


A decorative graphic consisting of a series of blue rectangular blocks of increasing height, arranged in a staircase pattern from left to right.

How to interpret and compare SPEC benchmark results

Swen Boehm, Mayara Gimenes, Robert Henschel,
Veronica G. Vergara Larrea, Junjie Li, Sandra Wienke



spec

<https://www.spec.org/hpg/publications>

<http://pages.iu.edu/~lijunj/sc19>

- Interpretation of Results
- Comparing Results
- Use Cases

- Interpretation of Results
- Comparing Results
- Use Cases

- Beyond creating run rule compliant results, how the results can be used is governed by SPEC
 - The source of the result must be clear (e.g. who produced the results)
 - The date of the result must be clear and correct
 - All SPEC trademarks must be referenced (e.g. SPEC Accel)
 - Metrics must be disclosed. (e.g. SPECaccel_acc, SPECspeed 2017 Integer, SPECspeed 2017 Floating Point)
 - Derived metrics may be used provided the SPEC metric is given. (e.g. score per \$)
 - Basis of comparison is disclosed (if applicable) (e.g. my result is 20% faster than xxx)

- Full fair use rules can be found at: <https://www.spec.org/fairuse.html>

- SPEC Score = geometric mean of all benchmark component ratios
- “Ratio” means: $\frac{\text{runtime on reference machine (given)}}{\text{median(runtime of your measurement)}}$
 - SPEC Score of reference machine is “1”
- Reference machines
 - For SPEC OMP2012
 - Sun Fire X4140, 2xAMD Opteron 2384, 8 cores, 2 chips, 4 cores/chip, 2.7 GHz
 - For SPEC Accel
 - SGI C3108-TY11, NVIDIA Tesla C2070, Intel Xeon E5620 2.4GHz

➤ Higher is better

Results - <http://spec.org/accel/results/accel.html>

OpenACC (26):

Test Sponsor	System Name	Accelerator Name	Results		Energy	
			Base	Peak	Base	Peak
Indiana University	Cray XC30 HTML CSV Text PDF PS Config	Intel Xeon E5-2697 v2	1.18	Not Run	--	--
Indiana University	Cray XK7 HTML CSV Text PDF PS Config	NVIDIA Tesla K20	1.71	Not Run	--	--
Indiana University	Cray XK7 HTML CSV Text PDF PS Config	NVIDIA Tesla K20	1.78	Not Run	--	--
Indiana University	Cray XK7 HTML CSV Text PDF PS Config	NVIDIA Tesla K20	2.00	Not Run	--	--
Indiana University	Cray XK7 HTML CSV Text PDF PS Config	NVIDIA Tesla K20	2.01	Not Run	--	--
Indiana University	Cray XK7 HTML CSV Text PDF PS Config	NVIDIA Tesla K20	2.07	Not Run	--	--
Indiana University	HP Z820 Workstation HTML CSV Text PDF PS Config	Intel Xeon E5-2640 v2	0.662	Not Run	1.10	--


SPEC Score


graphical representation

config file

NVIDIA Corporation	IBM Power Systems AC922 for High Performance Computing (8335-GTH) HTML CSV Text PDF PS Config	Tesla V100	11.9	11.9	--	--
Test Sponsor	System Name	Accelerator Name	Results		Energy	
			Base	Peak	Base	Peak
Oak Ridge National Laboratory	IBM POWER8 S822LC HTML CSV Text PDF PS Config	NVIDIA Tesla P100	8.25	Not Run	--	--
Oak Ridge National Laboratory	Cray XK7 system HTML CSV Text PDF PS Config	NVIDIA Tesla K20X	2.26	Not Run	--	--

as of 11/2019

- Lets take a look at **base vs. peak**: 
- <https://www.spec.org/accel/results/res2017q3/accel-20170726-00092.html>
- or go to Accel OpenMP results and search the entries for “LADMP00AP” (“system”), pick the result (“html”) from “Technische Universitaet Dresden”

- Lets take a look at **energy**: 
- <http://spec.org/accel/results/res2017q2/accel-20170515-00073.html>
- or go to Accel OpenMP results and search the page for “Pedestal”, pick the second result (“html”)

Note: This result is from SPEC Accel v1.1. Current version is v1.2. For demonstration purposes, this is fine!



