Random Early Detection Gateways for Congestion Avoidance*

* 5886 citations as of 19Sep11

Sally Floyd and Van Jacobson,

IEEE Transactions on Networking,
Vol.1, No. 4, (Aug 1993), pp.397-413.

Presented by Bob Kinicki

Outline

- Introduction
- Background: Definitions and Previous Work
- The **RED** Algorithm
- **RED** Parameters
- **RED** Simulation Results
- Evaluation of **RED** Design Goals
- Conclusions and Future Work

Introduction

Main idea :: *to provide **congestion control** at the router for TCP flows.*

- **RED Algorithm Goals**
  - The **primary goal** is to provide congestion avoidance by controlling the average queue size such that the router stays in a region of **low delay** and **high throughput**.
  - To avoid global synchronization (e.g., in Tahoe TCP).
  - To control misbehaving users (this is from a fairness context).
  - To seek a mechanism that is not biased against bursty traffic.

Definitions

- ***congestion avoidance*** – when impending congestion is indicated, take action to avoid congestion.
- ***incipient congestion*** – congestion that is beginning to be apparent.
- need to notify connections of congestion at the router by either ***marking*** the packet [ECN] or ***dropping*** the packet {This assumes a router drop is an implied signal to the source host.}

Previous Work

- Drop Tail (FIFO)
- Random Drop
- Early Random Drop
- Source Quench Messages
- DECbit Scheme

Drop Tail Router

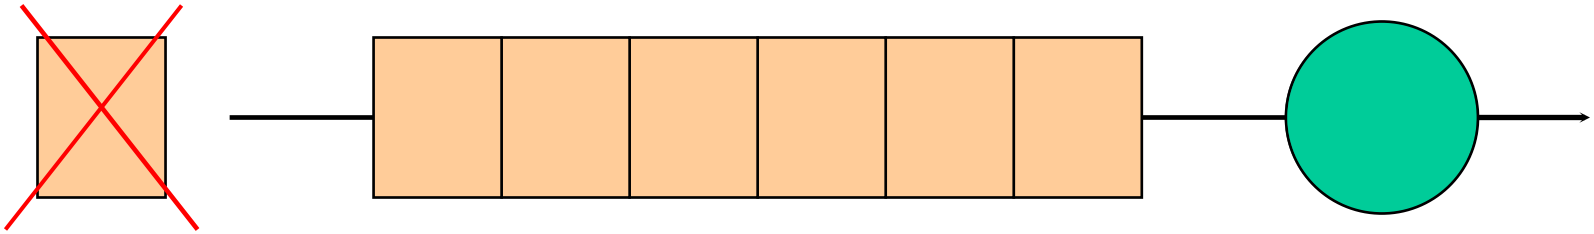

- FIFO queuing mechanism that drops packets from the tail when the queue overflows.
- Introduces ***global synchronization*** when packets are dropped from several connections.

Random Drop Router

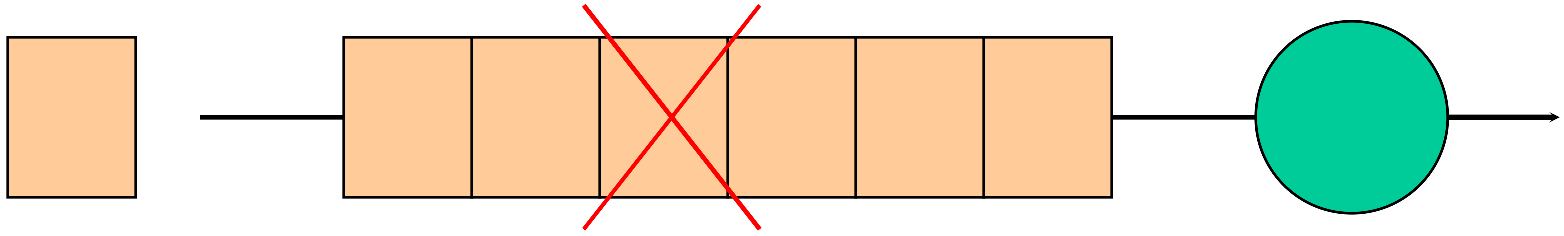

- When a packet arrives and the queue is full, randomly choose a packet from the queue to drop.

Early Random Drop Router

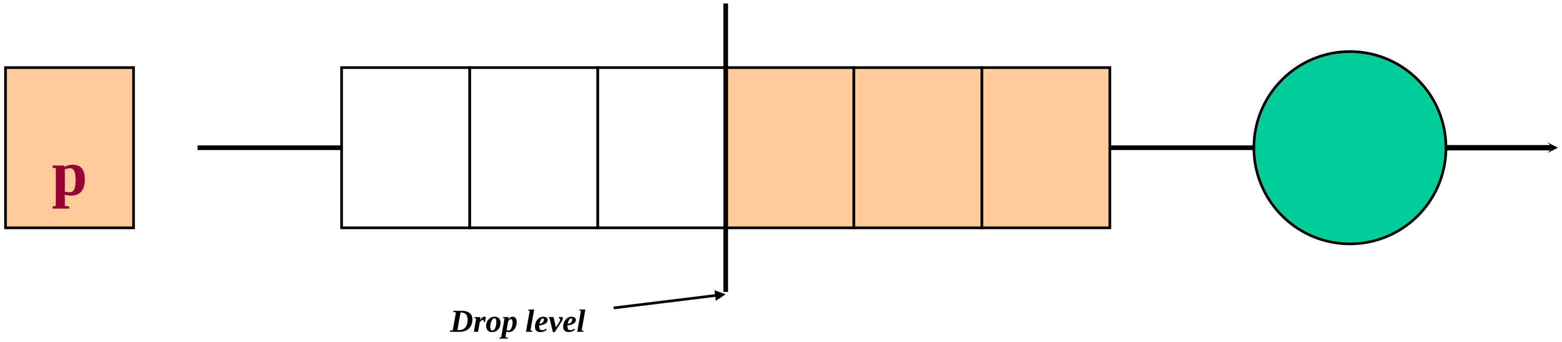

- If the queue length exceeds a drop level, then the router drops each arriving packet with a fixed *drop probability* p.
- Reduces global synchronization.
- Does **not** control misbehaving users (UDP flows).

Source Quench Messages

- Router sends ***source quench messages*** back to source <u>before</u> the queue reaches capacity.
- Complex solution that gets router involved in end-to-end protocol.
- This solution violates the end-to-end tenet of Internet architects!!

DECbit Scheme

- Uses a ***congestion-indication bit*** in packet header to provide feedback about congestion.
- Upon packet arrival, the average queue length is calculated for last (busy + idle) period plus current busy period.
- When the average queue length exceeds **1**, the router sets the congestion-indicator bit in arriving packet's header.

