Using the new SPEChpc 2021 Scientific Application Benchmark Suite

Sunita Chandrasekaran, Robert Henschel, Junjie Li, Verónica G. Melesse Vergara

https://www.spec.org/hpg/publications
https://www.olcf.ornl.gov/sc20-spec-hpg-tutorial/

Presenters

Junjie Li

Principal System Analyst @ Indiana University
HPG Secretary @ SPEC
lijunj@iu.edu

Robert Henschel

Director, Research Software and Solutions @Indiana University
HPG Chair @ SPEC
henschel@iu.edu

Veronica G. Melesse Vergara

Group Leader
User Assistance Pre-production Systems @ ORNL
vergaravg@ornl.gov

Sunita Chandrasekaran

Assistant Professor @ University of Delaware
Dept. of Computer & Information Sciences
schandra@udel.edu

Tutorial Overview

- Overview of SPEC and SPEC HPG (30 min)
- Benchmark selection and porting process (15 min)
- How to use SPEC benchmark results for decision making? (45 min)
- How to get and setup the HPC2020 benchmarks (15 min)
- Break (30 min)
- **How to interpret and publish SPEC benchmark results (20 min)**
- Hands-On (30 min)
- Conclusion and Wrap-Up (10 min)

Tutorial website:
https://www.olcf.ornl.gov/sc20-spec-hpg-tutorial

Why SPEC benchmarks?

spec

- Let's take a quick look at a published SPEC result: https://www.spec.org/mpi2007/results/res2017q4/mpi2007-20171011-00580.html

More details will be discussed in later sections.

- How much information can you obtain for other benchmark results?

- Benchmark reports contain critical **details for reproducibility**.
- Published SPEC results are **peer-reviewed**.
- All benchmarks are based on **real applications.**
- Rich database of published results.

- *You will discover more details in this tutorial*

Results Table

Benchmark	Base							Peak						
	Ranks	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Ranks	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
104.milc	640	15.4	102	14.9	105	**14.9**	**105**							
107.leslie3d	640	34.1	153	33.2	157	**33.4**	**156**							
113.GemsFDTD	640	187	33.8	**186**	**33.8**	186	33.9							
115.fds4	640	23.3	83.9	22.8	85.6	**23.2**	**84.0**							
121.pop2	640	77.5	53.2	**77.5**	**53.3**	77.3	53.4							
122.tachyon	640	31.4	89.0	**31.5**	**88.9**	32.1	87.2							
126.lammps	640	90.3	32.3	89.6	32.5	**89.7**	**32.5**							
127.wrf2	640	29.5	264	30.2	258	**29.6**	**264**							
128.GAPgeofem	640	8.10	255	8.31	249	**8.28**	**249**							
129.tera_tf	640	22.1	125	22.5	123	**22.3**	**124**							
130.socorro	640	30.7	124	**31.1**	**123**	31.8	120							
132.zeusmp2	640	19.8	157	**19.7**	**158**	19.7	158							
137.lu	640	19.1	192	18.9	195	**19.0**	**193**							

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

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

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

Contents

- Interpretation of Results
- Comparing Results
- Use Cases

Contents

- Interpretation of Results
- Comparing Results
- Use Cases

Fair Use

- 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

- SPEC Score = geometric mean of all benchmark component ratios
- “Ratio” : $\frac{\text{runtime on reference machine (given)}}{\text{median(runtime of your measurement)}}$
  - 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 Base	Results Peak	Energy Base	Energy 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	--
⋮						
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 Base**	**Results Peak**	**Energy Base**	**Energy 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	--	--

SPEC Score

graphical representation

config file

as of 11/2019

SPEC Score

- 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 Score: base vs. peak

Metrics

Ratios

Hardware disclosure

Detailed results

SPEC® ACCEL™ OMP Result

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

ACCEL license:	37A	Test date:	Jul-2017
Test sponsor:	Technische Universitaet Dresden	Hardware Availability:	Jun-2016
Tested by:	Technische Universitaet Dresden	Software Availability:	Dec-2016

