Discussing First Results of the SPEC ACCEL OpenMP Suite with Target Directives

All the SPEC ACCEL Contributors

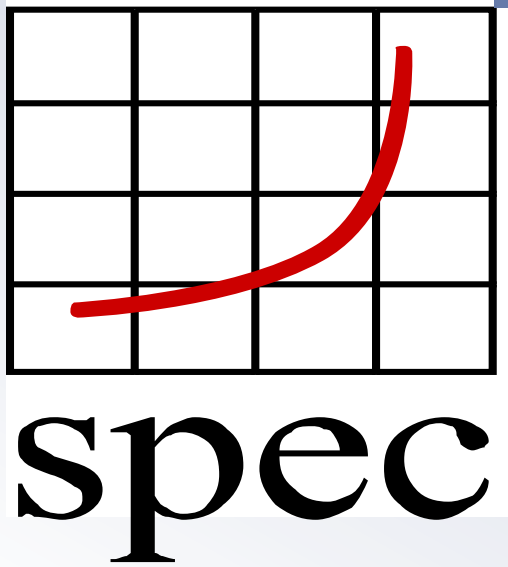

Guido Juckeland, Robert Henschel, Sunita Chandrasekaran, Sandra Wienke, Junjie Li, Alexander Bobyr, William Brantley, Mathew Colgrove, Oscar Hernandez, Arpith Jacob, Kalyan Kumaran, Dave Raddatz, Veronica Vergara Larrea, Bo Wang, Brian Whitney, Greg Rogers, Matthias Mueller and Andrey Naraikin

Contents

- Overview
- Intro to SPEC HPG
- The SPEC/HPG Benchmark Philosophy
- SPEC ACCEL 1.2 Benchmarks
- OpenMP experience
- Feedback

Standards Performance Evaluation Corporation (SPEC)

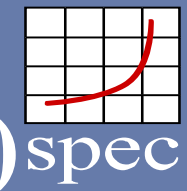

- SPEC is a non-profit corporation formed to "establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers"
- Composed of four groups
  - Graphics and Workstation Performance Group (GWPG)
  - High Performance Group (HPG)
  - Open Systems Group (OSG)
  - Research Group (RG)
- https://www.spec.org

SPEC Members

The SPEC Consortium: Members and Associates

SPEC Members:

*Acer Inc. * Action S.A. * Advanced Micro Devices * Amazon Web Services, Inc. * Apple Inc. * ARM * ASUSTeK Computer Inc. * Avere Systems * Bull S.A. * Cavium Inc. * Cisco Systems, Inc. * Dell, Inc. * E4 Computer Engineering SPA * EMC * Fujitsu * Gartner, Inc. * Hitachi Data Systems * Hitachi Ltd. * HP * Huawei Technologies Co. Ltd. * IBM * Inspur Corporation * Intel * Lenovo * Micron Technology, Inc. * Microsoft * NEC - Japan * NetApp * NVIDIA * Oracle * Panasas * Primary Data * Principled Technologies * Qualcomm Technologies Inc. * Quanta Computer Inc. * Red Hat * Samsung * SAP AG * Seagate * SGI * Sugon * Super Micro Computer, Inc. * SUSE * Symantec Corporation * Twitter, Inc. * Unisys * Via Technologies * VMware * Wipro Ltd. * ZTE Corporation **

SPEC Associates:

*Acadamia Sinica, Institute of Information Science * Argonne National Laboratory * Charles University * China Academy of Telecommunication Research * Department of Veterans Affairs - Corporate Data Center Operations * Dresden University of Technology ZIH * fortiss GmbH * Indiana University * Institute for Information Industry Taiwan * JAIST * Karlsruhe Institute of Technology * Leibniz Rechenzentrum - Germany * 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 * Tsinghua University * University of Aizu - Japan * University of California - Berkeley * University of Cologne * University of Houston * University of Illinois at Urbana-Champaign * University of Maryland * University of Miami * University of Pavia * University of Texas at Austin * University of Tsukuba * University of Wuerzburg * Virginia Polytechnic Institute and State University **

SPEC Research Group:

