



# Measuring the Performance of an XML-Based Communication Primitive

Gregory M. Kapfhammer  
Department of Computer Science  
Allegheny College

<http://cs.allegheny.edu/~gkapfham/>

# A Unique Invocation

*The London Times asks*


“What’s Wrong with the World?”

Dear Sirs,

I am.


Sincerely yours,  
G. K. Chesterton

# Storage and Communication Primitives



- How does object encoding impact performance?
- **Contribution:** A benchmarking framework to compare the performance of sockets and XML-RPC

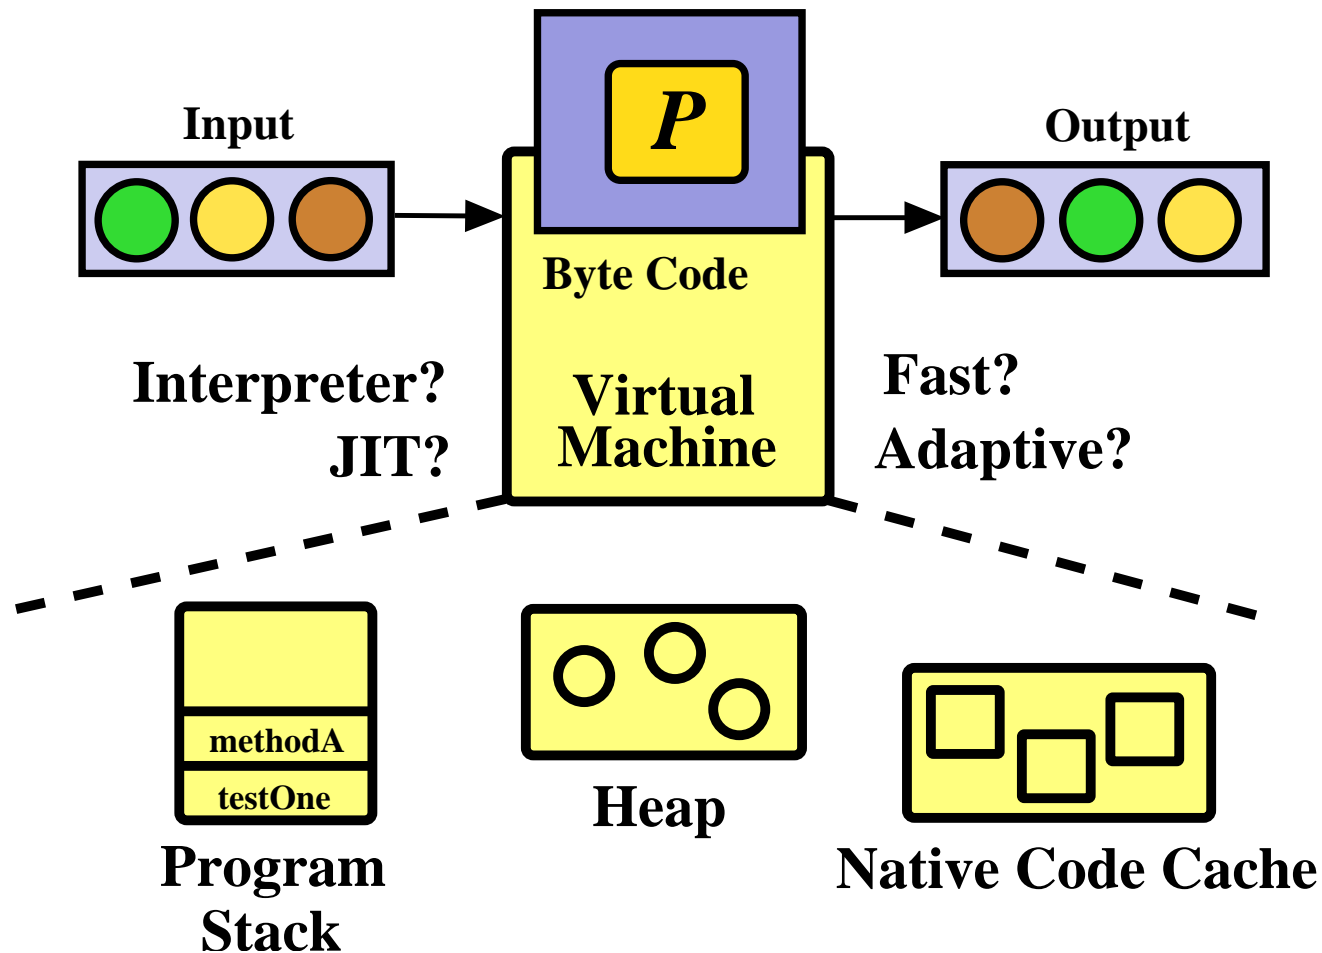
# Remote Communication and OpenDHT



## OpenDHT & PlanetLab

- Clients can `put` and `get` with Sun RPC or XML-RPC
- Does the communication primitive impact performance?
- How do we measure performance and/or correctness?

