Architecture of Parallel Computer Systems
- Performance Benchmarking -

SoSe 18

L.079.05810

www.uni-paderborn.de/pc2

Definition of Benchmark

The definition of *Benchmark* by Merriam-Webster:

a standardized problem or test that serves as a basis for evaluation or comparison (as of computer system performance)

- Tool to allow a quantitative comparison
- Benchmarking variants of an application program as well as computer systems
- Performance benchmarking often called performance evaluation
- Enables a progress in development by defining <u>measurable</u> and <u>repeatable</u> objectives

Sources of Benchmark Programs

- Standard Performance Evaluation Corporation (SPEC)
  https://www.spec.org
- NASA Ames Parallel Benchmarks
  https://www.nas.nasa.gov/publications/npb.html
- Intel MPI Benchmarks
  https://software.intel.com/en-us/articles/intel-mpi-benchmarks
- OSU Micro-Benchmarks (OSB)
- Top500 Supercomputer Sites
  https://www.top500.org
- High Performance Conjugate Gradients (HPCG)
  http://www.hpcg-benchmark.org/
- HPGMG
- Graph500
  https://graph500.org
- STREAM: Sustainable Memory Bandwidth in High Performance Computers
  https://www.cs.virginia.edu/stream/

Benchmarking of HPC Systems

- Where benchmarking can help
  - to know the performance of my favorite application on a certain system,
  - to find the fastest / most efficient HPC system for a set of applications,
  - to design/build a better HPC system.
- Why do the answers are depending on the application?
  - Applications have different requirements on compute, communication, file I/O, and other resources.
  - HPC systems are very different in their performance characteristics.
- How a benchmark program looks like?
  - Typically the execution time of an application is quite long and it also depends on the input instance.
  - A benchmark program is a representative of an application but with less runtime and a suitable input.

There is no progress if you can‘t measure it

- Quite simply, what gets measured gets done.
- Knowing that something is being monitored causes us to work harder and perform better - called the Hawthorne Effect.

- This is also valid for application development and computer system development.

How to build a benchmark program

- A Benchmark represents an application
  - Represents a real problem
  - Benchmark runtime correlates with runtime of the application (improvements can also been seen in the application)
- Keep the benchmark simple
  - Easy to compile, run, and understand.
  - Moderate requirements on main memory and file space (if it is not explicitly under investigation).
  - Runtime should be adjustable.
- Think in open source
  - Source code and reference performance data should be free to download and measured performance can be shared.
  - Avoid proprietary parts in the source and data

Make your life easier

- Avoid hard restrictions in required numbers of processes and threads (nodes and cores)
- Provide a test input and a small input to support functionality check and performance tuning
- Show progress information in case of expected long runtimes
- Clear and unique file naming of outputs

Which numbers are of interest?

- Basic performance metrics
  - Time (e.g. Floating-Point operations per second)
  - Energy (e.g. performance per Joule)
  - Cost (e.g. performance per €)
- Derived performance numbers
  - Performance related to reference performance (speedup, efficiency)
  - Sustained performance related to peak performance

Analysis and Reporting of Performance Data

- Statistic values
  - Arithmetic mean
    strongly influenced by large values
    $$\frac{1}{n}\sum_{i=0}^{n} a_i = \frac{a_1 + a_2 + \cdots + a_n}{n}$$
  - Harmonic mean
    strongly influenced by small values
    $$\left(\frac{\sum_{i=1}^{n} a_i^{-1}}{n}\right)^{-1} = \frac{n}{\frac{1}{a_1} + \frac{1}{a_2} + \cdots + \frac{1}{a_n}}$$
  - Geometric mean
    each value is equivalent important
    $$\left(\prod_{i=1}^{n} a_i\right)^{\frac{1}{n}} = \sqrt[n]{a_1 a_2 \cdots a_n}$$
- Tables, diagrams, …

Omni Path: MPI bandwidth (p-t-p)

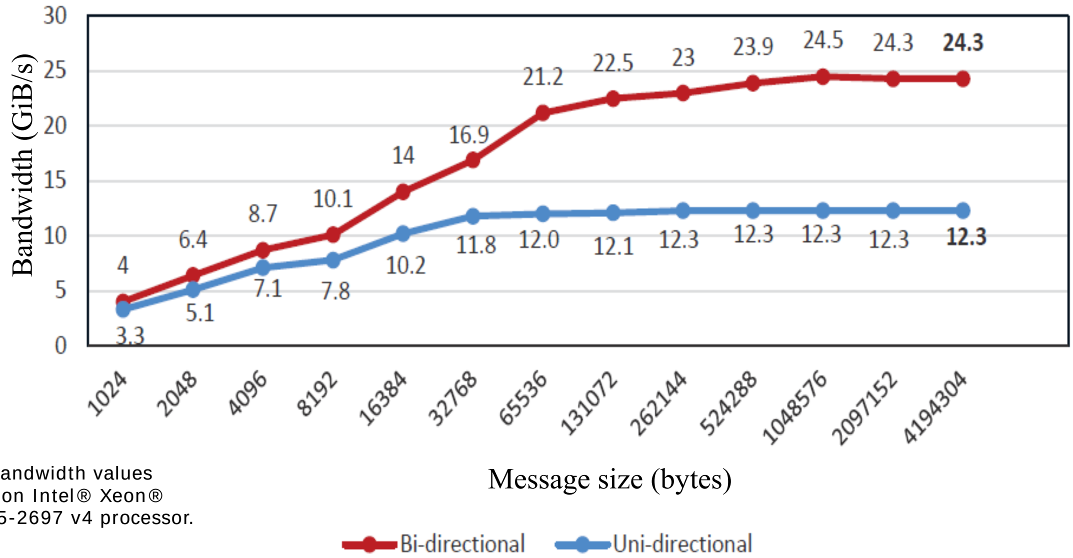

OSU Bandwidth values
based on Intel® Xeon®
CPU E5-2697 v4 processor.