*Advanced Strategic Technology LLC * bankmark UG * Barcelona Supercomputing Center * Charles University * Cisco Systems * Cloudera, Inc * Compilaflows * Delft University of Technology * Dell * fortiss GmbH * Friedrich-Alexander-University Erlangen-Nuremberg * Goethe University Frankfurt, Big Data Lab * Hewlett-Packard * Huawei * IBM * Imperial College London * Indian Institute of Technology, Bombay * Institute for Information Industry, Taiwan * Institute of Communication and Computer Systems/NTUA * Intel * Karlsruhe Institute of Technology * Kiel University * MIOsoft Corporation * Mitre Corporation * NICTA * Nova-Tec Consulting GmbH * Oracle * Purdue University * Red Hat * RWTH Aachen University * SalesForce.com * San Diego Supercomputing Center * San Francisco State University * SAP AG * Siemens Corporation * Technische Universität Darmstadt * Technische Universität Dresden * The MITRE Corporation * Umea University * University of Alberta * University of Coimbra * University of Florence * University of Lugano * University of Minnesota * University of North Florida * University of Paderborn * University of Pavia * University of Stuttgart * University of Texas at Austin * University of Wuerzburg * VMware **

SPEC High Performance Group (HPG)

- Develops benchmarks to represent high-performance computing applications for standardized, cross-platform performance evaluation.
- Benchmarks
  - SPEC OMP2012
  - SPEC MPI2007
  - SPEC ACCEL

SPEC Benchmark Philosophy

- The result of a SPEC benchmark is always a SPEC score.
  - Higher is better
- This score is always in relation to a reference machine.
  - Each benchmark has its own reference machine
  - For example, for SPEC ACCEL 1.2, the reference machine is:

SGI C 3108-TY11 with accelerator NVIDIA Tesla c2070

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 suits
  - Benchmark suite — i.e. SPEC ACCEL
  - Benchmark — i.e. OpenMP
  - Dataset size — i.e. Medium
  - Component — i.e. 550.md

SPEC Benchmark Philosophy cont'd

- Benchmarks support "Base" and "Peak" configuration
  - These yield separate SPEC scores.
  - "Peak" runs allow for more freedom.
- Base runs
  - The same compiler switches for all components
  - The same parallelism
  - Only portability switches allowed

SPEC ACCEL 1.2 OMP uses some portability switches

SPEC Power

- 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


```
me,03-10-2014 14:14:20.458,Watts,308.159000,Volts,120.442000,Amps,2.633360,PF,0.971602
me,03-10-2014 14:14:21.457,Watts,308.196000,Volts,120.513000,Amps,2.632220,PF,0.971564
me,03-10-2014 14:14:22.458,Watts,308.445000,Volts,120.520000,Amps,2.634020,PF,0.971622
me,03-10-2014 14:14:23.458,Watts,308.497000,Volts,120.530000,Amps,2.634310,PF,0.971602
me,03-10-2014 14:14:24.457,Watts,308.460000,Volts,120.521000,Amps,2.634150,PF,0.971614
-10-2014 14:14:29.833: Response to client sent: Time,03-10-2014 14:14:25.458,Watts,308.751000,Volts,120.
1000,Amps,2.636360,PF,0.971645
me,03-10-2014 14:14:26.458,Watts,308.834000,Volts,120.519000,Amps,2.637210,PF,0.971678
me,03-10-2014 14:14:27.457,Watts,308.958000,Volts,120.524000,Amps,2.638180,PF,0.971677
me,03-10-2014 14:14:28.458,Watts,309.092000,Volts,120.529000,Amps,2.639180,PF,0.971690
-10-2014 14:14:52.250: Go with mark '304.olbm'
-10-2014 14:14:52.250: Response to client sent: Starting untimed measurement, maximum 500000 samples at
00ms with 0 rampup samples
-10-2014 14:14:52.747,Watts,163.014000,Volts,120.827000,Amps,1.417630,PF,0.951701,Mark,304.olbm
-10-2014 14:14:53.746,Watts,153.447000,Volts,120.836000,Amps,1.336960,PF,0.949820,Mark,304.olbm
-10-2014 14:14:54.747,Watts,166.402000,Volts,120.858000,Amps,1.457690,PF,0.944537,Mark,304.olbm
-10-2014 14:14:55.747,Watts,199.812000,Volts,120.779000,Amps,1.730290,PF,0.956115,Mark,304.olbm
-10-2014 14:14:56.746,Watts,253.347000,Volts,120.582000,Amps,2.188800,PF,0.959902,Mark,304.olbm
-10-2014 14:14:57.747,Watts,265.319000,Volts,120.545000,Amps,2.278210,PF,0.966108,Mark,304.olbm
-10-2014 14:14:58.747,Watts,265.822000,Volts,120.550000,Amps,2.282260,PF,0.966183,Mark,304.olbm
-10-2014 14:14:59.746,Watts,261.949000,Volts,120.607000,Amps,2.253970,PF,0.963600,Mark,304.olbm
-10-2014 14:15:00.747,Watts,266.445000,Volts,120.558000,Amps,2.287140,PF,0.966313,Mark,304.olbm
-10-2014 14:15:01.747,Watts,261.673000,Volts,120.550000,Amps,2.252590,PF,0.963624,Mark,304.olbm
-10-2014 14:15:02.746,Watts,266.597000,Volts,120.577000,Amps,2.288130,PF,0.966299,Mark,304.olbm
```

