
**www.phoronix-test-suite.com**

# rtx3090

AMD Ryzen 9 7950X 16-Core testing with a ASUS ROG CROSSHAIR X670E HERO (0805 BIOS) and MSI NVIDIA GeForce RTX 4090 24GB on Linuxmint 21.1 via the Phoronix Test Suite.

**Automated Executive Summary**

*rtx4090 had the most wins, coming in first place for 85% of the tests.*

*Based on the geometric mean of all complete results, the fastest (rtx4090) was 1.544x the speed of the slowest (rtx3090-gpu).*

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

*MandelGPU (OpenCL Device: GPU) at 256.852x*
*NCNN (Target: Vulkan GPU - Model: googlenet) at 4.53x*
*SHOC Scalable HeterOgeneous Computing (Target: OpenCL - Benchmark: GEMM SGEMM_N) at 3.311x*
*SHOC Scalable HeterOgeneous Computing (Target: OpenCL - Benchmark: Reduction) at 2.639x*
*Hashcat (Benchmark: 7-Zip) at 2.54x*
*vkpeak (int16-vec4) at 2.445x*
*Hashcat (Benchmark: TrueCrypt RIPEMD160 + XTS) at 2.417x*
*clpeak (OpenCL Test: Single-Precision Float) at 2.394x*
*Hashcat (Benchmark: SHA-512) at 2.383x*

*clpeak (OpenCL Test: Integer Compute INT) at 2.374x.*

# Test Systems:

## rtx3090-gpu

Processor: AMD Ryzen 9 7950X 16-Core @ 5.88GHz (16 Cores / 32 Threads), Motherboard: ASUS ROG CROSSHAIR X670E HERO (0805 BIOS), Chipset: AMD Device 14d8, Memory: 62GB, Disk: 2000GB Sabrent Rocket 4.0 Plus + 2000GB CT2000P5PSSD8 + 2000GB Samsung SSD 980 PRO 2TB, Graphics: eVGA NVIDIA GeForce RTX 3090 24GB, Audio: NVIDIA GA102 HD Audio, Monitor: PHL 328E1 + LG ULTRAWIDE, Network: Intel I225-V + Intel Wi-Fi 6 AX210/AX211/AX411

OS: Linuxmint 21.1, Kernel: 5.15.0-58-generic (x86_64), Desktop: Cinnamon 5.6.7, Display Server: X Server 1.21.1.3, Display Driver: NVIDIA 525.60.11, OpenGL: 4.6.0, OpenCL: OpenCL 3.0 CUDA 12.0.89, Vulkan: 1.3.224, Compiler: GCC 11.3.0, File-System: ext4, Screen Resolution: 8960x2160

Kernel Notes: Transparent Huge Pages: madvise
Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --enable-cet --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++,m2 --enable-libphobos-checking=release --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-link-serialization=2 --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none=/build/gcc-11-xKiWfi/gcc-11-11.3.0/debian/tmp-nvptx/usr,amdgcn-amdhsa=/build/gcc-11-xKiWfi/gcc-11-11.3.0/debian/tmp-gcn/usr --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-build-config=bootstrap-lto-lean --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: amd-pstate performance (Boost: Enabled) - CPU Microcode: 0xa601203
Graphics Notes: BAR1 / Visible vRAM Size: 256 MiB - vBIOS Version: 94.02.42.c0.15
OpenCL Notes: GPU Compute Cores: 10496
Python Notes: Python 3.10.8
Security Notes: itlb_multihit: Not affected + l1tf: Not affected + mds: Not affected + meltdown: Not affected + mmio_stale_data: Not affected + retbleed: 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 Retpolines IBPB: conditional IBRS_FW STIBP: always-on RSB filling PBRSB-eIBRS: Not affected + srbds: Not affected + tsx_async_abort: Not affected

## rtx4090

Processor: AMD Ryzen 9 7950X 16-Core @ 5.88GHz (16 Cores / 32 Threads), Motherboard: ASUS ROG CROSSHAIR X670E HERO (0805 BIOS), Chipset: AMD Device 14d8, Memory: 62GB, Disk: 2000GB Sabrent Rocket 4.0 Plus + 2000GB CT2000P5PSSD8 + 2000GB Samsung SSD 980 PRO 2TB, Graphics: MSI NVIDIA GeForce RTX 4090 24GB, Audio: NVIDIA Device 22ba, Monitor: PHL 328E1 + LG ULTRAWIDE, Network: Intel I225-V + Intel Wi-Fi 6 AX210/AX211/AX411

OS: Linuxmint 21.1, Kernel: 5.15.0-58-generic (x86_64), Desktop: Cinnamon 5.6.7, Display Server: X Server 1.21.1.3, Display Driver: NVIDIA 525.78.01, OpenGL: 4.6.0, OpenCL: OpenCL 3.0 CUDA 12.0.133, Vulkan: 1.3.224, Compiler: GCC 11.3.0, File-System: ext4, Screen Resolution: 8960x2160

