

# Phys 7411: Measuring Program Performance

Karen Tomko and Juana Moreno

[ktomko@osc.edu](mailto:ktomko@osc.edu) or [moreno@phys.lsu.edu](mailto:moreno@phys.lsu.edu)

Spring 2009

# Today's topics



- Does performance matter?
- Guidelines/Rules of thumb
- Timers
- Profiling tools
- References:
  - Chap. 7, Kernighan and Pike, *The Practice of Programming*, Addison Wesley 1999.
  - Chap. 15, Dongarra et. all, *Sourcebook of Parallel Computing*, Morgan Kaufmann/Elsevier, 2003.

# Bottleneck

## **bot·tle·neck**

2. a place or stage in a process at which progress is impeded.

Dictionary.com Unabridged

Based on the Random House Dictionary, © Random House, Inc. 2009.

# Computing Bottlenecks

Typically restrict flow of data from

- Memory
- Other processors/computers
- Storage (Disk, data base, ..)

to the calculation units (floating point adders, multipliers, etc)

# Guidelines from Kernighan and Pike

- Automate timing measurements
- Use a profiler
- Concentrate on the hot spots
- Draw a picture
- Use a better algorithm or data structure
- Enable compiler optimizations

# Timing Measurement

- Whole program
  - Single processor: /usr/bin/time
  - In parallel, with qsub on glen: oprofile tool on the perfmon node
    - #PBS -l nodes=4:ppn=4:**perfmon**
- Time a section of code: MPI\_Wtime, C and fortran calls, see:  
[http://www.osc.edu/supercomputing/training/opt/opteron\\_mod2\\_0801.pdf](http://www.osc.edu/supercomputing/training/opt/opteron_mod2_0801.pdf) (slides 34-35)

# Timing library

- GPTL - General Purpose Timing Library  
A library which allows use to name timers for measuring many sections of code for parallel programs which use MPI and/or threads/OpenMP

See:

<http://www.burningserver.net/rosinski/gptl/>

[https://wiki.ucar.edu/display/ccsm/CAM-CLM+\(GPTL-based\)+Timers](https://wiki.ucar.edu/display/ccsm/CAM-CLM+(GPTL-based)+Timers)

# Compiler Optimization

- Beyond `-O2`
  - Enable inlining
  - Enable vector instructions
  - Enable loop unrolling
  - Seek advice

# Portland Group Compilers Common options

- -fast (“best” optimization for machine; equivalent to -O2 -Mnoframe -Munroll on IA32 systems)
- -O0 (no optimization)
- -O1 (light optimization; default)
- -O2 (heavy optimization)
- -Mnoframe (eliminates stack pointer operations)
- -Munroll (enables loop unrolling)
- -Mvect=assoc (enables vector optimizations)
- -Mvect=cachesize:1048576 (sets assumed L2 cache size to 1MB)
- -Mconcur=option (enables automatic parallelization)
- -tp amd64 (enables generation of code for the AMD Opteron or Athlon-64 in 64-bit mode)
- -tp x64 (enables generation of "fat binary" with code for both AMD Opteron/Athlon-64 and Intel EM64T in 64-bit mode)

# Profiling tools

- gprof (and its variants)
  - Single processor, time by program unit
- MPI tracing tools
  - Jumpshot (Glenn)
  - mpiP (<http://mpip.sourceforge.net/>)
- Oprofile (Glenn)
- Hardware Performance counters
  - Detailed chip statistics, such as cache misses and page faults
  - PAPI (<http://icl.cs.utk.edu/papi/>)
- Performance Analysis Systems
  - Tau (<http://www.cs.uoregon.edu/research/tau/home.php>)
- [http://www.osc.edu/supercomputing/training/opt/opteron\\_mod2\\_0801.pdf](http://www.osc.edu/supercomputing/training/opt/opteron_mod2_0801.pdf) (slides 36-51)