Benchmark Development Process

- Group effort, with lots of discussions
- Final decisions are by vote, even though we strive for consensus
- Technical and managerial parts
  - Find benchmark components and define run rules
- Using SPEC provided tools
  - SVN, harness, “common rules”
  - Websites, mailing lists, meeting venues

SPEC ACCEL

- Provide a benchmark that target accelerators
- SPEC ACCEL 1.2 just got released: June, 2017
- Use different programming models: OpenCL 1.1, OpenACC 1.0 and OpenMP 4.5
- Portable across architectures (host, XeonPhi, GPUs, etc)
- A benchmark that works with at least two compilers
- Program one accelerator
- Helps evaluate performance portability of compilers and directives/languages

Active members in SPEC ACCEL:

- NVIDIA, SGI, Intel, IBM, AMD, Argonne, ORNL, HZDR, Oracle, University of Houston, University of Virginia, Aachen University, University of Illinois, Indiana University, TU-Dresden

*Present at the DOE workshop

OpenMP 4.5 – Performance Portability

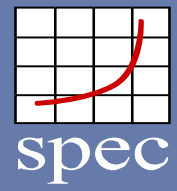

Consensus and community discussion on on how to write “performance portable” style in OpenMP 4+

- Initially members had different views
  - Experimented with different programming styles:
    - SIMD centric
    - Using nested parallelism
    - Using target combined directives with collapsed loops
  - Ported the benchmarks from OpenACC 1.0
- We agreed on some “guidelines” on how to write performance portable code
- We used these “guidelines” and successfully parallelized the 15 benchmarks with OpenMP 4.0
- While writing the benchmarks, OpenMP 4.5 came out
- We re-evaluated these “guidelines” to target OpenMP 4.5
- Tested different compilers, platforms, etc.

SPEC ACCEL 1.2: OpenMP 4.5

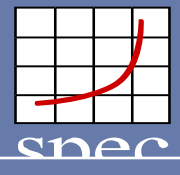

OpenMP Benchmarks	Language	Origin	Domain
503.ostencil	C	Parboil, University of Illinois	Thermodynamics
504.olbm	C	Parboil, University of Illinois	CFDm Lattice Boltzmann
514.omriq	C	Rodinia, University of Virginia	Medicine
550.md	Fortran	Indiana University	Molecular Dyn.
551.palm	Fortran	Leibniz University of Hannover	Large-eddy sim.
552.ep	C	NAS Parallel Benchmarks (NPB)	Embarrassing P.
553.clvrleaf	C, Fortran	Atomic Weapons Establishments	Hydrodynamics
554.cg	C	NPB	Conjugate Grad.
555.seismic	Fortran	GeoDynamics.org	Seismic Wave Modeling (PDE)
556.sp	Fortran	NPB	Scalar Peta-d solv
557.csp	C	NPB	Scalar Peta-d solv
559.miniGhost	C, Fortran	Sandia National Laboratory	Finite difference
560.ilbdc	Fortran	SPEC OMP2012	Fluid Mechanics
563.swim	Fortran	SPEC OMP2012	Weather
570.bt	C	NPB	BTS 3D PDE