Kernel Notes: Transparent Huge Pages: madvise
Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --enable-cet --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++,m2 --enable-libphobos-checking=release --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-link-serialization=2 --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none=/build/gcc-11-xKiWfi/gcc-11-11.3.0/debian/tmp-nvptx/usr,amdgcn-amdhsa=/build/gcc-11-xKiWfi/gcc-11-11.3.0/debian/tmp-gcn/usr --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-build-config=bootstrap-lto-lean --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: amd-pstate performance (Boost: Enabled) - CPU Microcode: 0xa601203
Graphics Notes: BAR1 / Visible vRAM Size: 256 MiB - vBIOS Version: 95.02.18.80.74
OpenCL Notes: GPU Compute Cores: 16384
Python Notes: Python 3.10.8

# rtx3090

Security Notes: itlb_multihit: Not affected + l1tf: Not affected + mds: Not affected + meltdown: Not affected + mmio_stale_data: Not affected + retbleed: 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 Retpolines IBPB: conditional IBRS_FW STIBP: always-on RSB filling PBRSB-eIBRS: Not affected + srbds: Not affected + tsx_async_abort: Not affected

| | rtx3090-gpu | rtx4090 |
|---|---|---|
| **NCNN - Vulkan GPU - googlenet (ms)** | **10.51** | **2.32** |
| **Normalized** | 22.07% | 100% |
| **Standard Deviation** | 2.3% | 4.8% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - GEMM SGEMM_N (GFLOPS)** | **8429** | **27908** |
| **Normalized** | 30.2% | 100% |
| **Standard Deviation** | 1.2% | 2.2% |
| **MandelGPU - GPU (Samples/sec)** | **3132110** | **804490080** |
| **Normalized** | 0.39% | 100% |
| **Standard Deviation** | 0% | 1.1% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - Reduction (GB/s)** | **395.854** | **1045** |
| **Normalized** | 37.89% | 100% |
| **Standard Deviation** | 0.2% | 0% |
| **Hashcat - 7-Zip (H/s)** | **1092733** | **2775433** |
| **Normalized** | 39.37% | 100% |
| **Standard Deviation** | 2.3% | 0.5% |
| **vkpeak - int16-vec4 (GIOPS)** | **16741** | **40926** |
| **Normalized** | 40.9% | 100% |
| **Standard Deviation** | 0% | 0.1% |
| **Hashcat - T.R.X (H/s)** | **790800** | **1911500** |
| **Normalized** | 41.37% | 100% |
| **Standard Deviation** | 2% | 2.4% |
| **clpeak - S.P.F (GFLOPS)** | **35530** | **85069** |
| **Normalized** | 41.77% | 100% |
| **Standard Deviation** | 1% | 0.1% |
| **Hashcat - SHA-512 (H/s)** | **2735933333** | **6520400000** |
| **Normalized** | 41.96% | 100% |
| **Standard Deviation** | 0.2% | 0.1% |
| **clpeak - I.C.I (GIOPS)** | **18146** | **43084** |
| **Normalized** | 42.12% | 100% |
| **Standard Deviation** | 1.5% | 1.6% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - Max SP Flops (GFLOPS)** | **39392** | **92405** |
| **Normalized** | 42.63% | 100% |
| **Standard Deviation** | 0.9% | 0.2% |
| **Hashcat - SHA1 (H/s)** | **21815700000** | **51093900000** |
| **Normalized** | 42.7% | 100% |
| **Standard Deviation** | 0.4% | 0.3% |
| **ViennaCL - OpenCL BLAS - dGEMM-TT (GFLOPs/s)** | **601** | **1390** |
| **Normalized** | 43.24% | 100% |
| **Standard Deviation** | | 0% |
| **Hashcat - MD5 (H/s)** | **68180633333** | **157233333333** |
| **Normalized** | 43.36% | 100% |
| **Standard Deviation** | 0.1% | 0.1% |
| **vkpeak - int16-scalar (GIOPS)** | **13692** | **30756** |