SPEC® ACCEL™ OMP Result

Copyright 2015-2017 Standard Performance Evaluation Corporation

Intel (Test Sponsor: Technische Universitaet Dresden)
Intel Xeon Phi 7210
Intel Server System LADMP00AP Family (Xeon Phi
7210, 1.3 GHz, 64 cores, 4 threads)


Metircs

SPECaccl_omp_base = 4.39

SPECaccl_omp_peak = 6.08

ACCEL license: 37A
Test sponsor: Technische Universitaet Dresden
Tested by: Technische Universitaet Dresden

Test date: Jul-2017
Hardware Availability: Jun-2016
Software Availability: Dec-2016



Hardware

SPEC Score: base vs. peak
(backup)

Maximum:
FPU: None
CPU(s) enabled: 64 cores, 1 chip, 64 cores/chip, 4 threads/core
CPU(s) orderable: 1 chip
Primary Cache: 32 KB I + 32 KB D on chip per core
Secondary Cache: 1 MB I+D on chip per 2 cores
L3 Cache: 16 GB I+D on chip per chip
Other Cache: None
Memory: 96 GB (6 x 16 GB 2Rx4 PC4-2400T-R, running at 1066 MHz)
Disk Subsystem: 275 GB INTEL SSDSC2BB30
Other Hardware: --

Accelerator

Model: Xeon Phi 7210
Vendor: Intel
Model Name: Intel Xeon Phi 7210
Type of Accel: CPU
Accel Connection: N/A
Does Accel Use ECC: yes
Accel Description: Intel Xeon Phi 7210, SMT ON, Turbo ON
Cluster Mode: Quadrant, Memory Mode: Cache
Accel Driver:
Software
Operating System: CentOS Linux release 7.3 3.10.0-514.21.2.el7.x86_64
Compiler: Intel Compiler C/C++/Fortran Version 17.0.1 20161005
File System: ext4
System State: Run level 3 (user-level)
Other Software: FFTW 3.3.6pl1

Results Table

Benchmark	Base						Peak					
	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
503.postencil	43.9	2.48	44.1	2.47	43.3	2.51	35.9	3.04	38.6	2.82	36.0	3.03
504.polbm	23.1	5.27	22.6	5.40	22.5	5.42	21.6	5.65	22.1	5.52	21.5	5.67
514.pomriq	294	2.11	293	2.12	291	2.13	233	2.67	230	2.70	230	2.70
550.pmd	77.9	3.09	77.5	3.11	77.4	3.12	55.0	4.38	55.1	4.38	54.9	4.39
551.ppaln	4660	0.117	4654	0.117	4645	0.117	251	2.17	251	2.17	252	2.16
552.pep	97.4	2.37	97.4	2.37	97.5	2.37	61.6	3.75	61.4	3.76	61.5	3.76
553.pclvrleaf	103	11.1	102	11.2	103	11.1	113	10.1	113	10.2	113	10.1
554.pcg	210	1.59	211	1.58	212	1.57	195	1.71	195	1.70	195	1.70
555.pseismic	50.6	5.57	51.2	5.51	51.3	5.50	52.2	5.41	52.3	5.40	52.3	5.39
556.psp	42.5	19.2	42.5	19.3	42.8	19.1	38.8	21.1	38.4	21.3	38.6	21.2
557.pcsp	56.1	15.3	55.6	15.4	55.9	15.4	46.3	18.5	44.4	19.3	45.3	19.0
559.pmniGhost	68.0	5.84	67.9	5.85	68.4	5.80	68.4	5.81	68.3	5.81	68.7	5.78
560.pilbdc	76.6	8.53	76.8	8.50	77.7	8.41	68.1	9.58	67.8	9.63	68.8	9.49
563.pswim	32.8	4.85	32.9	4.83	32.8	4.85	29.4	5.41	29.1	5.46	29.2	5.44
570.pbt	30.8	25.3	30.7	25.4	30.9	25.2	25.5	30.6	25.5	30.6	25.6	30.5

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



SPEC® ACCEL™ OMP Result

Copyright 2015-2017 Standard Performance Evaluation Corporation

Intel (Test Sponsor: Technische Universitaet Dresden)
Intel Xeon Phi 7210
Intel Server System LADMP00AP Family (Xeon Phi 7210, 1.3 GHz, 64 cores, 4 threads)


SPECaccel_omp_base = 4.39

SPECaccel_omp_peak = 6.08

SPEC Score: base vs. peak
(backup)