Guidelines – To write OpenMP 4.5 “Performance Portable Style”

Use OpenMP 4.5 “Accelerator Model”

Rely on compilers to generate implementation specific values for a given architecture:

- # of teams
- # thread_limit,
- # of threads – in parallel regions
- SIMD length
- dist_schedule – in distribute
- loop schedules – in parallel do

Compiler implementers should pick these values to enable performance portability. HPC compilers should generate platform specific optimizations.

OpenMP 4.5 "Performance Portable Style"

For level-1 parallel loop:

- #pragma target teams distribute parallel for simd

For perfectly nested parallel loops:

- Collapsed loops with the combined directive

```
#pragma omp target teams distribute parallel for collapse(N)
  for(i=0;….)
    for(j=0;….)
      for(k=0;…)
```

For loopnests that we could not collapse:,

- Parallelize the outer loops with teams distribute, parallel for
- Parallelize the inner loops always with SIMD

Using OpenMP 4.5 the user and programmer –

- We used combined directive and / collapse to specify the parallel potential on a given platform.

Experiences with OpenMP 4.5

Differences between OpenMP 4.0 and 4.5 that affected our porting:

Reductions

- In OpenMP 4.5 the default for scalars is firstprivate as a result

  reduction variables need to be mapped to/from

```
#pragma omp target map(tofrom:sum)
#pragma omp teams distribute parallel for reduction(+:sum)
    for(…. )
        sum = sum + ….
```

Privatization

- Avoid over-privatizing on multiple levels
- We should only privatize only at the parallelism level that is needed

```
#pragma omp teams distribute parallel for  // private(yy, zz)
for(i= …. )
   for(j= … )
#pragma omp simd private(yy,zz)
      for(z= …
            yy =
            zz =
```

Guidelines for OpenMP 4+

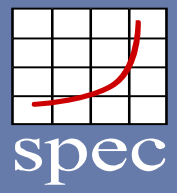

Only merge target regions when is safe to do so.

From:

```
#pragma omp target teams distribute parallel for
    for(i=…)
        c[i] =
#pragma omp target teams distribute parallel for
    for(i=…)
        b[i] =  a[i] + ….
```

To:

```
#pragma omp target teams
#pragma omp distribute
        for(i=…)
            c[i] =
#pragma omp distribute
        for(i=…)
            b[i] = a[i] + ….   // if c were a, then we will have a race condition.
```

Test your code using multiple platforms / compilers

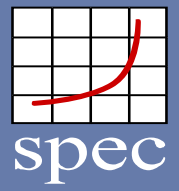

Implementation specific issues:

- Some compiler implementations may pick num_teams equal to one.
  - E.g. for some platforms: e.g. Xeon Phi, etc
- Other implementations may pick num_threads equal to one, when a parallel for is nested inside a teams distribute and when it contains “simd”
- Compiler implementations may have different strategies to implement “simd” on a given platform
  - Some compiler implementations may be ignoring “simd” or tuning on autovec

This makes debugging hard

- You need to test on multiple platforms to guarantee correctness

SPEC ACCEL 1.2 OpenMP results:

All ACCEL Results Published by SPEC

These results have been submitted to SPEC; see the disclaimer before studying any results.

Last update: *Wednesday, 6 September 2017, 18:07*

OpenMP (16):