DECbit Scheme

- The source updates its window every two RTTs.
- If at least half of packets in source's last window have the bit set, decrease the congestion window exponentially.
- Otherwise, the window is increased linearly.

RED Algorithm

for each packet arrival

calculate the average queue size avg

if $min_{th} \leq avg < max_{th}$

calculate the probability p_a

with probability p_a:

mark the arriving packet

else if $max_{th} \leq avg$

mark the arriving packet.

RED Drop Probability (p_a)

$$p_b = max_p \times (avg - min_{th})/(max_{th} - min_{th}) \quad [1]$$

where

$$p_a = p_b / (1 - count \times p_b) \quad [2]$$

Note: this calculation assumes queue size is measured in packets. If queue is in bytes, we need to add [1.a] between [1] and [2]

$$p_b = p_b \times \text{PacketSize/MaxPacketSize} \quad [1.a]$$

avg - Average Queue Length

$$avg = (1 - w_q) \times avg + w_q \times q$$

where q is the newly measured queue length.

This ***Exponential Weighted Moving Average* (EWMA)** is designed such that short-term increases in queue size from bursty traffic or transient congestion do not significantly increase average queue size.

RED/ECN Router Mechanism

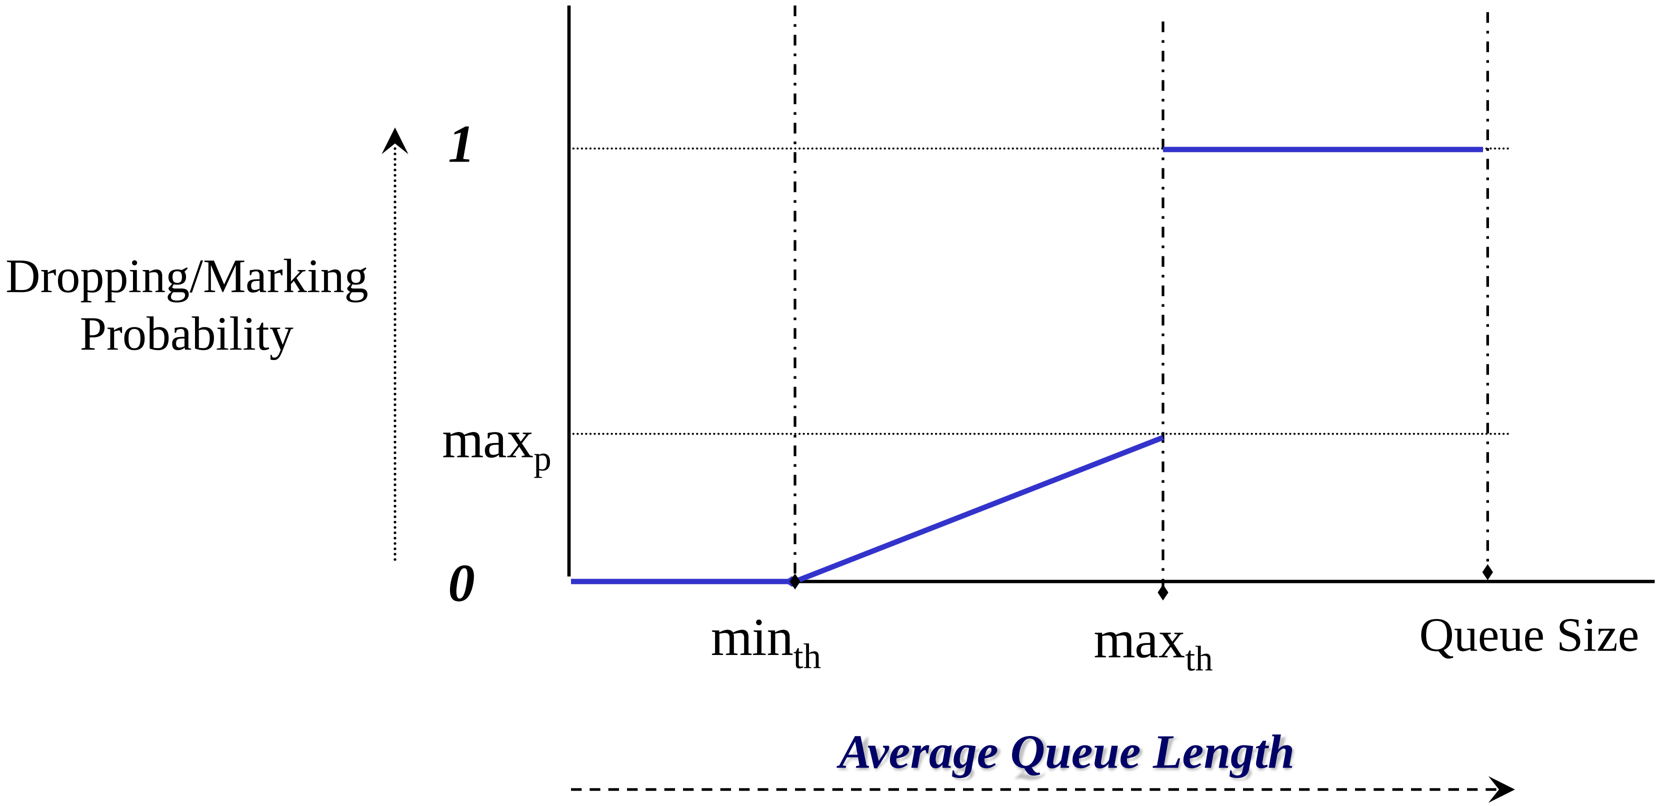

'Gentle' RED

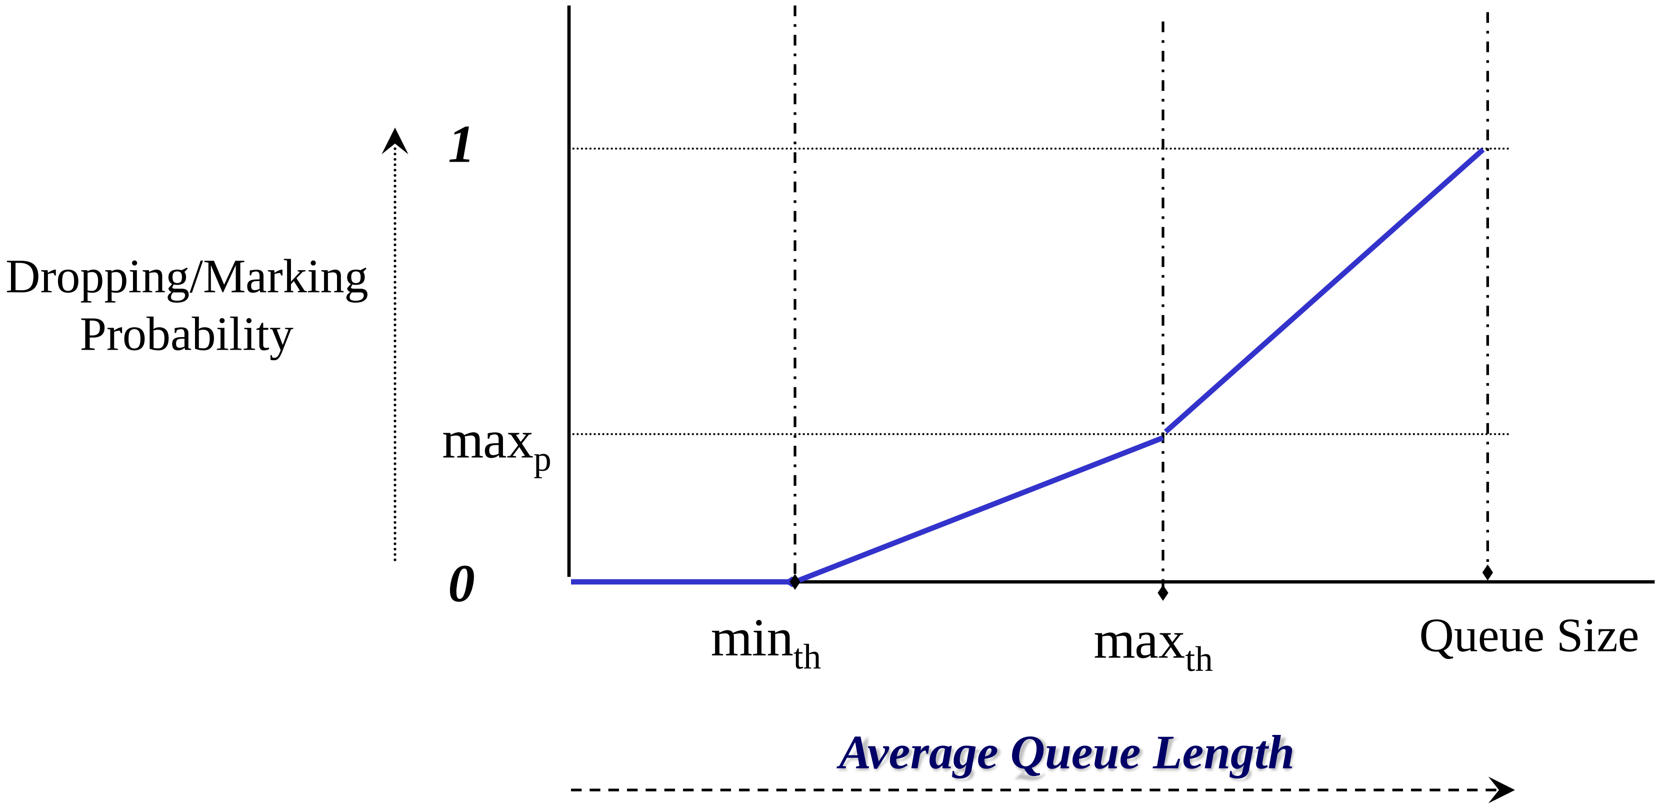

Dropping/Marking Probability
1
max$_p$
0
min$_{th}$
max$_{th}$
Queue Size
Average Queue Length