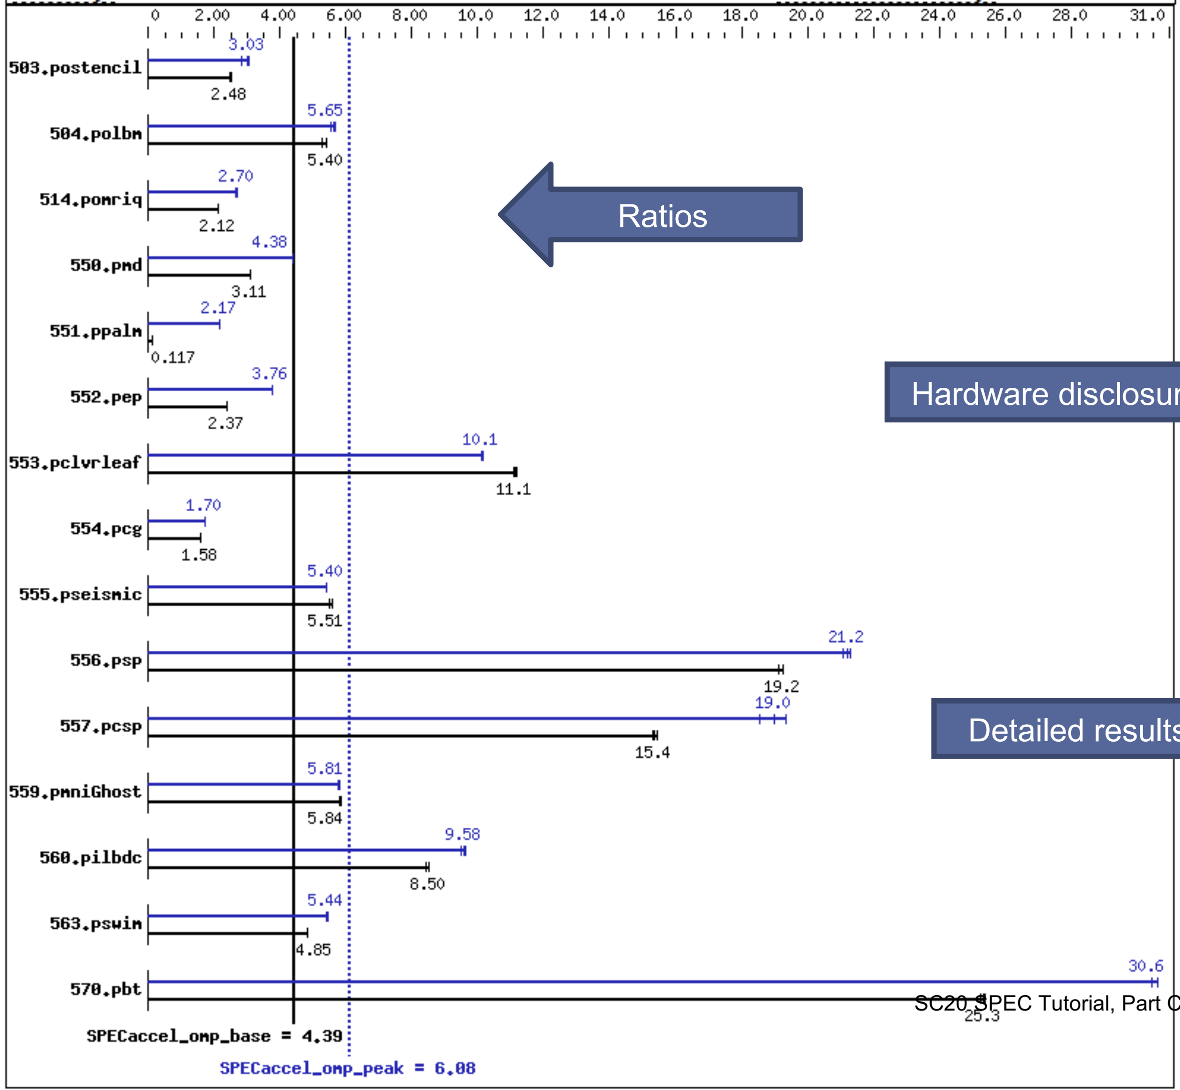

Hardware

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

l Model	Xeon Phi 7210
e:	
l Vendor:	Intel
l Name:	Intel Xeon Phi 7210
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.ppalm	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 Score: base vs. peak

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.pep:	-DSPEC_USE_INNER_SIMD
553.pclvrleaf:	-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

all the same

Fortran benchmarks:

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

Benchmarks using both Fortran and C:

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

Peak Compiler Invocation

C benchmarks:

...nd 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.pep:	-DSPEC_USE_INNER_SIMD
553.pclvrleaf:	-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.pep:	-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
.3.6pl1-gcc5.3-intelmpi5.1/include
-L/sw/taurus/libraries/fftw/3.3.6pl1-gcc5.3-intelmpi5.1/lib

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

Benchmarks using both Fortran and C:

553.pclvrleaf:	-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

SPEC Score: energy (backup)

spec

SPEC® ACCEL™ OMP Result

energy score

Colfax International (Test Sponsor: Indiana University)
Xeon Phi 7210
Ninja Developer Platform Pedestal: Liquid Cooled

SPECaccel_omp_base =	3.40
SPECaccel_omp_energy_base =	4.54
SPECaccel_omp_peak =	Not Run
SPECaccel_omp_energy_peak =	--

ACCEL license:	3440A	Test date:	May-2017
Test sponsor:	Indiana University	Hardware Availability:	Aug-2016
Tested by:	Indiana University	Software Availability:	Jan-2017

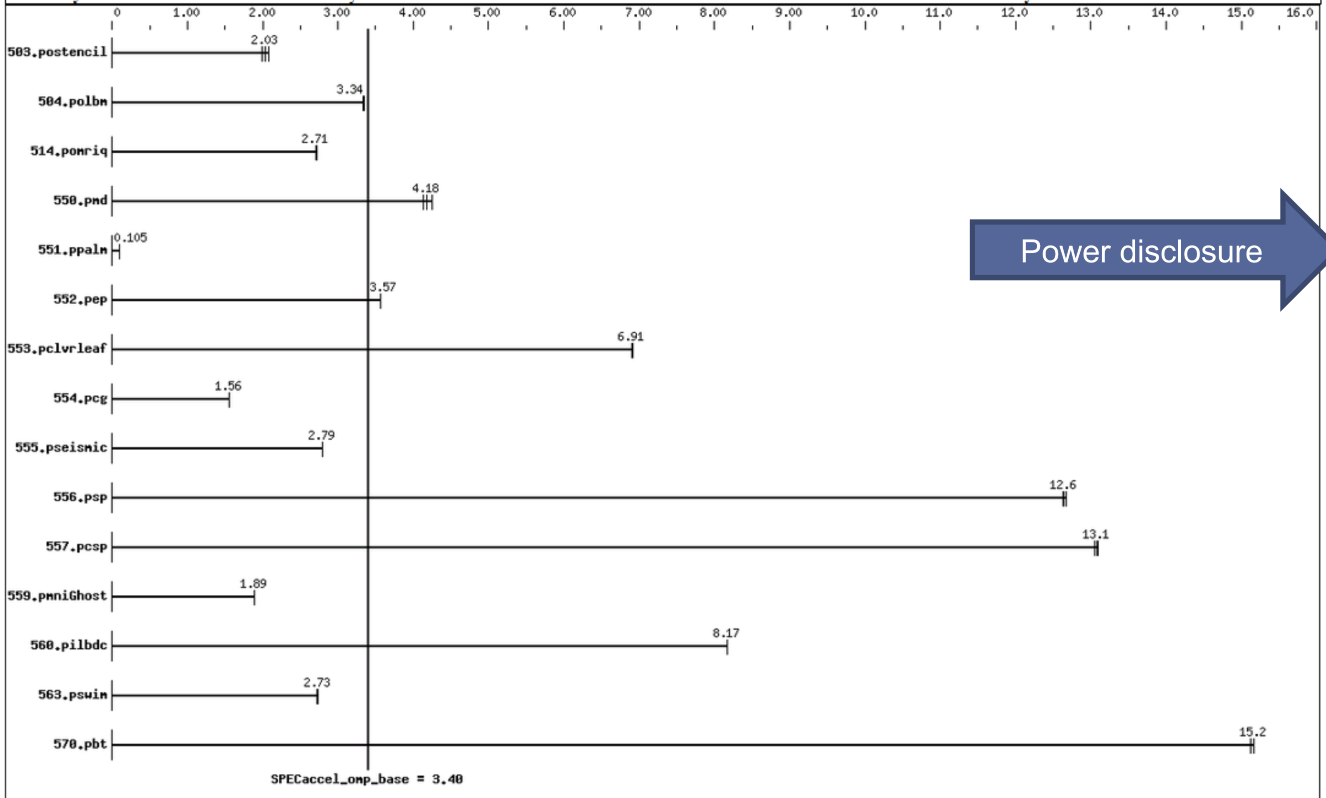

Power disclosure

Hardware

CPU Name:	Intel Xeon Phi 7210
CPU Characteristics:	Simultaneous multithreading (SMT) on, Turbo off.
CPU MHz:	1300
CPU MHz Maximum:	1300
FPU:	Integrated
CPU(s) enabled:	64 cores, 1 chip, 64 cores/chip, 4 threads/core
CPU(s) orderable:	1 to 1 chip
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per tile (2 cores)
L3 Cache:	None
Other Cache:	None
Memory:	96 GB (6 x 16 GB 2Rx8 PC4-2400T-REB-11, ECC) + 16 GB MCDRAM
Disk Subsystem:	Intel S3510 SSD 800GB, SATA3
Other Hardware:	None

