



Metrics for Success: Performance Analysis 101

February 21, 2008

Kuldip Oberoi
Developer Tools
Sun Microsystems, Inc.



Agenda




- Application Performance
- Compiling for performance
- Profiling for performance
- Closing remarks

Sun Studio software

C/C++/Fortran tooling for the multi-core era



- Adoption increased 100+% over 2 years. Large footprint in enterprise/technical accts & growth in open src usage
- #1 IDE in the Evans survey for Performance of Resulting Applications category
- Intel, AMD, Sun & Fujitsu Partnerships



- **Parallelism** – feature-rich toolchain (auto-parallelizing compilers, thread analysis / debugging / profiling, OpenMP support, ...) & MPI support via Sun HPC ClusterTools
- **Performance** – dozens of industry benchmark records in the past year over Intel, AMD, Sun, & Fujitsu architectures
- **Productivity** – NetBeans-based IDE, code & memory debuggers, application profiler
- **Platforms** – Simplified dev across architectures & OSs (Solaris OS, OpenSolaris OS, Linux)



Sun Studio Software Overview



Integrated Toolchain



- Record-setting parallelizing C/C++/Fortran Compilers with autopar
- NetBeans-based IDE
- Stable, Scriptable, Multilingual Debugger (dbx)
- Memory Debugger- leak, access, usage (RTC)
- Application Profiling Tools (Performance Analyzer)
- Multi-core Optimizations, Multithreaded High Performance Libraries
- OpenMP API Support
- Multithreading Tools- Thread Analysis

<http://developers.sun.com/sunstudio>



Agenda



- Application Performance
- Compiling for performance
- Profiling for performance
- Closing remarks

Why Care about Performance?

- Because your company cares
 - > Faster code => greater productivity => lower cost
- Because your peers and customers care
 - > Better performance => less HW => lower cost
- Because it's interesting and suprising
 - > Coding is based on assumptions about behavior
 - > Performance problems arise from disconnects
- Because it's easy to do
 - > Simple runs, automated runs

What's a Performance Problem?

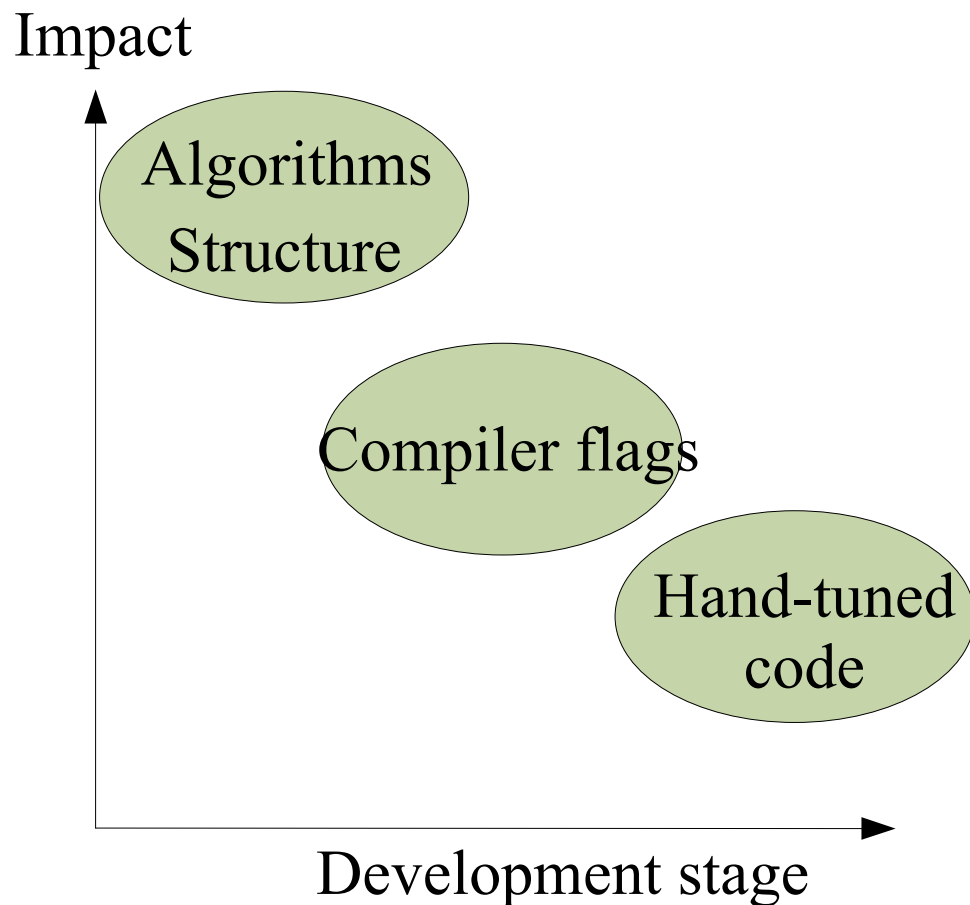
- Subjective criteria:
 - > It takes too long to finish
 - > It responds too slowly
- Objective criteria:
 - > It can't handle the required load
 - > It consumes too many resources to do its work
- Is it worth fixing?
 - > Cost of fixing vs. aggregate cost of problem
- Most untuned codes have low-hanging fruit!

