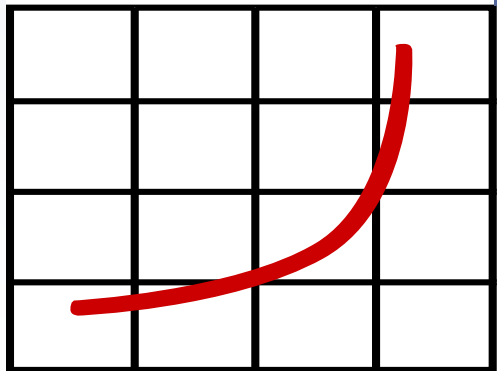




Using the SPEC HPG Benchmarks for Better Analysis and Evaluation of HPC Systems

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

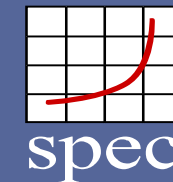


spec

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

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

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. Vergara Larrea

HPC Engineer @ Oak Ridge National Lab
vergaravg@ornl.gov

Mayara Gimenes

Research Assistant @ University of Delaware
mayara@udel.edu

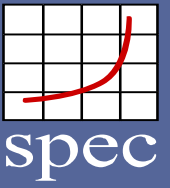
Sandra Wienke

Deputy HPC Group Lead @ RWTH Aachen University
wienke@itc.rwth-aachen.de

Swen Boehm

Computer Science Researcher @ Oak Ridge National Lab
boehms@ornl.gov

Tutorial Overview



- Overview of SPEC and SPEC HPG (40 min)
- How to get, setup and run the SPEC benchmarks (90 min)
- Break (10 AM) (30 min)
- How to interpret and compare SPEC benchmark results (40 min)
- Conclusion and Wrap-Up (10 min)

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

Get RSA token from Veronica & Swen, it's needed to access Summit for the hands-on session.
(ID is required)

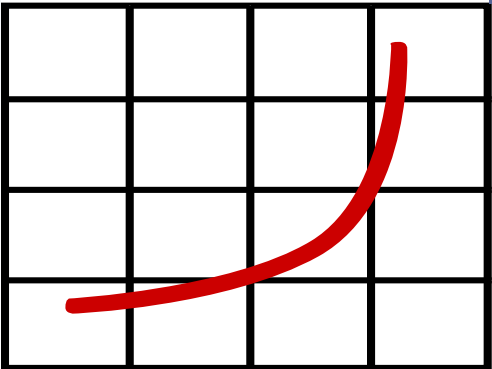
Please return the RSA tokens when you leave.





Overview of SPEC and the SPEC High Performance Group

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

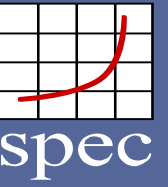


spec

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

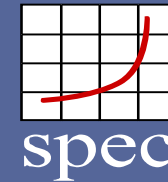
<http://pages.iu.edu/~lijunj/pearc/>

Contents



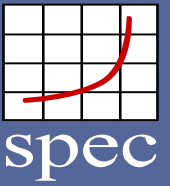
- Why SPEC benchmark?
- Intro to SPEC and SPEC HPG
- The SPEC Benchmark Philosophy
- SPEC HPG Benchmarks

Contents



- Why SPEC benchmark?
- Intro to SPEC and SPEC HPG
- The SPEC Benchmark Philosophy
- SPEC HPG Benchmarks

Why SPEC benchmark?