Accelerator

Accel Model Name:	Xeon Phi 7210
Accel Vendor:	Intel
Accel Name:	Xeon Phi 7210
Type of Accel:	CPU
Accel Connection:	N/A
Does Accel Use ECC:	Yes
Accel Description:	Second generation Xeon Phi self-bootable CPU, SMT on, Turbo off, flat DDR4+MCDRAM
Accel Driver:	N/A

Software

Operating System:	CentOS Linux release 7.2.1511 (Core) 3.10.0-327.13.1.el7.xppsl_1.3.3.151.x86_64
Compiler:	Intel Parallel Studio XE 2017 Update 1 for Linux, Version 17.0.1.132 Build 20161005
File System:	ext4
System State:	Run level 3 (multi-user with networking)
Other Software:	None

Power

Power Supply:	750W
Power Supply Details:	Seasonic SSR-750RM Active PFC F3
Max. Power (W):	286.39
Idle Power (W):	91.01
Min. Temperature (C):	21.69

Power Analyzer

Power Analyzer:	156.56.179.146:8888
Hardware Vendor:	ZES Zimmer
Model:	ZES LMG450:4-Channel
Serial Number:	01001849
Input Connection:	RS232 USB adapter
Metrology Institute:	NIST (National Institute of Standards and Technology)
Calibration By:	ZES Zimmer
Calibration Label:	3783190001e
Calibration Date:	02.20.2017
PTDaemon Version:	1.8.1 (a497ea15; 2016-12-20)
Setup Description:	connected to the single power supply that powers the system
Current Ranges Used:	0.0A
Voltage Range Used:	130V

Temperature Meter

Temperature Meter:	156.56.179.146:8889
Hardware Vendor:	Digi
Model:	Watchport/H
Serial Number:	W40236768
Input Connection:	USB
PTDaemon Version:	1.8.1 (a497ea15; 2016-12-20)
Setup Description:	positioned in front of intake fan

Contents

- 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 on the SPEC HPG website as well

- Note that you cannot compare results between versions or between data sets!