Omni Path: NAMD Benchmark

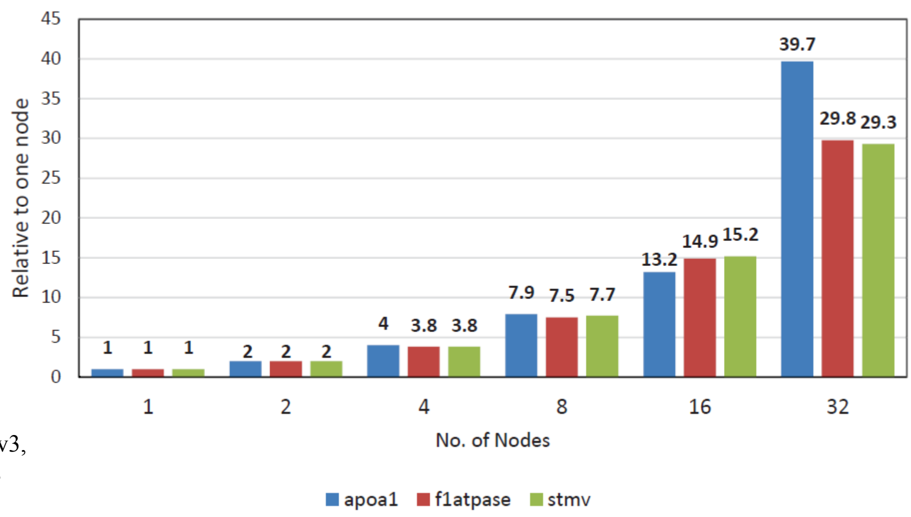

each 2 x E5-2697v3,
2.6GHz, 14 Cores

12 ways to fool the scientist (with performance evaluation)

- use benchmark applications unknown to others; give no reference
- use applications that have the same name as known benchmarks, but that show better performance of your innovation
- use only those benchmarks out of a suite that show good performance on your novel technique
- use only the benchmarks out of the suite that don’t break your technique
- modify the benchmark source code
- change data set parameters
- use the “debug” data set
- use a few loops out of the full programs only
- measure loop performance but label it as the full application
- don’t mention in your paper why you have chosen the benchmarks in this way and what changes you have made
- time the interesting part of the program only; exclude overheads
- measure interesting overheads only, exclude large unwanted items

Spec CPU

- SPEC CPU®2006 (retired, replaced by SPEC CPU®2017)
- Includes SPECint® and SPECfp® benchmark suite
- Several programs are included in each benchmark suite
- C99, C++98 compilers and for SPECfp additionally Fortran 95 compiler
- SPECspeed® and SPECrate® metrics for single and multiple tasks measurements
- Reports the geometric mean of the relative runtimes
- Designed to fit within about 2 GiB physical memory for each task you plan to run

Spec CPU

SPECint: Integer Benchmarks

Benchmark	Lang.	Application Domain
400.Perlbench	C	PERL Programming Language
401.bzip2	C	Compression
403.gcc	C	C Compiler
429.mcf	C	Combinatorial Optimization
445.gobmk	C	Artificial Intelligence: go
456.hmmer	C	Search Gene Sequence
458.sjeng	C	Artificial Intelligence: chess
462.libquantum	C	Physics: Quantum Computing
464.h264ref	C	Video Compression
471.omnetpp	C++	Discrete Event Simulation
473.astar	C++	Path-finding Algorithms
483.xalancbmk	C++	XML Processing

SPEC: Floating-Point Benchmarks

Benchmark	Language	Application Domain
410.Bwaves	Fortran	Fluid Dynamics
416.gamess	Fortran	Quantum Chemistry
433.milc	C	Physics: Quantum Chromodynamics
434.zeusmp	Fortran	Physics / CFD
435.gromacs	C/Fortran	Biochemistry/Molecular Dynamics
436.cactusADM	C/Fortran	Physics / General Relativity
437.leslie3d	Fortran	Fluid Dynamics
444.namd	C++	Biology / Molecular Dynamics
447.dealII	C++	Finite Element Analysis
450.soplex	C++	Linear Programming, Optimization
453.povray	C++	Image Ray-tracing
454.calculix	C/Fortran	Structural Mechanics
459.GemsFDTD	Fortran	Computational Electromagnetics
465.tonto	Fortran	Quantum Chemistry
470.lbm	C	Fluid Dynamics
481.wrf	C/Fortran	Weather Prediction
482.sphinx3	C	Speech recognition

SPEC CFP2006: Intel Gold 6148

SPEC® CFP2006 Result

Copyright 2006-2017 Standard Performance Evaluation Corporation

Hewlett Packard Enterprise (Test Sponsor: HPE)
ProLiant DL360 Gen10
(2.40 GHz, Intel Xeon Gold 6148)

SPECfp®_rate2006 = Not Run
SPECfp_rate_base2006 = 1440

CPU2006 license:	3	**Test date:**	Oct-2017
Test sponsor:	HPE	**Hardware Availability:**	Oct-2017
Tested by:	HPE	**Software Availability:**	Apr-2017

Hardware

CPU Name:	Intel Xeon Gold 6148
CPU Characteristics:	Intel Turbo Boost Technology up to 3.70 GHz
CPU MHz:	2400
FPU:	Integrated
CPU(s) enabled:	40 cores, 2 chips, 20 cores/chip, 2 threads/core
CPU(s) orderable:	1, 2 chip(s)
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per core
L3 Cache:	27.5 MB I+D on chip per chip
Other Cache:	None
Memory:	192 GB (24 x 8 GB 2Rx8 PC4-2666V-R)
Disk Subsystem:	1 x 480 GB SATA SSD, RAID 0
Other Hardware:	None

