Comparison of Web Server Architectures: a Measurement Study

Enrico Gregori, IIT-CNR, enrico.gregori@iit.cnr.it
Joint work with
Marina Buzzi, Marco Conti and Davide Pagnin

Workshop
Qualità del Servizio nei Sistemi Geograficamente Distribuiti
9-10 Giugno 2004 Palazzo Mattei - Villa Celimontana, Roma

Motivation

- We started a study an geographically replicated WEB servers during the sabbatical year that Prof. Fabio Panzieri spent in CNR
- The Quality of Service perceived by users is the dominant factor for the success of an Internet-based web service and the target of geographical replication is to improve it by bringing the information closer to the user
- Geographical replication has significant impact on the server load and policies for geographical replication cannot be evaluated without a precise characterization of the web server behavior

Selected Previous papers

- Conti, M. Gregori, E. Panzieri, *F. Load Distribution among Replicated Web Servers: A QoS-based Approach*, Proc. Second ACM Workshop on Internet Server Performance (WISP'99), Atlanta, Georgia, May 1, 1999.

- Conti, M. Gregori, E. Panzieri, F. QoS-based Architecture for Geographically Replicated Web Servers, Cluster Computing Journal, 4, 2001, pp. 105-116.

- Conti, M. Gregori, E., Lapenna, W. “Content Delivery Policies in Replicated Web Services: Client-side vs Server-side, Cluster Computing Journal (to appear).

- Conti, M. Gregori, E., Lapenna, W. “Client-Side Content Delivery Policies in Replicated WebServices: Parallel Access vs. Single Server Approach”, Performance Evaluation Journal (to appear).

Motivation for Replicated WEB Servers

Performing data distribution structure closer to Internet structure

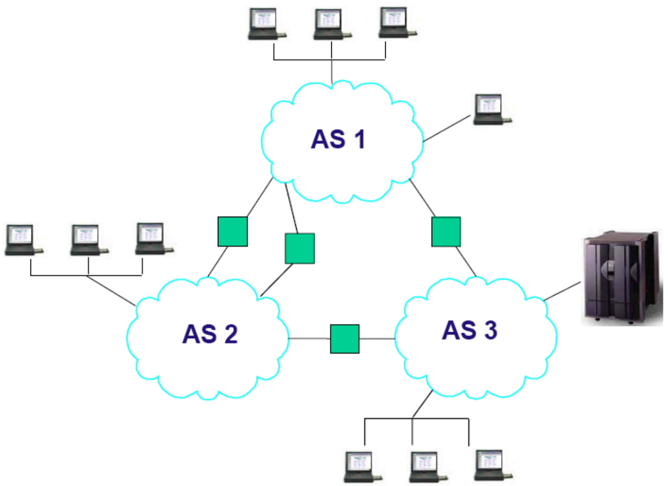

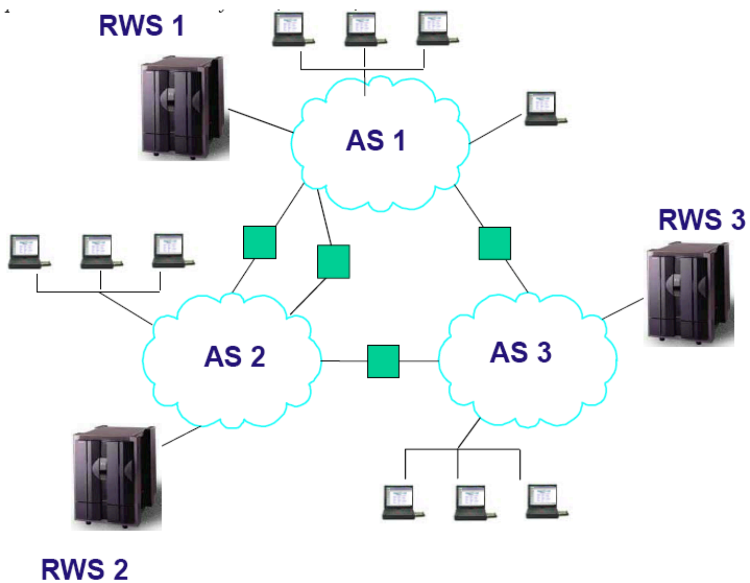

- Client side vs Server side

Objective of this study

- An accurate model of the web server behavior is essential for evaluating policies for replicated web server web servers
- We initially started using simple queueing models to characterize the WEB servers but we need an environment for measuring performance of real servers
- To understand the key elements in web server performance
  - creation of a controlled test-bed environment which allows to analyze web server performance in a simple and effective way
  - analysis of two different web server architectures and their performances, with particular focus on discovering bottlenecks