Test Sponsor	System Name	Accelerator Name	Results		Energy	
			Base	Peak	Base	Peak
HZDR	ASUS ESC4000 G3 Series HTML | CSV | Text | PDF | PS | Config	Intel Xeon CPU E5-2630 v3	1.97	Not Run	--	--
Indiana University	Ninja Developer Platform Pedestal: Liquid Cooled HTML | CSV | Text | PDF | PS | Config	Xeon Phi 7210	0.186	Not Run	0.356	--
Indiana University	Ninja Developer Platform Pedestal: Liquid Cooled HTML | CSV | Text | PDF | PS | Config	Xeon Phi 7210	3.40	Not Run	4.54	--
Indiana University	Ninja Developer Platform Pedestal: Liquid Cooled HTML | CSV | Text | PDF | PS | Config	Xeon Phi 7210	1.62	Not Run	2.26	--
Indiana University	HP Z820 Workstation HTML | CSV | Text | PDF | PS | Config	Intel Xeon E5-2640 v2	0.462	Not Run	0.673	--
Indiana University	HP Z820 Workstation HTML | CSV | Text | PDF | PS | Config	Intel Xeon E5-2640 v2	1.04	Not Run	1.42	--
Indiana University	Lenovo NeXtScale nx360 M5 HTML | CSV | Text | PDF | PS | Config	Intel Xeon E5-2680 v3	2.32	Not Run	--	--
Indiana University	Cray XC30 HTML | CSV | Text | PDF | PS | Config	Intel Xeon E5-2697 v2	1.75	Not Run	--	--
Intel	Endeavour Node(Intel Xeon E5-2697 v4, 2.3GHz, DDR4-2400 MHz, SMT ON, Turbo OFF) HTML | CSV | Text | PDF | PS | Config	Intel Xeon E5-2697 v4	3.29	3.56	--	--
Intel	Endeavour Node(Intel Xeon E5-2697 v4, 2.3GHz, DDR4-2400 MHz, SMT ON, Turbo ON) HTML | CSV | Text | PDF | PS | Config	Intel Xeon E5-2697 v4	3.49	3.76	--	--
Intel	Endeavour Node(Intel Xeon Phi CPU 7250F, 1.40 GHz, SMT ON, Turbo ON, flat MCDRAM) HTML | CSV | Text | PDF | PS | Config	Intel Xeon Phi CPU 7250F	6.14	7.55	--	--
Intel	Intel Server System R2208WFTZS (2 x Intel Xeon Gold 6148, 2.40 GHz, SMT ON, Turbo ON) Endeavour Node HTML | CSV | Text | PDF | PS | Config	Intel Xeon Gold 6148	5.16	5.61	--	--
Test Sponsor	**System Name**	**Accelerator Name**	**Results**		**Energy**	
			Base	**Peak**	**Base**	**Peak**
Intel	Intel Server System R2208WFTZS (2 x Intel Xeon Platinum 8180, 2.50 GHz, SMT ON, Turbo ON) HTML | CSV | Text | PDF | PS | Config	Intel Xeon Platinum 8180	5.89	6.45	--	--
Oak Ridge National Laboratory	Cray XC40 Intel Xeon Phi series HTML | CSV | Text | PDF | PS | Config	Intel Xeon Phi CPU 7230	3.59	Not Run	--	--
RWTH Aachen	Intel Xeon SandyBridge EN/EP processor (8 core, 2.00 GHz, Intel Xeon CPU E5-2650) HTML | CSV | Text | PDF | PS | Config	Intel Xeon Phi 5110p	0.991	Not Run	0.989	--
Technische Universitaet Dresden	Intel Server System LADMP00AP Family (Xeon Phi 7210, 1.3 GHz, 64 cores, 4 threads) HTML | CSV | Text | PDF | PS | Config	Intel Xeon Phi 7210	4.39	6.08	--	--

Last update: *Wednesday, 6 September 2017, 18:07*

Example of a published result:

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 sponsor: Technische Universitaet Dresden
Tested by: Technische Universitaet Dresden