Software

Operating System:	SUSE Linux Enterprise Server 12 (x86_64) SP2 Kernel 4.4.21-69-default
Compiler:	C/C++: Version 17.0.3.191 of Intel C/C++ Compiler for Linux; Fortran: Version 17.0.3.191 of Intel Fortran Compiler for Linux
Auto Parallel:	No
File System:	xfs
System State:	Run level 3 (multi-user)
Base Pointers:	32/64-bit
Peak Pointers:	Not Applicable
Other Software:	None

Results Table

Benchmark	Base							Peak						
	Copies	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Copies	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
410.bwaves	80	1032	1050	**1033**	**1050**	1034	1050							
416.gamess	80	966	1620	**965**	**1620**	964	1630							
433.milc	80	704	1040	704	1040	**704**	**1040**							
434.zeusmp	80	413	1760	**415**	**1750**	420	1730							
435.gromacs	80	**282**	**2030**	282	2030	281	2030							
436.cactusADM	80	495	1930	**493**	**1940**	493	1940							
437.leslie3d	80	953	789	**952**	**790**	951	790							
444.namd	80	476	1350	474	1350	**474**	**1350**							
447.dealII	80	363	2520	**360**	**2540**	359	2550							
450.soplex	80	**836**	**798**	837	797	836	798							
453.povray	80	189	2250	**189**	**2260**	189	2260							
454.calculix	80	267	2470	267	2470	**267**	**2470**							
459.GemsFDTD	80	1157	734	1157	734	**1157**	**734**							
465.tonto	80	**485**	**1620**	480	1640	488	1610							
470.lbm	80	748	1470	**748**	**1470**	749	1470							
481.wrf	80	685	1300	**685**	**1300**	685	1310							
482.sphinx3	80	1125	1390	**1125**	**1390**	1125	1390							

Results appear in the order in which they were run. Bold underlined text indicates a median measurement.

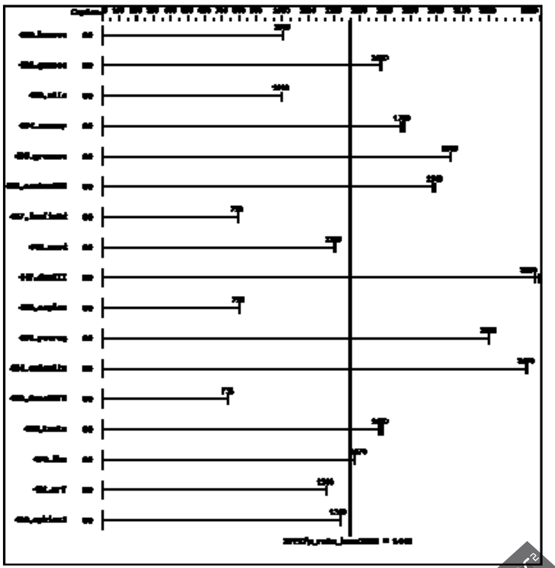

Paderborn Center for Parallel Computing

SPEC CFP2006: Intel Gold 6148 (cont.)

Operating System Notes

```
Stack size set to unlimited using "ulimit -s unlimited"
Transparent Huge Pages enabled by default
Filesystem page cache cleared with:
 shell invocation of 'sync; echo 3 > /proc/sys/vm/drop_caches' prior to run
runspec command invoked through numactl i.e.:
 numactl --interleave=all runspec <etc>
irqbalance disabled with "service irqbalance stop"
tuned profile set wtih "tuned-adm profile throughput-performance"
VM Dirty ratio was set to 40 using "echo 40 > /proc/sys/vm/dirty_ratio"
Numa balancing was disabled using "echo 0 > /proc/sys/kernel/numa_balancing"
```

Platform Notes

```
BIOS Configuration:
 Thermal Configuration set to Maximum Cooling
 Memory Patrol Scrubbing set to Disabled
 LLC Prefetcher set to Enabled
 LLC Dead Line Allocation set to Disabled
 Workload Pofile set to General Throughput Compute
  Minimum Processor Idle Power Core C-State set to C1E

Sysinfo program /home/specuser/cpu2006/config/sysinfo.rev6993
Revision 6993 of 2015-11-06 (b5e8d4b4eb51ed28d7f98696cbe290c1)
running on linux-lvuc Fri Oct  6 08:21:20 2017

This section contains SUT (System Under Test) info as seen by
some common utilities.  To remove or add to this section, see:
  http://www.spec.org/cpu2006/Docs/config.html#sysinfo

From /proc/cpuinfo
   model name : Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz
      2 "physical id"s (chips)
      80 "processors"
   cores, siblings (Caution: counting these is hw and system dependent.  The
   following excerpts from /proc/cpuinfo might not be reliable.  Use with
   caution.)
      cpu cores : 20
      siblings  : 40
      physical 0: cores 0 1 2 3 8 9 10 11 12 16 17 18 19 20 24 25 26 27 28
      physical 1: cores 0 1 2 3 8 9 10 11 12 16 17 18 19 20 24 25 26 27 28
   cache size : 28160 KB

From /proc/meminfo
   MemTotal:       197547324 kB
   HugePages_Total:       0
   Hugepagesize:       2048 kB

From /etc/*release* /etc/*version*
   SuSE-release:
      SUSE Linux Enterprise Server 12 (x86_64)
      VERSION = 12
      PATCHLEVEL = 2
      # This file is deprecated and will be removed in a future service pack or
      release.
      # Please check /etc/os-release for details about this release.
   os-release:
      NAME="SLES"
      VERSION="12-SP2"
      VERSION_ID="12.2"
```

