SPEChpc 2021 Benchmark Suites for Modern HPC Systems

Junjie Li
TACC
United States
jli@tacc.utexas.edu

Alexander Bobyr
Intel
Russia

Swen Boehm
ORNL
United States

William Brantley
AMD
United States

Holger Brunst
TU Dresden
Germany

Aurelien Cavelan
U of Basel
Switzerland

Sunita Chandrasekaran
U of Delaware
United States

Jimmy Cheng
Lenovo
United States

Florina M. Ciorba
U of Basel
Switzerland

Mathew Colgrove
NVIDIA
United States

Tony Curtis
Stony Brook U
United States

Christopher Daley
LBNL
United States

Mauricio Ferrato
U of Delaware
United States

Mayara Gimenes de Souza
U of Delaware
United States

Nick Hagerty
ORNL
United States

Robert Henschel
Indiana U
United States

Guido Juckeland
HZDR
Germany

Jeffrey Kelling
HZDR
Germany

Kelvin Li
IBM
Canada

Ron Lieberman
AMD
United States

Kevin McMahon
HPE
United States

Egor Melnichenko
Intel
Russia

Mohamed Ayoub Neggaz
U of Basel
Switzerland

Hiroshi Ono
HPE
United States

Carl Ponder
NVIDIA
United States

Dave Raddatz
HPE
United States

Severin Schueller
RWTH Aachen U
Germany

Robert Searles
U of Delaware
United States

Fedor Vasilev
Intel
Russia

Veronica Melesse Vergara
ORNL
United States

Bo Wang
RWTH Aachen U
Germany

Bert Wesarg
TU Dresden
Germany

Sandra Wienke
RWTH Aachen U
Germany

Miguel Zavala
U of Delaware
United States

Abstract

The SPEChpc™ 2021 suites are application-based benchmarks designed to measure performance of modern HPC systems. The benchmarks support MPI, MPI+OpenMP, MPI+OpenMP target offload, MPI+OpenACC and are portable across all major HPC platforms.

ACM Reference Format:

Junjie Li, Alexander Bobyr, Swen Boehm, William Brantley, Holger Brunst, Aurelien Cavelan, Sunita Chandrasekaran, Jimmy Cheng, Florina M. Ciorba, Mathew Colgrove, Tony Curtis, Christopher Daley, Mauricio Ferrato, Mayara Gimenes de Souza, Nick Hagerty, Robert Henschel, Guido Juckeland, Jeffrey Kelling, Kelvin Li, Ron Lieberman, Kevin McMahon, Egor Melnichenko, Mohamed Ayoub Neggaz, Hiroshi Ono, Carl Ponder, Dave Raddatz, Severin Schueller, Robert Searles, Fedor Vasilev, Veronica Melesse Vergara, Bo Wang, Bert Wesarg, Sandra Wienke, and Miguel Zavala. 2022. SPEChpc 2021 Benchmark Suites for Modern HPC Systems . In *Companion of the 2022 ACM/SPEC International Conference on Performance Engineering (ICPE '22 Companion), April 9–13, 2022, Bejing, China.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3491204.3527498

ICPE '22 Companion, April 9–13, 2022, Bejing, China

ACM ISBN 978-1-4503-9159-7/22/04.
https://doi.org/10.1145/3491204.3527498

1 About SPEC High Performance Group

The Standard Performance Evaluation Corporation (SPEC®) is a non-profit corporation that develops standard computer benchmarks and oversees the corresponding fair usage. The High Performance Group (HPG) is the team at SPEC responsible for designing and maintaining benchmarks for High-Performance Computing (HPC). Since 1996, the group has released a series of HPC benchmarks, and now actively maintains four benchmark suites including the MPI suite – SPEC MPI2007[3], the classical OpenMP suite – SPEC OMP2012[2], the single accelerator suites – SPEC Accel[1] supporting OpenCL, OpenACC and OpenMP target offload, and the newly released SPEChpc 2021[5] designed for both CPU-only and heterogeneous HPC systems with multiple accelerators.

SPEC and its benchmark names are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

2 Design of SPEChpc 2021

HPC architectures are evolving at a fast pace along with the programming models that support the hardware. It is highly desired to have a portable application-based benchmark to evaluate real-world hardware performance and maturity of software eco-system for the various architectures. SPEChpc 2021 is designed to approach this goal. From a pool of applications solicited from the search program, 9 production HPC codes and mini-apps were selected for the final release. They cover a wide range of scientific domains and multiple programming languages. A full list of the applications are to be found in SPEChpc online document[4].Different from the prior benchmarks offered by SPEC HPG where each suite supports one programming model, all applications in SPEChpc 2021 were ported to multiple node level host and accelerator models on top of MPI. Architecture specific tuning were removed for portable general performance and optimization relies on compilers and runtimes. The suites have been extensively tested on major hardware platforms using popular software stacks including different combinations of compilers, MPI libraries, OS, etc. Similar to any other SPEC benchmarks, SPEChpc 2021 is equipped with the SPEC harness that guards the source code from tampering and handles benchmark execution, output validation, result collection and reporting.

