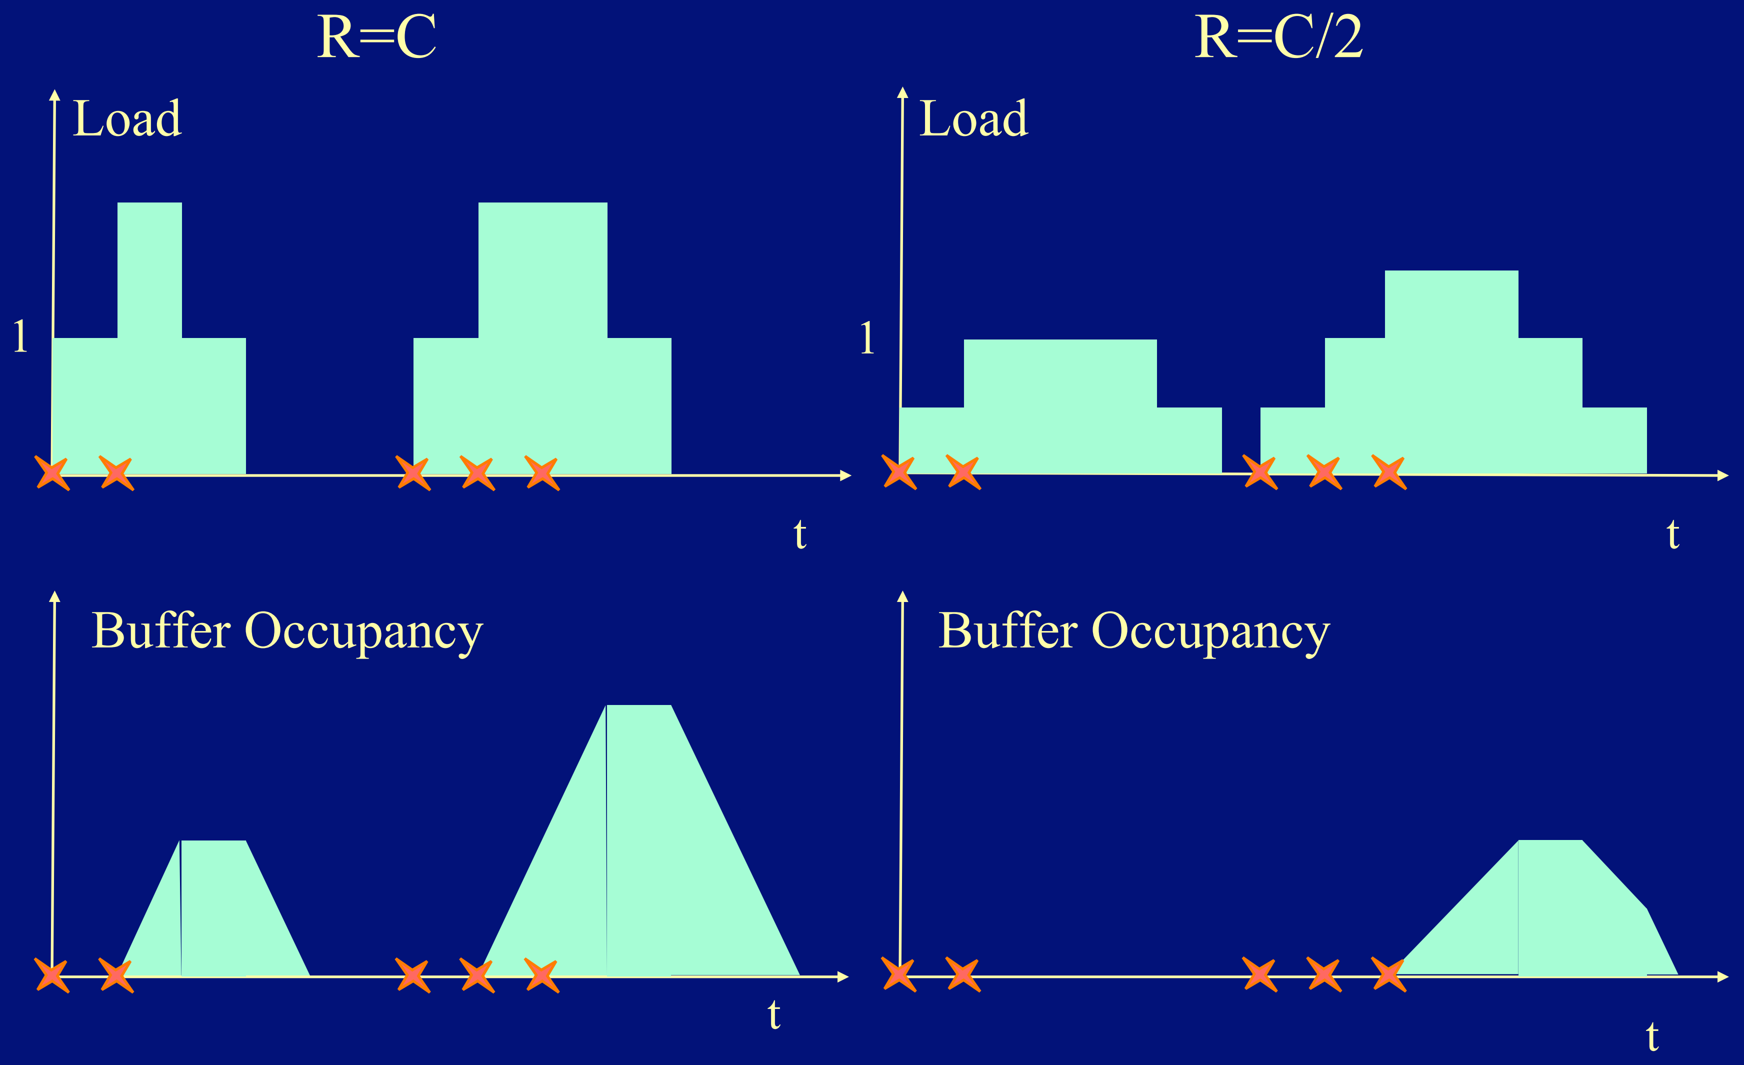

Rate Control Protocol (RCP)