…

Base Compiler Invocation

C benchmarks:
icc -m64

C++ benchmarks:
icpc -m64

Fortran benchmarks:
ifort -m64

Benchmarks using both Fortran and C:
icc -m64 ifort -m64

Base Portability Flags

410.bwaves: -DSPEC_CPU_LP64
416.gamess: -DSPEC_CPU_LP64
433.milc: -DSPEC_CPU_LP64
434.zeusmp: -DSPEC_CPU_LP64
435.gromacs: -DSPEC_CPU_LP64 -nofor_main
436.cactusADM: -DSPEC_CPU_LP64 -nofor_main
437.leslie3d: -DSPEC_CPU_LP64
444.namd: -DSPEC_CPU_LP64
447.dealII: -DSPEC_CPU_LP64
450.soplex: -DSPEC_CPU_LP64
453.povray: -DSPEC_CPU_LP64
454.calculix: -DSPEC_CPU_LP64 -nofor_main
459.GemsFDTD: -DSPEC_CPU_LP64
465.tonto: -DSPEC_CPU_LP64
470.lbm: -DSPEC_CPU_LP64
481.wrf: -DSPEC_CPU_LP64 -DSPEC_CPU_CASE_FLAG -DSPEC_CPU_LINUX
482.sphinx3: -DSPEC_CPU_LP64

Base Optimization Flags

C benchmarks:
-xCORE-AVX2 -ipo -O3 -no-prec-div -qopt-prefetch -auto-p32 -qopt-mem-layout-trans=3

C++ benchmarks:
-xCORE-AVX2 -ipo -O3 -no-prec-div -qopt-prefetch -auto-p32 -qopt-mem-layout-trans=3

Fortran benchmarks:
-xCORE-AVX2 -ipo -O3 -no-prec-div -qopt-prefetch

Benchmarks using both Fortran and C:
-xCORE-AVX2 -ipo -O3 -no-prec-div -qopt-prefetch -auto-p32 -qopt-mem-layout-trans=3

Paderborn Center for Parallel Computing

Spec OMP

- SPEC OMP®2012
- Based on OpenMP3.1
- C99, C++98, and Fortran 95 compilers
- Optional metric for power measurement
- Only one class of benchmark - SPEC OMPG2012
- Designed to fit within about 28 GiB physical memory
- Benchmarks emphasize the performance of
  - the processors and the interconnect between the processor
  - the memory architecture
  - the parallel support libraries
  - the compilers

SPEC omp

Benchmark	Language	Application Domain
350.Md	Fortran	Physics: Molecular Dynamics
351.bwaves	Fortran	Physics: Computational Fluid Dynamics (CFD)
352.nab	C	Molecular Modeling
357.bt331	Fortran	Physics: Computational Fluid Dynamics (CFD)
358.botsalgn	C	Protein Alignment
359.botsspar	C	Sparse LU
360.ilbdc	Fortran	Lattic Boltzmann
362.fma3d	Fortran	Mechanical Response Simulation
363.swim	Fortran	Weather Prediction
367.imagick	C	Image Processing
370.mgrid331	Fortran	Physics: Computational Fluid Dynamics (CFD)
371.applu331	Fortran	Physics: Computational Fluid Dynamics (CFD)
372.smithwa	C	Optimal Pattern Matching
376.kdtree	C++	Sorting and Searching

SPEC OMPG2012: Intel Gold 6148

SPEC® OMPG2012 Result

Copyright 2012-2017 Standard Performance Evaluation Corporation

Intel Intel Server System R2208WFTZS (Intel Xeon Gold 6148, DDR4-2666 Turbo OFF, SMT ON) Endeavour Node	SPECompG_base2012 = 15.0 SPECompG_peak2012 = 17.3

OMP2012 license:	13	**Test date:**	Jul-2017
Test sponsor:	Intel	**Hardware Availability:**	Jul-2017
Tested by:	Intel	**Software Availability:**	Sep-2017

Hardware

CPU Name:	Intel Xeon Gold 6148
CPU Characteristics:	Intel Turbo Boost Technology Disabled
CPU MHz:	2400
CPU MHz Maximum:	3700
FPU:	Integrated
CPU(s) enabled:	40 cores, 2 chips, 20 cores/chip, 2 threads/core
CPU(s) orderable:	1,2 Chips
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per core
L3 Cache:	27.5 MB I+D on chip per chip
Other Cache:	None
Memory:	192 GB (12 x 16 GB 2Rx4 DDR4-2666 ECC Registered)
Disk Subsystem:	Panasas ActiveStor 14 (Size 124TB)
Other Hardware:	--
Base Threads Run:	80
Minimum Peak Threads:	40
Maximum Peak Threads:	80

Software

Operating System:	Oracle Linux Server release 7.3
Compiler:	C/C++/Fortran: Version 18.0.0.082 of Intel Compiler for Linux Beta Build 20170510
Auto Parallel:	No
File System:	Linux ext3
System State:	Default
Base Pointers:	64-bit
Peak Pointers:	64-bit
Other Software:	None