RED Parameter Settings

- w_q authors suggest $0.001 <= w_q <= 0.0042$
  authors use $w_q = 0.002$ for simulations
- min_{th}, max_{th} depend on desired average queue size
  - bursty traffic ➔ increase min_{th} to maintain link utilization.
  - max_{th} depends on the maximum average delay allowed.
  - **RED** is most effective when $max_{th} - min_{th}$ is larger than typical increase in calculated average queue size in one round-trip time.
  - **"parameter setting rule of thumb":** max_{th} at least twice min_{th} . However, max_{th} = 3 times min_{th} is used in some of the experiments shown.

Packet-Marking Probability

- The goal is to uniformly spread out the *marked* packets. This reduces global synchronization.

Method 1: Geometric random variable

When each packet is marked with probability p_b,
the packet inter-marking time, X, is a geometric random variable with $E[X] = 1/p_b$.

- This distribution will both cluster packet drops and have some long intervals between drops!!

Packet-Marking Probability

Method 2: Uniform random variable

Mark packet with probability $p_b / (1 - count \times p_b)$ where *count* is the number of unmarked packets that have arrived since last marked packet.

$$E[X] = 1/(2\,p_b) + 1/2$$

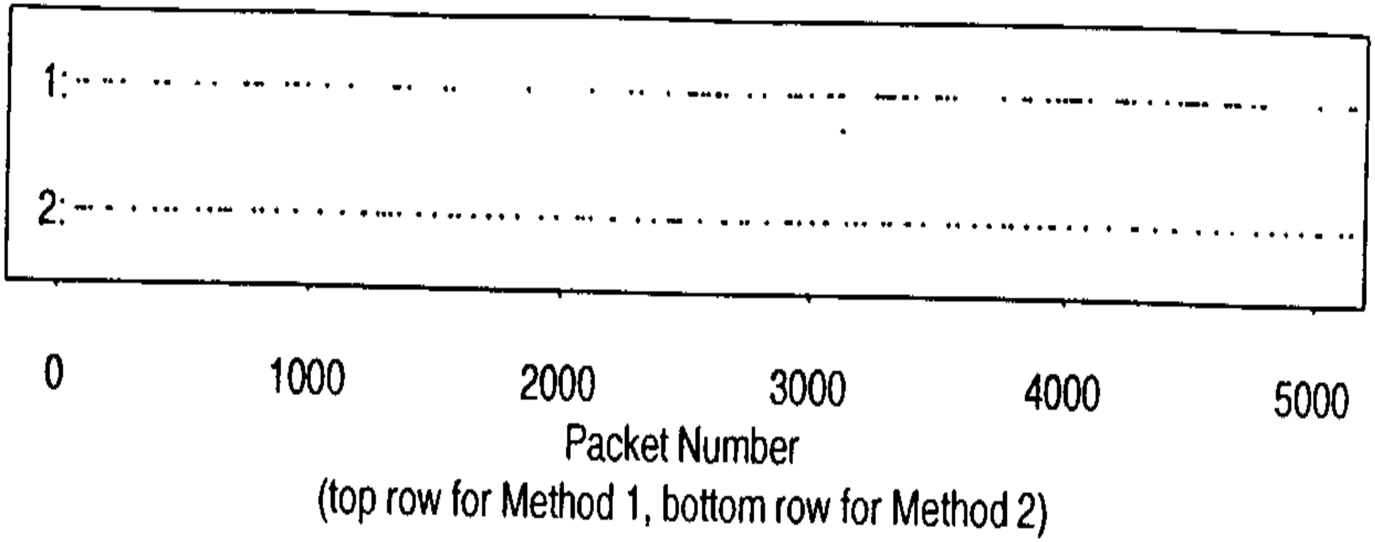

Figure 8: Randomly-marked packets, comparing two packet-marking methods.