- 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 do you see by reading 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*

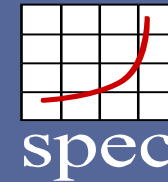
Results Table													
Benchmark	Base						Peak						
	Ranks	Seconds	Ratio	Seconds	Ratio		Ranks	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
l04.mile	640	15.4	102	14.9	105	14.9	105						
l07.leslie3d	640	34.1	153	33.2	157	33.4	156						
l13.GemsFDTD	640	187	33.8	186	33.8	186	33.9						
l15.fds4	640	23.3	83.9	22.8	85.6	23.2	84.0						
l21.pop2	640	77.5	53.2	77.5	53.3	77.3	53.4						
l22.tachyon	640	31.4	89.0	31.5	88.9	32.1	87.2						
l26.lammps	640	90.3	32.3	89.6	32.5	89.7	32.5						
l27.wrf2	640	29.5	264	30.2	258	29.6	264						
l28.GAPgeofem	640	8.10	255	8.31	249	8.28	249						
l29.tera_tf	640	22.1	125	22.5	123	22.3	124						
l30.socorro	640	30.7	124	31.1	123	31.8	120						
l32.zeusmp2	640	19.8	157	19.7	158	19.7	158						
l37.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		Software Summary	
Type of System:	Homogeneous	C Compiler:	Intel C Composer XE for Linux, Version 18.0.0.128 Build 20170811
Compute Node:	HPE XA730i Gen10 Server Node	C++ Compiler:	Intel C++ Composer XE for Linux, Version 18.0.0.128 Build 20170811
Interconnect:	InfiniBand (MPI and I/O)	Fortran Compiler:	Intel Fortran Composer XE for Linux, Version 18.0.0.128 Build 20170811
File Server Node:	Lustre FS	Base Pointers:	64-bit
Total Compute Nodes:	16	Peak Pointers:	Not Applicable
Total Chips:	32	MPI Library:	HPE Performance Software - Message Passing Interface 2.17
Total Cores:	640	Other MPI Info:	OFED 3.2.2
Total Threads:	1280	Pre-processors:	None
Total Memory:	3 TB	Other Software:	None
Base Ranks Run:	640		
Minimum Peak Ranks:	--		
Maximum Peak Ranks:	--		

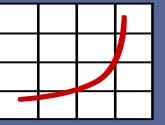
Node Description: HPE XA730i Gen10 Server Node	
Hardware	Software
Number of nodes:	16
Uses of the node:	compute
Vendor:	Hewlett Packard Enterprise
Model:	SGI 8600 (Intel Xeon Gold 6148, 2.40 GHz)
CPU Name:	Intel Xeon Gold 6148
CPU(s) orderable:	1-2 chips
Chips enabled:	2
Cores enabled:	40
Cores per chip:	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
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

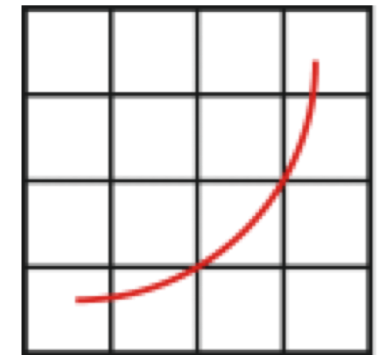


- Why SPEC benchmark?
- **Intro to SPEC and SPEC HPG**
- The SPEC Benchmark Philosophy
- SPEC HPG Benchmarks

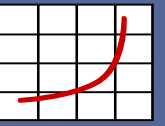
Standards Performance Evaluation Corporation (SPEC)^{spec}



- **SPEC** is a non-profit corporation formed in 1988 to establish, maintain and endorse standardized benchmarks and tools to evaluate performance and energy efficiency for the newest generation of computing systems.
- Composed of four groups
 - Graphics and Workstation Performance Group (GWPG)
 - High Performance Group (HPG)
 - Open Systems Group (OSG)
 - Research Group (RG)
- <https://www.spec.org>
- <https://www.spec.org/hpg/>



spec[®]

