www.phoronix-test-suite.com

RTX 3070 Compute

AMD Ryzen 9 5900X 12-Core testing with a ASUS ROG CROSSHAIR VIII HERO (3402 BIOS) and NVIDIA GeForce RTX 3070 8GB on Ubuntu 20.04 via the Phoronix Test Suite.

Automated Executive Summary

1 had the most wins, coming in first place for 71% of the tests.

Based on the geometric mean of all complete results, the fastest (2) was 1.003x the speed of the slowest (3). 1 was 1x the speed of 2 and 3 was 0.997x the speed of 1.

The results with the greatest spread from best to worst included:

NCNN (Target: Vulkan GPU - Model: squeezenet_ssd) at 1.066x
NCNN (Target: Vulkan GPU - Model: mobilenet) at 1.055x
ViennaCL (Test: CPU BLAS - dAXPY) at 1.033x
NCNN (Target: Vulkan GPU - Model: blazeface) at 1.033x
ViennaCL (Test: CPU BLAS - dCOPY) at 1.027x
ViennaCL (Test: CPU BLAS - sCOPY) at 1.023x
ViennaCL (Test: CPU BLAS - dGEMM-NN) at 1.022x
ViennaCL (Test: CPU BLAS - dGEMM-TN) at 1.022x
ViennaCL (Test: CPU BLAS - dDOT) at 1.018x

ViennaCL (Test: OpenCL BLAS - dGEMM-TN) at 1.018x.

Test Systems:

1

2

3

Processor: AMD Ryzen 9 5900X 12-Core @ 3.70GHz (12 Cores / 24 Threads), Motherboard: ASUS ROG CROSSHAIR VIII HERO (3402 BIOS), Chipset: AMD Starship/Matisse, Memory: 16GB, Disk: 1000GB Sabrent Rocket 4.0 Plus + 2000GB, Graphics: NVIDIA GeForce RTX 3070 8GB, Audio: NVIDIA Device 228b, Monitor: ASUS VP28U, Network: Realtek RTL8125 2.5GbE + Intel I211

OS: Ubuntu 20.04, Kernel: 5.8.0-48-generic (x86_64), Desktop: GNOME Shell 3.36.7, Display Server: X Server 1.20.9, Display Driver: NVIDIA 460.67, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 11.2.162, Vulkan: 1.2.155, Compiler: GCC 9.3.0 + CUDA 11.2, File-System: ext4, Screen Resolution: 3840x2160

Kernel Notes: Transparent Huge Pages: madvise
Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --enable-checking=release --enable-clocale=gnu --enable-default-pie --enable-gnu-unique-object --enable-languages=c,ada,c++,go,brig,d,fortran,objc,obj-c++,gm2 --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none=/build/gcc-9-HskZEa/gcc-9-9.3.0/debian/tmp-nvptx/usr,hsa --enable-plugin --enable-shared --enable-threads=posix --host=x86_64-linux-gnu --program-prefix=x86_64-linux-gnu- --target=x86_64-linux-gnu --with-abi=m64 --with-arch-32=i686 --with-default-libstdcxx-abi=new --with-gcc-major-version-only --with-multilib-list=m32,m64,mx32 --with-target-system-zlib=auto --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: acpi-cpufreq performance (Boost: Enabled) - CPU Microcode: 0xa201009
OpenCL Notes: GPU Compute Cores: 5888
Security Notes: itlb_multihit: Not affected + l1tf: Not affected + mds: Not affected + meltdown: Not affected + spec_store_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre_v1: Mitigation of usercopy/swapgs barriers and __user pointer sanitization + spectre_v2: Mitigation of Full AMD retpoline IBPB: conditional IBRS_FW STIBP: always-on RSB filling + srbds: Not affected + tsx_async_abort: Not affected

	1	2	3
SHOC Scalable HeterOgeneous Computing - OpenCL - Max SP Flops (GFLOPS)	**23118**	**23179**	
Normalized	99.74%	100%	
Standard Deviation	0.2%	0.4%	
Blender - Barbershop - CUDA (sec)	**509.33**	**508.38**	
Normalized	99.81%	100%	
Standard Deviation	0.1%	0.1%	
Blender - Barbershop - NVIDIA OptiX (sec)	**465.14**	**464.01**	
Normalized	99.76%	100%	
Standard Deviation	0.7%	0.5%	
LeelaChessZero - OpenCL (Nodes/s)	**28914**	**29117**	
Normalized	99.3%	100%	
Standard Deviation	0.4%	1.4%	
VkFFT (Benchmark Score)	**32004**	**32323**	

Normalized	99.01%	100%
Standard Deviation	1.3%	2.3%
RedShift Demo (sec)	228	228
Standard Deviation	0.7%	0.8%
Blender - Pabellon Barcelona - CUDA (sec)	**190.29**	**190.98**
Normalized	100%	99.64%
Standard Deviation	0%	0%
OctaneBench - Total Score (Score)	**410.514043**	**411.111675**
Normalized	99.85%	100%
GROMACS - Water Benchmark (Ns/Day)	**8.138**	**8.083**
Normalized	100%	99.32%
Standard Deviation	0.5%	0.4%
FAHBench (Ns/Day)	**267.0901**	**265.8324**
Normalized	100%	99.53%
Standard Deviation	0%	0.1%
Blender - Pabellon Barcelona - NVIDIA OptiX (sec)	**76.77**	**77.16**
Normalized	100%	99.49%
Standard Deviation	0.1%	0.1%
Blender - Classroom - CUDA (sec)	**76.10**	**76.49**
Normalized	100%	99.49%
Standard Deviation	0%	0%
Chaos Group V-RAY - NVIDIA CUDA GPU (vpaths)	**1342**	**1337**
Normalized	100%	99.63%
Standard Deviation	0.2%	0.1%
Chaos Group V-RAY - NVIDIA RTX GPU (vrays)	**1713**	**1710**
Normalized	100%	99.82%
Standard Deviation	0.2%	
LuxCoreRender OpenCL - Food (M samples/sec)	3.39	3.39
Standard Deviation	1.4%	1.2%
LuxCoreRender OpenCL - LuxCore Benchmark (M samples/sec)	**6.51**	**6.50**
Normalized	100%	99.85%
Standard Deviation	0%	0.3%
IndigoBench - OpenCL GPU - Bedroom (M samples/s)	**12.917**	**12.882**
Normalized	100%	99.73%
Standard Deviation	0.2%	0.2%
LuxCoreRender OpenCL - DLSC (M samples/sec)	**7.97**	**7.94**
Normalized	100%	99.62%
Standard Deviation	0.1%	0.1%
IndigoBench - OpenCL GPU - Supercar (M samples/s)	**37.257**	**37.163**
Normalized	100%	99.75%
Standard Deviation	0.1%	0.1%
Blender - Fishy Cat - CUDA (sec)	**54.20**	**54.44**
Normalized	100%	99.56%
Standard Deviation	0%	0.1%
NCNN - Vulkan GPU - regnety_400m (ms)	**16.75**	**16.87**
Normalized	100%	99.29%

Standard Deviation	2.6%	3.3%	
NCNN - Vulkan GPU - squeezenet_ssd (ms)	**14.49**	**15.45**	
Normalized	100%	93.79%	
Standard Deviation	1.8%	1.3%	
NCNN - Vulkan GPU - yolov4-tiny (ms)	**22.20**	**23.73**	
Normalized	100%	93.55%	
Standard Deviation	0.1%	8.7%	
NCNN - Vulkan GPU - resnet50 (ms)	**24.47**	**24.75**	
Normalized	100%	98.87%	
Standard Deviation	2.4%	2.6%	
NCNN - Vulkan GPU - alexnet (ms)	**11.08**	**11.22**	
Normalized	100%	98.75%	
Standard Deviation	0.5%	1%	
NCNN - Vulkan GPU - resnet18 (ms)	**13.95**	**14.03**	
Normalized	100%	99.43%	
Standard Deviation	0.1%	0.9%	
NCNN - Vulkan GPU - vgg16 (ms)	**55.71**	**55.89**	
Normalized	100%	99.68%	
Standard Deviation	0.6%	0.3%	
NCNN - Vulkan GPU - googlenet (ms)	**13.27**	**13.17**	
Normalized	99.25%	100%	
Standard Deviation	1.1%	3.6%	
NCNN - Vulkan GPU - blazeface (ms)	**1.84**	**1.90**	
Normalized	100%	96.84%	
Standard Deviation	0.6%	5.8%	
NCNN - Vulkan GPU - efficientnet-b0 (ms)	**5.58**	**5.62**	
Normalized	100%	99.29%	
Standard Deviation	3.1%	3.8%	
NCNN - Vulkan GPU - mnasnet (ms)	**4.05**	**4.10**	
Normalized	100%	98.78%	
Standard Deviation	4.8%	0.2%	
NCNN - Vulkan GPU - shufflenet-v2 (ms)	**4.85**	**4.92**	
Normalized	100%	98.58%	
Standard Deviation	2.2%	0.7%	
NCNN - Vulkan GPU-v3-v3 - mobilenet-v3	**4.22**	**4.24**	
Normalized	100%	99.53%	
Standard Deviation	2.6%	6.1%	
NCNN - Vulkan GPU-v2-v2 - mobilenet-v2	**4.39**	**4.35**	
Normalized	99.09%	100%	
Standard Deviation	1.6%	1.8%	
NCNN - Vulkan GPU - mobilenet (ms)	**12.85**	**13.56**	
Normalized	100%	94.76%	
Standard Deviation	1.1%	0.2%	
RealSR-NCNN - 4x - Yes (sec)	50.058	**49.994**	**50.181**
Normalized	99.87%	100%	99.63%
Standard Deviation	0.3%	0.3%	0.3%
Blender - Classroom - NVIDIA OptiX (sec)	**48.17**	**48.26**	
Normalized	100%	99.81%	
Standard Deviation	0.1%	0.1%	
Blender - Fishy Cat - NVIDIA OptiX (sec)	**35.70**	**35.84**	
Normalized	100%	99.61%	
Standard Deviation	0.1%	0.1%	
VkResample - 2x - Double (ms)	**220.452**	**221.037**	
Normalized	100%	99.74%	
Standard Deviation	0.1%	0.1%	

RTX 3070 Compute

clpeak - D.P.D (GFLOPS)	**360.90**	**364.99**
Normalized	98.88%	100%
Standard Deviation	0%	0%
Blender - BMW27 - CUDA (sec)	**29.07**	**29.03**
Normalized	99.86%	100%
Standard Deviation	0.2%	0%
ViennaCL - CPU BLAS - dGEMM-TT (GFLOPs/s)	**55.3**	**54.5**
Normalized	100%	98.55%
Standard Deviation	0.5%	3.1%
ViennaCL - CPU BLAS - dGEMM-TN (GFLOPs/s)	**57.0**	**55.8**
Normalized	100%	97.89%
Standard Deviation	0.4%	3.8%
ViennaCL - CPU BLAS - dGEMM-NT (GFLOPs/s)	**53.4**	**52.7**
Normalized	100%	98.69%
Standard Deviation	0.5%	2.3%
ViennaCL - CPU BLAS - dGEMM-NN (GFLOPs/s)	**54.7**	**53.5**
Normalized	100%	97.81%
Standard Deviation	0.3%	3.2%
ViennaCL - CPU BLAS - dGEMV-T (GB/s)	**81.3**	**81.9**
Normalized	99.27%	100%
Standard Deviation	3.1%	1.2%
ViennaCL - CPU BLAS - dGEMV-N (GB/s)	**76.5**	**77.0**
Normalized	99.35%	100%
Standard Deviation	1%	1.1%
ViennaCL - CPU BLAS - dDOT (GB/s)	**44.3**	**43.5**
Normalized	100%	98.19%
Standard Deviation	1.6%	2.5%
ViennaCL - CPU BLAS - dAXPY (GB/s)	**34.5**	**33.4**
Normalized	100%	96.81%
Standard Deviation	1.1%	3.4%
ViennaCL - CPU BLAS - dCOPY (GB/s)	**23.2**	**22.6**
Normalized	100%	97.41%
Standard Deviation	0.7%	5.7%
ViennaCL - CPU BLAS - sDOT (GB/s)	**141**	**139**
Normalized	100%	98.58%
Standard Deviation	1.4%	4.4%
ViennaCL - CPU BLAS - sAXPY (GB/s)	**94.0**	**92.8**
Normalized	100%	98.72%
Standard Deviation	1.6%	4.2%
ViennaCL - CPU BLAS - sCOPY (GB/s)	**62.0**	**63.4**
Normalized	97.79%	100%
Standard Deviation	2.3%	1.6%
ViennaCL - OpenCL BLAS - dGEMM-TN (GFLOPs/s)	**342**	**336**
Normalized	100%	98.25%
ViennaCL - OpenCL BLAS - dGEMM-NT (GFLOPs/s)	**343**	**340**
Normalized	100%	99.13%
Standard Deviation	0.6%	

RTX 3070 Compute

ViennaCL - OpenCL BLAS - dGEMM-NN (GFLOPs/s)	**342**	**338**	
Normalized	100%	98.83%	
Standard Deviation	0.5%	0.3%	
ViennaCL - OpenCL BLAS - dGEMV-T (GB/s)	**334**	**332**	
Normalized	100%	99.4%	
ViennaCL - OpenCL BLAS - dGEMV-N (GB/s)	**222**	**220**	
Normalized	100%	99.1%	
ViennaCL - OpenCL BLAS - dDOT (GB/s)	**397**	**396**	
Normalized	100%	99.75%	
ViennaCL - OpenCL BLAS - dAXPY (GB/s)	**396**	**395**	
Normalized	100%	99.75%	
ViennaCL - OpenCL BLAS - dCOPY (GB/s)	**365**	**364**	
Normalized	100%	99.73%	
Standard Deviation		0.2%	
ViennaCL - OpenCL BLAS - sDOT (GB/s)	**325**	**324**	
Normalized	100%	99.69%	
ViennaCL - OpenCL BLAS - sAXPY (GB/s)	**358**	**357**	
Normalized	100%	99.72%	
ViennaCL - OpenCL BLAS - sCOPY (GB/s)	293	293	
Standard Deviation		0.3%	
Blender - BMW27 - NVIDIA OptiX (sec)	**16.11**	**16.17**	
Normalized	100%	99.63%	
Standard Deviation	0.4%	0.4%	
SHOC Scalable HeterOgeneous Computing - OpenCL - T.R.B (GB/s)	**2121**	**2131**	
Normalized	99.5%	100%	
Standard Deviation	0.6%	0.4%	
VkResample - 2x - Single (ms)	**17.489**	**17.532**	
Normalized	100%	99.75%	
Standard Deviation	0.1%	0.1%	
NAMD CUDA - ATPase Simulation - 327,506 Atoms (days/ns)	**0.13273**	**0.13388**	
Normalized	100%	99.14%	
Standard Deviation	0.7%	1.9%	
LuxCoreRender OpenCL - R.C.a.P (M samples/sec)	**19.18**	**19.19**	
Normalized	99.95%	100%	
Standard Deviation	0.4%	0.2%	
RealSR-NCNN - 4x - No (sec)	**8.438**	**8.427**	8.437
Normalized	99.87%	100%	99.88%
Standard Deviation	1.7%	1.9%	2%
Betsy GPU Compressor - ETC2 RGB - Highest (sec)	**6.039**	**6.051**	
Normalized	100%	99.8%	
Standard Deviation	0.7%	0.4%	
Rodinia - O.P.F (sec)	**5.958**	**6.016**	
Normalized	100%	99.04%	
Standard Deviation	0.4%	0.7%	
Hashcat - SHA-512 (H/s)	**1664800000**	**1669566667**	
Normalized	99.71%	100%	
Standard Deviation	0.1%	0.1%	
Hashcat - SHA1 (H/s)	**13120133333**	**13144266667**	
Normalized	99.82%	100%	

RTX 3070 Compute

Standard Deviation	0.3%	0.1%	
Hashcat - MD5 (H/s)	**38778833333**	**38839033333**	
Normalized	99.85%	100%	
Standard Deviation	0.2%	0.2%	
cl-mem - Read (GB/s)	**393.6**	**393.2**	
Normalized	100%	99.9%	
Standard Deviation	0%	0.2%	
cl-mem - Write (GB/s)	**380.2**	**379.9**	
Normalized	100%	99.92%	
Standard Deviation	0%	0.1%	
cl-mem - Copy (GB/s)	**297.6**	**296.9**	
Normalized	100%	99.76%	
Standard Deviation	0.1%	0.1%	
Waifu2x-NCNN Vulkan - 2x - 3 - Yes (sec)	**4.297**	4.303	**4.321**
Normalized	100%	99.86%	99.44%
Standard Deviation	0.1%	0.2%	0%
Betsy GPU Compressor - ETC1 - Highest	**4.309**	**4.300**	
Normalized	99.79%	100%	
Standard Deviation	0.4%	0.8%	
Hashcat - T.R.X (H/s)	**501033**	**502833**	
Normalized	99.64%	100%	
Standard Deviation	0.5%	0.4%	
MandelGPU - GPU (Samples/sec)	**319970659**	**319688588**	
Normalized	100%	99.91%	
Standard Deviation	0.4%	0.6%	
SHOC Scalable HeterOgeneous Computing - OpenCL - GEMM SGEMM_N (GFLOPS)	**3807**	**3769**	
Normalized	100%	99.02%	
Standard Deviation	0.5%	1%	
ArrayFire - C.G.O (ms)	**2.086**	**2.094**	
Normalized	100%	99.62%	
Standard Deviation	0%	0%	
SHOC Scalable HeterOgeneous Computing - OpenCL - Bus Speed Readback (GB/s)	**26.3950**	**26.3909**	
Normalized	100%	99.98%	
Standard Deviation	0%	0.1%	
Hashcat - 7-Zip (H/s)	**686733**	**686700**	
Normalized	100%	100%	
Standard Deviation	0.3%	0.3%	
SHOC Scalable HeterOgeneous Computing - OpenCL - Triad (GB/s)	**24.6721**	**24.6714**	
Normalized	100%	100%	
Standard Deviation	0.1%	0.1%	
clpeak - I.C.I (GIOPS)	**10264**	**10202**	
Normalized	100%	99.4%	
Standard Deviation	2.3%	1.5%	
SHOC Scalable HeterOgeneous Computing - OpenCL - Reduction (GB/s)	**325.671**	**325.801**	
Normalized	99.96%	100%	
Standard Deviation	0.3%	0.2%	
SHOC Scalable HeterOgeneous Computing - OpenCL - FFT SP (GFLOPS)	**1135**	**1134**	
Normalized	100%	99.9%	
Standard Deviation	0.1%	0.2%	

RTX 3070 Compute

clpeak - G.M.B (GBPS)	**389.57**	**389.61**
Normalized	99.99%	100%
Standard Deviation	0%	0%
SHOC Scalable HeterOgeneous Computing - OpenCL - MD5 Hash (GHash/s)	**25.4712**	**25.4940**
Normalized	99.91%	100%
Standard Deviation	0.1%	0.3%
SHOC Scalable HeterOgeneous Computing - OpenCL - Bus Speed Download (GB/s)	**26.2814**	**26.3097**
Normalized	99.89%	100%
Standard Deviation	0.2%	0.2%
FinanceBench - B.S.O (ms)	**10.483**	**10.496**
Normalized	100%	99.88%
Standard Deviation	0.2%	0.6%
Mixbench - OpenCL - Integer (GIOPS)	**11434**	**11337**
Normalized	100%	99.15%
Standard Deviation	0%	0.8%
clpeak - S.P.F (GFLOPS)	**20100**	**19992**
Normalized	100%	99.46%
Standard Deviation	0%	0.9%
Mixbench - OpenCL - Double Precision (GFLOPS)	**295.11**	**299.65**
Normalized	98.48%	100%
Standard Deviation	1.1%	0.4%
Mixbench - OpenCL - Single Precision (GFLOPS)	**22082**	**21966**
Normalized	100%	99.47%
Standard Deviation	0.1%	0.9%
SHOC Scalable HeterOgeneous Computing - OpenCL - S3D (GFLOPS)	**218.440**	**218.605**
Normalized	99.92%	100%
Standard Deviation	0.2%	0.2%