2.1 Supported Programming Models

Since each programming model only supports a subset of computing devices, SPEChpc 2021 is ported to support MPI+X parallelism, where X is chosen by user and can be one of the three models: OpenACC (2.6 standard), OpenMP thread and task (5.0 standard) or OpenMP target offload (5.0 standard). An MPI only execution model is also available. X is designed to be only these directive based models due to their good portability and the benefit of maintaining a single code path without algorithmic difference. Considering the standards and compilers are rapidly evolving, modification to the directives is allowed in peak run (explained in Section 2.3).

2.2 Workload Sizes and Design Limits

SPEChpc 2021 is strong scaled with fixed workload sizes. In order to properly benchmark machine of different sizes, 4 workloads, *Tiny*, *Small*, *Medium* and *Large*, were developed. All together, they are able cover from a single node to a few hundred nodes. The designed applicable range using MPI-only parallelism and the corresponding memory requirement can be found in the SPEChpc 2021 online document[4] under *System Requirements* page. While only the rank-counts in the suggested range of were thoroughly tested, each workload will likely also work outside these ranges. In such cases, the performance characteristics may vary significantly, so users should be cautious when exploring beyond the design limits.

2.3 Metrics

SPEChpc 2021 reports SPEC scores based on time-to-solution. Instead of reporting raw timing, a ratio of the run time on the reference system to the run time on the system-under-test is calculated for each application, then the SPEC score is calculated as geometric mean of the ratios of all applications, so a higher score indicates faster speed. The Taurus HPC system (Haswell & infiniband based) hosted at Technische Universität Dresden was chosen as the reference machine. The reference runs were performed in MPI-only using 24, 240, 2040, and 7200 ranks for *tiny*, *small*, *medium* and *large* workloads respectively. Details of the reference runs can be found in the published results [6].

Two metrics are available for each workload – base score and optionally peak score. The philosophy of base run is to provide an idea of basic tuning for a given system. In a base run, programming models, compiler flags, runtime setups and other settings need to be the same across all codes, a.k.a the "same for all" rule. More optimizations and tunings are allowed in peak run, runtime parameters like number of MPI ranks/OpenMP threads, compiler flags, environmental settings, programming models, etc. are allowed to differ on a application-by-application basis. In addition, tuning of OpenMP/OpenACC directives are allowed in peak, but these changes must be fully disclosed and subject to approval by SPEC HPG committee. Given the complexity involved in large scale HPC facilities, measurement of power efficiency isn't supported.

3 Result Disclosure

A formal SPEC run needs to be *reportable* where all apps in a chosen workload are required to run for a minimum of 2 iterations. The median of the ratios for odd number of iterations or the lower median for even number of iterations are used to calculate the SPEC score. A report is generated after a valid benchmark execution. In a reportable run, the SPEC harness will automatically populate the report with scores along with raw timing and ratios. To provide enough information for reproducibility, a formal SPEChpc 2021 report requires documentation of all the performance relevant software and hardware configurations. Same as other SPEC benchmarks, it is encouraged to submit results to SPEC for publication, all submissions will go through a two-week review period to ensure compliance with the run rules and improve consistency of documentation. All results published through SPEC are hosted in a public repository [6]. Self-publish is allowed but the results need to follow the same run rules, and SPEC reserves the rights to take action on non-compliant results. More details on result disclosure are to be found in the Run Rule part of the online document [4].

When browsing results in the repository, in addition to the score, it is noteworthy to review the machine details to learn what is driving the score as the computing resource used for the results can be vastly different.

References

[1] G. Juckeland, W. Brantley, S. Chandrasekaran, B. Chapman, S. Che, M. Colgrove, H. Feng, A. Grund, R. Henschel, W.-M. W. Hwu, et al. SPEC ACCEL: A standard application suite for measuring hardware accelerator performance. In *PMBS*, pages 46–67. Springer, 2014.

[2] M. S. Müller, J. Baron, W. C. Brantley, H. Feng, D. Hackenberg, R. Henschel, G. Jost, D. Molka, C. Parrott, J. Robichaux, et al. SPEC OMP2012—an application benchmark suite for parallel systems using OpenMP. In *IWOMP*, pages 223–236. Springer, 2012.

[3] M. S. Müller, M. Van Waveren, R. Lieberman, B. Whitney, H. Saito, K. Kumaran, J. Baron, W. C. Brantley, C. Parrott, T. Elken, et al. SPEC MPI2007—an application benchmark suite for parallel systems using MPI. *Concurrency and Computation: Practice and Experience*, 22:191–205, 2010.

[4] SPEC. SPEChpc 2021 documentation, 2021.

[5] SPEC. Spechpc 2021 press release, 2021.

[6] SPEC. SPEChpc 2021 results repository, 2021.