Overview of SPEC HPC Benchmarks and Details of the SPEChpc 2021 Benchmark

Robert Henschel
Project Director Research Engagement
Indiana University

Verónica Melesse Vergara
Group Leader, System Acceptance and User Environment
Oak Ridge Leadership Computing Facility

spec®

Content

- SPEC and SPEC HPG
- SPEC HPG Benchmark Suites
- Overview SPEChpc 2021
- First Results
- ORNL Experiences

Standard Performance Evaluation Corporation

- Worldwide non-profit consortium formed in 1988
- Develops industry-standard performance and energy efficiency benchmarks, mainly for servers and workstations
- Creates benchmarks through member collaboration with a focus on real-world applicability
- Sponsors research and international conferences addressing diverse aspects of performance
- Its membership comprises more than 127 leading computer hardware and software vendors, educational institutions, research organizations, and government agencies worldwide.
- https://www.spec.org/consortium/

High Performance Group

- Membership
  - Industry: 14   Academia: 23
- Active Participants
  - AMD, ATOS, HPE, Intel, Lenovo, NextSilicon, NVIDIA, Siemens
  - Argonne NL, Brookhaven NL, Indiana University, Lawrence Berkeley NL, Oak Ridge NL, RWTHA Aachen University, Stony Brook University, Texas Advanced Computing Center, Technische Universität Dresden, University of Basel
- Benchmarks
  - Represent large, real applications, in scientific and technical computing,
  - Use industry standard parallel application programming interfaces (APIs), OpenACC, OpenMP and MPI
  - Support shared-memory and message passing programming paradigms,
  - Can evaluate shared-memory computers, distributed-memory computers and workstation clusters as well as traditional massively parallel processor computers,
  - Come in several data sets sizes (from a few minutes to days of execution time),

Prior SPEC HPG Benchmarks

- **SPEC MPI®2007**: Performance of compute intensive applications using the Message-Passing Interface (MPI)
- **SPEC OMP®2012**: Measuring performance using applications based on the OpenMP 3.1 standard for shared-memory parallel processing.
- **SPEC ACCEL®**: Performance with computationally intensive parallel applications running under the OpenCL, OpenACC, and OpenMP 4 target offloading APIs.
- **SPEChpc™ 2021:** Performance of hybrid applications using MPI plus OpenACC, OpenMP, OpenMP target offload or pure MPI.

SPEChpc Benchmark Suites

- Combines elements of previous SPEC HPG suites to create an application-based benchmark which can be run using MPI and optionally hybrid with a node-level parallel model
- Four suites, Tiny, Small, Medium, and Large, with increasing workload sizes, allows for appropriate evaluation of different sized HPC systems, ranging from a single node to many hundreds of nodes.
- The suites contain 9 full and proxy scientific applications from various domains written in C, C++, or Fortran.
- Comprehensive support for multiple programming models, including MPI, MPI+OpenACC, MPI+OpenMP, and MPI+OpenMP with target offload.
- Able to run on either purely CPUs or offloaded to accelerators
- **4 parallel models x 4 dataset sizes**

SPEChpc Design Choices (highlights)

- Focus on Portable General Performance rather than allowing architecture specific application tuning
  - Rely on compiler rather than application engineer
  - Though researcher are encouraged to investigate code modifications and optimization provided the results are marked as an estimate.
- Split OpenMP into two ports
  - Thread/Task based targeting multicore-CPU
  - Target based targeting accelerators
  - Because of potential bias, directive modification is allowed in peak
- Benchmark selection based on availability of the code, portability, scalability, and performance characteristics.

SPEChpc 2021 Benchmarks

Benchmark	Domain	Submitter
LBM D2Q37	Computational Fluid Dynamics	Sebastiano Fabio Schifano, University of Ferrara and INFN
SOMA	Polymeric Systems	Ludwig Schneider for the SOMA collaboration
Tealeaf	High Energy Physics	Simon McIntosh-Smith, University of Bristol
Cloverleaf	High Energy Physics	Simon McIntosh-Smith, University of Bristol
MiniSweep	Radiation Transport	Wayne Joubert, Oak Ridge National Laboratory
POT3D	Solar Physics	Ron Caplan, Predictive Science
SPH-EXA	Astrophysics and Cosmology	Florina Ciorba, University of Basel
HPGMG-FV	Cosmology and Combustion	Christopher Daley, Lawrence Berkely National Laboratory
miniWeather	Weather Modeling	Matt Norman, Oak Ridge National Laboratory

Results: https://www.spec.org/hpc2021/results/hpc2021.html

spec®

SPEChpc2021 Medium (10):

Test Sponsor	System Name	System Configuration				Results	
		Node-level Parallelization Model	Compute Nodes Used	MPI Ranks	Base Threads Per Rank	Base	Peak
Intel	Endeavour: Intel Server M50CYP2UR208 (Intel Xeon Platinum 8360Y) HTML | CSV | Text | PDF | PS | Config	OMP	16	192	6	0.682	Not Run
Intel	Endeavour: Intel Server M50CYP2UR208 (Intel Xeon Platinum 8360Y) HTML | CSV | Text | PDF | PS | Config	OMP	32	256	9	1.36	1.47
Intel	Endeavour: Intel Server M50CYP2UR208 (Intel Xeon Platinum 8360Y) HTML | CSV | Text | PDF | PS | Config	OMP	64	1152	4	2.79	2.96
Intel	Endeavour: Intel Server M50CYP2UR208 (Intel Xeon Platinum 8360Y) HTML | CSV | Text | PDF | PS | Config	OMP	128	1536	6	5.62	5.97
Oak Ridge National Laboratory	Summit: IBM Power System AC922 (IBM Power9, Tesla V100-SXM2-16GB) HTML | CSV | Text | PDF | PS | Config	ACC	700	4200	1	41.3	Not Run
RWTH Aachen University	CLAIX-2018: Intel Compute Module HNS2600BPM (Intel Xeon Platinum 8160) HTML | CSV | Text | PDF | PS | Config	MPI	100	4800	1	2.00	2.32
Technische Universitaet Dresden	Taurus: bullx DLC B720 (Intel Xeon E5-2680 v3) HTML | CSV | Text | PDF | PS | Config	MPI	85	2040	1	1.04	Not Run
Texas Advanced Computing Center	Frontera: PowerEdge C6420 (Intel Xeon Platinum 8280) HTML | CSV | Text | PDF | PS | Config	OMP	512	1024	27	15.8	Not Run
Texas Advanced Computing Center	Frontera: PowerEdge C6420 (Intel Xeon Platinum 8280) HTML | CSV | Text | PDF | PS | Config	OMP	1024	2048	27	24.3	Not Run
Texas Advanced Computing Center	Frontera: PowerEdge C6420 (Intel Xeon Platinum 8280) HTML | CSV | Text | PDF | PS | Config	OMP	2048	4096	27	30.8	Not Run