RTX 3070 Compute

SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Max SP Flops

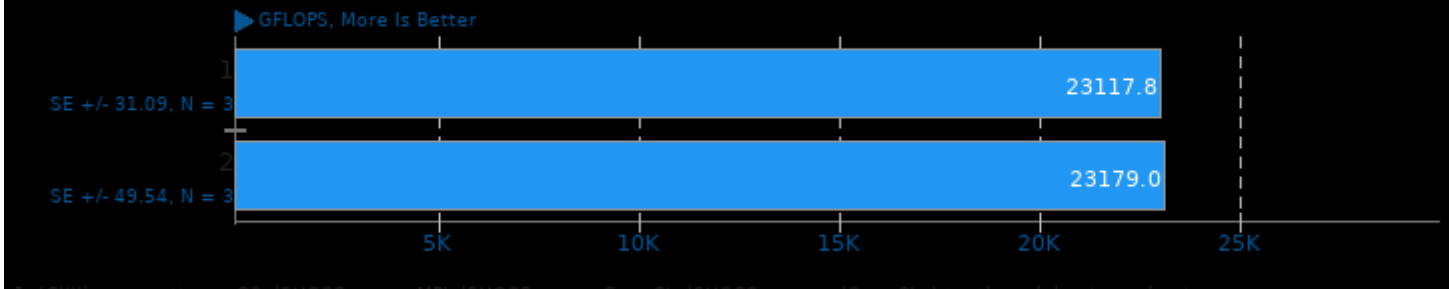

1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

Blender 2.92

Blend File: Barbershop - Compute: CUDA

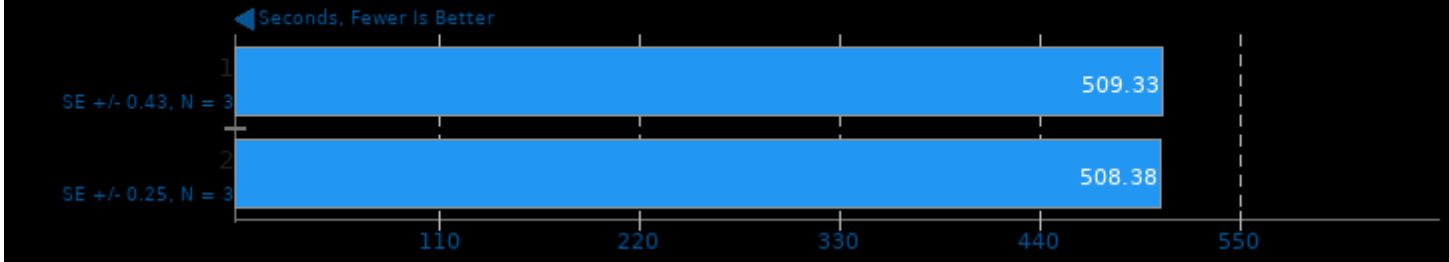

Blender 2.92

Blend File: Barbershop - Compute: NVIDIA OptiX

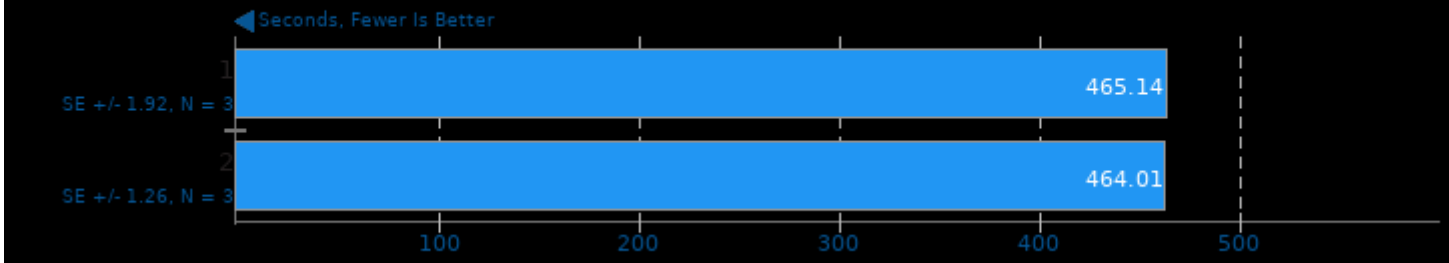

LeelaChessZero 0.26

Backend: OpenCL

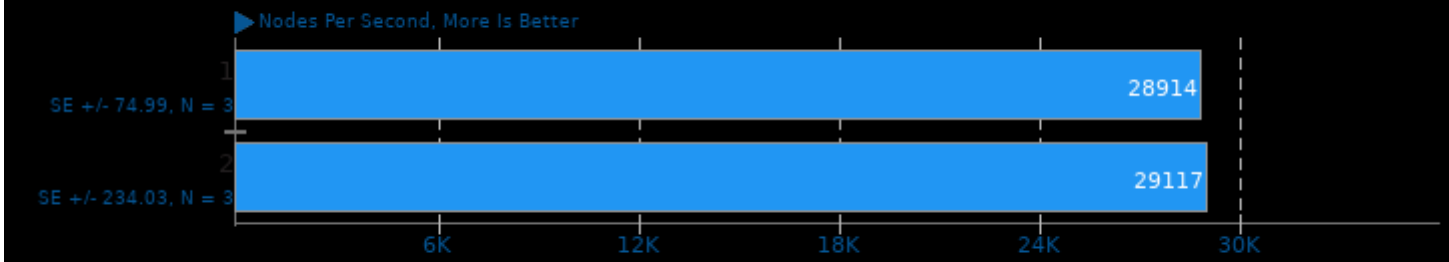

1. (CXX) g++ options: -flto -pthread

RTX 3070 Compute

VkFFT 1.1.1

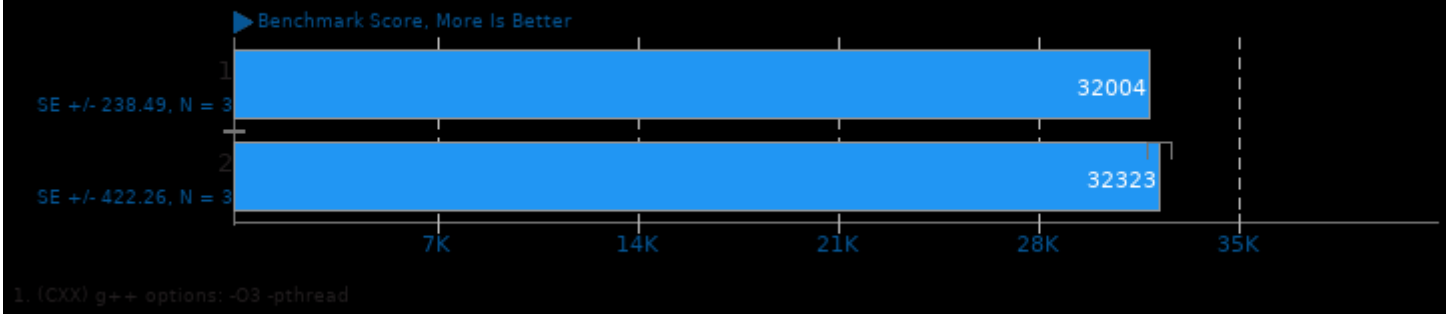

1. (CXX) g++ options: -O3 -pthread

RedShift Demo 3.0

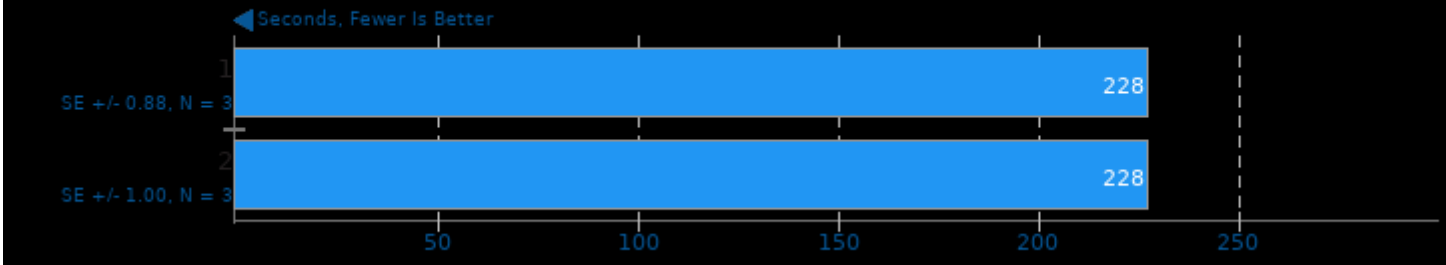

Blender 2.92

Blend File: Pabellon Barcelona - Compute: CUDA

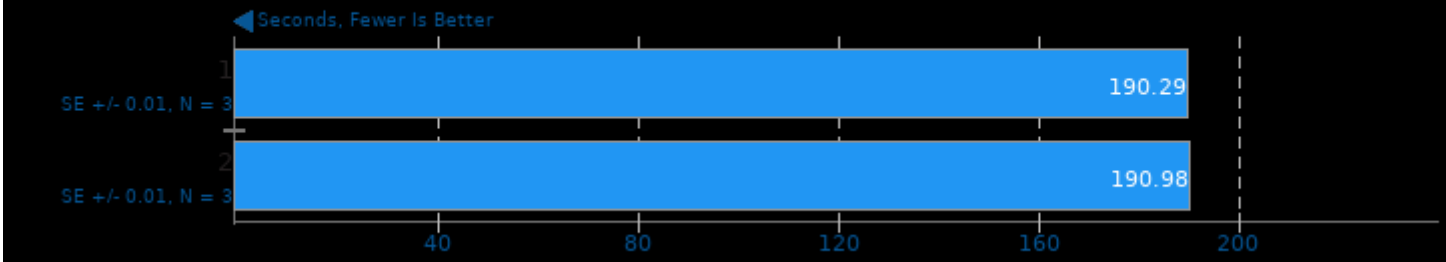

OctaneBench 2020.1

Total Score

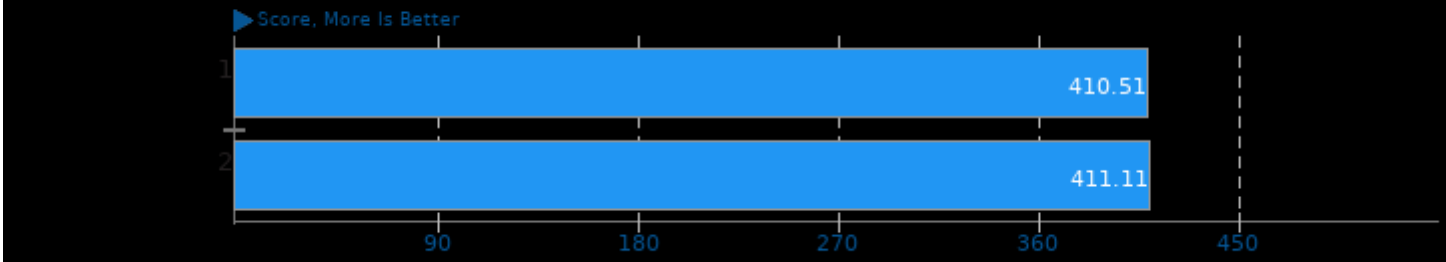

RTX 3070 Compute

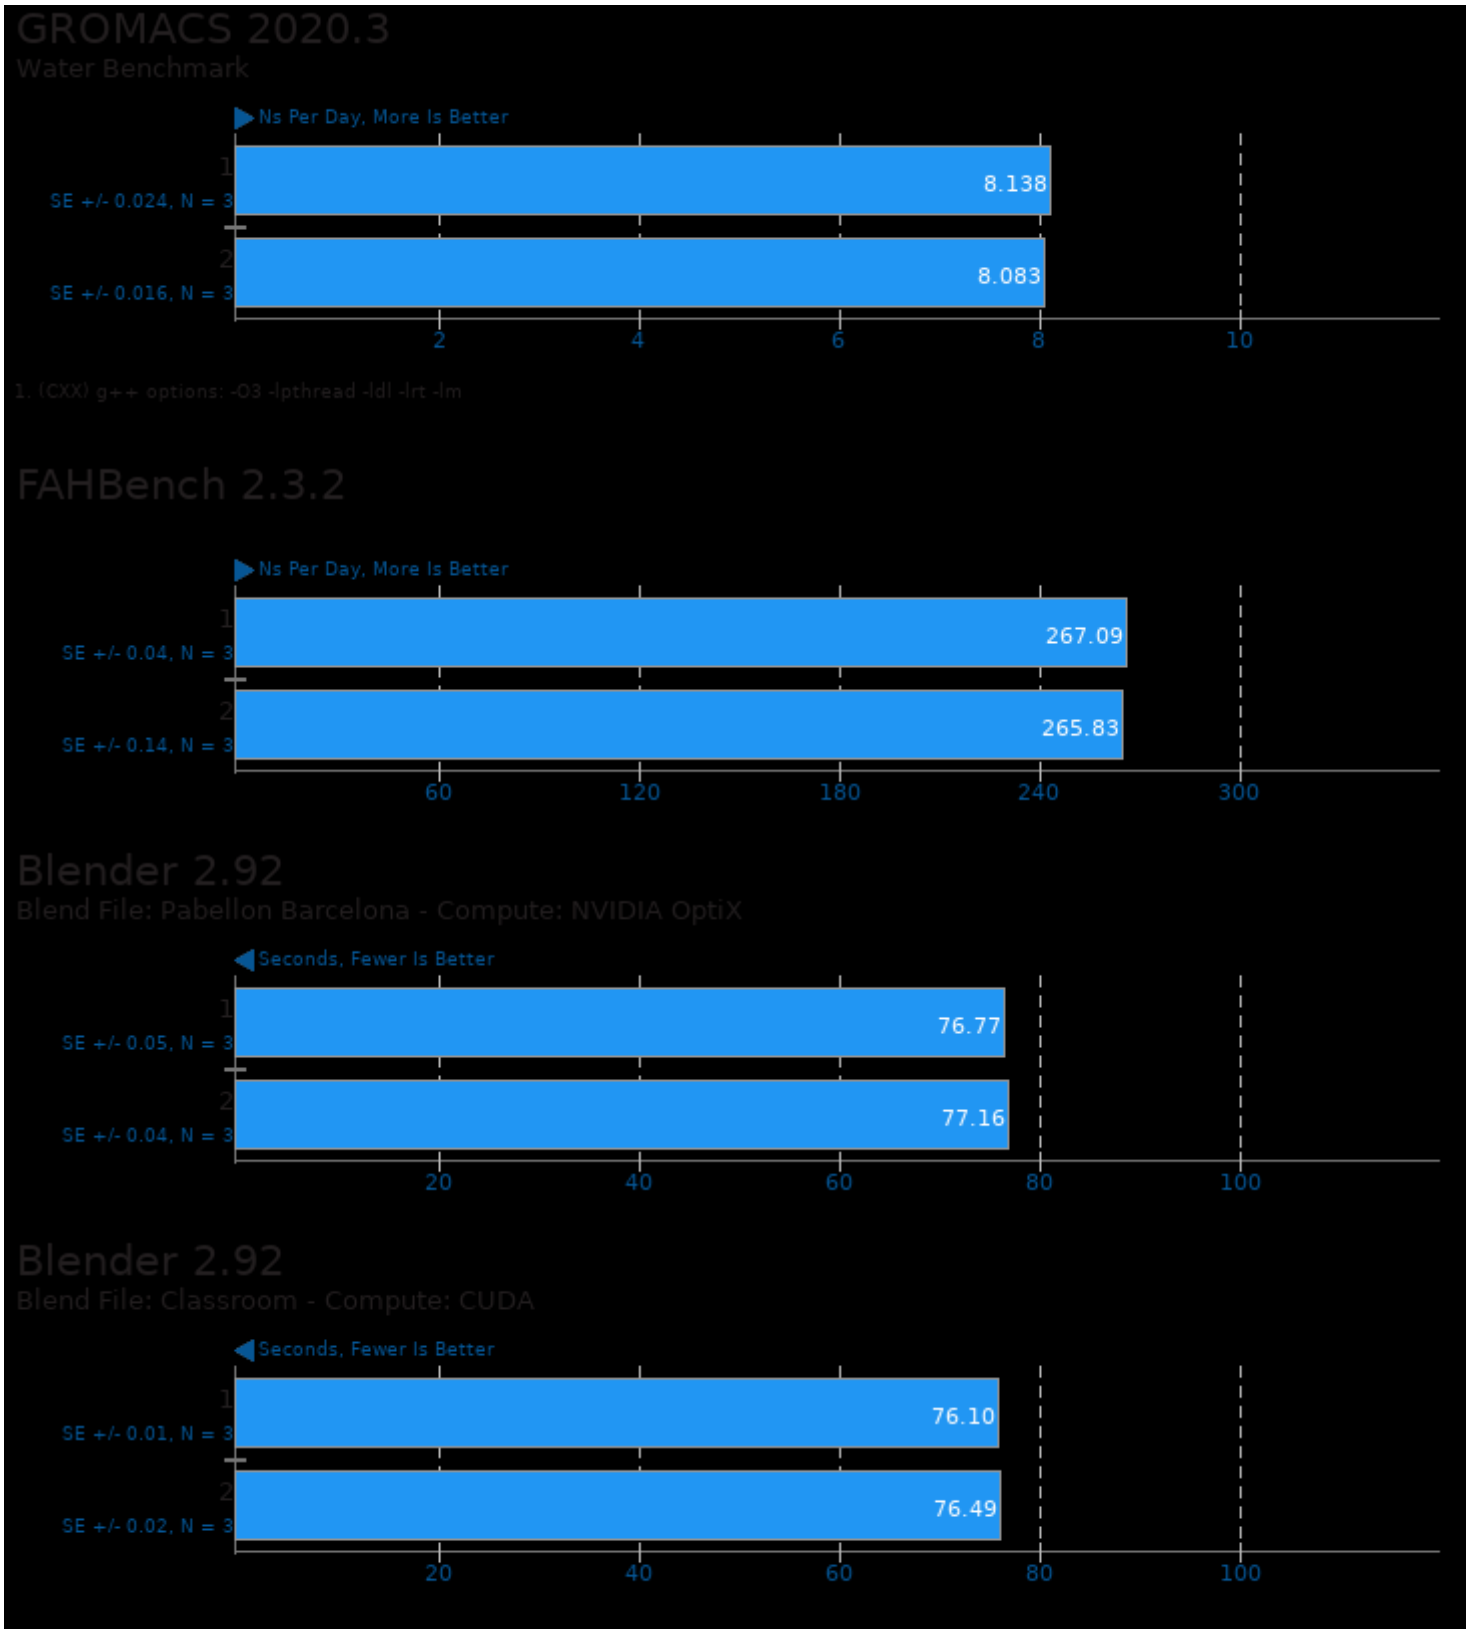

RTX 3070 Compute

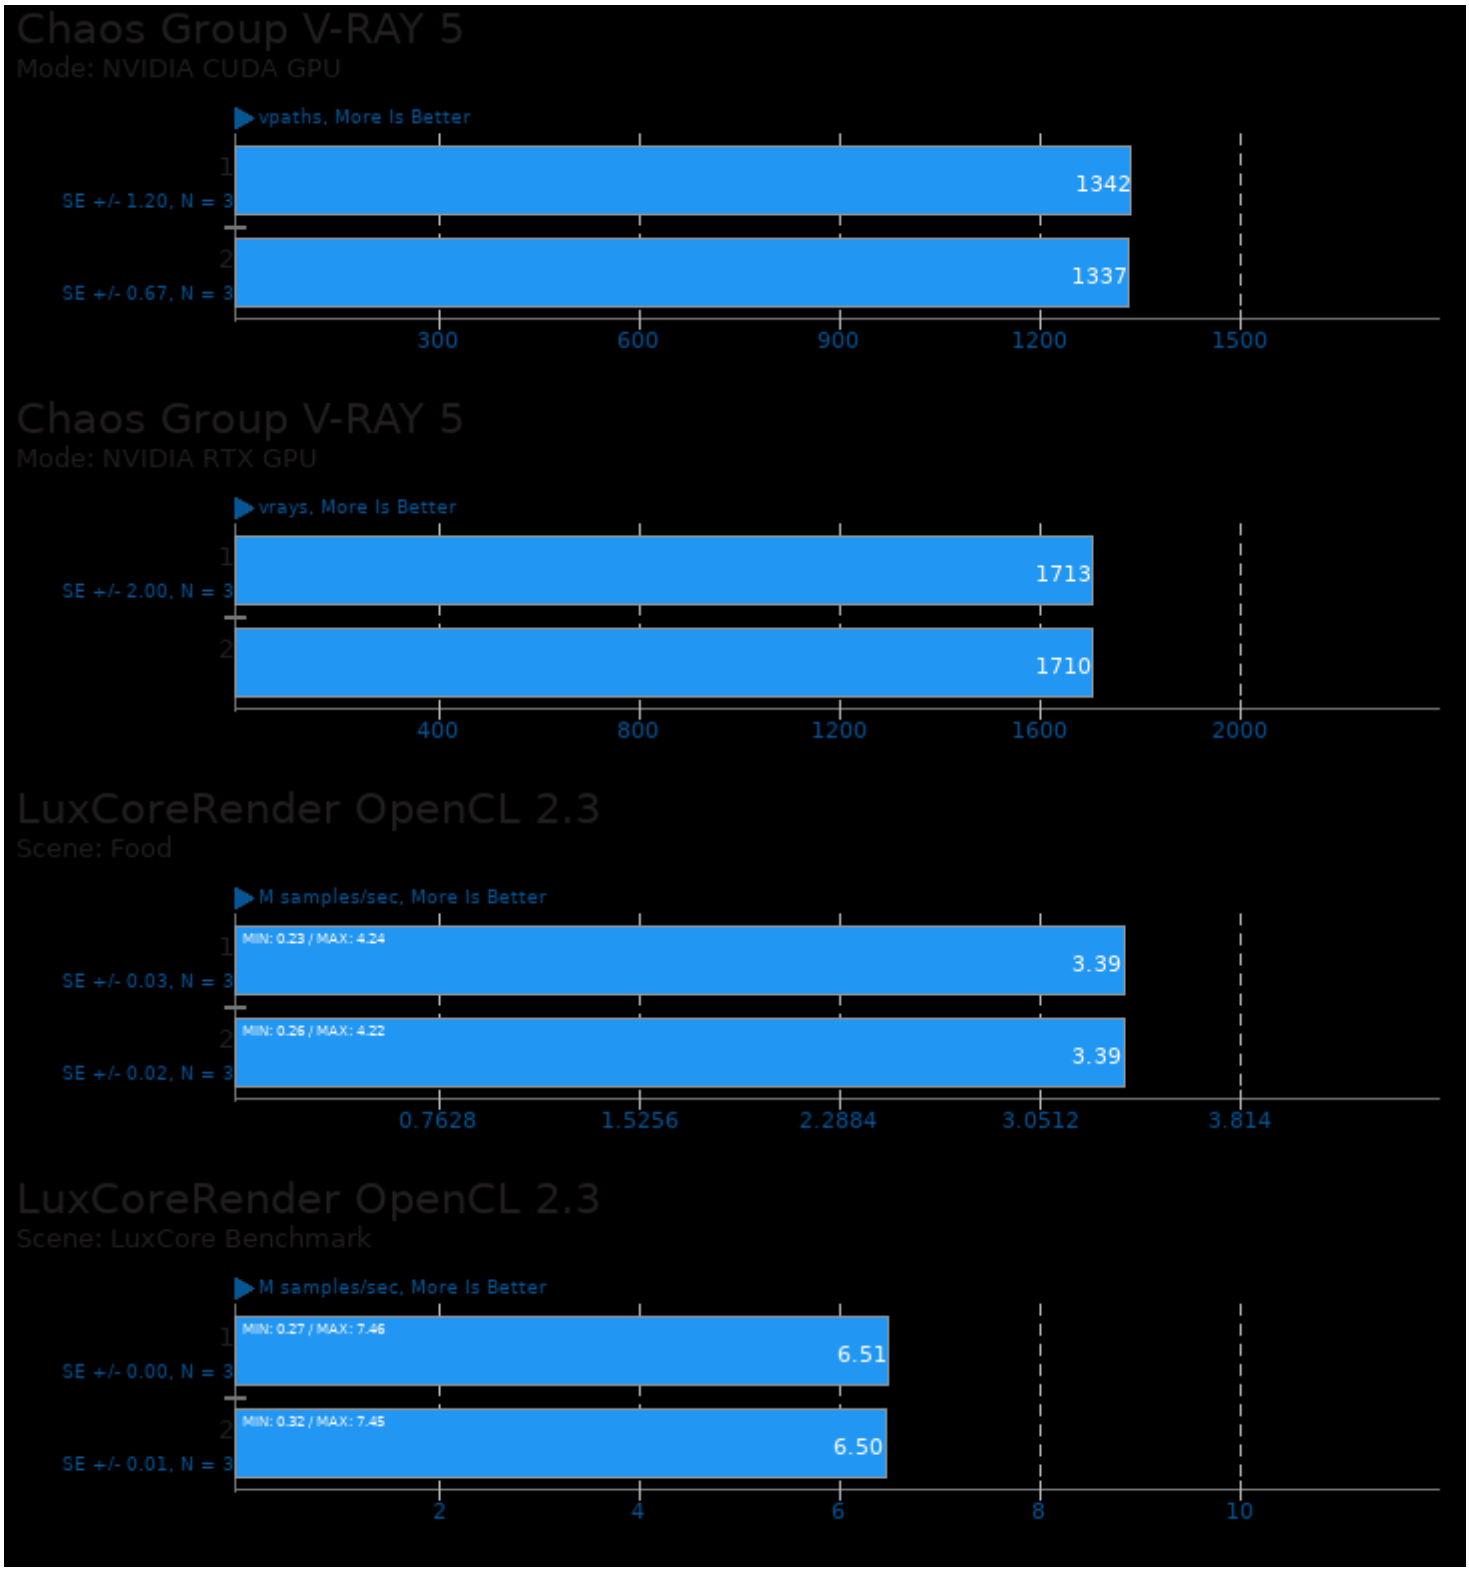

RTX 3070 Compute

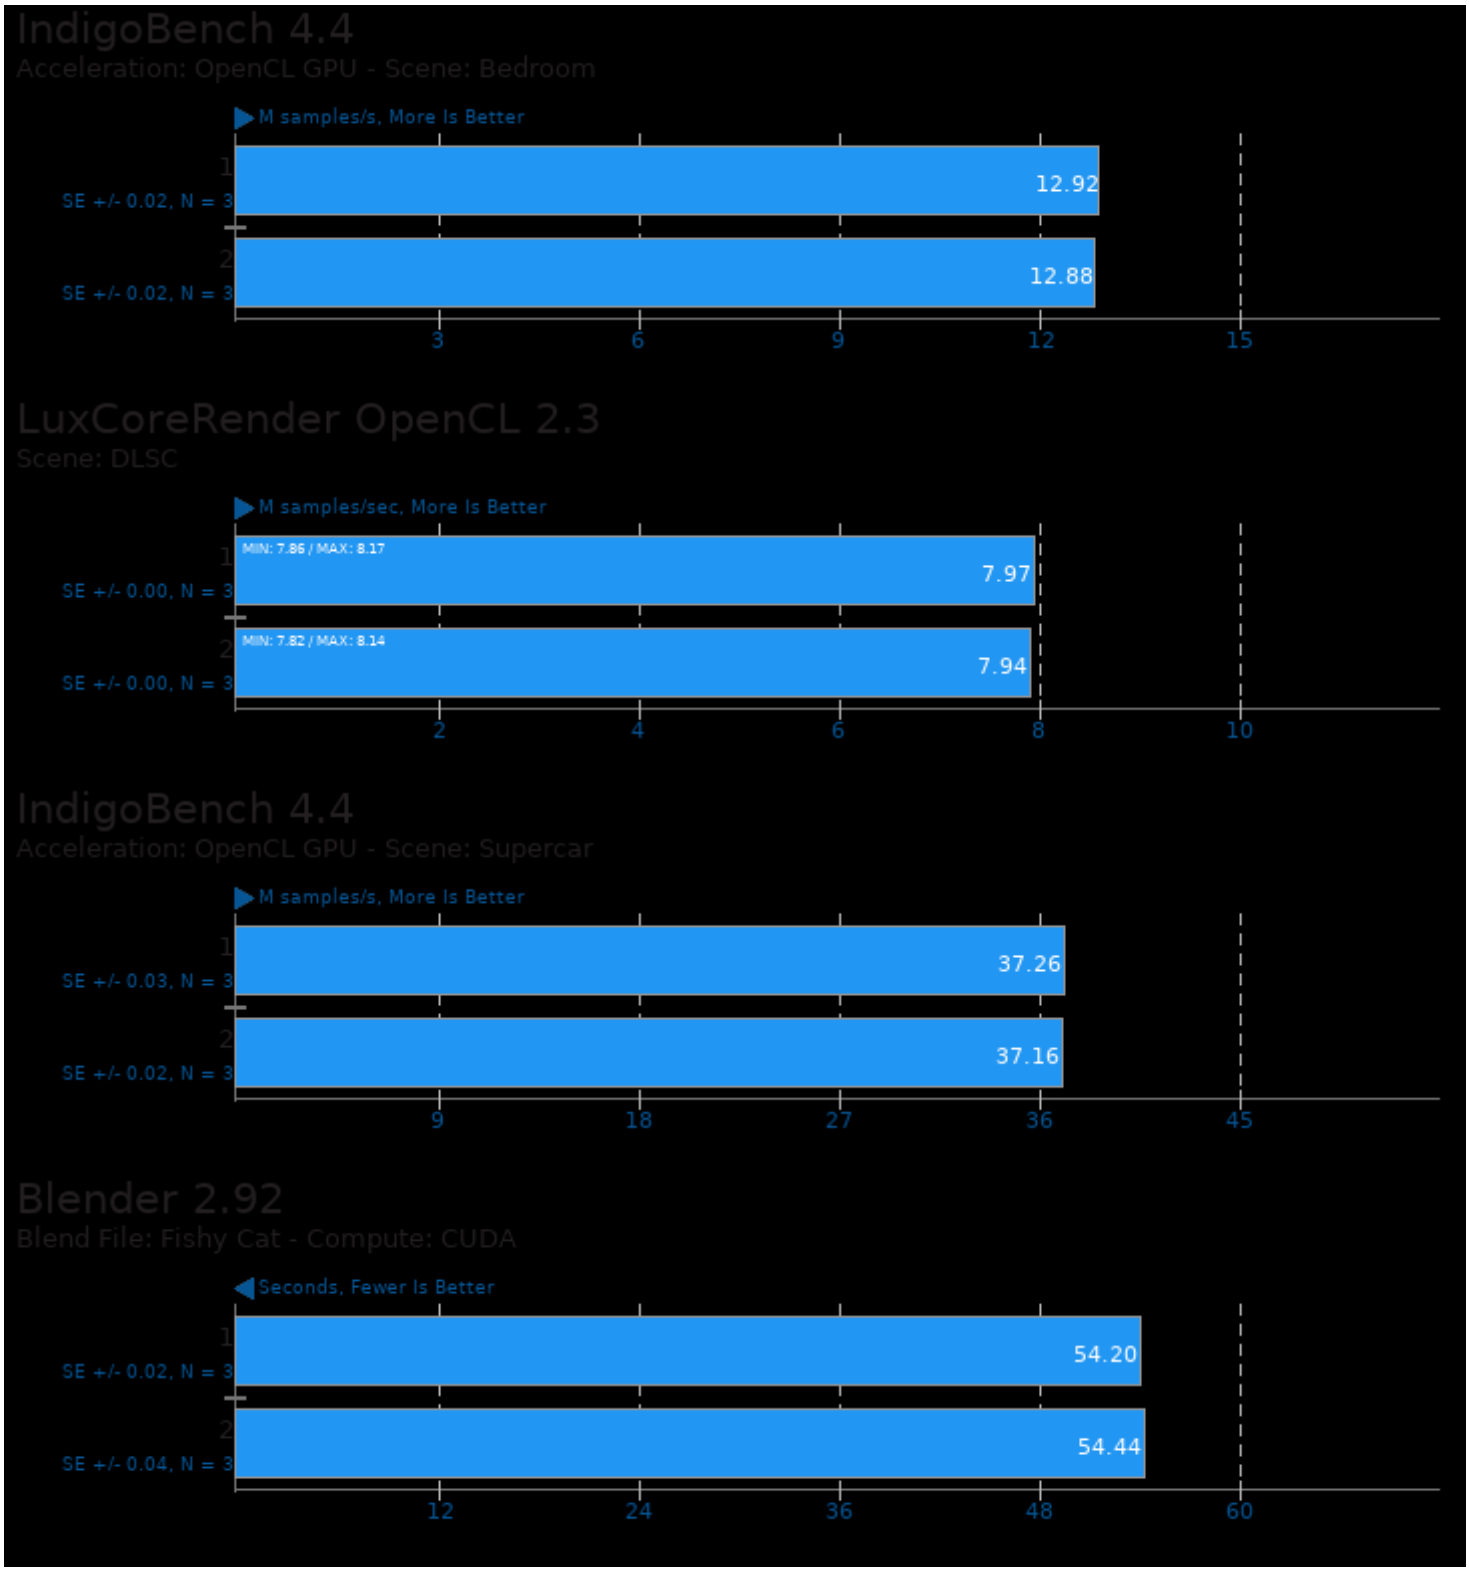

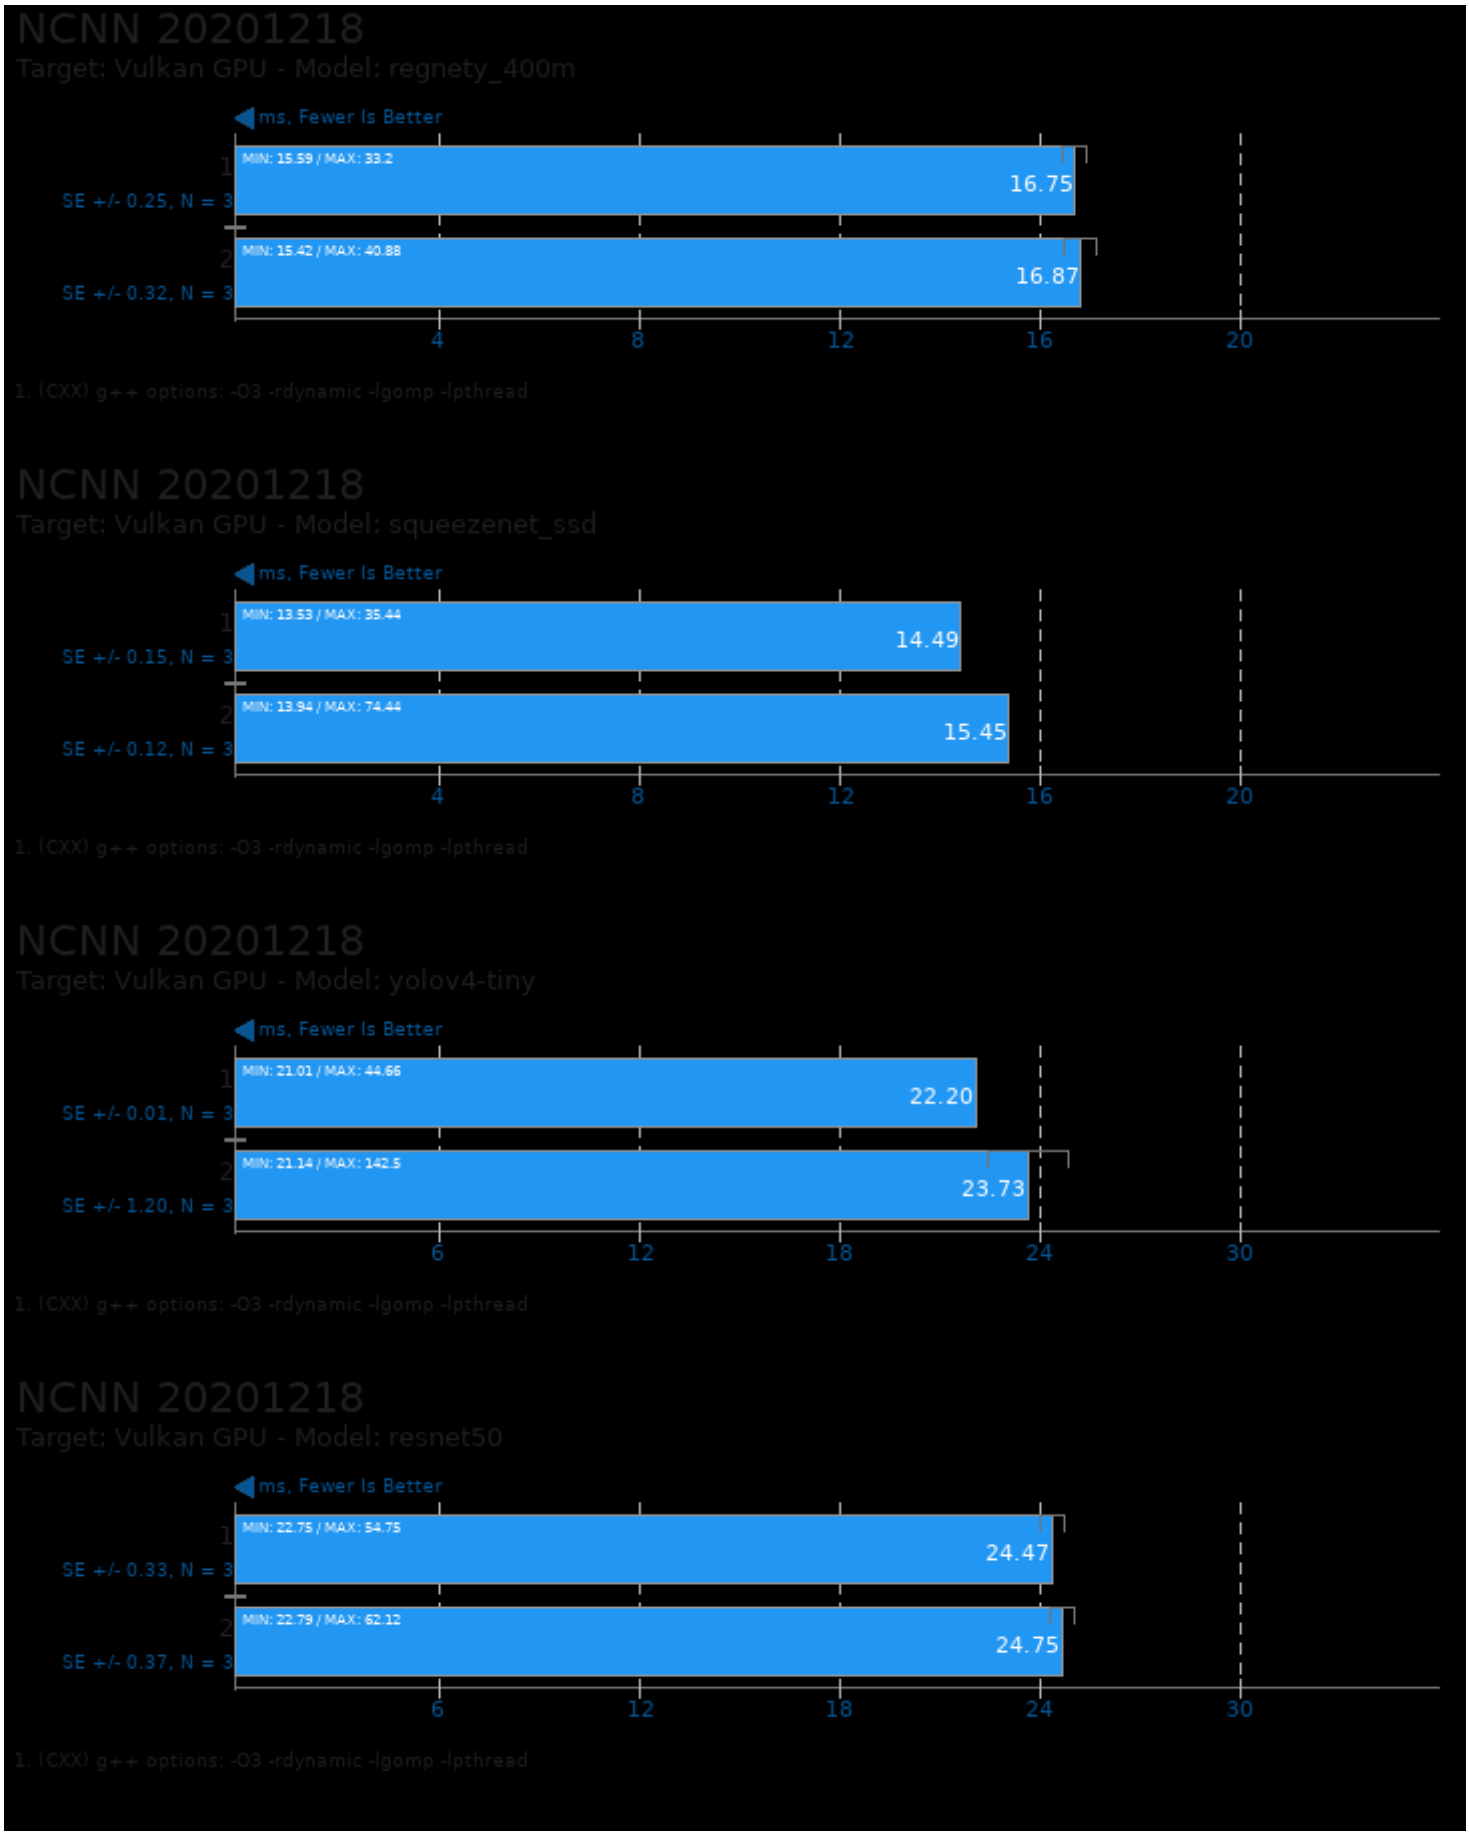
NCNN 20201218
Target: Vulkan GPU - Model: regnety_400m
ms, Fewer Is Better
MIN: 15.59 / MAX: 33.2
16.75
SE +/- 0.25, N = 3
MIN: 15.42 / MAX: 40.88
16.87
SE +/- 0.32, N = 3
4
8
12
16
20
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU - Model: squeezenet_ssd
ms, Fewer Is Better
MIN: 13.53 / MAX: 35.44
14.49
SE +/- 0.15, N = 3
MIN: 13.94 / MAX: 74.44