| | | |
|---|---|---|
| Normalized | 44.52% | 100% |
| Standard Deviation | 0.1% | 0.1% |
| **vkpeak - fp64-vec4 (GFLOPS)** | **652.88** | **1455** |
| Normalized | 44.87% | 100% |
| Standard Deviation | 0.1% | 0% |
| **vkpeak - fp64-scalar (GFLOPS)** | **652.97** | **1455** |
| Normalized | 44.88% | 100% |
| Standard Deviation | 0.1% | 0% |
| **vkpeak - int32-scalar (GIOPS)** | **20755** | **46168** |
| Normalized | 44.96% | 100% |
| Standard Deviation | 0.1% | 0% |
| **vkpeak - fp32-vec4 (GFLOPS)** | **27458** | **61075** |
| Normalized | 44.96% | 100% |
| Standard Deviation | 0.4% | 0.1% |
| **vkpeak - int32-vec4 (GIOPS)** | **20661** | **45917** |
| Normalized | 45% | 100% |
| Standard Deviation | 0.1% | 0.2% |
| **ViennaCL - OpenCL BLAS - dGEMM-TN (GFLOPs/s)** | **603** | **1340** |
| Normalized | 45% | 100% |
| Standard Deviation | 0.4% | 0% |
| **vkpeak - fp16-scalar (GFLOPS)** | **20781** | **46167** |
| Normalized | 45.01% | 100% |
| Standard Deviation | 0.1% | 0.2% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - MD5 Hash (GHash/s)** | **43.1315** | **95.7989** |
| Normalized | 45.02% | 100% |
| Standard Deviation | 0.2% | 2.2% |
| **vkpeak - fp16-vec4 (GFLOPS)** | **41153** | **91396** |
| Normalized | 45.03% | 100% |
| Standard Deviation | 0.1% | 0.2% |
| **clpeak - D.P.D (GFLOPS)** | **653.69** | **1450** |
| Normalized | 45.09% | 100% |
| Standard Deviation | 0.2% | 0.1% |
| **vkpeak - fp32-scalar (GFLOPS)** | **20868** | **46188** |
| Normalized | 45.18% | 100% |
| Standard Deviation | 0.5% | 0.3% |
| **VkResample - 2x - Double (ms)** | **118.989** | **53.815** |
| Normalized | 45.23% | 100% |
| Standard Deviation | 0.2% | 0.1% |
| **ViennaCL - OpenCL BLAS - dGEMM-NT (GFLOPs/s)** | **602** | **1320** |
| Normalized | 45.61% | 100% |
| Standard Deviation | 1.3% | 0% |
| **FinanceBench - B.S.O (ms)** | **5.853** | **2.874** |
| Normalized | 49.1% | 100% |
| Standard Deviation | 0.1% | 2.5% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - Bus Speed Download (GB/s)** | **13.3594** | **26.8302** |
| Normalized | 49.79% | 100% |
| Standard Deviation | 0% | 0% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - Triad (GB/s)** | **13.1025** | **26.2868** |
| Normalized | 49.84% | 100% |
| Standard Deviation | 0.1% | 0.1% |

# rtx3090

| | | |
|---|---|---|
| **SHOC Scalable HeterOgeneous Computing - OpenCL - Bus Speed Readback (GB/s)** | **13.1913** | **26.3504** |
| Normalized | 50.06% | 100% |
| Standard Deviation | 0% | 0% |
| **Blender - Pabellon Barcelona - NVIDIA OptiX (sec)** | **15.43** | **7.74** |
| Normalized | 50.16% | 100% |
| Standard Deviation | 0.1% | 0.1% |
| **ViennaCL - OpenCL BLAS - dGEMM-NN (GFLOPs/s)** | **601** | **1197** |
| Normalized | 50.21% | 100% |
| Standard Deviation | 0.5% | 0.5% |
| **Blender - Classroom - NVIDIA OptiX (sec)** | **14.00** | **7.03** |
| Normalized | 50.21% | 100% |
| Standard Deviation | 0.3% | 0.5% |
| **OctaneBench - Total Score (Score)** | **686.766135** | **1366** |
| Normalized | 50.26% | 100% |
| **LuxCoreRender - Danish Mood - GPU (M samples/sec)** | **9.17** | **18.16** |
| Normalized | 50.5% | 100% |
| Standard Deviation | 0.3% | 0.2% |
| **LuxCoreRender - DLSC - GPU (M samples/sec)** | **11.28** | **21.90** |
| Normalized | 51.51% | 100% |
| Standard Deviation | 0.1% | 0.2% |
| **Blender - Fishy Cat - NVIDIA OptiX (sec)** | **10.15** | **5.28** |
| Normalized | 52.02% | 100% |
| Standard Deviation | 0.2% | 3.8% |
| **ArrayFire - C.G.O (ms)** | **1.598** | **0.8427** |
| Normalized | 52.73% | 100% |
| Standard Deviation | 0.6% | 0.2% |
| **LeelaChessZero - OpenCL (Nodes/s)** | **40156** | **74185** |
| Normalized | 54.13% | 100% |
| Standard Deviation | 2% | 1.3% |
| **IndigoBench - OpenCL GPU - Bedroom (M samples/s)** | **18.643** | **33.701** |
| Normalized | 55.32% | 100% |
| Standard Deviation | 0.4% | 0.3% |
| **Rodinia - O.P.F (sec)** | **3.767** | **2.091** |
| Normalized | 55.51% | 100% |
| Standard Deviation | 2.4% | |
| **Blender - Barbershop - NVIDIA OptiX (sec)** | **51.61** | **29.68** |
| Normalized | 57.51% | 100% |
| Standard Deviation | 0.4% | 0.3% |
| **LuxCoreRender - Orange Juice - GPU (M samples/sec)** | **10.18** | **17.42** |
| Normalized | 58.44% | 100% |
| Standard Deviation | 0.4% | 1.3% |
| **LuxCoreRender - LuxCore Benchmark - GPU (M samples/sec)** | **11.14** | **19.05** |
| Normalized | 58.48% | 100% |
| Standard Deviation | 0.4% | 0.3% |
| **IndigoBench - OpenCL GPU - Supercar (M samples/s)** | **49.238** | **77.657** |
| Normalized | 63.4% | 100% |
| Standard Deviation | 0.5% | 0% |
| **RealSR-NCNN - 4x - Yes (sec)** | **29.671** | **18.886** |
| Normalized | 63.65% | 100% |
| Standard Deviation | 0.3% | 0.1% |