Benchmark	Base							Peak						
	Threads	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Threads	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
350.md	80	67.6	68.5	**67.6**	**68.5**	68.2	67.8	80	68.5	67.6	**68.6**	**67.5**	69.1	67.0
351.bwaves	80	263	17.2	**263**	**17.2**	265	17.1	40	256	17.7	253	17.9	**253**	**17.9**
352.nab	80	356	10.9	358	10.9	**358**	**10.9**	80	349	11.2	**398**	**9.78**	399	9.76
357.bt331	80	258	18.4	**257**	**18.4**	256	18.5	80	258	18.4	**262**	**18.1**	263	18.0
358.botsalgn	80	392	11.1	394	11.0	**394**	**11.0**	80	394	11.0	481	9.04	**481**	**9.04**
359.botsspar	80	639	8.22	641	8.19	**640**	**8.21**	40	82.8	63.4	92.8	56.6	**92.7**	**56.6**
360.ilbdc	80	**308**	**11.5**	308	11.6	308	11.5	80	**313**	**11.4**	311	11.4	314	11.4
362.fma3d	80	**391**	**9.72**	395	9.62	389	9.78	80	**385**	**9.87**	379	10.0	388	9.80
363.swim	80	347	13.0	**347**	**13.0**	347	13.1	40	329	13.8	**328**	**13.8**	328	13.8
367.imagick	80	**458**	**15.4**	458	15.4	460	15.3	40	**413**	**17.0**	412	17.1	414	17.0
370.mgrid331	80	361	12.3	360	12.3	**360**	**12.3**	40	346	12.8	**345**	**12.8**	344	12.8
371.applu331	80	244	24.8	283	21.4	**258**	**23.5**	80	264	23.0	**261**	**23.2**	244	24.9
372.smithwa	80	**288**	**18.6**	287	18.6	293	18.3	40	228	23.5	226	23.7	**227**	**23.6**
376.kdtree	80	**411**	**11.0**	409	11.0	411	11.0	80	442	10.2	442	10.2	**442**	**10.2**

Results appear in the order in which they were run. Bold underlined text indicates a median measurement.

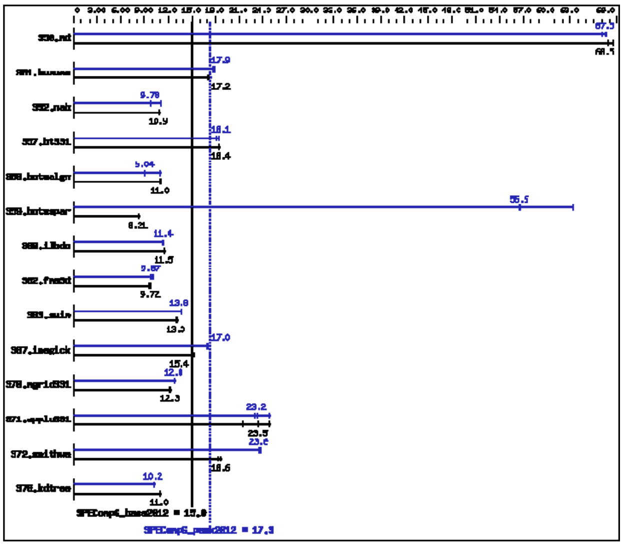

Paderborn Center for Parallel Computing

SPEC OMPG2012: Intel Gold 6148

Platform Notes

```
Sysinfo program /global/panfs02/innl/aknyazel/OMP2012/1.1/Docs/sysinfo
Revision 563 of 2016-06-10 (097295389cf6073d8c3b03fa376740a5)
running on epb331 Tue Jul  4 08:09:28 2017

This section contains SUT (System Under Test) info as seen by
some common utilities.  To remove or add to this section, see:
  http://www.spec.org/omp2012/Docs/config.html#sysinfo

From /proc/cpuinfo
   model name : Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz
      2 "physical id"s (chips)
      80 "processors"
   cores, siblings (Caution: counting these is hw and system dependent.  The
   following excerpts from /proc/cpuinfo might not be reliable.  Use with
   caution.)
      cpu cores : 20
      siblings  : 40
      physical 0: cores 0 1 2 3 4 8 9 10 11 12 16 17 18 19 20 24 25 26 27 28
      physical 1: cores 0 1 2 3 4 8 9 10 11 12 16 17 18 19 20 24 25 26 27 28
   cache size : 28160 KB

From /proc/meminfo
   MemTotal:       196699188 kB
   HugePages_Total:       0
   Hugepagesize:       2048 kB

From /etc/*release* /etc/*version*
   oracle-release: Oracle Linux Server release 7.3
   os-release:
      NAME="Oracle Linux Server"
      VERSION="7.3"
      ID="ol"
      VERSION_ID="7.3"
      PRETTY_NAME="Oracle Linux Server 7.3"
      ANSI_COLOR="0;31"
      CPE_NAME="cpe:/o:oracle:linux:7:3:server"
      HOME_URL="https://linux.oracle.com/"
   redhat-release: Red Hat Enterprise Linux Server release 7.3 (Maipo)
   system-release: Oracle Linux Server release 7.3
   system-release-cpe: cpe:/o:oracle:linux:7:3:server
uname -a:
   Linux epb331 3.10.0-514.6.2.0.1.el7.x86_64.knl1 #1 SMP Thu Mar 2 10:19:17 MST
   2017 x86_64 x86_64 x86_64 GNU/Linux

run-level 3 Jun 30 10:26

SPEC is set to: /global/panfs02/innl/aknyazel/OMP2012/1.1
   Filesystem                Type   Size  Used Avail Use% Mounted on
   panfs://36.101.212.1/innl panfs  108T   40T   69T  37% /global/panfs02/innl
Additional information from dmidecode:

   Warning: Use caution when you interpret this section. The 'dmidecode' program
   reads system data which is "intended to allow hardware to be accurately
   determined", but the intent may not be met, as there are frequent changes to
   hardware, firmware, and the "DMTF SMBIOS" standard.
```