SPEChpc2021 Large (5):

Test Sponsor	System Name	System Configuration				Results	
		Node-level Parallelization Model	Compute Nodes Used	MPI Ranks	Base Threads Per Rank	Base	Peak
Oak Ridge National Laboratory	Summit: IBM Power System AC922 (IBM Power9, Tesla V100-SXM2-16GB) HTML | CSV | Text | PDF | PS | Config	ACC	1400	8400	1	41.0	Not Run
Technische Universitaet Dresden	Taurus: bullx DLC B720 (Intel Xeon E5-2680 v3) HTML | CSV | Text | PDF | PS | Config	MPI	300	7200	1	0.983	Not Run
Texas Advanced Computing Center	Frontera: PowerEdge C6420 (Intel Xeon Platinum 8280) HTML | CSV | Text | PDF | PS | Config	OMP	2048	4096	27	31.2	Not Run
Texas Advanced Computing Center	Frontera: PowerEdge C6420 (Intel Xeon Platinum 8280) HTML | CSV | Text | PDF | PS | Config	OMP	1024	2048	27	17.3	Not Run
Texas Advanced Computing Center	Frontera: PowerEdge C6420 (Intel Xeon Platinum 8280) HTML | CSV | Text | PDF | PS | Config	OMP	512	1024	27	8.47	Not Run

Last update: *Monday, 21 March 2022, 09:05*

Result Details

SPEChpc™ 2021 Medium Result

NEC Corporation
(Test Sponsor: RWTH Aachen University)

CLAIX-2018: Intel Compute Module HNS2600BPM (Intel Xeon Platinum 8160)

SPEChpc 2021_med_base = 2.00

SPEChpc 2021_med_peak = 2.32

hpc2021 License: 055A
Test Sponsor: RWTH Aachen University
Tested by: RWTH Aachen University
Test Date: Sep-2021
Hardware Availability: Nov-2018
Software Availability: Sep-2020

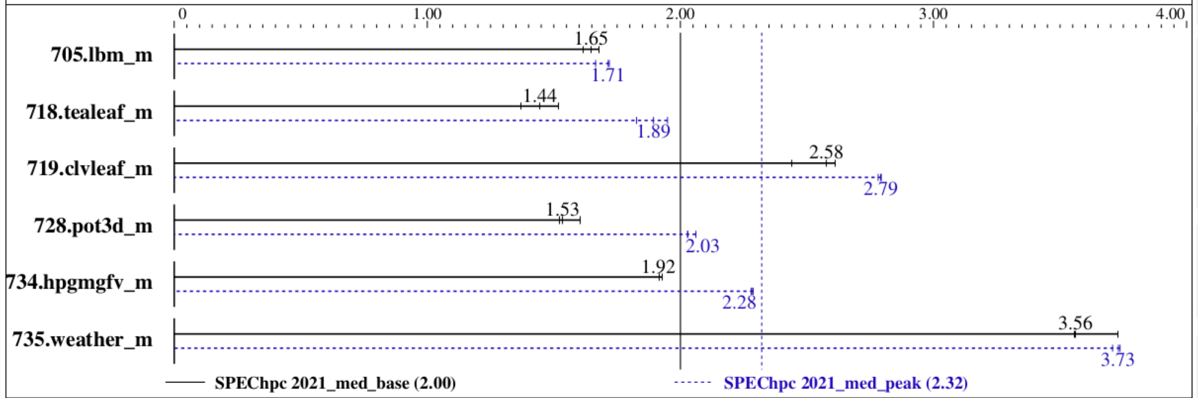

Results Table

Benchmark	Base									Peak								
	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
705.lbm_m	MPI	4800	1	**744**	**1.65**	759	1.61	731	1.68	MPI	4800	1	736	1.67	713	1.72	**715**	**1.71**
718.tealeaf_m	MPI	4800	1	890	1.52	**936**	**1.44**	987	1.37	MPI	4400	1	**713**	**1.89**	739	1.83	693	1.95
719.clvleaf_m	MPI	4800	1	709	2.61	759	2.44	**718**	**2.58**	MPI	4400	1	662	2.79	**663**	**2.79**	665	2.78
728.pot3d_m	MPI	4800	1	1154	1.60	**1208**	**1.53**	1216	1.52	MPI	4400	1	912	2.03	898	2.06	**912**	**2.03**
734.hpgmgfv_m	MPI	4800	1	518	1.93	**522**	**1.92**	522	1.92	MPI	4400	1	439	2.28	437	2.29	**438**	**2.28**
735.weather_m	MPI	4800	1	643	3.73	674	3.56	**674**	**3.56**	MPI	4800	1	642	3.74	647	3.71	**643**	**3.73**

SPEChpc 2021_med_base = 2.00

SPEChpc 2021_med_peak = 2.32

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

SPEChpc™ 2021 Medium Result

Copyright 2021 Standard Performance Evaluation Corporation

NEC Corporation
(Test Sponsor: RWTH Aachen University)

CLAIX-2018: Intel Compute Module HNS2600BPM (Intel Xeon Platinum 8160)

SPEChpc 2021_med_base = 2.00

SPEChpc 2021_med_peak = 2.32

hpc2021 License: 055A
Test Sponsor: RWTH Aachen University
Tested by: RWTH Aachen University
Test Date: Sep-2021
Hardware Availability: Nov-2018
Software Availability: Sep-2020

Hardware Summary

Type of System:	Homogenous
Compute Node:	Intel HNS2600BPB
Interconnect:	Intel Omni-Path 100 Series
Compute Nodes Used:	100
Total Chips:	200
Total Cores:	4800
Total Threads:	4800
Total Memory:	19200 GB
Max. Peak Threads:	1

Software Summary

Compiler:	C/C++/Fortran: Intel Compilers for Linux 2021.3.0
MPI Library:	Intel MPI Library for Linux 2018.4.274
Other MPI Info:	None
Other Software:	None
Base Parallel Model:	MPI
Base Ranks Run:	4800
Base Threads Run:	1
Peak Parallel Models:	MPI
Minimum Peak Ranks:	4400
Maximum Peak Ranks:	4800
Max. Peak Threads:	1
Min. Peak Threads:	1