Same Programming Model on Different Hardware

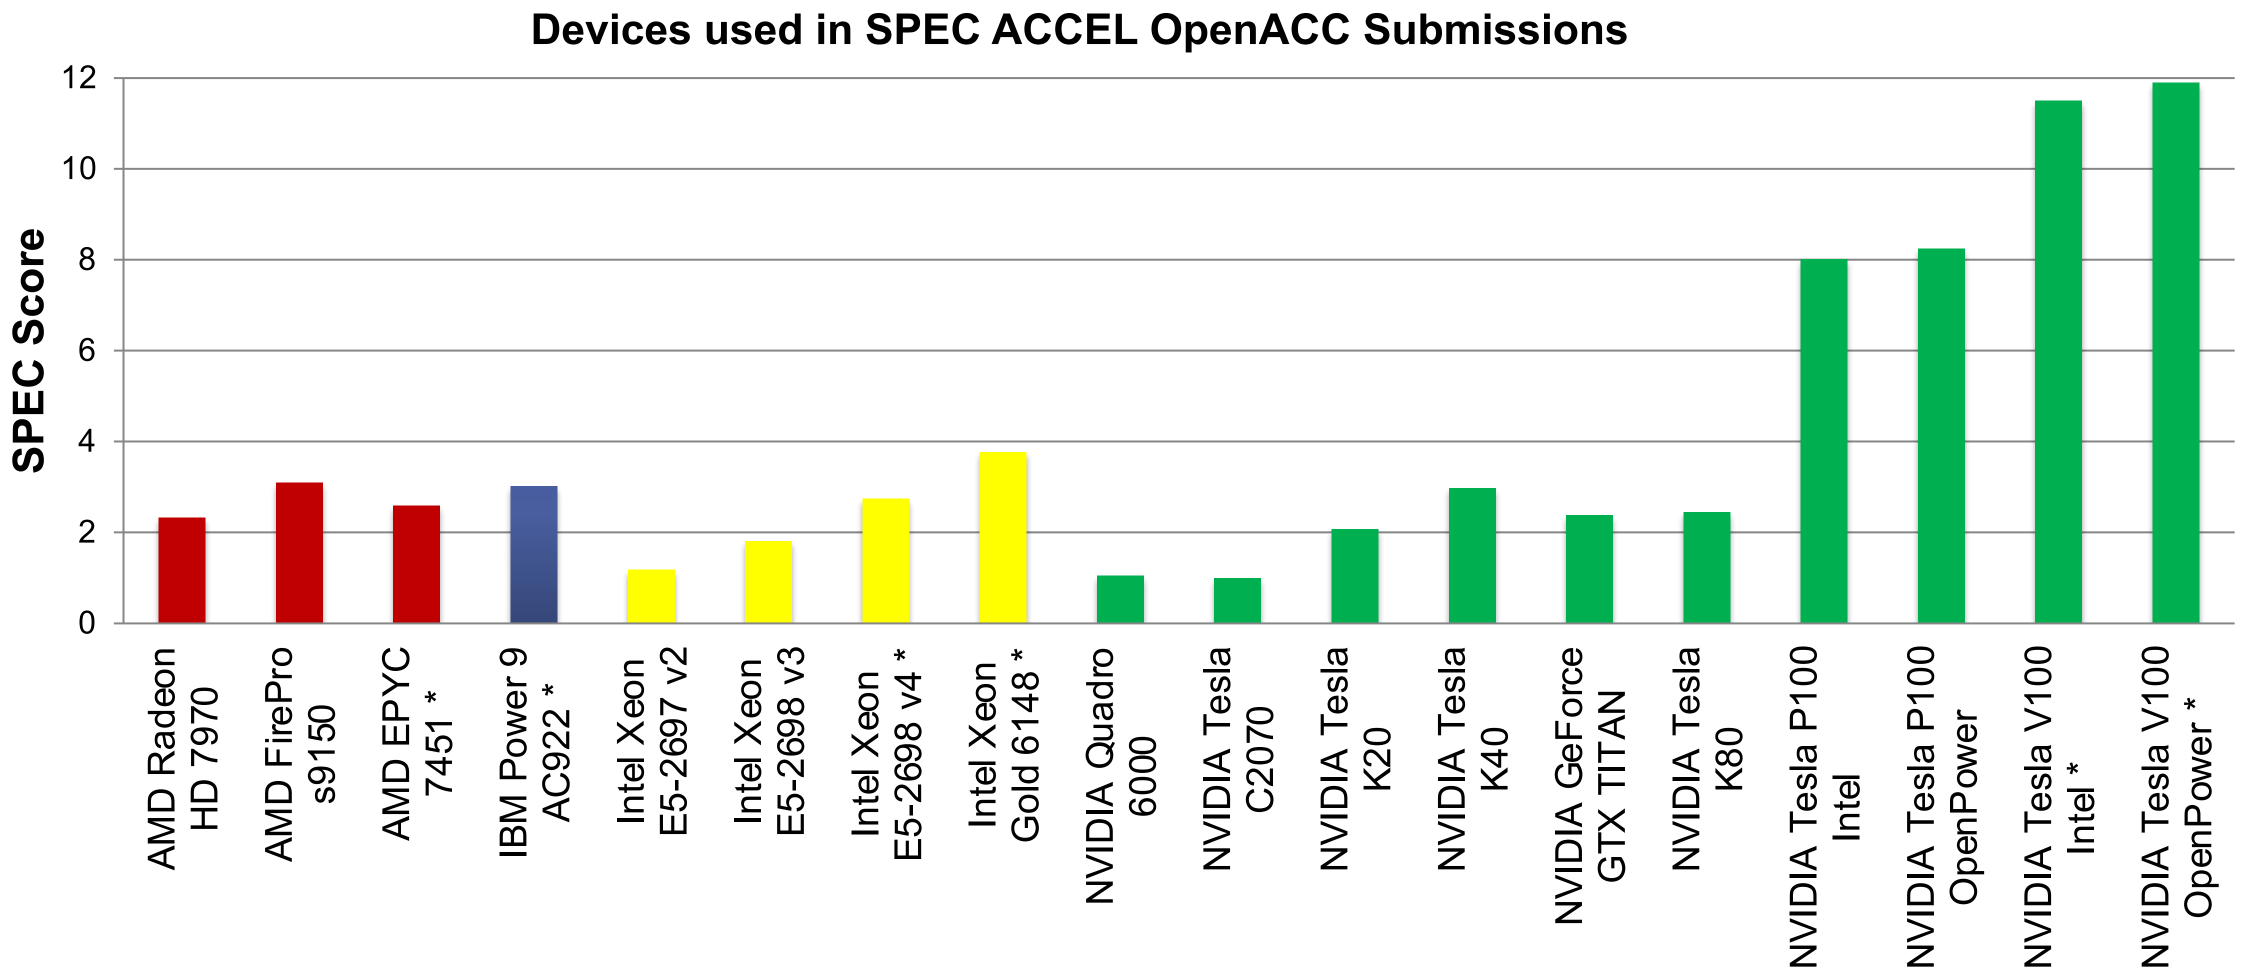

* Results from Version 1.2 of the SPEC ACCEL benchmark while all other results are from version 1.1.

