

# **The Effects of Active Queue Management on Web Performance**



**From the dept of CS, UNC, Chapel Hill  
Sigcomm, Sep 2003**

**Offense : Santa**



# Preparation Outline

---

- Found out: Sigcomm 2003 best paper
  - Best Paper in Best Conference
- Tried to talk Andreas to trade places
  - Santa: Everyone does 1 offense and 1 defense
  - Andreas: Says who? Eugene assigned
  - Santa: Rice tradition. Both of you are new
- Started reading papers and googling



# I Digress...

---

- But, I have a point to make
- Sigcomm 2001: Outrageous Opinion Session
  - Tips from a Networking Insider: Stefan Savage



# Presentation Outline

---

- The major results
  - And my grudge
- My other grudges



# Major Result - I

---

- For offered loads up to 80% of bottleneck capacity, no AQM scheme provides better response times than simple drop-tail FIFO queue

# 80% load means uncongested

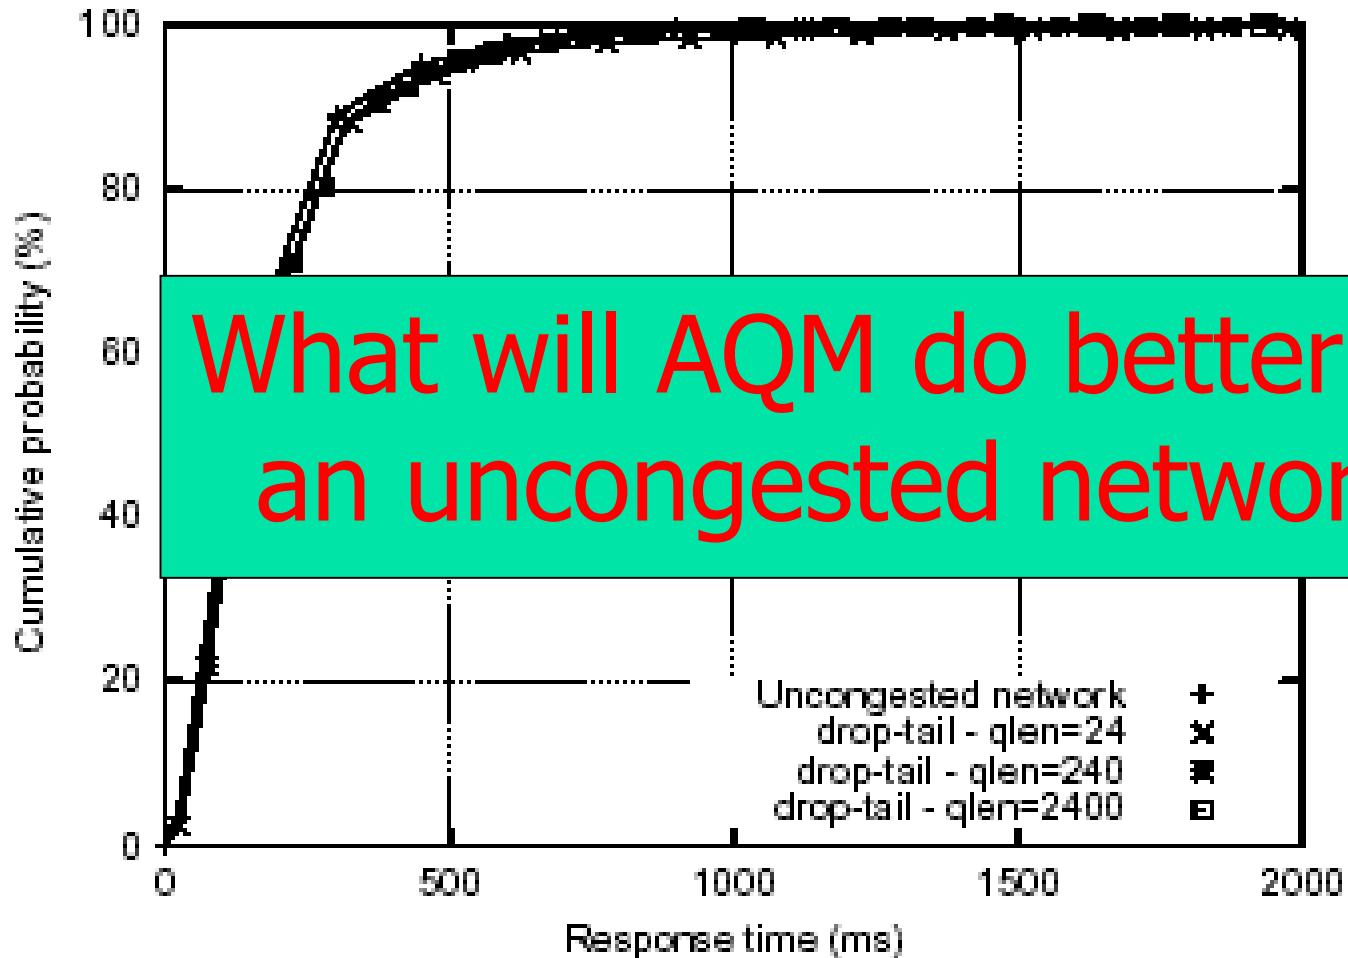


Figure 5: Drop-tail performance, 80% offered load.