General Notes

```
================================================================================
Intel Turbo Boost Technology Disabled
================================================================================
General base OMP Library Settings
  ENV_KMP_AFFINITY=compact,0,verbose

================================================================================
General peak OMP Library Settings
  ENV_KMP_AFFINITY=compact,0,verbose

================================================================================
Per benchmark peak OMP Library Settings

================================================================================
System settings notes:

================================================================================
General OMP Library Settings
  KMP_LIBRARY=turnaround
  KMP_STACKSIZE=292M
  KMP_BLOCKTIME=infinite
  OMP_DYNAMIC=FALSE
  OMP_NESTED=FALSE
  OMP_SCHEDULE=static
================================================================================
351.bwaves:peak:
  ENV_KMP_AFFINITY=compact,1,verbose

================================================================================
359.botsspar:peak:
  ENV_KMP_AFFINITY=compact,1,verbose

================================================================================
363.swim:peak:
  ENV_KMP_AFFINITY=compact,1,verbose

================================================================================
367.imagick:peak:
  ENV_KMP_AFFINITY=compact,1,verbose

================================================================================
370.mgrid331:peak:
  ENV_KMP_AFFINITY=compact,1,verbose

================================================================================
372.smithwa:peak:
  ENV_KMP_AFFINITY=compact,1,verbose
```


Paderborn Center for Parallel Computing

SPEC OMPG2012: Intel Gold 6148

Base Compiler Invocation

C benchmarks:
icc

C++ benchmarks:
icpc

Fortran benchmarks:
ifort

Base Portability Flags

350.md: -FR
357.bt331: -mcmodel=medium
363.swim: -mcmodel=medium
367.imagick: -std=c99

Base Optimization Flags

C benchmarks:
-O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -ansi-alias -no-prec-div -no-prec-sqrt

C++ benchmarks:
-O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -ansi-alias -no-prec-div -no-prec-sqrt

Fortran benchmarks:
-O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -ansi-alias -no-prec-div -no-prec-sqrt -align all

Peak Compiler Invocation

C benchmarks:
icc

C++ benchmarks:
icpc

Fortran benchmarks:
ifort

Peak Portability Flags

350.md: -FR
357.bt331: -mcmodel=medium
363.swim: -mcmodel=medium
367.imagick: -std=c99

Peak Optimization Flags

C benchmarks:
352.nab: -O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -no-prec-div -no-prec-sqrt -fno-alias -qopt-malloc-options=1 -qopt-calloc
358.botsalgn: -O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -fno-alias -no-prec-div -no-prec-sqrt
359.botsspar: Same as 358.botsalgn
367.imagick: -O3 -qopenmp -xCORE-AVX2 -fp-model fast=2 -fno-alias -no-prec-div -no-prec-sqrt
372.smithwa: Same as 367.imagick

C++ benchmarks:
-O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -fno-alias -no-prec-div -no-prec-sqrt

Fortran benchmarks:
350.md: -O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -fno-alias -no-prec-div -no-prec-sqrt -align all
351.bwaves: -O3 -qopenmp -xCORE-AVX2 -fp-model fast=2 -fno-alias -no-prec-div -no-prec-sqrt -align all
357.bt331: Same as 350.md
360.ilbdc: -O3 -qopenmp -xCORE-AVX512 -fp-model fast=2 -ansi-alias -no-prec-div -no-prec-sqrt -align all
362.fma3d: Same as 350.md
363.swim: -O3 -qopenmp -xCORE-AVX2 -fp-model fast=2 -no-prec-div -no-prec-sqrt -fno-alias -qopt-malloc-options=3 -align all
370.mgrid331: Same as 363.swim
371.applu331: Same as 351.bwaves

Spec MPI

- SPEC MPI®2007
- Focus on MPI-parallel, floating-point, and compute intensive applications
- Fortran 95 or later, C99, C++98, and MPI library is required
- MPI®M2007 medium size benchmark
- MPI®L2007 large size benchmark
  - designed scaling up to 2048 ranks
  - fit within about 2 GiB physical memory per rank
  - minimum officially support rank count is 64

Benchmark Application of Spec MPI

Benchmark	Suite	Language	Application Domain
104.Milc	medium	C	Physics: Quantum Chromodynamics (QCD)
107.leslie3d	medium	Fortran	Computational Fluid Dynamics (CFD)
113.GemsFDTD	medium	Fortran	Computational Electromagnetics (CEM)
115.fds4	medium	C/Fortran	Computational Fluid Dynamics (CFD)
121.pop2	medium, large	C/Fortran	Ocean Modeling
122.tachyon	medium, large	C	Graphics: Parallel Ray Tracing
125.RAxML	Large	C	DNA Matching
126.lammps	medium, large	C++	Molecular Dynamics Simulation
127.wrf2	Medium	C/Fortran	Weather Prediction
128.GAPgeofem	medium, large	C/Fortran	Heat Transfer using Finite Element Methods (FEM)
129.tera_tf	medium, large	Fortran	3D Eulerian Hydrodynamics
130.socorro	Medium	C/Fortran	Molecular Dynamics using Density-Functional Theory (DFT)
132.zeusmp2	medium, large	C/Fortran	Physics: Computational Fluid Dynamics (CFD)
137.lu	medium, large	Fortran	Computational Fluid Dynamics (CFD)
142.dmilc	large	C	Physics: Quantum Chromodynamics (QCD)
143.dleslie	large	Fortran	Computational Fluid Dynamics (CFD)
145.lGemsFDTD	large	Fortran	Computational Electromagnetics (CEM)
147.l2wrf2	large	C/Fortran	Weather Prediction

Spec MPIL2007: Intel Gold 6148

SPEC® MPIL2007 Result

Hewlett Packard Enterprise
SGI 8600
(Intel Xeon Gold 6148, 2.40 GHz)

SPECmpiL_peak2007 = Not Run
SPECmpiL_base2007 = 37.4

MPI2007 license:	1	Test date:	Oct-2017
Test sponsor:	HPE	Hardware Availability:	Jul-2017
Tested by:	HPE	Software Availability:	Nov-2017

