

Specification and Implementation of Dynamic Web Site Benchmarks

Sameh Elnikety


Department of Computer Science

Rice University

Dynamic Content Is Common


1

THINKPAD A22M P3-1GHZ 30GB 128M



Our Price: \$1,740.95

Type: System
Platform: Pentium III
Mfg part#: 2628TTU



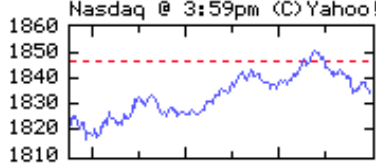
Quantity Remaining: 96

Qty [Buy now](#)

2

Market Summary [Edit](#)

Nasdaq @ 3:59pm (C) Yahoo!



12-Feb 10am 12pm 2pm

Dow	9863.74	-21.04 (-0.21%)
Nasdaq	1834.21	-12.45 (-0.67%)
S&P 500	1107.50	-4.44 (-0.40%)
10-Yr Bond	4.911%	+0.012
NYSE Volume	1,090,619,000	
Nasdaq Volume	1,621,515,000	

3

amazon.com. [VIEW CART](#) | [WISH LIST](#) | [YOUR ACCOUNT](#) | [HELP](#)

WELCOME [SAMEH'S STORE](#) | [BOOKS](#) | [ELECTRONICS](#) | [TOYS & GAMES](#) | [HEALTH & BEAUTY](#) | [COMPUTER & VIDEO GAMES](#) | [KITCHEN & HOUSEWARES](#) | [SEE MORE STORES](#)

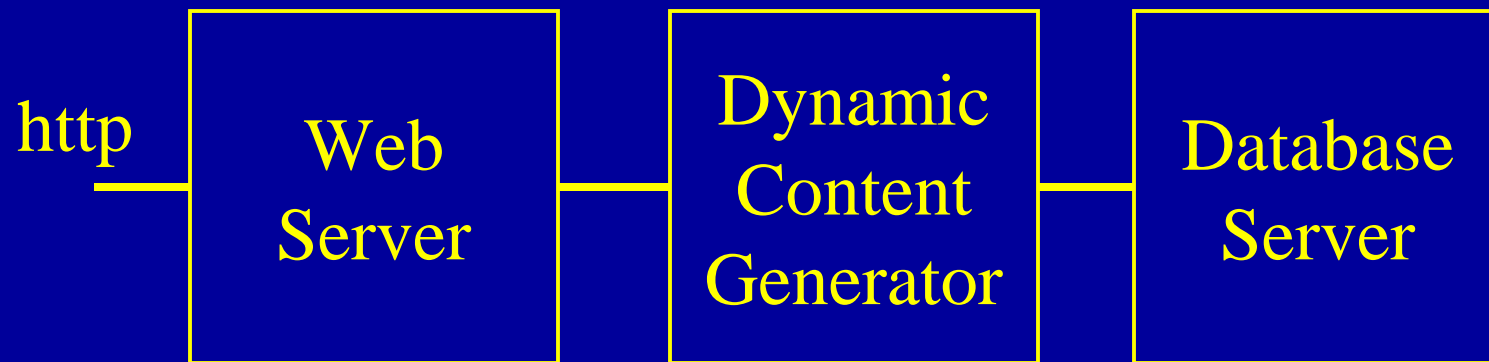
[INTERNATIONAL TOP SELLERS](#) | [TARGET](#) | [FRIENDS & FAVORITES](#) | [FREE E-CARDS](#)

Valentine's Day Roses for \$39.99, romantic books, and more
[Order flowers](#) [Find gift](#)

Hello, Sameh M. Elnikety. We have [recommendations](#) for you. (If you're not Sameh M. Elnikety, [click here](#).)

[Your Shopping Cart](#)

Generating Dynamic Content



DynaServer Project

- Scale Performance of Dynamic Content Sites
 - Caching
 - Clustering
 - Replication
 - Dynamic content delivery networks

Motivation

- We Need Workloads to Evaluate Different Architectures
- No Workloads for Dynamic Content

Benchmarks

- **Online Bookstore**
 - Models amazon.com
- **Auction Site**
 - Models ebay.com
- **Bulletin Board**
 - Models slashdot.org

Common Applications

**Stress Different Parts
of the System**

Contributions

- **Online Bookstore**
 - Followed TPC-W specification
 - Wrote implementation
- **Auction Site**
 - Developed specification & implementation
- **Bulletin Board**
 - Developed specification & implementation