# Major Result - II

---

- For loads of 90% link capacity or greater when ECN is not used, PI results in modest improvement over drop-tail and other AQM techniques

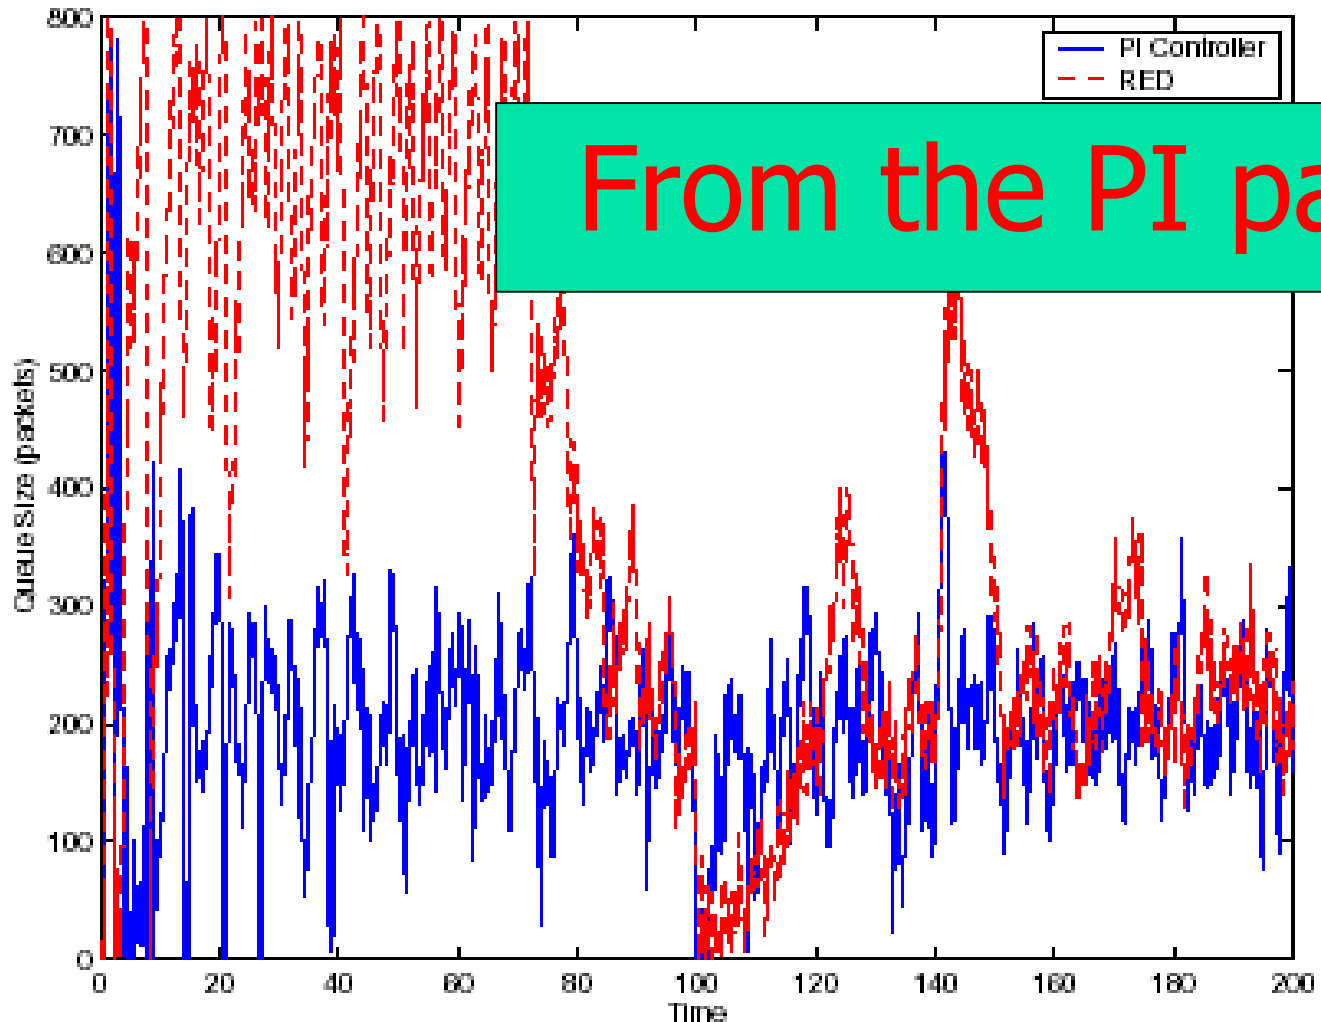


# PI is cool

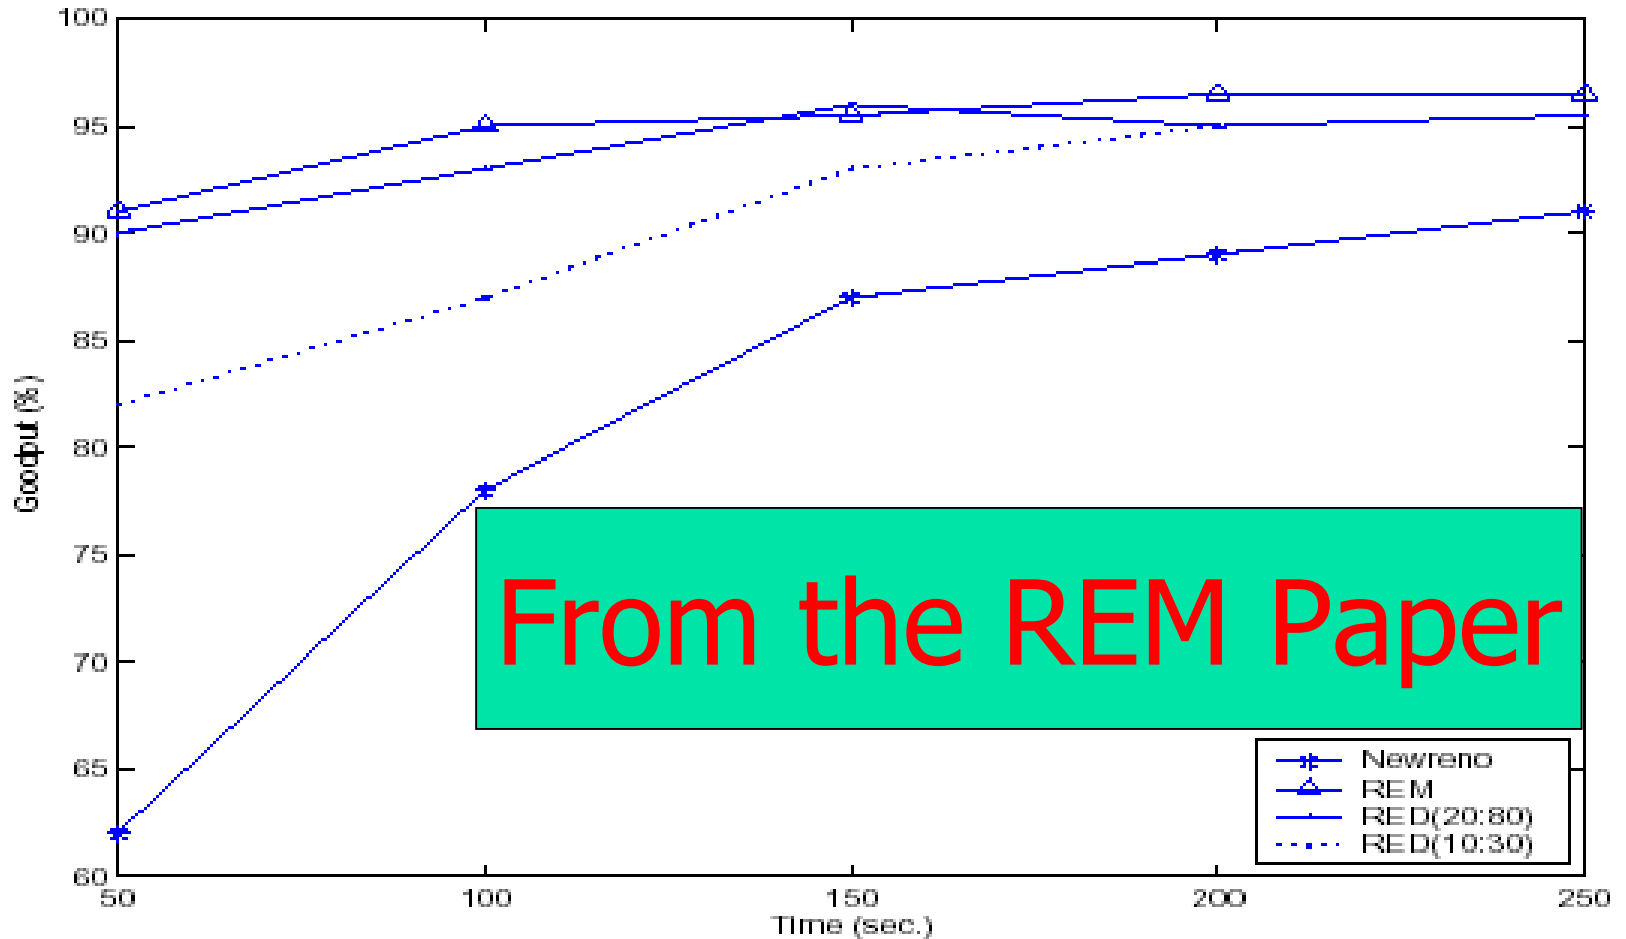
---

- This unfortunately was analyzed, simulated, implemented and shown in every way possible by the PI designers in Infocom, 2001
- **Why bother?**

# Comparing: PI with RED



# Comparing: REM with RED





# Major Result - III

---

- With ECN, both PI and REM provide significant response time improvement at offered loads above 90% link capacity



# Importance of ECN

## *A. The importance of ECN*

It is critical for the success of any AQM scheme that attempts to control the router queue that it be used in conjunction with ECN [17]. For instance, the PI controller can regulate queue length to a low level. This results in a lower delay than a corresponding drop-tail system. However, when dropping instead of marking packets, this may not result in more efficient performance, especially in the case of short lived flows.

**From the PI paper**



# Major Result - IV

---

- ARED with recommended parameter settings consistently resulted in the poorest response times



# Goals of AQM

---

- Primary goals
  - Controlled average queuing delay
  - Maintain high link utilization
- Secondary goals
  - Improving fairness
  - Reducing global synchronization
  - Accommodating transient congestion