HTTP &TCP and Overload

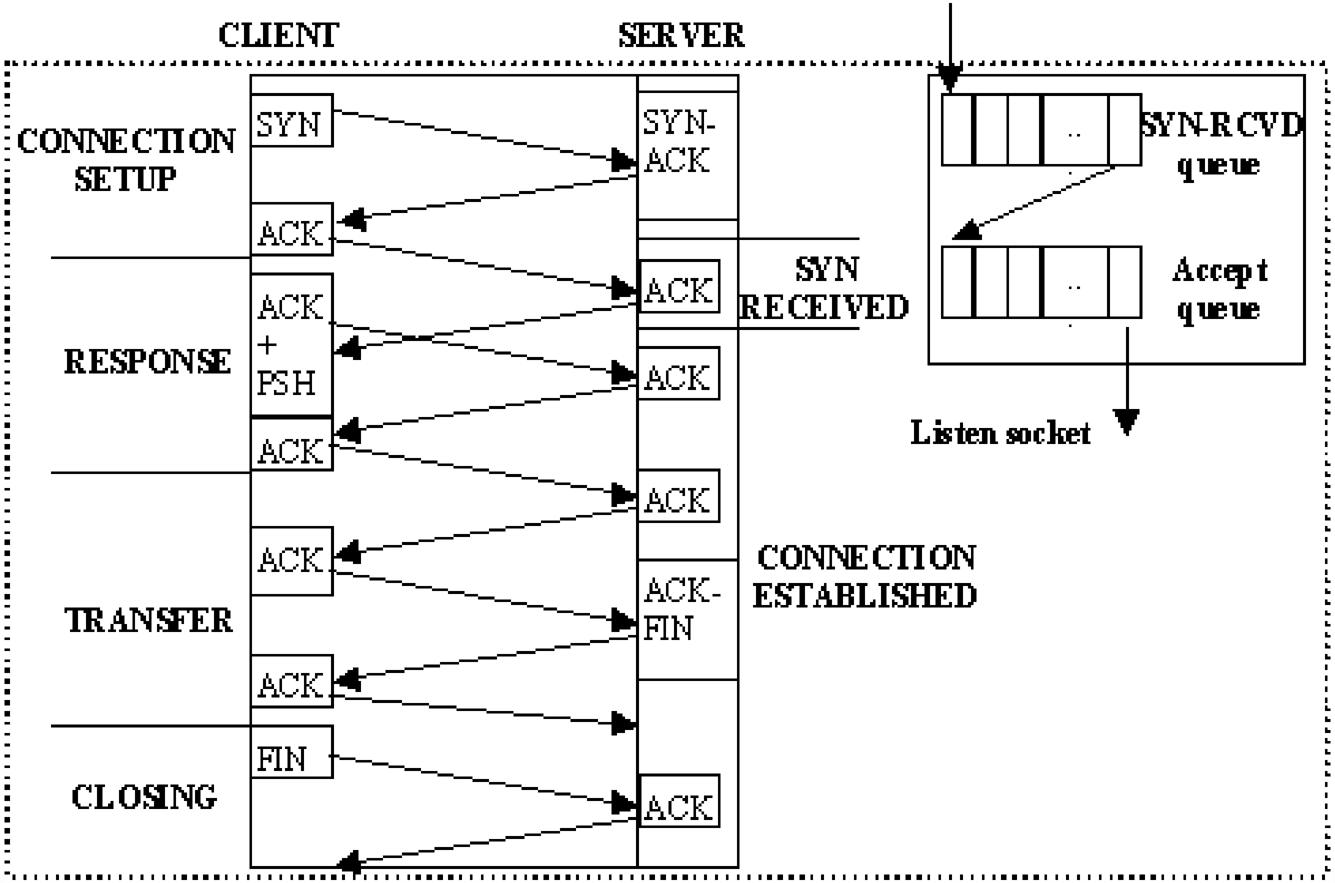

Measurement of a live system

- The approach of directly evaluating performances of a real web server suffers from difficulties related to
  - irreproducibility of real workloads
  - highly unpredictable behavior of the Internet
  - need for non-intrusive measurement of a live system

