SPECaccel®2023 Result

Supermciro
(Test Sponsor: NVIDIA Corporation)
Tesla A100 PCIe 80GB
120GQ-TNRT

SPECaccel 2023_base = 2.08
SPECaccel 2023_peak = 2.08

accel2023 License: 9045
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Oct-2023
Hardware Availability: Mar-2023
Software Availability: Nov-2023

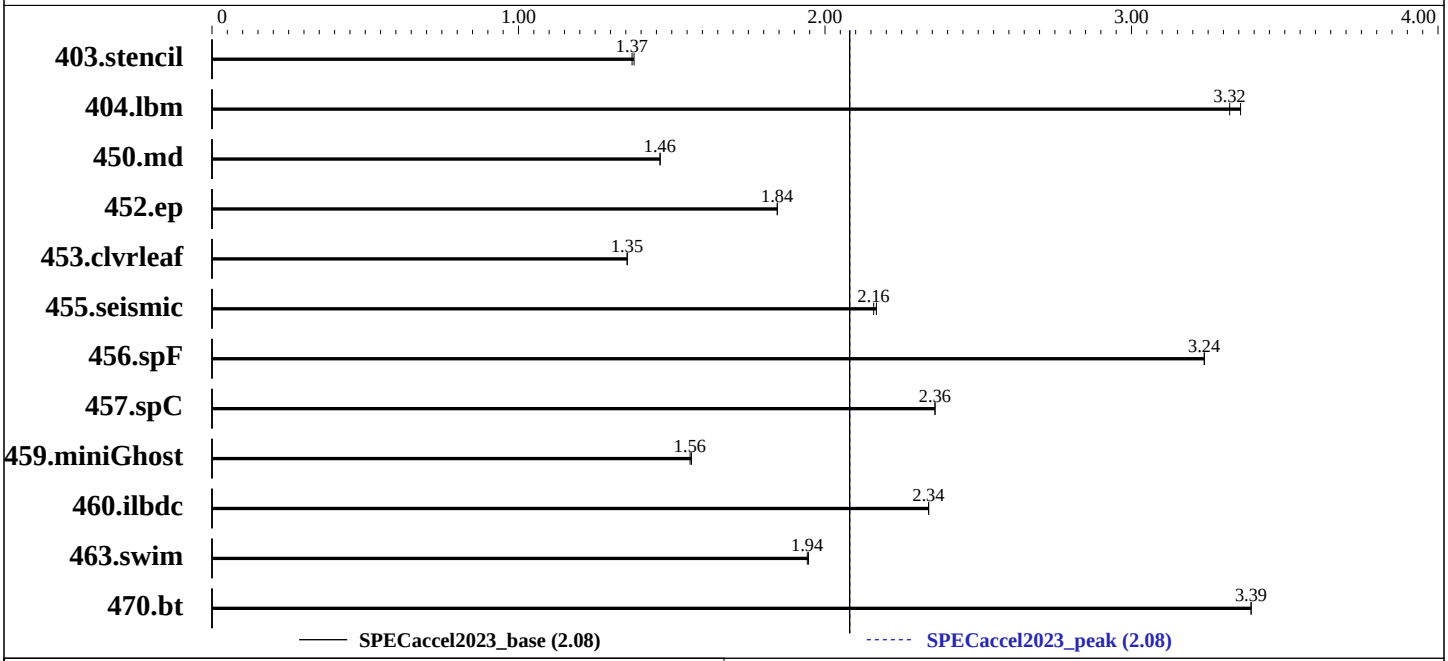

Hardware

CPU Name:	Intel Xeon Gold 6338
Max MHz.:	3400
Nominal:	2000
Enabled:	64 cores, 2 chips, 2 threads/core
Orderable:	2 chips
Cache L1:	32 KB I + 48 KB D on chip per core
L2:	1280 KB I+D on chip per core
L3:	48 MB I+D on chip per chip
Other:	None
Memory:	512 GB (16x 16GB, PC3200 CL3 DDR4)
Storage:	1TB SATA
Other:	None
Base Threads Run:	1
Min. Peak Threads:	1
Max. Peak Threads:	1

Accelerator

Accel Model Name:	A100 PCIe 80GB
Accel Vendor:	NVIDIA
Accel Name:	Tesla A100 PCIe 80GB
Type of Accel:	GPU
Accel Connection:	PCIe 4.0 16x
Does Accel Use ECC:	Yes
Accel Description:	See Notes
Accel Driver:	NVIDIA UNIX x86_64 Kernel Module 525.60.13