ACCEL license: 37A
Test sponsor: Technische Universitaet Dresden
Tested by: Technische Universitaet Dresden

Test date: Jul-2017
Hardware Availability: Jun-2016
Software Availability: Dec-2016



Hardware

Maximum: None
FPU: None
CPU(s) enabled: 64 cores, 1 chip, 64 cores/chip, 4 threads/core
CPU(s) orderable: 1 chip
Primary Cache: 32 KB I + 32 KB D on chip per core
Secondary Cache: 1 MB I+D on chip per 2 cores
L3 Cache: 16 GB I+D on chip per chip
Other Cache: None
Memory: 96 GB (6 x 16 GB 2Rx4 PC4-2400T-R, running at 1066 MHz)
Disk Subsystem: 275 GB INTEL SSDSC2BB30
Other Hardware: --

Accelerator

Model: Xeon Phi 7210
Vendor: Intel
Model Name: Intel Xeon Phi 7210
Type of Accel: CPU
Accel Connection: N/A
Does Accel Use ECC: yes
Accel Description: Intel Xeon Phi 7210, SMT ON, Turbo ON
Cluster Mode: Quadrant, Memory Mode: Cache
Accel Driver:
Operating System: CentOS Linux release 7.3 3.10.0-514.21.2.el7.x86_64
Compiler: Intel Compiler C/C++/Fortran Version 17.0.1 20161005
File System: ext4
System State: Run level 3 (user-level)
Other Software: FFTW 3.3.6pl1

Software

Results Table

Benchmark	Base						Peak					
	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
503.postencil	43.9	2.48	44.1	2.47	43.3	2.51	35.9	3.04	38.6	2.82	36.0	3.03
504.polbm	23.1	5.27	22.6	5.40	22.5	5.42	21.6	5.65	22.1	5.52	21.5	5.67
514.pomriq	294	2.11	293	2.12	291	2.13	233	2.67	230	2.70	230	2.70
550.pmd	77.9	3.09	77.5	3.11	77.4	3.12	55.0	4.38	55.1	4.38	54.9	4.39
551.ppaln	4660	0.117	4654	0.117	4645	0.117	251	2.17	251	2.17	252	2.16
552.pep	97.4	2.37	97.4	2.37	97.5	2.37	61.6	3.75	61.4	3.76	61.5	3.76
553.pclvrleaf	103	11.1	102	11.2	103	11.1	113	10.1	113	10.2	113	10.1
554.pcg	210	1.59	211	1.58	212	1.57	195	1.71	195	1.70	195	1.70
555.pseismic	50.6	5.57	51.2	5.51	51.3	5.50	52.2	5.41	52.3	5.40	52.3	5.39
556.psp	42.5	19.2	42.5	19.3	42.8	19.1	38.8	21.1	38.4	21.3	38.6	21.2
557.pcsp	56.1	15.3	55.6	15.4	55.9	15.4	46.3	18.5	44.4	19.3	45.3	19.0
559.pmniGhost	68.0	5.84	67.9	5.85	68.4	5.80	68.4	5.81	68.3	5.81	68.7	5.78
560.pilbdc	76.6	8.53	76.8	8.50	77.7	8.41	68.1	9.58	67.8	9.63	68.8	9.49
563.pswim	32.8	4.85	32.9	4.83	32.8	4.85	29.4	5.41	29.1	5.46	29.2	5.44
570.pbt	30.8	25.3	30.7	25.4	30.9	25.2	25.5	30.6	25.5	30.6	25.6	30.5

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

Base Compiler Invocation

C benchmarks:

icc

Fortran benchmarks:

ifort

Benchmarks using both Fortran and C:

icc ifort

Base Portability Flags

503.postencil: -DSPEC USE INNER SIMD
504.polbm: -DSPEC USE INNER SIMD
514.pomriq: -DSPEC USE INNER SIMD
550.pmd: -DSPEC USE INNER SIMD -80
551.ppalm: -DSPEC USE INNER SIMD
552.pcp: -DSPEC USE INNER SIMD
553.pclvleaf: -DSPEC USE INNER SIMD
554.pcg: -DSPEC USE INNER SIMD
555.pseismic: -DSPEC USE INNER SIMD
556.psp: -DSPEC USE INNER SIMD
557.pcsp: -DSPEC USE INNER SIMD
559.pmniGhost: -DSPEC USE INNER SIMD -nofor_main
560.pilbdc: -DSPEC USE INNER SIMD
563.pswim: -DSPEC USE INNER SIMD
570.pbt: -DSPEC USE INNER SIMD