Method 1: geometric p = 0.02

Method 2: uniform p = 0.01

Result :: *marked packets more clustered for method 1*

➔ *uniform is better at eliminating "bursty drops"*

Setting max_p

- "**RED** performs best when packet-marking probability changes fairly slowly as the average queue size changes."
  - This is a **stability argument** in that the claim is that **RED** with small ***max_p*** will reduce oscillations in ***avg*** and actual marking probability.
- They recommend that ***max_p*** never be greater than 0.1

{This is not a robust recommendation.}

RED Simulation* Results

- Figure 4: Four heterogeneous **FTP** sources
- Figure 6: Two homogeneous **FTP** sources
- Figure 10: 41 Two-way, short **FTP** and **TELNET** flows
- Figure 11: Four **FTP non-bursty** flows and **one bursty FTP** flow

*As direct predecessor of NS2, Real simulator is paper's biggest contribution.

Simple Simulation
Four Heterogeneous FTP Sources

TCP Tahoe
1KB packet size
$w_q = 0.002$
$max_p = 1/50$
$min_{th} = 5$
$max_{th} = 15$
max $cwnd = bdp$
Large buffer with no drop-tail packet drops.

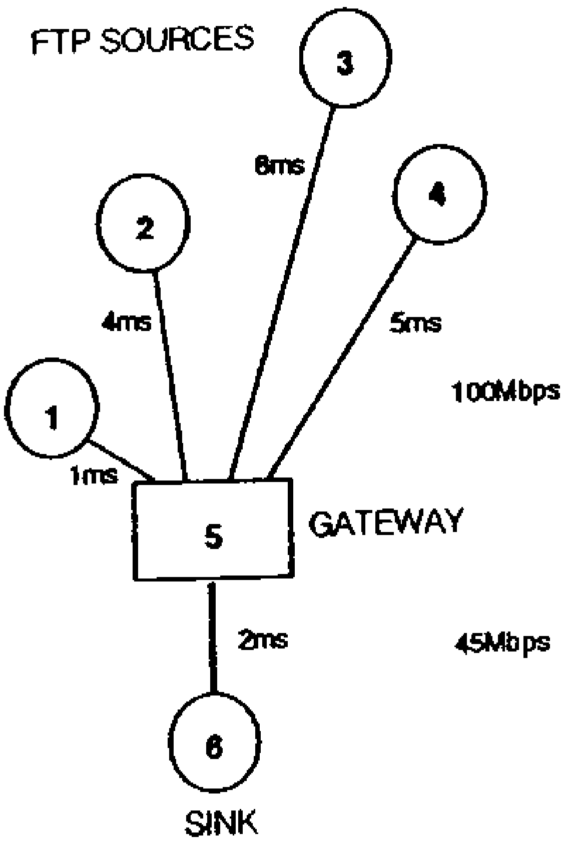

Figure 4: Simulation network.

Note: staggered start times and uneven throughputs

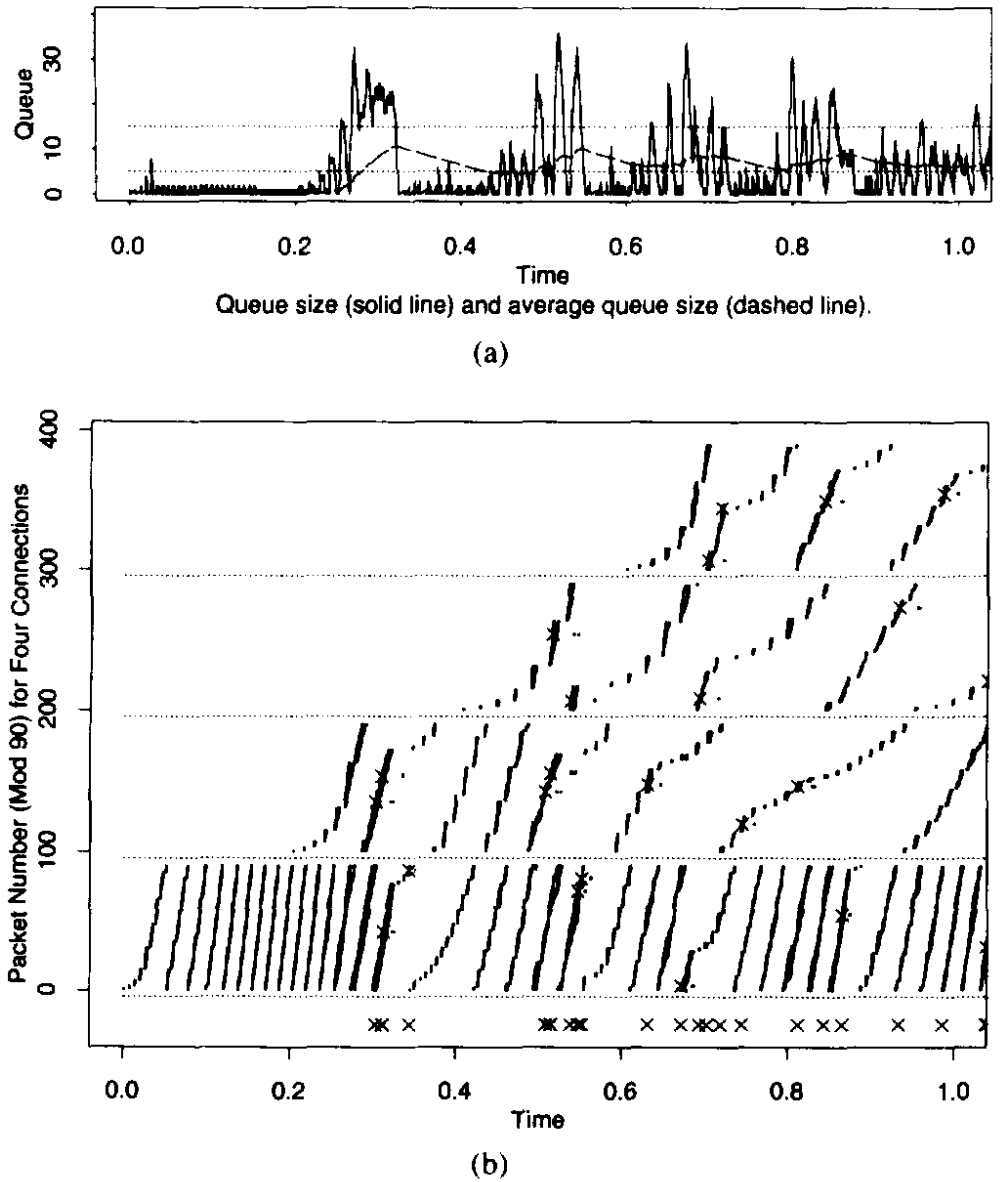

Fig. 3. A simulation with four FTP connections with staggered start times.