Compiler Evolution – PGI and Cray OpenACC

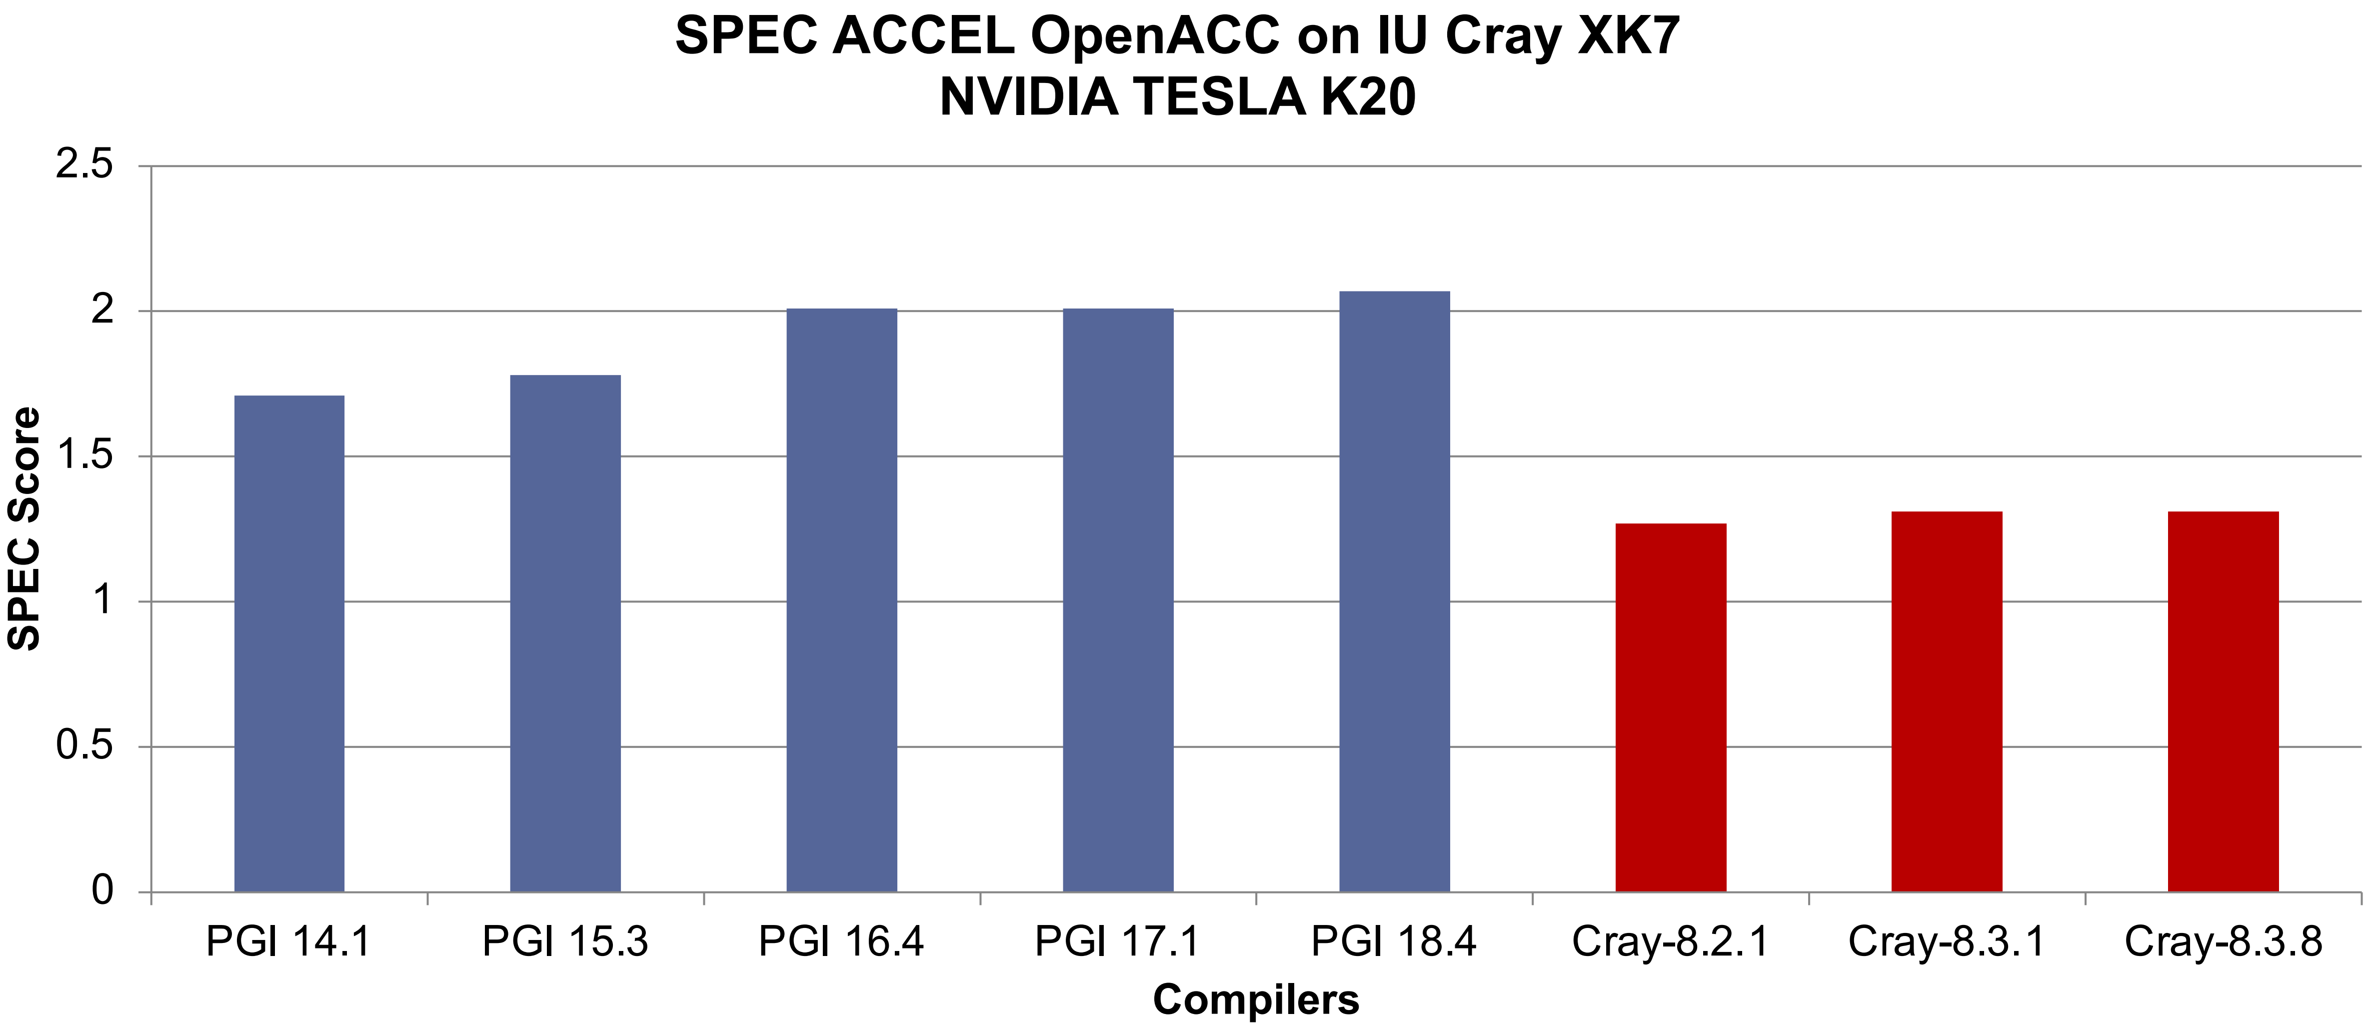

SPEC OMP2012
Performance and Energy