Base Optimization Flags

C benchmarks:

-O3 -g -qopenmp -xMIC-AVX512 -qopenmp-offload=host

Fortran benchmarks:

-O3 -g -qopenmp -xMIC-AVX512 -qopenmp-offload=host

Benchmarks using both Fortran and C:

-O3 -g -qopenmp -xMIC-AVX512 -qopenmp-offload=host

all the same

SPEC Score: base vs. peak
(backup)

Peak Compiler Invocation

C benchmarks:

and C:

Peak Portability Flags

503.postencil: -DSPEC USE INNER SIMD
504.polbm: -DSPEC USE INNER SIMD
514.pomriq: -DSPEC USE INNER SIMD
550.pmd: -DSPEC USE INNER SIMD -80
551.ppalm: -DSPEC USE INNER SIMD -DSPEC_HOST_FFTW3
552.pcp: -DSPEC USE INNER SIMD
553.pclvleaf: -DSPEC USE INNER SIMD
554.pcg: -DSPEC USE INNER SIMD
555.pseismic: -DSPEC USE INNER SIMD
556.psp: -DSPEC USE INNER SIMD
557.pcsp: -DSPEC USE INNER SIMD
559.pmniGhost: -DSPEC USE INNER SIMD -nofor_main
560.pilbdc: -DSPEC USE INNER SIMD
563.pswim: -DSPEC USE INNER SIMD
570.pbt: -DSPEC USE INNER SIMD

Peak Optimization Flags

C benchmarks:

503.postencil: -O3 -xCORE-AVX2 -g -qopenmp -qopenmp-offload=host -qopt-prefetch=3
504.polbm: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host -qopt-prefetch=5
514.pomriq: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host -qopt-prefetch=2
552.pcp: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host -qopt-streaming-stores always
554.pcg: -O3 -xCORE-AVX2 -g -qopenmp -qopenmp-offload=host -qopt-prefetch=2 -qopt-streaming-stores always
557.pcsp: Same as 504.polbm
570.pbt: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host

may be different per benchmark component

-qopenmp -qopenmp-offload=host -qopt-prefetch=3 -no-prec-div
-qopenmp -qopenmp-offload=host -no-prec-sqrt
-L/sw/taurus/libraries/ftw/3.3.6p11-gcc5.3-intelmpi5.1/lib
-L/sw/taurus/libraries/ftw/3.3.6p11-gcc5.3-intelmpi5.1/lib

555.pseismic: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host
556.psp: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host -qopt-prefetch=2
560.pilbdc: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host -qopt-prefetch=3
563.pswim: Same as 555.pseismic


Benchmarks using both Fortran and C:

553.pclvleaf: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host -qopt-streaming-stores always
559.pmniGhost: -O3 -xMIC-AVX512 -g -qopenmp -qopenmp-offload=host -qopt-prefetch=3 -qopt-streaming-stores always

shing SPEC results

SPEC Score: base vs. peak (

SPEC Score: base vs. peak
(*backup*)



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

Submit Notes

The config file option 'submit' was used.
submit = numactl -p 0 \$command

Platform Notes

Sysinfo program /tmp/spec-accel/1.2/Docs/sysinfo
\$Rev: 6965 \$ \$Date:: 2015-04-21 # \$ c05a7f14b1b1765e3fe1df68447e8a35
running on tauruskn128.taurus.hrsk.tu-dresden.de Mon Jul 24 10:45:36 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/accel/Docs/config.html#sysinfo>

From /proc/cpuinfo
model name : Intel(R) Xeon Phi(TM) CPU 7210 @ 1.30GHz
1 "physical id"s (chips)
256 "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 : 64
siblings : 256
physical 0: cores 0 1 2 3 6 7 10 11 12 13 14 15 18 19 20 21 22 23 24 25 26
27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51
52 53 56 57 58 59 60 61 62 63 64 65 68 69 70 71 72 73
cache size : 1024 KB

From /proc/meminfo
MemTotal: 98707216 kB
HugePages_Total: 0
Hugepagesize: 2048 kB

