SPEC ACCEL™ ACC Result

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)

Tesla V100-SMX2-16GB
PowerEdge C4140 Server

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 13.3

ACCEL license: 3440A
Test sponsor: Indiana University
Tested by: Indiana University

Test date: Aug-2019
Hardware Availability: May-2019
Software Availability: Apr-2019

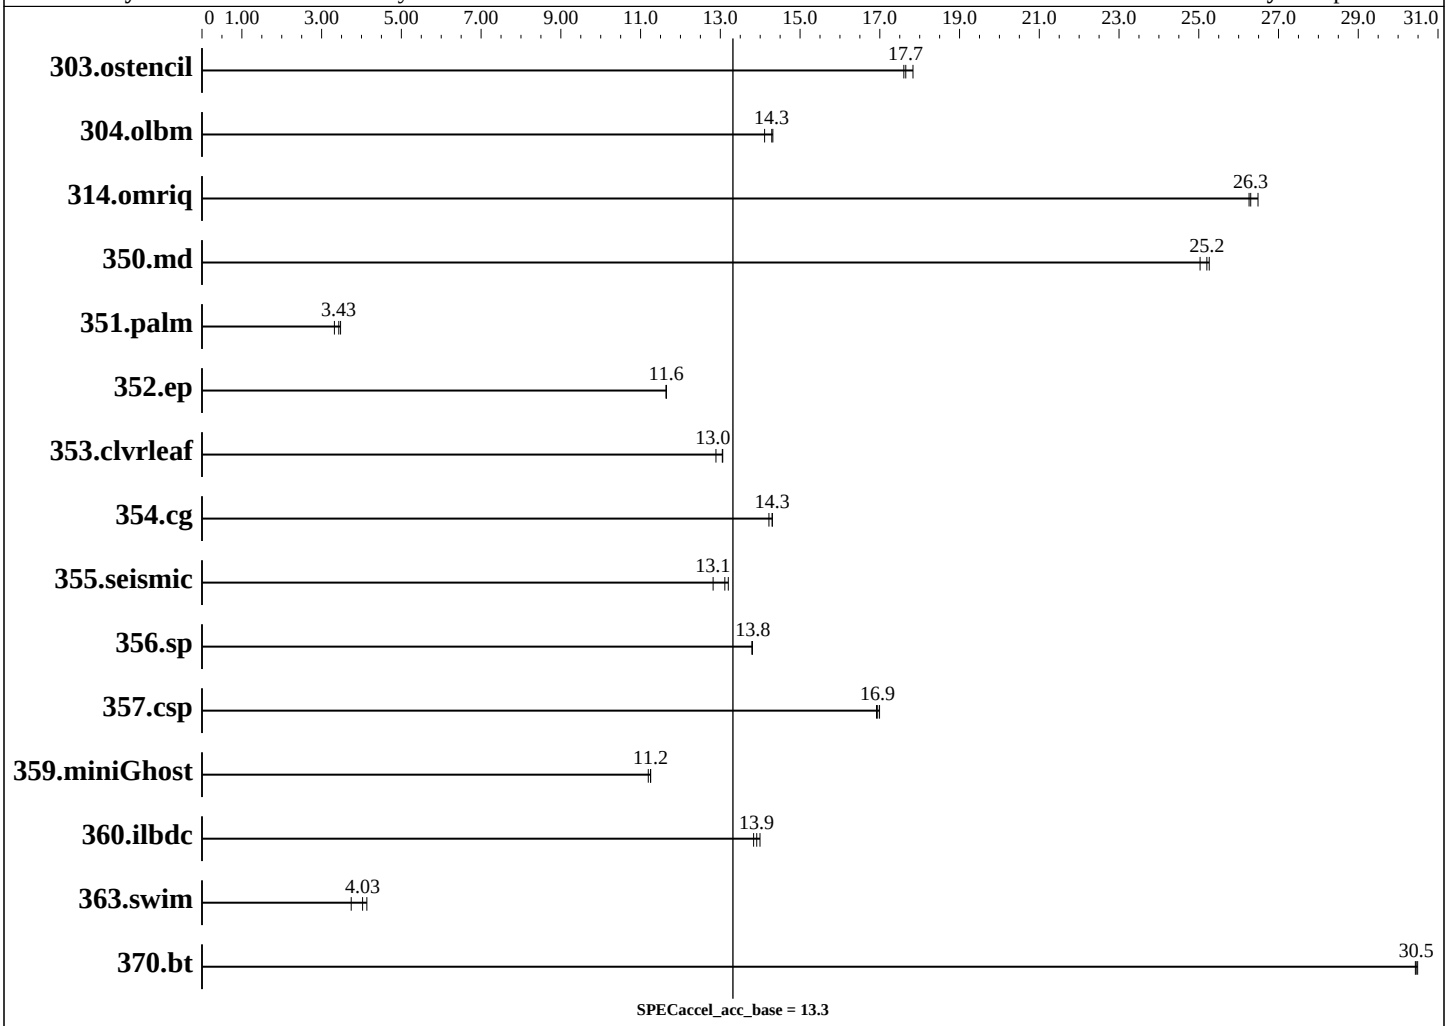

Hardware

CPU Name:	Intel Xeon Gold 6130
CPU Characteristics:	Intel Turbo Boost on, SMT off
CPU MHz:	2100
CPU MHz Maximum:	3700
FPU:	Integrated
CPU(s) enabled:	32 cores, 2 chips, 16 cores/chip
CPU(s) orderable:	1,2 chips
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per core
L3 Cache:	22 MB I+D on chip per chip
Other Cache:	None