15.45
SE +/- 0.12, N = 3
4
8
12
16
20
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU - Model: yolov4-tiny
ms, Fewer Is Better
MIN: 21.01 / MAX: 44.66
22.20
SE +/- 0.01, N = 3
MIN: 21.14 / MAX: 142.5
23.73
SE +/- 1.20, N = 3
6
12
18
24
30
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU - Model: resnet50
ms, Fewer Is Better
MIN: 22.75 / MAX: 54.75
24.47
SE +/- 0.33, N = 3
MIN: 22.79 / MAX: 62.12
24.75
SE +/- 0.37, N = 3
6
12
18
24
30
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

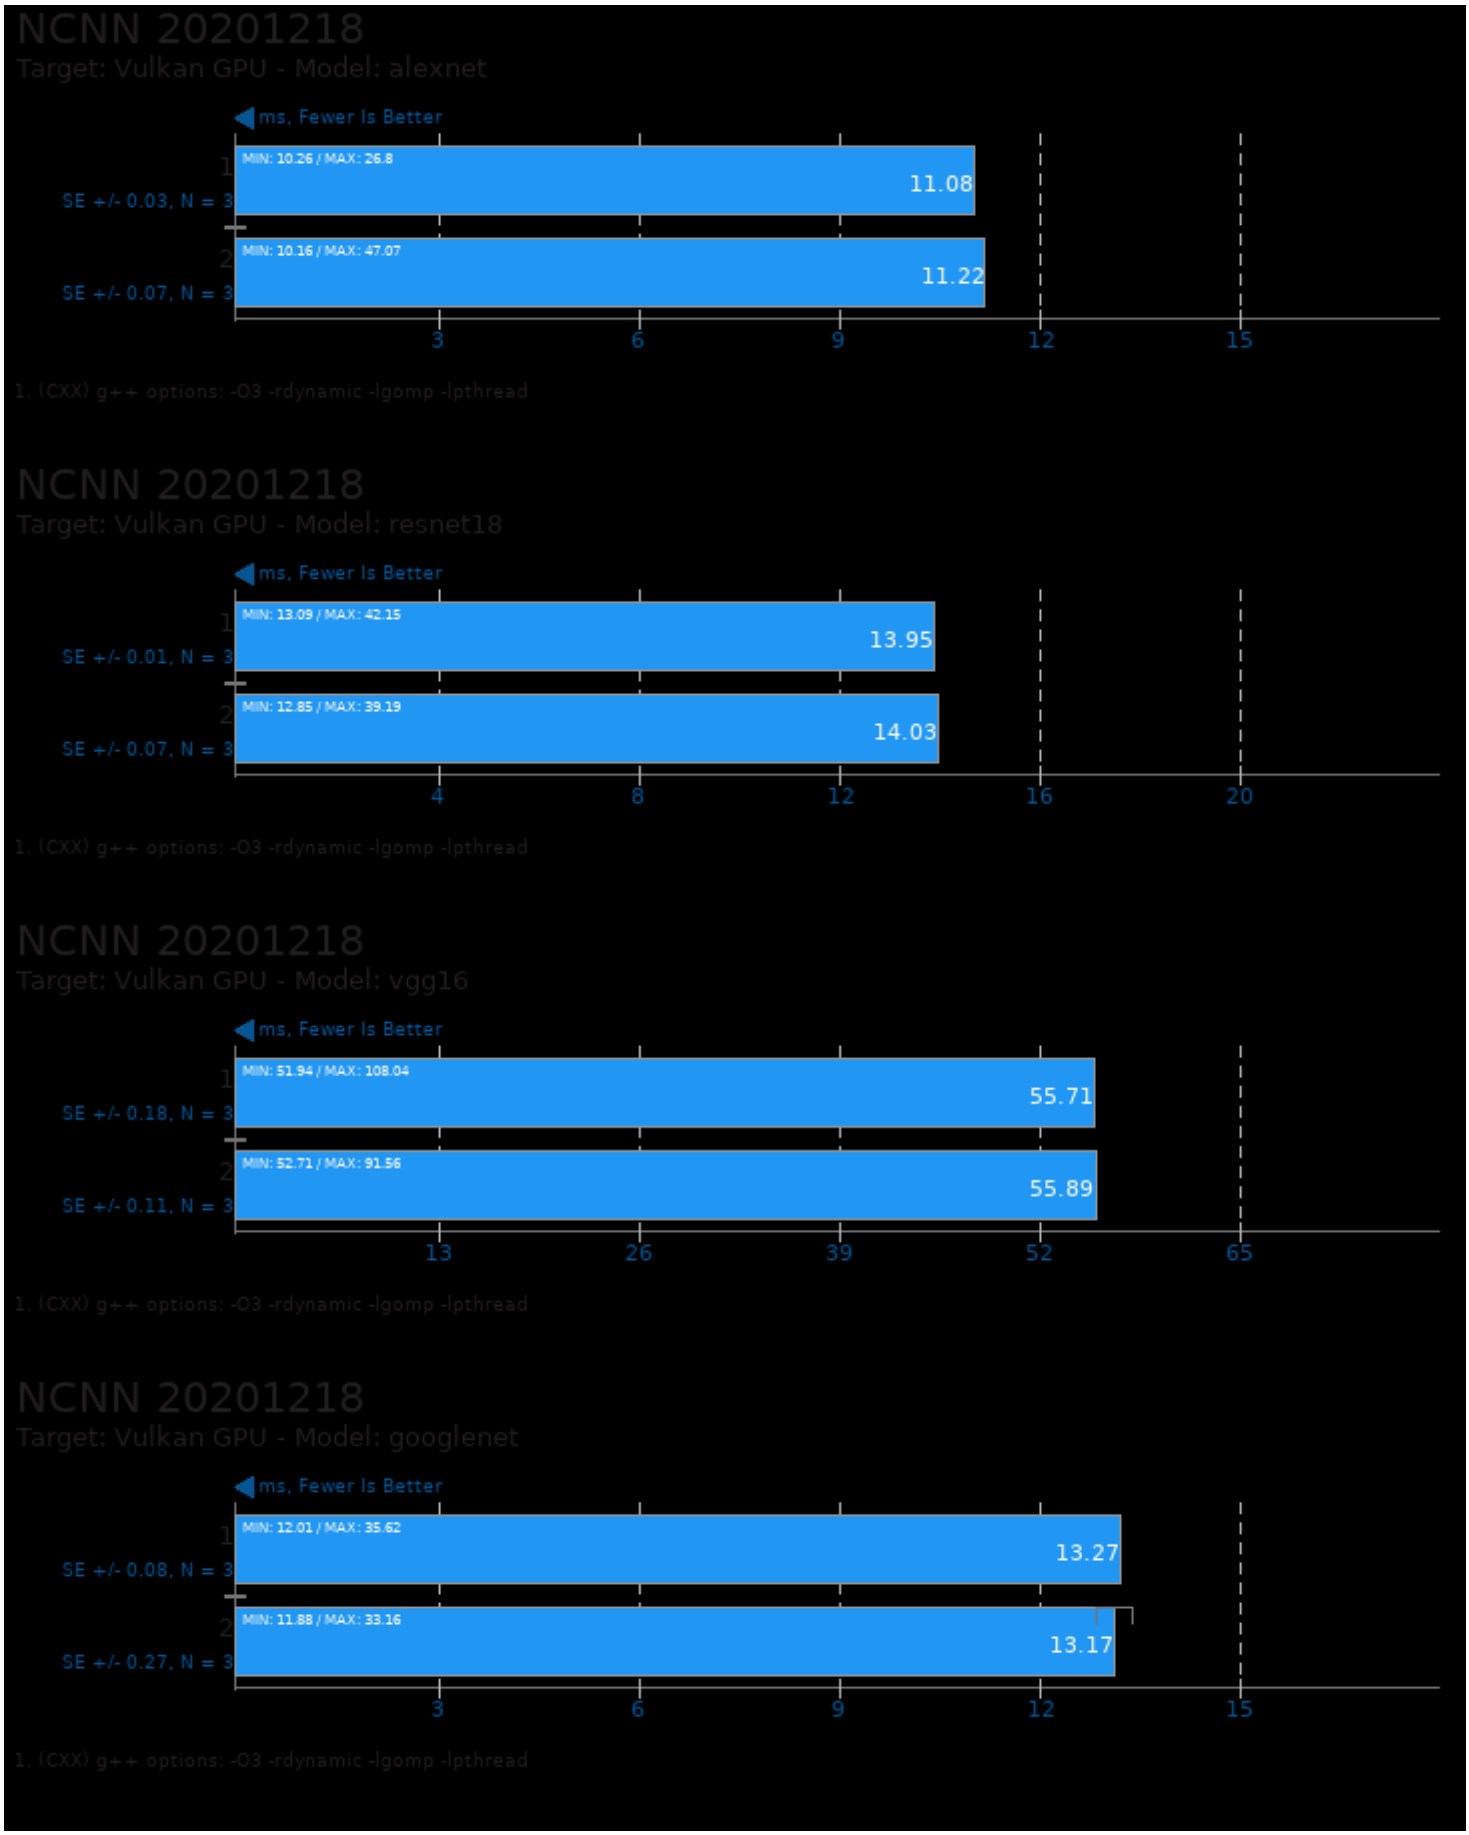
NCNN 20201218
Target: Vulkan GPU - Model: alexnet
ms, Fewer Is Better
1
MIN: 10.26 / MAX: 26.8
11.08
SE +/- 0.03, N = 3
2
MIN: 10.16 / MAX: 47.07
11.22
SE +/- 0.07, N = 3
3
6
9
12
15
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU - Model: resnet18
ms, Fewer Is Better
1
MIN: 13.09 / MAX: 42.15
13.95
SE +/- 0.01, N = 3
2
MIN: 12.85 / MAX: 39.19
14.03
SE +/- 0.07, N = 3
4
8
12
16
20
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU - Model: vgg16
ms, Fewer Is Better
1
MIN: 51.94 / MAX: 108.04
55.71
SE +/- 0.18, N = 3
2
MIN: 52.71 / MAX: 91.56
55.89
SE +/- 0.11, N = 3
13
26
39
52
65
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU - Model: googlenet
ms, Fewer Is Better
1
MIN: 12.01 / MAX: 35.62
13.27
SE +/- 0.08, N = 3
2
MIN: 11.88 / MAX: 33.16
13.17
SE +/- 0.27, N = 3
3
6
9
12
15
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