/usr/bin/lsb_release -d
CentOS Linux release 7.3.1611 (Core)

```
From /etc/*release* /etc/*version*
centos-release: CentOS Linux release 7.3.1611 (Core)
centos-release-upstream: Derived from Red Hat Enterprise Linux 7.3 (Source)
os-release:
  NAME="CentOS Linux"
  VERSION="7 (Core)"
  ID="centos"
  ID_LIKE="rhel fedora"
  VERSION_ID="7"
  PRETTY_NAME="CentOS Linux 7 (Core)"
  ANSI_COLOR="0;31"
  CPE_NAME="cpe:/o:centos:centos:7"
redhat-release: CentOS Linux release 7.3.1611 (Core)
system-release: CentOS Linux release 7.3.1611 (Core)
system-release-cpe: cpe:/o:centos:centos:7
```

```
uname -a:
Linux tauruskn128.taurus.hrsk.tu-dresden.de 3.10.0-514.21.2.el7.x86_64 #1 SMP
Tue Jun 20 12:24:47 UTC 2017 x86_64 x86_64 x86_64 GNU/Linux
```

run-level 3 Jun 30 13:46

SPEC is set to: /tmp/spec-accel/1.2
Filesystem Type Size Used Avail Use% Mounted on
/dev/sda1 ext4 275G 6.5G 255G 3% /
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.