Outline

- **Benchmarks**
 - Online Bookstore
 - Auction Site
 - Bulletin Board

- **Bottleneck Characterization**
 - Experimental Setup
 - Measurement
 - Performance Results

Online Bookstore Benchmark


- Models amazon.com
- Follows TPC-W (TPC Feb 2000)
- Activities
 - Browsing, buying, administrative updates
- Persistent Data
 - Images in Web server file system
 - All others in database

Online Bookstore Benchmark

Address <http://sp18.cs.ice.edu:8080/TPCW/ShoppingCart.php>

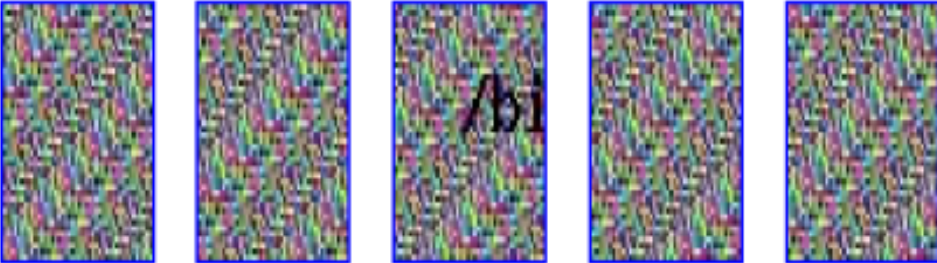
Search | Yahoo! Companion: Sign in | My Yahoo! | News | Entertainment | Sports | Shopping | Finance | Yahoo! | Yahoo! Mail

TPC Web Commerce Benchmark (TPC-W)

Image → 

Shopping Cart Page

Promotion → *Click on one of our latest books to find out more!*



Shopping Cart →

Qty	Product
<input type="text" value="1"/>	Title: <i>YEvbyKbF-bQOAh LMXc(OGaZhbMh,?tyX,/-</i> - Backing: PAPERBACK SRP. \$316.56 Your Price: \$208.9296

Subtotal price: 208.929600

Next Interaction → [Checkout](#) [Home](#)

Online Bookstore Benchmark

- 14 Interactions (as specified in TPC-W), e.g.
 - Home
 - Best sellers
 - Secure payment
 - Shopping cart
- Workload Mixes (as specified in TPC-W)
 - Browsing (95% read-only)
 - Shopping (80% read-only)
 - Ordering (50% read-only)

Auction Site Benchmark

- Models Ebay.com
- Activities
 - Browsing, selling, bidding
- User Sessions
 - Visitor, seller, buyer
- Persistent Data
 - Images in Web server file system
 - All others in database

Auction Site Benchmark

- 26 interactions
 - Browsing by category
 - Bidding
 - Buying
 - Selling
 - Leaving comments
 - User page
- Workload Mixes
 - Browsing (100% read-only)
 - Bidding (85% read-only)

Bulletin Board Benchmark

- Models Slashdot.org
- Activities
 - Discussion thread (story & comments)
 - Ratings
- User sessions
 - Visitor, moderator
- Persistent Data
 - Images in Web server file system
 - All others in database



Bulletin Board Benchmark

- 24 interactions
 - Browsing by category
 - Submit new story
 - Search story titles
 - Display story & filter comments
- No Full Text Search
- Workload Mixes
 - Browsing (100% read-only)
 - Submission (85% read-only)


Bottleneck Characterization

- Experimental Setup
- Measurement
- Performance Results

Software



Hardware



CPU	AMD Athlon 1.33GHz
Memory	768MB
Disk	60GB, 9msec, 5400rpm
Network	100Mbps Ethernet

