



SPEC ACCEL™ OCL Result

Copyright 2014-2017 Standard Performance Evaluation Corporation

Cray

(Test Sponsor: Indiana University)

NVIDIA Tesla K20

Cray XK7

SPECaccel_ocl_peak = Not Run

SPECaccel_ocl_base = 1.70

ACCEL license: 3440A


Test sponsor: Indiana University

Tested by: Indiana University

Test date: Mar-2017

Hardware Availability: Apr-2013

Software Availability: Jan-2017





SPEC ACCEL OCL Result

Copyright 2014-2017 Standard Performance Evaluation Corporation

Cray

(Test Sponsor: Indiana University)

NVIDIA Tesla K20

Cray XK7

SPECaccel_ocl_peak = Not Run

SPECaccel_ocl_base = 1.70

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

Test date: Mar-2017
Hardware Availability: Apr-2013
Software Availability: Jan-2017

Hardware

CPU Name: AMD Opteron 6276
CPU Characteristics: AMD Turbo CORE Technology up to 3.2GHz, Turbo CORE off
CPU MHz: 2300
CPU MHz Maximum: 3200
FPU: Integrated
CPU(s) enabled: 16 cores, 1 chip, 16 cores/chip
CPU(s) orderable: 1 chip
Primary Cache: 32 KB I + 16 KB D on chip per core
Secondary Cache: 16 MB I+D on chip per chip, 2 MB shared / 2 cores
L3 Cache: 16 MB I+D on chip per chip, 8 MB shared / 8 cores
Other Cache: None
Memory: 32 GB (4 x 8 GB 2Rx4 PC3L-12800R-11, ECC)
Disk Subsystem: None
Other Hardware: None

Accelerator

Accel Model Name: Tesla K20
Accel Vendor: NVIDIA
Accel Name: NVIDIA Tesla K20
Type of Accel: GPU
Accel Connection: PCIe 2.0 16x
Does Accel Use ECC: yes
Accel Description: NVIDIA Tesla K20m GPU, 2496 CUDA cores, 706MHz, 5 GB GDDR5 RAM
Accel Driver: NVIDIA UNIX x86_64 Kernel Module 352.68

Software

Operating System: SUSE Linux Enterprise Server 11 (x86_64), Cray Linux Environment 5.2
3.0.101-0.46.1_1.0502.8871-cray_gem_c
Compiler: GNU Fortran (GCC) 5.3.0
File System: NFSv3 (DDN SFA12KE) over 10GB Ethernet
System State: Run level 3 (Multi-user)
Other Software: NVIDIA CUDA 7.5.18



SPEC ACCEL OCL Result

Copyright 2014-2017 Standard Performance Evaluation Corporation

Cray

(Test Sponsor: Indiana University)

NVIDIA Tesla K20

Cray XK7

SPECaccel_ocl_peak = Not Run

SPECaccel_ocl_base = 1.70

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

Test date: Mar-2017
Hardware Availability: Apr-2013
Software Availability: Jan-2017

Results Table

Benchmark	Base						Peak					
	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
101.tpacf	77.0	1.39	<u>77.0</u>	<u>1.39</u>	77.0	1.39						
103.stencil	65.3	1.92	65.3	1.91	<u>65.3</u>	<u>1.91</u>						
104.lbm	47.0	2.38	<u>47.2</u>	<u>2.37</u>	47.3	2.37						
110.fft	63.2	1.76	<u>63.2</u>	<u>1.76</u>	63.2	1.76						
112.spmv	89.8	1.64	89.6	1.64	<u>89.7</u>	<u>1.64</u>						
114.mriq	22.6	4.83	<u>22.5</u>	<u>4.84</u>	22.5	4.84						
116.histo	<u>114</u>	<u>1.00</u>	118	0.962	113	1.01						
117.bfs	69.7	1.68	69.6	1.68	<u>69.7</u>	<u>1.68</u>						
118.cutcp	43.5	2.28	43.5	2.28	<u>43.5</u>	<u>2.28</u>						
120.kmeans	<u>122</u>	<u>0.819</u>	122	0.820	123	0.816						
121.lavamd	21.1	5.17	<u>21.1</u>	<u>5.16</u>	21.4	5.09						
122.cfd	79.5	1.59	79.7	1.58	<u>79.5</u>	<u>1.58</u>						
123.nw	<u>81.6</u>	<u>1.41</u>	81.5	1.41	81.6	1.41						
124.hotspot	46.7	2.44	<u>46.8</u>	<u>2.44</u>	46.9	2.43						
125.lud	<u>111</u>	<u>1.07</u>	111	1.07	111	1.07						
126.ge	51.4	3.01	<u>51.4</u>	<u>3.01</u>	51.6	3.01						
127.srad	76.3	1.49	<u>76.3</u>	<u>1.49</u>	76.4	1.49						
128.heartwall	<u>157</u>	<u>0.676</u>	157	0.676	157	0.675						
140.bplustree	113	0.955	<u>113</u>	<u>0.954</u>	113	0.953						