Measurement of a live system

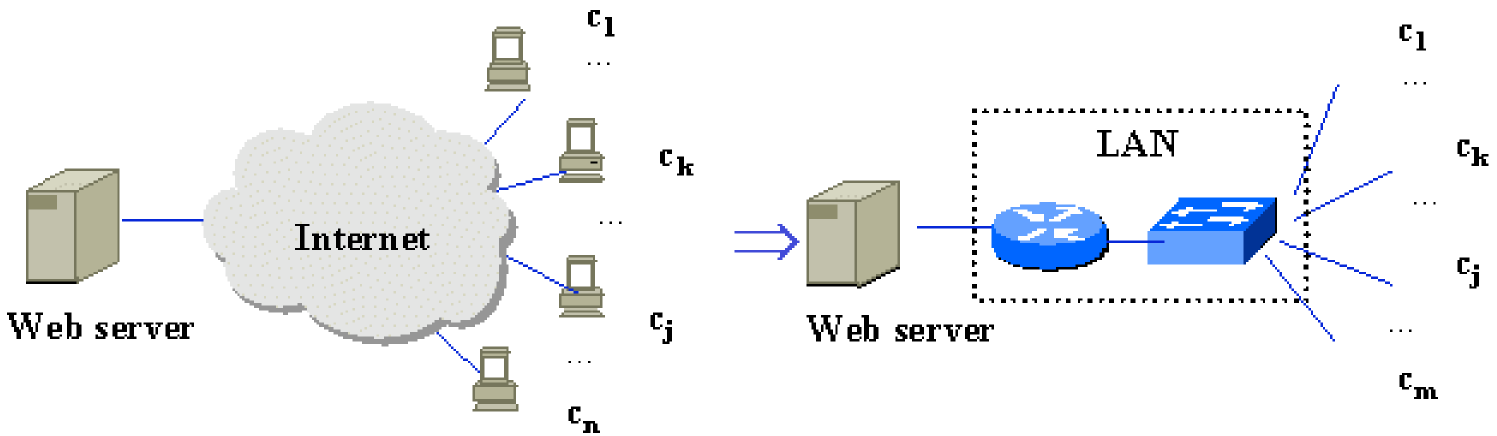

- The alternative is evaluation through generation of synthetic HTTP client traffic in a controlled environment requires great care as there several parameters limiting the performance.

Synthetic Workload Generation Tool

- Httperf
  - is capable of sustaining server overload
  - leaves a nearly total freedom as to the kind of workload and measurements to be performed
- Parameters
  - number of HTTP requests/sec
  - total number of requests to perform
  - users' think time
  - request timeout
  - ....

It is crucial to avoid exhausting the available sockets of the load generator.This means that starting from about 60,000 total available sockets and considering that the TCP TIME_WAIT status lasts 60 seconds in many TCP implementations, we have about 1,000 sockets available per second.

Furthermore, we have to consider the chosen request timeout and divide those 1,000 available sockets by the request timeout, and then subtract the mean number of established connections to get an idea of the limit of maximum number of requests per second that a single client machine can generate

Introducing WAN impact

- Dummynet for introducing WAN delays in the router machine
  - Flexible tool to modify delay and throughput, i.e., good modeling of the Internet bottleneck