Performance Metric


- **Throughput**
 - Number of interactions per minute

Resources Monitored

- **Web Server**
- **Database**

- CPU
- Memory
- Disk
- Network

Emulated Clients

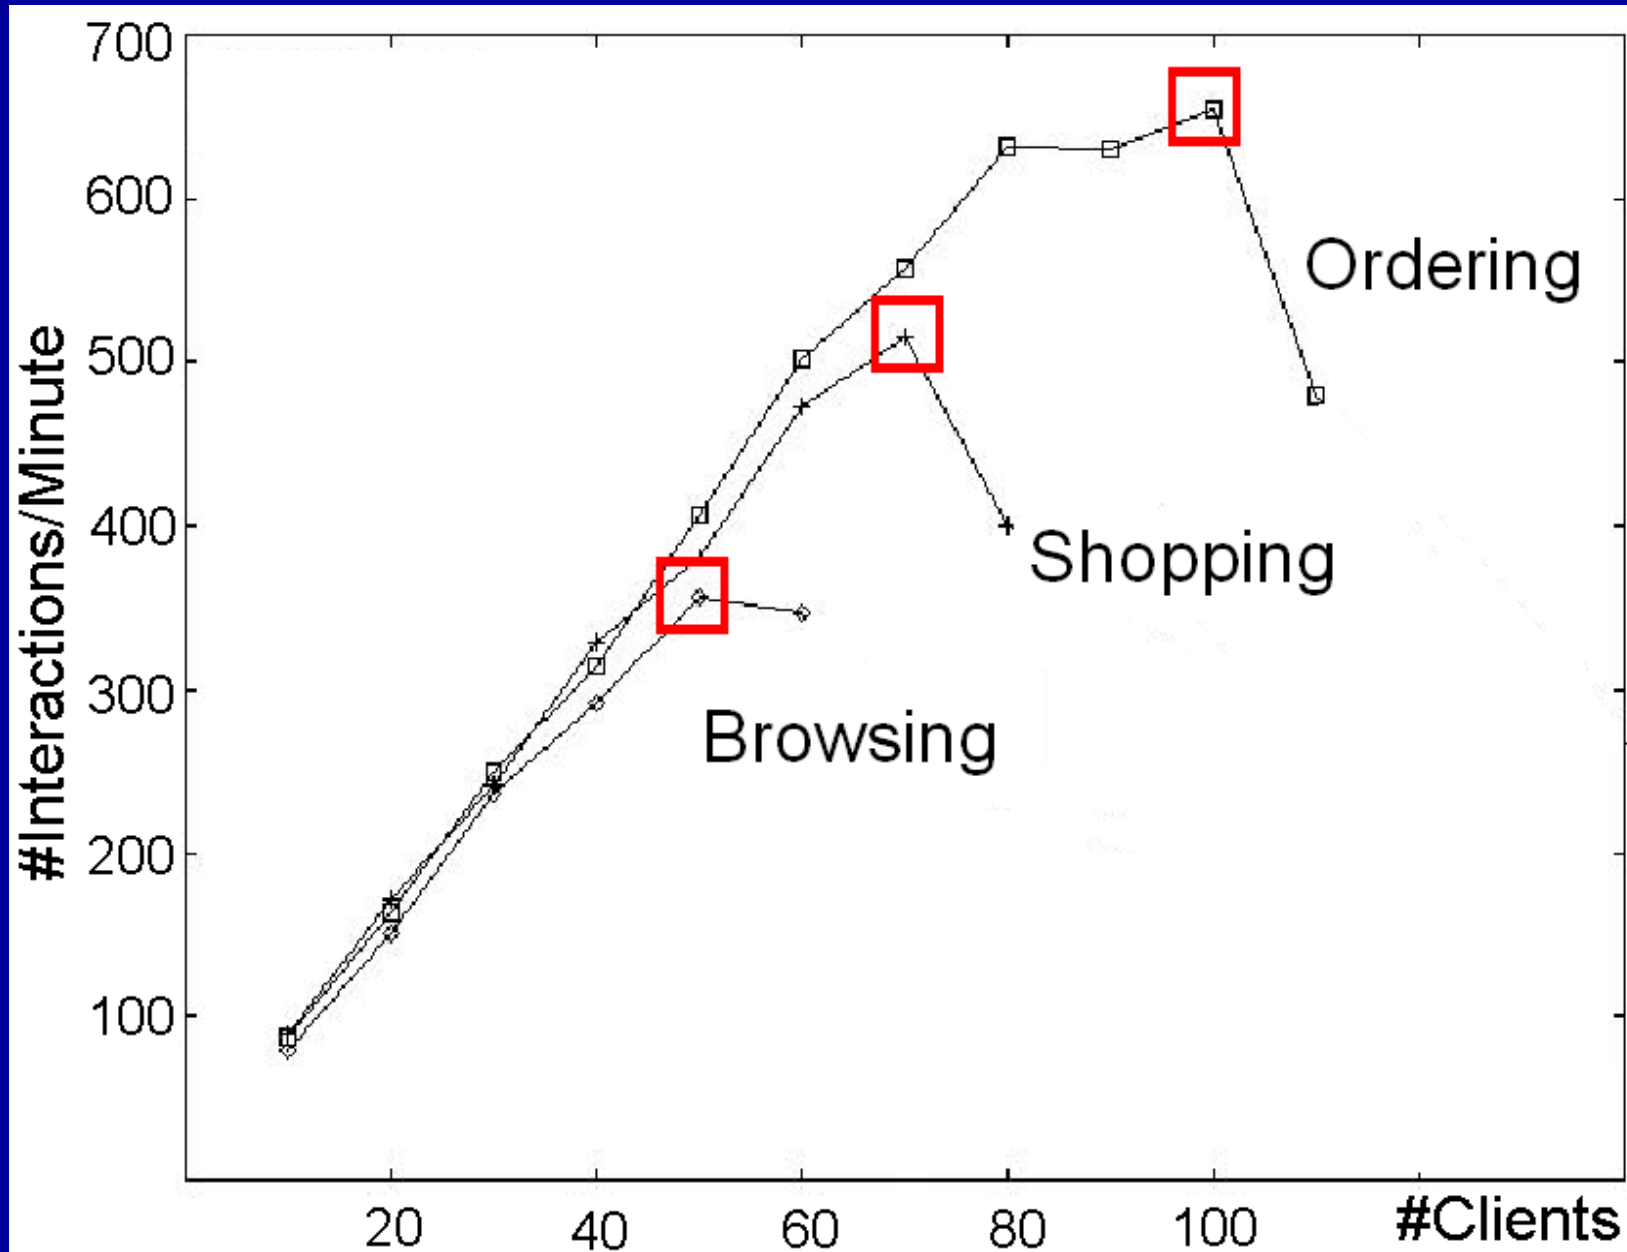


- Java Emulator
 - Session duration
 - Think time
 - Markov model

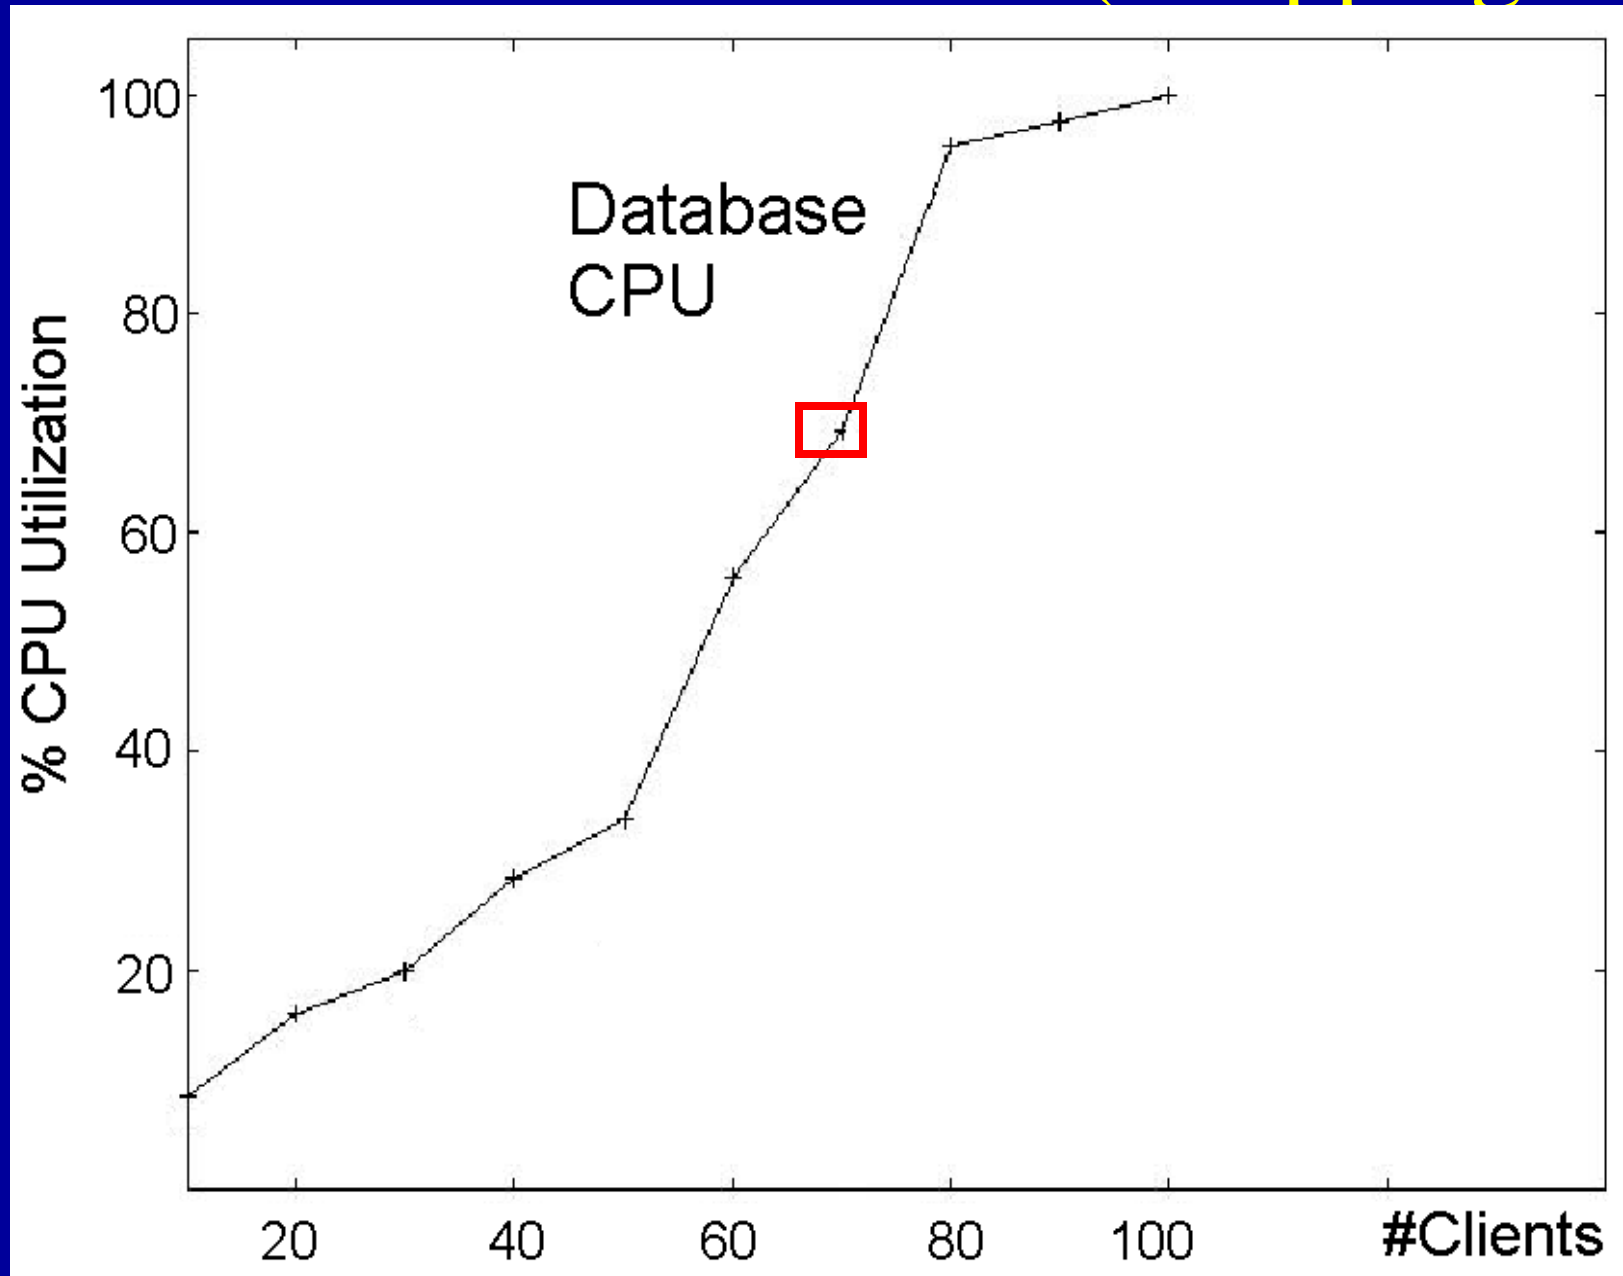
Measurement

- Driving the System
 - Several clients
 - Ramp-up time
 - Steady-state measurement

Online Bookstore Benchmark Throughput



Bookstore CPU Utilization (Shopping Mix)