RTX 3070 Compute

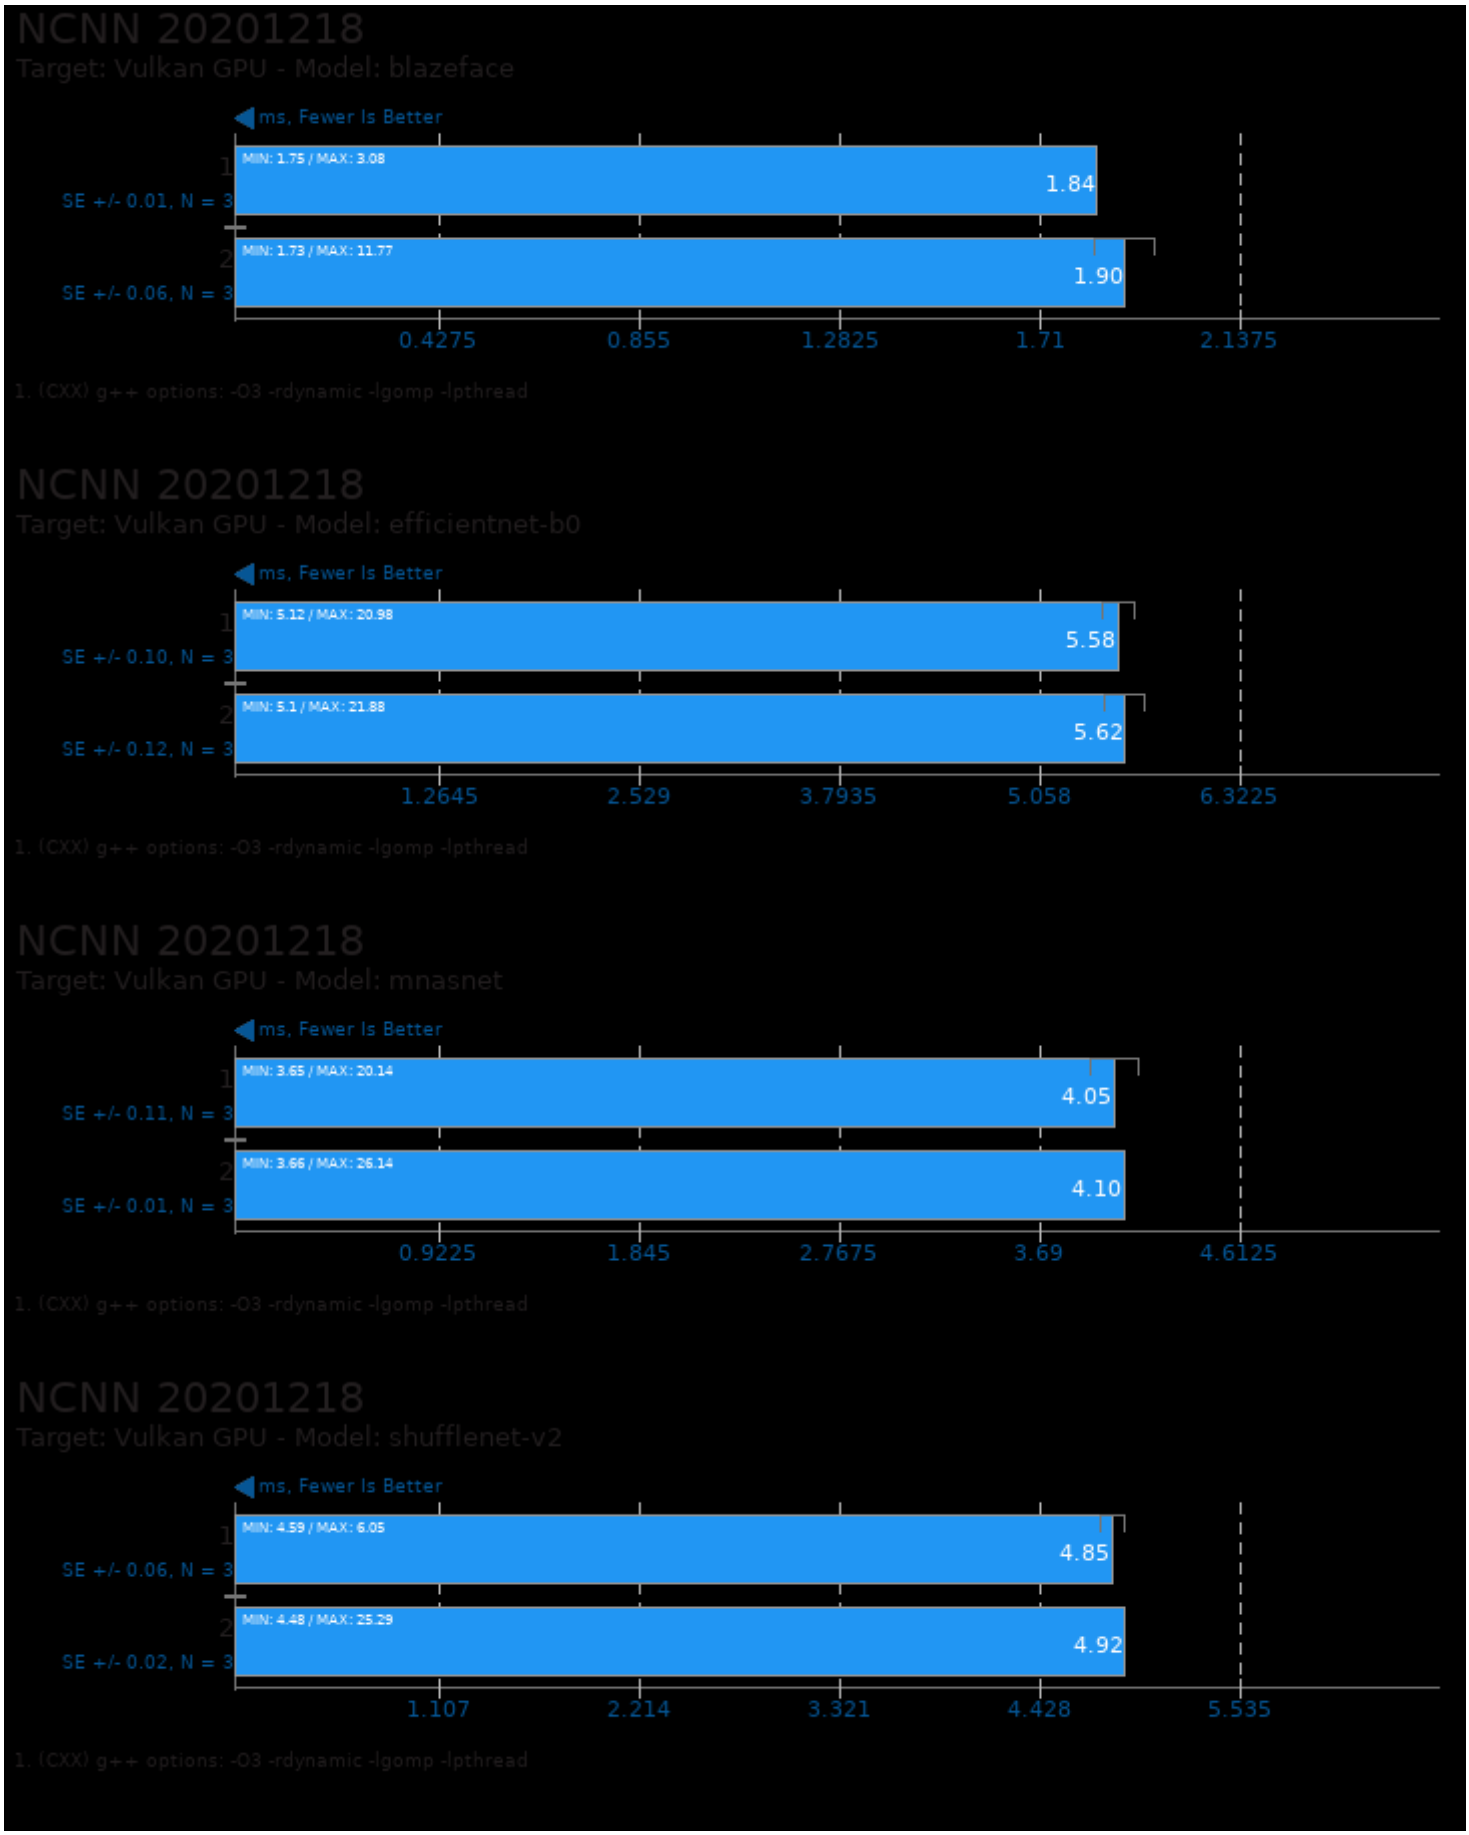

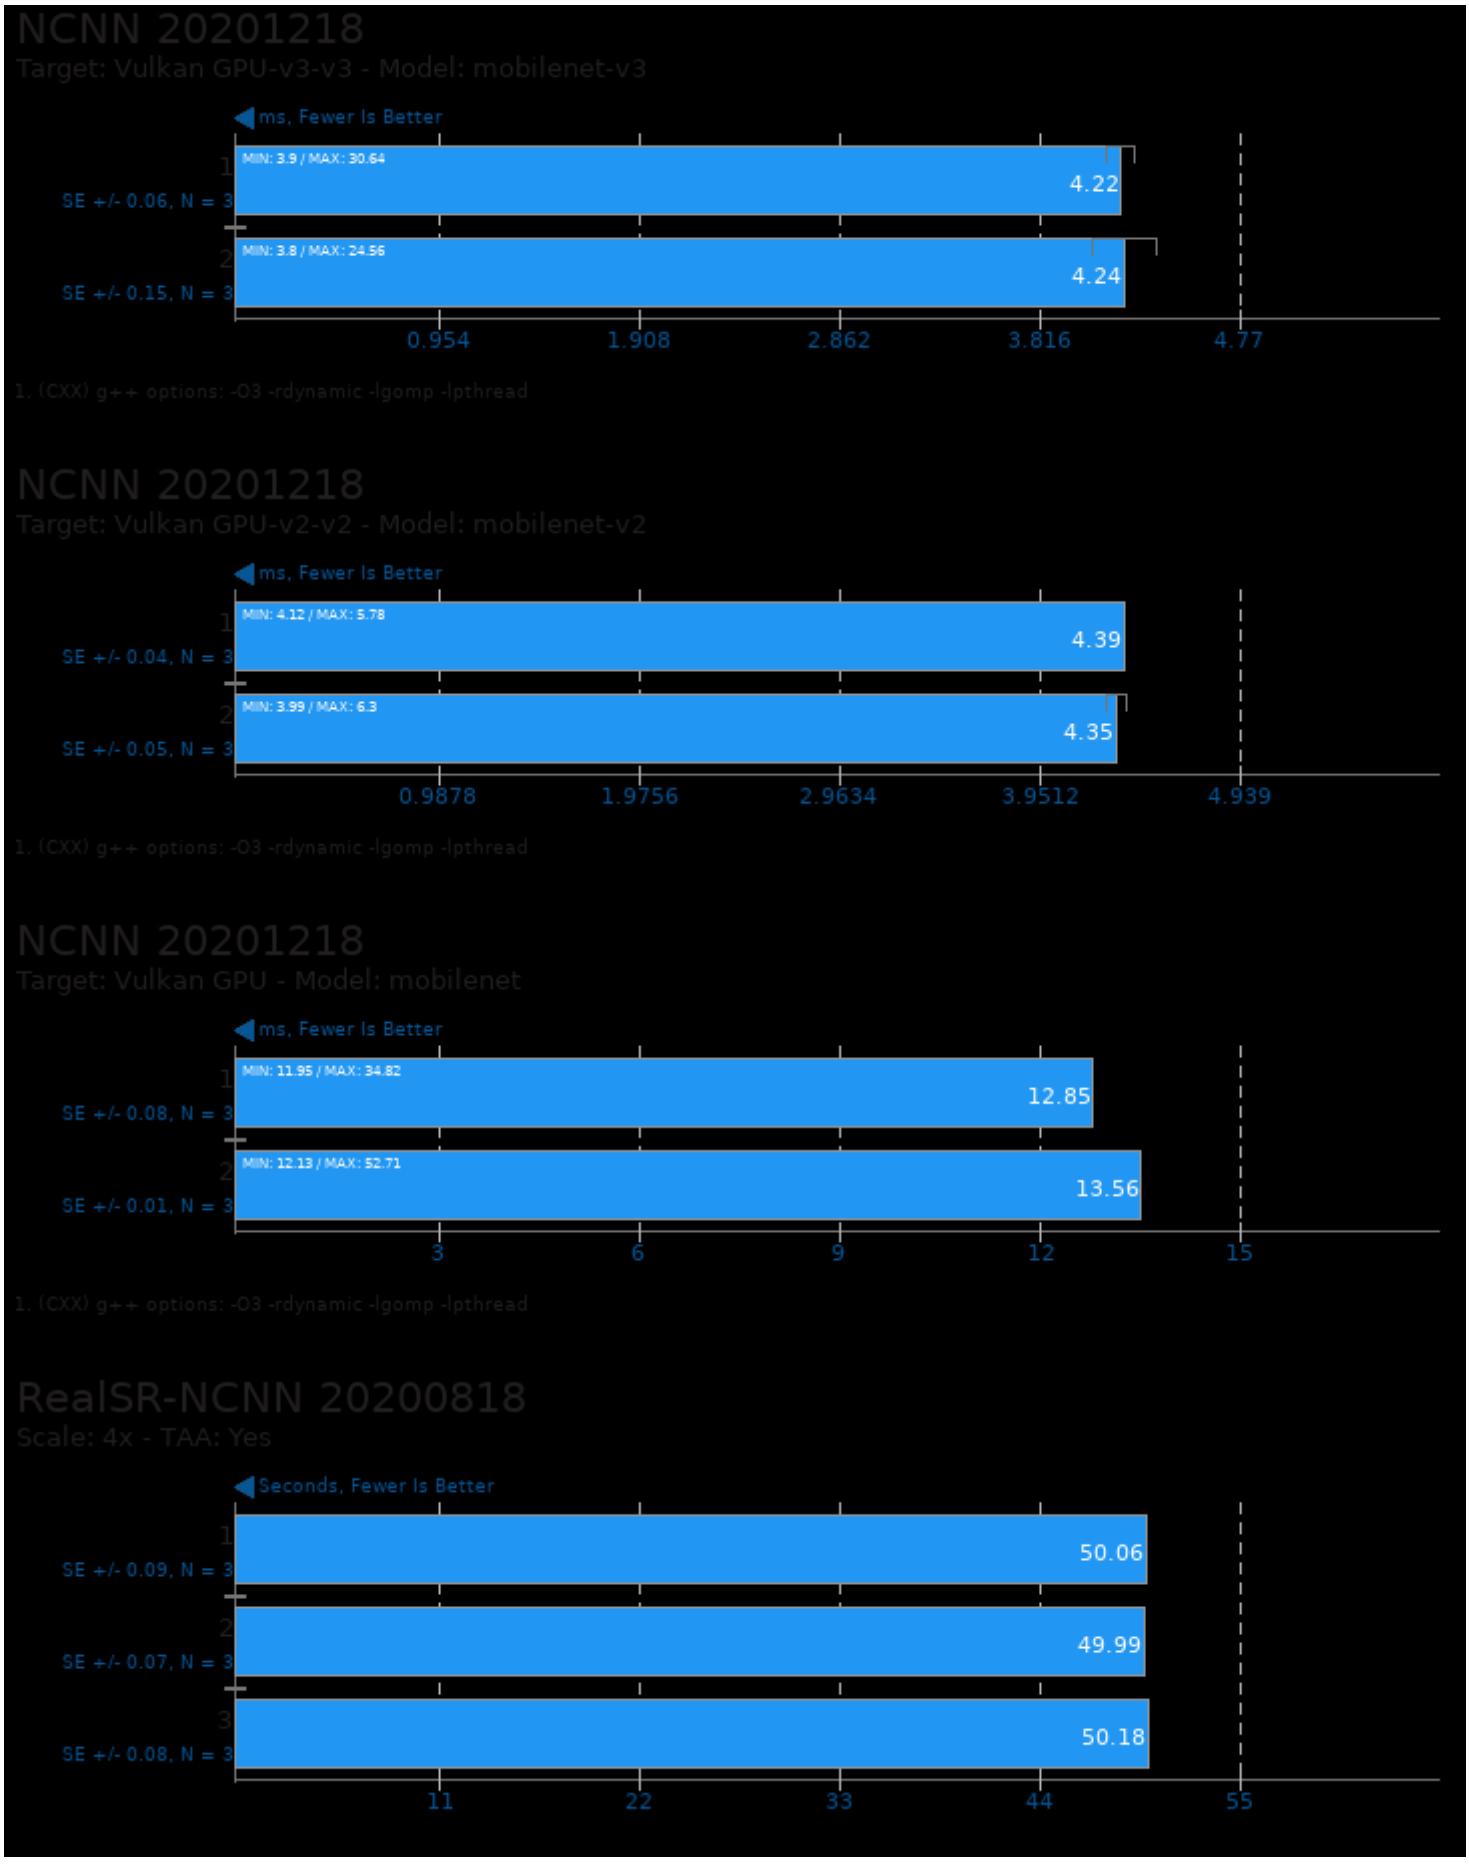
NCNN 20201218
Target: Vulkan GPU-v3-v3 - Model: mobilenet-v3
ms, Fewer Is Better
1
SE +/- 0.06, N = 3
MIN: 3.9 / MAX: 30.64
4.22
2
SE +/- 0.15, N = 3
MIN: 3.8 / MAX: 24.56
4.24
0.954
1.908
2.862
3.816
4.77
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU-v2-v2 - Model: mobilenet-v2
ms, Fewer Is Better
1
SE +/- 0.04, N = 3
MIN: 4.12 / MAX: 5.78
4.39
2
SE +/- 0.05, N = 3
MIN: 3.99 / MAX: 6.3
4.35
0.9878
1.9756
2.9634
3.9512
4.939
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20201218
Target: Vulkan GPU - Model: mobilenet
ms, Fewer Is Better
1
SE +/- 0.08, N = 3
MIN: 11.95 / MAX: 34.82
12.85
2
SE +/- 0.01, N = 3
MIN: 12.13 / MAX: 52.71
13.56
3
6
9
12
15
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
RealSR-NCNN 20200818
Scale: 4x - TAA: Yes
Seconds, Fewer Is Better
1
SE +/- 0.09, N = 3
50.06
2
SE +/- 0.07, N = 3
49.99
3
SE +/- 0.08, N = 3
50.18
11
22
33
44
55