Where does performance come from?

- Not the real question...
 - > Every application has a maximum theoretical performance
 - > Design decisions and implementation can deliver lower actual performance
- Alternative: *Where has performance been lost?*
 - > Need to identify where the time is spent
 - > Then determine what can be done about it

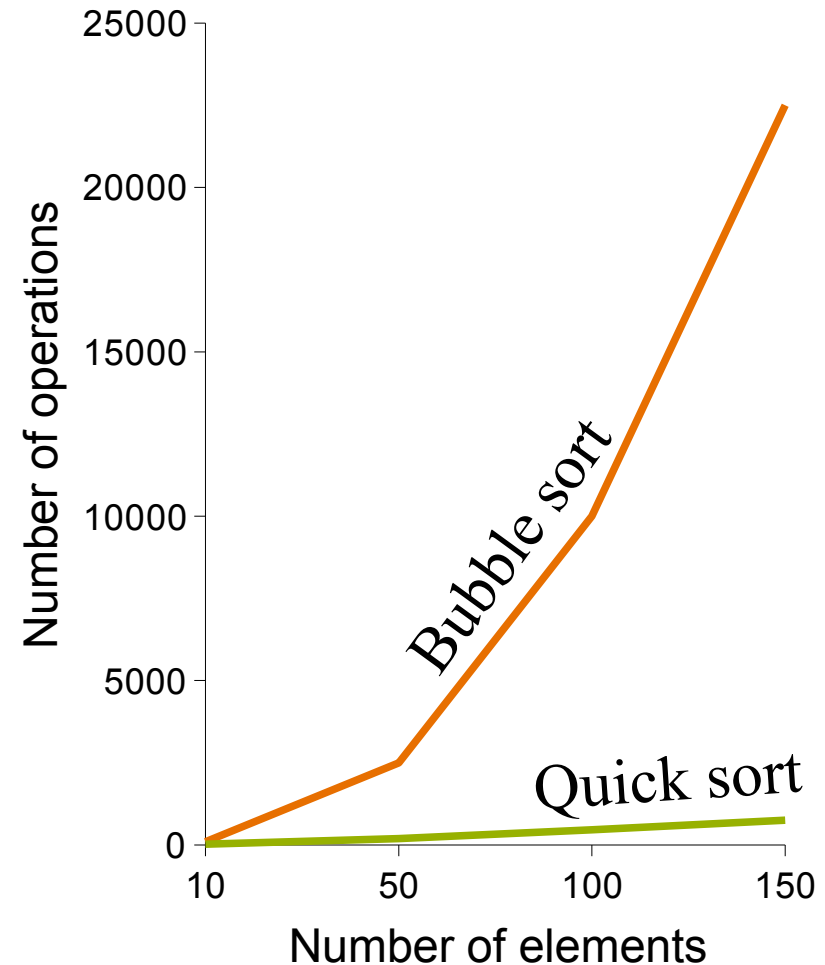
Where performance goes to

- Performance opportunities:
 - > Algorithms
 - > Structure
 - > Compiler flags
 - > Hand-tuned code