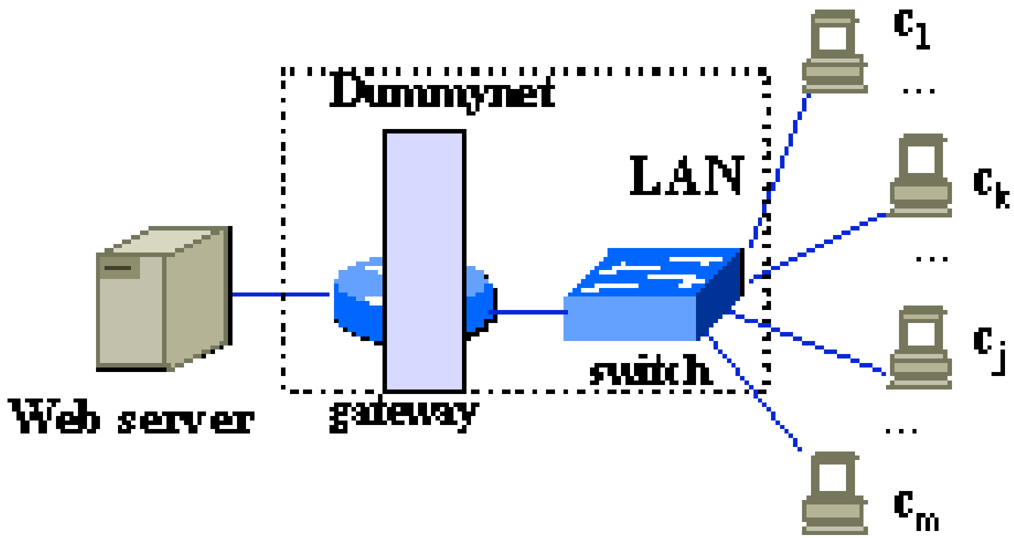

Web Server Architectures

- Single-process architecture: Boa 0.94.13
  - has a very efficient and lightweight implementation with very low memory requirements
- Multi-thread architecture: Apache2
  - the newest implementation of the most popular general-purpose web server

Experimental test-bed environment

- Web server
  - ASUS TX97E motherboard with Intel 430TX chipset and with 128Mb of memory, OS Linux Redhat 7.3
  - to evaluate the influence of the processor, we switched from a Pentium 133Mhz to a Pentium MMX 200Mhz

- Low speed processor are needed to be able to reach saturation condition

Experimental test-bed environment

- Clients
  - 5-10 clients running httperf: Linux Redhat 7.3, with 2.4.18-19.7.x kernel rpm version and no active firewall
  - HW ranged from K6 III at 450Mhz to Athlon XP 2000+
  - all systems were equipped with full duplex Fast-Ethernet network card and system memory between 256Mb and 512Mb

Experimental test-bed environment

- WAN simulator gateway: a dual Pentium III at 1GHz with 1Gb of memory and 2 Fast-Ethernet network adapters, with FreeBSD 4.7 release and dummynet option enabled
- Dummynet configured to simulate a population of 20% ADSL (128/256Kbit/sec) and 80% MODEM (56Kbit) connections with a RTT delay of 200msec

Performance indexes

- Average connection setup time
- Average response time
- Average transfer time

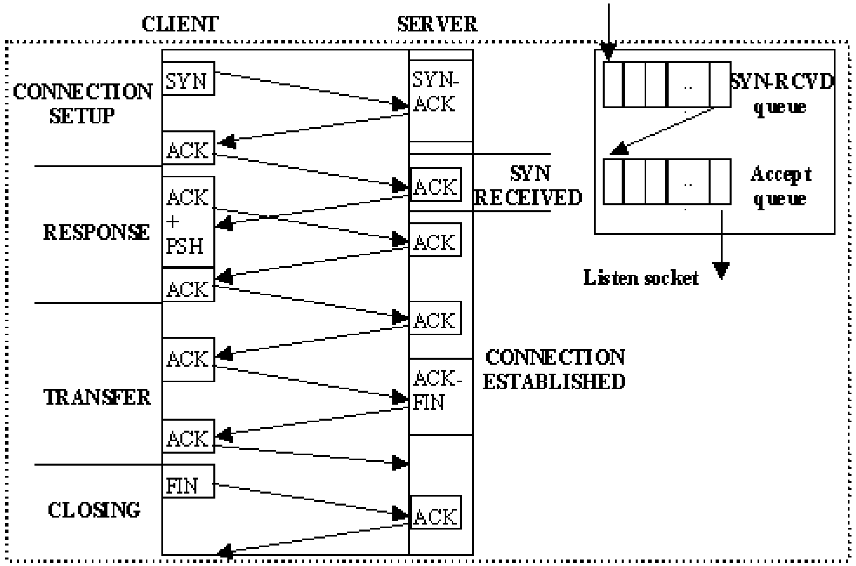

- Server throughput
  - calculated by summing up all the bytes transferred by a connection divided by the duration of the experiment (summed up for all the httperf instances)
- Completed requests per second