RTX 3070 Compute

Blender 2.92

Blend File: Classroom - Compute: NVIDIA OptiX

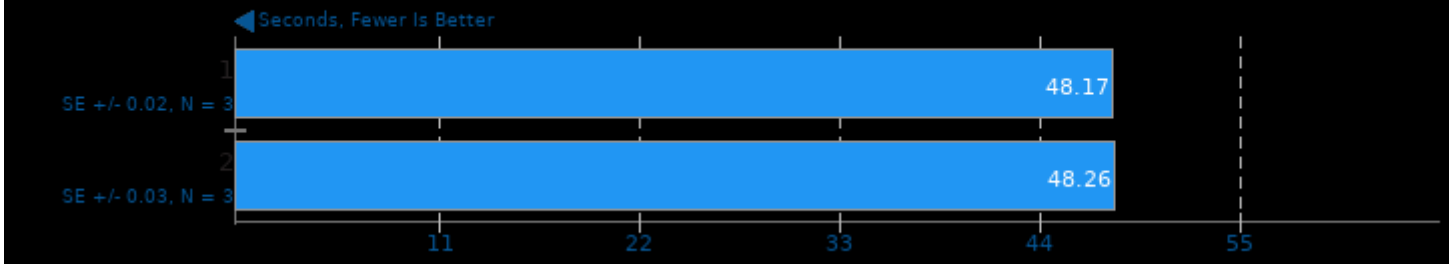

Blender 2.92

Blend File: Fishy Cat - Compute: NVIDIA OptiX

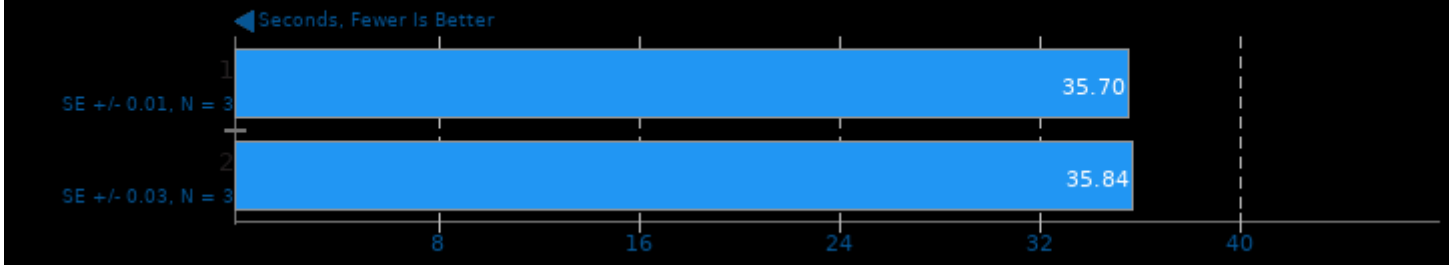

VkResample 1.0

Upscale: 2x - Precision: Double

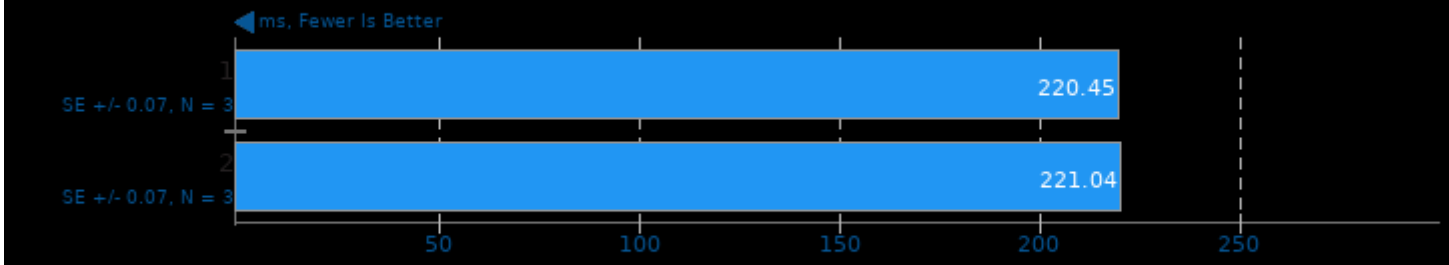

1. (CXX) g++ options: -O3 -pthread

clpeak

OpenCL Test: Double-Precision Double

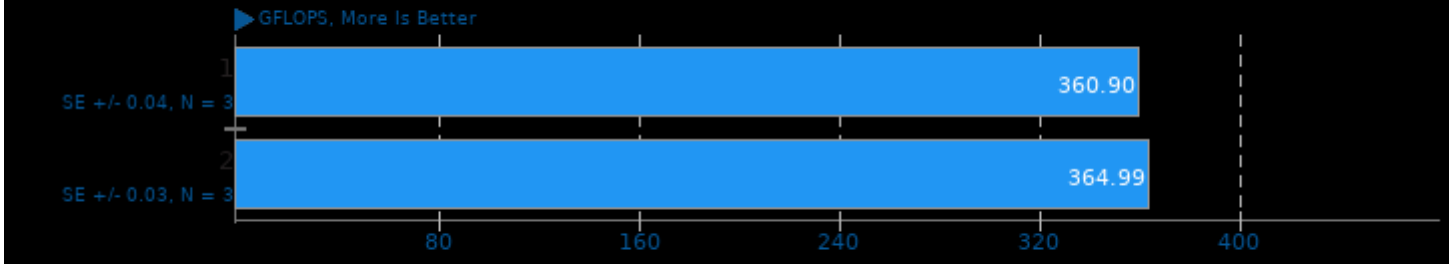

1. (CXX) g++ options: -O3 -rdynamic -lOpenCL

RTX 3070 Compute

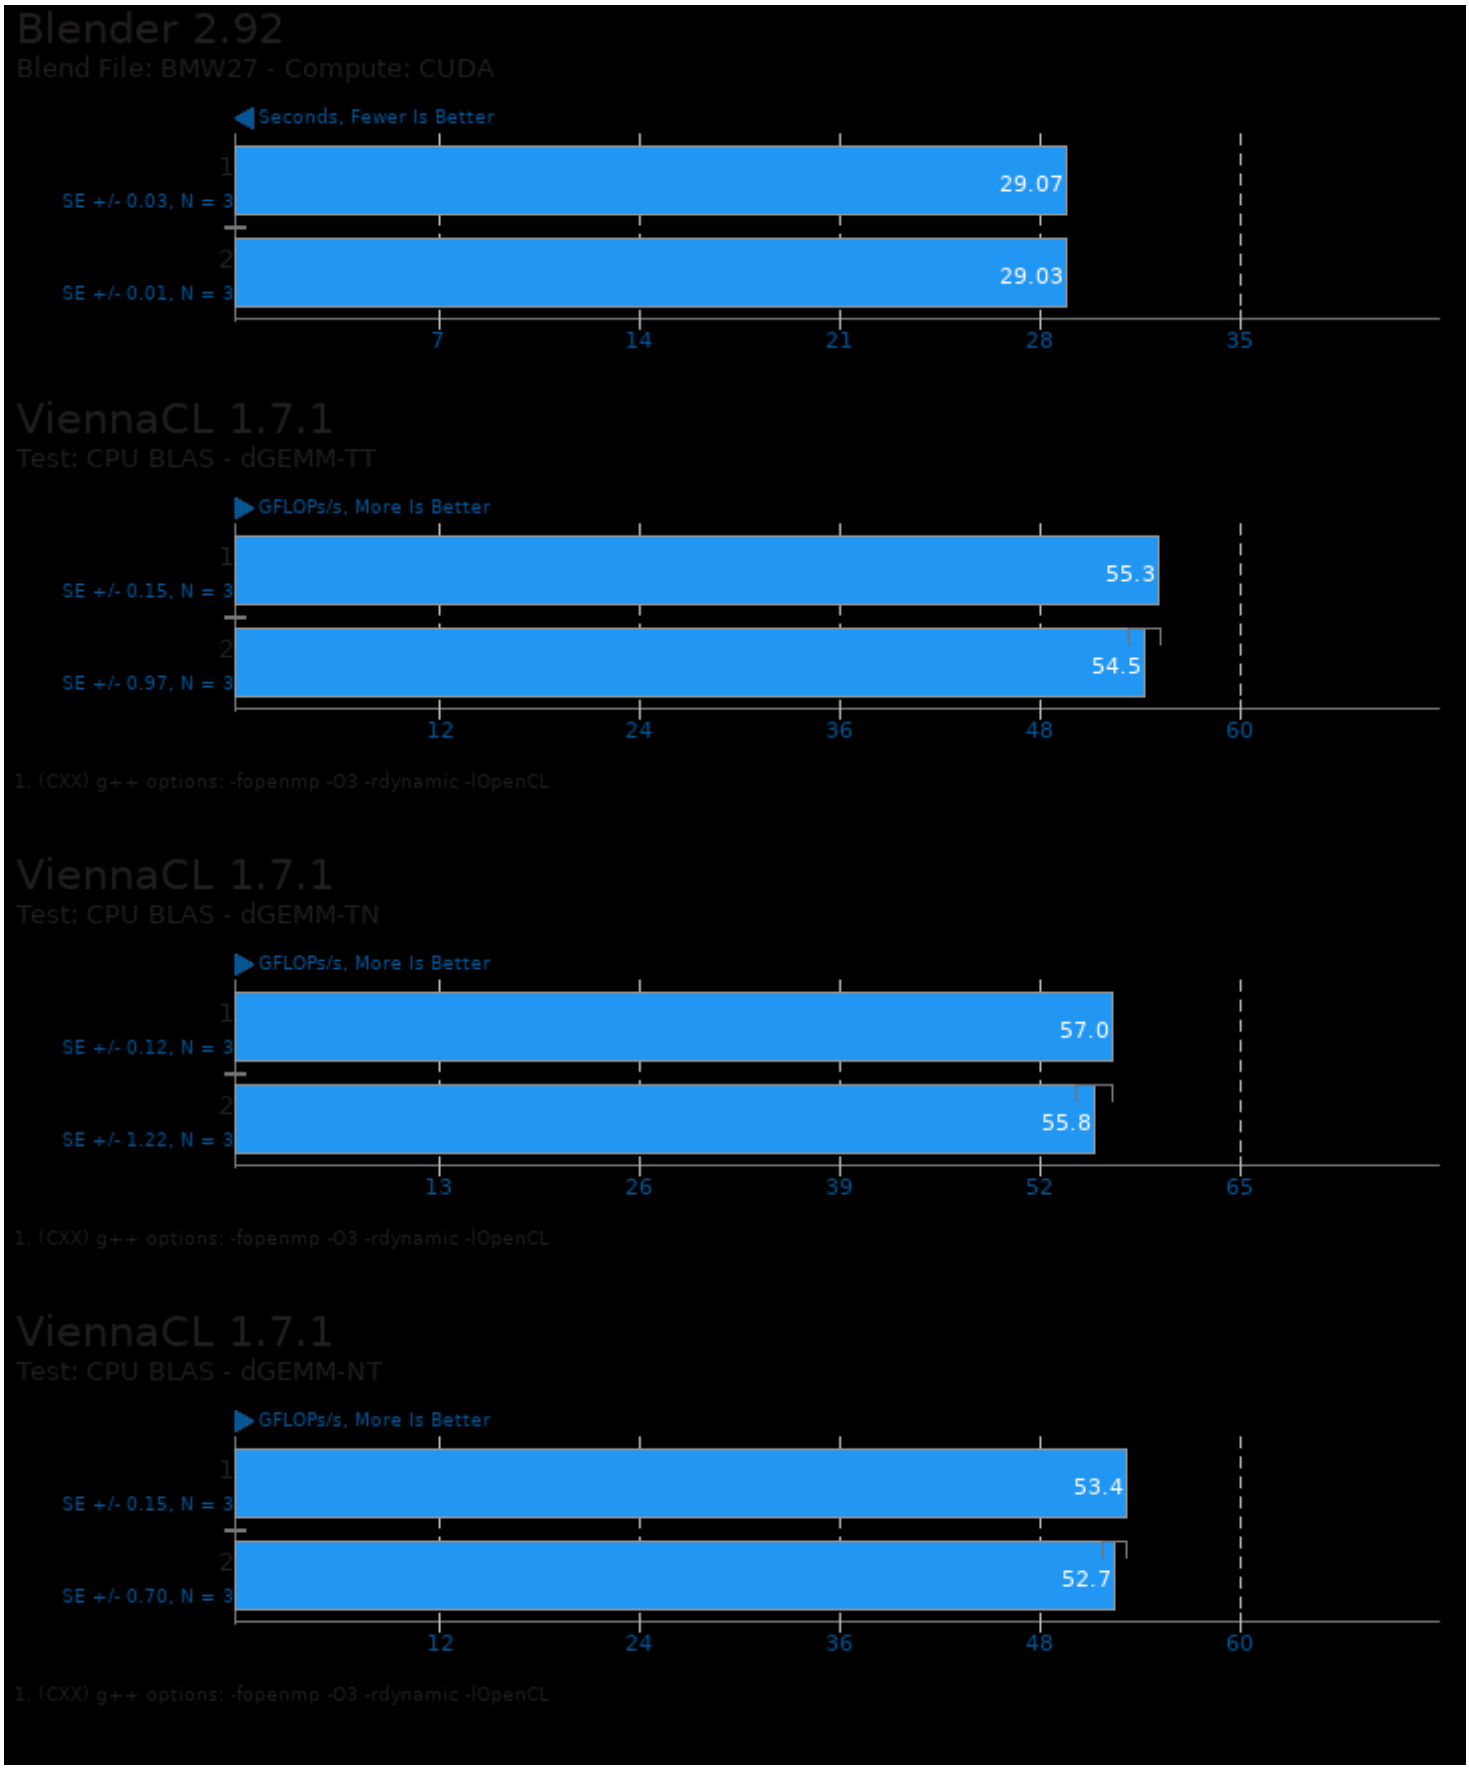

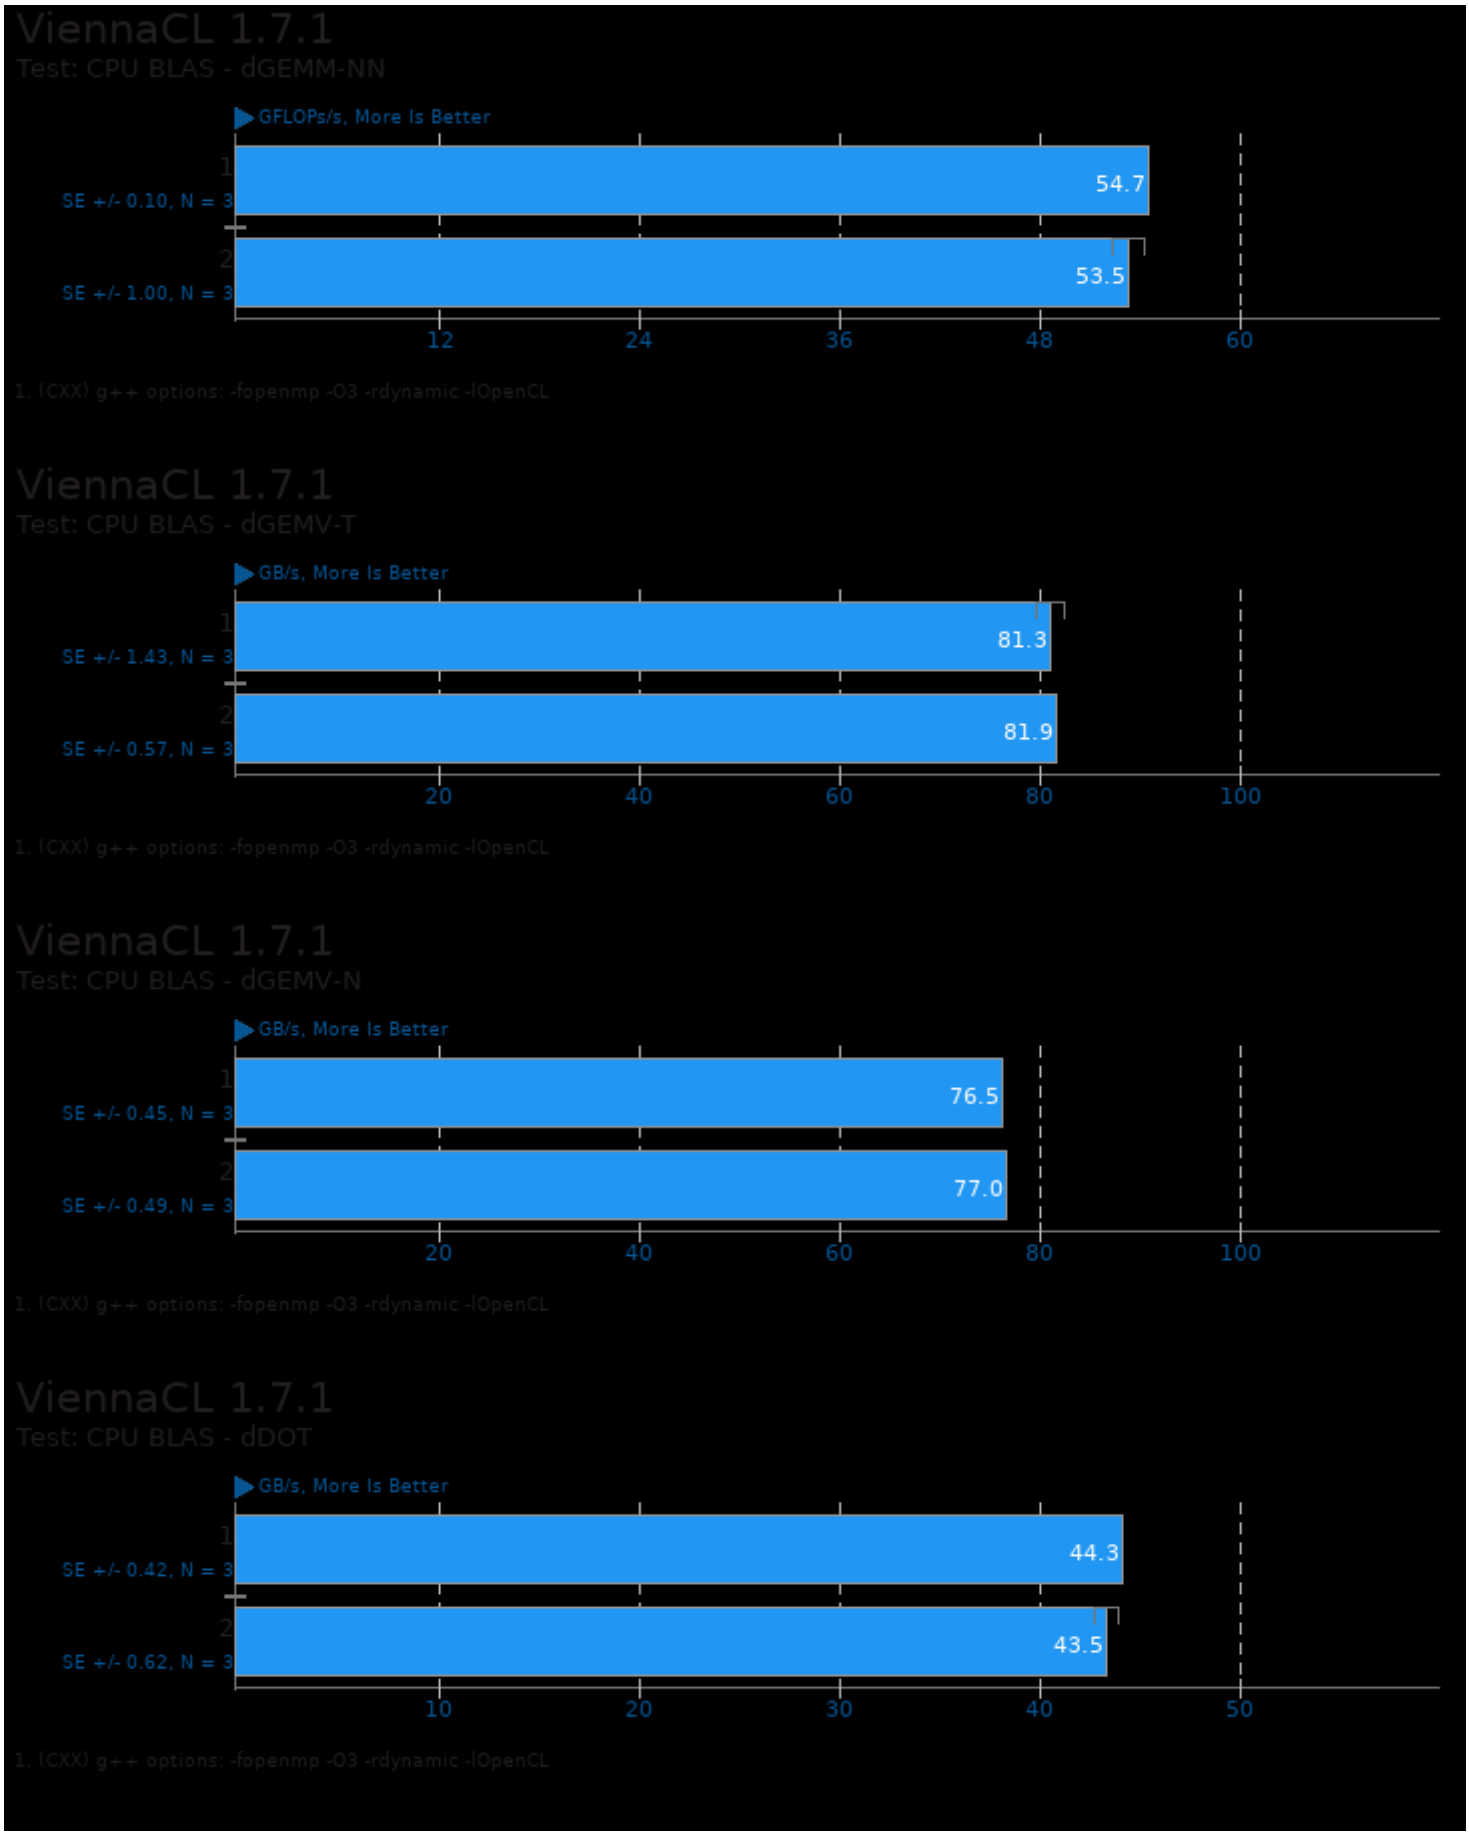
ViennaCL 1.7.1
Test: CPU BLAS - dGEMM-NN
GFLOPs/s, More Is Better
1
SE +/- 0.10, N = 3
54.7
2
SE +/- 1.00, N = 3
53.5
12
24
36
48
60
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dGEMV-T
GB/s, More Is Better
1
SE +/- 1.43, N = 3
81.3
2
SE +/- 0.57, N = 3
81.9
20
40
60
80
100
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dGEMV-N
GB/s, More Is Better
1
SE +/- 0.45, N = 3
76.5
2
SE +/- 0.49, N = 3
77.0
20
40
60
80
100
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dDOT
GB/s, More Is Better
1
SE +/- 0.42, N = 3
44.3
2
SE +/- 0.62, N = 3
43.5
10
20
30
40
50
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

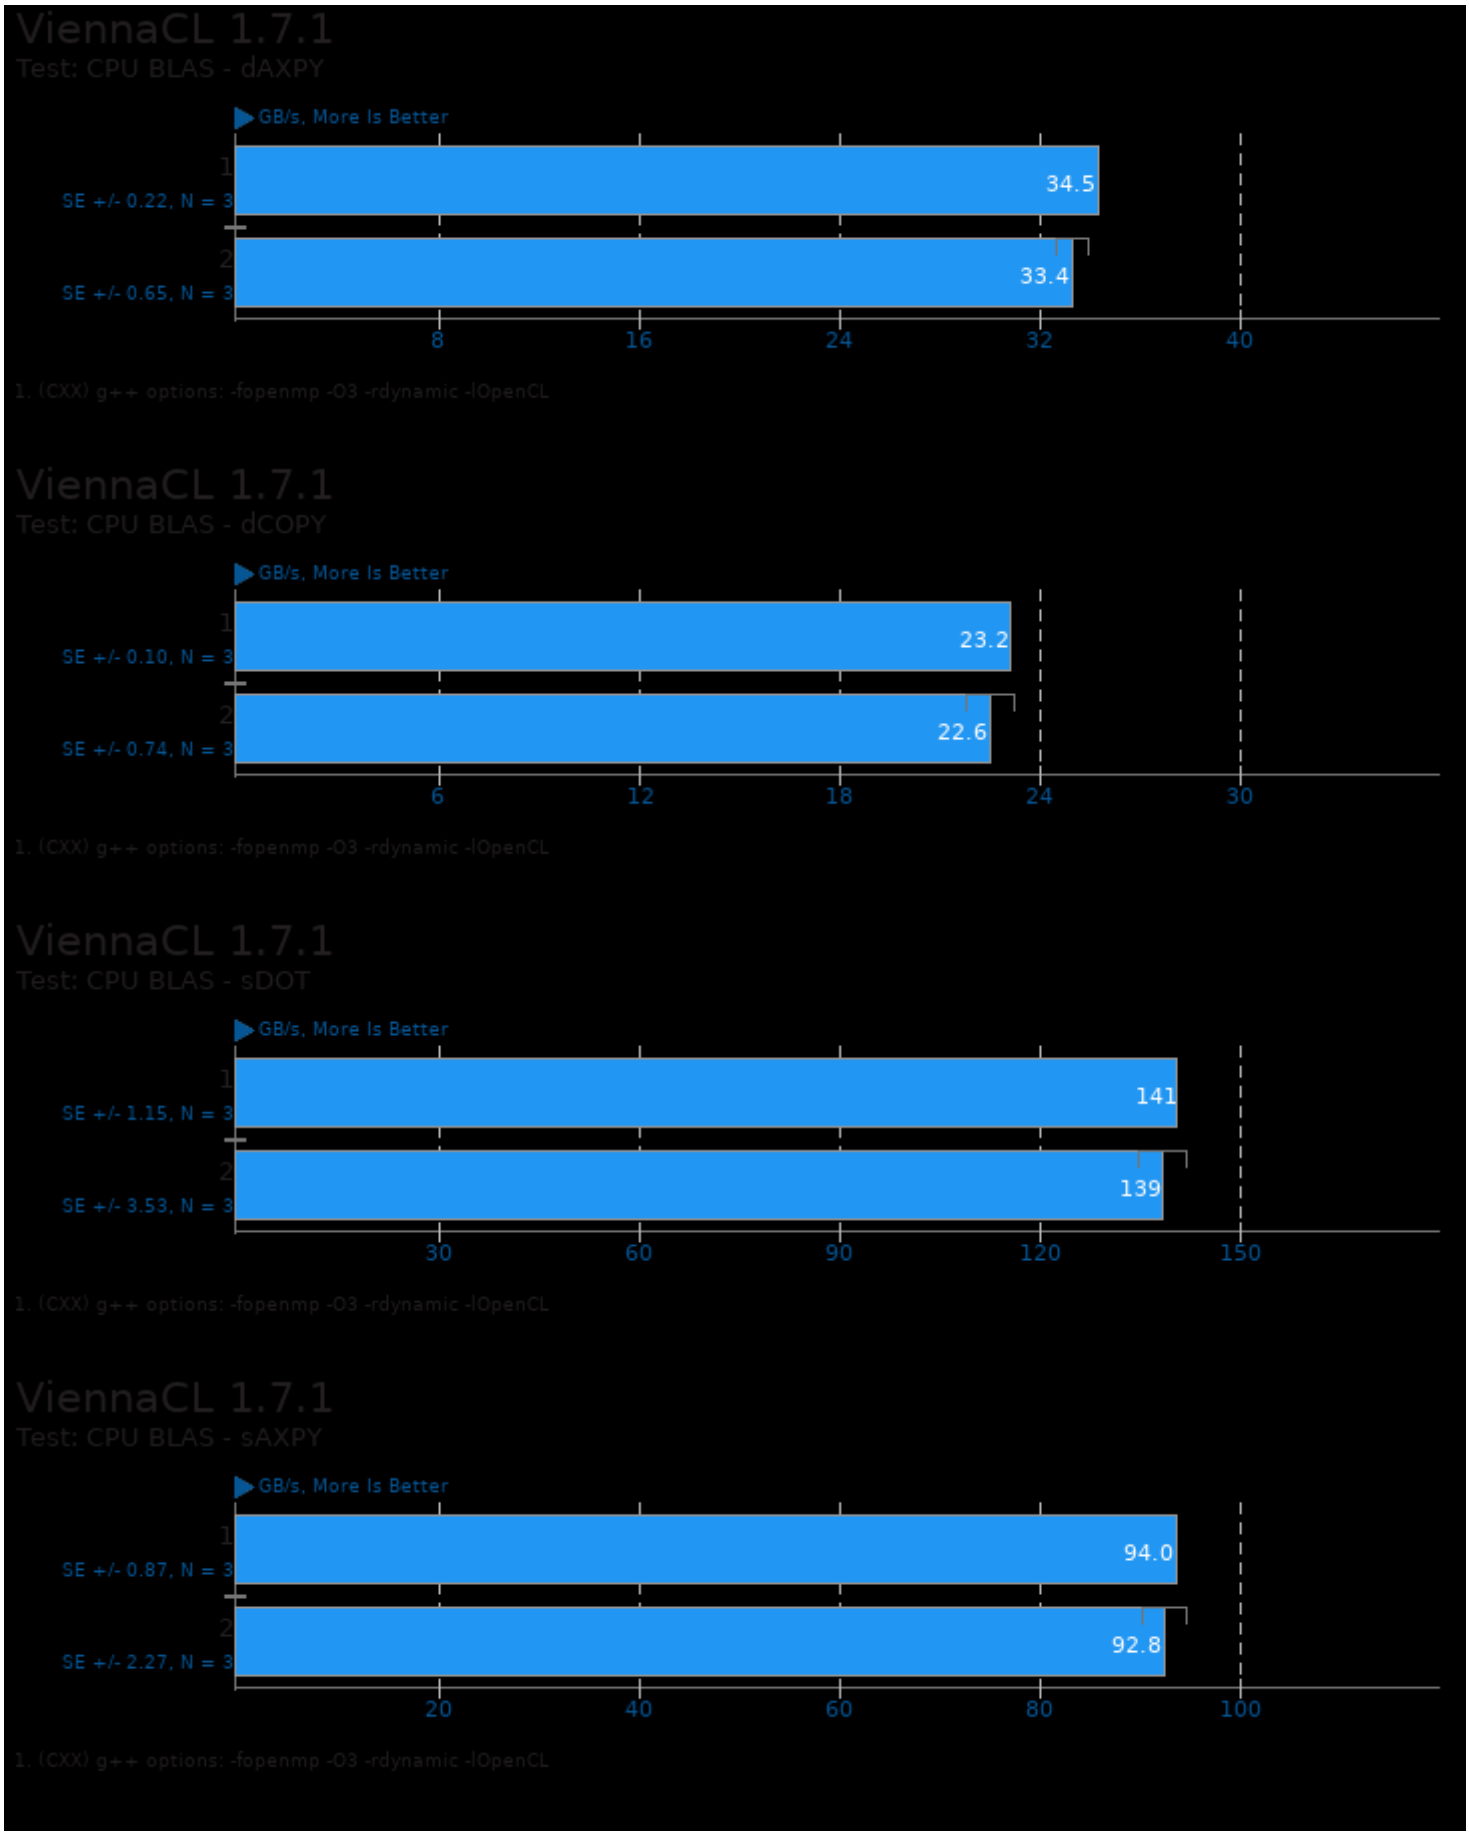
ViennaCL 1.7.1
Test: CPU BLAS - dAXPY
GB/s, More Is Better
1
SE +/- 0.22, N = 3
34.5
2
SE +/- 0.65, N = 3
33.4
8
16
24
32
40
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dCOPY
GB/s, More Is Better
1
SE +/- 0.10, N = 3
23.2
2
SE +/- 0.74, N = 3
22.6
6
12
18
24
30
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - sDOT
GB/s, More Is Better
1
SE +/- 1.15, N = 3
141
2
SE +/- 3.53, N = 3
139
30
60
90
120
150
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - sAXPY
GB/s, More Is Better
1
SE +/- 0.87, N = 3
94.0
2
SE +/- 2.27, N = 3
92.8
20
40
60
80
100
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

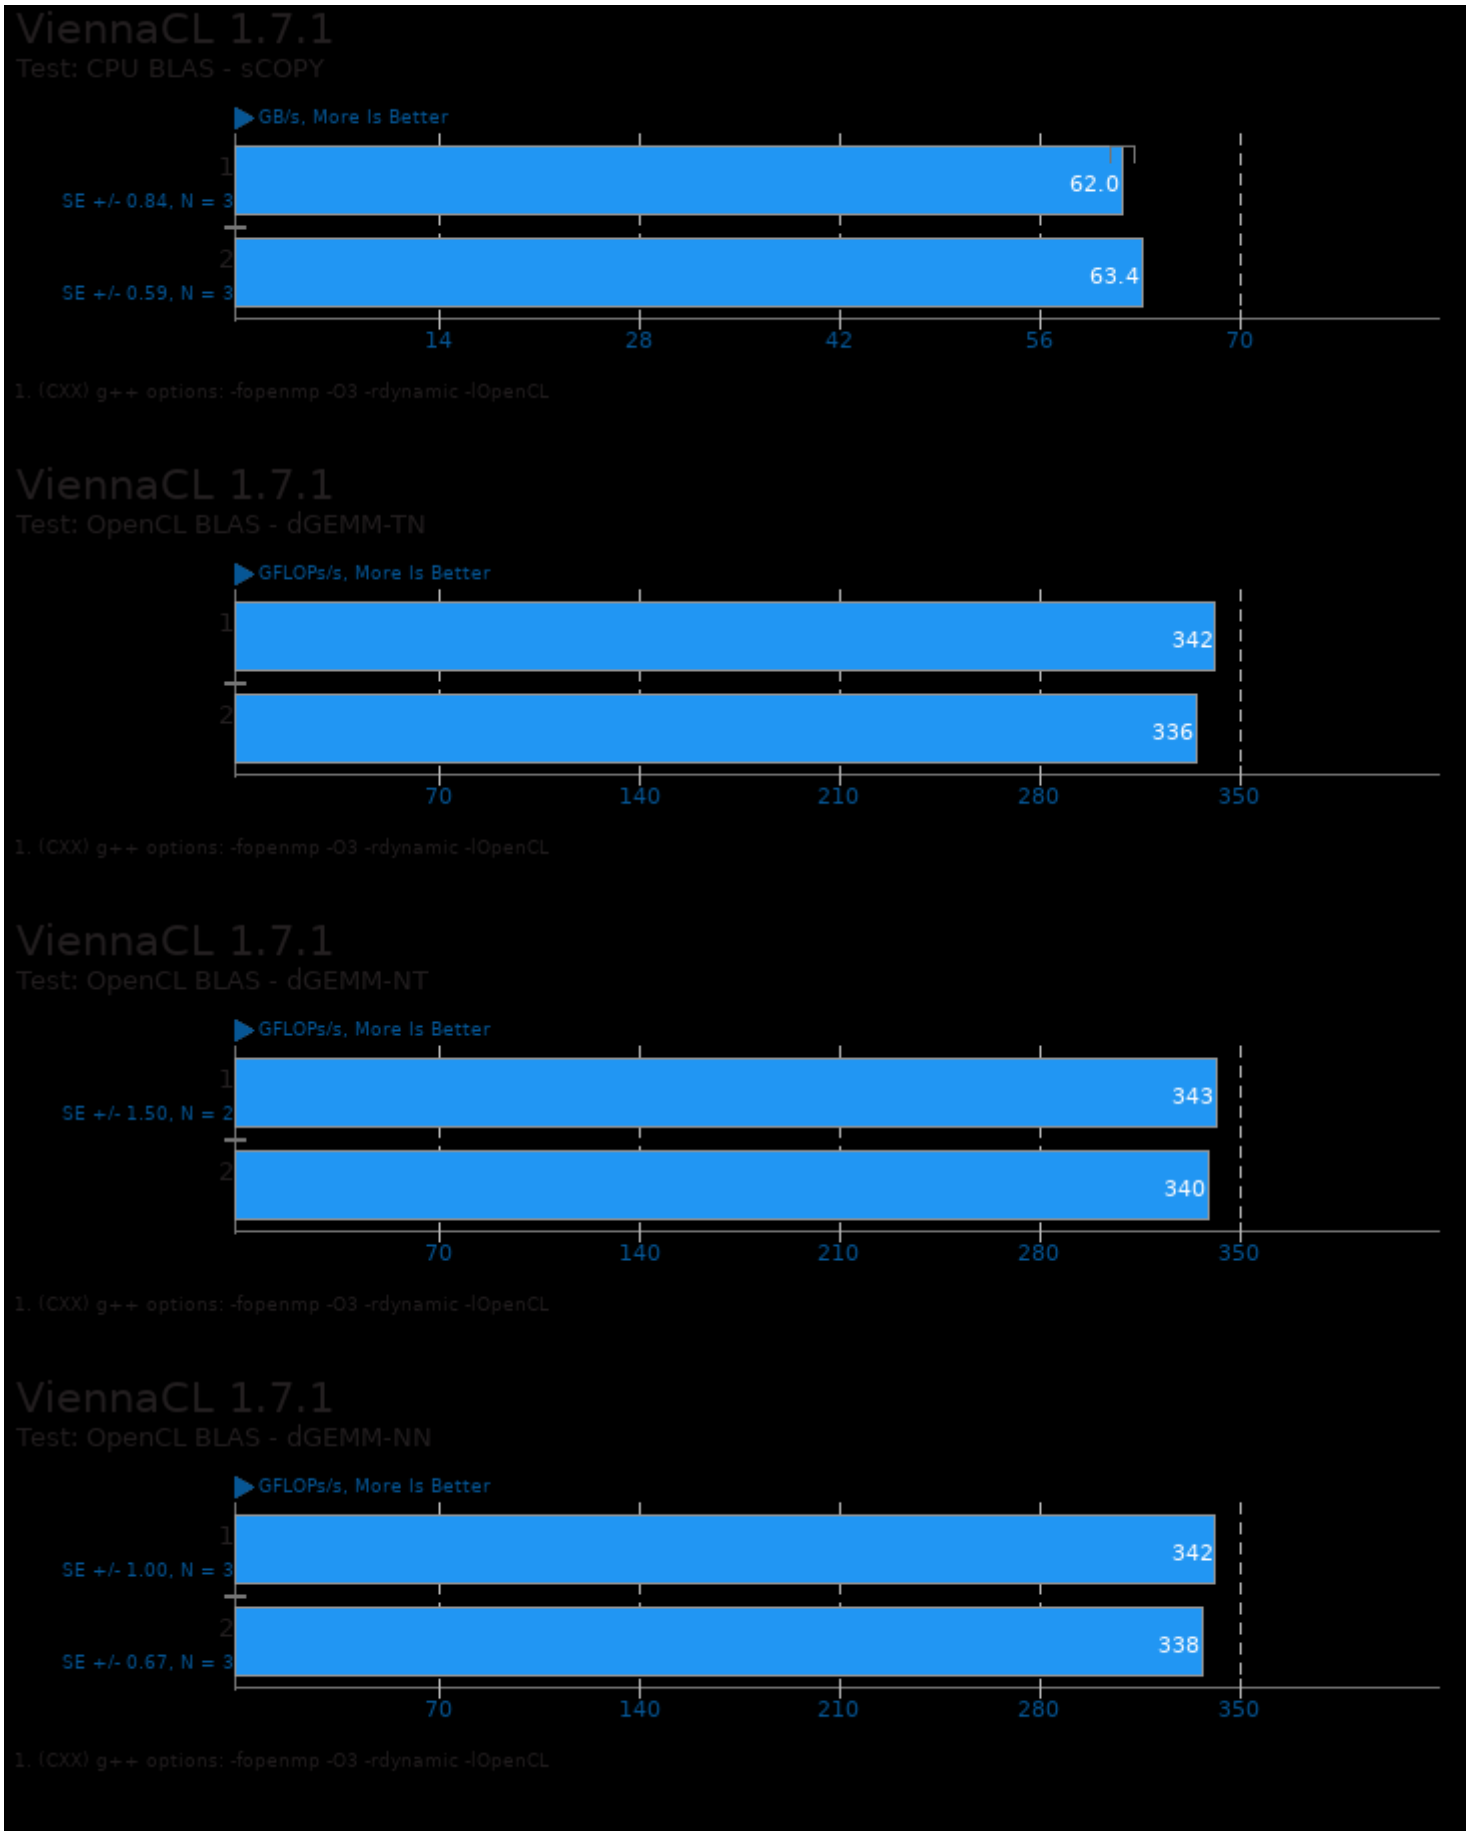
ViennaCL 1.7.1
Test: CPU BLAS - sCOPY
GB/s, More Is Better
1
SE +/- 0.84, N = 3
62.0
2
SE +/- 0.59, N = 3
63.4
14
28
42
56
70
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: OpenCL BLAS - dGEMM-TN
GFLOPs/s, More Is Better
1
342
2
336
70
140
210
280
350
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: OpenCL BLAS - dGEMM-NT
GFLOPs/s, More Is Better
1
SE +/- 1.50, N = 2
343
2
340
70
140
210
280
350
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: OpenCL BLAS - dGEMM-NN
GFLOPs/s, More Is Better
1
SE +/- 1.00, N = 3
342
2
SE +/- 0.67, N = 3
338
70
140
210
280
350
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - dGEMV-T