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

Platform Notes

Sysinfo program

```
/N/dc2/projects/hpc/lijunj/SPEC/accel-1.1-run/bigred2/Docs/sysinfo
$Rev: 6965 $ $Date:: 2015-04-21 #$ c05a7f14b1b1765e3fe1df68447e8a35
running on nid00221 Fri Mar 10 13:54:43 2017
```

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>

From /proc/cpuinfo

```
model name : AMD Opteron(TM) Processor 6276
1 "physical id"s (chips)
16 "processors"
```

cores, siblings (Caution: counting these is hw and system dependent. The following excerpts from /proc/cpuinfo might not be reliable. Use with

Continued on next page



SPEC ACCEL OCL Result

Copyright 2014-2017 Standard Performance Evaluation Corporation

Cray

(Test Sponsor: Indiana University)

NVIDIA Tesla K20

Cray XK7

SPECaccel_ocl_peak = Not Run

SPECaccel_ocl_base = 1.70

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

Test date: Mar-2017
Hardware Availability: Apr-2013
Software Availability: Jan-2017

Platform Notes (Continued)

```
caution.)
  cpu cores : 8
  siblings  : 16
  physical 0: cores 0 1 2 3 4 5 6 7
  cache size : 2048 KB
```

```
From /proc/meminfo
MemTotal:      33083764 kB
HugePages_Total: 0
Hugepagesize:  2048 kB
```

```
/usr/bin/lsb_release -d
SUSE Linux Enterprise Server 11 (x86_64)
```

```
From /etc/*release* /etc/*version*
SuSE-release:
SUSE Linux Enterprise Server 11 (x86_64)
VERSION = 11
PATCHLEVEL = 3
```

```
uname -a:
Linux nid00221 3.0.101-0.46.1_1.0502.8871-cray_gem_c #1 SMP Sat Oct 22
15:26:43 UTC 2016 x86_64 x86_64 x86_64 GNU/Linux
```

```
SPEC is set to: /N/dc2/projects/hpc/lijunj/SPEC/accel-1.1-run/bigred2
Filesystem      Type      Size  Used Avail Use% Mounted on
10.10.0.171@o2ib:/dc2 lustre  5.3P  5.0P  223T  96% /N/dc2
```

Cannot run dmidecode; consider saying 'chmod +s /usr/sbin/dmidecode'

(End of data from sysinfo program)

(End of data from sysinfo program)

Base Runtime Environment

C benchmarks:

OpenCL Platform: NVIDIA CUDA, OpenCL 1.2 CUDA 7.5.23
OpenCL Device #0: Tesla K20, v 352.68

C++ benchmarks:

OpenCL Platform: NVIDIA CUDA, OpenCL 1.2 CUDA 7.5.23
OpenCL Device #0: Tesla K20, v 352.68



SPEC ACCEL OCL Result

Copyright 2014-2017 Standard Performance Evaluation Corporation

Cray

(Test Sponsor: Indiana University)

NVIDIA Tesla K20

Cray XK7

SPECaccel_ocl_peak = Not Run

SPECaccel_ocl_base = 1.70

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

Test date: Mar-2017
Hardware Availability: Apr-2013
Software Availability: Jan-2017

Base Compiler Invocation

C benchmarks:
gcc

C++ benchmarks:
g++

Base Optimization Flags

C benchmarks:
-O2 -I/opt/nvidia/cudatoolkit7.5/7.5.18-1.0502.10743.2.1/include
-L/opt/nvidia/cudatoolkit7.5/7.5.18-1.0502.10743.2.1/lib64 -lOpenCL

C++ benchmarks:
-O2 -I/opt/nvidia/cudatoolkit7.5/7.5.18-1.0502.10743.2.1/include
-L/opt/nvidia/cudatoolkit7.5/7.5.18-1.0502.10743.2.1/lib64 -lOpenCL

Base Other Flags

C benchmarks:
116.histo: -DSPEC_LOCAL_MEMORY_HEADROOM=1

The flags file that was used to format this result can be browsed at
<http://www.spec.org/accel/flags/flags-advanced.20170426.html>

You can also download the XML flags source by saving the following link:
<http://www.spec.org/accel/flags/flags-advanced.20170426.xml>

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 v1.1.
Report generated on Wed Apr 26 11:41:23 2017 by SPEC ACCEL PS/PDF formatter v1290.
Originally published on 26 April 2017.