Two Homogeneous FTP Sources

- RED varies min_{th} from 3 to 50 packets with fixed buffer of 100 packets.
- max_{th} set to 3 min_{th}
- Drop Tail varies buffer from 15 to 140 packets
- max cwnd = 240 packets

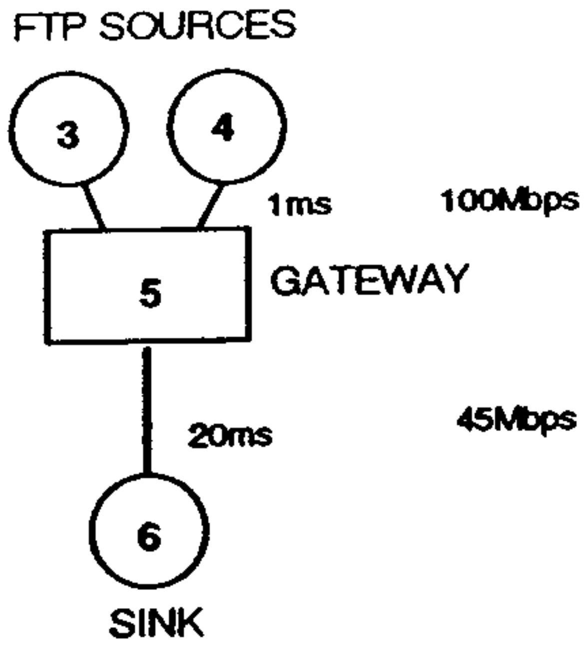

Figure 6: Simulation network.

Two Homogeneous FTP Sources

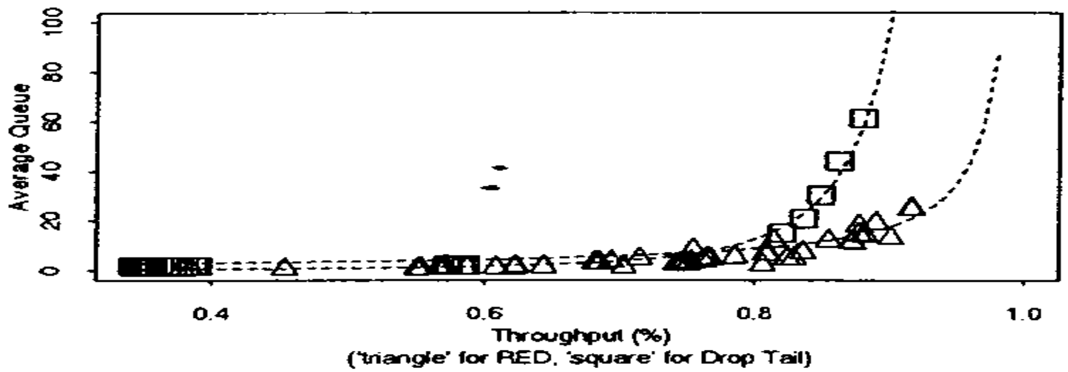

Figure 5: Comparing Drop Tail and RED gateways.

Figure 5 represents many 5-sec. simulation experiments.

RED yields lower queuing delay as utilization improves by increasing min_{th} from 3 to 50 packets.

Drop-tail yields *unacceptable delay* at high utilization.

The power measure is better for **RED** !

Network with 41 Short Duration Connections

Traffic going in both directions.

Mostly FTP and a few TELNET connections.

Total packets varies from

20 to 400 packets per connection.

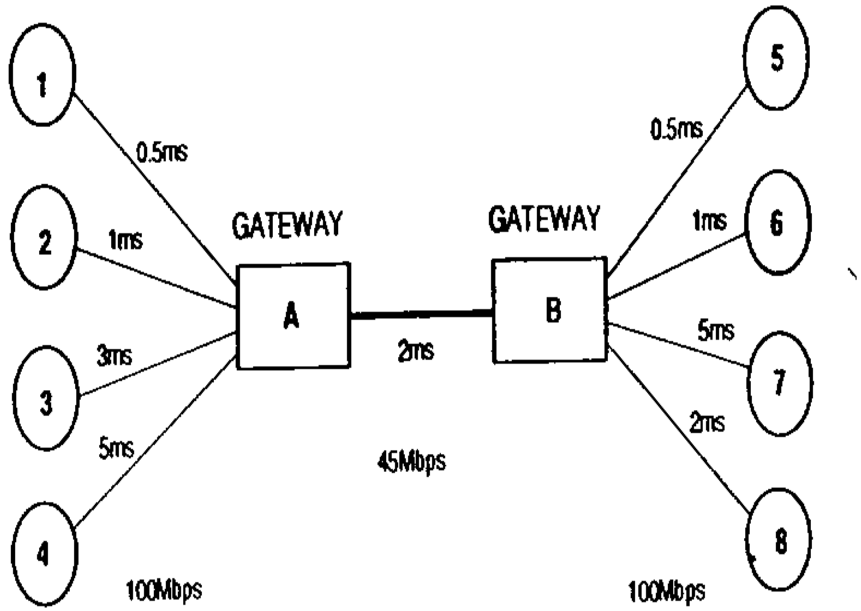

Figure 10: A network with many short connections.

Short, Two-Way
FTP and TELNET Flows

- **RED** controls the average queue size in both directions.
- Flows have small **cwnd maximums (8 or 16).**
- Packet dropping is higher and bursty.
- Low utilization: **(61% and 59%).**
- **ACK-compression** contributes to bursty packet arrivals.

.