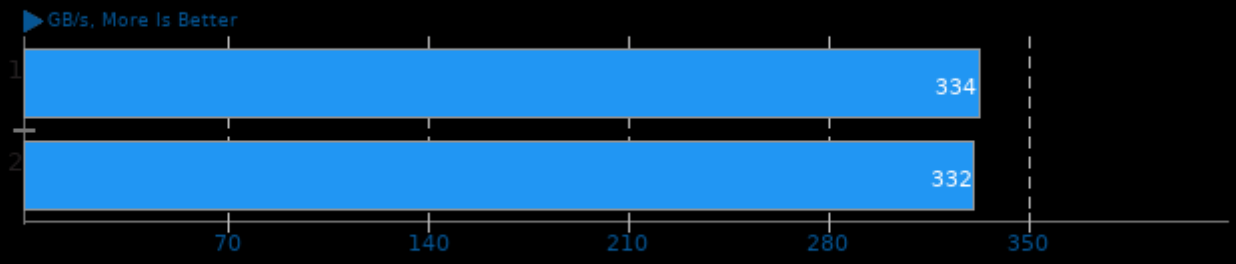

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - dGEMV-N

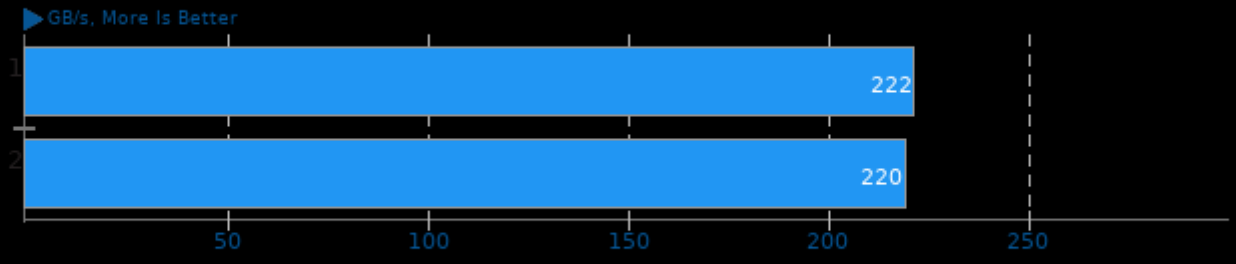

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - dDOT

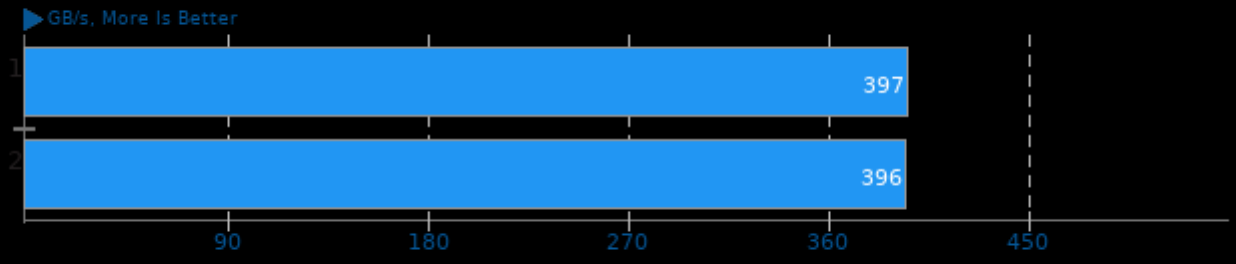

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - dAXPY

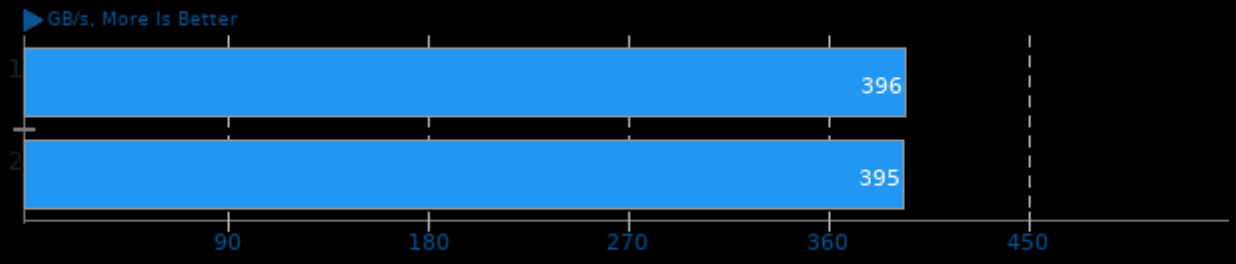

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - dCOPY

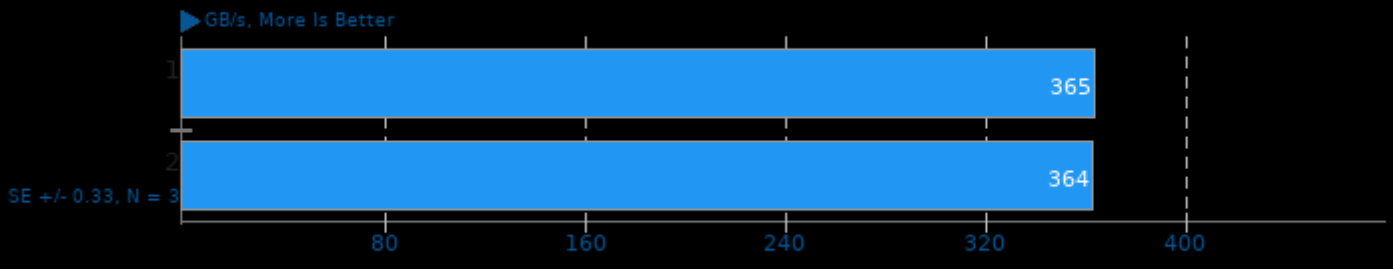

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - sDOT

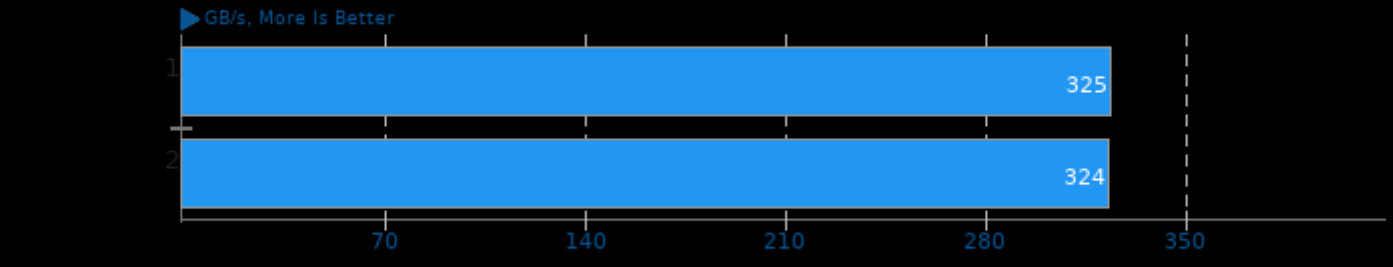

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - sAXPY

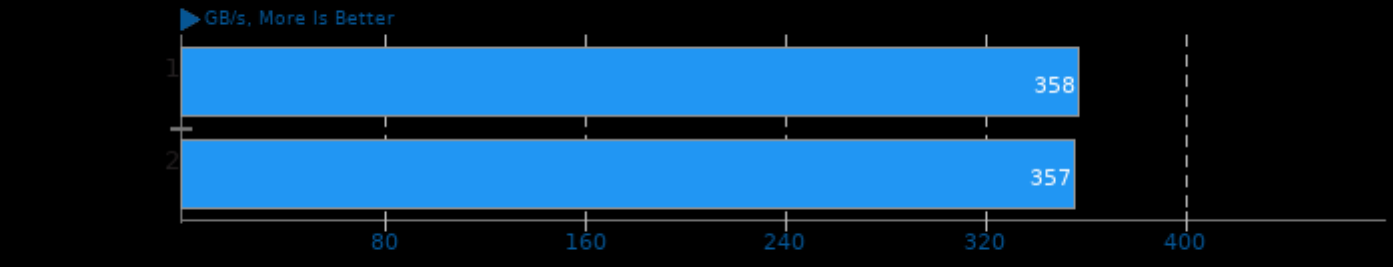

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

ViennaCL 1.7.1

Test: OpenCL BLAS - sCOPY

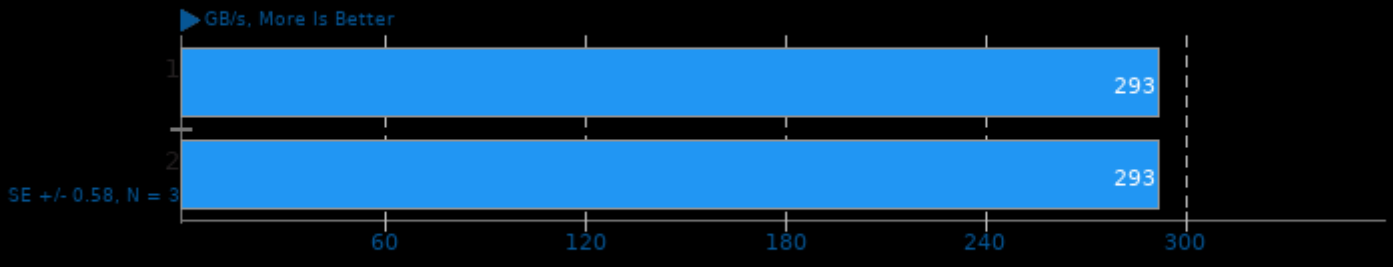

1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

Blender 2.92

Blend File: BMW27 - Compute: NVIDIA OptiX

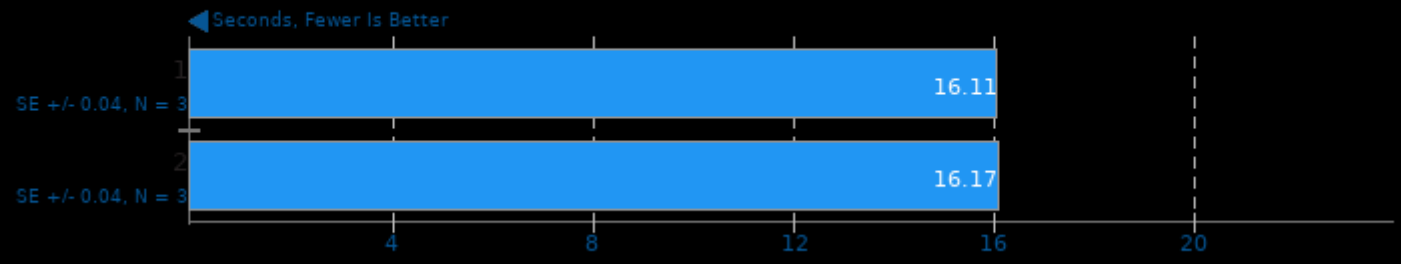

SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Texture Read Bandwidth

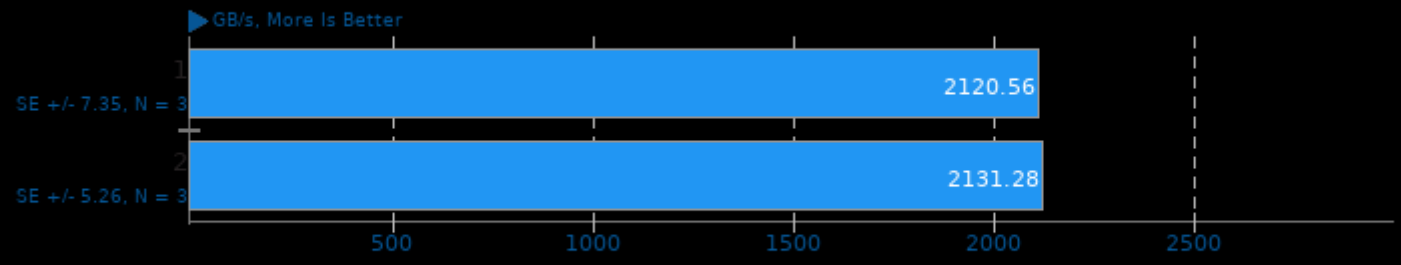

1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

VkResample 1.0

Upscale: 2x - Precision: Single

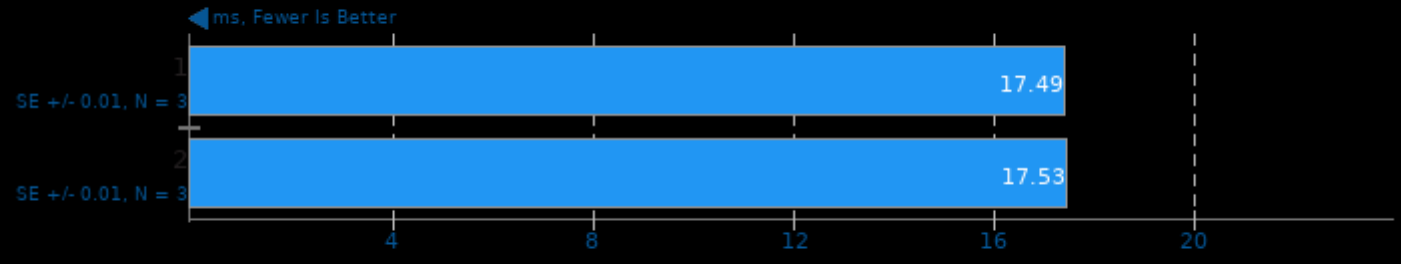

1. (CXX) g++ options: -O3 -pthread

NAMD CUDA 2.14

ATPase Simulation - 327,506 Atoms

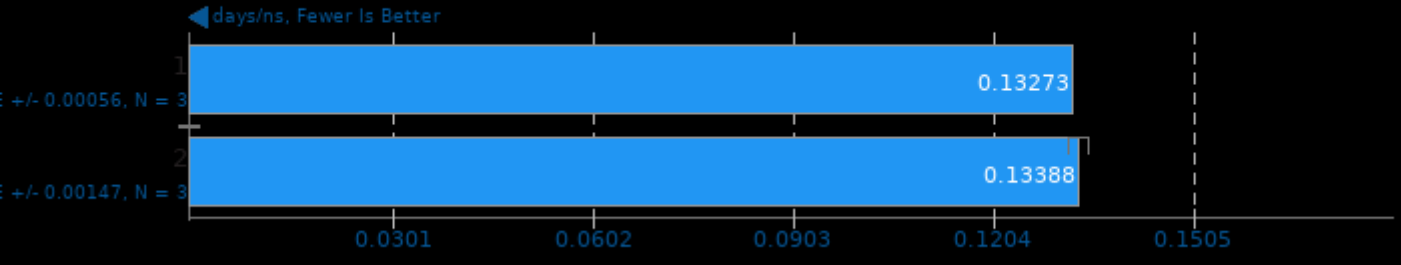

RTX 3070 Compute

LuxCoreRender OpenCL 2.3

Scene: Rainbow Colors and Prism

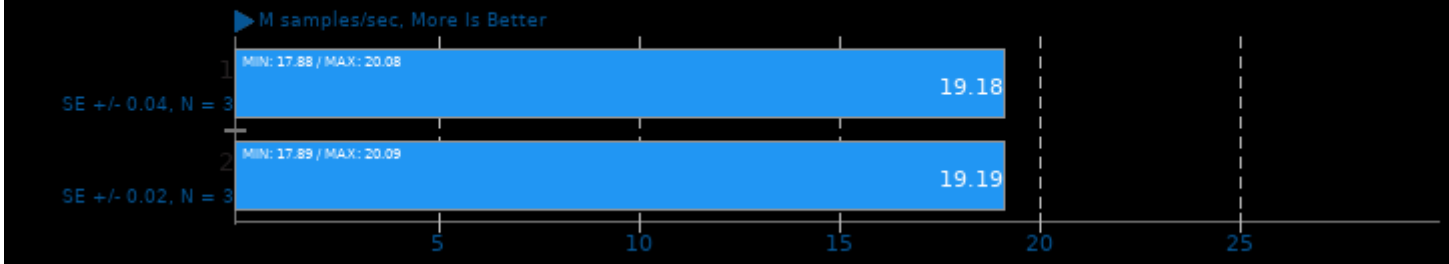

RealSR-NCNN 20200818

Scale: 4x - TAA: No

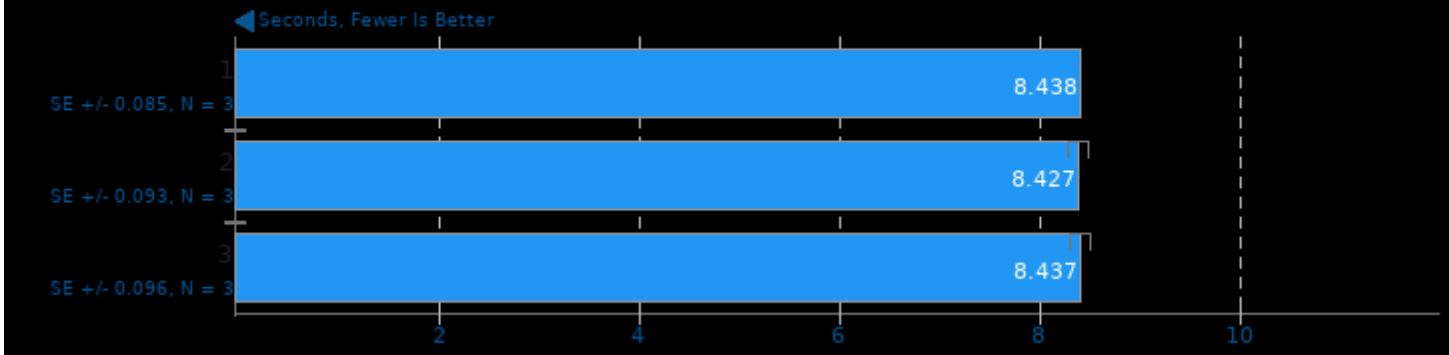

Betsy GPU Compressor 1.1 Beta

Codec: ETC2 RGB - Quality: Highest

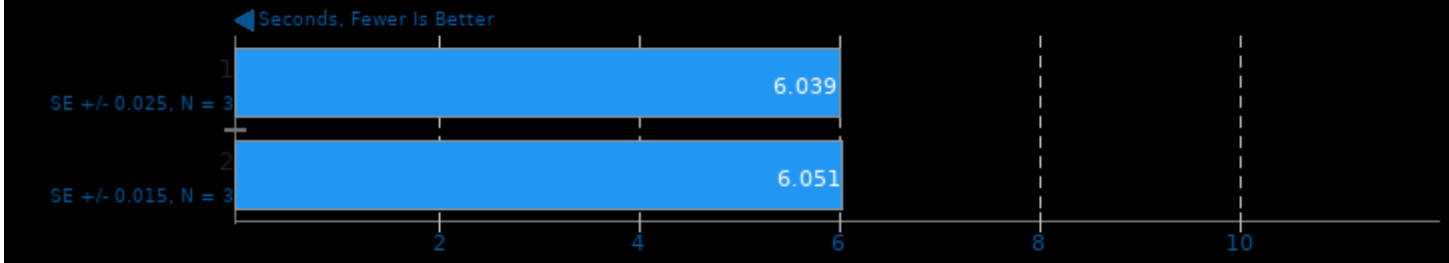

1. (CXX) g++ options: -O3 -O2 -lpthread -ldl

Rodinia 3.1

Test: OpenCL Particle Filter

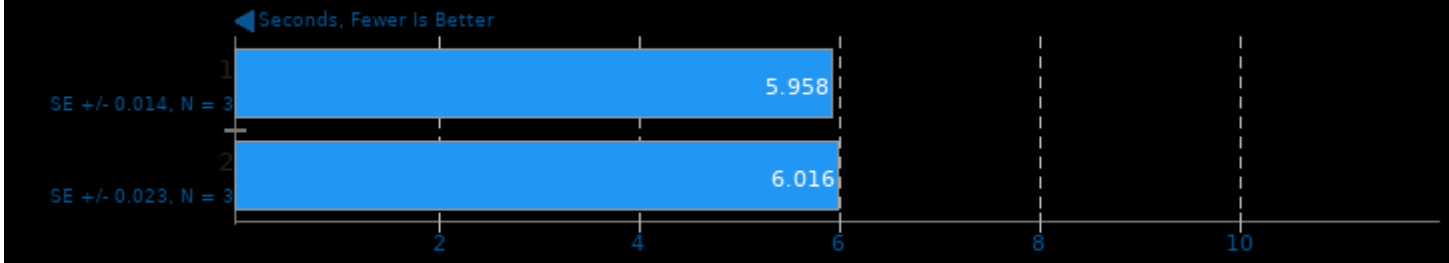

1. (CXX) g++ options: -m64 -lm -lcuda -lcudart -lcudadevrt -lcudart_static -lrt -lpthread -ldl