Software

OS:	Rocky Linux release 8.8 (Green Obsidian) 4.18.0-477.15.1.el8_8.x86_64
Compiler:	C/Fortran: Version 23.11 of NVHPC SDK
Firmware:	1.1b 11/01/2021
File System:	xfs
System State:	Run level 3 (multi-user)
Other:	None
Base Parallel Model:	TGT
Base Threads Run:	1
Peak Parallel Models:	TGT

(Continued on next page)

Software (Continued)

Max. Peak Threads: 1
Min. Peak Threads: 1

Results Table

Benchmark	Base							Peak						
	Model	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Model	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
403.stencil	TGT	320	1.38	**321**	**1.37**			TGT	320	1.38	**321**	**1.37**		
404.lbm	TGT	**137**	**3.32**	136	3.36			TGT	**137**	**3.32**	136	3.36		
450.md	TGT	410	1.46	**411**	**1.46**			TGT	410	1.46	**411**	**1.46**		
452.ep	TGT	225	1.84	**225**	**1.84**			TGT	225	1.84	**225**	**1.84**		
453.clvrleaf	TGT	738	1.35	**738**	**1.35**			TGT	738	1.35	**738**	**1.35**		
455.seismic	TGT	**361**	**2.16**	360	2.17			TGT	**361**	**2.16**	360	2.17		
456.spF	TGT	**147**	**3.24**	147	3.24			TGT	**147**	**3.24**	147	3.24		
457.spC	TGT	**229**	**2.36**	229	2.36			TGT	**229**	**2.36**	229	2.36		
459.miniGhost	TGT	377	1.56	**378**	**1.56**			TGT	377	1.56	**378**	**1.56**		
460.ilbdc	TGT	**237**	**2.34**	237	2.34			TGT	**237**	**2.34**	237	2.34		
463.swim	TGT	**227**	**1.94**	226	1.95			TGT	**227**	**1.94**	226	1.95		
470.bt	TGT	**311**	**3.39**	311	3.39			TGT	**311**	**3.39**	311	3.39		

SPEC accel2023_base = 2.08

SPEC accel2023_peak = 2.08

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

Operating System Notes

```
Shell stacksize set to unlimited via "limit stacksize unlimited"
```

Platform Notes

```
Information from nvaccelinfo
CUDA Driver Version:           12000
NVRM version:                  NVIDIA UNIX x86_64 Kernel Module  525.60.13  Wed Nov 30 06:39:21 UTC 2022
Device Number:                 0
Device Name:                   NVIDIA A100 80GB PCIe
Device Revision Number:        8.0
Global Memory Size:            85024112640
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
```

(Continued on next page)

Platform Notes (Continued)

```
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:             1512 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
Cooperative Launch:            Yes
Default Target:                cc80

Sysinfo program /local/home/mcolgrove/ACCELV2/bin/sysinfo
Rev: r6622 of 2021-04-07 b1a7d5f8f71be5aff70a755cad7211a0
running on ice2 Wed Oct 25 10:37:27 2023

SUT (System Under Test) info as seen by some common utilities.
For more information on this section, see
   https://www.spec.org/cpu2017/Docs/config.html#sysinfo

From /proc/cpuinfo
   model name : Intel(R) Xeon(R) Gold 6338 CPU @ 2.00GHz
      2  "physical id"s (chips)
      128 "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 : 32
      siblings  : 64
      physical 0: cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
      25 26 27 28 29 30 31
      physical 1: cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
      25 26 27 28 29 30 31

From lscpu from util-linux 2.34:
     Architecture:                    x86_64
     CPU op-mode(s):                  32-bit, 64-bit
     Byte Order:                      Little Endian
     Address sizes:                   46 bits physical, 57 bits virtual
```

(Continued on next page)