Continued on next page

Accelerator

Accel Model Name:	Tesla V100
Accel Vendor:	NVIDIA Corporation
Accel Name:	Tesla V100-SMX2-16GB
Type of Accel:	GPU
Accel Connection:	PCIe
Does Accel Use ECC:	Yes
Accel Description:	V100-SMX2-16GB with NVLink (Persistence Mode enabled)
Accel Driver:	NVIDIA UNIX x86_64 Kernel Module 418.67

SPEC ACCEL ACC Result

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)

Tesla V100-SMX2-16GB
PowerEdge C4140 Server

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 13.3

ACCEL license: 3440A
Test sponsor: Indiana University
Tested by: Indiana University

Test date: Aug-2019
Hardware Availability: May-2019
Software Availability: Apr-2019

Hardware (Continued)

Memory: 256 GB (16 x 16 GB 2Rx8 PC4-2666V-R)
Disk Subsystem: None
Other Hardware: None

Software

Operating System: Red Hat Enterprise Linux Server release 7.6 (Maipo) 3.10.0-957.21.3.el7.x86_64
Compiler: PGI Community Edition, Release 19.4
File System: cephfs nfsv4 (ganesha) over 100Gbits/s Ethernet
System State: Run level 3 (multi-user)
Other Software: CUDA 10.1

Results Table

Benchmark	Base						Peak					
	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
303.ostencil	**8.21**	**17.7**	8.13	17.8	8.24	17.6						
304.olbm	32.2	14.1	**31.8**	**14.3**	31.8	14.3						
314.omriq	36.1	26.5	36.4	26.3	**36.3**	**26.3**						
350.md	**10.0**	**25.2**	9.98	25.3	10.1	25.0						
351.palm	106	3.48	**108**	**3.43**	111	3.32						
352.ep	45.5	11.6	45.5	11.6	**45.5**	**11.6**						
353.clvrleaf	34.1	13.1	**34.1**	**13.0**	34.5	12.9						
354.cg	28.7	14.2	28.5	14.3	**28.5**	**14.3**						
355.seismic	28.0	13.2	**28.2**	**13.1**	28.9	12.8						
356.sp	20.0	13.8	**20.0**	**13.8**	20.0	13.8						
357.csp	16.0	16.9	**15.9**	**16.9**	15.9	17.0						
359.miniGhost	33.0	11.2	32.8	11.3	**32.8**	**11.2**						
360.ilbdc	26.5	13.8	26.2	14.0	**26.4**	**13.9**						
363.swim	55.6	4.13	**57.1**	**4.03**	61.5	3.74						
370.bt	7.33	30.4	7.31	30.5	**7.32**	**30.5**						

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

Platform Notes

```
Sysinfo program /home/lijunj/junjie_benchmarks/spec/accel-1.2-bm/Docs/sysinfo
$Rev: 6965 $ $Date:: 2015-04-21 #$ c05a7f14b1b1765e3fe1df68447e8a35
running on r08g05 Wed Aug  7 02:05:24 2019

This section contains SUT (System Under Test) info as seen by
some common utilities.  To remove or add to this section, see:
  http://www.spec.org/accel/Docs/config.html#sysinfo
```

Continued on next page

Platform Notes (Continued)

```
From /proc/cpuinfo
   model name : Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz
      2 "physical id"s (chips)
      32 "processors"
   cores, siblings (Caution: counting these is hw and system dependent.  The
   following excerpts from /proc/cpuinfo might not be reliable.  Use with
   caution.)
      cpu cores : 16
      siblings  : 16
      physical 0: cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
      physical 1: cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
   cache size : 22528 KB

From /proc/meminfo
   MemTotal:       196261708 kB
   HugePages_Total:       0
   Hugepagesize:       2048 kB

From /etc/*release* /etc/*version*
   os-release:
      NAME="Red Hat Enterprise Linux Server"
      VERSION="7.6 (Maipo)"
      ID="rhel"
      ID_LIKE="fedora"
      VARIANT="Server"
      VARIANT_ID="server"
      VERSION_ID="7.6"
      PRETTY_NAME="Red Hat Enterprise Linux Server 7.6 (Maipo)"
   redhat-release: Red Hat Enterprise Linux Server release 7.6 (Maipo)
   system-release: Red Hat Enterprise Linux Server release 7.6 (Maipo)
   system-release-cpe: cpe:/o:redhat:enterprise_linux:7.6:ga:server

uname -a:
   Linux r08g05 3.10.0-957.21.3.el7.x86_64 #1 SMP Fri Jun 14 02:54:29 EDT 2019
   x86_64 x86_64 x86_64 GNU/Linux

run-level 3 Jul 15 10:51

SPEC is set to: /home/lijunj/junjie_benchmarks/spec/accel-1.2-bm
   Filesystem                                                      Type
   Size  Used Avail Use% Mounted on
   172.16.129.177:/volumes/_nogroup/24ec4401-f96d-40a5-99a5-e96f73257d2f nfs4
   128G  102G   27G  79% /mnt
Additional information from dmidecode:

   Warning: Use caution when you interpret this section. The 'dmidecode' program
   reads system data which is "intended to allow hardware to be accurately
   determined", but the intent may not be met, as there are frequent changes to
```