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

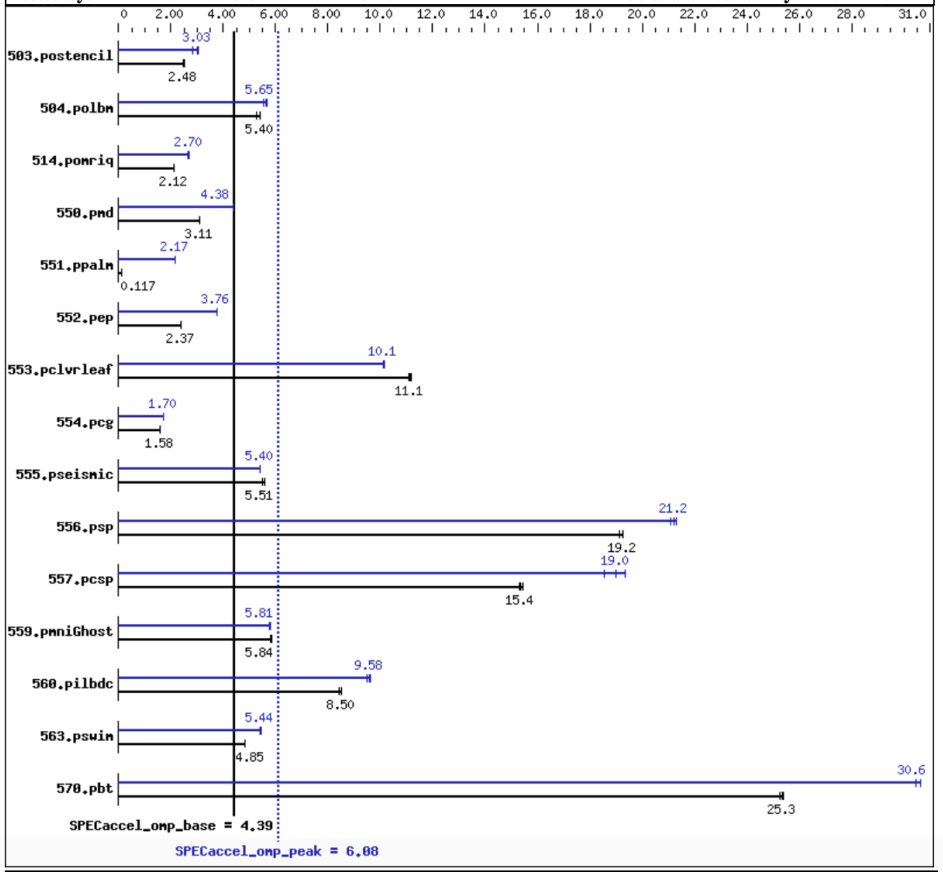

Hardware

CPU Name:	Intel Xeon Phi 7210
CPU Characteristics:	Intel Turbo Boost 2 Technology up to 1.50 GHz
CPU MHz:	1300
CPU MHz Maximum:	1500
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

Accel Model Name:	Xeon Phi 7210
Accel Vendor:	Intel
Accel 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

Preliminary results are showing

- Benchmark is very new
- Only 16 reportable results (compared to 100s from other SPEC/HPG benchamrks)
- No official results yet for OpenMP 4.5 offload on GPUs
  - Major milestone may be hit by the end of the year where we will have reportable results
  - IBM OpenMP 4.5 compiler are still beta
  - Some vendors got access to the benchmarks recently (outside the HPG team) and are fixing compiler bugs.
- Compilers need optimize for each target architecture – unless they expect the user to auto-tune
- Compilers are still working on their OpenMP optimizations and few support multiple architectures.
  - e.g. We need OpenMP compiler that supports a variety of architectures: e.g. multicore, Xeon Phi and GPUs using target regions to do better performance portable studies.

Result Submission Process

- Obtain and install the benchmark
- Perform a valid run
- Supply hardware and software description
- Submit result for review (and publication) to SPEC HPG
  - 2 week review process
  - (Define embargo period)
- Use the result as you would like

SPEC MPI Accelerator Search Program

- Combine SPEC MPI and SPEC ACCEL benchmark suites into one.
  - Facilitate support for homogeneous and heterogeneous multi-node systems
  - MPI+X, where "X" is a node-level parallel programming model such as OpenACC, OpenMP, CUDA, OpenCL, Kokkos, TBB, etc.
  - Compute intensive
- Focus on hybrid programming with accelerators
  - Being able to run on a single nodes or up to N nodes.
  - Support one or M accelerators per MPI rank (multiple-accelerator programming)
- Represents the state of the art for the given field
  - Results needs to be validated numerically
- Submissions that are in included in the benchmark suite will receive up to $5,000 and a licensed copy of the new benchmark suite when released.
- More information: https://www.spec.org/hpg/search

Thank you!

Questions?