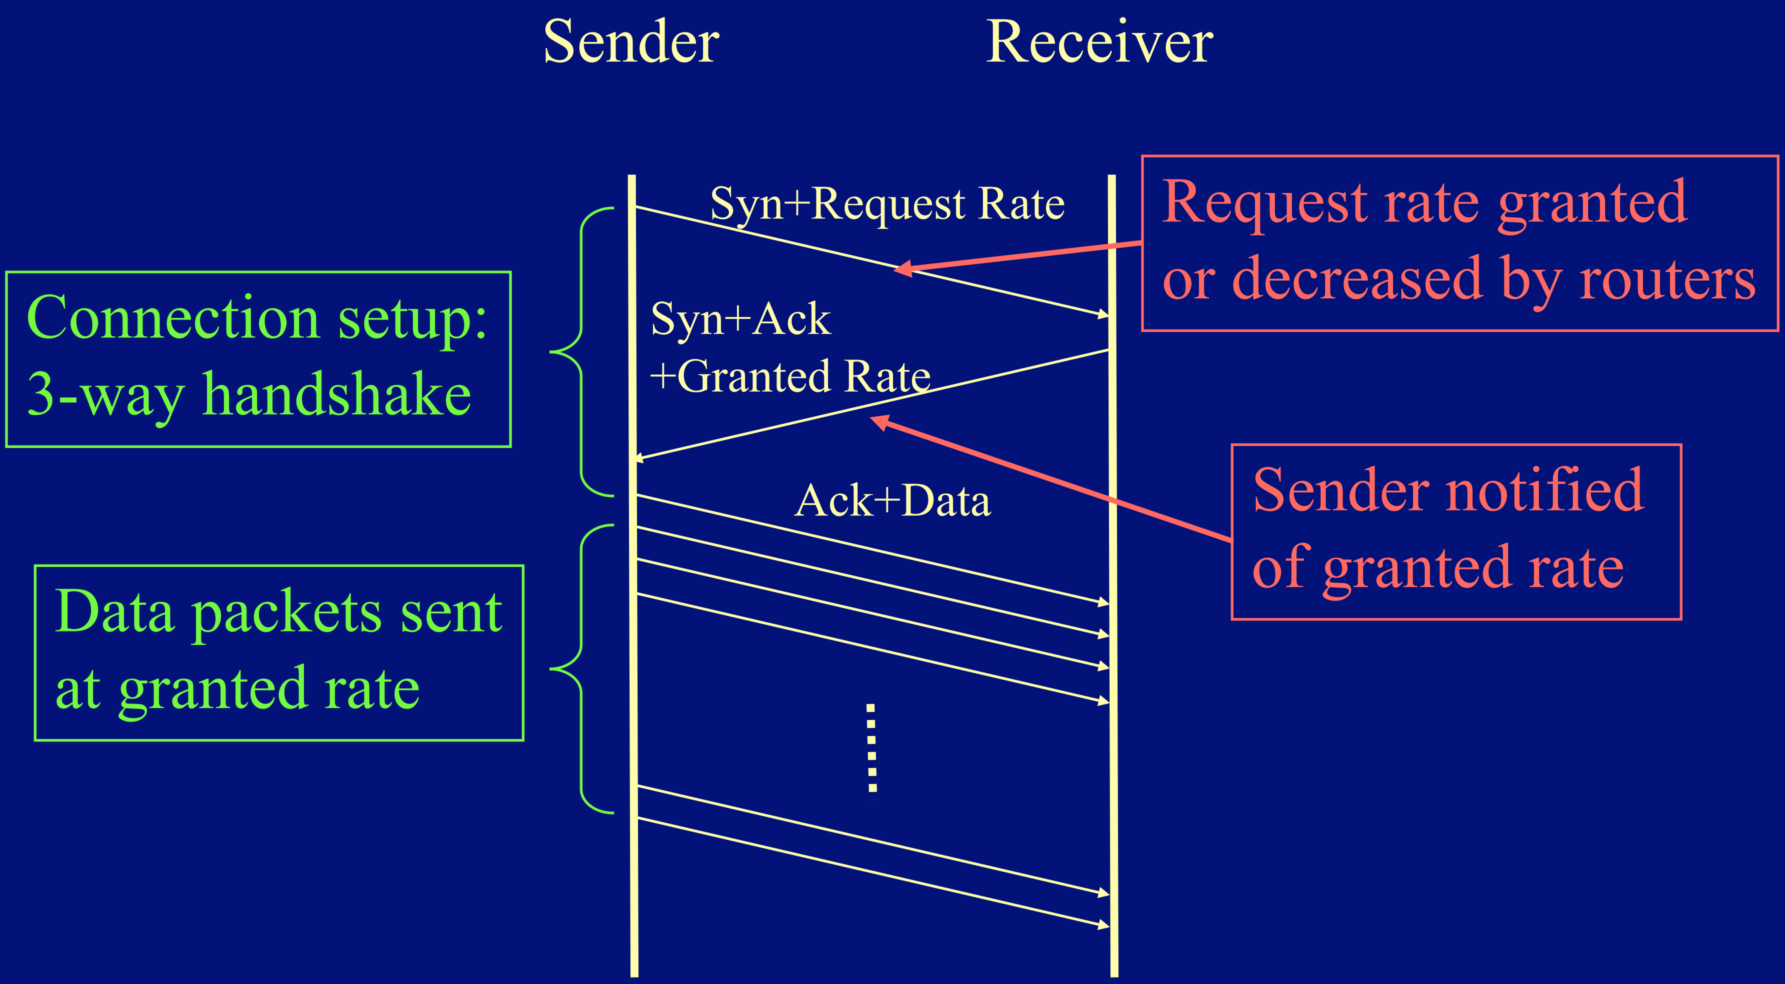

Request rate granted or decreased by routers

Connection setup: 3-way handshake

Sender notified of granted rate

Data packets sent at granted rate

Rate Control Protocol: Router

- Router determines "*n*"
- Router grants rate C/*n* to each flow
- This provides fairness amongst flows
- Choosing *n* involves considerations of
  - Buffer requirement
  - And flow response time