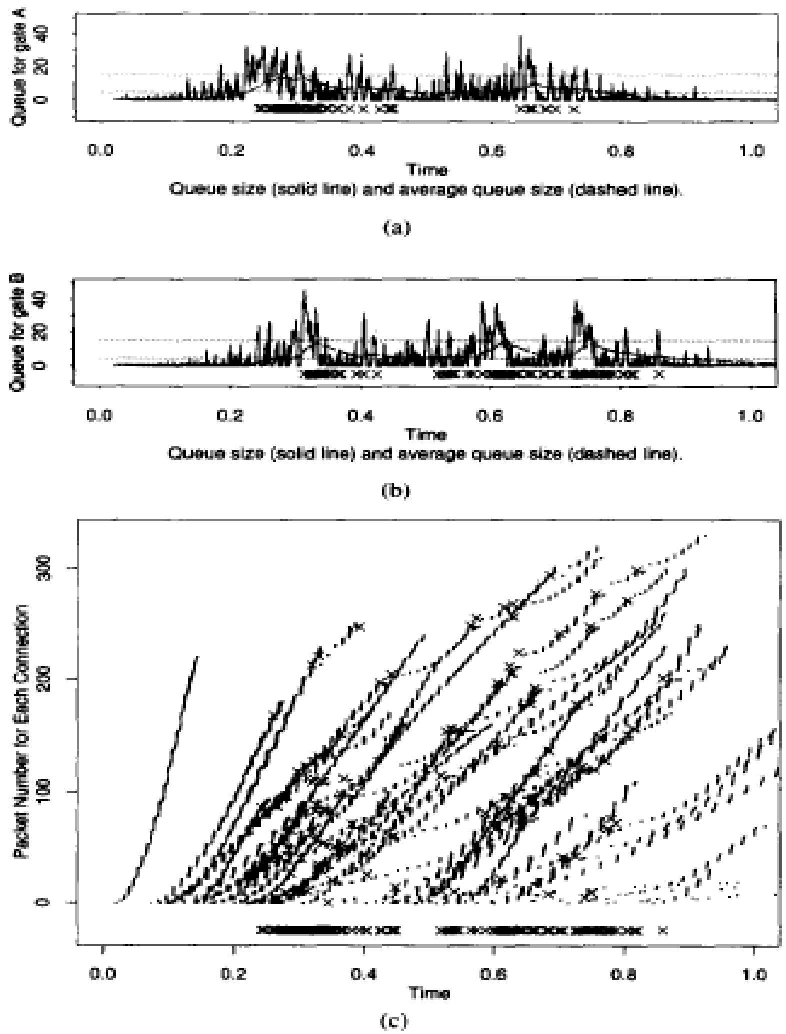

Fig. 9. A RED gateway simulation with heavy congestion, two-way traffic, and many short FTP and TELNET connections.

Five FTP Flows
Including One Bursty Flow

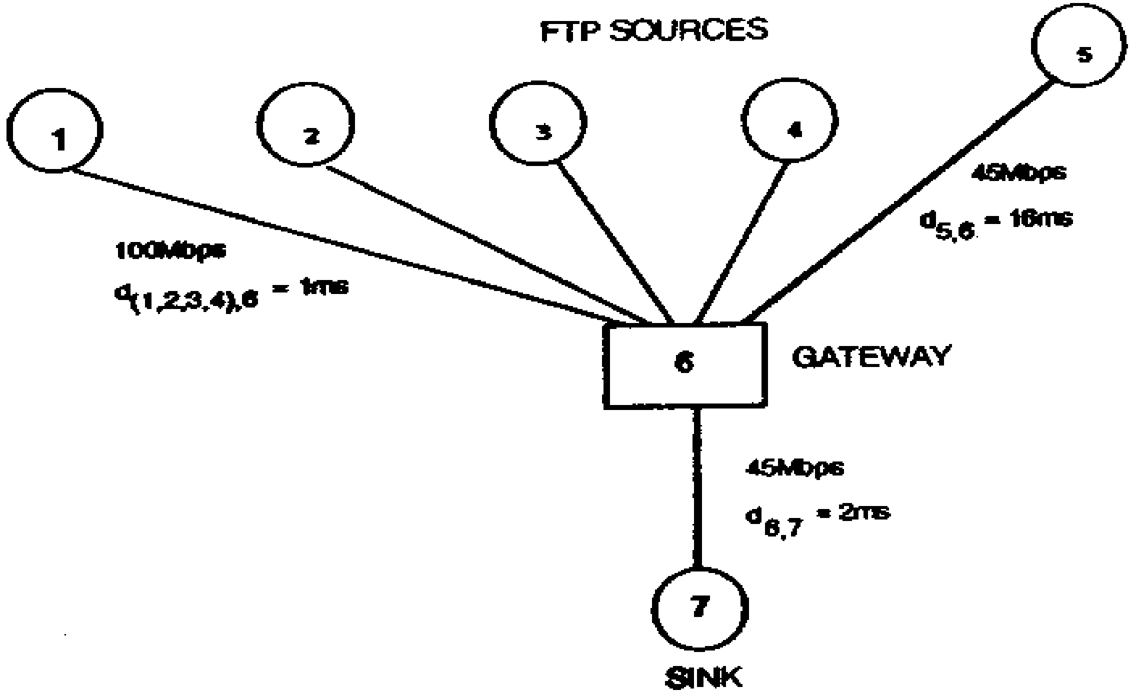

Figure 11: A simulation network with five FTP connections.

Simulation Details

- Flow 5 is a bursty flow due to large RTT, small cwnd (8 packets).
- First four flows are **robust** due small RTT, small cwnd of 12 packets.
- Gateway buffer varies from 8 to 22 packets for “drop” queues.
- Each simulation runs for 10 seconds and each mark in the figures represents one second (i.e., 10 throughput data points per cwnd size).
- Graphs show flow 5 utilizations.