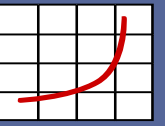


Standards Performance Evaluation Corporation (SPEC)_{spec}

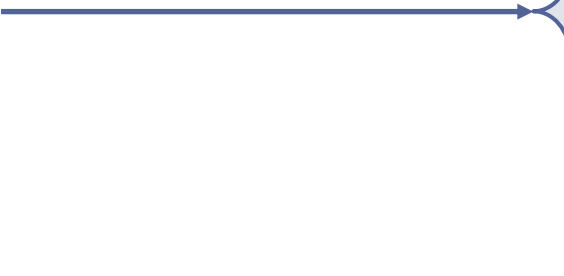
- **SPEC** is a non-profit corporation formed in 1988 to establish, maintain and endorse standardized benchmarks and tools to evaluate performance and energy efficiency for the newest generation of computing systems.
- Composed of four groups
 - Graphics and Workstation Performance Group (GWPG)
 - High Performance Group (HPG)
 - Open Systems Group (OSG) —————→
 - Research Group (RG)
- <https://www.spec.org>
- <https://www.spec.org/hpg/>

Largest & Oldest Group

- Cloud
- CPU
- Java
- Power
- Virtual Machine
- File Server



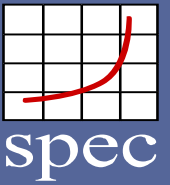
Standards Performance Evaluation Corporation (SPEC)_{spec}

- **SPEC** is a non-profit corporation formed in 1988 to establish, maintain and endorse standardized benchmarks and tools to evaluate performance and energy efficiency for the newest generation of computing systems.
- Composed of four groups
 - Graphics and Workstation Performance Group (GWPG)
 - High Performance Group (HPG) 
 - Open Systems Group (OSG)
 - Research Group (RG)
- <https://www.spec.org>
- <https://www.spec.org/hpg/>

HPC benchmarks


- MPI
- OpenMP
- Accelerator
 - OpenCL
 - OpenACC
 - OpenMP 4.5

SPEC Members



132 member organizations as of July-2019, including:

- 95 companies
- 37 academic institutions

**Standard Performance Evaluation Corporation**

Home Benchmarks Tools Results Contact Site Map Search Help

Benchmarks

- Cloud
- CPU
- Graphics/Workstations
- ACCEL/MPI/OMP
- Java Client/Server
- Mail Servers
- Storage
- Power
- Virtualization
- Web Servers

Results Search

Submitting Results

- Cloud/CPU/Java/Power
- SFS/Virtualization
- ACCEL/MPI/OMP
- SPECapc/SPECviewperf/SPECwpc

Tools

- SERT
- PTDaemon
- Chauffeur WDK

Order Benchmarks

- Order Form
- Downloads

SPEC

- About SPEC
 - GWFG
 - HPG
 - OSG
 - RG
- Membership
 - Member organizations

The SPEC Consortium: Members and Associates

SPEC Members:

Acer Inc. * Action S.A. * Advanced Micro Devices * Amazon Web Services, Inc. * Apple Inc. * ARM * Avere Systems * Bull SAS * Cavium Inc. * Ciara Technologies Inc. * Cisco Systems, Inc. * Dell, Inc. * Digital Ocean * E4 Computer Engineering SPA * Fujitsu * Gartner, Inc. * Guizhou Huaxintong Semiconductor Technology Co. Ltd * Hitachi Data Systems * Hitachi Ltd. * Hewlett Packard Enterprise * HP Inc. * Huawei Technologies Co. Ltd. * IBM * Inspur Corporation * Intel * Lenovo * M Computers s.r.o. * Microsoft * NEC - Japan * NetApp * New H3C Technologies Co., Ltd. * NVIDIA * Oracle * OVH SAS * Primary Data * Principled Technologies * Pure Storage * Qualcomm Technologies Inc. * Quanta Computer Inc. * Red Hat * Samsung * SAP AG * Seagate * Sugon * Super Micro Computer, Inc. * SUSE * Taobao (China) Software Co. Ltd. * Unisys * Veritas Technologies * Via Technologies * VMware * WekaIO *

SPEC Associates:

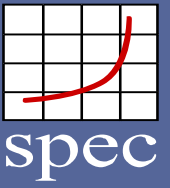
Academia Sinica, Institute of Information Science * Argonne National Laboratory * Charles University * China Academy of Telecommunication Research * Dresden University of Technology ZIH * fortiss GmbH * Helmholtz-Zentrum Dresden Rossendorf (HZDR) * Indiana University * JAIST * Karlsruhe Institute of Technology * Leibniz Rechenzentrum - Germany * Linaro Limited * National University of Singapore * Oak Ridge National Laboratory * Ohio State University * Pennsylvania State University * Purdue University * RWTH Aachen University * Technische Universität Darmstadt * Technische Universität Dresden * Telecommunications Technology Association * Tsinghua University * University of Aizu - Japan * University of Basel * University of California - Berkeley * University of Cologne * University of Delaware * University of Illinois at Urbana-Champaign * University of Maryland * University of Miami * University of Texas at Austin * University of Tsukuba * University of Wuerzburg * Virginia Polytechnic Institute and State University *

SPEC Research Group:

Advanced Strategic Technology LLC * Apple Inc. * ARM * bankmark UG * Barcelona Supercomputing Center * BEZNet * Charles University * Cisco Systems * Cloudera, Inc * Compilaflores * Delft University of Technology * Dell * Escuela Superior Politécnica del Litoral * fortiss GmbH * Friedrich-Alexander-University Erlangen-Nuremberg * Goethe University Frankfurt, Big Data Lab * Hewlett Packard Enterprise * Huawei * IBM * Imperial College London * Institute for Information Industry, Taiwan * Intel * Karlsruhe Institute of Technology * Kiel University * Linköping University * Lund University * Microsoft * NICTA * NovaTec Consulting GmbH * Oracle * Purdue University * Queen's University * Red Hat * RETIT GmbH * RWTH Aachen University * Salesforce.com * San Diego Supercomputing Center * San Francisco State University * SAP AG * Stiftung University * SINTEF * Software Performance and Scalability Consulting * Tata Consultancy Services * Technica Corporation * Technische Universität Darmstadt * Technische Universität Dresden * The MITRE Corporation * Umea University * University of Alberta * University of Coimbra * University of Lugano * University of Minnesota * University of North Florida * University of Paderborn * University of Stuttgart * University of Texas at Austin * University of Wuerzburg * University Politehnica of Bucharest * VMware * York University *

To learn about SPEC Membership, please read the [SPEC FAQ](#).

SPEC High Performance Group (HPG)

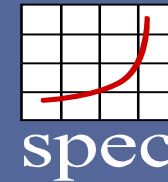


HPG develops benchmarks for high-performance computing systems, using real world applications.

- 30 member organizations as of July-2019
- 10 companies
- 20 academic

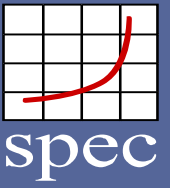


Contents



- Why SPEC benchmark?
- Intro to SPEC and SPEC HPG
- **The SPEC Benchmark Philosophy**
- SPEC HPG Benchmarks

SPEC Benchmark Philosophy

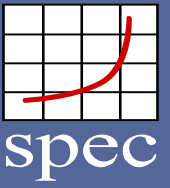


- The result of a SPEC benchmark suite is always a SPEC score.
 - Higher is better
 - Some benchmarks also have a power score, in addition to a performance score
- This score is always in relation to a reference machine.
 - Each benchmark has its own reference machine

SPEC Benchmark Philosophy cont'd

- SPEC (HPG) benchmarks are full applications.
 - Including all the overhead of a real application
- SPEC harness ensures correctness of results.
 - To detect “overly aggressive optimization”
 - To guard against tampering
- Each benchmark suite has a set of run rules.

SPEC Benchmark Philosophy cont'd



- Hierarchy within benchmark suites
 - Benchmark suite i.e. SPEC MPI2007
 - Dataset size -> i.e. Large
 - Component -> i.e. 104.milc

SPEC Benchmark Philosophy cont'd

- Benchmarks support “Base” and “Peak” configuration
 - These yield separate SPEC scores.
- Base runs
 - The same compiler switches for all components
 - The same parallelism
- “Peak” runs allow for more freedom of optimization.