Platform Notes (Continued)

```
CPU(s):                          128
On-line CPU(s) list:             0-127
Thread(s) per core:              2
Core(s) per socket:              32
Socket(s):                       2
NUMA node(s):                    2
Vendor ID:                       GenuineIntel
CPU family:                      6
Model:                           106
Model name:                      Intel(R) Xeon(R) Gold 6338 CPU @ 2.00GHz
Stepping:                        6
Frequency boost:                 enabled
CPU MHz:                         2600.000
CPU max MHz:                     3200.0000
CPU min MHz:                     800.0000
BogoMIPS:                        4000.00
Virtualization:                  VT-x
L1d cache:                       3 MiB
L1i cache:                       2 MiB
L2 cache:                        80 MiB
L3 cache:                        96 MiB
NUMA node0 CPU(s):               0-31,64-95
NUMA node1 CPU(s):               32-63,96-127
Vulnerability Itlb multihit:     Not affected
Vulnerability L1tf:              Not affected
Vulnerability Mds:               Not affected
Vulnerability Meltdown:          Not affected
Vulnerability Mmio stale data:   Mitigation; Clear CPU buffers; SMT vulnerable
Vulnerability Retbleed:          Not affected
Vulnerability Spec store bypass: Mitigation; Speculative Store Bypass disabled via
prctl and seccomp
Vulnerability Spectre v1:        Mitigation; usercopy/swapgs barriers and __user
pointer sanitization
Vulnerability Spectre v2:        Mitigation; Enhanced IBRS, IBPB conditional, RSB
filling, PBRSB-eIBRS SW sequence
Vulnerability Srbds:             Not affected
Vulnerability Tsx async abort:   Not affected
Flags:                           fpu vme de pse tsc msr pae mce cx8 apic sep mtrr
pge mca cmov pat pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx
pdpe1gb rdtscp lm constant_tsc art arch_perfmon pebs bts rep_good nopl xtopology
nonstop_tsc cpuid aperfmperf pni pclmulqdq dtes64 monitor ds_cpl vmx smx est tm2
ssse3 sdbg fma cx16 xtpr pdcm pcid dca sse4_1 sse4_2 x2apic movbe popcnt
tsc_deadline_timer aes xsave avx f16c rdrand lahf_lm abm 3dnowprefetch cpuid_fault
epb cat_l3 invpcid_single ssbd mba ibrs ibpb stibp ibrs_enhanced tpr_shadow vnmi
flexpriority ept vpid ept_ad fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms invpcid
cqm rdt_a avx512f avx512dq rdseed adx smap avx512ifma clflushopt clwb intel_pt
avx512cd sha_ni avx512bw avx512vl xsaveopt xsavec xgetbv1 xsaves cqm_llc
```

(Continued on next page)

Platform Notes (Continued)

```
      cqm_occup_llc cqm_mbm_total cqm_mbm_local wbnoinvd dtherm ida arat pln pts
      avx512vbmi umip pku ospke avx512_vbmi2 gfni vaes vpclmulqdq avx512_vnni
      avx512_bitalg tme avx512_vpopcntdq rdpid md_clear pconfig flush_l1d
      arch_capabilities

  From lscpu --cache:
      NAME ONE-SIZE ALL-SIZE WAYS TYPE        LEVEL
      L1d       48K       3M   12 Data            1
      L1i       32K       2M    8 Instruction     1
      L2       1.3M      80M   20 Unified         2
      L3        48M      96M   12 Unified         3

  /proc/cpuinfo cache data
     cache size : 49152 KB

  From numactl --hardware
  WARNING: a numactl 'node' might or might not correspond to a physical chip.
    available: 2 nodes (0-1)
    node 0 cpus: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27
    28 29 30 31 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88
    89 90 91 92 93 94 95
    node 0 size: 257622 MB
    node 0 free: 206061 MB
    node 1 cpus: 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56
    57 58 59 60 61 62 63 96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112
    113 114 115 116 117 118 119 120 121 122 123 124 125 126 127
    node 1 size: 257990 MB
    node 1 free: 244551 MB
    node distances:
    node   0   1
      0:  10  20
      1:  20  10

  From /proc/meminfo
     MemTotal:       527987596 kB
     HugePages_Total:       0
     Hugepagesize:       2048 kB

  /sys/devices/system/cpu/cpu*/cpufreq/scaling_governor has
     performance

  /usr/bin/lsb_release -d
     Ubuntu 20.04.6 LTS

  From /etc/*release* /etc/*version*
     debian_version: bullseye/sid
     os-release:
```

(Continued on next page)

Platform Notes (Continued)

```
      NAME="Ubuntu"
      VERSION="20.04.6 LTS (Focal Fossa)"
      ID=ubuntu
      ID_LIKE=debian
      PRETTY_NAME="Ubuntu 20.04.6 LTS"
      VERSION_ID="20.04"
      HOME_URL="https://www.ubuntu.com/"
      SUPPORT_URL="https://help.ubuntu.com/"

uname -a:
   Linux ice2 5.4.0-153-generic #170-Ubuntu SMP Fri Jun 16 13:43:31 UTC 2023 x86_64
   x86_64 x86_64 GNU/Linux

Kernel self-reported vulnerability status:

CVE-2018-12207 (iTLB Multihit):                          Not affected
CVE-2018-3620 (L1 Terminal Fault):                       Not affected
Microarchitectural Data Sampling:                        Not affected
CVE-2017-5754 (Meltdown):                                Not affected
mmio_stale_data:                                        Mitigation: Clear CPU buffers; SMT
                                                          vulnerable
retbleed:                                                Not affected
CVE-2018-3639 (Speculative Store Bypass):                Mitigation: Speculative Store
                                                         Bypass disabled via prctl and
                                                         seccomp
CVE-2017-5753 (Spectre variant 1):                       Mitigation: usercopy/swapgs
                                                         barriers and __user pointer
                                                         sanitization
CVE-2017-5715 (Spectre variant 2):                      Mitigation: Enhanced IBRS, IBPB:
                                                         conditional, RSB filling,
                                                         PBRSB-eIBRS: SW sequence
CVE-2020-0543 (Special Register Buffer Data Sampling): Not affected
CVE-2019-11135 (TSX Asynchronous Abort):                 Not affected

run-level 3 Aug 10 17:47

SPEC is set to: /local/home/mcolgrove/ACCELV2
   Filesystem     Type  Size  Used Avail Use% Mounted on
   /dev/nvme1n1p1 ext4  916G   25G  845G   3% /local

From /sys/devices/virtual/dmi/id
    Vendor:         Supermicro
    Product:        SYS-120GQ-TNRT
    Product Family: SMC X12

Cannot run dmidecode; consider saying (as root)
   chmod +s /usr/sbin/dmidecode
```