Drop Tail Gateways

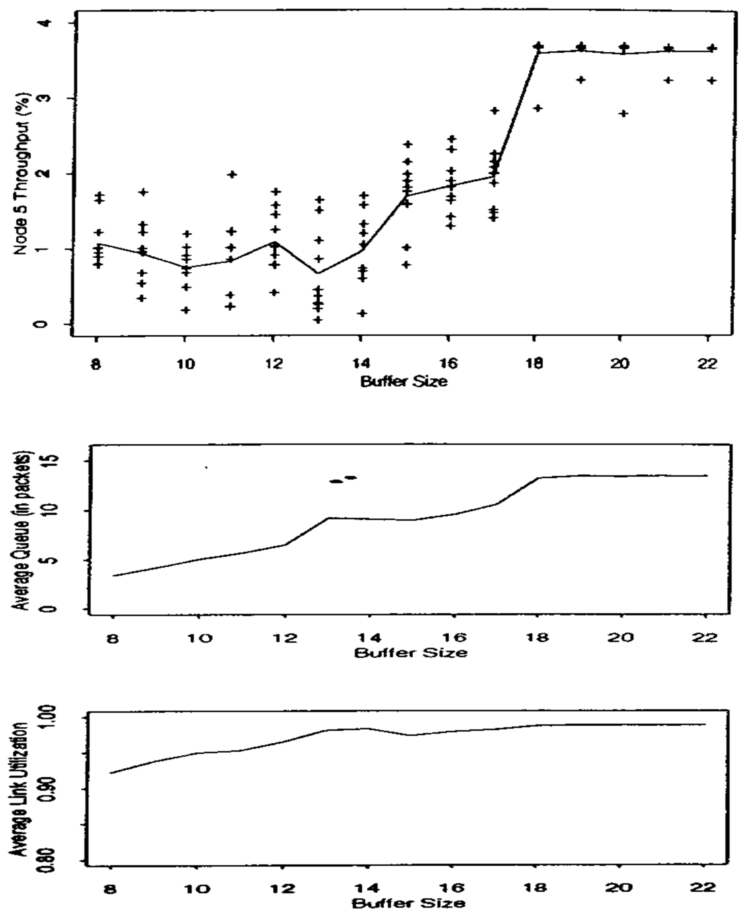

Figure 12: Simulations with Drop Tail gateways.

Random Drop Gateways

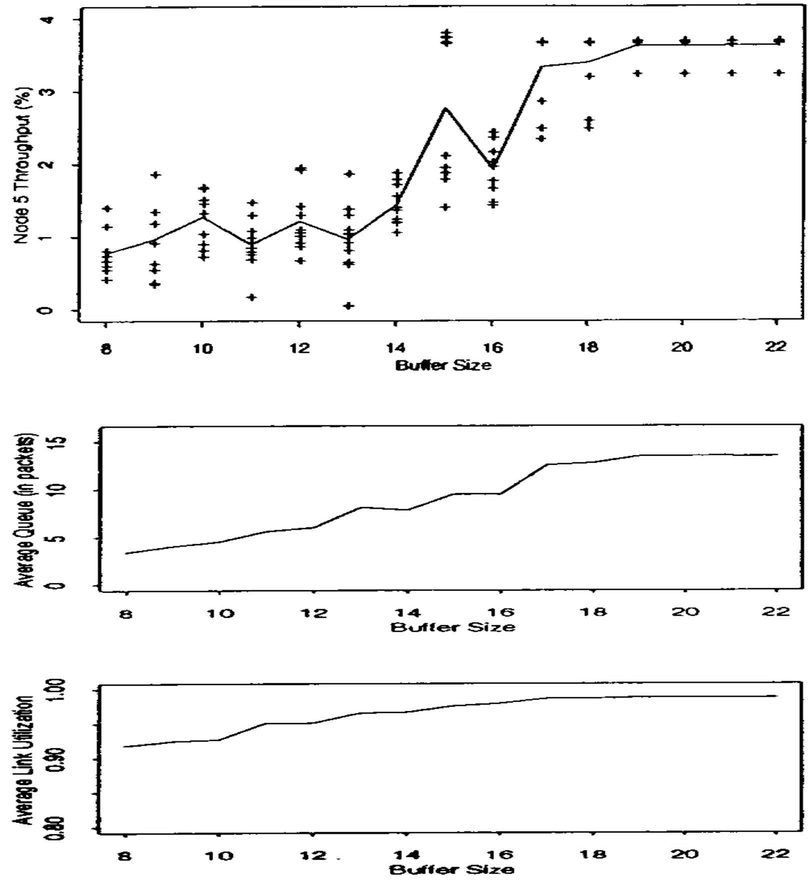

Figure 13: Simulations with Random Drop gateways.