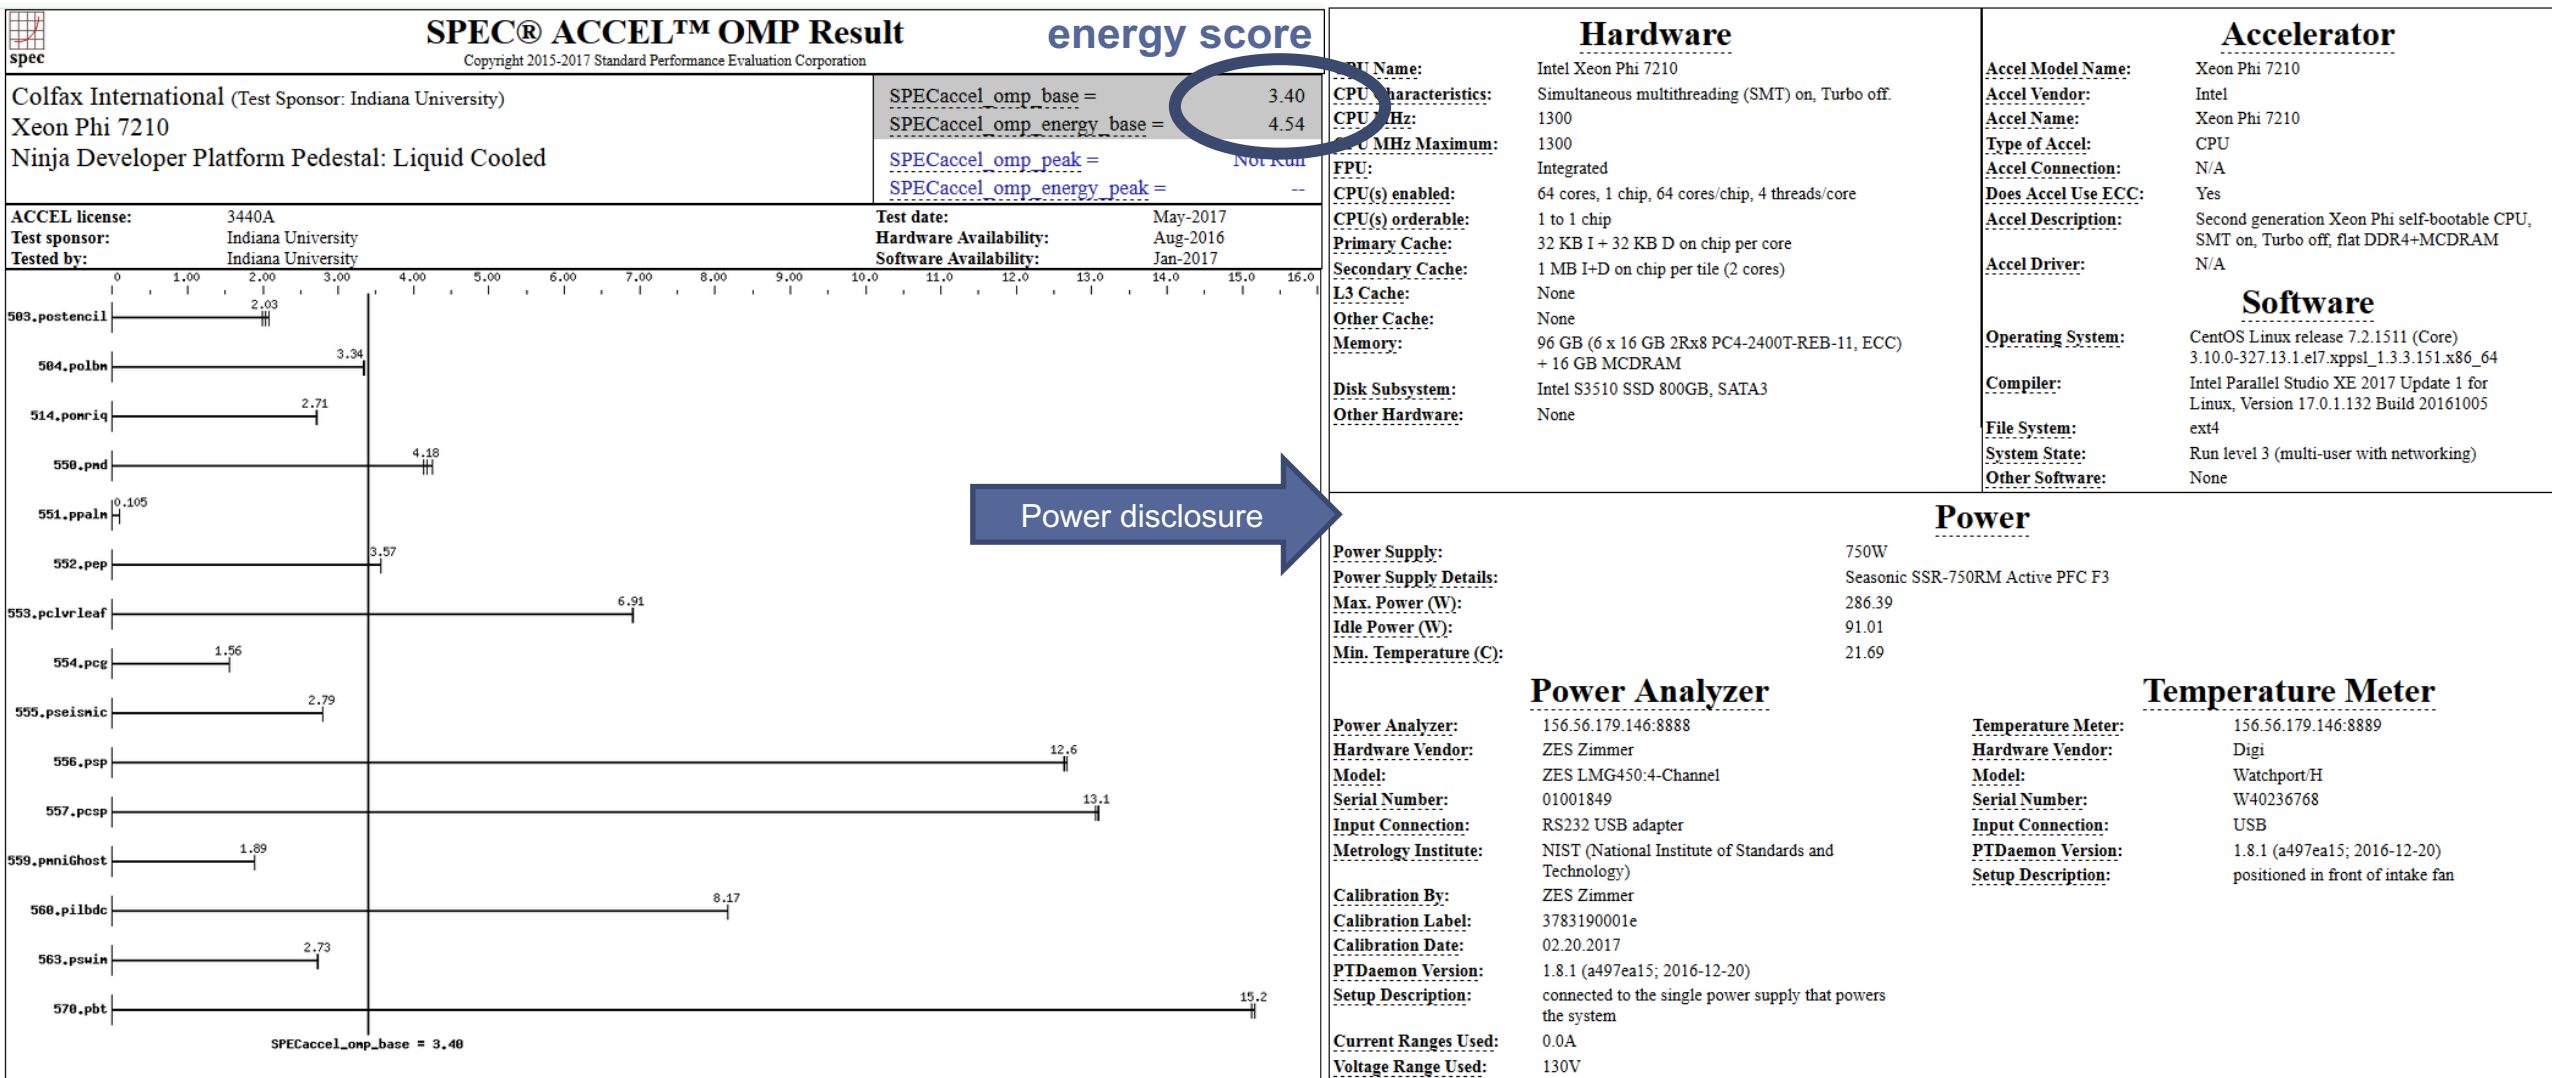

(End of data from sysinfo program)

General Notes

Used Environment Variables:


ENV_KMP_AFFINITY=compact,0 - assign OpenMP Threads continuously
ENV_OMP_NUM_THREADS=128 - limits number of Threads to be started to 128
ENV_KMP_HW_SUBSET=1S,64C,2T - control Thread distribution accross sockets, cores and hw threads
ENV_FORT_BUFFERED=true - enables buffered I/O for Fortran
ENV_OMP_DYNAMIC - Enable or disable the dynamic adjustment of the number of threads within a task
ENV_KMP_LIBRARY - Selects the OpenMP runtime library throughput. The options for the variable value
ENV_KMP_BLOCKTIME - Sets the time, in milliseconds, that a thread should wait, after completing the


SPEC Score: energy (backup)





- Interpretation of Results
- **Comparing results**
- Use Cases

Comparing Results





- For both SPEC OpenMP2012 and SPEC MPI2007, most results are part of a scalability or comparison study.
 - Increasing MPI ranks
 - Testing different compilers
 - ...
- The next charts are created from published results!
- Take a look online.... 
