SPEChpc™ 2021 Small Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

NVIDIA Corporation

DGX A100 (AMD EPYC 7742, Tesla A100-SXM-80GB)

SPEChpc 2021_sml_base = 7.78

SPEChpc 2021_sml_peak = 8.30

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jul-2020
Software Availability: Sep-2021

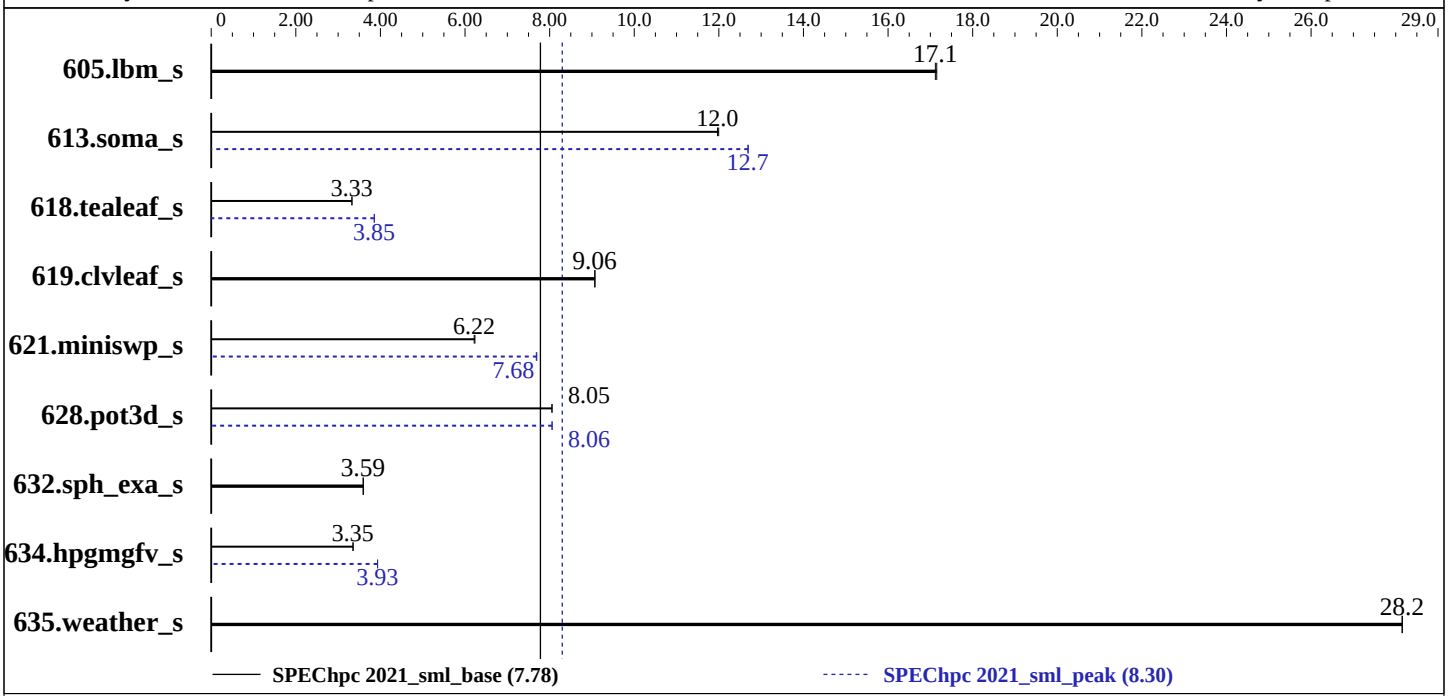

Results Table

Benchmark	Base									Peak								
	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
605.lbm_s	ACC	8	1	**90.5**	**17.1**	90.4	17.2			ACC	8	1	**90.5**	**17.1**	90.4	17.2		
613.soma_s	ACC	8	1	133	12.0	**134**	**12.0**			ACC	8	1	**126**	**12.7**	126	12.7		
618.tealeaf_s	ACC	8	1	**616**	**3.33**	616	3.33			ACC	8	1	**532**	**3.85**	532	3.86		
619.clvleaf_s	ACC	8	1	182	9.07	**182**	**9.06**			ACC	8	1	182	9.07	**182**	**9.06**		
621.miniswp_s	ACC	8	1	176	6.23	**177**	**6.22**			ACC	8	1	143	7.70	**143**	**7.68**		
628.pot3d_s	ACC	8	1	208	8.06	**208**	**8.05**			ACC	8	1	208	8.06	**208**	**8.06**		
632.sph_exa_s	ACC	8	1	639	3.60	**641**	**3.59**			ACC	8	1	639	3.60	**641**	**3.59**		
634.hpgmgfv_s	ACC	8	1	291	3.36	**291**	**3.35**			ACC	8	1	248	3.93	**248**	**3.93**		
635.weather_s	ACC	8	1	92.3	28.2	**92.4**	**28.2**			ACC	8	1	92.3	28.2	**92.4**	**28.2**		