- SPEC provides a standard methodology to measure and report power usage which can be incorporated into a SPEC benchmark.
- Normalizes the power usage across the full run of the suite

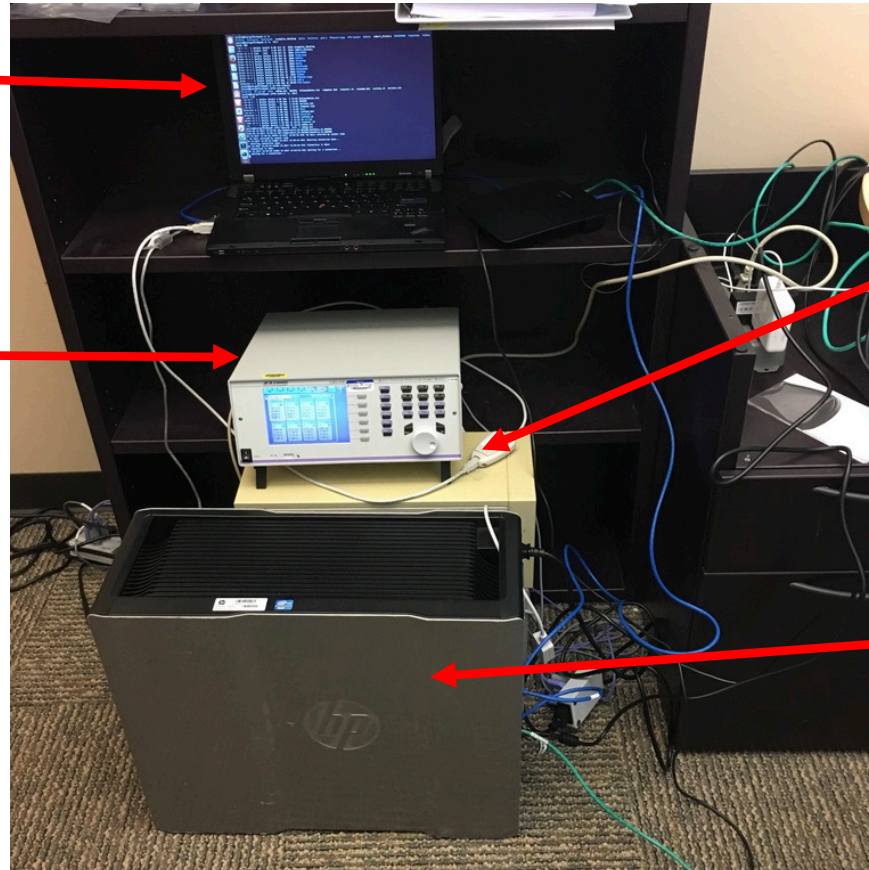


Power and
temperature
daemon

Power meter

Temperature sensor

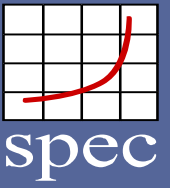
System Under Test
(SUT)



Benchmark Development Process

- Group effort, with lots of discussions from members
- Final decisions are by vote, even though we strive for consensus
- Find benchmark candidates
- Port into SPEC framework
- Port codes to additional programming models
- Define run rules
- Members meet in person multiple times a year, and online weekly.

Result Submission Process

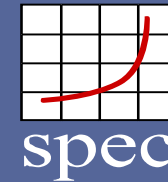


- Perform a valid run
- Supply hardware and software description
- Submit result to SPEC HPG for review (and publication)
 - 2 week review process
- Use the published result as you like, respecting the SPEC fair use guidelines.
(you can access the results even if you are not a member)

The Value of a Curated Result Repository

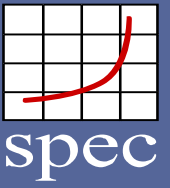
- Given appropriate hardware.... a published result should be reproducible just with the information available in the submission.
- Peer reviewed results are so much better than “everyone can upload a result”!
- The value of a benchmark suite lies in public results, their correctness and the ability to compare them.