Node Description: Intel HNS2600BPB

Hardware

Number of nodes:	100
Uses of the node:	compute
Vendor:	Intel Corporation
Model:	Intel Compute Module HNS2600BPB
CPU Name:	Intel Xeon Platinum 8160
CPU(s) orderable:	1-2 chips
Chips enabled:	2
Cores enabled:	48
Cores per chip:	24
Threads per core:	1
CPU Characteristics:	Intel Turbo Boost Technology up to 3.7 GHz
CPU MHz:	2100
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per core
L3 Cache:	33 MB I+D on chip per chip
Other Cache:	None
Memory:	192 GB (12 x 16 GB 2RX4 PC4-2666V-R)
Disk Subsystem:	Intel SSDSC2KG48, 480GB, SATA
Other Hardware:	None
Accel Count:	--
Accel Model:	--
Accel Vendor:	--
Accel Type:	--
Accel Connection:	--
Accel ECC enabled:	--
Accel Description:	--
Adapter:	Omni-Path HFI Silicon 100 Series
Number of Adapters:	1
Slot Type:	PCI Express Gen3 x16
Data Rate:	100Gbits/s
Ports Used:	1

Software

Accelerator Driver:	--
Adapter:	Omni-Path HFI Silicon 100 Series
Adapter Driver:	ib_ipoib 1.0.0
Adapter Firmware:	1.27.0
Operating System:	CentOS Linux release 7.9.2009
Local File System:	xfs
Shared File System:	1.4 PB NFS (Concat EMC Isilon X410) over Omni-Path
System State:	Multi-user, run level 3
Other Software:	None

Overview – Results

- Result example

SPEChpc™ 2021 Medium Result

Copyright 2021 Standard Performance Evaluation Corporation

NEC Corporation
(Test Sponsor: RWTH Aachen University)

CLAIX-2018: Intel Compute Module HNS2600BPM (Intel Xeon Platinum 8160)

SPEChpc 2021_med_base = 2.00

SPEChpc 2021_med_peak = 2.32

hpc2021 License:	055A	**Test Date:**	Sep-2021
Test Sponsor:	RWTH Aachen University	**Hardware Availability:**	Nov-2018
Tested by:	RWTH Aachen University	**Software Availability:**	Sep-2020

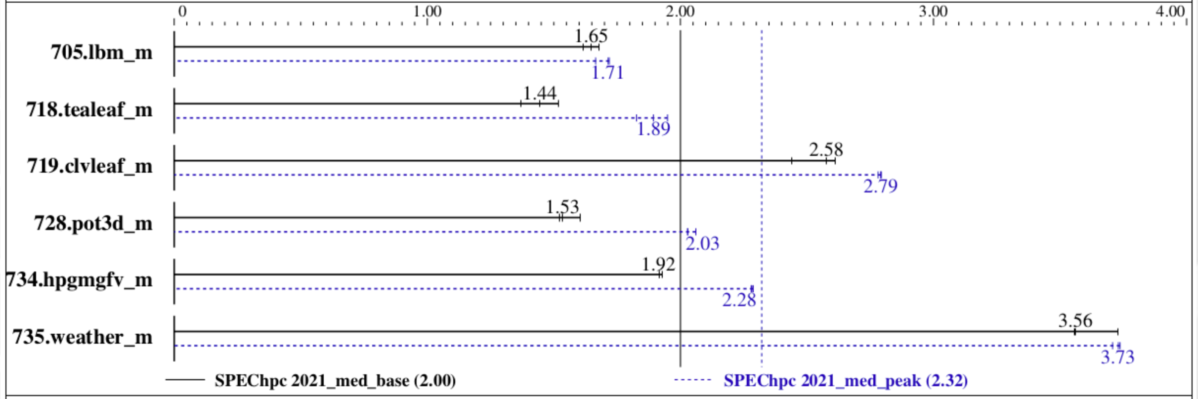

Results Table

Benchmark	Base Model	Base Ranks	Base Thrds/Rnk	Base Seconds	Base Ratio	Base Seconds	Base Ratio	Base Seconds	Base Ratio	Peak Model	Peak Ranks	Peak Thrds/Rnk	Peak Seconds	Peak Ratio	Peak Seconds	Peak Ratio	Peak Seconds	Peak Ratio
705.lbm_m	MPI	4800	1	**744**	**1.65**	759	1.61	731	1.68	MPI	4800	1	736	1.67	713	1.72	**715**	**1.71**
718.tealeaf_m	MPI	4800	1	890	1.52	**936**	**1.44**	987	1.37	MPI	4400	1	**713**	**1.89**	739	1.83	693	1.95
719.clvleaf_m	MPI	4800	1	709	2.61	759	2.44	**718**	**2.58**	MPI	4400	1	662	2.79	**663**	**2.79**	665	2.78
728.pot3d_m	MPI	4800	1	1154	1.60	**1208**	**1.53**	1216	1.52	MPI	4400	1	912	2.03	898	2.06	**912**	**2.03**
734.hpgmgfv_m	MPI	4800	1	518	1.93	**522**	**1.92**	522	1.92	MPI	4400	1	439	2.28	437	2.29	**438**	**2.28**
735.weather_m	MPI	4800	1	643	3.73	674	3.56	**674**	**3.56**	MPI	4800	1	642	3.74	647	3.71	**643**	**3.73**