RTX 3070 Compute

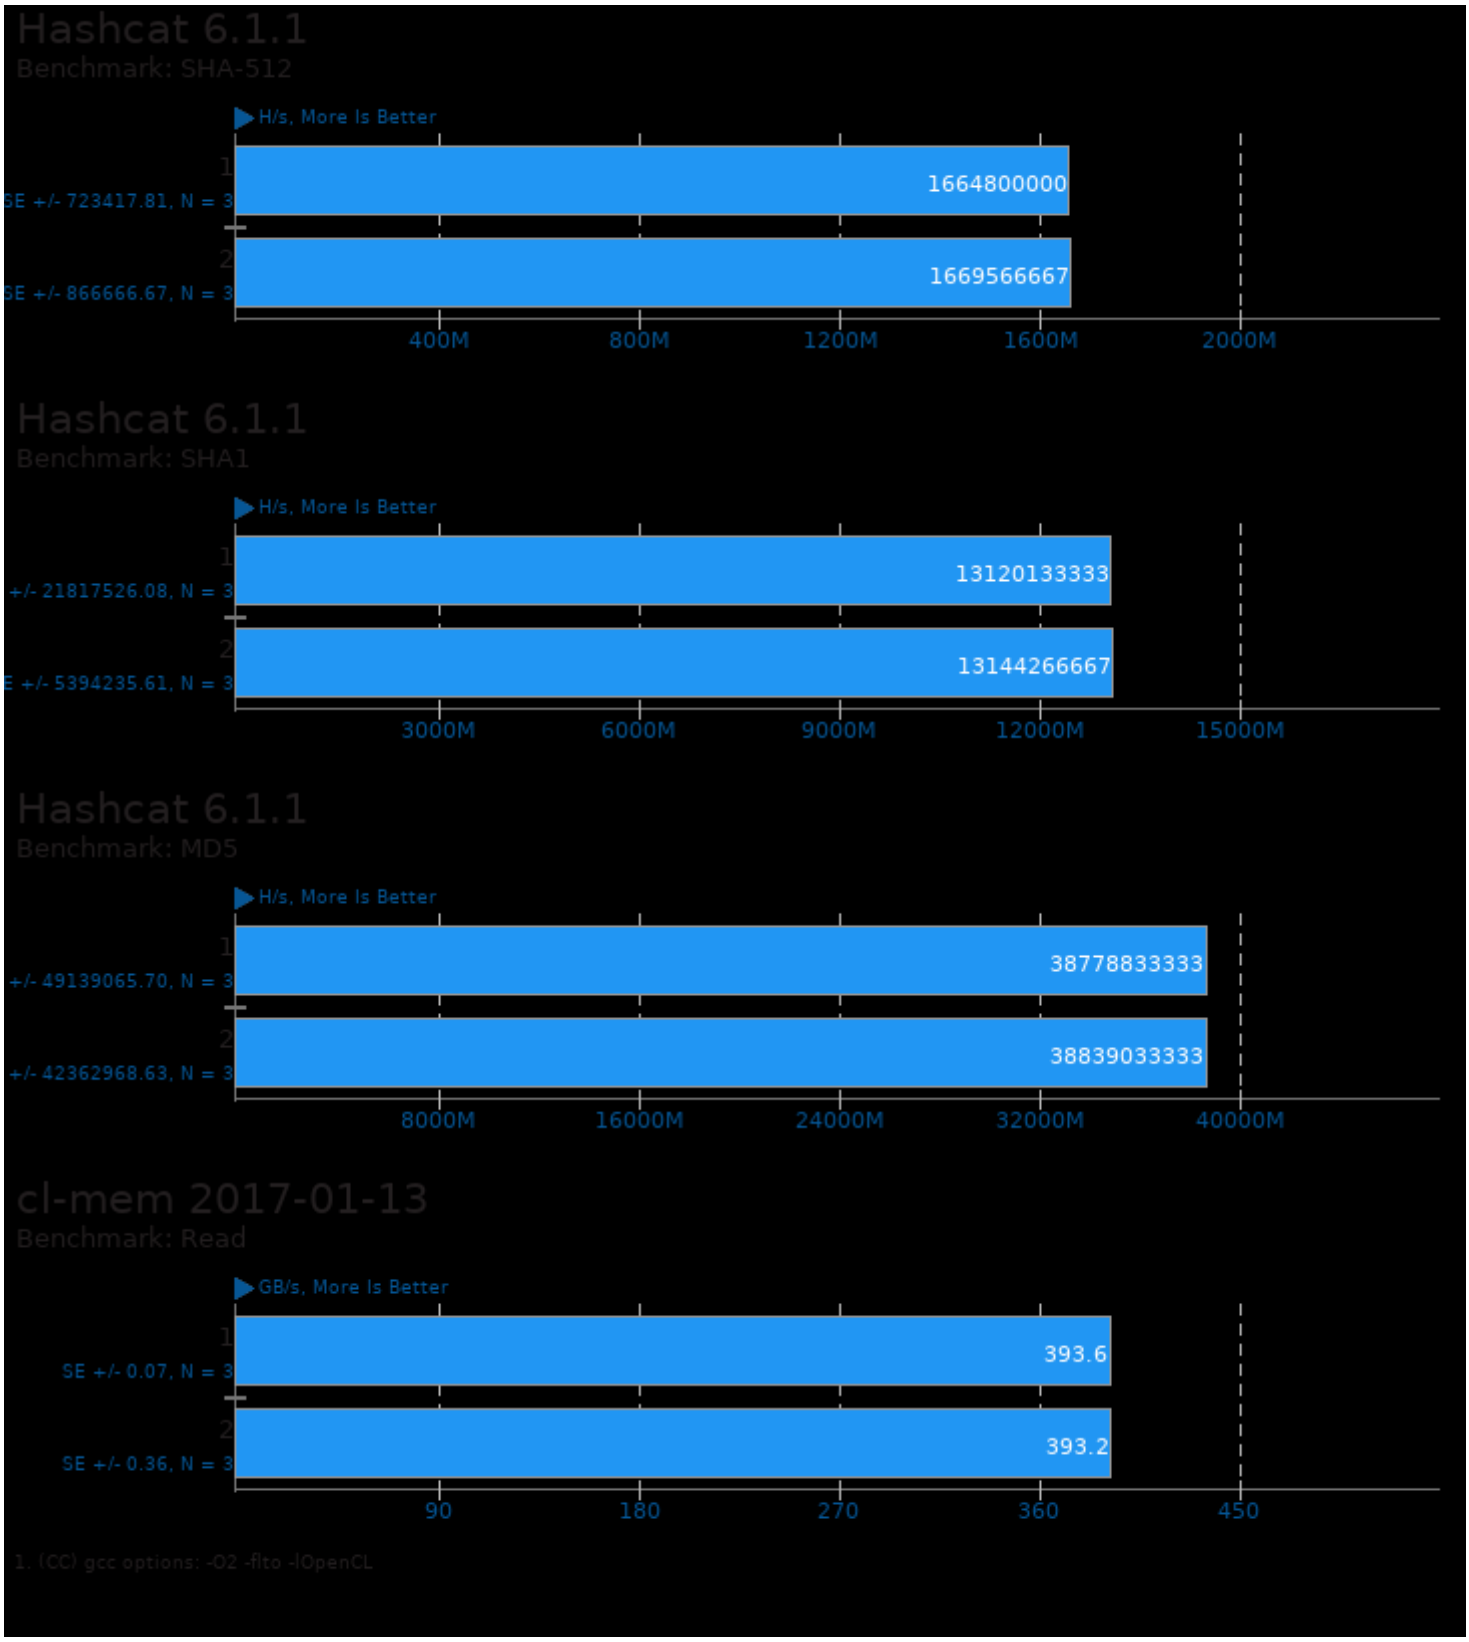

cl-mem 2017-01-13

Benchmark: Write

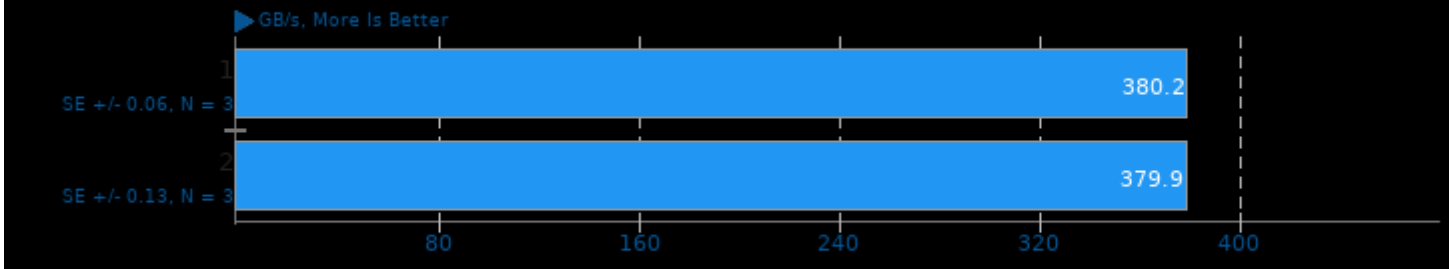

1. (CC) gcc options: -O2 -flto -lOpenCL

cl-mem 2017-01-13

Benchmark: Copy

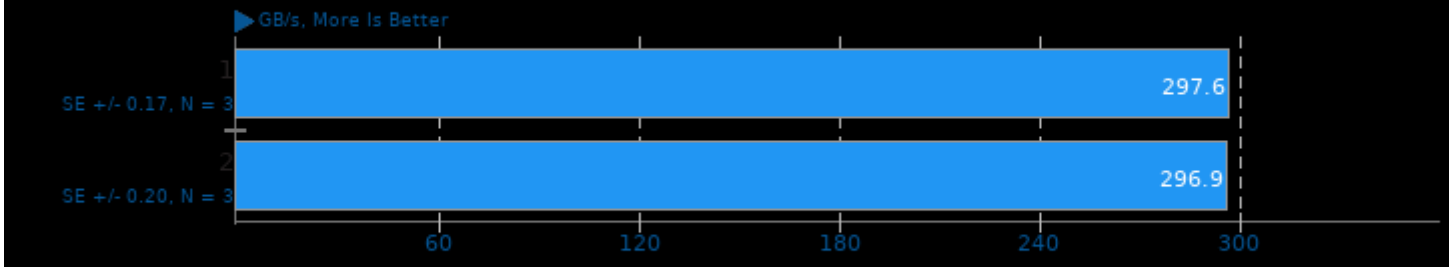

1. (CC) gcc options: -O2 -flto -lOpenCL

Waifu2x-NCNN Vulkan 20200818

Scale: 2x - Denoise: 3 - TAA: Yes

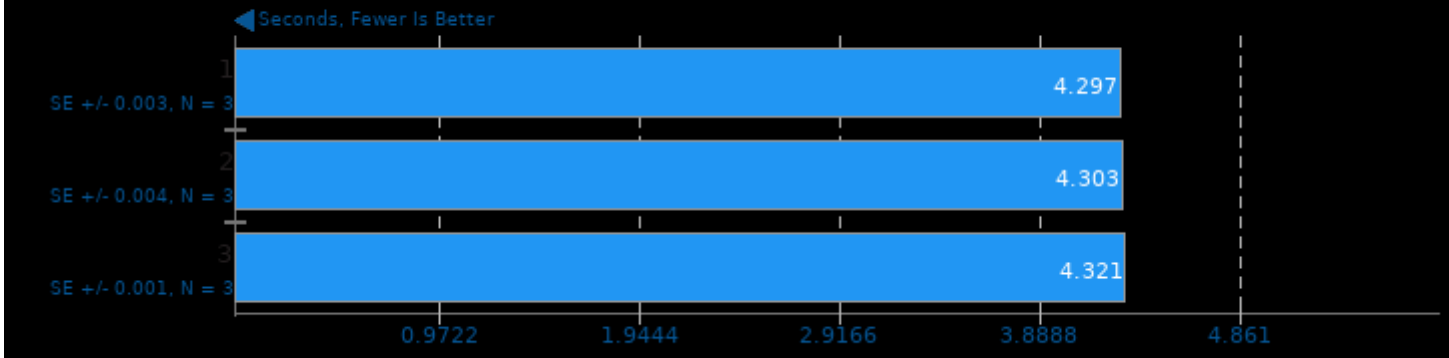

Betsy GPU Compressor 1.1 Beta

Codec: ETC1 - Quality: Highest

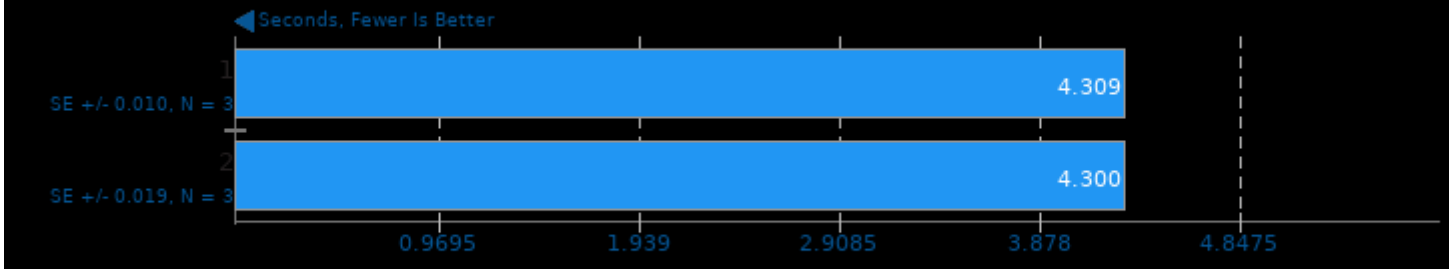

1. (CXX) g++ options: -O3 -O2 -lpthread -ldl

RTX 3070 Compute

Hashcat 6.1.1

Benchmark: TrueCrypt RIPEMD160 + XTS

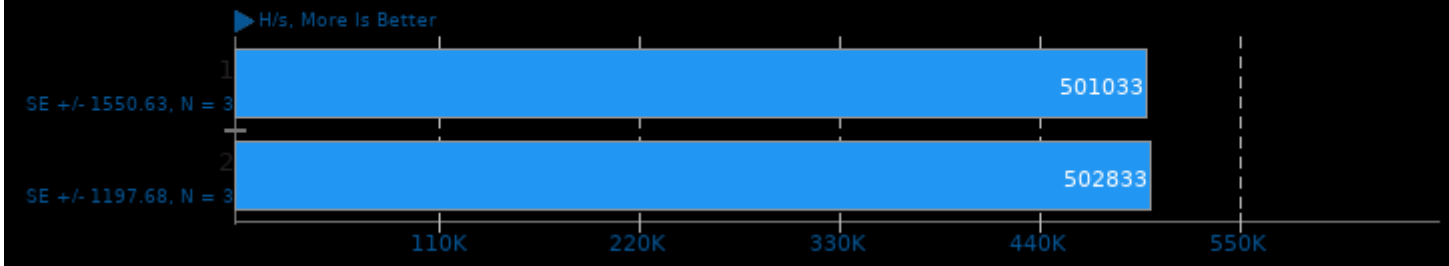

MandelGPU 1.3pts1

OpenCL Device: GPU

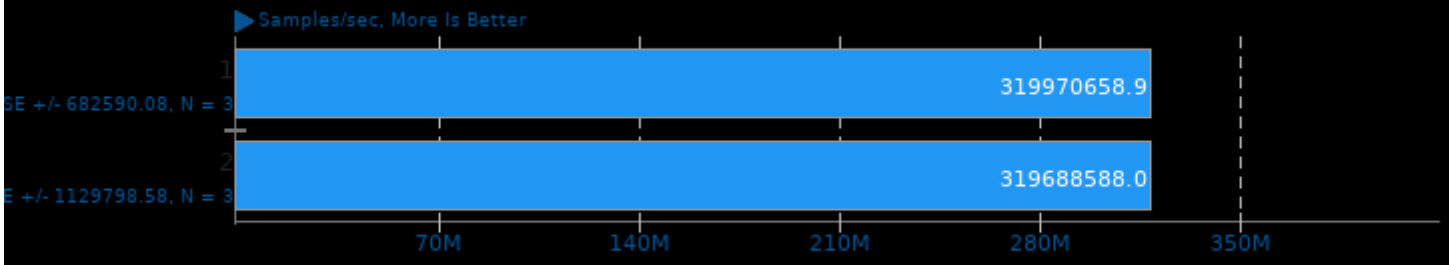

1. (CC) gcc options: -O3 -lm -ftree-vectorize -funroll-loops -lglut -lOpenCL -lGL

SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: GEMM SGEMM_N

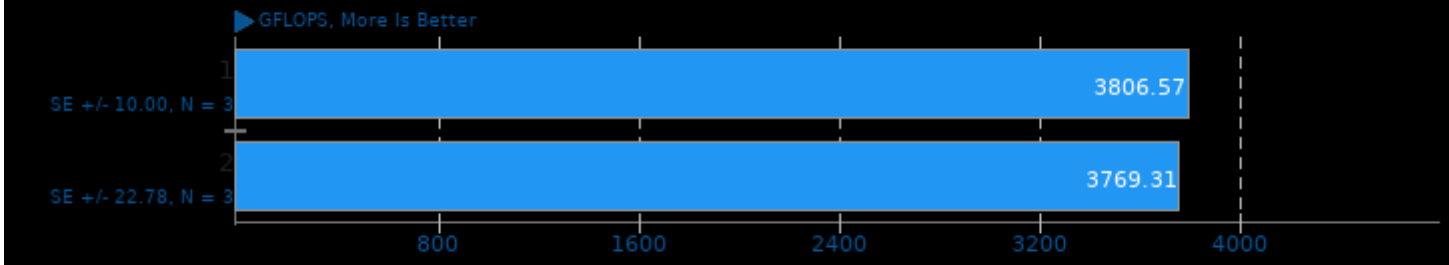

1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

ArrayFire 3.7

Test: Conjugate Gradient OpenCL

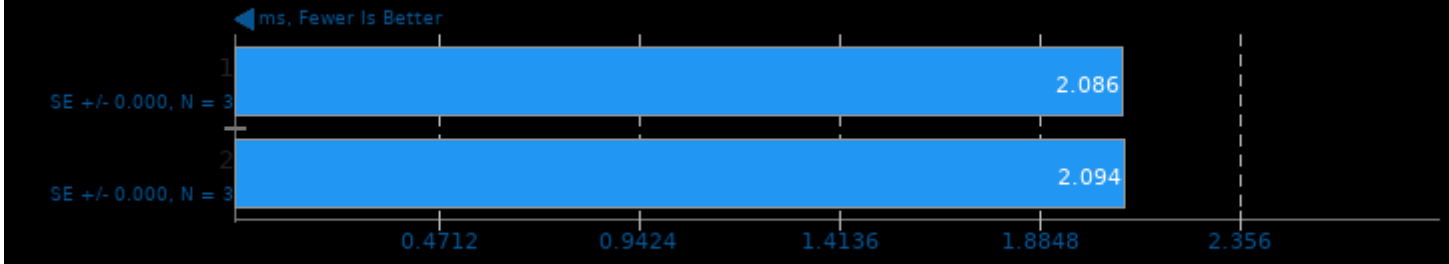

1. (CXX) g++ options: -rdynamic

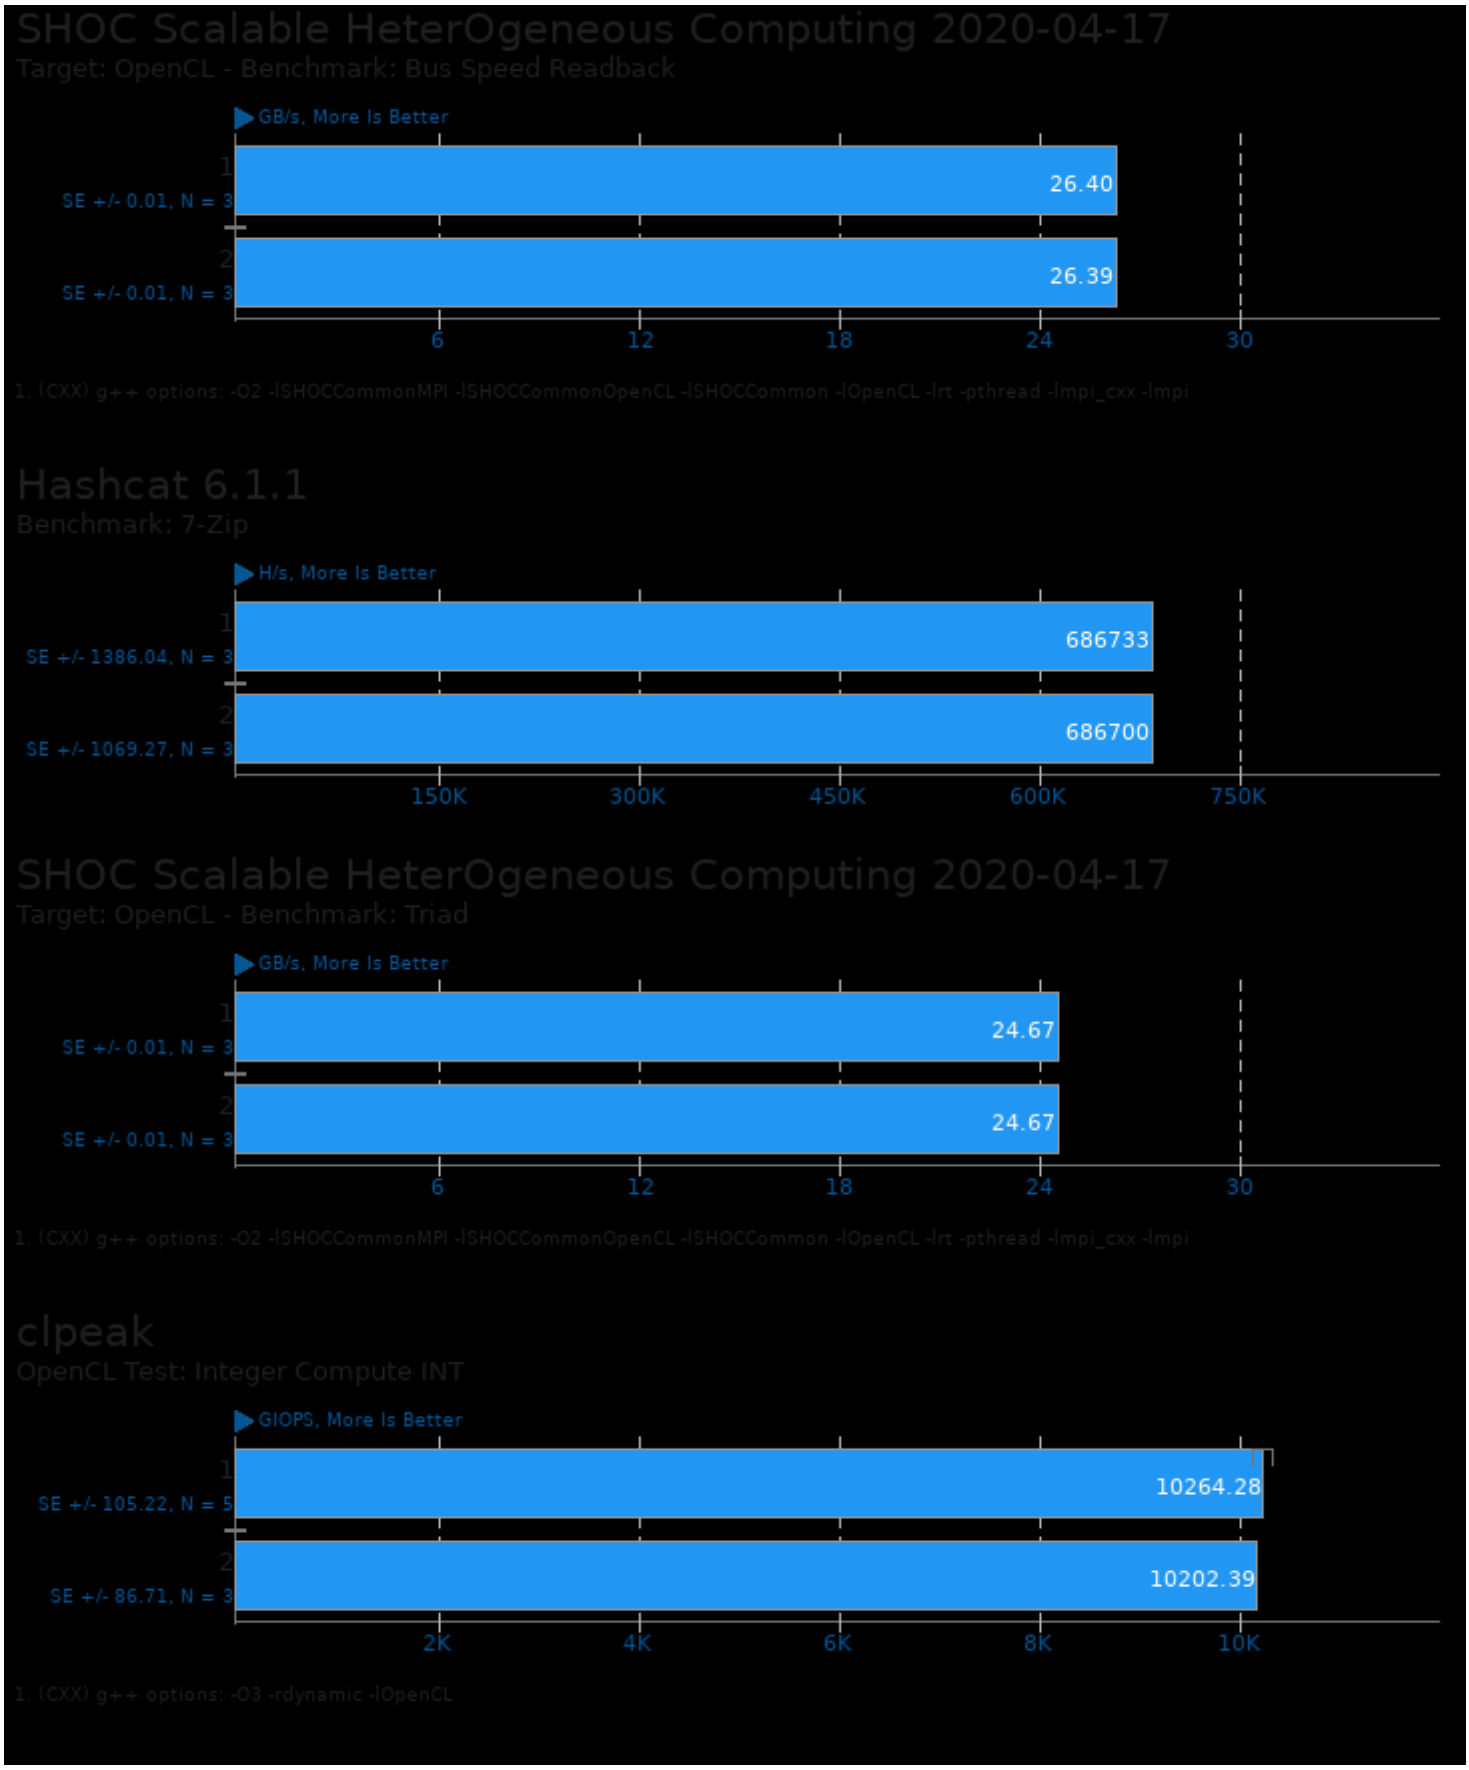
SHOC Scalable HeterOgeneous Computing 2020-04-17
Target: OpenCL - Benchmark: Bus Speed Readback
GB/s, More Is Better
1
SE +/- 0.01, N = 3
26.40
2
SE +/- 0.01, N = 3
26.39
6
12
18
24
30
1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi
Hashcat 6.1.1
Benchmark: 7-Zip
H/s, More Is Better
1
SE +/- 1386.04, N = 3
686733
2
SE +/- 1069.27, N = 3
686700
150K
300K
450K
600K
750K
SHOC Scalable HeterOgeneous Computing 2020-04-17
Target: OpenCL - Benchmark: Triad
GB/s, More Is Better
1
SE +/- 0.01, N = 3
24.67
2
SE +/- 0.01, N = 3
24.67
6
12
18
24
30
1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi
clpeak
OpenCL Test: Integer Compute INT
GIOPS, More Is Better
1
SE +/- 105.22, N = 5
10264.28
2
SE +/- 86.71, N = 3
10202.39
2K
4K
6K
8K
10K
1. (CXX) g++ options: -O3 -rdynamic -lOpenCL

