



# Measuring Web Performance in the Wide Area

Paul Barford  
and  
Mark Crovella

Boston University

October 23, 1999

Wide Area Web Measurements



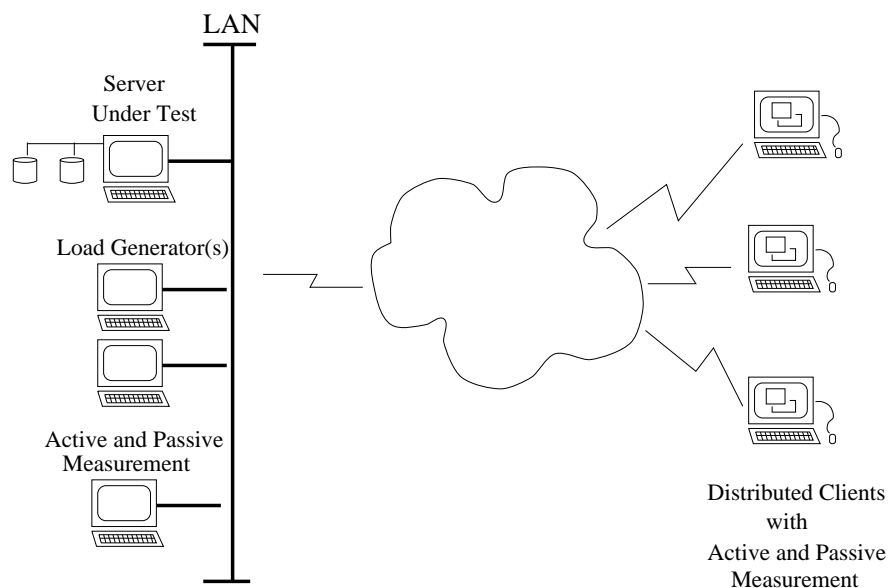
## Motivation and Introduction

- What are the root causes of long response times in the Web?
- Most research to date has focused on measuring servers or the network in isolation.
- The Wide Area Web Measurement (WAWM) project uses an infrastructure distributed across the Internet to study Web performance.
- Key ideas:
  1. *Integrated* server/network performance measurements.
  2. Both *active* and *passive* performance measurements.
  3. A *critical path* analysis of HTTP/TCP transactions.
- A prototype study has been completed.



## WAWM Infrastructure

- PC clients with GPS time cards are distributed across the Internet
- Software used consists of: Linux/Apache, MS WinNT/IIS, SURGE, tcpdump, Poisson Ping (Zing), Traceroute, xntpd, kernel performance monitors.




General Architecture of WAWM Hardware Components




# Critical Path Analysis

- Transfer latency is determined by a transaction's critical path (CP).
- CP is constructed by determining which data packets are liberated by each ACK.
- The CP profile enables the determination relative latency due to server, client, network or time-out's.
- tcpeval enables the automated construction of the critical path from tcpdump traces.



Entire file




Critical path

Time line diagrams for transfer of 500KB file



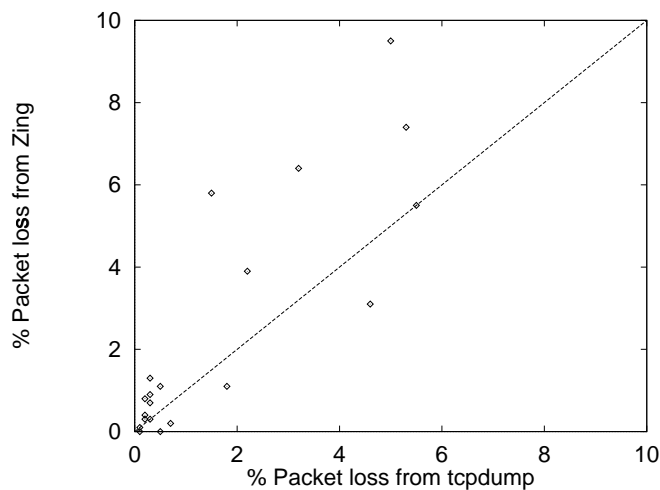
## Sample Preliminary Results

- **File Latency:** Latency is often lower when server is heavily loaded.



500KB file transfers latency under heavy network load

- **Active vs. passive measurement:** Poisson Ping (Zing) is not necessarily a good predictor of packet loss seen by TCP



Scatter plot of packet loss rate in TCP vs. Zing.