SPEChpc 2021_med_base = 2.00

SPEChpc 2021_med_peak = 2.32

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

applications →

MPI Characteristics

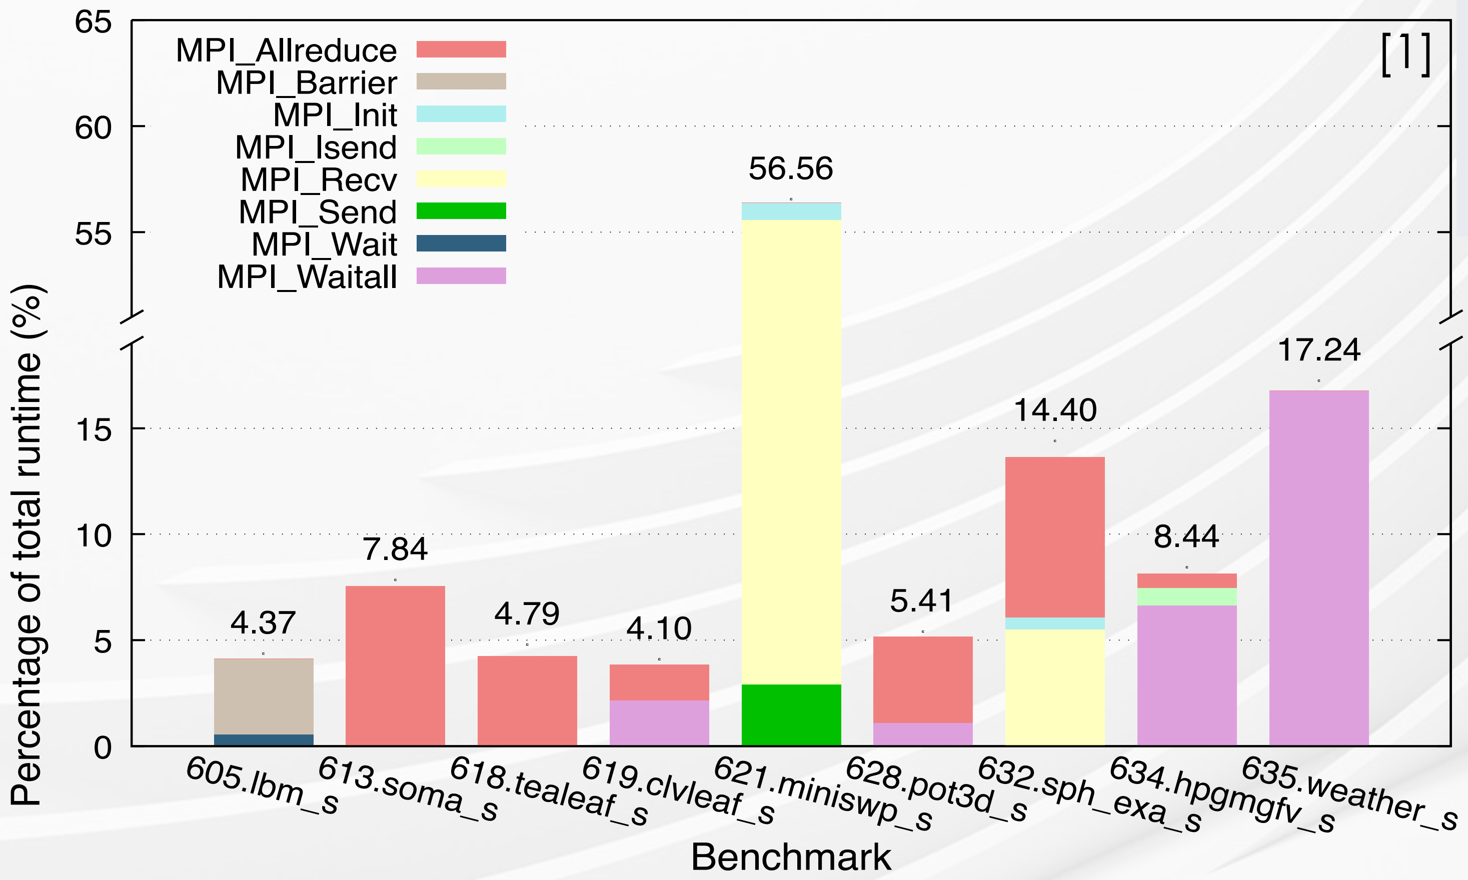

Setup

- Frontera@TACC: 2xIntel Xeon Platinum 8280 (Cascade Lake)
- Intel Compiler, Intel MPI
- Pmodel: **MPI-only**
- Workload: small suite
- #ranks: 224 (4 nodes w/ 56 ranks/node)

Relevant MPI functions

- MPI_Allreduce (red)
  SOMA, Tealeaf, Cloverleaf, Pot3d, SPH-EXA
- P2P communication (yellow, green)
  Minisweep, SPH-EXA, Hpgmg
- MPI_Waitall (purple)
  Cloverleaf, Pot3d, Hpgmg, weather

Code Characteristics

Setup
- Frontera@TACC:
  2xIntel Xeon Platinum 8280
  (Cascade Lake)
- Intel Compiler, Intel MPI
- Pmodel: **MPI-only**
- Workload: small suite
- #ranks: 224
  (4 nodes w/ 56 ranks/node)

spec®

- Instruction mix
  - Mix FP and non-FP ops
  - Mostly FP64-heavy codes (just SOMA some FP32 ops)
  - Mostly high vectorization rate

Benchmark	FP32 (% of uOps)	FP64 (% of uOps)	Non-FP (% of uOps)	Vectorization of FP (% of uOps)
605.lbm_s	0.00	51.98	48.02	86.80
613.soma_s	0.20	23.43	76.17	1.18
618.tealeaf_s	0.00	42.20	57.80	2.67
619.clvleaf_s	0.00	21.93	78.08	86.65
621.miniswp_s	0.00	8.92	91.07	57.90
628.pot3d_s	0.00	17.70	82.30	97.90
632.sph_exa_s	0.00	36.27	63.70	49.75
634.hpgmgfv_s	0.00	22.30	77.70	81.22
635.weather_s	0.00	26.32	73.67	3.45

[1]

Diversified Instructions similar for the tiny, medium and large suites.