Contents



- Why SPEC benchmark?
- Intro to SPEC and SPEC HPG
- The SPEC Benchmark Philosophy
- **SPEC HPG Benchmarks**

SPEC CPU – Not an HPG Benchmark!!



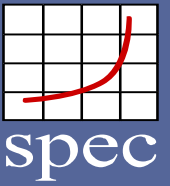
- SPEC CPU (2006 and 2017) is the most well known SPEC benchmark.
- Created by the Open Systems Group of SPEC
- HPG uses the same framework, if you are familiar with running SPEC CPU, you can run SPEC HPG benchmarks (And the other way around!).

- Follow on to SPEC OMP2001
- 14 applications
- Scales up to 512 threads
- **Support for power measurement**

- Citation:

Matthias S. Müller, John Baron, William C. Brantley, Huiyu Feng, Daniel Hackenberg, Robert Henschel, Gabriele Jost, Daniel Molka, Chris Parrott, Joe Robichaux, Pavel Shelepugin, Matthijs van Waveren, Brian Whitney, and Kalyan Kumaran. 2012. SPEC OMP2012 -- an application benchmark suite for parallel systems using OpenMP. In Proceedings of the 8th international conference on OpenMP in a Heterogeneous World (IWOMP'12), Barbara M. Chapman, Federico Massaioli, Matthias S. Müller, and Marco Rorro (Eds.). Springer-Verlag, Berlin, Heidelberg, 223-236. DOI=http://dx.doi.org/10.1007/978-3-642-30961-8_17

SPEC OMP2012

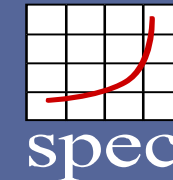


Code	Memory MB	LOC	Language	OMP call sites	OMP direc- tives	Area
350.md	5	1,768	Fortran	14	3	Molecular Dynamics
351.bwaves	22,800	876	F77	29	1	Computational Fluid Dynamics
352.nab	618	11,485	C	60	5	Molecular Modeling
357.bt331	11,188	2,331	Fortran	44	5	Computational Fluid Dynamics
358.botsalgn	156	1,277	C	4	3	Sequence Alignment
359.botsspar	7,179	209	C	8	4	LU factorization
360.ilbdc	16,482	978	Fortran	7	1	Lattice Boltzmann
362.fma3d	5,205	19,681	F90	142	5	Finite Element Method
363.swim	6,490	212	Fortran	14	3	Finite Difference
367.imagick	1,733	96,810	C	312	6	Image Processing
370.mgrid331	13,972	806	Fortran	20	5	Multi-Grid Solver
371.applu331	14,884	1,782	Fortran	81	9	PDE/SSOR
372.smithwa	177	2,561	C	22	3	Optimal Pattern Matching
376.kdtree	119	287	C++	4	3	Sorting and Searching

