

Measuring the startup time of a Java Virtual Machine

Bruce Merry and Carl Hultquist

Why?

- Programming contests use timing to reward efficient algorithms
- Fairness between languages
- Fairness between runs within a language

Terminology

- Execution time: total time used by the process
- Startup time: time prior to any user code
- Solution time: time used by user code

Terminology

$\text{Solution time} = \text{Execution time} - \text{Startup time}$

Simple approach

- Run an empty Java program N times
- Call the average time the startup time
- Subtract that constant from execution time

Simple approach: limitations

- Benchmark is only valid for one combination of hardware and software
- Only measures *average* startup time
- Does not guarantee fairness between runs

Our approach

- Measure the startup time of each run as it happens
- The startup time is fed back into the framework and used to compute solution time

Our approach: implementation

1. Launch the JVM on a wrapper class
2. Load the user's class
3. Measure the startup time
4. Report the startup time to the framework
5. Execute the user's main method

Loading the user's class

- Done using Java's reflection mechanisms.
- Class is *not* initialized, merely loaded.
- Allows the wrapper to be written once.
- See paper for sample code

Measuring time elapsed

- Want to measure CPU time, not wall-clock
- Done using a native method
- Native method also reports to the parent
 - Writes elapsed time to file descriptor 3

Security

- Startup time reported before any user code has a chance to run (including static initializer)
- FD 3 is closed after reporting

Test setup

- 64-bit Gentoo Linux with Sun JDK 1.6
- Laptop with Core 2 Duo 2.16GHz
- Test system uses a 1000Hz tick
- Timing thus quantised to 1ms

Results

- JVM startup time is 70ms
- Wrapper increases execution time by <2ms
- Startup time is surprisingly stable
 - Std dev 0.5ms (below quantisation noise)
 - Our method does not improve std dev of solution time over simple method

Conclusions

- Startup time much less than 5-10 years ago
- Startup time much more consistent than 5-10 years ago
- Startup time can be reliably measured
- Startup time should not be a barrier to use of Java (and a JVM) in programming contests
 - Execution speed is outside the scope of this paper

Questions

?