# rtx3090

| | | |
|---|---|---|
| **SHOC Scalable HeterOgeneous Computing - OpenCL - S3D (GFLOPS)** | **427.911** | **653.058** |
| **Normalized** | 65.52% | 100% |
| **Standard Deviation** | 0.2% | 0.1% |
| **VkFFT (Benchmark Score)** | **42747** | **63465** |
| **Normalized** | 67.36% | 100% |
| **Standard Deviation** | 1.6% | 5.3% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - T.R.B (GB/s)** | **2199** | **3099** |
| **Normalized** | 70.97% | 100% |
| **Standard Deviation** | 0.2% | 0.1% |
| **FAHBench (Ns/Day)** | **328.4820** | **442.4502** |
| **Normalized** | 74.24% | 100% |
| **Standard Deviation** | 0.2% | 0.2% |
| **NeatBench - GPU (FPS)** | **3090** | **4090** |
| **Normalized** | 75.55% | 100% |
| **Standard Deviation** | 0% | 0% |
| **SHOC Scalable HeterOgeneous Computing - OpenCL - FFT SP (GFLOPS)** | **2141** | **2813** |
| **Normalized** | 76.12% | 100% |
| **Standard Deviation** | 0% | 0.1% |
| **ViennaCL - OpenCL BLAS - sCOPY (GB/s)** | **368** | **482** |
| **Normalized** | 76.35% | 100% |
| **Standard Deviation** | 0.3% | 0.2% |
| **Waifu2x-NCNN Vulkan - 2x - 3 - Yes (sec)** | **3.263** | **2.492** |
| **Normalized** | 76.37% | 100% |
| **Standard Deviation** | 0.2% | 0.6% |
| **RealSR-NCNN - 4x - No (sec)** | **5.613** | **4.289** |
| **Normalized** | 76.41% | 100% |
| **Standard Deviation** | 0.3% | 1.1% |
| **LuxCoreRender - R.C.a.P - GPU (M samples/sec)** | **31.03** | **40.30** |
| **Normalized** | 77% | 100% |
| **Standard Deviation** | 0.8% | 1.1% |
| **ViennaCL - OpenCL BLAS - sAXPY (GB/s)** | **501** | **610** |
| **Normalized** | 82.13% | 100% |
| **Standard Deviation** | 0.6% | 0.2% |
| **ViennaCL - OpenCL BLAS - sDOT (GB/s)** | **379** | **460** |
| **Normalized** | 82.39% | 100% |
| **Standard Deviation** | 0.3% | 0.1% |
| **VkResample - 2x - Single (ms)** | **9.350** | **7.746** |
| **Normalized** | 82.84% | 100% |
| **Standard Deviation** | 1.2% | 0.4% |
| **ViennaCL - OpenCL BLAS - dGEMV-T (GB/s)** | **375** | **450** |
| **Normalized** | 83.33% | 100% |
| **Standard Deviation** | 0% | 0.1% |
| **NAMD CUDA - ATPase Simulation - 327,506 Atoms (days/ns)** | **0.09694** | **0.08095** |
| **Normalized** | 83.51% | 100% |
| **Standard Deviation** | 5.5% | 4.4% |
| **ViennaCL - OpenCL BLAS - dGEMV-N (GB/s)** | **189** | **225** |
| **Normalized** | 84% | 100% |
| **Standard Deviation** | 0.3% | 0.3% |
| **cl-mem - Copy (GB/s)** | **362.3** | **416.7** |
| **Normalized** | 86.95% | 100% |
| **Standard Deviation** | 0.2% | 0.2% |