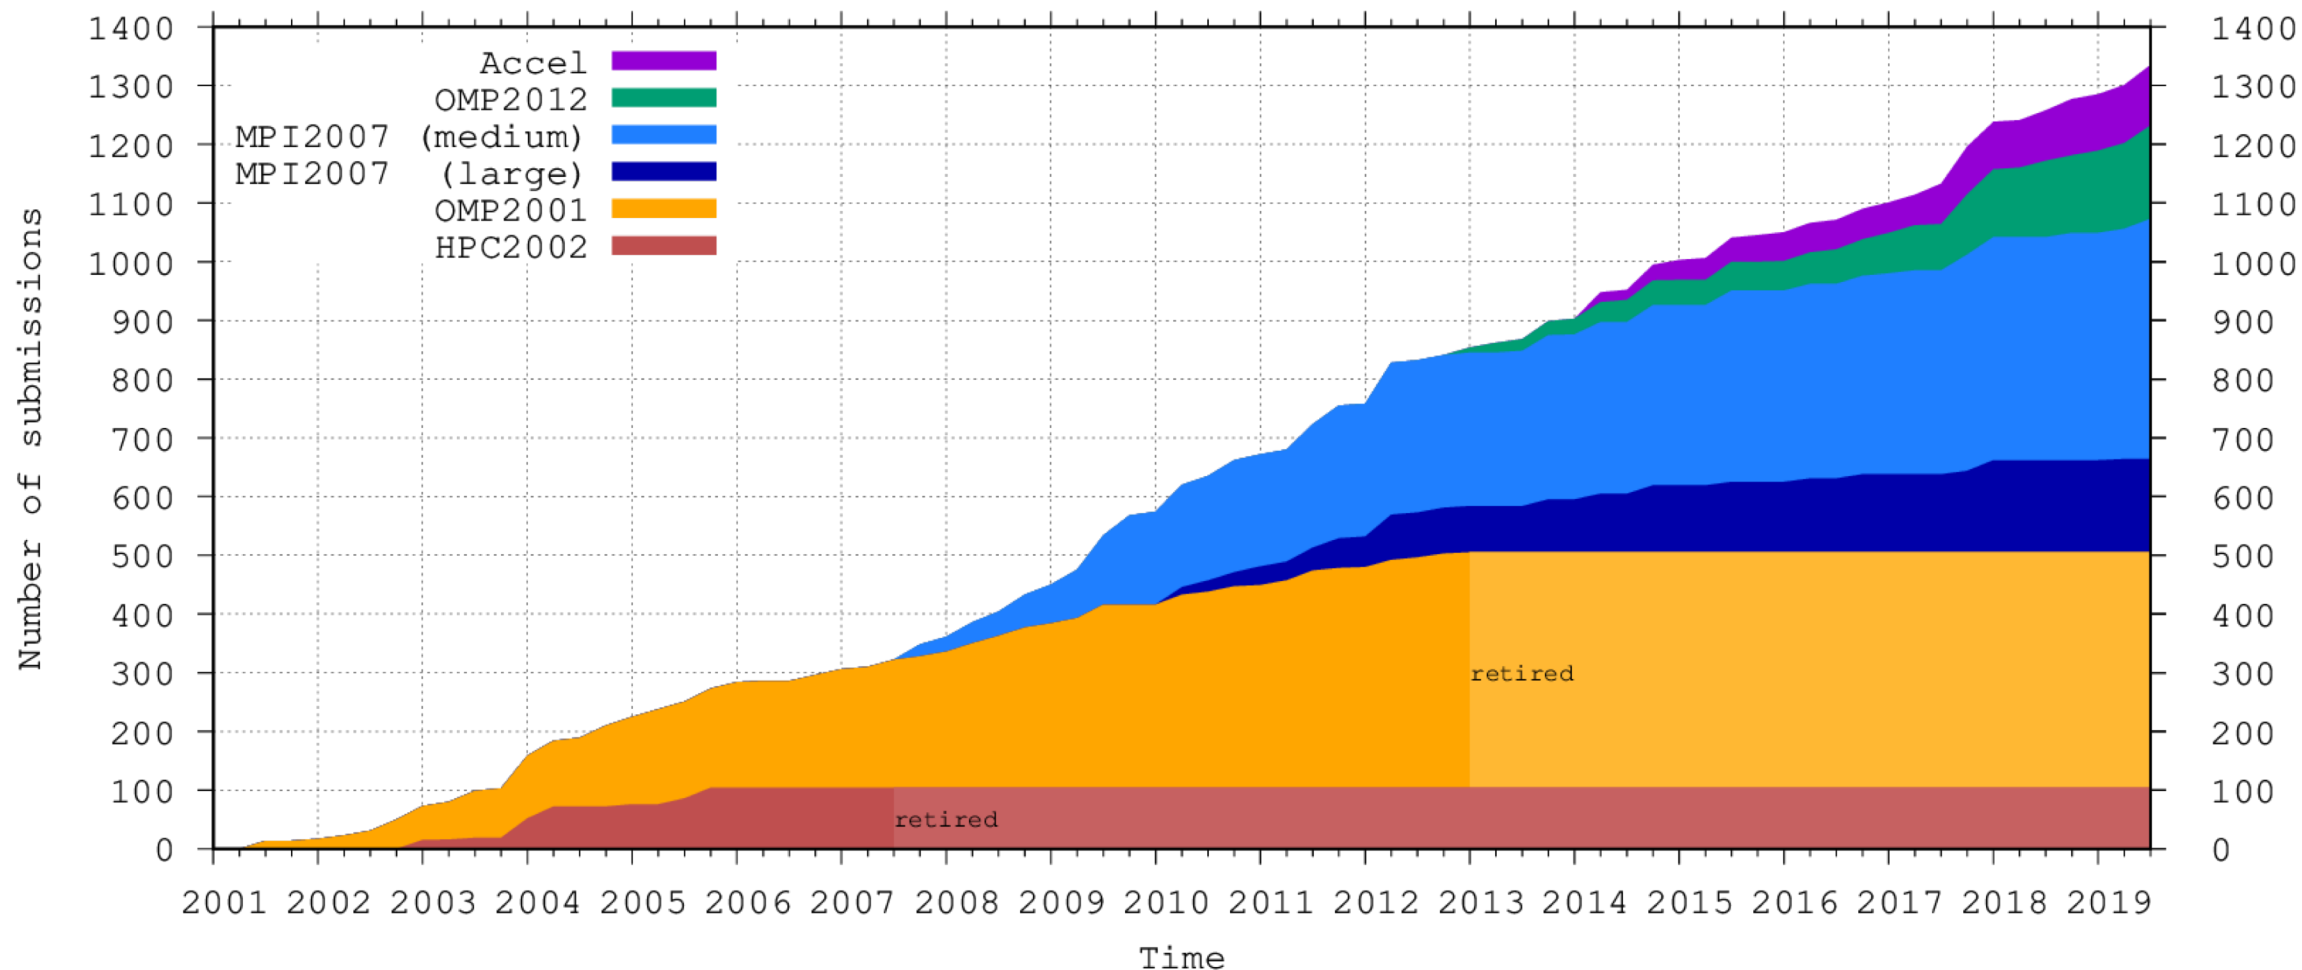
- SPEC Accel provides a comparative performance measure of
 - Hardware accelerator devices (GPU, Co-processors, etc.)
 - Supporting software tool chains (Compilers, Drivers, etc.)
 - Host systems and accelerator interface (CPU, PCIe, etc.)
- Computationally-intensive parallel HPC applications and mini-apps
- Portable across multiple accelerators
- Three distinct benchmarks
 - OpenACC
 - OpenCL
 - OpenMP 4.5 (first OpenMP 4.x benchmark supporting target offload)
- **Support for power measurement**

- Large and medium data set
- 13 applications
- Scales to 2048 MPI processes
- Power not supported

Result Submissions by Benchmark



- 1300+ published results, include all latest hardware
- Rich database for performance study

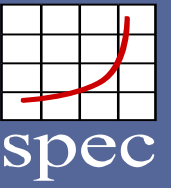


SPEC HPC2020

- First hybrid benchmark to address larger heterogeneous systems
- Scale up to 8000 MPI ranks
- Supports MPI, and MPI + multiple node level parallelisms (e.g. OpenMP, OpenACC, and possibly others)
- Expect to ship in 2020

- Submit results
- Become a member (\$800 for academic member)
- Contribute benchmark components
- Help with benchmark suite development
- Test release candidates

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