SPEChpc 2021_sml_base = 7.78

SPEChpc 2021_sml_peak = 8.30

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

SPEChpc™ 2021 Small Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

NVIDIA Corporation

DGX A100 (AMD EPYC 7742, Tesla A100-SXM-80GB)

SPEChpc 2021_sml_base = 7.78

SPEChpc 2021_sml_peak = 8.30

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jul-2020
Software Availability: Sep-2021

Hardware Summary

Type of System:	SMP
Compute Node:	DGX A100
Interconnect:	None
Compute Nodes Used:	1
Total Chips:	2
Total Cores:	128
Total Threads:	256
Total Memory:	2 TB
Max. Peak Threads:	1

Software Summary

Compiler:	C/C++/Fortran: Version 21.9 of NVIDIA HPC SDK for Linux
MPI Library:	OpenMPI Version 4.0.5
Other MPI Info:	None
Other Software:	None
Base Parallel Model:	ACC
Base Ranks Run:	8
Base Threads Run:	1
Peak Parallel Models:	ACC
Minimum Peak Ranks:	8
Maximum Peak Ranks:	8
Max. Peak Threads:	1
Min. Peak Threads:	1

Node Description: DGX A100

Hardware

Number of nodes:	1
Uses of the node:	compute
Vendor:	NVIDIA Corporation
Model:	DGX A100
CPU Name:	AMD EPYC 7742
CPU(s) orderable:	2 chips
Chips enabled:	2
Cores enabled:	128
Cores per chip:	64
Threads per core:	2
CPU Characteristics:	Turbo Boost up to 3400MHz
CPU MHz:	2250
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	512 KB I+D on chip per core
L3 Cache:	256 MB I+D on chip per chip 16 MB shared / 4 cores
Other Cache:	None
Memory:	2 TB (32 x 64 GB 2Rx8 PC4-3200AA-R)
Disk Subsystem:	OS: 2TB U.2 NVMe SSD drive Internal Storage: 30TB (8x 3.84TB U.2 NVMe SSD drives)
Other Hardware:	None
Accel Count:	8
Accel Model:	Tesla A100-SXM-80GB
Accel Vendor:	NVIDIA Corporation
Accel Type:	GPU
Accel Connection:	NVLINK 3.0, NVSWITCH 2.0 600GB/s
Accel ECC enabled:	Yes
Accel Description:	See Notes
Adapter:	None
Number of Adapters:	0
Slot Type:	None

(Continued on next page)

Software

Accelerator Driver:	NVIDIA UNIX x86_64 Kernel Module 470.57.02
Adapter:	None
Adapter Driver:	None
Adapter Firmware:	None
Operating System:	Ubuntu 20.04 4.12.14-94.41-default
Local File System:	xfs
Shared File System:	None
System State:	Run level 3 (multi-user)
Other Software:	None

Node Description: DGX A100

Hardware (Continued)

Data Rate:	None
Ports Used:	0
Interconnect Type:	None

Interconnect Description: None

Hardware

Vendor:	N/A
Model:	N/A
Switch Model:	N/A
Number of Switches:	0
Number of Ports:	0
Data Rate:	0
Firmware:	0
Topology:	N/A
Primary Use:	N/A

Software

: --

Compiler Invocation Notes

```
Binaries built and run within a NVHPC SDK 21.9 CUDA 11.4 Ubuntu 20.04
 Container available from NVIDIA's NGC Catalog:
 https://ngc.nvidia.com/catalog/containers/nvidia:nvhpc
```

Submit Notes

```
The config file option 'submit' was used.
 MPI startup command:
   mpirun command was used to start MPI jobs.

 Indiviual Ranks were bound to the CPU cores on the same NUMA node as
  the GPU using 'numactl' within the following "bindACC.pl" perl script:
---- Start bindACC.pl ------
my %core_map = (
  0=>48, 1=>56, 2=>16, 3=>24, 4=>112, 5=>120, 6=>80, 7=>88
);
my %mem_map = (
  0=>3, 1=>3, 2=>1, 3=>1, 4=>7, 5=>7, 6=>5, 7=>5,
);
my $rank = $ENV{OMPI_COMM_WORLD_LOCAL_RANK};
my $mrank = $rank % 8;
my $cplus = int($rank/8);
```

(Continued on next page)