- Note that you cannot compare results between versions or between data sets!

System and Interconnect Comparison




- Cray XC30:
 - 2x Xeon E5-2697 v2 (24C)
 - Cray Aries interconnect
 - Cray MPI
 - Dragonfly
- NEC HPC1812Rg-2 node:
 - 2x Xeon E5-2650 v4 (24C)
 - Intel Omni-Path interconnect
 - Intel MPI
 - Fat tree
- HPE SGI 8600 node:
 - 2x Xeon Gold 6148 (40C)
 - Dual-rail InfiniBand 4X EDR
 - HPE SGI MPI
 - Enhanced hypercube


Same Programming Model on Different Hardware




Same Programming Model on Different Hardware




Devices used in SPEC ACCEL OpenACC Submissions




Compiler Evolution – PGI and Cray OpenACC




SPEC ACCEL OpenACC on IU Cray XK7 NVIDIA TESLA K20



Comparing Different MPI Libraries and OSes







**SPEC MPI2007 Medium
IBM iDataPlex (Intel Xeon L5420)**



SPEC OMP2012

Performance and Energy





Comparing Partial Results – Academic Fair Use

- E.g., Cray and IBM compilers support OpenMP 4.5 offload to GPUs. This shows the potential of the Cray compiler but at that time only 6 of 15 benchmarks worked!
- Non-published results are “**SPEC estimate**”



