



PERFORMANCE

YOU WANT IT WHEN?

Performance Benchmarking

- Jigsaw + Clojure
- Performance Benchmarking
 - Everyone needs it
 - Peer Review (getting what we expect)
 - OpenJDK performance 'gate'
- Working workshop



e 2-5, 2009

About Tom

- MSEE (I like hardware hacking too :)
- Worked at Sun 1997-2008
 - Chief Technologist for the dot-com sales area
 - Java Performance Engineer
 - Re-licensed Java for Debian non-free (DLJ)
 - OpenJDK Ambassador (core team)
- Founder in automotive electronics startup
- Independent Consultant



Update on Jigsaw in Debian

- Both take interdependency mgmt seriously
- Discussed at DebConf10
- Lintian improvements
- GSoC project, 2011
 - Detour into jtharness, jtreg
 - Understanding, merge IcedTea + Debian OpenJDK delta patches
 - Thinking of how best to handle the module library
 - Can't wait for build refactoring!

Thoughts about Clojure on Jigsaw

- STM / concurrency now, but we could do better
 - “par” branch (David Edgar Liebke)
- Better Startup
- Clojure library modularity
- New Features
 - Get Fork/Join
 - Tail Calls?
 - Invokedynamic?
- How *fast* could it be, could it beat Scala?



Recent work: **courir**

- Needed a harness to test “daemonization” of the PRISM model checker (using curl, jetty)
- Made simplifying assumptions: Linux, shell
 - Easy for job control, file redirection, “plugins”
 - Not as good at floating point
- Executables (scripts, other)
 - Batch pre/post, Trial pre/post, Run script
- Unit types, Statistics, Comparison
- Surprising results of running parallel jobs



Tsutomu Takasu

Proposal: Open Source Performance Benchmark Tools

- Open Source
 - Everyone needs this
 - Tool to facilitate peer review
 - Good for OpenJDK
- Cross-platform
- Doesn't depend on shell scripts (not Cygwin)
- Can measure what we want to see in the JVM
- Has basic statistics and comparison primitives

I'm going to use... Java?

- Our best known cross platform choice
- Some disadvantages
 - Job control
 - Parsing + string manipulation
 - Harness discoloration of benchmark results (?)
- The Devil you know
- Select interfaces (that could be implemented in other languages)
- Clear, independent data format (disk layout)
 - XML? Flat files? Database (ready)?



Arthur Chapman

Easy to install, repeat results

- Just like CPAN, but for Java
- Hey, how about using Jigsaw with a public module repository?
- Use your OS' native installer which collaborates with Jigsaw (vs. maven)
- Facilitates peer review
 - Is that code really doing what we *think* it's doing?


Easy to upload + share results

- Data repository like OpenStreetMap or Wikipedia
- Simple publication mechanism (zip file?)
- Compare results (and show caveats)
- Let data lead to visualization creativity
 - Data mining
 - Projections (high dimension reductions)
- Note: careful to not leak private/proprietary info

Continuous Performance Analysis

- Part of CI:
 - Build, Functional Test, Performance Test, Release
- Pool of resources
 - Slaves by OS/Architecture
 - EC2: Linux, Windows, Solaris? Not Mac :(
- Leverage Jenkins' features
 - Plugins, Notification (e-mail, IM, emacs)
 - Simple RPC, CLI, Web Services
- Going further
 - Public pools (OpenJDK qualification?)
 - BOINC style performance at home?





Choice of License

- Copyleft vs. Permissive
- Really want to get broad adoption
- Certain actors in the OpenJDK ecosystem have very strong feelings about this choice
- Compatibility with other projects
 - (e.g. Jenkins is MIT)
- Listen to jra on faif.us 0x0F

Statistics

- Student's T-Test
- Run to confidence
- Distribution of latency
- Comparison (bigger/smaller is better)
- Derived results
 - e.g. `sys + user => cpu`, `cpu/elapsed => % cpu`

Platform introspection

- CPU's/cores
- RAM
- Disk
- Graphics hardware
- OS version

Load Generation, Instrumentation

- Simulated clients (http only, http+ javascript)
- Jstat
- OS statistics
- Stress tools (inside the JVM process)

X3

Eventos Esportivos

... É A VIDA SEM DROGAS

AZ



Job Control

- Need to manage children and grandchildren
- Could terminate early due to:
 - Interactive quit (^C)
 - Sending SIGTERM
 - Benchmark setup/execution fails
 - Timeout
- Discussions on toolkits (ProcessBuilder is not enough)
 - JNR
 - NIO2

More Challenges + Opportunities

- Memory accounting
 - Shared vs. resident
- GUI automation
 - Will Windows GUI's, you know, like start?
 - T-Plan scripting for VNC?
- Ergonomics, hill climbing...
- Leverage existing tools
 - Jenkins, JNR, Caliper, other ?
 - “Why would a microbenchmark result be fictional?”
 - “Java code itself does not have intrinsic performance properties, any more than a cocktail-napkin sketch of an airplane has intrinsic aerodynamic properties. “

Discussion

- Would like to have such a beast?
- What's missing from the wish list?
- Let's collaborate!

@tmarble
tmarble@info9.net

Informatique

Copyright © 2011 Tom Marble, Licensed [CC by-sa](#)