Code Characteristics

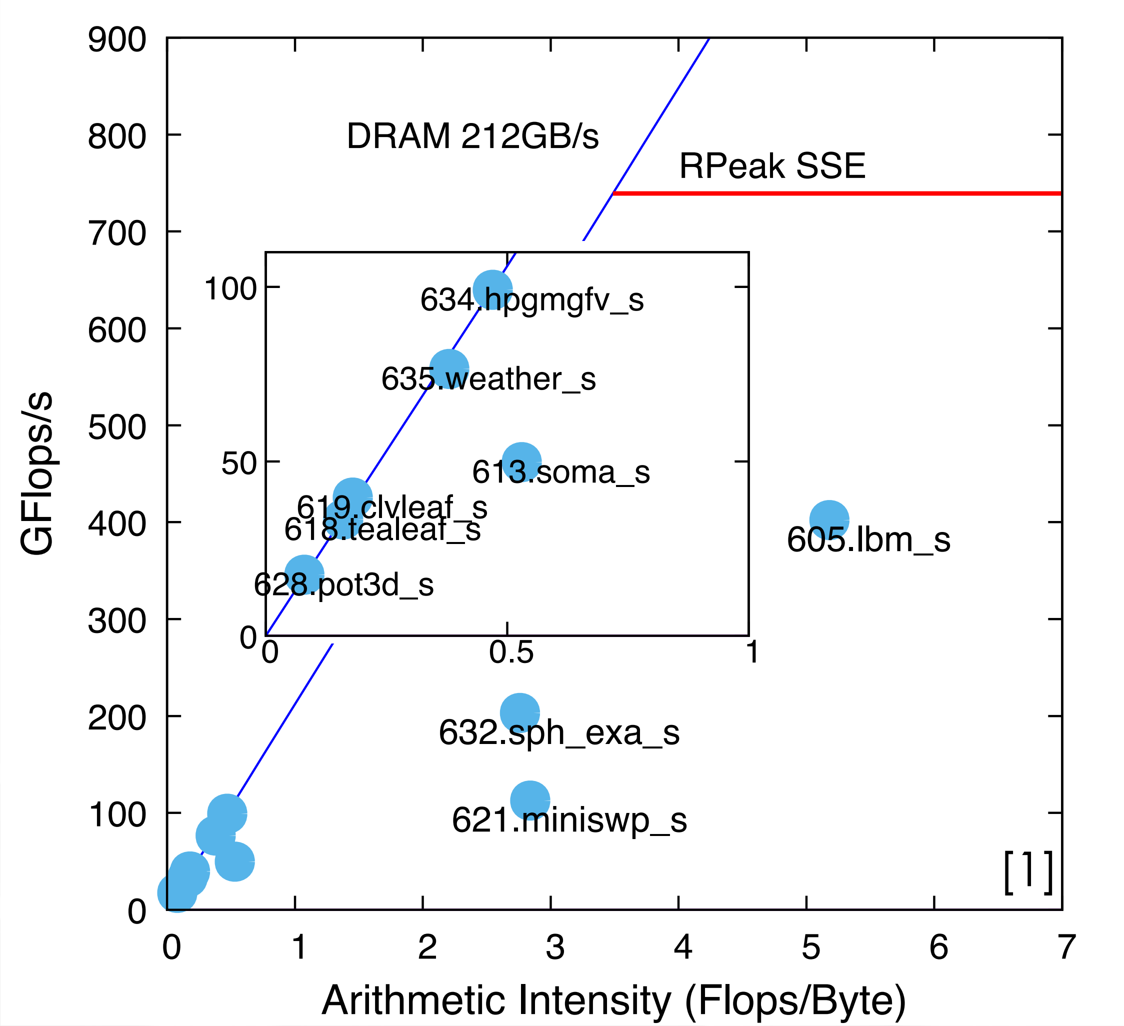

Setup

- Frontera@TACC:
  2xIntel Xeon Platinum 8280
  (Cascade Lake)
- Intel Compiler, Intel MPI
- Pmodel: **MPI-only**
- Workload: small suite
- #ranks: 224
  (4 nodes w/ 56 ranks/node)

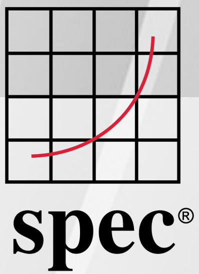

- Roofline Models
  - Most applications are clearly memory-bound
    - Tealeaf, Cloverleaf, Pot3d, Hpgmg, Weather
  - Some codes become less memory-bound with more nodes
    - Tealeaf, Weather
  - LBM: most compute-intensive code
    - Benefits most from vectorization

Roofline plots similar for the tiny, medium and large suites.
Arithmetic intensity collected for entire duration of each code.

Scalability: CPUs

Setup
- Frontera@TACC:
  2xIntel Xeon Platinum 8280 (Cascade Lake)
- Intel Compiler, Intel MPI
- Pmodel: **MPI-only**

- Scalability runs w/ all workloads (4 - 1024 nodes)
  - From a few nodes to a few hundreds
  - All suites scale well within their design limit

More results in [1]