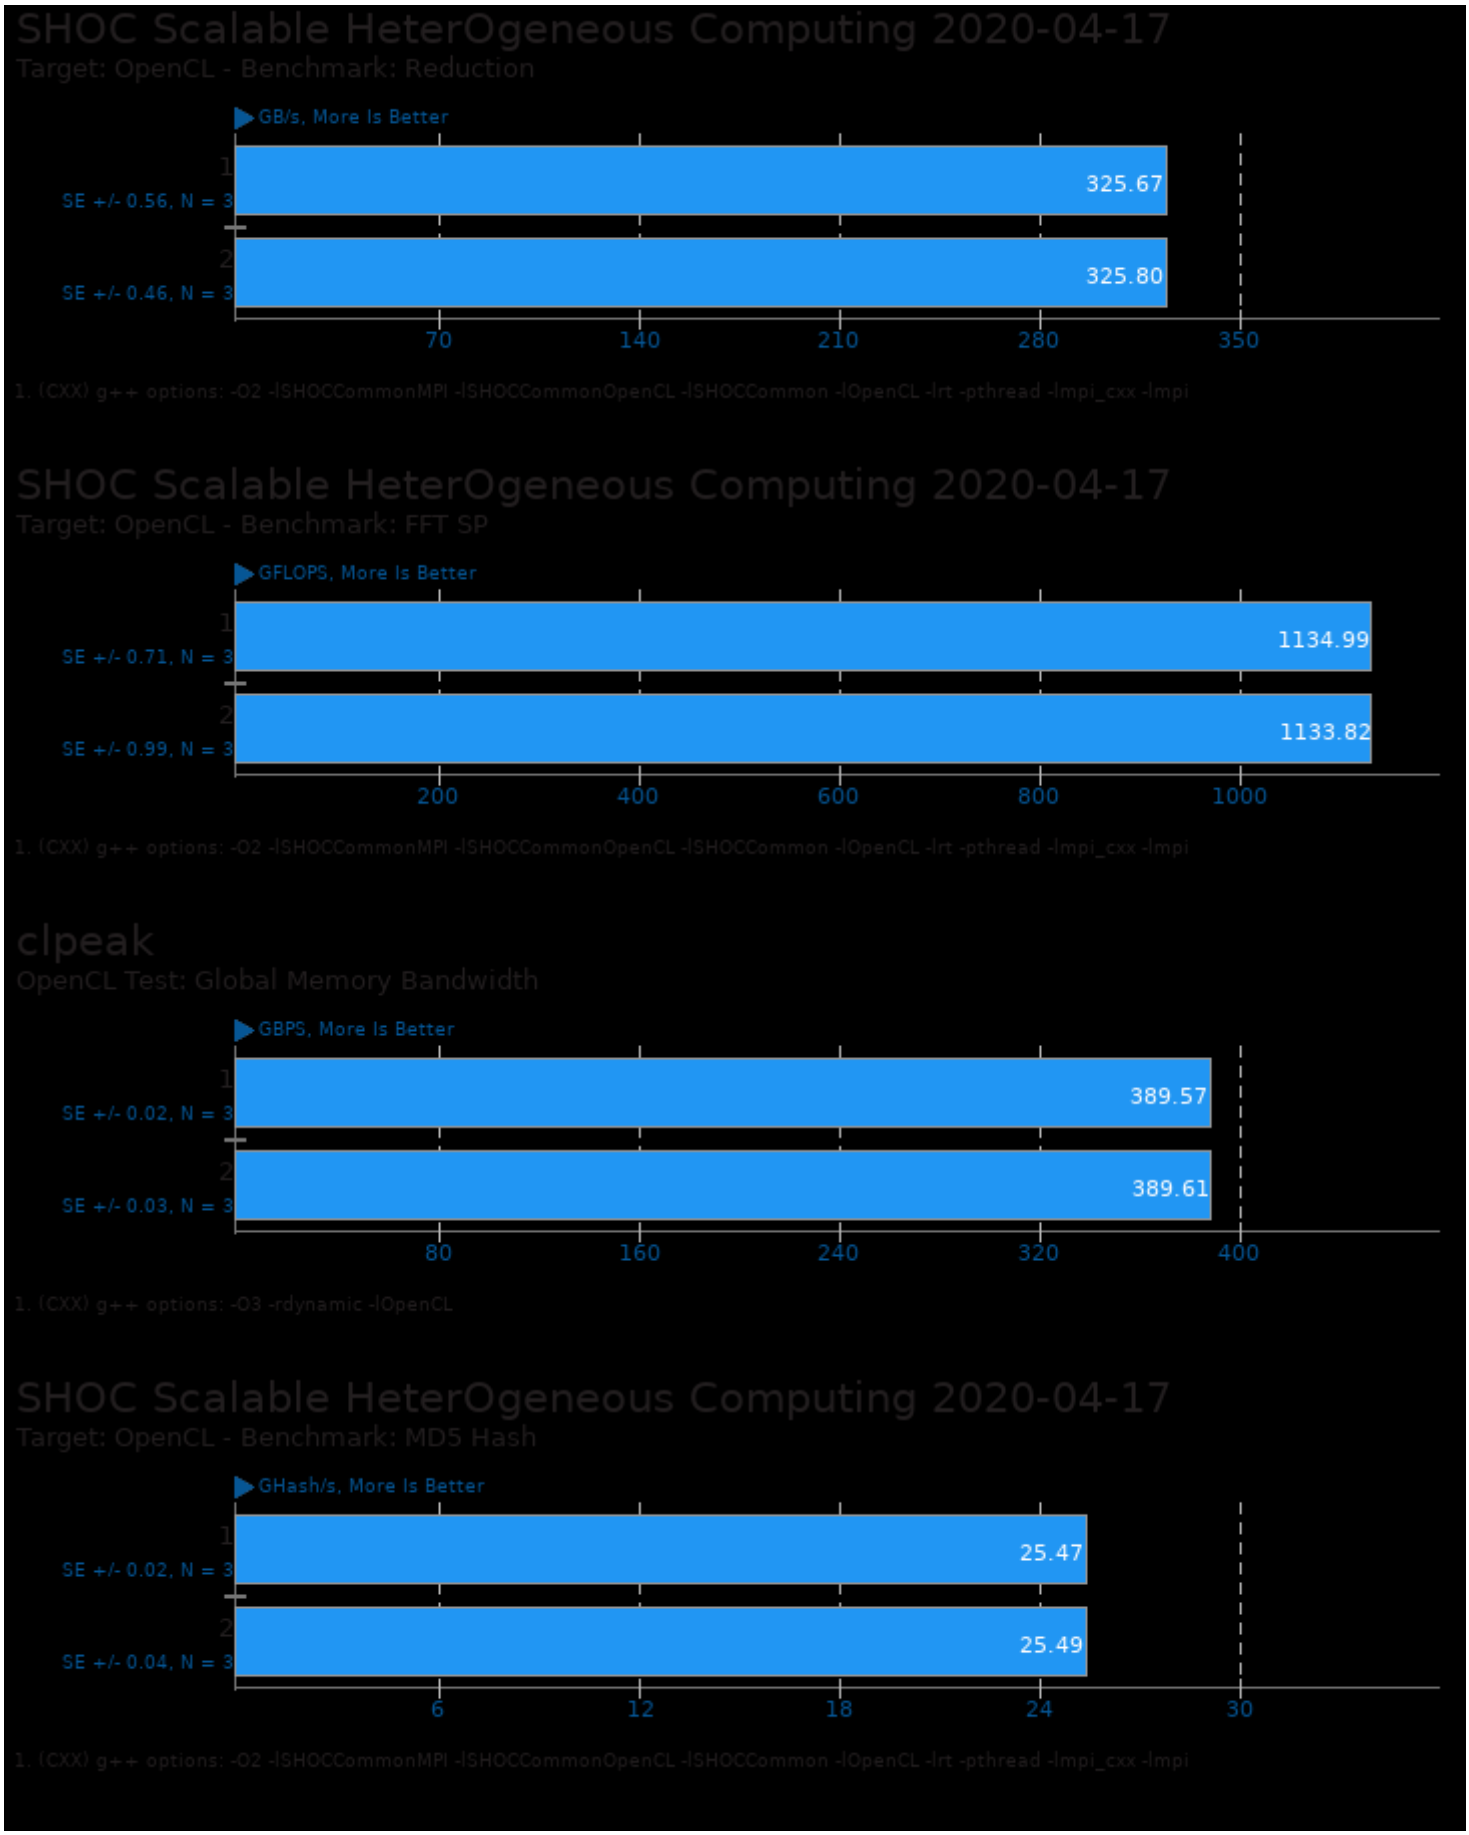
SHOC Scalable HeterOgeneous Computing 2020-04-17
Target: OpenCL - Benchmark: Reduction
GB/s, More Is Better
1
SE +/- 0.56, N = 3
325.67
2
SE +/- 0.46, N = 3
325.80
70
140
210
280
350
1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi
SHOC Scalable HeterOgeneous Computing 2020-04-17
Target: OpenCL - Benchmark: FFT SP
GFLOPS, More Is Better
1
SE +/- 0.71, N = 3
1134.99
2
SE +/- 0.99, N = 3
1133.82
200
400
600
800
1000
1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi
clpeak
OpenCL Test: Global Memory Bandwidth
GBPS, More Is Better
1
SE +/- 0.02, N = 3
389.57
2
SE +/- 0.03, N = 3
389.61
80
160
240
320
400
1. (CXX) g++ options: -O3 -rdynamic -lOpenCL
SHOC Scalable HeterOgeneous Computing 2020-04-17
Target: OpenCL - Benchmark: MD5 Hash
GHash/s, More Is Better
1
SE +/- 0.02, N = 3
25.47
2
SE +/- 0.04, N = 3
25.49
6
12
18
24
30
1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Bus Speed Download

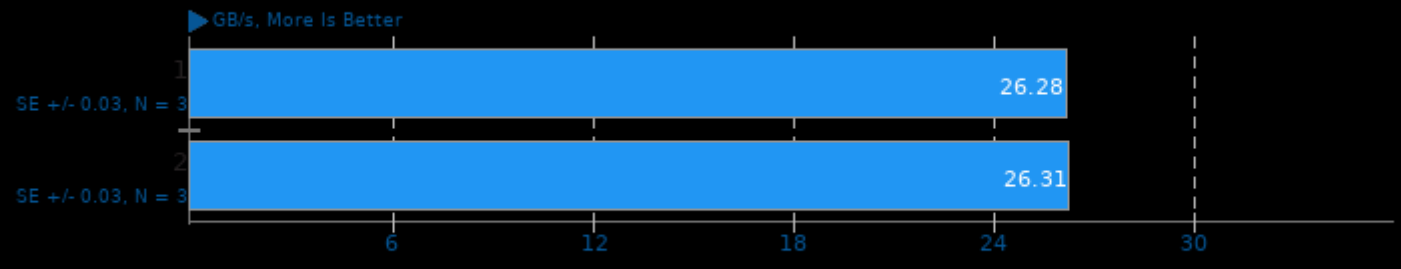

1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

FinanceBench 2016-07-25

Benchmark: Black-Scholes OpenCL

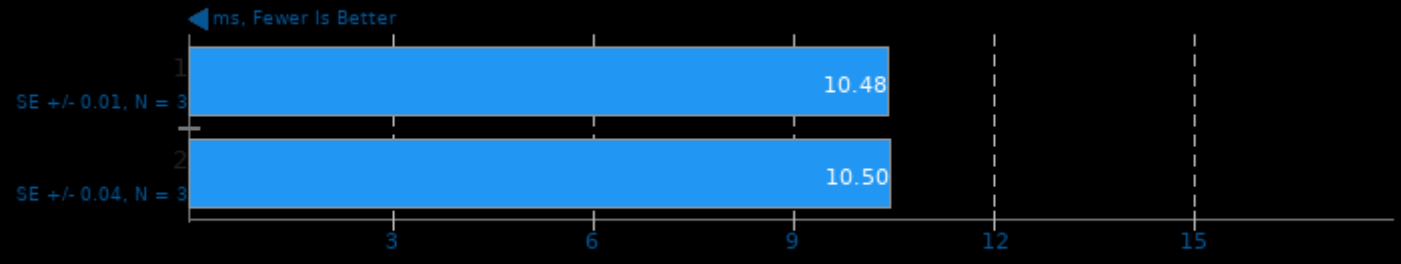

1. (CXX) g++ options: -O3 -march=native -fopenmp

Mixbench 2020-06-23

Backend: OpenCL - Benchmark: Integer

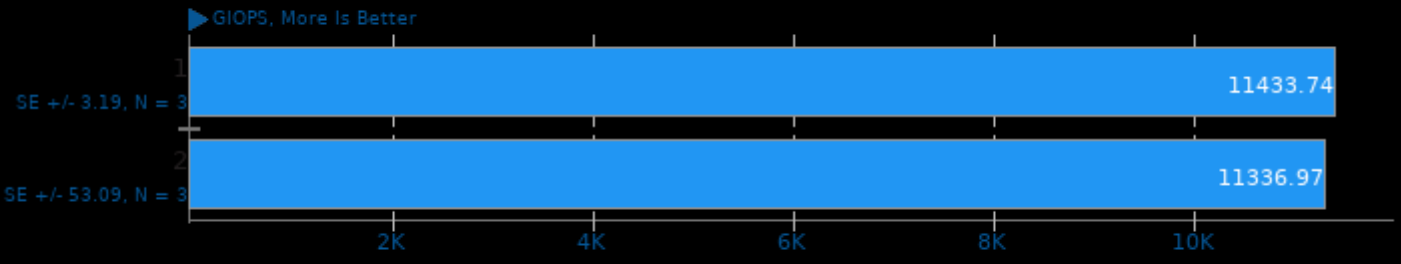

1. (CXX) g++ options: -lm -lstdc++ -lOpenCL -lrt -O2

clpeak

OpenCL Test: Single-Precision Float

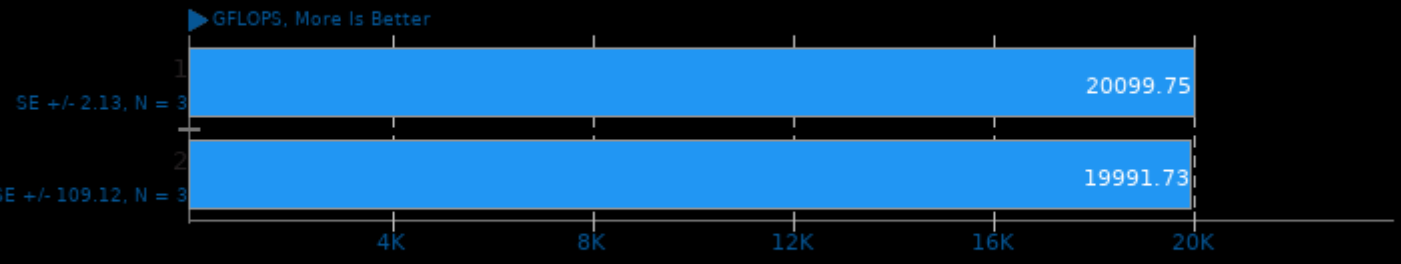

1. (CXX) g++ options: -O3 -rdynamic -lOpenCL

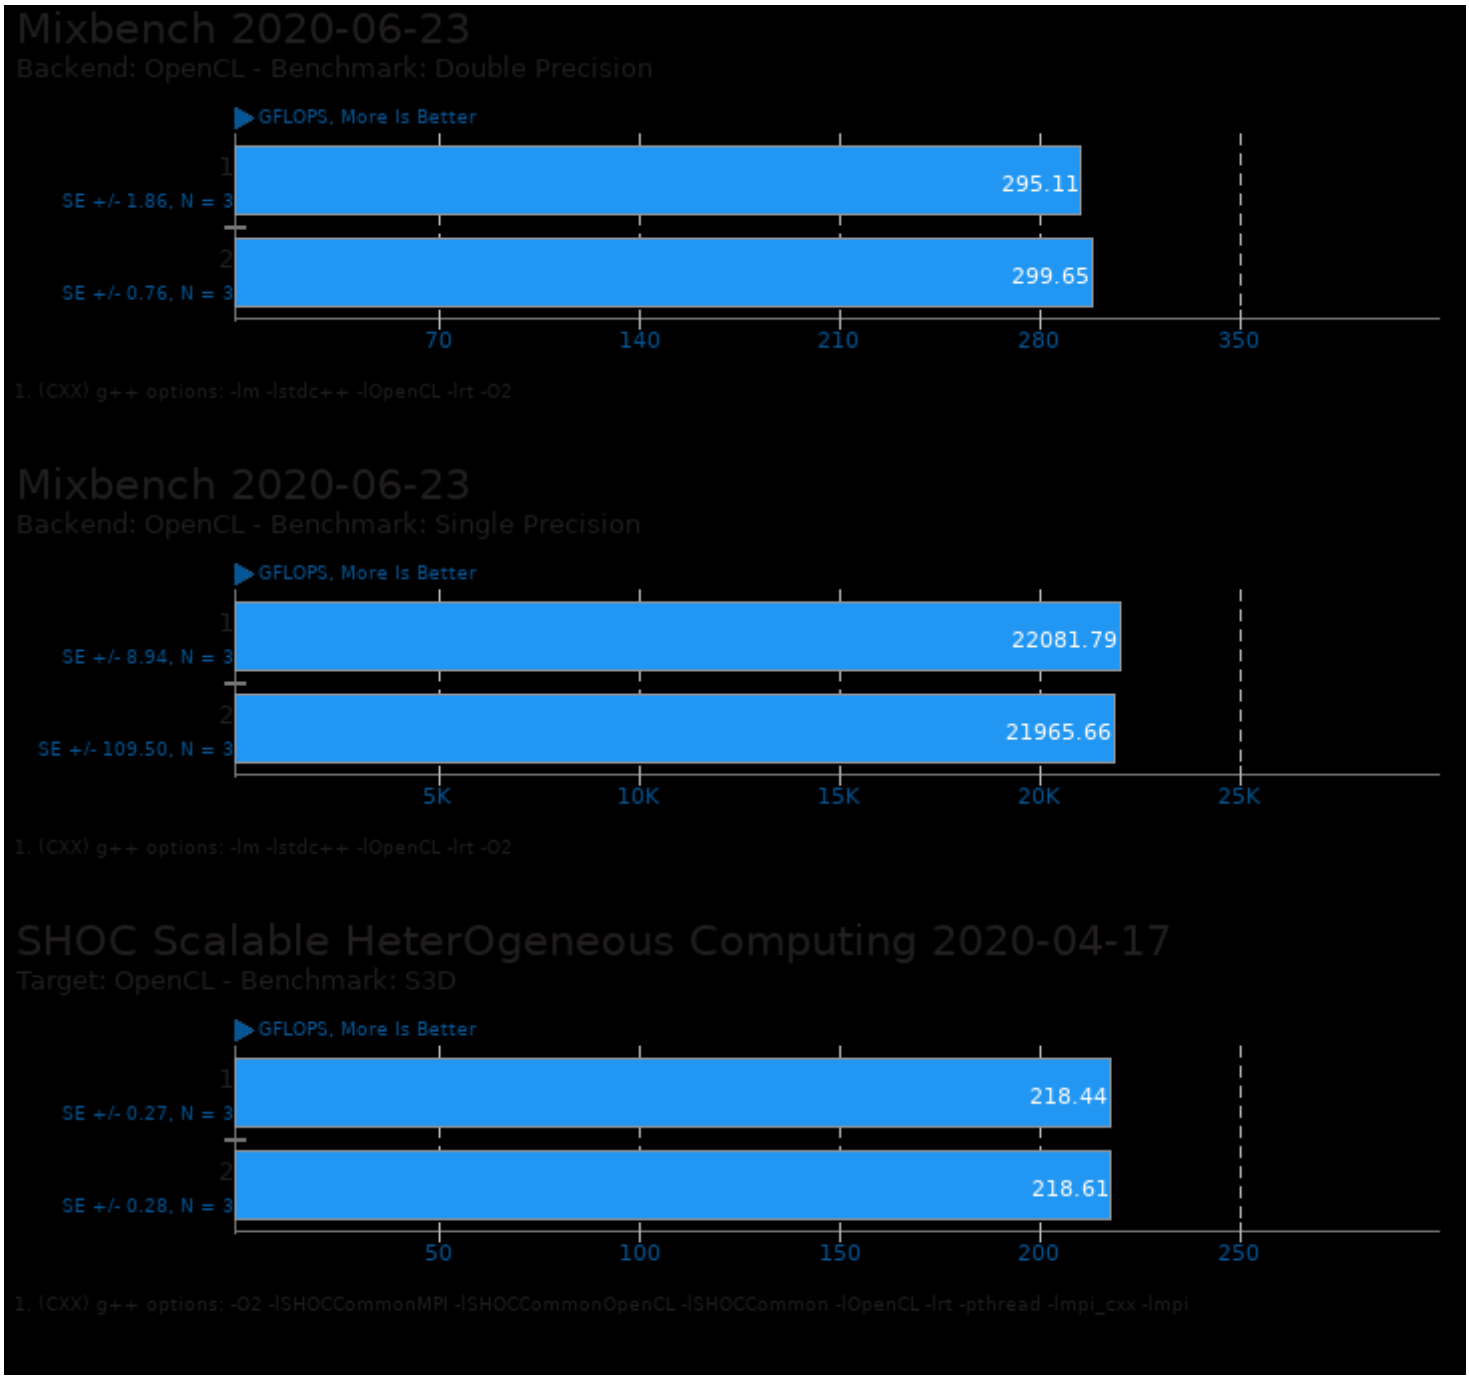
Mixbench 2020-06-23
Backend: OpenCL - Benchmark: Double Precision
GFLOPS, More Is Better
1
SE +/- 1.86, N = 3
295.11
2
SE +/- 0.76, N = 3
299.65
70
140
210
280
350
1. (CXX) g++ options: -lm -lstdc++ -lOpenCL -lrt -O2
Mixbench 2020-06-23
Backend: OpenCL - Benchmark: Single Precision
GFLOPS, More Is Better
1
SE +/- 8.94, N = 3
22081.79
2
SE +/- 109.50, N = 3
21965.66
5K
10K
15K
20K
25K
1. (CXX) g++ options: -lm -lstdc++ -lOpenCL -lrt -O2
SHOC Scalable HeterOgeneous Computing 2020-04-17
Target: OpenCL - Benchmark: S3D
GFLOPS, More Is Better
1
SE +/- 0.27, N = 3
218.44
2
SE +/- 0.28, N = 3
218.61
50
100
150
200
250
1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

RTX 3070 Compute

These geometric means are based upon test groupings / test suites for this result file.

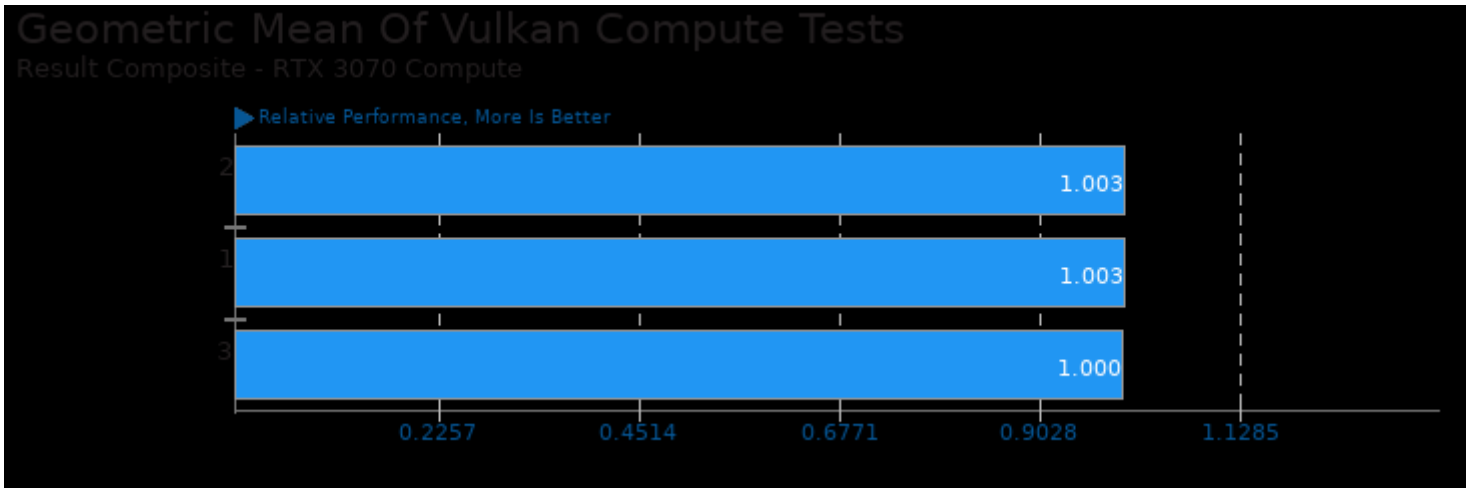

Geometric mean based upon tests: pts/vkfft, pts/vkresample, pts/ncnn, pts/realsr-ncnn, pts/waifu2x-ncnn and pts/betsy

This file was automatically generated via the Phoronix Test Suite benchmarking software on Thursday, 28 March 2024 05:37.