Algorithmic complexity

- How many operations?
- Classic example:
 - > Bubble sort $O(n^2)$
 - > Quick sort $O(n \log(n))$
- Sun Studio libraries have optimized code
 - > perflib (BLAS, FFTs, etc)
 - > medialib (images, codecs)
 - > Optimized maths libraries



Compiler flags

- The compiler's job:
 - > Produce the best code
 - > Given little knowledge of developer's intent
 - > Best code regardless of coding style
- Increasing optimization
 - > Leads to improved performance
 - > Relies on standard conforming source code
- Does a better job with
 - > Visibility of more of the code
 - > More information about the intent of the developer

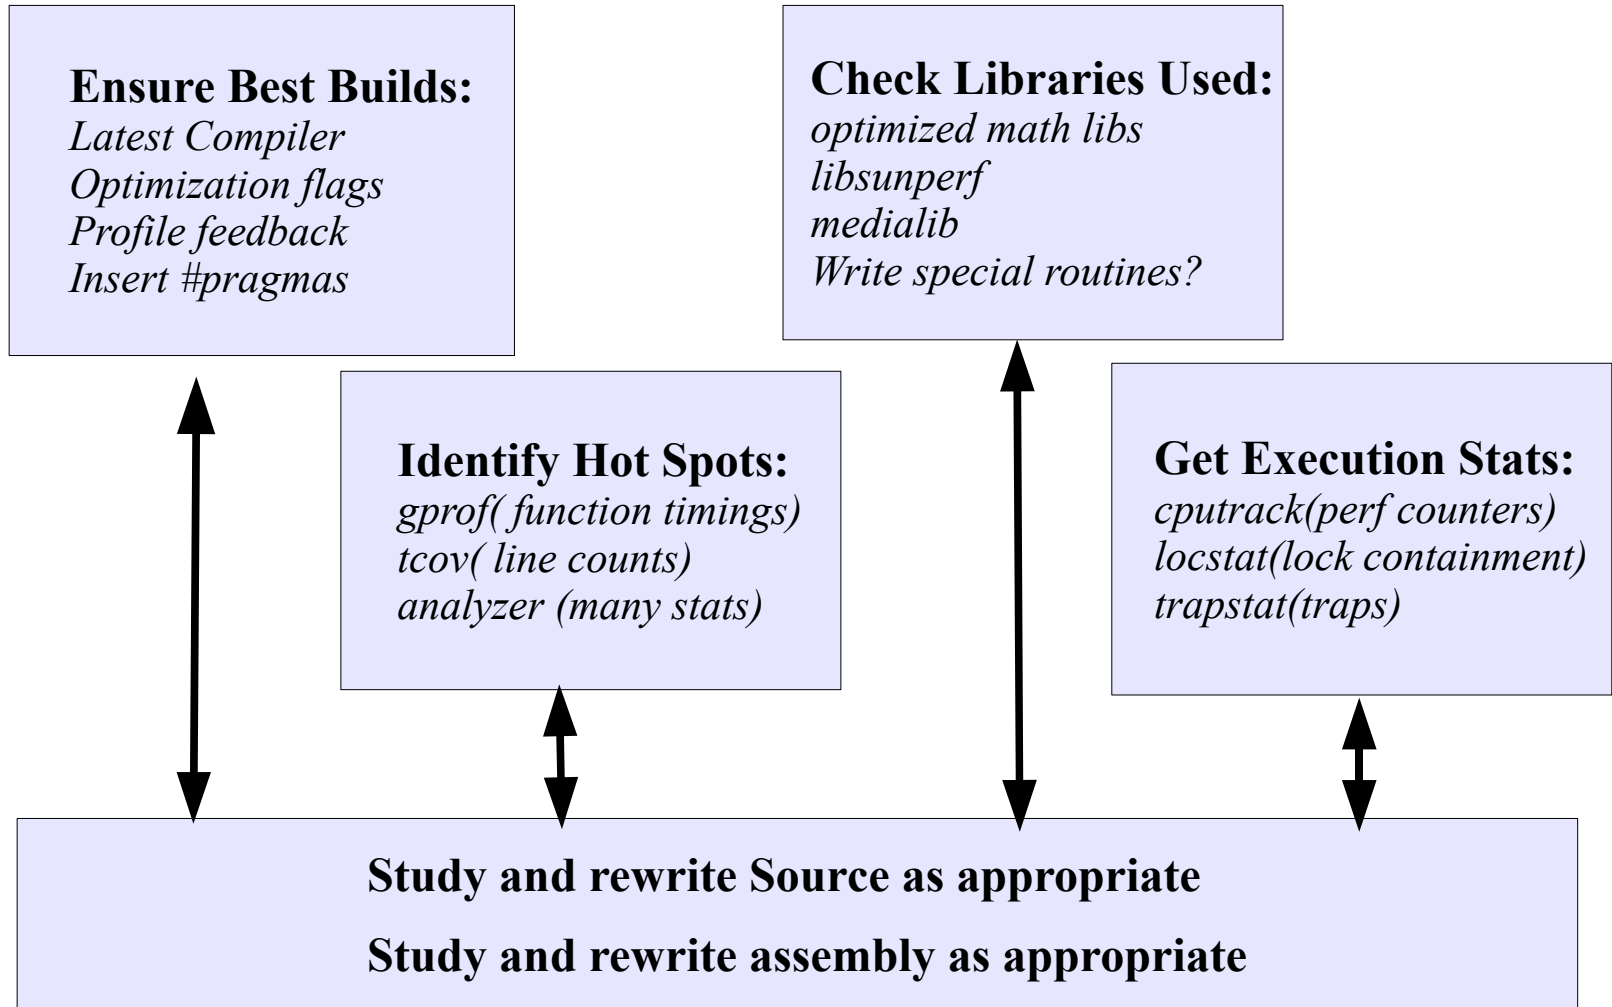
Hand-tuned code

- Examples:
 - > Special case code for the common situation
 - > Assembly language versions of key routines
- Cons:
 - > Time consuming
 - > Inflexible (e.g. workload specific)
 - > Platform specific
- Good examples:
 - > Use of language standards to provide compiler with more information (e.g. **restrict** keyword)

Perspective on optimization

- Algorithms, structures, and compiler flags
 - > High-level change
 - > Useful for all platforms
 - > Improve the application for all workloads
- Tweaks, hand-code
 - > Low-level (localised) change
 - > Platform specific
 - > Often already done by the compiler
- So:
 - > Focus on the high-level
 - > Unless the low-level gains are clear

Methodology / Tools Used



Agenda



- Application Performance