# rtx3090

| | | |
|---|---|---|
| **ViennaCL - CPU BLAS - sDOT (GB/s)** | **190** | **218** |
| **Normalized** | 87.16% | 100% |
| **Standard Deviation** | 4.5% | 5.3% |
| **ViennaCL - OpenCL BLAS - dDOT (GB/s)** | **654** | **732** |
| **Normalized** | 89.34% | 100% |
| **Standard Deviation** | 0.2% | 0.1% |
| **ViennaCL - CPU BLAS - sAXPY (GB/s)** | **223** | **245** |
| **Normalized** | 91.02% | 100% |
| **Standard Deviation** | 3.9% | 3.6% |
| **ViennaCL - OpenCL BLAS - dCOPY (GB/s)** | **606** | **664** |
| **Normalized** | 91.27% | 100% |
| **Standard Deviation** | 0.2% | 0% |
| **cl-mem - Write (GB/s)** | **751.5** | **813.7** |
| **Normalized** | 92.36% | 100% |
| **Standard Deviation** | 0.1% | 0.3% |
| **cl-mem - Read (GB/s)** | **827.1** | **887.5** |
| **Normalized** | 93.19% | 100% |
| **Standard Deviation** | 0.1% | 0% |
| **ViennaCL - OpenCL BLAS - dAXPY (GB/s)** | **724** | **776** |
| **Normalized** | 93.3% | 100% |
| **Standard Deviation** | 0.1% | 0.1% |
| **NCNN - Vulkan GPU - blazeface (ms)** | **0.91** | **0.85** |
| **Normalized** | 93.41% | 100% |
| **Standard Deviation** | 3.1% | 1.7% |
| **clpeak - G.M.B (GBPS)** | **814.22** | **869.72** |
| **Normalized** | 93.62% | 100% |
| **Standard Deviation** | 0% | 0.1% |
| **ViennaCL - CPU BLAS - sCOPY (GB/s)** | **152** | **162** |
| **Normalized** | 93.83% | 100% |
| **Standard Deviation** | 2.6% | 2% |
| **ViennaCL - CPU BLAS - dDOT (GB/s)** | **71.5** | **73.0** |
| **Normalized** | 97.95% | 100% |
| **Standard Deviation** | 4.4% | 1.7% |
| **ViennaCL - CPU BLAS - dCOPY (GB/s)** | **47.6** | **48.5** |
| **Normalized** | 98.14% | 100% |
| **Standard Deviation** | 3% | 0.8% |
| **ViennaCL - CPU BLAS - dAXPY (GB/s)** | **71.5** | **72.4** |
| **Normalized** | 98.76% | 100% |
| **Standard Deviation** | 3% | 3.7% |
| **ViennaCL - CPU BLAS - dGEMM-TN (GFLOPs/s)** | **117** | **116** |
| **Normalized** | 100% | 99.15% |
| **Standard Deviation** | 2% | 0.9% |
| **ViennaCL - CPU BLAS - dGEMM-TT (GFLOPs/s)** | 110 | 110 |
| **Standard Deviation** | 1.8% | 0.9% |
| **ViennaCL - CPU BLAS - dGEMM-NT (GFLOPs/s)** | 108 | 108 |
| **Standard Deviation** | 1.2% | 0.5% |
| **ViennaCL - CPU BLAS - dGEMM-NN (GFLOPs/s)** | 111 | 111 |
| **Standard Deviation** | 1.7% | 1.4% |
| **Blender - BMW27 - NVIDIA OptiX (sec)** | **5.78** | **13.06** |
| **Normalized** | 100% | 44.26% |
| **Standard Deviation** | 0.3% | 257.3% |
| **NCNN - Vulkan GPU - FastestDet (ms)** | **2.22** | **3.72** |
| **Normalized** | 100% | 59.68% |
| **Standard Deviation** | 6.8% | 116.8% |
| **NCNN - Vulkan GPU - vision_transformer (ms)** | **158.72** | **270.65** |

| | | |
|---|---|---|
| Normalized | 100% | 58.64% |
| Standard Deviation | 23.5% | 141.5% |
| **NCNN - Vulkan GPU - regnety_400m (ms)** | **2.11** | **1.48** |
| Normalized | 70.14% | 100% |
| Standard Deviation | 20.1% | 4.7% |
| **NCNN - Vulkan GPU - squeezenet_ssd (ms)** | **34.29** | **8.16** |
| Normalized | 23.8% | 100% |
| Standard Deviation | 8.3% | 157.1% |
| **NCNN - Vulkan GPU - yolov4-tiny (ms)** | **12.78** | **9.46** |
| Normalized | 74.02% | 100% |
| Standard Deviation | 67.7% | 124.7% |
| **NCNN - Vulkan GPU - resnet50 (ms)** | **2.56** | **3.75** |
| Normalized | 100% | 68.27% |
| Standard Deviation | 8.5% | 128.3% |
| **NCNN - Vulkan GPU - alexnet (ms)** | **1.17** | **2.14** |
| Normalized | 100% | 54.67% |
| Standard Deviation | 1.8% | 23.7% |
| **NCNN - Vulkan GPU - resnet18 (ms)** | **1.88** | **1.62** |
| Normalized | 86.17% | 100% |
| Standard Deviation | 13.8% | 7.3% |
| **NCNN - Vulkan GPU - vgg16 (ms)** | **3.10** | **3.37** |
| Normalized | 100% | 91.99% |
| Standard Deviation | 16.1% | 115.9% |
| **NCNN - Vulkan GPU - efficientnet-b0 (ms)** | **5.35** | **3.78** |
| Normalized | 70.65% | 100% |
| Standard Deviation | 43.8% | 113.4% |
| **NCNN - Vulkan GPU - mnasnet (ms)** | **1.23** | **2.50** |
| Normalized | 100% | 49.2% |
| Standard Deviation | 2.7% | 168.7% |
| **NCNN - Vulkan GPU - shufflenet-v2 (ms)** | **1.53** | **2.89** |
| Normalized | 100% | 52.94% |
| Standard Deviation | 8.2% | 156.4% |
| **NCNN - Vulkan GPU-v3-v3 - mobilenet-v3 (ms)** | **1.53** | **2.93** |
| Normalized | 100% | 52.22% |
| Standard Deviation | 4.6% | 121.3% |
| **NCNN - Vulkan GPU-v2-v2 - mobilenet-v2 (ms)** | **1.79** | **2.66** |
| Normalized | 100% | 67.29% |
| Standard Deviation | 4.6% | 147.9% |
| **NCNN - Vulkan GPU - mobilenet (ms)** | **3.22** | **6.98** |
| Normalized | 100% | 46.13% |
| Standard Deviation | 2.3% | 174.5% |
| **ViennaCL - CPU BLAS - dGEMV-T (GB/s)** | **90.6** | **98.8** |
| Normalized | 91.7% | 100% |
| Standard Deviation | 10.9% | 1.7% |
| **ViennaCL - CPU BLAS - dGEMV-N (GB/s)** | **77.7** | **81.5** |
| Normalized | 95.34% | 100% |
| Standard Deviation | 24.5% | 16% |