Continued on next page

Platform Notes (Continued)

```
    hardware, firmware, and the "DMTF SMBIOS" standard.

  Information from pgaccelinfo:
  CUDA Driver Version:            10010
  NVRM version:                   NVIDIA UNIX x86_64 Kernel Module  418.67
  Device Number:                  0
  Device Name:                    Tesla V100-SXM2-16GB
  Device Revision Number:         7.0
  Global Memory Size:             16914055168
  Number of Multiprocessors:      80
  Concurrent Copy and Execution:  Yes
  Total Constant Memory:          65536
  Total Shared Memory per Block:  49152
  Registers per Block:            65536
  Warp Size:                      32
  Maximum Threads per Block:      1024
  Maximum Block Dimensions:       1024, 1024, 64
  Maximum Grid Dimensions:        2147483647 x 65535 x 65535
  Maximum Memory Pitch:           2147483647B
  Texture Alignment:              512B
  Clock Rate:                     1530 MHz
  Execution Timeout:              No
  Integrated Device:              No
  Can Map Host Memory:            Yes
  Compute Mode:                   default
  Concurrent Kernels:             Yes
  ECC Enabled:                    No
  Memory Clock Rate:              877 MHz
  Memory Bus Width:               4096 bits
  L2 Cache Size:                  6291456 bytes
  Max Threads Per SMP:            2048
  Async Engines:                  5
  Unified Addressing:             Yes
  Managed Memory:                 Yes
  Concurrent Managed Memory:      Yes
  Preemption Supported:           Yes
  Cooperative Launch:             Yes
    Multi-Device:                 Yes
  PGI Default Target:             -ta=tesla:cc70
```

General Notes

```
  Four V100-SMX2-16GB GPUs were installed on the system,
  three of them were idle while only one is active for running SPEC Accel.

  CPUs and GPUs are connected via PCIe, while the four GPUs are connected
  through NVLink.  The NVlink connection is not used in this test.
```

Continued on next page

General Notes (Continued)

```
Stacksize set to 'unlimited':
     ulimit -s unlimited

Spectre & Meltdown:
  Yes: The test sponsor attests, as of date of publication, that CVE-2017-5754 (Meltdown)
   is mitigated in the system as tested and documented.
  Yes: The test sponsor attests, as of date of publication, that CVE-2017-5753 (Spectre variant 1)
   is mitigated in the system as tested and documented.
  Yes: The test sponsor attests, as of date of publication, that CVE-2017-5715 (Spectre variant 2)
   is mitigated in the system as tested and documented.
```

Base Compiler Invocation

C benchmarks:
```
pgcc
```

Fortran benchmarks:
```
pgfortran
```

Benchmarks using both Fortran and C:
```
pgcc pgfortran
```

Base Optimization Flags

C benchmarks:
```
-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1
```

Fortran benchmarks:
```
-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1
```

Benchmarks using both Fortran and C:

353.clvrleaf: `-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1`

359.miniGhost: `-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1 -Mnomain`

The flags file that was used to format this result can be browsed at
https://www.spec.org/accel/flags/pgi2019_flags.html

You can also download the XML flags source by saving the following link:
https://www.spec.org/accel/flags/pgi2019_flags.xml

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)
Tesla V100-SMX2-16GB
PowerEdge C4140 Server

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 13.3

ACCEL license:	3440A	**Test date:**	Aug-2019
Test sponsor:	Indiana University	**Hardware Availability:**	May-2019
Tested by:	Indiana University	**Software Availability:**	Apr-2019

SPEC ACCEL is a trademark of the Standard Performance Evaluation Corporation. All other brand and product names appearing in this result are trademarks or registered trademarks of their respective holders.

For questions about this result, please contact the tester.
For other inquiries, please contact webmaster@spec.org.

Tested with SPEC ACCEL v1.2.
Report generated on Thu Oct 24 12:55:48 2019 by SPEC ACCEL PS/PDF formatter v1290.
Originally published on 24 October 2019.