- **Compiling for performance**
- Profiling for performance
- Closing remarks

Choosing compiler flags

- Use the flags
 - > That you understand
 - > That you need (i.e. make a difference)
- Don't use flags
 - > That you don't understand
 - > That don't have an impact

Optimization

- Rule:
 - > No optimization flags means no optimization
- Suggestions:
 - > Use at least `-O`
 - > Try `-fast`
- Notes:
 - > Compile and link with the same flags

Debug information

- **Always** generate debug information
 - > **-g** for C/Fortran
 - > **-g0** for C++
- Also useful for profiling
- No/minimal performance impact

Exploring **-fast**

- **-fast** is a macro-flag:
 - > Enables a number of potentially useful optimisations
 - > May not be suitable for all situations
- Assumes build machine = run machine
 - > Use **-xtarget=** to specify otherwise
- Enables floating point simplification
 - > Use **-fsimple=0 -fns=no** otherwise
- Assumes basic pointer types do not alias
 - > Use **-xalias_level=any** otherwise
- Flags are parsed from left to right
 - > Override by placing flags on right

Target hardware

- Not all processors implement the same instructions
 - > Application will not run if instructions are not implemented
- If build machine is machine that will run binary:
 - > **-xtarget=native**
- For binaries that will run on a wide range of machines:
 - > SPARC: **-xtarget=generic -xarch=sparcvis2**
 - > x86: **-xtarget=generic -xarch=sse2**