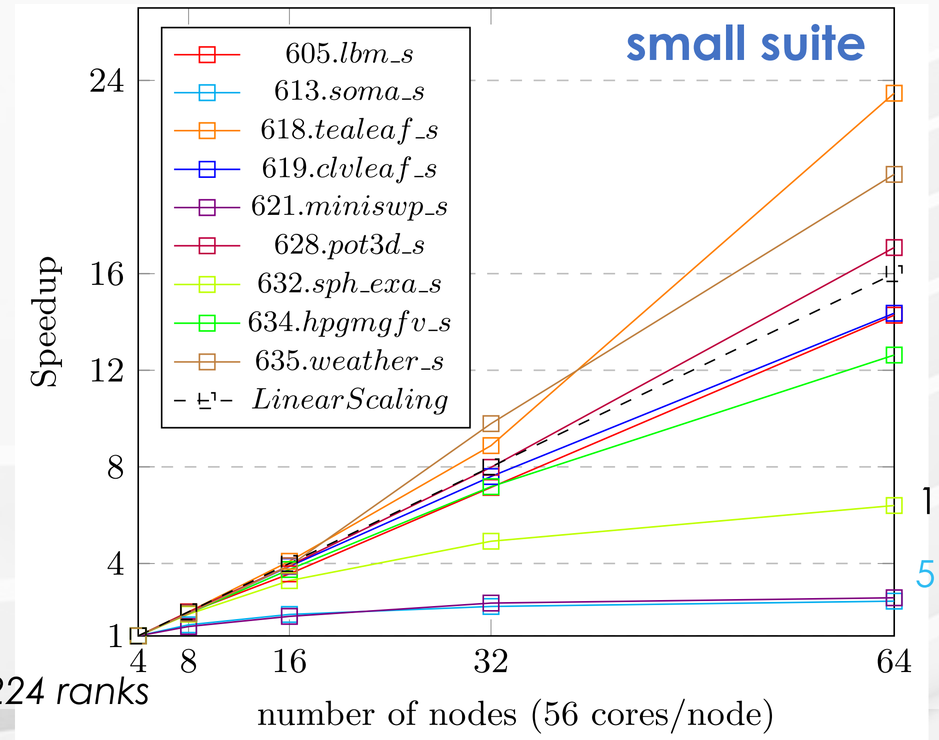

224 ranks

OMP approach improves runtime and scalability for soma + miniswp

14% Allreduce + Recv

5.5% Allreduce + Waitall

8% Allreduce

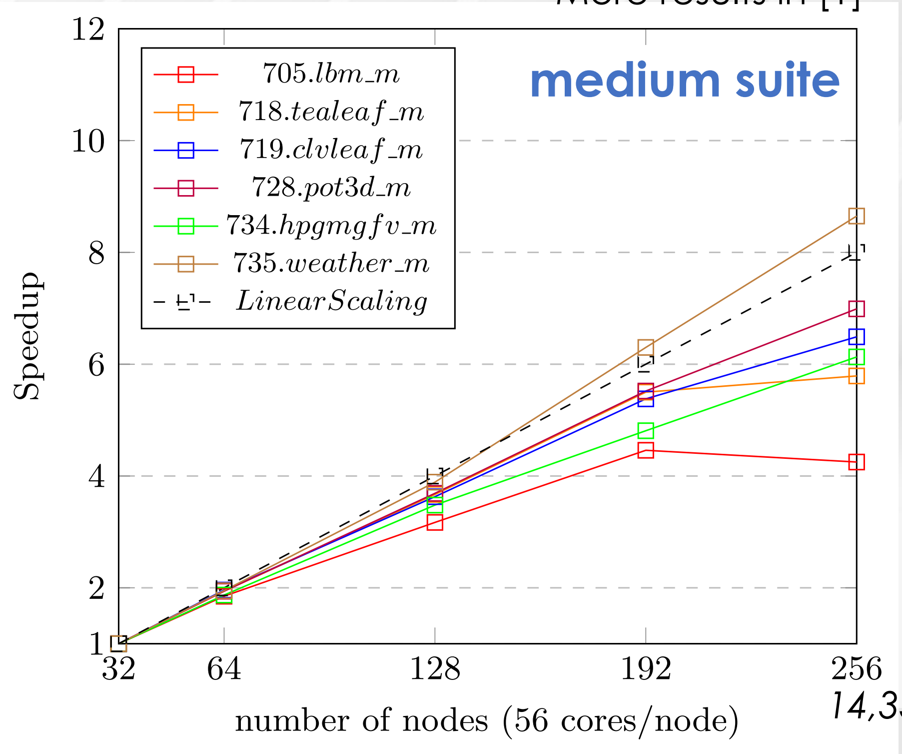

14,336 ranks

Scalability: GPUs

Setup
- Juwels Booster@JSC: 4x NVIDIA A100 GPUs
- GCC compiler, NVHPC, ParaStation MPI
- 1 MPI rank per GPU
- Workload: **large** suite

- Scalability runs on GPUs (ACC, TGT)
  - ACC and TGT scalability mostly good
  - ACC runtimes faster than TGT (except for Tealeaf)