# Program Execution with a JVM



→ JVM implementation and configuration impacts performance

# Micro Benchmarks

Experiment	Sent by client	Received by client
SS	Single primitive	Single primitive
SV	Single primitive	Vector
VS	Vector	Single primitive
VV	Vector	Vector

- Use benchmarks similar to those proposed by Allman et al.
- Implement the benchmarks in the Java language
- *ExperimentCampaign* framework uses Perl and Mathematica

# Micro Benchmarks II

Experiment	Sent by client	Received by client
FIND (SS)	Single primitive	Single primitive
FACT (SV)	Single primitive	Vector
GCD (VS)	Vector	Single primitive
REV (VV)	Vector	Vector

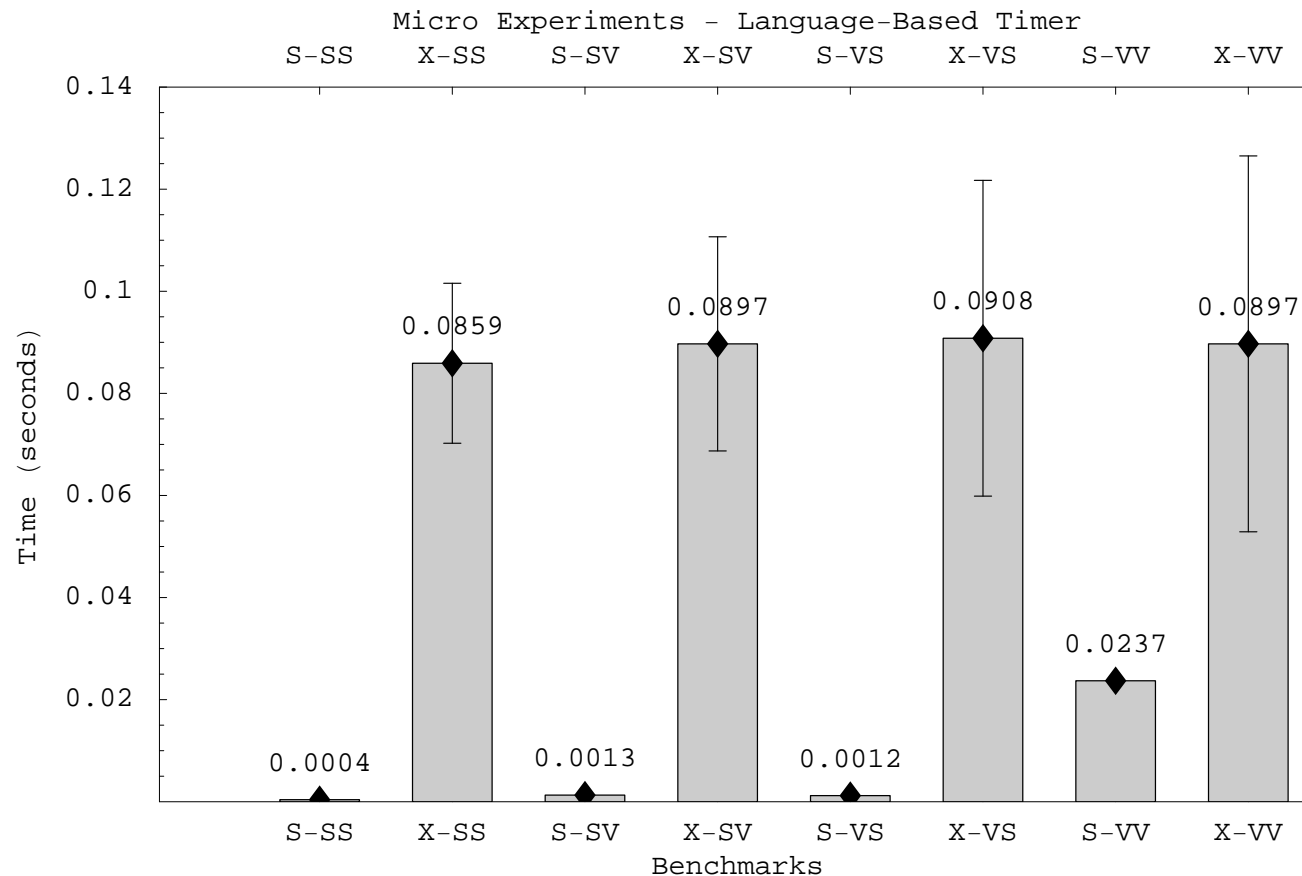
- Benchmarks use sockets and Apache XML-RPC
- Benchmarks perform a simple computation on the server
- Configure the client and server to execute on same node