- Rpeak

- KNL-7210: 2.60 Tflops
 - K20: 1.17 TFlops
- Ratio: 2.2x

	SPEC Score (<i>estimate</i>)			Speedup	
Benchmarks	KNL(MCDRAM) Intel	KNL(DDR4) Intel	K20 Cray	KNL(MCDRAM) vs K20	KNL(DDR4) vs K20
503.postencil	1.99	0.700	1.26	1.6x	0.6x
504.polbm	3.42	0.754	0.898	3.8x	0.8x
514.pomriq	2.71	2.72	1.11	2.4x	2.4x
555.pseismic	2.83	1.06	1.43	2.0x	0.7x
560.pilbdc	8.43	1.97	4.61	1.8x	0.4x
570.pbt	27.4	20.2	18.2	1.5x	1.1x
Geometric Mean				2.1x	0.8x

And Now for Something Totally Different

- HPL vs. SPEC OpenMP 2012
- How much information is publically available about Top500 results?




Comparing Results – Advanced Search

- Let's use the search function on the home page:
 - Advanced Search:
 - <https://www.spec.org/cgi-bin/osgresults?conf=omp2012>
 - Indiana University, Power8 results, showing compiler, sort by compiler first, thread count second.




- Yes, you could have done that using copy and paste, but imagine doing this with SPEC CPU2006 results!


Comparing Results – Advanced Search



SPEC OpenMP2012 on IBM S822LC (Power 8 + NVLink)



Comparing Results



- Dump all Records as CSV

- <https://www.spec.org/cgi-bin/osgresults?conf=omp2012>

- Interpretation of Results
- Comparing results
- **Use Cases**

- System, accelerator and software vendors
- Application developers
- Users and HPC centers
- Researchers
- HPC tool developers


- Marketing
- Drive benchmark development
 - To utilize state of the art hardware/software features
- Internal validation suite
 - Compiler
 - OMP / MPI runtime libraries
- Prepare for RFPs (procurement)

Use Cases – Application Developers

- Include their application in the benchmark suite
 - See results on a lot of different systems.
- Compare hardware and software stack
 - Compilers
 - Parallel runtimes
 - Different versions of processors
 - Different interconnects

- Include the benchmarks in the RFP (procurement) process
 - ORNL used the SPEC benchmark suites (OMP, ACCEL) during Summit acceptance
 - Test compliance of vendor and open source compilers
- Use them for performance regression testing
 - Hardware
 - Software
- System configuration and tuning
- Power consumption

Use Cases – Researchers





- Scalability studies
- Novel implementations of parallel runtime libraries
- Detailed power consumption studies
- Comparison of parallel programming paradigms



Use Cases – HPC Tool Developers

- MUST
 - Implements MPI runtime correctness analysis and reports deadlocks, mismatches in types or collective arguments and scales to more than 16k MPI ranks.
 - SPEC MPI L2007 v2 (ref) up to 2k ranks used to
 - Evaluation of general tool runtime overhead, i.e., (runtime with tool) / (runtime without tool)
 - Evaluation of the influence of specific changes in the analysis or tool infrastructure (e.g. guarantee to provide complete results when the application crashes).
 - Publications: <http://www.itc.rwth-aachen.de/go/id/fddi/lidx/1/file/540356>
- Archer / ThreadSanitizer
 - Data race analysis for OpenMP programs
 - SPEC OMP (train) up to 12 threads used to evaluate tool runtime overhead for data race detection.
 - Publication: <http://www.itc.rwth-aachen.de/go/id/fddi/lidx/1/file/706852>
- OMPT Interface of Intel/LLVM OMP runtime
 - OMPT (OpenMP tools interface) implementation in the LLVM/Intel OpenMP runtime
 - Requirement by Intel: negligible overhead in the absence of a OMPT tool
 - SPEC OMP 2012 (ref) was used for evaluation of the overhead of the OMPT implementation and acceptance test of the OMPT implementation.


Use Cases – MUST (MPI Correctness and Deadlocks)



Use Cases – Archer / ThreadSanitizer (Data Race Det)



Thank you!



Questions?

Contact

SPEC Headquarters:

Swen Boehm:

Mayara Gimenes:

Robert Henschel:

Veronica G. Vergara Larrea:

Junjie Li:

Sandra Wienke:

info@spec.org

boehms@ornl.gov

mayara@udel.edu

henschel@iu.edu

vergaravg@ornl.gov

lijunj@iu.edu

wienke@itc.rwth-aachen.de