Maximum Server Throughput: Cached Files

- Boa server, Pentium MMX 200Mhz, 7K single file

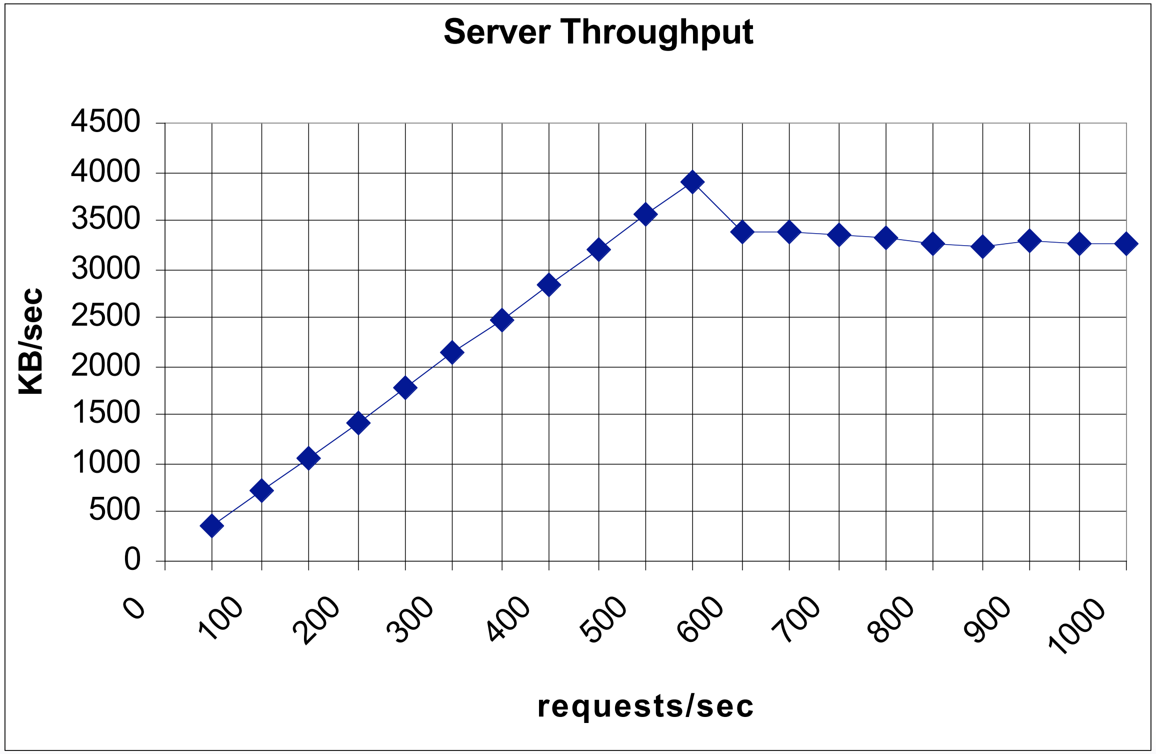

Maximum Server Throughput: Cached Files

- Boa server, Pentium MMX 200Mhz, 7K single file, Modem access

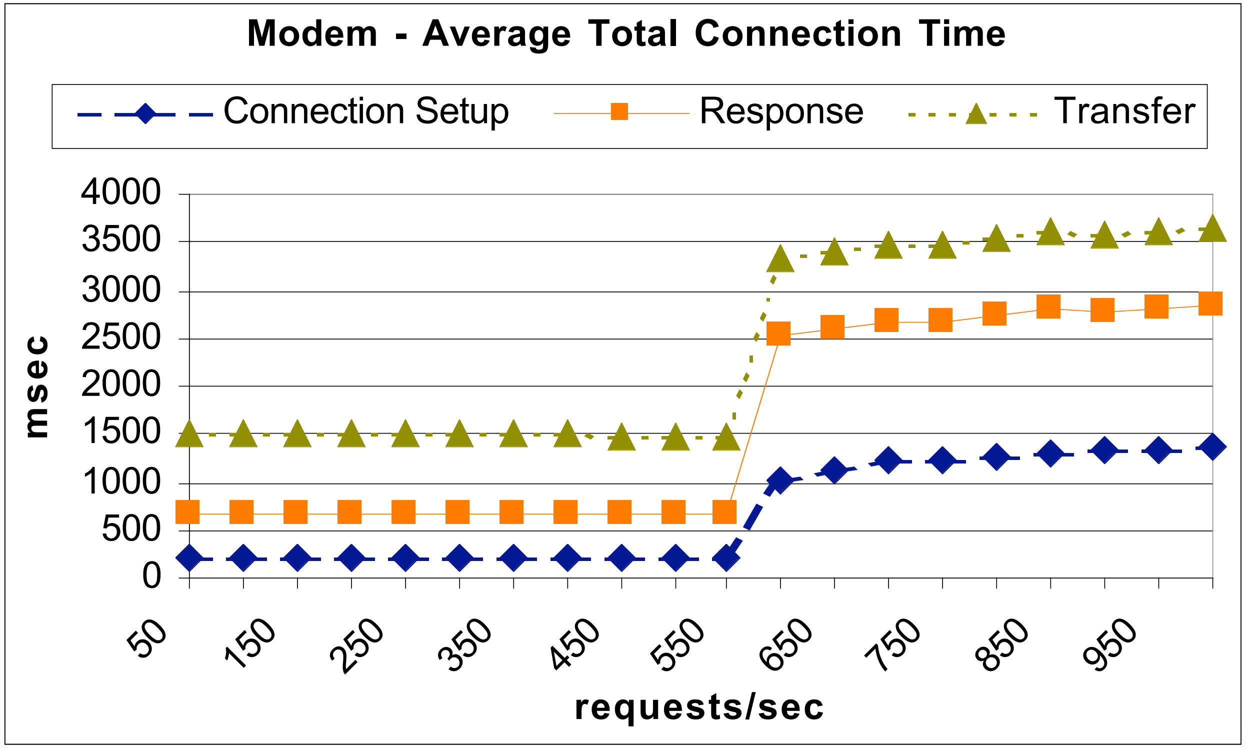

From Cached to Memory-mapped Files