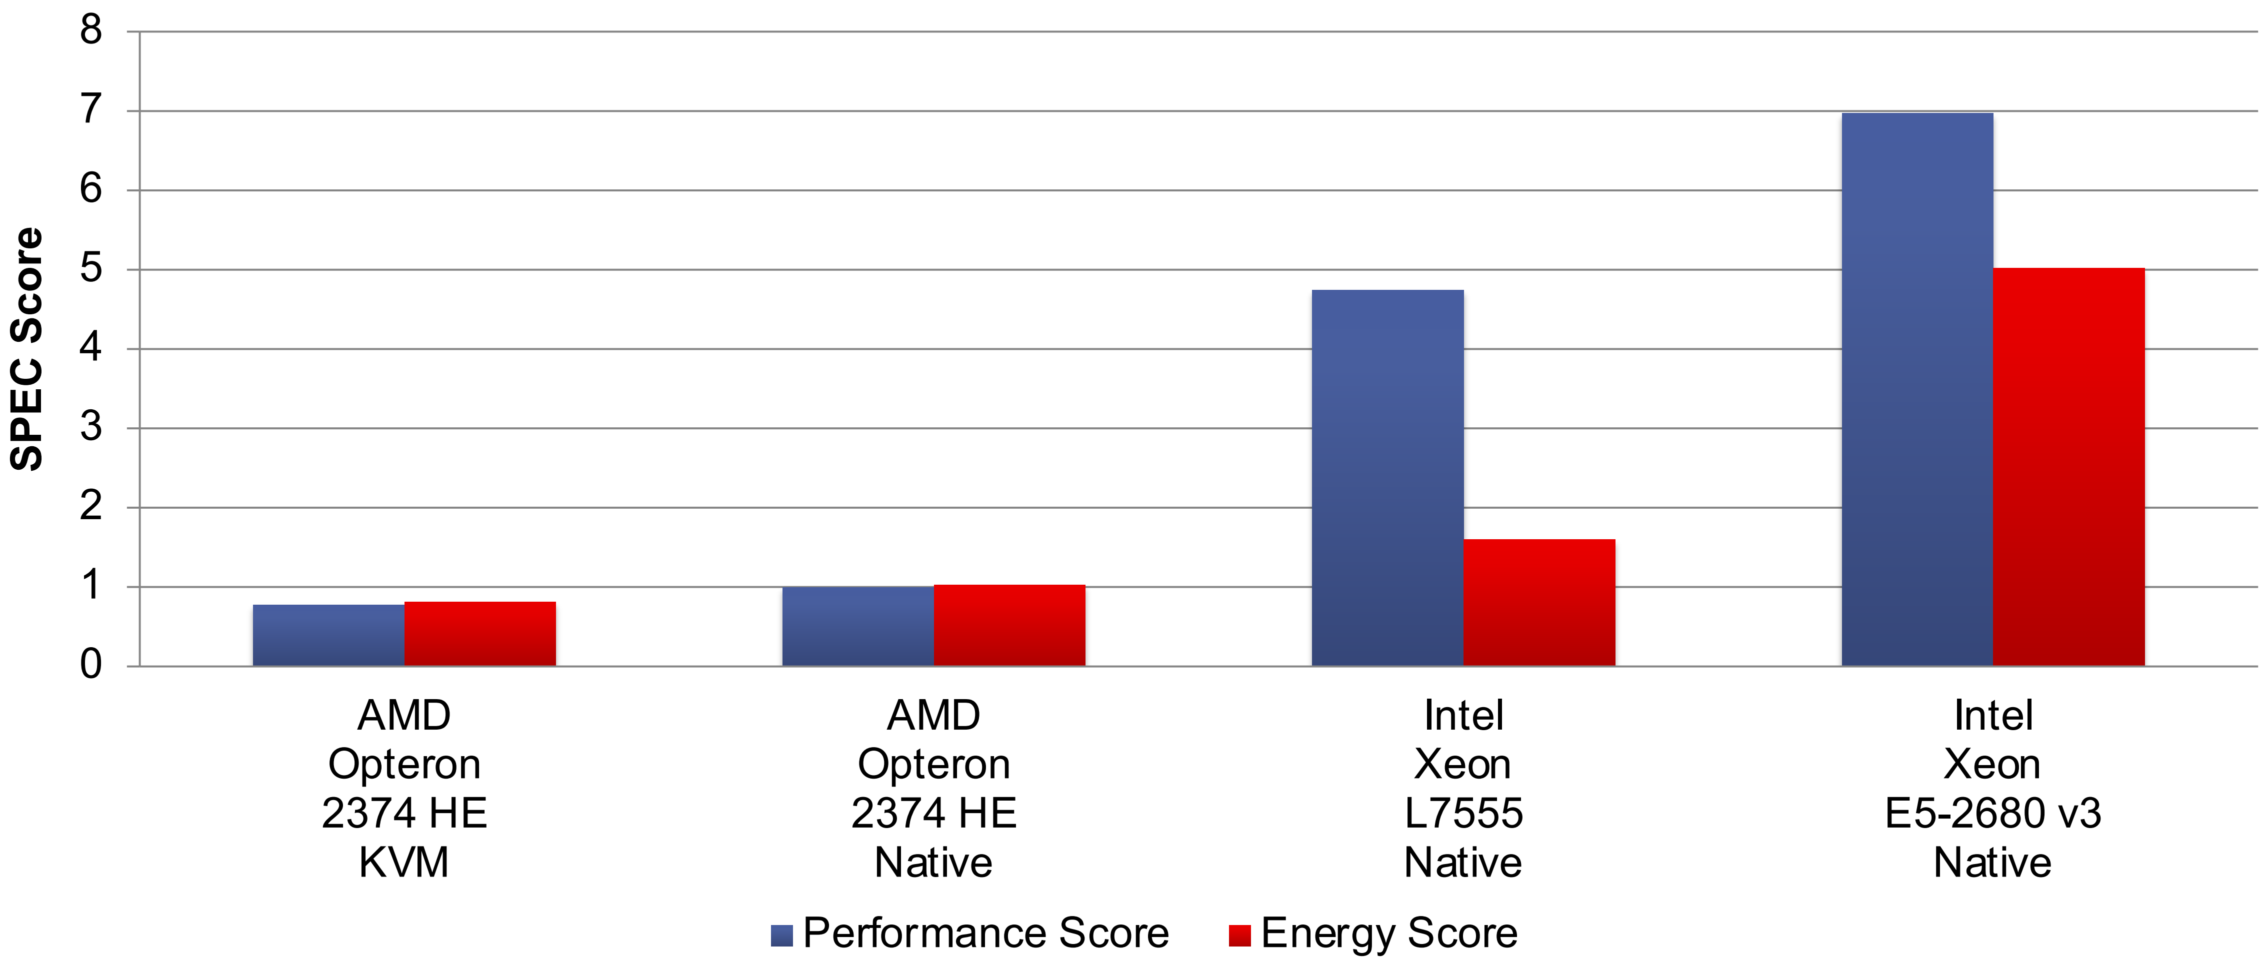

SPEC ACCEL OpenCL
The effect of ECC (Using results for NVIDIA K40c, base)