# Experiment Design

- Select Java 1.5.0, GNU/Linux with kernel 2.6.12, 3 GHz P4, 1 GB main memory, 1 MB L1 Cache, CPU hyperthreading
- Use operating system and language-based timers to calculate  $R(B, P)$ ,  $R_{\Delta}(B, P, P')$ , and  $R_{\Delta}^{\%}(B, P, P')$
- Replace the socket communication primitive with XML-RPC
- Execute ten trials and calculate arithmetic means, standard deviations, and confidence intervals
- Formulate the null hypothesis as  $H_0 : \mu_{R(B,P)} = \mu_{R(B,P')}$
- Use the Welch's approximate t-test with  $\alpha = .01$

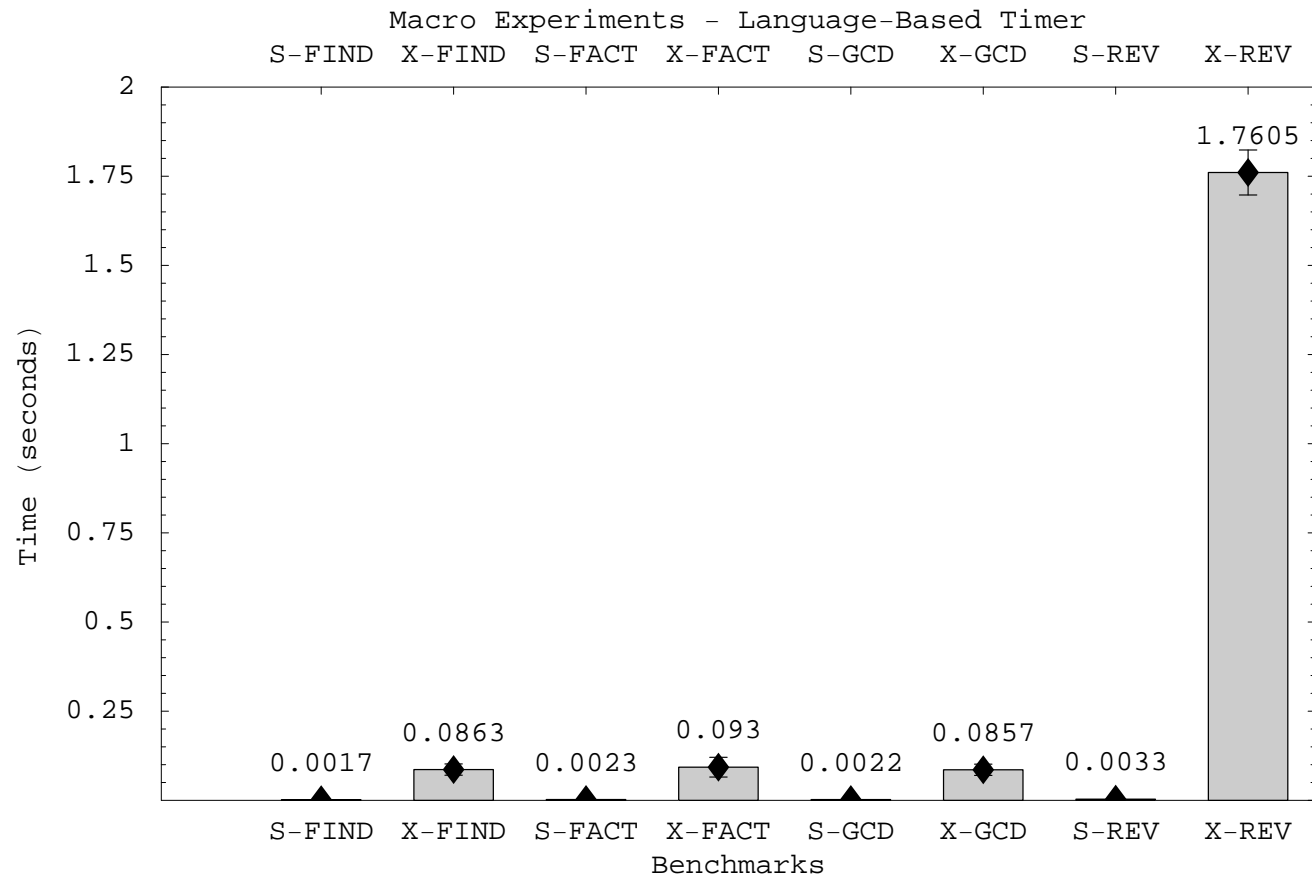


# Micro Benchmark I



→ XML-RPC shows greater response time with more dispersion

# Micro Benchmark II



→ X-REV exhibits high response time due to string parsing

# Using Very Large Vectors

$size(V)$	$size(V)$ (bytes)	$R(VV, S)$ (sec)	$R(VV, X)$ (sec)
5000	80,520	0.298	0.347
10000	161,000	0.598	0.523
50000	927,720	18.784	1.697

- At smaller vector sizes sockets demonstrate slightly better response times
- XML-RPC shows better response time when  $size(V) = 50000$  : *why?*

# Explanatory Power of GC

<i>size(V)</i>	YGC (count)	YGC (sec)	FGC (count)	FGC (sec)
5000	16	.008	0	0