Results Table

Benchmark	Base							Peak						
	Ranks	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Ranks	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
121.pop2	640	108	36.0	107	36.3	107	36.2							
122.tachyon	640	109	17.8	108	18.1	108	17.9							
125.RAxML	640	129	22.7	128	22.7	128	22.9							
126.lammps	640	96.8	25.4	97.2	25.3	97.8	25.2							
128.GAPgeofem	640	102	58.2	102	58.0	103	57.8							
129.tera_tf	640	57.7	19.0	58.3	18.9	58.1	18.9							
132.zeusmp2	640	62.5	33.9	60.8	34.9	61.0	34.8							
137.lu	640	53.1	79.1	53.1	79.2	53.1	79.1							
142.dmilc	640	65.1	56.6	65.3	56.4	65.3	56.5							
143.dleslie	640	44.2	70.1	44.4	69.9	45.9	67.5							
145.lGemsFDTD	640	116	38.1	116	38.0	116	38.0							
147.l2wrf2	640	188	43.7	188	43.7	188	43.6							

Results appear in the order in which they were run. Bold underlined text indicates a median measurement.

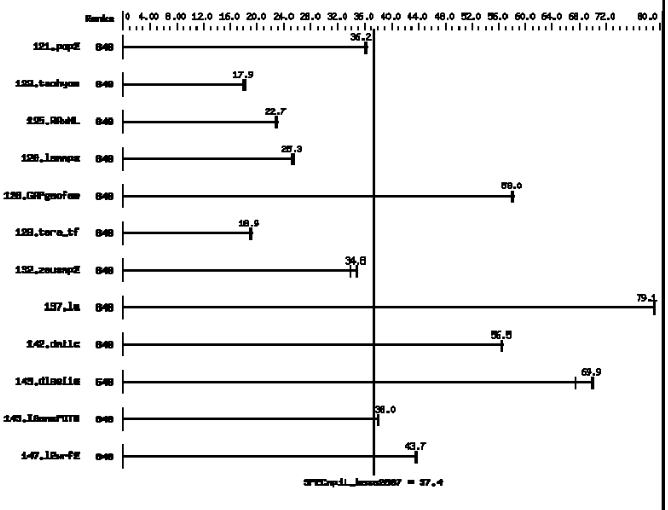

Spec MPIL2007: Intel Gold 6148 (cont.)

Hardware Summary

Type of System:	Homogeneous
Compute Node:	HPE XA730i Gen10 Server Node
Interconnect:	InfiniBand (MPI and I/O)
File Server Node:	Lustre FS
Total Compute Nodes:	16
Total Chips:	32
Total Cores:	640
Total Threads:	1280
Total Memory:	3 TB
Base Ranks Run:	640
Minimum Peak Ranks:	--
Maximum Peak Ranks:	--

Software Summary

C Compiler:	Intel C Composer XE for Linux, Version 18.0.0.128 Build 20170811
C++ Compiler:	Intel C++ Composer XE for Linux, Version 18.0.0.128 Build 20170811
Fortran Compiler:	Intel Fortran Composer XE for Linux, Version 18.0.0.128 Build 20170811
Base Pointers:	64-bit
Peak Pointers:	Not Applicable
MPI Library:	HPE Performance Software - Message Passing Interface 2.17
Other MPI Info:	OFED 3.2.2
Pre-processors:	None
Other Software:	None

Node Description: HPE XA730i Gen10 Server Node

Hardware

Number of nodes:	16
Uses of the node:	compute
Vendor:	Hewlett Packard Enterprise
Model:	SGI 8600 (Intel Xeon Gold 6148, 2.40 GHz)
CPU Name:	Intel Xeon Gold 6148
CPU(s) orderable:	1-2 chips
Chips enabled:	2
Cores enabled:	40
Cores per chip:	20
Threads per core:	2
CPU Characteristics:	Intel Turbo Boost Technology up to 3.70 GHz
CPU MHz:	2400
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per core
L3 Cache:	27.5 MB I+D on chip per chip
Other Cache:	None
Memory:	192 GB (12 x 16 GB 2Rx4 PC4-2666V-R)
Disk Subsystem:	None
Other Hardware:	None
Adapter:	Mellanox MT27700 with ConnectX-4 ASIC
Number of Adapters:	2
Slot Type:	PCIe x16 Gen3 8GT/s
Data Rate:	InfiniBand 4X EDR
Ports Used:	1
Interconnect Type:	InfiniBand

Software

Adapter:	Mellanox MT27700 with ConnectX-4 ASIC
Adapter Driver:	OFED-3.4-2.1.8.0
Adapter Firmware:	12.18.1000
Operating System:	Red Hat Enterprise Linux Server 7.3 (Maipo), Kernel 3.10.0-514.2.2.el7.x86_64
Local File System:	LFS
Shared File System:	LFS
System State:	Multi-user, run level 3
Other Software:	SGI Management Center Compute Node 3.5.0, Build 716r171.rhel73-1705051353

Hardware

Number of nodes:	4
Uses of the node:	fileserver
Vendor:	Hewlett Packard Enterprise
Model:	Rackable C1104-GP2 (Intel Xeon E5-2690 v3, 2.60 GHz)
CPU Name:	Intel Xeon E5-2690 v3
CPU(s) orderable:	1-2 chips
Chips enabled:	2
Cores enabled:	24
Cores per chip:	12
Threads per core:	1
CPU Characteristics:	Intel Turbo Boost Technology up to 3.50 GHz Hyper-Threading Technology disabled
CPU MHz:	2600
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	256 KB I+D on chip per core
L3 Cache:	30 MB I+D on chip per chip
Other Cache:	None
Memory:	128 GB (8 x 16 GB 2Rx4 PC4-2133P-R)
Disk Subsystem:	684 TB RAID 6 48 x 8+2 2TB 7200 RPM
Other Hardware:	None
Adapter:	Mellanox MT27700 with ConnectX-4 ASIC
Number of Adapters:	2
Slot Type:	PCIe x16 Gen3
Data Rate:	InfiniBand 4X EDR
Ports Used:	1
Interconnect Type:	InfiniBand