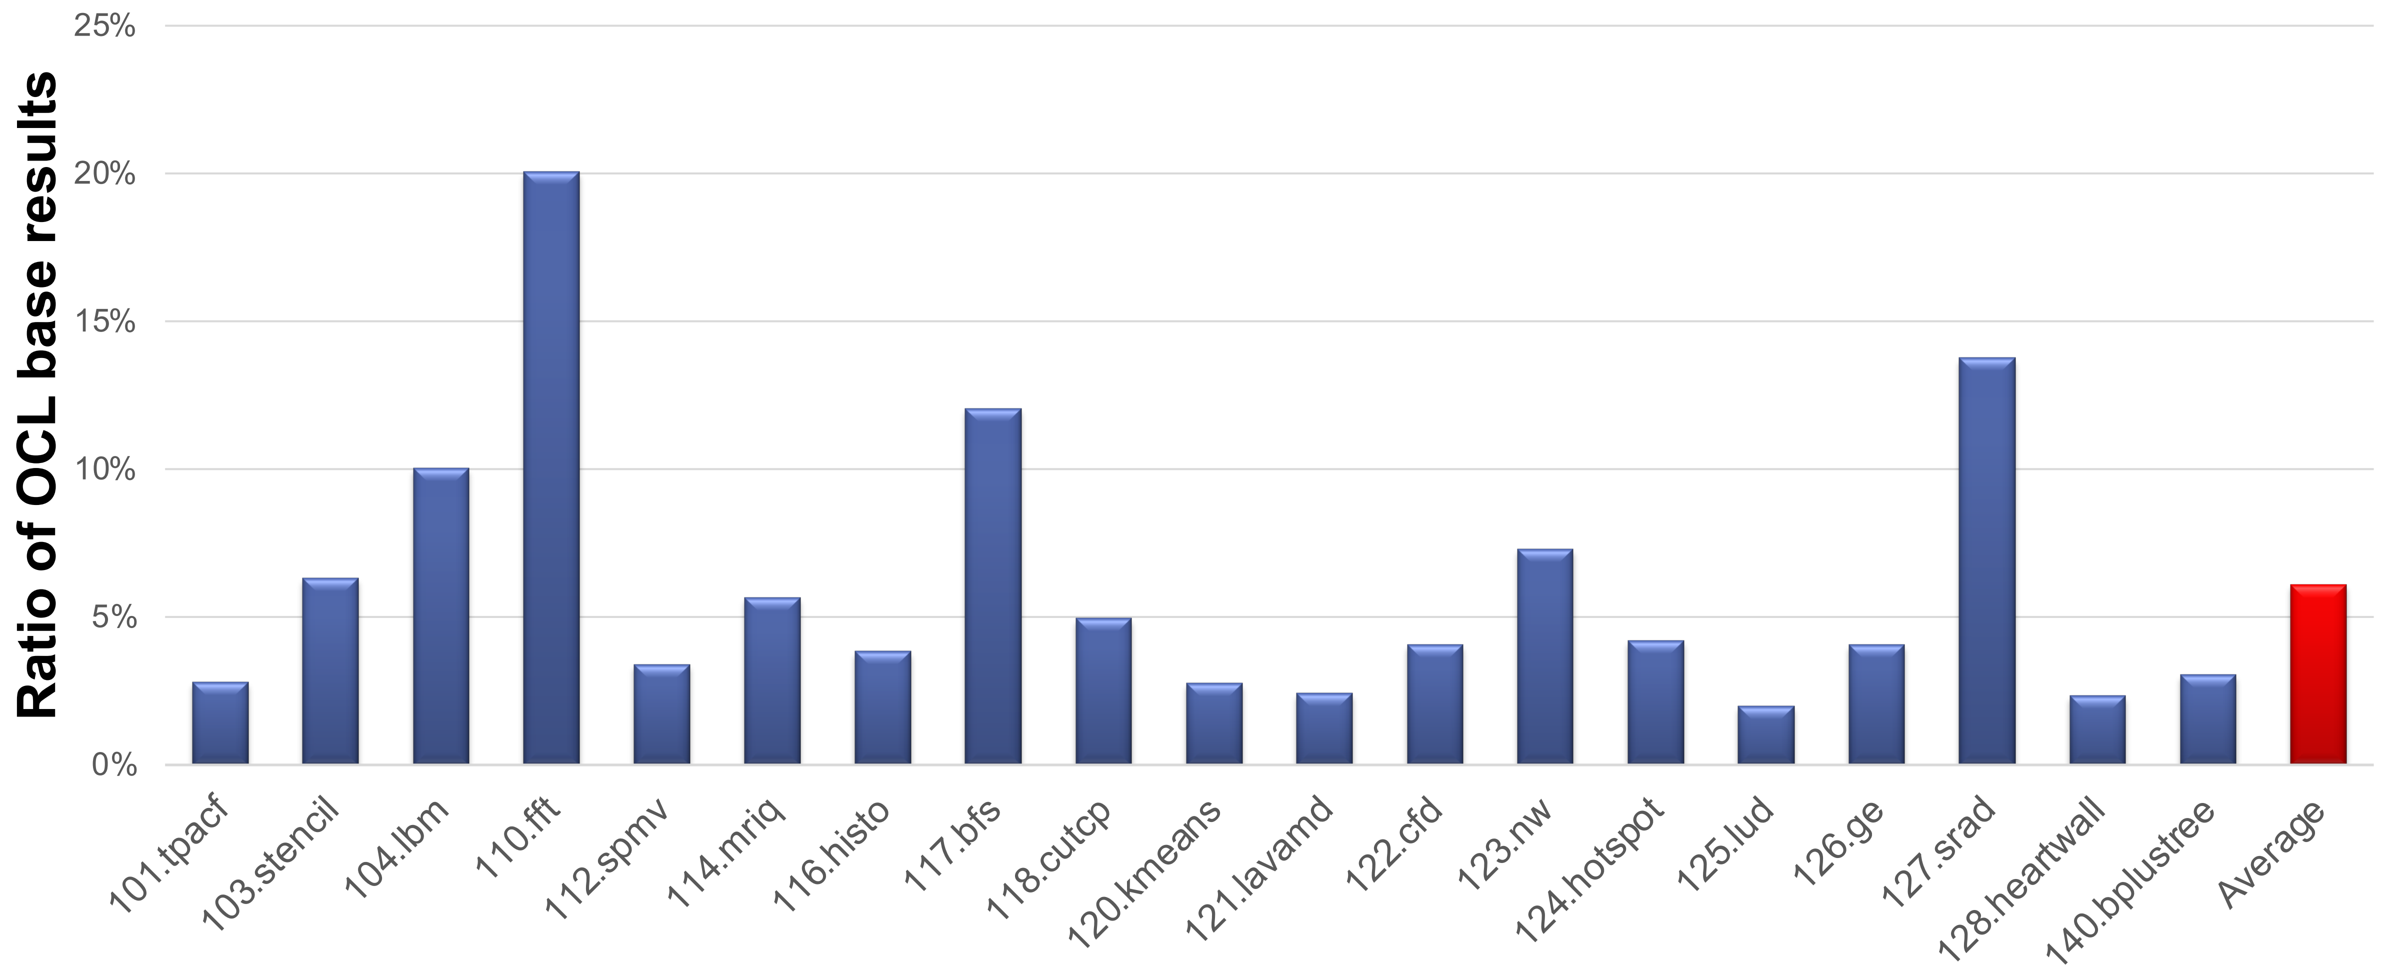

And Now for Something Totally Different

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

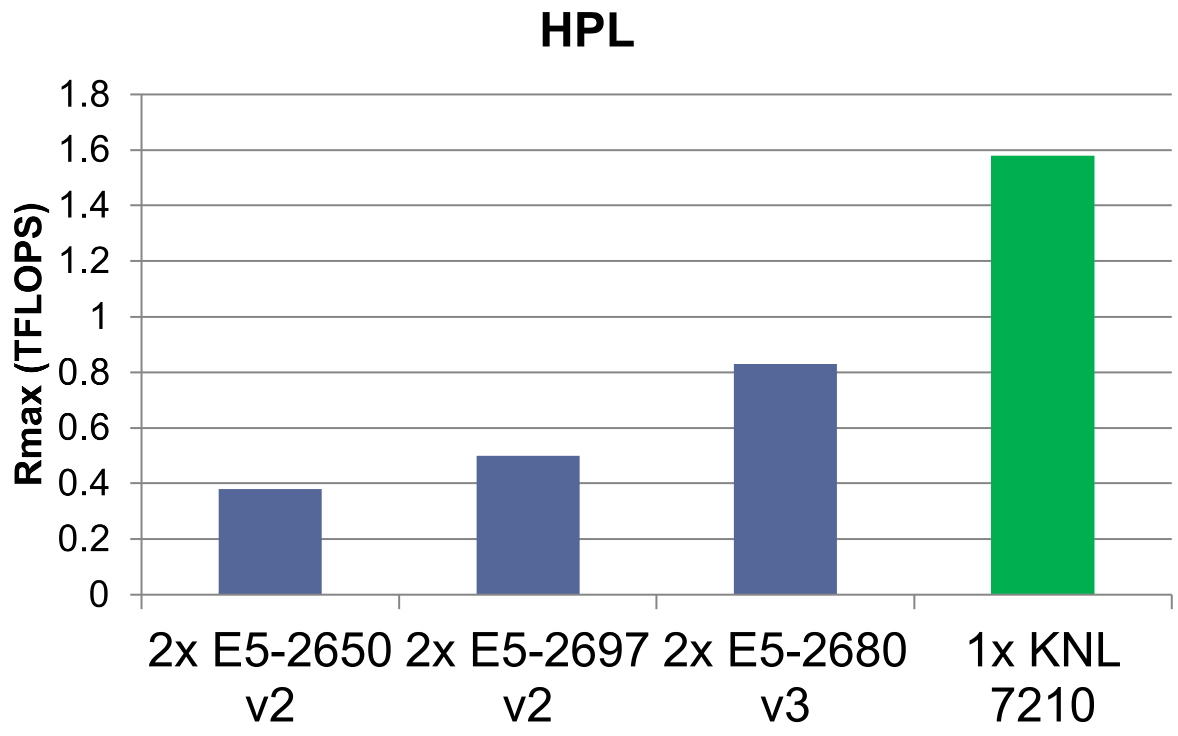

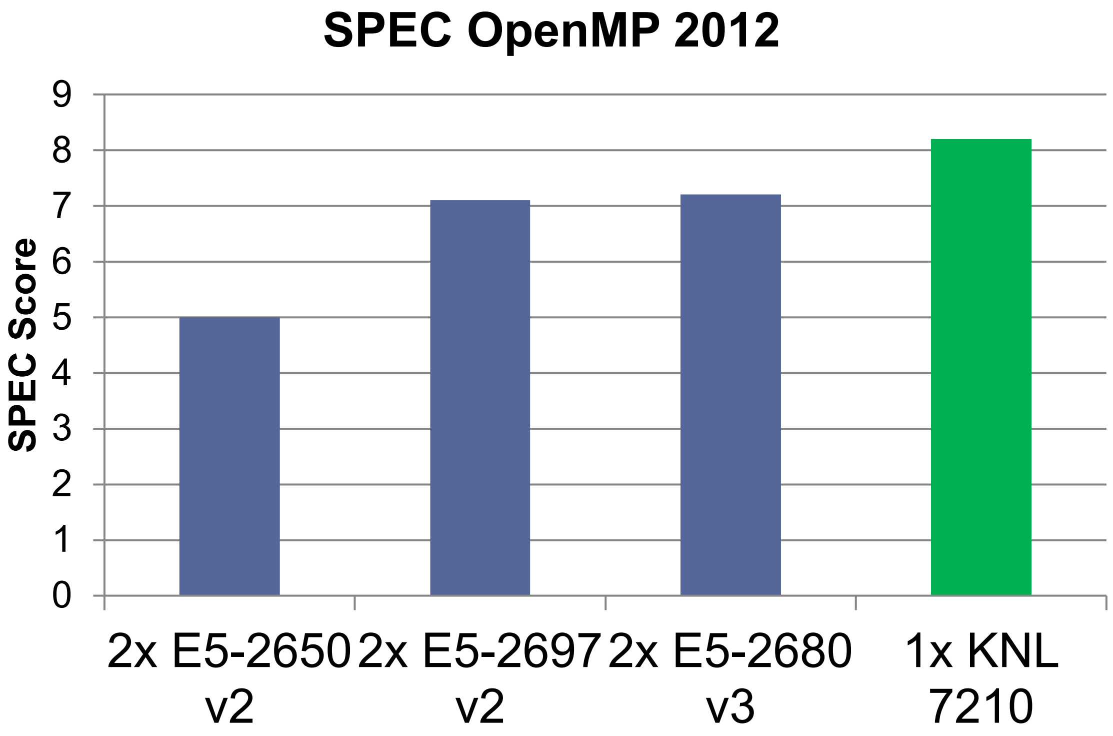

Comparing Results – Advanced Search

- 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

- Dump all Records as CSV
  - https://www.spec.org/cgi-bin/osgresults?conf=omp2012

Contents

- Interpretation of Results
- Comparing results
- Use Cases

Use Cases

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

Use Cases - Academia

spec

- Identify use cases for compiler validation and verification
  - Part of the ECP SOLLVE Validation and Verification testsuite
- Stress test internal UD systems with the new benchmark suite – SPEChpc2021
  - Push boundaries of both software and hardware
  - Procurement
- Train students to learn to use benchmark, scripts, queue jobs, ssh and so on
  - Submit results for SPEChpc2021
- Integrate large scientific code within the SPEC harness
  - Endless, tiring hours of integration process
  - Go thru steep learning curve of Makefile, CMAKE, cluster usage, dataset sizes, work with application developers and so on

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.

Thank you!

Questions?

Contact

SPEC Headquarters: info@spec.org
Sunita Chandrasekaran: schandra@udel.edu
Robert Henschel: henschel@iu.edu
Veronica G. Vergara Larrea: vergaravg@ornl.gov
Junjie Li: lijunj@iu.edu