Instruction set extensions (SSE/x86)

- Instruction set extensions on x86
- Single Instruction Multiple Data (SIMD)
 - > e.g. two parallel add operations in a single instruction
- Enable generation with:
 - > **-xtarget=generic -xarch=sse2**
-xvector=simd

Target hardware 32-bit or 64-bit

- 32-bit (**-m32**) can address 4GB of memory
- 64-bit (**-m64**) can address >>4GB of memory
- 64-bit: pointers and longs are 8 bytes
 - > => larger memory footprint
 - > => slower
- x86 64-bit has
 - > More registers, Better ABI
 - > => faster
 - > For x86 64-bit faster except when dominated by increased memory footprint

Inlining and cross-file optimisation

- Inlining:
 - > Avoids call overhead
 - > Provides more opportunities to code optimisation
 - > Increases code size
- Within file inlining at **-xO4**
- Crossfile inlining at **-xipo**
 - > Inlines between source files
 - > Reduces impact of source code structuring
 - > Can increase compile time

Profile feedback

- Profile feedback enables the compiler
 - > To see the runtime behavior of the application
 - > To make better code layout decision
 - > To make the right inlining decisions
- Three step process:
 - > Compile with
-xprofile=collect:/dir/profile
 - > Run with training workload
 - > Compile with **-xprofile=use:/dir/profile**
- Very useful for applications containing lots of decision logic

Agenda



- Application Performance
- Compiling for performance
- Profiling for performance
- Closing remarks

Good practices

- **Always** profile your application
 - > Is the time being spent in the important code?
 - > Are there obvious hot-spots to improve?
 - > How does the profile change with the workload?
- Amdahl's law
 - > Limit on performance gain is the time spent in the slow code
- Fix performance issues
 - > But make the fixes at the highest possible level of abstraction

Why use Sun Studio Tools? (I)

- The work for production code, production runs
 - > Runs from tens of seconds through hours
- They measure real behavior
 - > Fully optimized and parallelized applications
 - > Java HotSpot enabled
- They have minimal dilation and distortion
 - > ~5% for typical apps, ~10% for Java apps
- Supports code compile with Sun Studio & GNU compilers
- They're **FREE** for Solaris & Linux

Why use Sun Studio Tools? (II)

- They make things as simple as possible
 - > Show data in the user's source model
 - > Including OpenMP, MPI, Java, threads, etc.
- ..., but no simpler
 - > Show exactly what the compiler did
 - > Inlines, outlines, clones, parallel routines
 - > Show what JVM did
 - > Interpreted methods, HotSpot-compiled methods
 - > GC & HotSpot-compiler activities

Questions

- What can I change to improve performance?
- Which resources are being used?
- Where are they being used?
- Single-threaded
 - > Is the CPU being used efficiently?
 - > Memory subsystem delays? (TLBs, caches)
 - > I/O subsystem problems? (disk, network, paging)
- Multi-threaded
 - > Similar to single-threaded &
 - > Load unbalanced? Lock contention? memory/cache contention?

Gathering profiles

- Use **-g/-g0** for attribution of time to source line
- Gather profile with:
 - > **collect** <app> <params>
 - > **collect -P** <pid>
- Analyse profile with:
 - > **analyzer** test.<N>.er
 - > **er_print** test.<N>.er

Application profile

Sun Studio Analyzer [test.2.er]

File View Timeline Help

Functions Callers-Callees Source Disassembly Timeline Experiments

| User CPU (sec.) | User CPU (sec.) | Name |
|-----------------|-----------------|------------------|
| 11.838 | 11.838 | <Total> |
| 11.818 | 11.838 | main |
| 0.020 | 0.020 | _brk_unlocked |
| 0. | 0.020 | malloc |
| 0. | 0.020 | _malloc_unlocked |
| 0. | 0.020 | _morecore |
| 0. | 0.020 | sbrk |
| 0. | 0.020 | _sbrk_unlocked |
| 0. | 11.838 | _start |