Software

Adapter:	Mellanox MT27700 with ConnectX-4 ASIC
Adapter Driver:	OFED-3.3-1.0.0.0
Adapter Firmware:	12.14.2036
Operating System:	Red Hat Enterprise Linux Server 7.3 (Maipo), Kernel 3.10.0-514.2.2.el7.x86_64
Local File System:	ext3
Shared File System:	LFS
System State:	Multi-user, run level 3
Other Software:	None

Interconnect Description: InfiniBand (MPI and I/O)

Hardware

Vendor:	Mellanox Technologies and SGI
Model:	SGI P0002145
Switch Model:	SGI P0002145
Number of Switches:	2
Number of Ports:	36
Data Rate:	InfiniBand 4X EDR
Firmware:	11.0350.0394
Topology:	Enhanced Hypercube
Primary Use:	MPI and I/O traffic

Paderborn Center for Parallel Computing

Spec MPIL2007: Intel Gold 6148 (cont.)

Interconnect Description: InfiniBand (MPI and I/O)

Hardware

Vendor:	Mellanox Technologies and SGI
Model:	SGI P0002145
Switch Model:	SGI P0002145
Number of Switches:	2
Number of Ports:	36
Data Rate:	InfiniBand 4X EDR
Firmware:	11.0350.0394
Topology:	Enhanced Hypercube
Primary Use:	MPI and I/O traffic

Base Tuning Notes

```
src.alt used: 143.dleslie->integer_overflow
```

Submit Notes

```
The config file option 'submit' was used.
```

General Notes

```
Software environment:
  export MPI_REQUEST_MAX=65536
  export MPI_TYPE_MAX=32768
  export MPI_IB_RAILS=2
  export MPI_IB_IMM_UPGRADE=false
  export MPI_IB_DCIS=2
  export MPI_IB_HYPER_LAZY=false
  export MPI_CONNECTIONS_THRESHOLD=0
  ulimit -s unlimited

BIOS settings:
  AMI BIOS version SAED7177, 07/17/2017

Job Placement:
  Each MPI job was assigned to a topologically compact set
  of nodes.

Additional notes regarding interconnect:
  The Infiniband network consists of two independent planes,
  with half the switches in the system allocated to each plane.
  I/O traffic is restricted to one plane, while MPI traffic can
  use both planes.
```

Base Compiler Invocation

C benchmarks:
icc

C++ benchmarks:
126.lammps: icpc

Fortran benchmarks:
ifort

Benchmarks using both Fortran and C:
icc ifort

Base Portability Flags

121.pop2: -DSPEC_MPI_CASE_FLAG

Base Optimization Flags

C benchmarks:
-O3 -xCORE-AVX512 -no-prec-div -ipo

C++ benchmarks:
126.lammps: -O3 -xCORE-AVX512 -no-prec-div -ansi-alias -ipo

Fortran benchmarks:
-O3 -xCORE-AVX512 -no-prec-div -ipo

Benchmarks using both Fortran and C:
-O3 -xCORE-AVX512 -no-prec-div -ipo

Base Other Flags

C benchmarks:
-lmpi

C++ benchmarks:
126.lammps: -lmpi

Fortran benchmarks:
-lmpi

Benchmarks using both Fortran and C:
-lmpi

Paderborn Center for Parallel Computing

Memory Bandwith: STREAM Benchmark

- Created by John D. McCalpin in 1991
- Intended to be an extremely simplified representation of low-compute-intensity, long-vector operations characteristic of applications
- Vector operations are (*i=1..N*)
  - o Copy: C[i] = A[i]
  - o Scale: B[i] = scalar * C[i]
  - o Add: C[i] = A[i] + B[i]
  - o Triad: A[i] = B[i] + scalar * C[i]
- Length of the vectors *N* can vary
- Choose *N* to make each array >> cache size
- Each operation is separately timed (output in MByte/s)
- OpenMP version available

NAS Parallel Benchmarks (NPB)

- Benchmarks derived from computational fluid dynamics (CFD) applications
- NPB 1 consists of five kernels and three pseudo-applications
- Extended to include benchmarks for unstructured adaptive mesh, parallel I/O, multi-zone applications, and computational grids
- Reference implementations of NPB are available
  - – Serial version
  - – Hybrid MPI and OpenMP version
  - – Nested OpenMP version
- Reported measuring unit is MFlop/s
- Different problem classes / problem sizes (S, W, A, .., F)

NPB Suite

- Multi-zone versions
  - Multiple levels of parallism in applications
  - Effectivenes of multi-level and hybrid parallelization paradigms
  - Derived from single-zone pseudo application of NPB 1

  BT-MZ – uneven-size zones within a problem class, increased number of zones as problem class grows

  SP-MZ – even-size zones within a problem class, increased number of zones as problem class grows

  LU-MZ – even-size zones within a problem class, a fixed number of zones for all problem classes

Exercises

- STREAM-Benchmark
  - Use the OpenMP version of the benchmark to analyze the performance of the ARMINIUS system.
  - What are the advantages and disadvantages of processor binding?
  - How to get the most performance out of a single ARMINIUS node?
- NAS Parallel Benchmark
  - Compile and run the OpenMP version of the NPB Multi-Zone on ARMINIUS