# RED on Web Traffic

---

- M. Christiansen, K. Jeffay, D. Ott, and F.D. Smith, Tuning RED for Web Traffic, ACM SIGCOMM, August 2000.
  - "We conclude that for links carrying only web traffic, RED queue management appears to provide no clear advantage over tail-drop FIFO for end-user response times"
- Same as Major Result #4, albeit 3 years before

# A Recommendation

- RFC-2309: Recommendations on Queue Management and Congestion Avoidance in the Internet
  - Authors: B. Braden, D. Clark, J. Crowcroft, B. Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, L. Zhang
- Internet routers should implement some active queue management mechanism to manage queue lengths, reduce end-to-end latency, reduce packet dropping, and avoid lock-out phenomena within the Internet. The **default mechanism for managing queue lengths** to meet these goals in FIFO queues is **Random Early Detection (RED) [RED93]**. **Unless a developer has reasons to provide another equivalent mechanism, we recommend that RED be used.**



# RED Deployment

---

- RED is deployment in a lot of today's routers
  - Most simple and efficient scheme
- “Most current core routers are enabled with RED queue management algorithms”
  - Cisco Systems, “Technical specification from cisco, *random early detection on the cisco routers*”.
- Would Juniper & Cisco deploy a new technology without convincing proof of benefit?

# RED Parameter Setting

- May M., Bolot J., Diot C., and Lyles B., Reasons not to deploy RED, TR-June '99.