RCP in Action

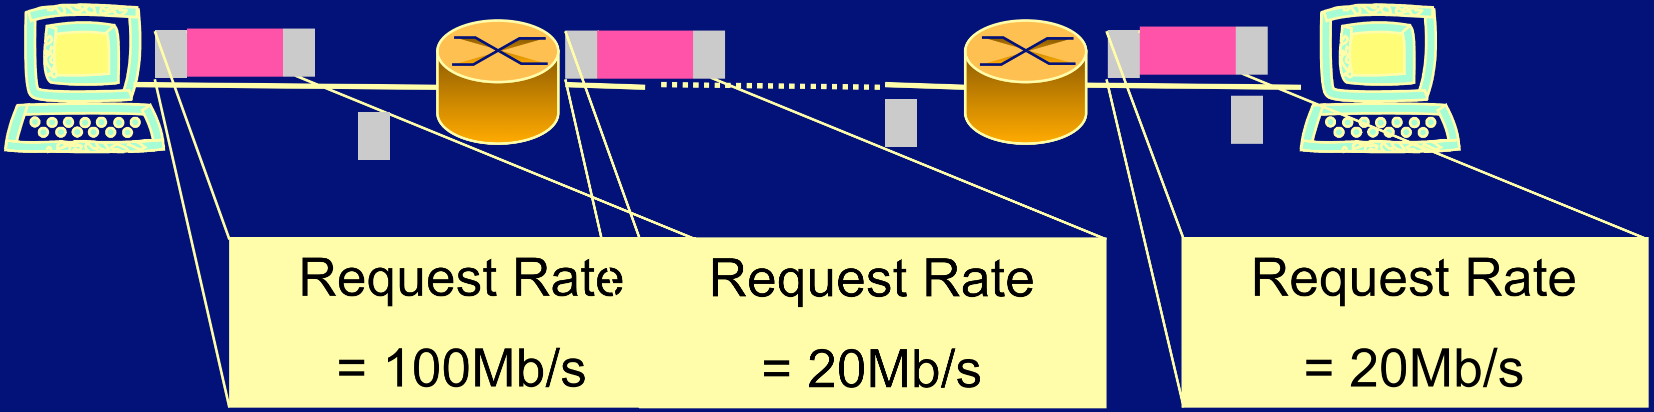

Request Rate

= 100Mb/s

Request Rate

= 20Mb/s

Request Rate

= 20Mb/s

M/G/1 Input Process

- Arrival process (flow start times) and departure process (flow finish times) are both Poisson(□)
- Number of active flows $N(t)$ is Poisson(□E□)=Poisson(n□)
- Pr{load(t)>1}=1 - PoissonCDF(n)
- Given Pr{load(t)>1}, n only depends on □!

How to Choose “*n*”