Caller-Callee

Performance Analyzer [test.1.er]

File View Timeline Help

Find Text:

Functions Callers-Callees Source Lines Disassembly PCs Timeline LeakList Statistics Experiments

| ↓ ↑ User CPU (sec.) | ↓ ↑ User CPU (sec.) | ↓ ↑ User CPU (sec.) | Name |
|---------------------|---------------------|---------------------|-------------------------|
| 53.237 | 5.574 | 84.299 | Quiesce |
| 7.785 | 15.181 | 138.347 | Search |
| 42.400 | 42.400 | 61.023 | Evaluate |
| 9.216 | 9.216 | 9.216 | EvaluatePawns |
| 8.346 | 8.346 | 8.346 | EvaluatePassedPawns |
| 0.520 | 0.570 | 0.570 | FirstOne |
| 0.230 | 13.720 | 13.720 | Swap |
| 0.200 | 0.400 | 0.400 | LastOne |
| 0.100 | 0.080 | 0.100 | EvaluatePassedPawnRaces |
| 0.010 | 0.010 | 0.010 | EvaluateDraws |

Source level profile

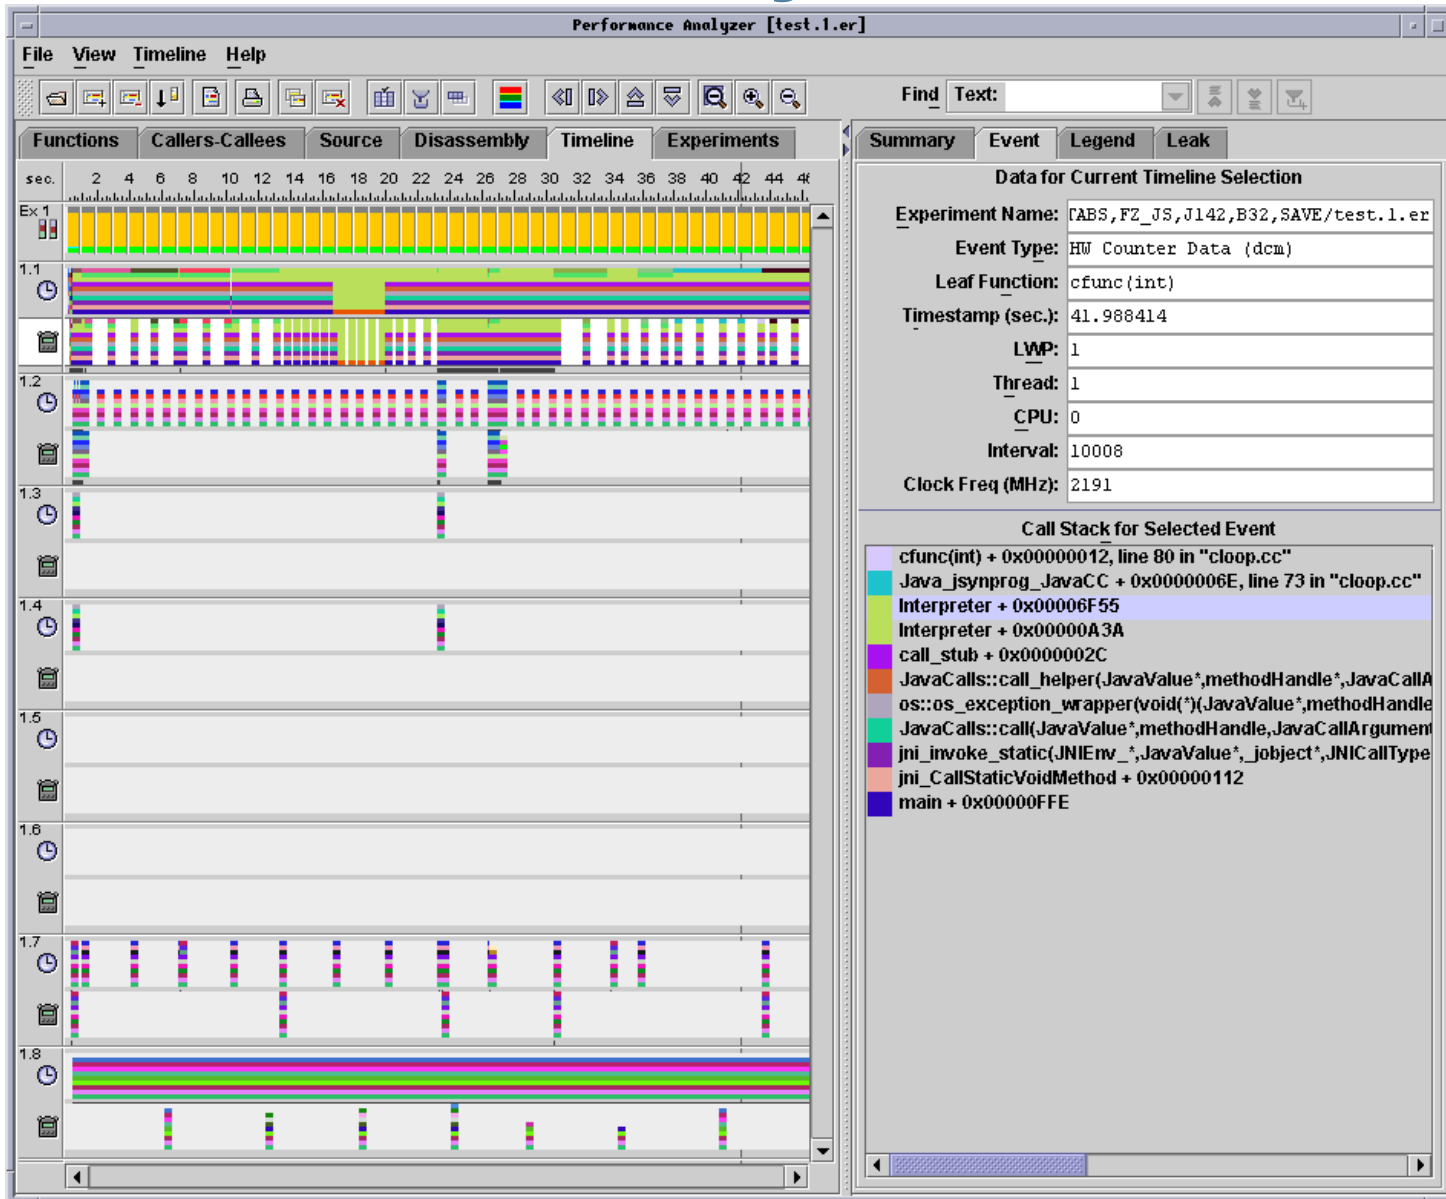
Sun Studio Analyzer [test.2.er]