- Boa server, single cached file vs 28K mmap-ed files

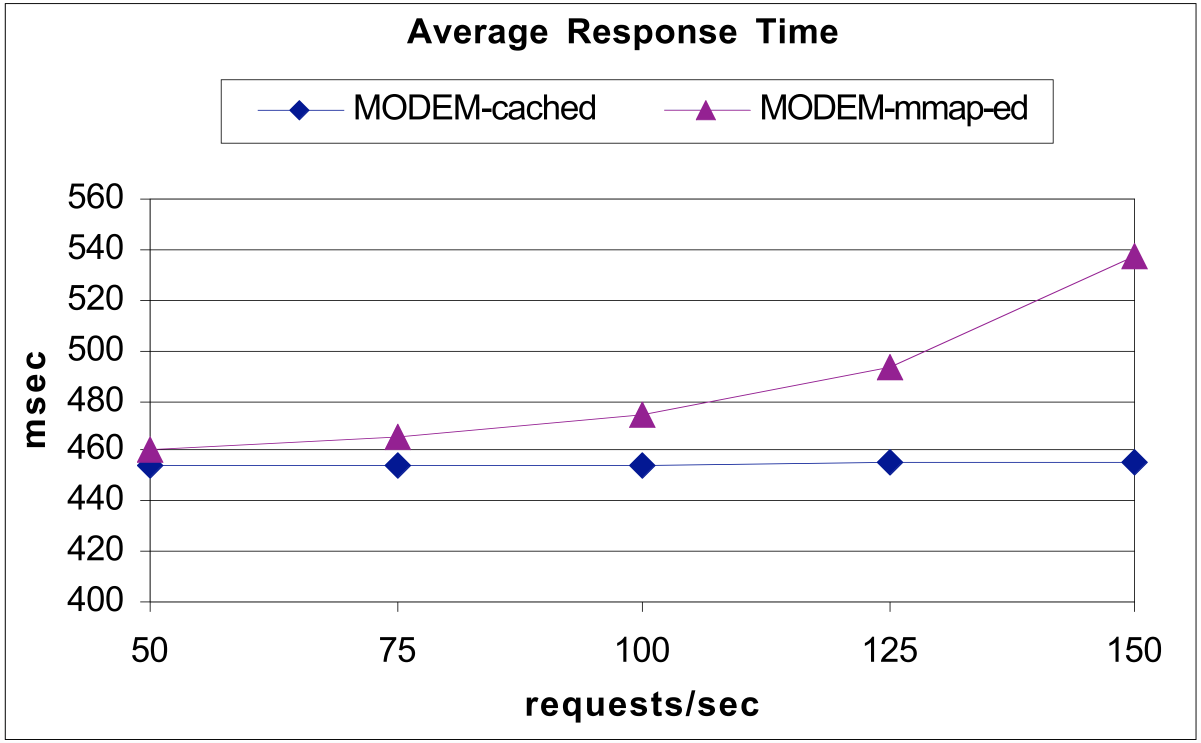

Server Architectures: Single-process versus Multi-threaded

- Pentium 133Mhz, 2500 different mmap-ed files

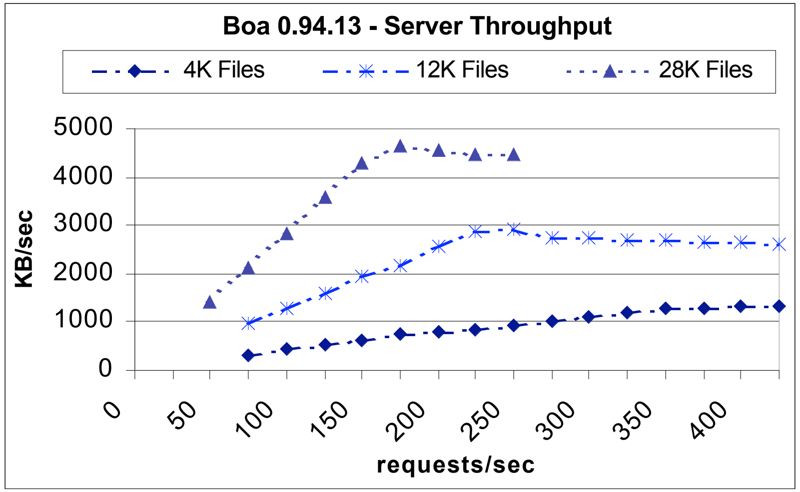

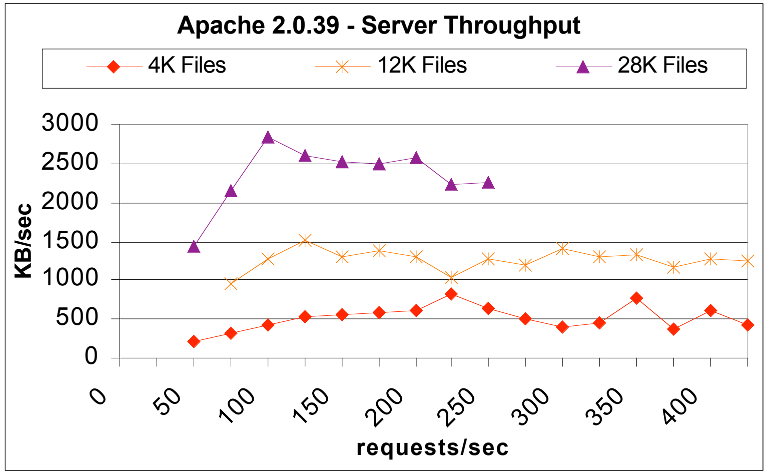

Impact of the Processor