  - Parameter tuning in RED remains an inexact science."
- Floyd, S., RED: Discussions of Setting Parameters
- Showing bad performance of RED for some parameter setting does not prove anything
- More research may be needed

# RED parameters in this paper



- How did the authors come to the shown parameter setting if RED?
  - Admittedly an inexact science
- What about byte mode? Why did they not try it?
  - Were they just out to prove RED is BAD



# Presentation Outline

---

- Other grudges



# Evaluation Criterion


---

- “These conclusions are based on a premise that user-perceived response times are the primary yardstick of performance”
- Not the primary or secondary goals of RED

# Experimental Methodology

- 2 ISP Networks. 1 peering link

- A realistic topology



- Carries solely web traffic between sources and destinations on both sides
- Equally balanced in both directions



# Traffic Scenario

---

- Experiments using only HTTP Traffic model
  - Why not a realistic mix non-HTTP traffic
- Specially, as RED was previously shown to work not so well with web traffic



# Queue Sizes

---

- Viola... Some magic numbers
  - 24 & 240
- What about a range of numbers maybe?



# Conclusion

---

- Sigcomm decided to accept different sort of papers than it traditionally accepts
  - Savage jokes hurt too much
- What better way to prove than give best paper award to a “RED is bad” paper
- Paper is a well-written good comparison paper..... But, Best Paper Award?