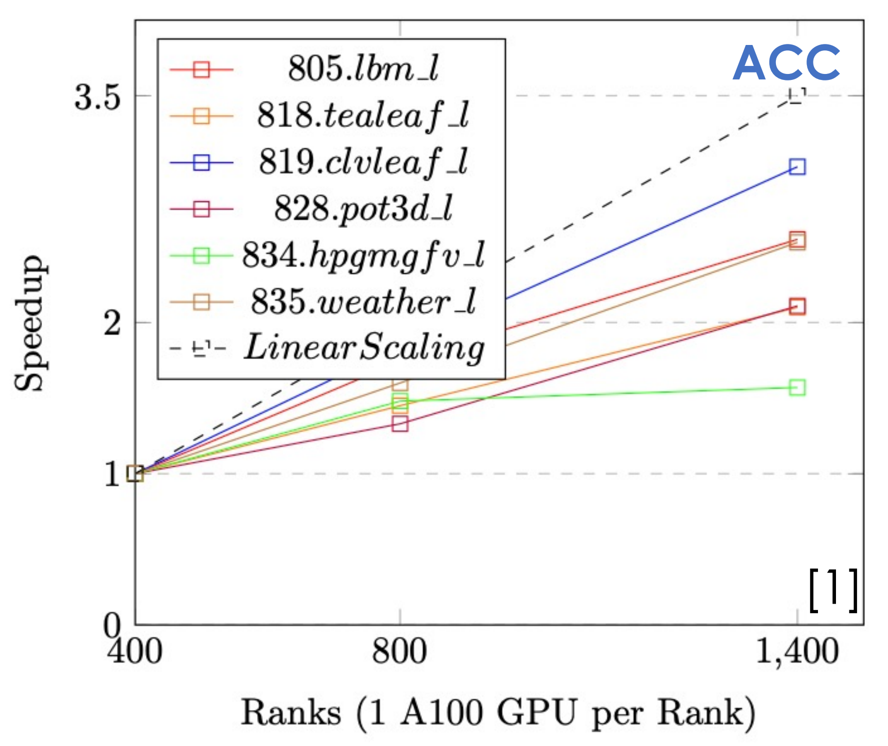

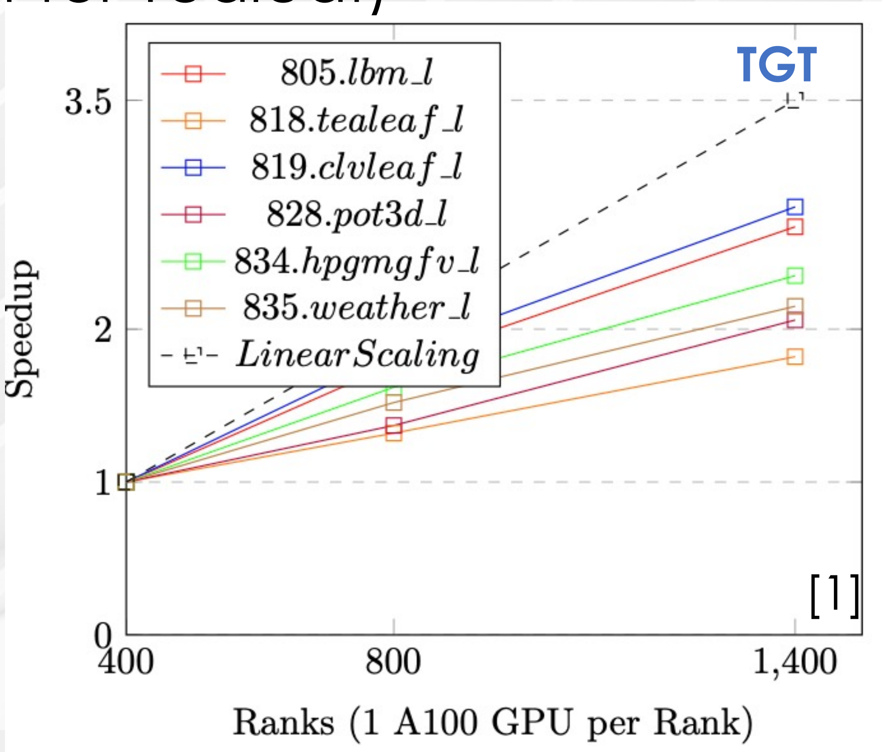

IU BigRed200, Slingshot 10

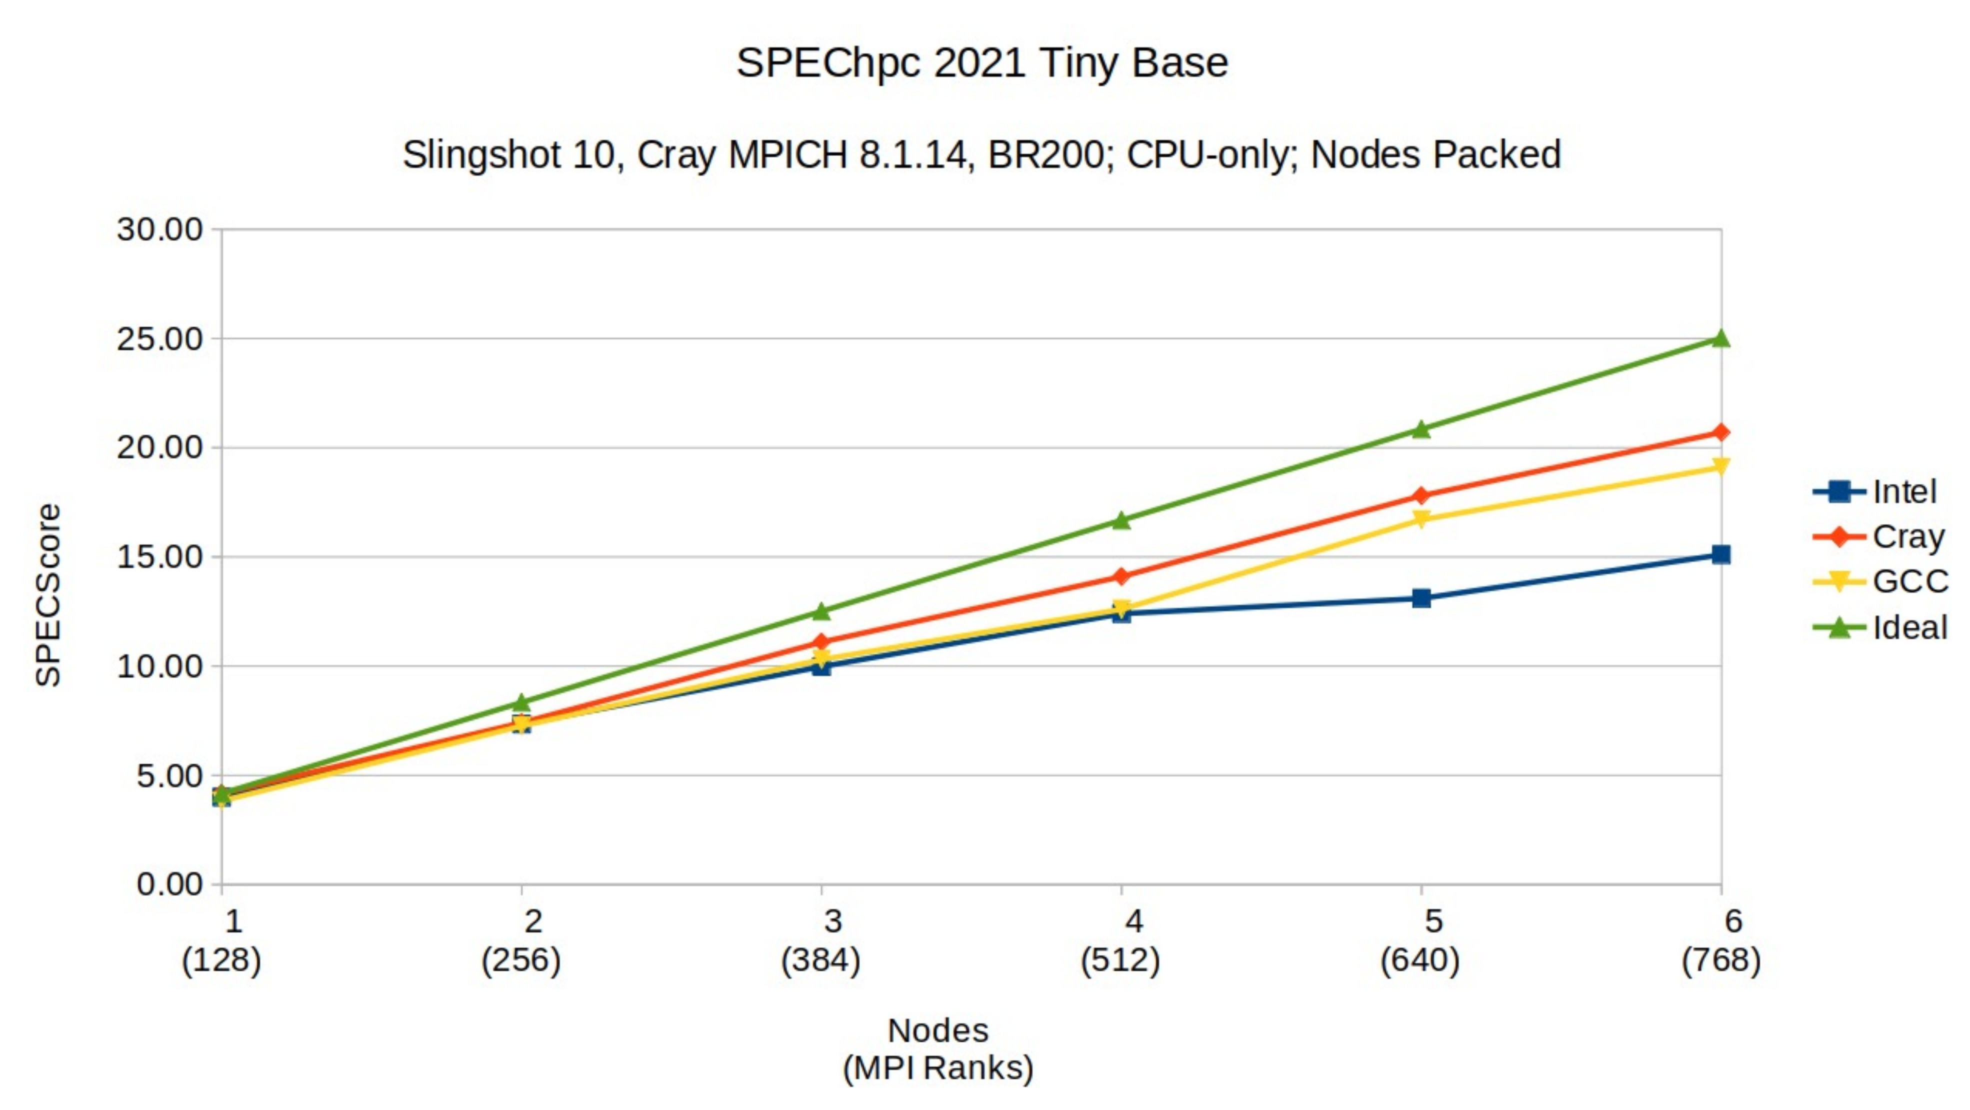

Community Impact

- SPEChpc 2021 benchmark suite instrumental to
  - Identifying compiler implementation inconsistencies
  - Determining ambiguities in OpenMP specification
  - Identifying compiler/runtime bugs
  - Critical for ECP SOLLVE (Scaling OpenMP With LLVm for Exascale Performance and Portability) project
  - Identifying of non-performing HPC nodes in large clusters in universities and centers
  - Comparing/contrasting machines and procurement using SPEC scores
  - Building the next-generation workforce who learn to use large clusters, job scheduling, build roofline modeling, created scalability plots etc.,

SPEC HPG at ORNL

- At ORNL, we have been using SPEC HPG benchmarks for:
  - Evaluating maturity of compilers on our production and pre-production systems
  - Evaluating readiness of the programming environment for acceptance testing of pre-production systems
  - Track performance of individual benchmarks over time and as the system software stack evolves
- SPEC ACCEL OMP and ACC were used for acceptance testing of Summit
- SPEChpc 2021 has been used for Summit's regression testing and for Frontier's acceptance testing