10000	63	.023	4	.050
<b>50000</b>	<b>1645</b>	<b>.697</b>	<b>663</b>	<b>10.375</b>

<i>size(V)</i>	YGC (count)	YGC (sec)	FGC (count)	FGC (sec)
5000	14	.016	0	0
10000	27	.022	1	.020
<b>50000</b>	<b>123</b>	<b>.695</b>	<b>5</b>	<b>.143</b>

→ Varying the heap size of socket configuration yields similar results

# GC Allocation Rate

- S-VV allocates 710,374,184 bytes and X-VV only allocates 54,101,312 bytes
- At benchmark termination, S-VV has 4,773,224 bytes and X-VV has 7,234,520 bytes of live objects
- Sockets use `char[]` and XML-RPC uses `java.nio.CharBuffer`
- Can we use past GC behavior to predict future program performance?

# Conclusions

- A suite of micro benchmarks to measure the performance of communication primitives
- A comparison of sockets and XML-RPC that we can extend to other primitives
- Experiments reveal a trade-off in the performance of the two primitives
- Extend the study to new primitives and JVMs
- Focus on remote communication, long running benchmarks, and the measurement of throughput
- **What are your suggestions?**

# An Invitation to Participate



→ I value your comments, suggestions, and participation!

<http://cs.alleghey.edu/~gkapfham/research/>