Submit Notes (Continued)

```
my $core = $core_map{$mrank} + $cplus;
my $mem = $mem_map{$mrank};
my $cmd = "numactl -C $core -m $mem ";
while (my $arg = shift) {
       $cmd .= "$arg ";
}
system($cmd);
---- End bindACC.pl ------
```

Platform Notes

```
 Detailed A100 Information from nvaccelinfo
 CUDA Driver Version:           11040
 NVRM version:                  NVIDIA UNIX x86_64 Kernel Module 470.57.02
 Device Number:                 0
 Device Name:                   NVIDIA A100-SXM-80GB
 Device Revision Number:        8.0
 Global Memory Size:            85198045184
 Number of Multiprocessors:     108
 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:                    1410 MHz
 Execution Timeout:             No
 Integrated Device:             No
 Can Map Host Memory:           Yes
 Compute Mode:                  default
 Concurrent Kernels:            Yes
 ECC Enabled:                   Yes
 Memory Clock Rate:             1593 MHz
 Memory Bus Width:              5120 bits
 L2 Cache Size:                 41943040 bytes
 Max Threads Per SMP:           2048
 Async Engines:                 3
 Unified Addressing:            Yes
 Managed Memory:                Yes
 Concurrent Managed Memory:     Yes
 Preemption Supported:          Yes
```

(Continued on next page)

Platform Notes (Continued)

```
 Cooperative Launch:            Yes
   Multi-Device:                Yes
 Default Target:                cc80
```

Compiler Version Notes

```
==============================================================================
 CC  605.lbm_s(base, peak) 613.soma_s(base, peak) 618.tealeaf_s(base, peak)
     621.miniswp_s(base, peak) 634.hpgmgfv_s(base, peak)
------------------------------------------------------------------------------
nvc 21.9-0 64-bit target on x86-64 Linux -tp zen
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
 CXXC 632.sph_exa_s(base, peak)
------------------------------------------------------------------------------
nvc++ 21.9-0 64-bit target on x86-64 Linux -tp zen
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
 FC  619.clvleaf_s(base, peak) 628.pot3d_s(base, peak) 635.weather_s(base,
     peak)
------------------------------------------------------------------------------
nvfortran 21.9-0 64-bit target on x86-64 Linux -tp zen
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------
```

Base Compiler Invocation

C benchmarks:
`mpicc`

C++ benchmarks:
`mpicxx`

Fortran benchmarks:
`mpif90`

Copyright 2021-2023 Standard Performance Evaluation Corporation

Base Portability Flags

621.miniswp_s: `-DUSE_KBA -DUSE_ACCELDIR`
632.sph_exa_s: `-DSPEC_USE_LT_IN_KERNELS --c++17`

Base Optimization Flags

C benchmarks:
`-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu`

C++ benchmarks:
`-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu`

Fortran benchmarks:
`-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu`

Base Other Flags

C benchmarks:
`-w`

C++ benchmarks:
`-w`

Fortran benchmarks:
`-w`

Peak Compiler Invocation

C benchmarks:
`mpicc`

C++ benchmarks:
`mpicxx`

Fortran benchmarks:
`mpif90`

Peak Portability Flags

621.miniswp_s: `-DUSE_KBA -DUSE_ACCELDIR`
632.sph_exa_s: `-DSPEC_USE_LT_IN_KERNELS --c++17`

Peak Optimization Flags

C benchmarks:

605.lbm_s: `basepeak = yes`

613.soma_s: `-fast -O3 -acc=gpu -gpu=pinned`

618.tealeaf_s: `-fast -Msafeptr -acc=gpu`

621.miniswp_s: `-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu -gpu=pinned`

634.hpgmgfv_s: `-fast -acc=gpu -gpu=pinned -static-nvidia`

C++ benchmarks:

632.sph_exa_s: `basepeak = yes`

Fortran benchmarks:

619.clvleaf_s: `basepeak = yes`

628.pot3d_s: `-Mstack_arrays -fast -acc=gpu`

635.weather_s: `basepeak = yes`

Peak Other Flags

C benchmarks:
`-w`

C++ benchmarks:
`-w`

Fortran benchmarks:
`-w`

The flags file that was used to format this result can be browsed at
http://www.spec.org/hpc2021/flags/nv2021_flags_v1.0.3.html

You can also download the XML flags source by saving the following link:
http://www.spec.org/hpc2021/flags/nv2021_flags_v1.0.3.xml

SPEChpc 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 info@spec.org.

Tested with SPEChpc2021 v1.0.2 on 2021-09-13 22:48:33-0400.
Report generated on 2023-08-25 18:58:36 by hpc2021 PS formatter v1.0.3.
Originally published on 2021-10-20.