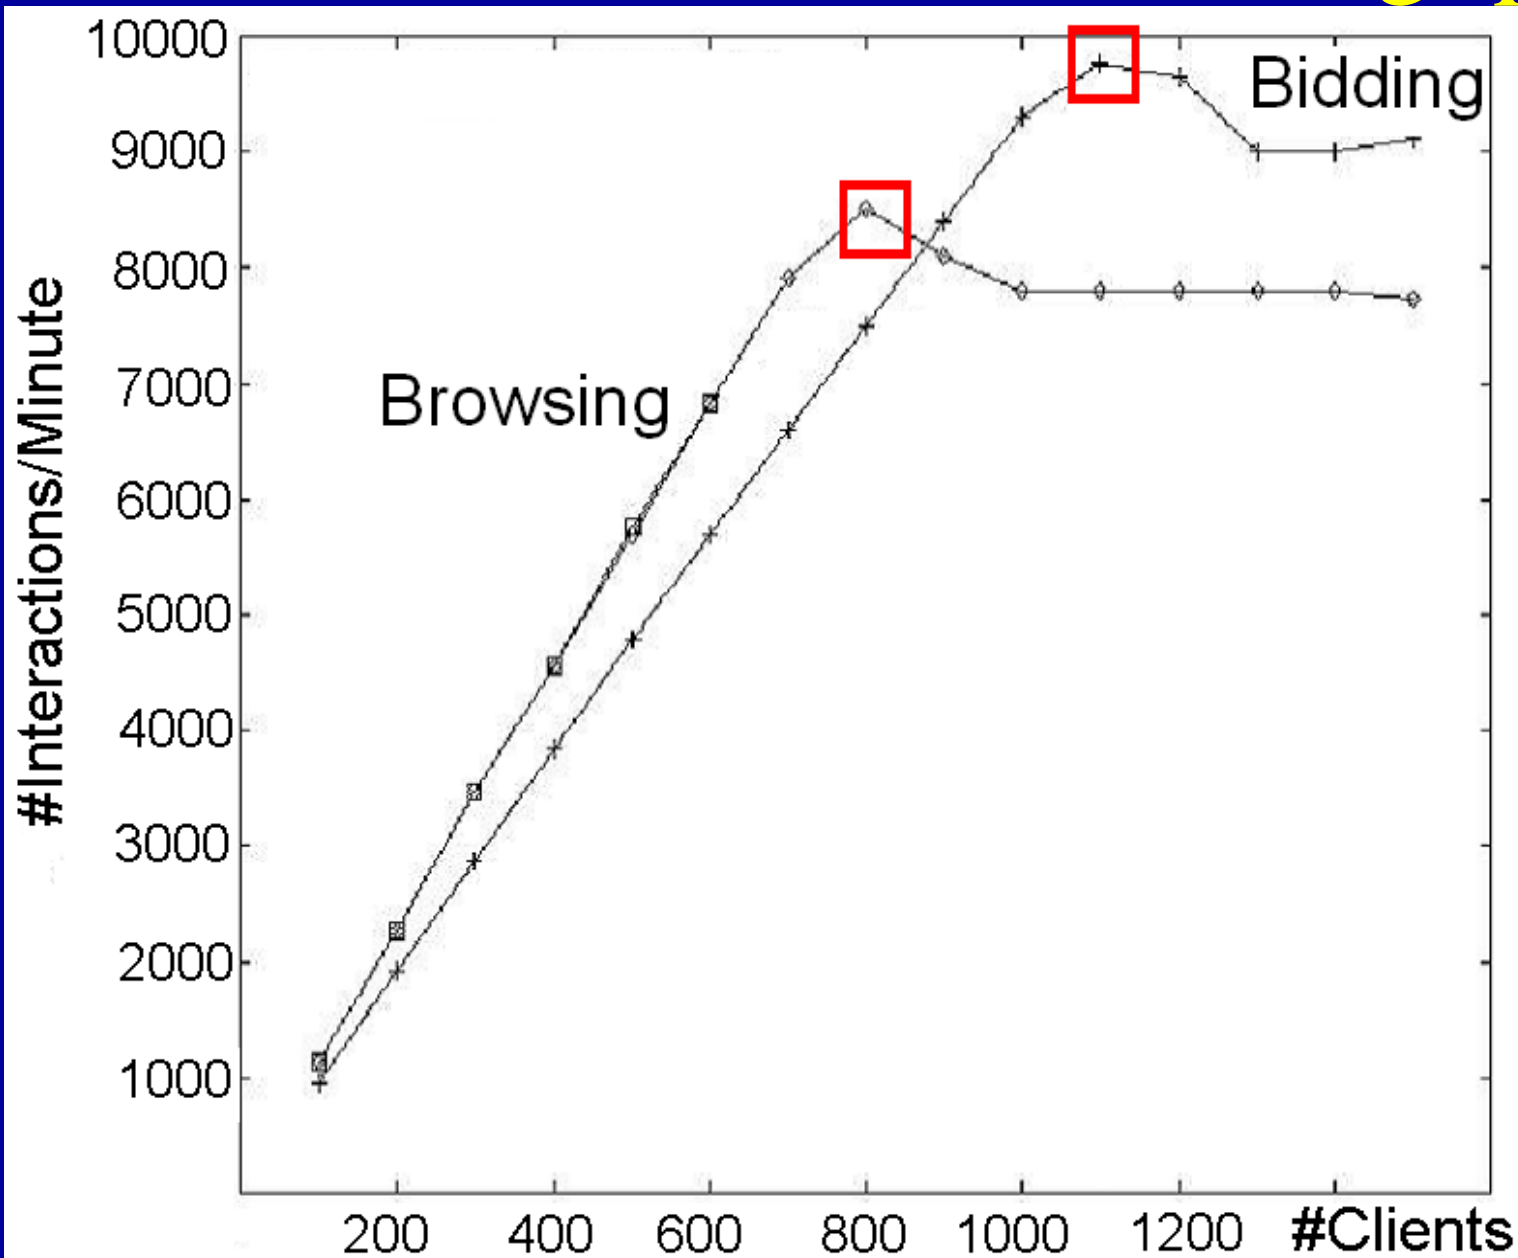


Bookstore Resources (Shopping Mix)

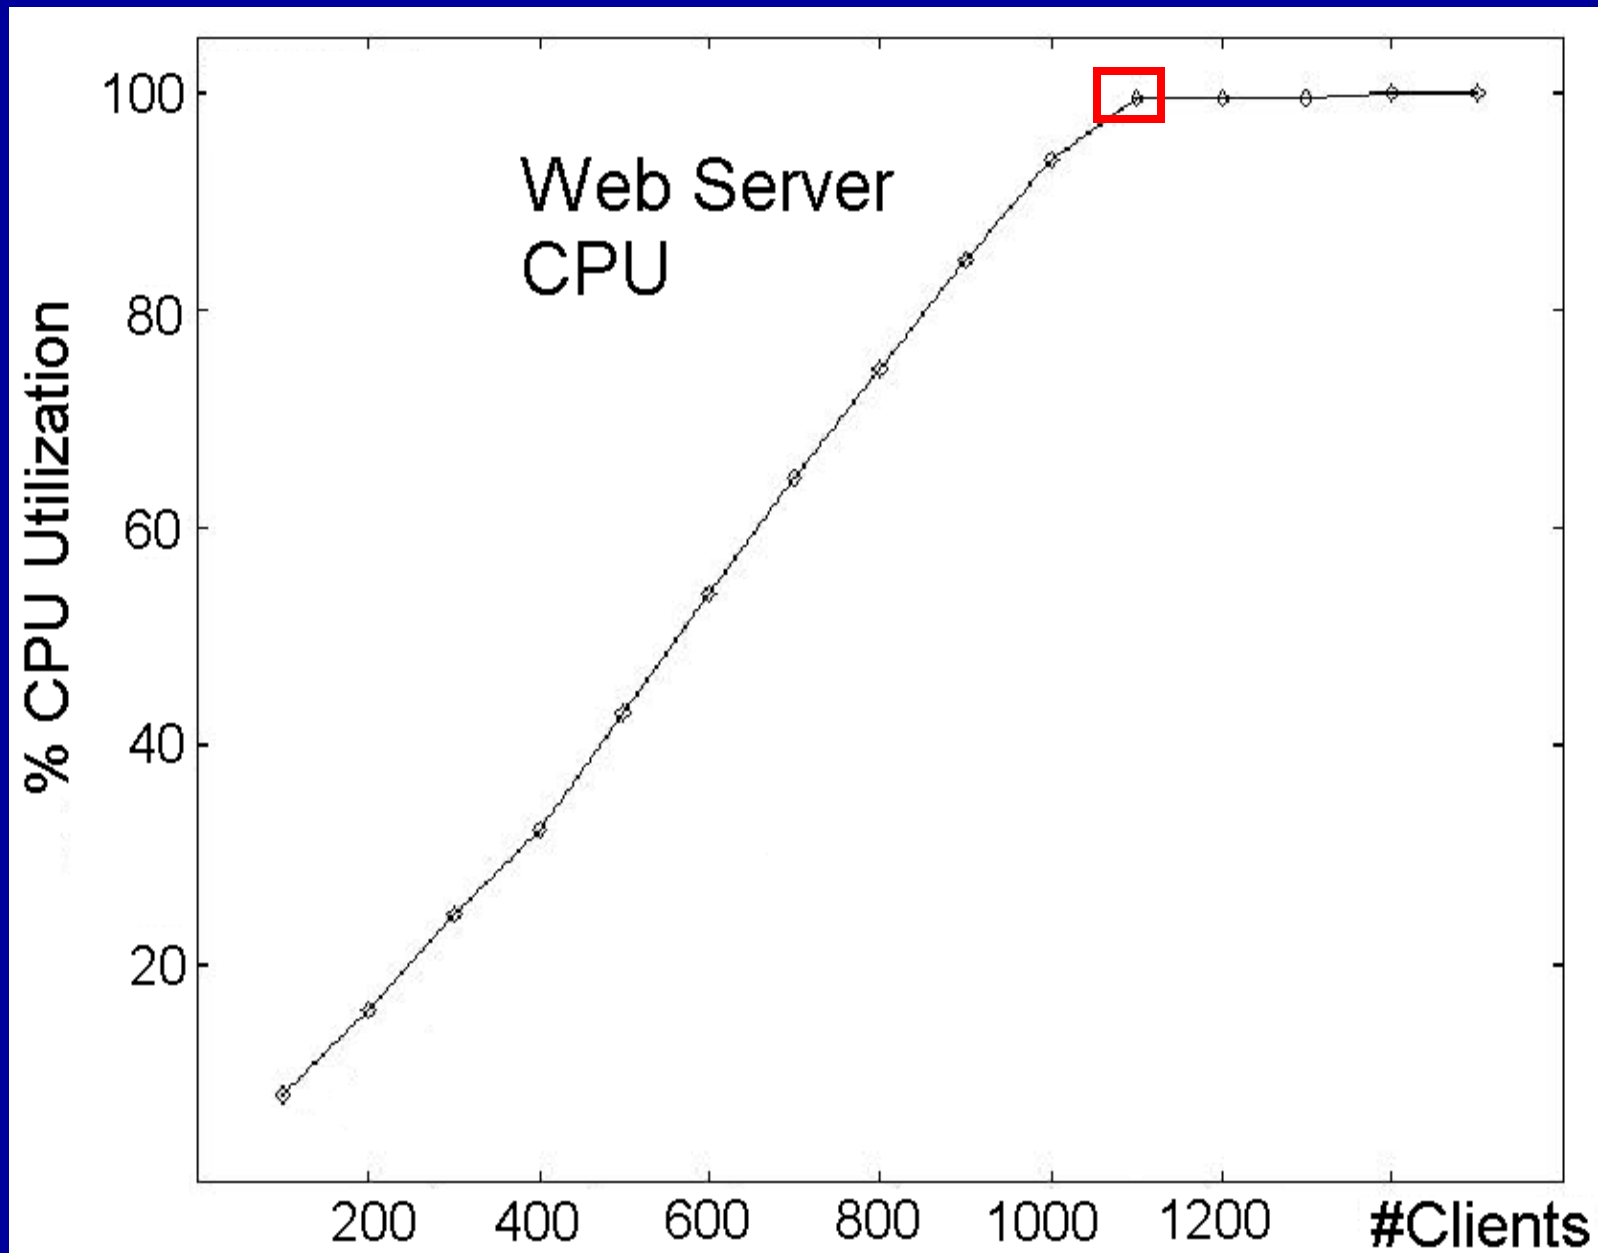
Resource	Web Server	Database
CPU	30%	71%
Memory	200 MB	390 MB
Disk	<20 block/s	<20 block/s
Network	3.2 + 0.6 Mb/s	0.6 Mb/s