RED Gateways

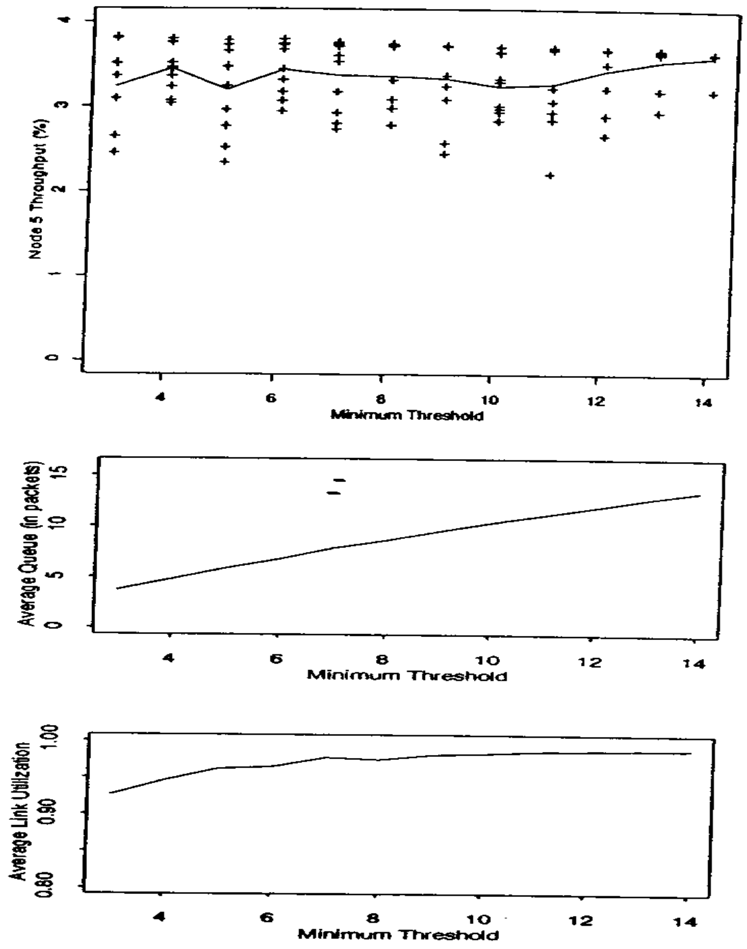

Figure 14: Simulations with RED gateways

Bursty Flow Packet Drop Bias

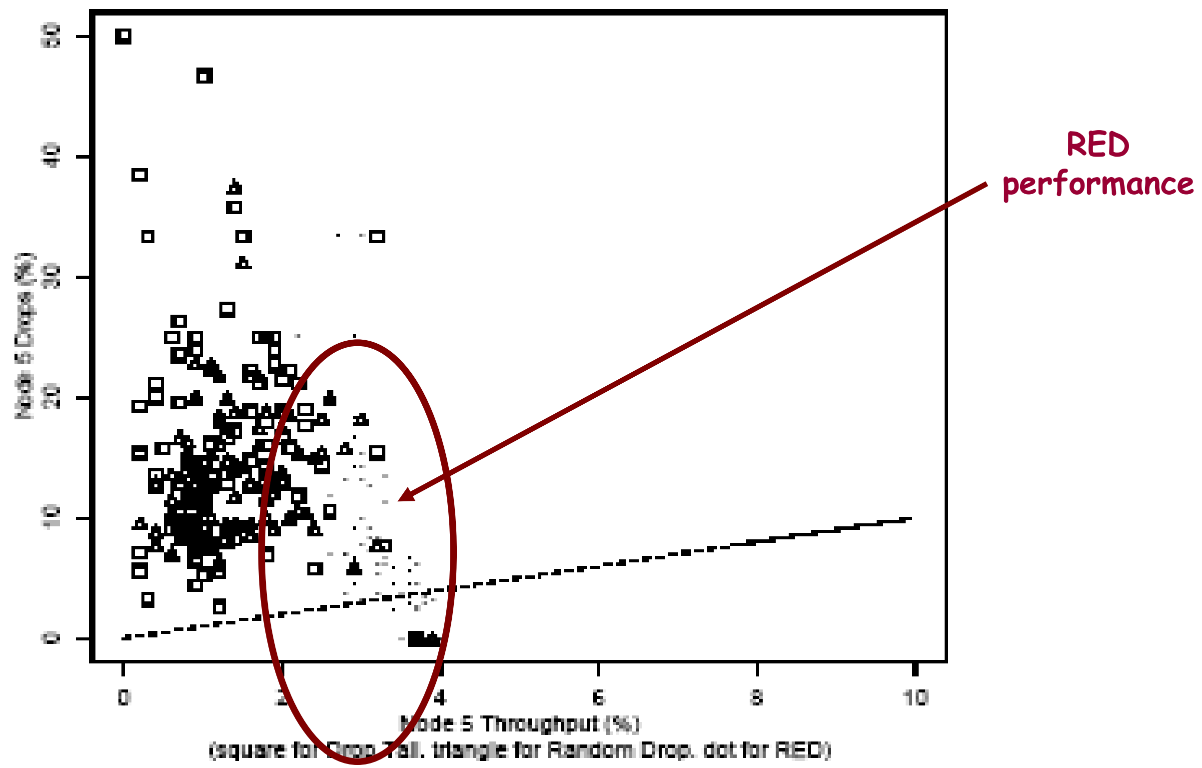

Figure 15: Scatter plot, packet drops vs. throughput

Identifying Misbehaving Flows

The assumption is marked packets matches the flows' share of the bandwidth.

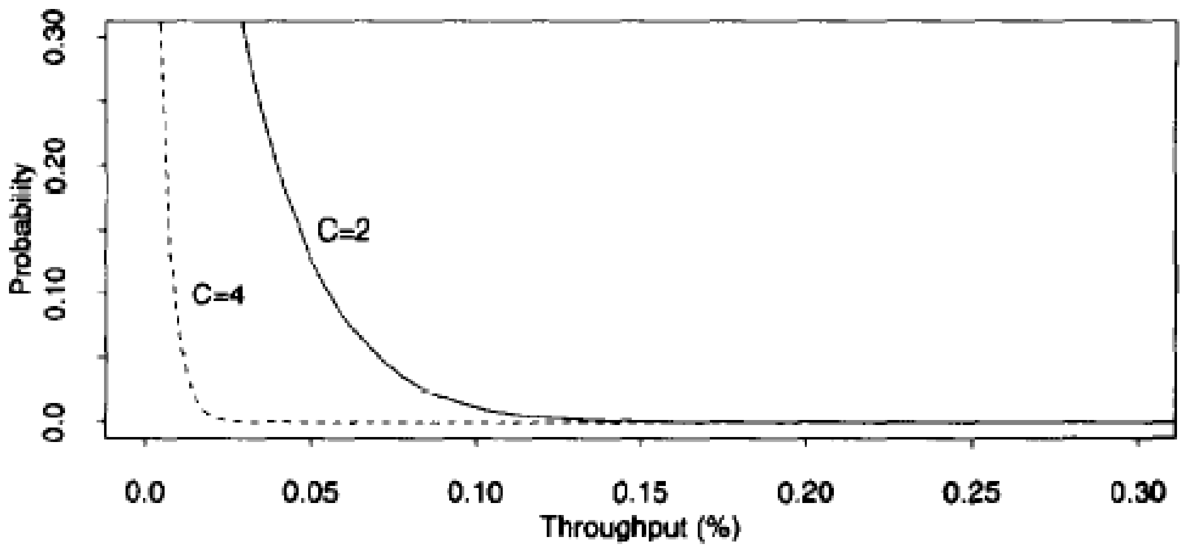

Fig. 16. Upper bound on probability that a connection's fraction of marked packets is more than C times the expected number, given 100 total marked packets.

Evaluation of RED Design Goals

- **congestion avoidance**
  - If **RED** *drops* packets, this guarantees the calculated average queue size does not exceed the max threshold. If w_q is set properly, **RED** controls the *actual* average queue size.
  - If **RED** *marks* packets when ***avg*** exceeds max_{th}, the router relies on source cooperation to control the average queue size. **{not part of RED, this is ECN.}**

Evaluation of RED Design Goals

- **appropriate time scales**
  - claim:: The detection time scale *roughly matches* time scale of source's response to congestion.
  - **RED** does not notify connections during transient congestion at the router.

{This argument is weak here.}

Evaluation of RED Design Goals

- **no global synchronization**
  - RED avoids global synchronization by marking at as low a rate as possible with marking distribution spread out.
- **simplicity**
  - detailed argument about how to cheaply implement in terms of adds and shifts.
- {Historically, the **simplicity** of **RED** has been strongly refuted because **RED** has too many parameters to make it robust.}

Evaluation of RED Design Goals

- **maximizing global power**
  - ***power*** *is ratio of* ***throughput*** *to* ***delay*****.**
  - see Figure 5 for comparison against drop tail.
- **fairness**
  - The authors' claim fairness is not well-defined.
  - {This is an obvious side-step of this issue.}
  - [later this becomes a **big deal -** see FRED paper.]

Evaluation of RED Design Goals

- **Appropriate for a wide range of environments**
  - Discussion is weak.
  - Shifts into parameter sensitivity discussion.

Conclusions

- **RED** is an effective mechanism for congestion avoidance at the router in cooperation with TCP.
- By controlling the calculated average queue size, **RED** provides an upper bound on the average delay at the gateway.
- claim:: The probability that **RED** chooses a particular connection to notify during congestion is roughly proportional to that connection's share of the bandwidth.

Future Work (circa 1993)

- Is **RED** really fair?
- How do we tune **RED**?
- Is there a way to optimize power?
- What happens with other versions of TCP?
- How does **RED** work when mixed with drop tail routers?
- How robust is **RED**?
- What happens when there are many flows?