(Continued on next page)

Platform Notes (Continued)

```
 BIOS:
    BIOS Vendor:       American Megatrends International, LLC.
    BIOS Version:      1.1b
    BIOS Date:         11/01/2021

 (End of data from sysinfo program)
```

Compiler Version Notes

```
==============================================================================
C          | 457.spC(base)
------------------------------------------------------------------------------
/usr/bin/ld: /usr/lib/x86_64-linux-gnu/crt1.o: in function `_start':
(.text+0x24): undefined reference to `main'
pgacclnk: child process exit status 1: /usr/bin/ld
nvc 23.11-0 64-bit target on x86-64 Linux -tp icelake-server
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C          | 403.stencil(base) 404.lbm(base) 452.ep(base) 470.bt(base)
------------------------------------------------------------------------------
nvc 23.11-0 64-bit target on x86-64 Linux -tp icelake-server
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C          | 457.spC(base)
------------------------------------------------------------------------------
/usr/bin/ld: /usr/lib/x86_64-linux-gnu/crt1.o: in function `_start':
(.text+0x24): undefined reference to `main'
pgacclnk: child process exit status 1: /usr/bin/ld
nvc 23.11-0 64-bit target on x86-64 Linux -tp icelake-server
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C          | 403.stencil(base) 404.lbm(base) 452.ep(base) 470.bt(base)
------------------------------------------------------------------------------
nvc 23.11-0 64-bit target on x86-64 Linux -tp icelake-server
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
```

(Continued on next page)

Compiler Version Notes (Continued)

```
------------------------------------------------------------------------------

==============================================================================
Fortran    | 450.md(base) 455.seismic(base) 456.spF(base) 460.ilbdc(base)
           | 463.swim(base)
------------------------------------------------------------------------------
nvfortran 23.11-0 64-bit target on x86-64 Linux -tp icelake-server
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
Fortran, C | 453.clvrleaf(base) 459.miniGhost(base)
------------------------------------------------------------------------------
nvfortran 23.11-0 64-bit target on x86-64 Linux -tp icelake-server
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
nvc 23.11-0 64-bit target on x86-64 Linux -tp icelake-server
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------
```

Base Compiler Invocation

C benchmarks:
nvc

Fortran benchmarks:
nvfortran

Benchmarks using both Fortran and C:
nvfortran nvc

Base Portability Flags

403.stencil: -DSPEC_NO_NOTHING
457.spC: -mcmodel=medium -Wl,--no-relax

Base Optimization Flags

C benchmarks:
`-Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia`

Fortran benchmarks:
`-Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia`

Benchmarks using both Fortran and C:

453.clvrleaf: `-Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia`

459.miniGhost: `-Mnomain -Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia`

Peak Optimization Flags

C benchmarks:

403.stencil: `basepeak = yes`

404.lbm: `basepeak = yes`

452.ep: `basepeak = yes`

457.spC: `basepeak = yes`

470.bt: `basepeak = yes`

Fortran benchmarks:

450.md: `basepeak = yes`

455.seismic: `basepeak = yes`

456.spF: `basepeak = yes`

460.ilbdc: `basepeak = yes`

463.swim: `basepeak = yes`

Benchmarks using both Fortran and C:

453.clvrleaf: `basepeak = yes`

(Continued on next page)

Peak Optimization Flags (Continued)

459.miniGhost: `basepeak = yes`

The flags file that was used to format this result can be browsed at
`http://www.spec.org/accel2023/flags/nv2023_flags_v2.html`

You can also download the XML flags source by saving the following link:
`http://www.spec.org/accel2023/flags/nv2023_flags_v2.xml`

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

Tested with SPECaccel2023 v2.0.17 on 2023-10-25 13:37:27-0400.
Report generated on 2023-12-06 13:07:34 by accel2023 PDF formatter v112.
Originally published on 2023-11-08.