Conclusion Database CPU is the bottleneck

Auction Site Benchmark Throughput



Auction CPU Utilization (Bidding Mix)



Auction Resources (Bidding Mix)

Resource	Web Server	Database
CPU	100%	58%
Memory	70 MB	250 MB
Disk	<20 block/s	<20 block/s
Network	55 + 2 Mb/s	2 Mb/s

Conclusion Web Server CPU is the bottleneck

Bottlenecks

- Online Bookstore
 - CPU of Database
 - Models amazon.com
- Bidding
 - CPU of Web Server
 - Models ebay.com
- Bulletin Boards
 - CPU of Web Server
 - Models slashdot.org

Conclusions

- 3 Benchmarks
 - Online bookstore, auction site, bulletin board
 - Open-source software
 - Available implementations (PHP, Servlets, EJB)
- Capacity Planning
- Research on Dynamic Content

Collaborators

**Cristiana Amza, Anupam Chanda, Alan Cox,
Romer Gil, Willy Zwaenepoel**

Department of Computer Science - Rice University

Emmanuel Cecchet, Julie Marguerite

INRIA Rhône - Alpes

Karthick Rajamani

IBM Austin Research Lab

Download the Benchmarks

Benchmark Implementations in

- PHP
- Servlets
- EJB

<http://www.cs.rice.edu/CS/Systems/DynaServer/>