SPEC HPG at ORNL (cont'd)

- Have added SPEChpc 2021 to the OLCF Test Harness
- Using CPU only and GPU offloading to track performance
- Running with offloading using AMD and CCE toolchains

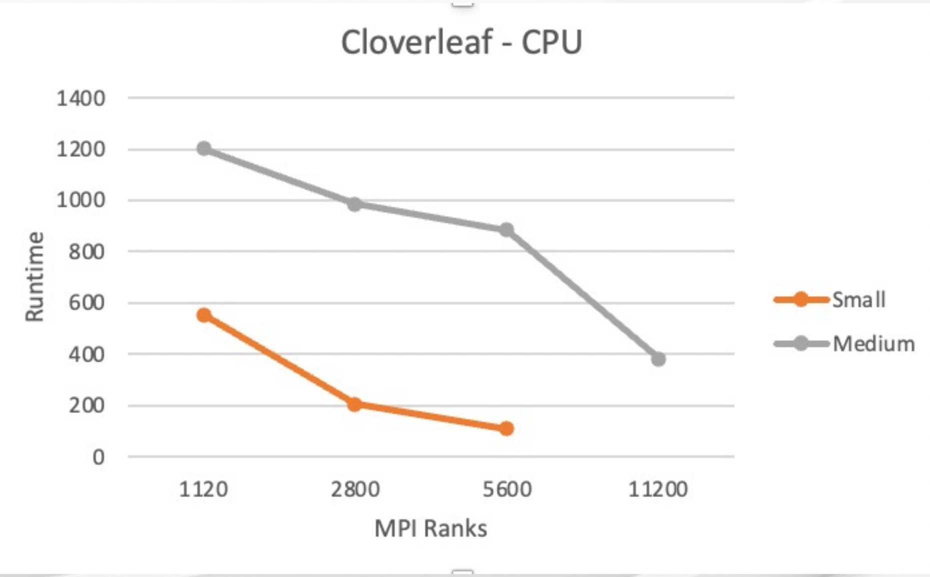

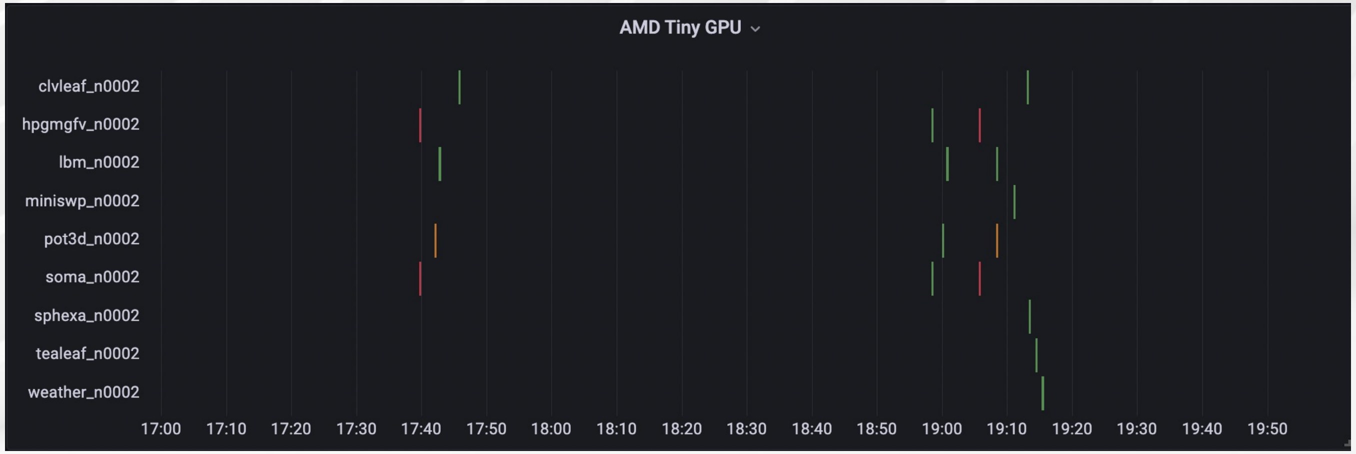

How to get SPEChpc and join the team

spec®

- Licenses are free for non-commercial:
  https://www.spec.org/hpgdownload.html
- Commercial Licenses available, full information found at:
  http://www.spec.org/hpc20

Future Developments

- SPEChpc Weak Scaling Suite (looking for codes now!)
- Refresh of SPEC ACCEL

Join SPEC HPG!

https://www.spec.org/spec/membership.html

Disclaimer & Attribution

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change. SPEC assumes no obligation to update or otherwise correct or revise this information. However, SPEC reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of SPEC to notify any person of such revisions or changes.

SPEC makes no representations or warranties with respect to the included contents and does not assume responsibility for any inaccuracies, errors, or omissions. In not event will SPEC be liable for any direct or indirect consequential damage from the use of any information contained herein, even if SPEC is expressly advised of the possibility of such damages.