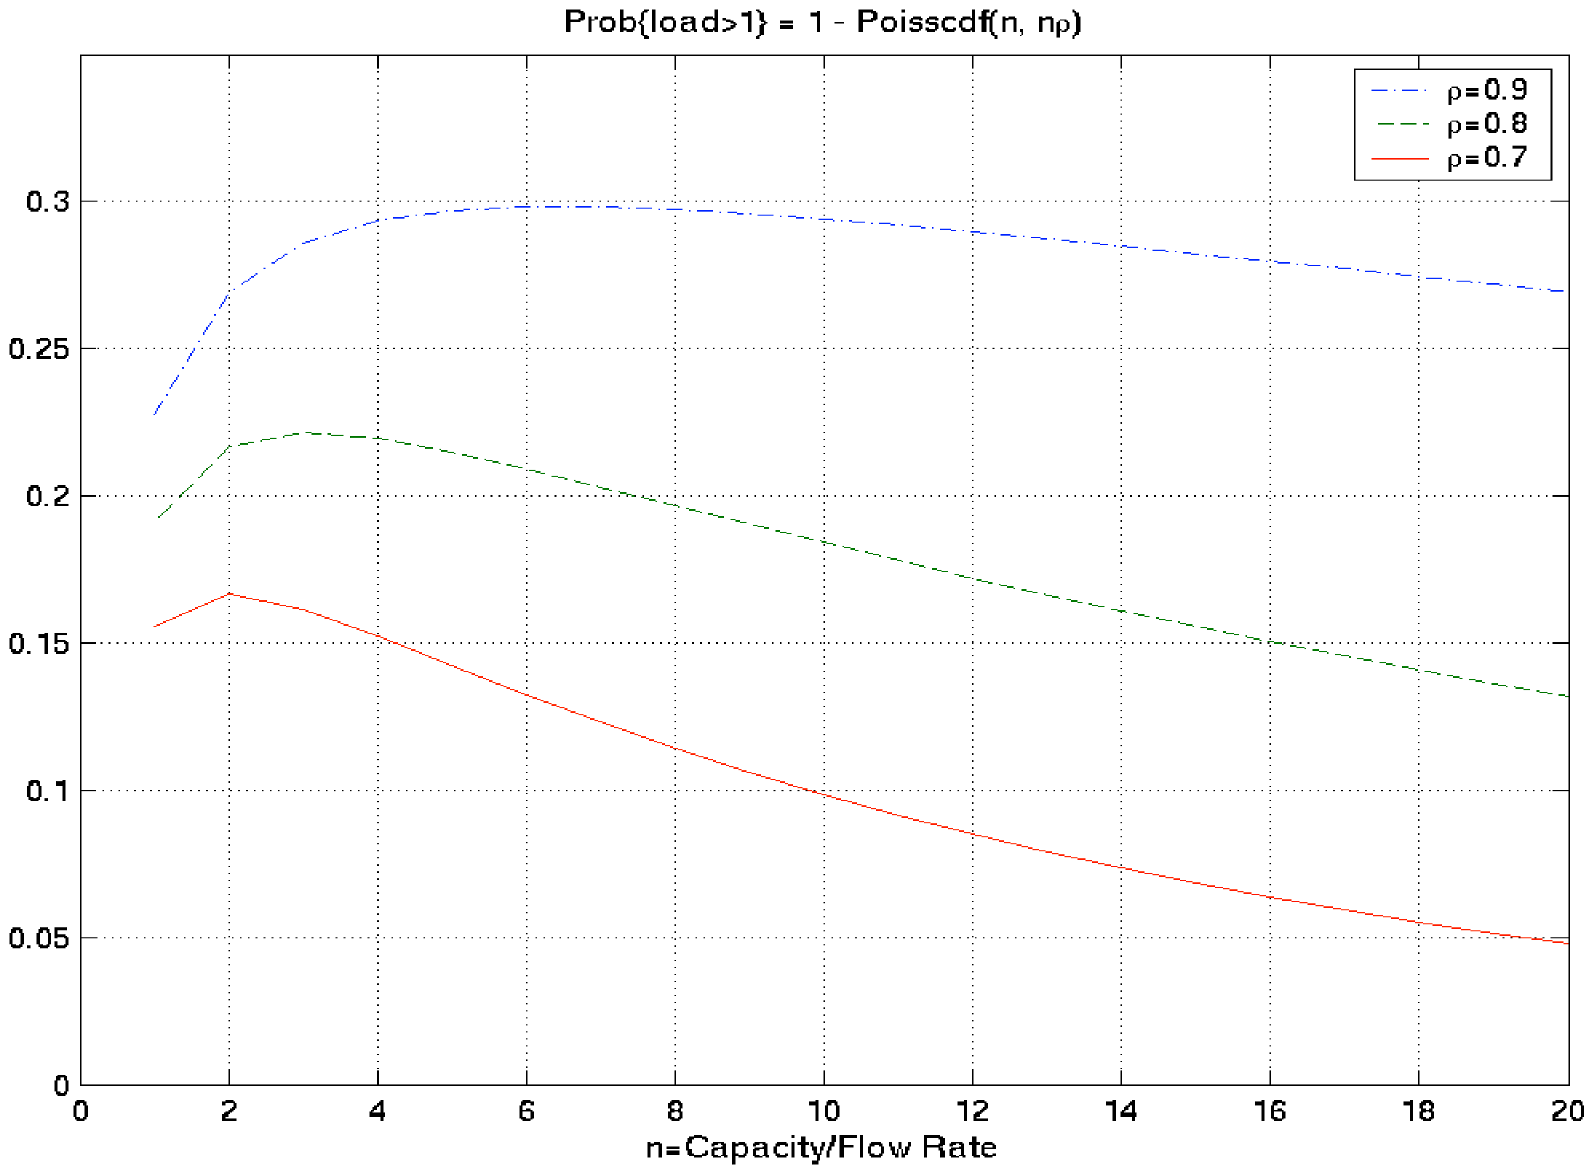

Simulation Comparison of TCP and RCP

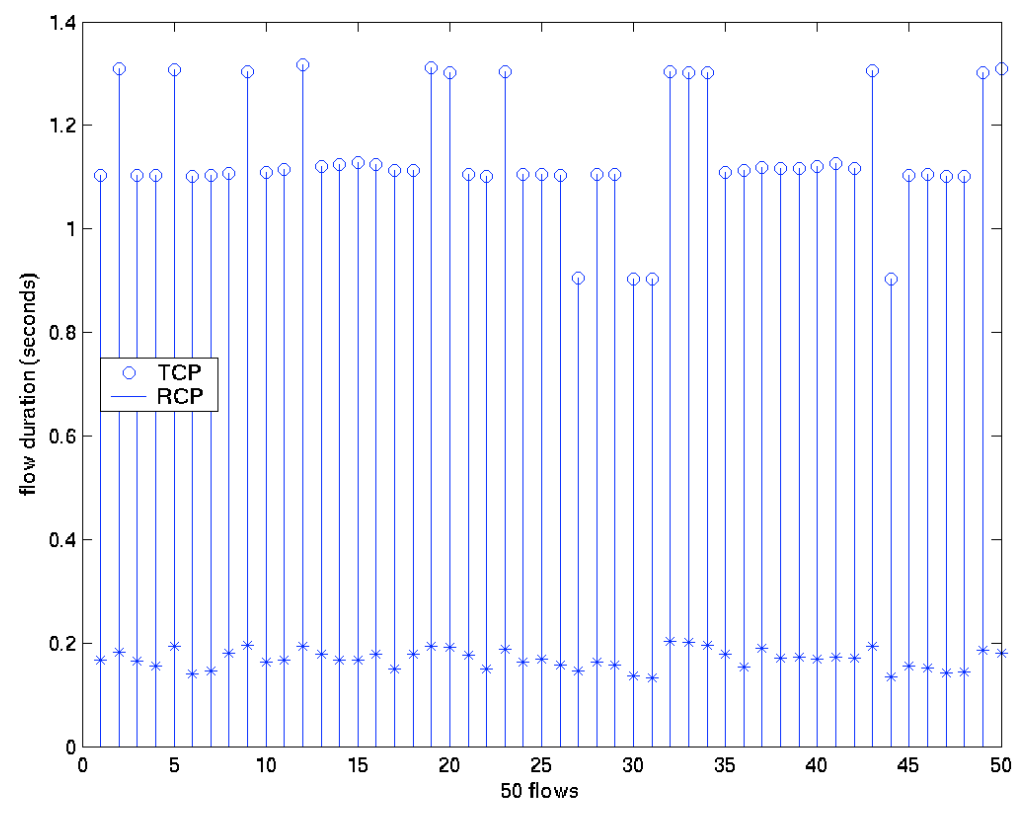

Single bottleneck link of 100Mb/s; RTT=200ms; □=0.8; *n*=8; flow size~Uniform[50, 150]kB

*Pr{load(*t*)>1} from simulation matches theoretical values

Future Work

- Derive buffer distribution with different “G”
- Investigate effects of buffer size on RCP
- Intuitively, RCP should also work for long flows
  - Performance of RCP with long-tailed traffic, etc.
- User models: Interaction of load and network
- Find optimal *n* which minimizes average delay