File View Timeline Help

Functions Callers-Callees Source Disassembly Timeline Experiments

| User CPU (sec.) | User CPU (sec.) | Source File: ./ml.c Object File: ./ml Load Object: <ml> |
|-----------------|-----------------|---|
| 0. | 0. | 20. int inset(double ix, double iy) |
| 0. | 0. | 21. { |
| 0. | 0. | <Function: inset> |
| 0. | 0. | 22. int iterations=0; |
| 0. | 0. | 23. double x=ix, y=iy, x2=x*x, y2=y*y; |
| 6.885 | 6.885 | 24. while ((x2+y2<4) && (iterations<1000)) |
| 2.141 | 2.141 | 25. { |
| 0.430 | 0.430 | 26. y = 2 * x * y + iy; |
| 0.480 | 0.480 | 27. x = x2 - y2 + ix; |
| 1.411 | 1.411 | 28. x2 = x * x; |
| 0. | 0. | 29. y2 = y * y; |
| 0. | 0. | 30. iterations++; |
| 0. | 0. | 31. } |
| 0. | 0. | 32. return iterations; |
| | | 33. } |

Performance Analyzer - Timeline



DEMO

Agenda



- Application Performance
- Compiling for performance
- Profiling for performance
- Closing remarks

The checklist

- Build with optimization
 - > At least `-O`
- Build with debug enabled
 - > `-g` (`-g0` for C++)
- Profile
 - > **collect** `<app>` `<params>`

Optimization: Increasing optimisation

- Increased optimization (**-fast**)
 - > Typically improved performance
 - > Be aware of the optimisations enabled
- Use crossfile optimization (**-xipo**)
 - > Typically good for all codes

Optimization: Increasing information

- Profile feedback to give more information
 - > `-xprofile=[collect:|use:]`
 - > Good for all codes
 - > Particularly helpful for inlining and branches

Optimization: Leveraging libraries

- If time is spent in library code:
 - > Supplied optimized maths functions
-xlibmil -xlibmopt
 - > Optimized STL for C++
-library=stlport4
 - > The performance library
-library=sunperf

Summary

- Always profile
- Always use optimization

Gains from Tuning Categories

| <u>Tuning Category</u> | <u>Typical Range of Gain</u> |
|--|------------------------------|
| <i>Source Change</i> | <i>25-100%</i> |
| <i>Compiler Flags</i> | <i>5-20%</i> |
| <i>Use of libraries</i> | <i>25-200%</i> |
| <i>Assembly coding / tweaking</i> | <i>5-20%</i> |
| <i>Manual prefetching</i> | <i>5-30%</i> |
| <i>TLB thrashing/cache</i> | <i>20-100%</i> |
| <i>Using vis/inlines/micro-vectorization</i> | <i>100-200%</i> |

Sun Studio C, C++ and Fortran Compilers and Tools

http://developers.sun.com/sunstudio/ RSS Google

Apple Yahoo! Google Maps YouTube Wikipedia News (221) Popular

Sun Java Solaris Communities My SDN Account Join SDN


Sun Developer Network (SDN) search tips Search

APIs Downloads Products Support Training Participate

Developers Home > Sun Studio

Sun Studio

C, C++ & Fortran Compilers and Tools



PERFORMANCE MATTERS
Download record-setting compilers and tools for FREE or
[Order a Free Media Kit Now](#)

» Site Index

Sun Studio
» on Twitter
» on Facebook

Search Sun Studio

Sun Studio Add-ons
Cool Tools work together to support the stages of porting, building, tuning, and debugging an application on SPARC systems.
» Download Now

Free GlassFish Training
Download and register your GlassFish server to get free training and more. [FEEDBACK](#)

Overview Features Documentation Community Support Downloads

At a Glance | What's New | Topics | Videos | Heroes | Partners

Sun Studio software delivers a high-performance, optimizing C, C++, and Fortran developer toolchain for Solaris, **OpenSolaris**, and Linux operating systems, including support for multicore x86- and SPARC-based systems. The toolchain includes parallelizing compilers, code-level and memory debuggers, performance and thread analysis tools, OpenMP support as well as optimized math libraries. With a next-generation **NetBeans**-based IDE, development of multicore applications has never been easier. Get started with Sun Studio 12 or Sun Studio Express 11/08.

[» Download](#) **Sun Studio 12** [» Download](#) **Sun Studio Express 11/08**
(contains the latest features and enhancements)

What's New

February 20, 2009
Sun Studio Webinar: Maximizing Application Performance
To increase performance and assure scalability in your applications you need compilers that optimize your code and profiling tools that identify bottlenecks, hot spots, and memory access issues. Learn how the Sun Studio Thread Analyzer, Performance Analyzer, and D-Light can help you tune your application for maximum performance.

Spotlight
Sun Studio Software Registration Sweepstakes Register your Sun Studio Express 07/08 or Sun Studio Express 11/08 product for a chance to win Bose QuietComfort headphones, a Wii game system, or a Guitar Hero World Tour Band Kit!

Podcasts and Videos
Sun Studio Express 11/08
Join Ikroop Dhillon, Sun Studio Product Marketing Manager, and Vijay Tatkar, Sun Senior Engineer.



Metrics for Success: Performance Analysis 101

Thanks!!

Kuldip Oberoi
koberoi@sun.com
<http://koberoi.com>