- Files size 12K, Pentium 133 and 200MMX processors

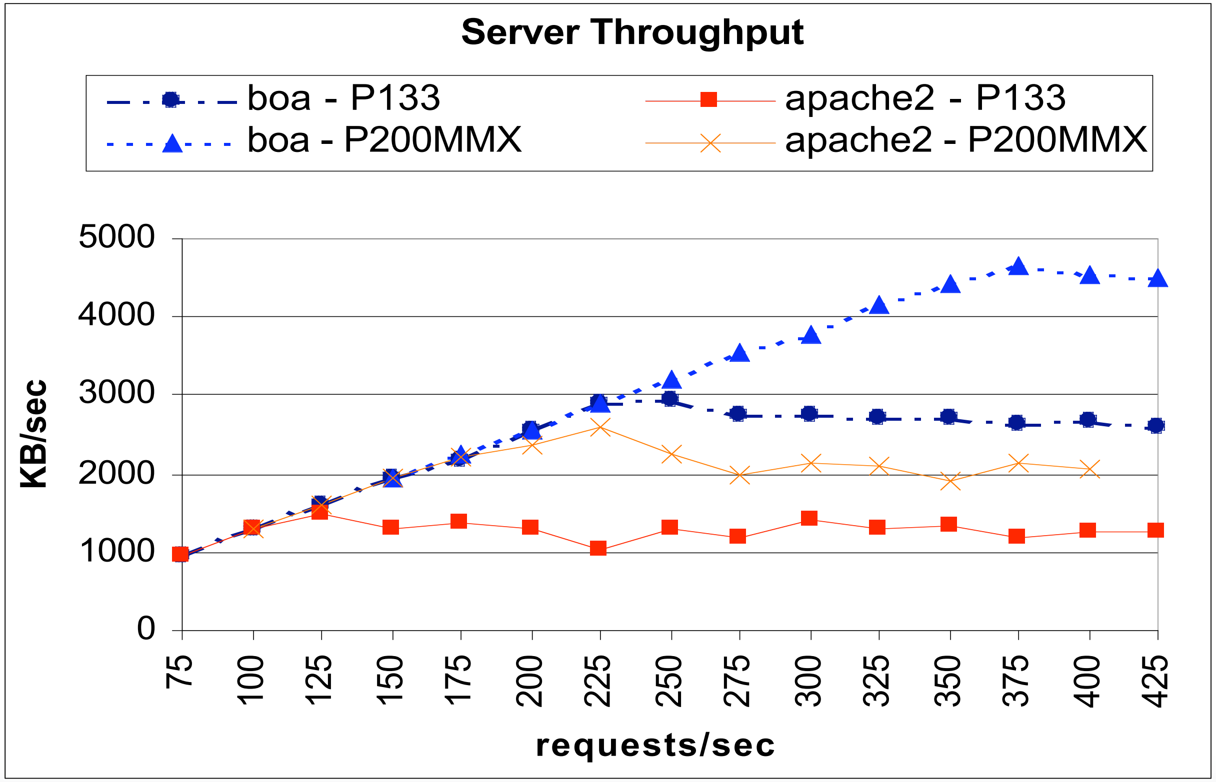

DISK Accessed Files

- Files size 12K, Pentium 133 and 200MMX processors

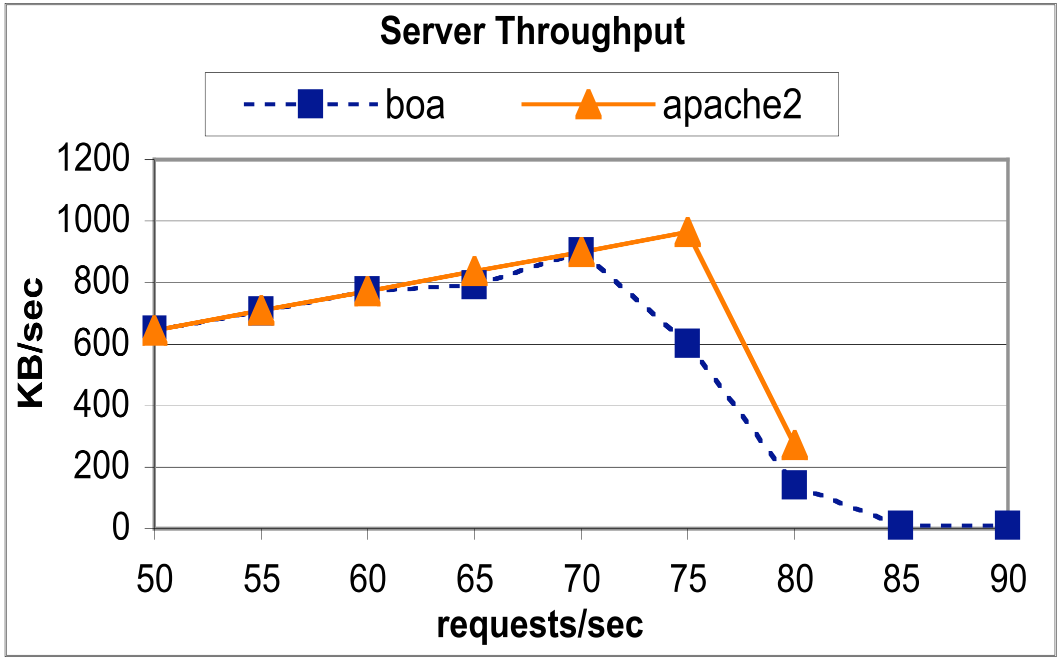

Summary

- Performance measurements of real systems with controllable network and load parameters are extremely important

- Preliminary results indicate:
  - The multi-thread architecture (i.e. Apache2) is CPU intensive but is less penalized by disk access
  - Disk access is the most limiting factor for the server throughput
  - User perceived QoS degrades greatly when the server is overloaded