## NCNN 20220729

Target: Vulkan GPU - Model: googlenet

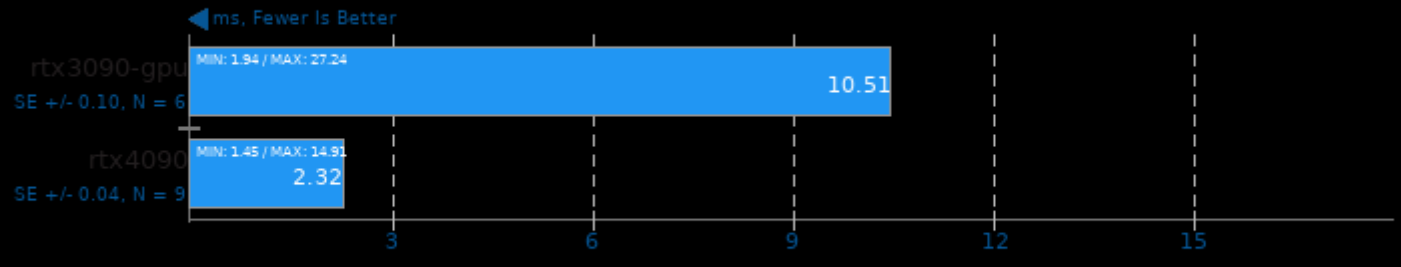


1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: GEMM SGEMM_N

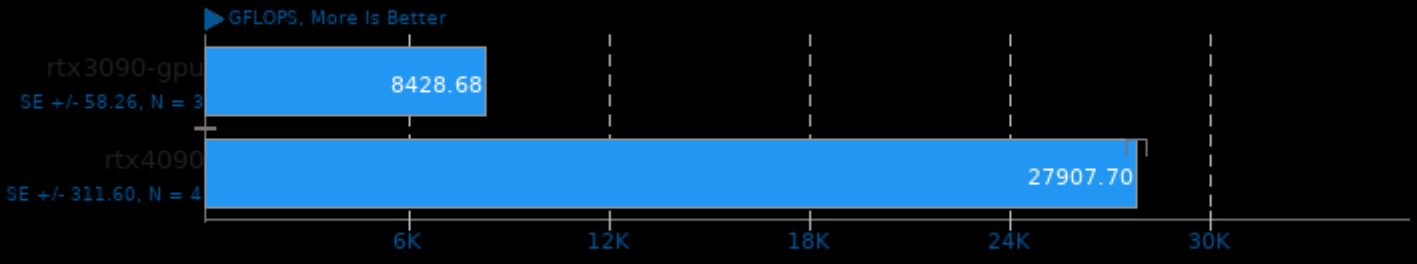


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

## MandelGPU 1.3pts1

OpenCL Device: GPU

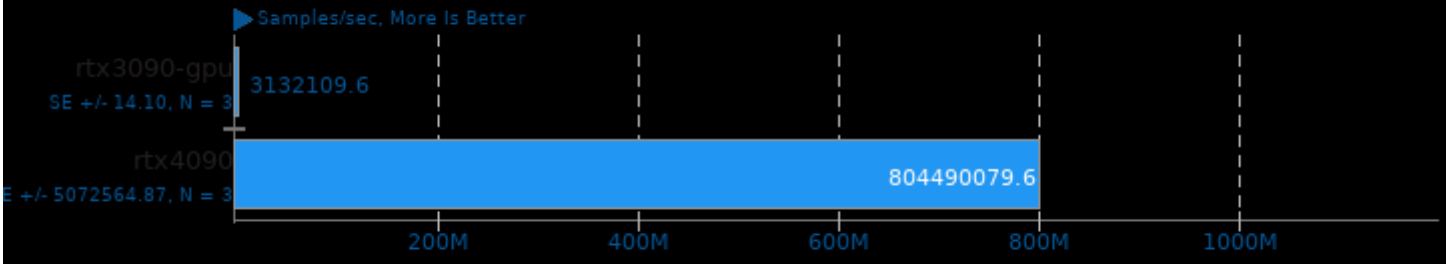


1. (CC) gcc options: -O3 -lm -ftree-vectorize -funroll-loops -lglut -lOpenCL -lGL

## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Reduction

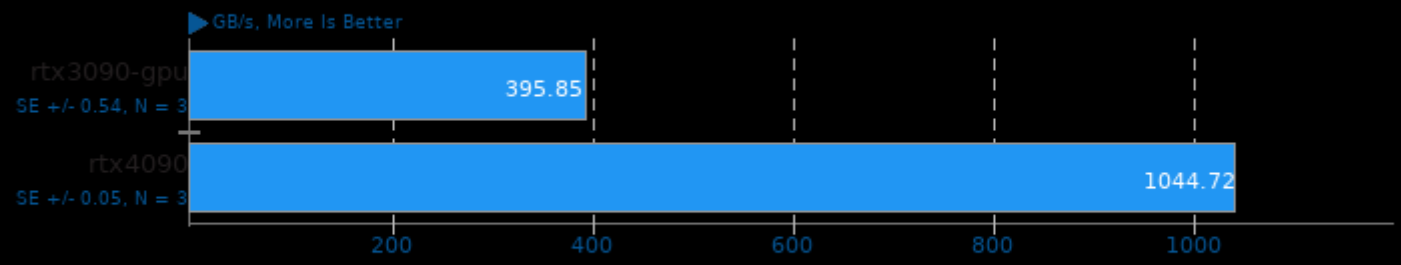


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

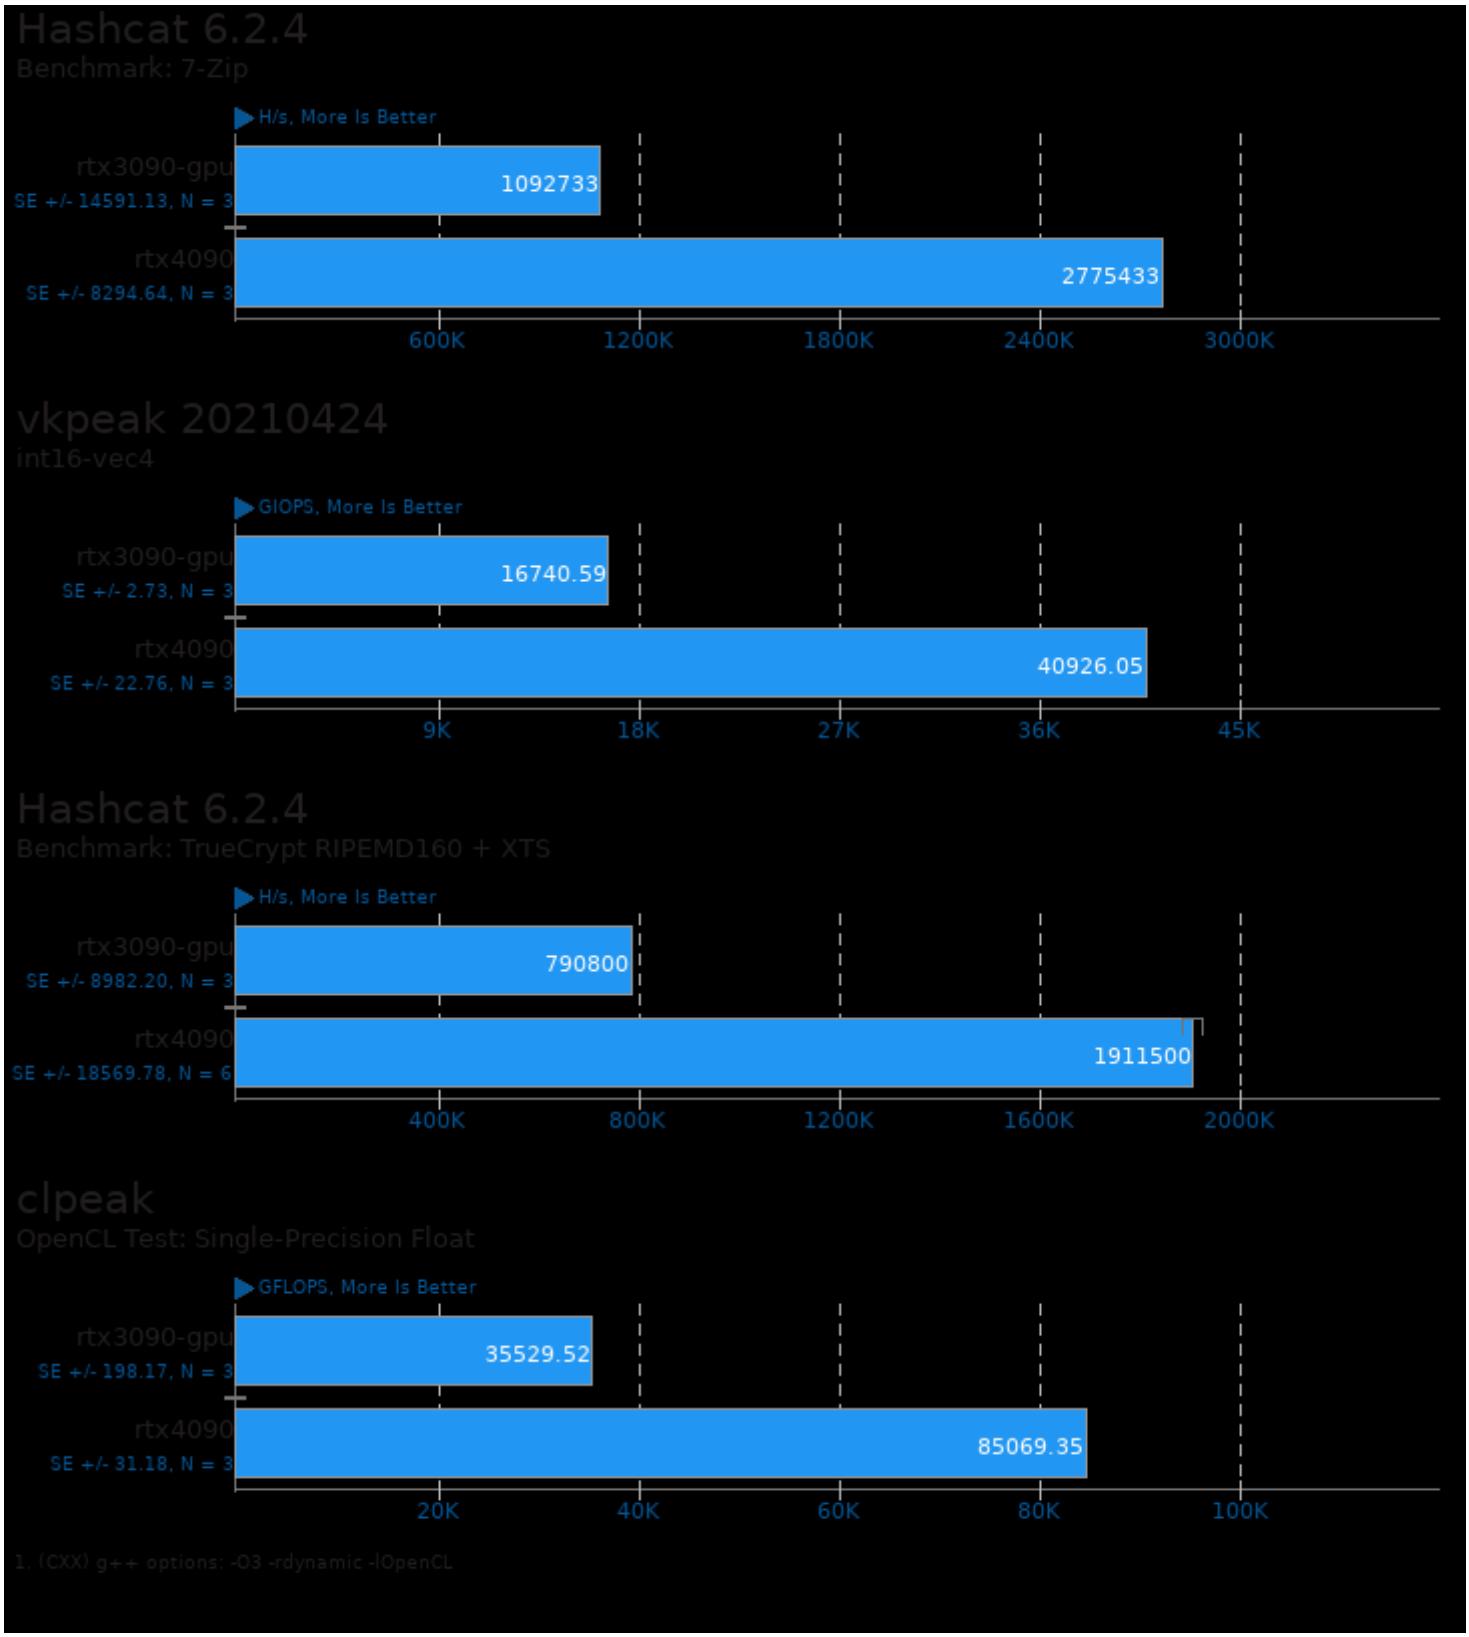
Hashcat 6.2.4
Benchmark: 7-Zip
H/s, More Is Better
rtx3090-gpu
SE +/- 14591.13, N = 3
1092733
rtx4090
SE +/- 8294.64, N = 3
2775433
600K
1200K
1800K
2400K
3000K
vkpeak 20210424
int16-vec4
GIOPS, More Is Better
rtx3090-gpu
SE +/- 2.73, N = 3
16740.59
rtx4090
SE +/- 22.76, N = 3
40926.05
9K
18K
27K
36K
45K
Hashcat 6.2.4
Benchmark: TrueCrypt RIPEMD160 + XTS
H/s, More Is Better
rtx3090-gpu
SE +/- 8982.20, N = 3
790800
rtx4090
SE +/- 18569.78, N = 6
1911500
400K
800K
1200K
1600K
2000K
clpeak
OpenCL Test: Single-Precision Float
GFLOPS, More Is Better
rtx3090-gpu
SE +/- 198.17, N = 3
35529.52
rtx4090
SE +/- 31.18, N = 3
85069.35
20K
40K
60K
80K
100K
1. (CXX) g++ options: -O3 -rdynamic -lOpenCL

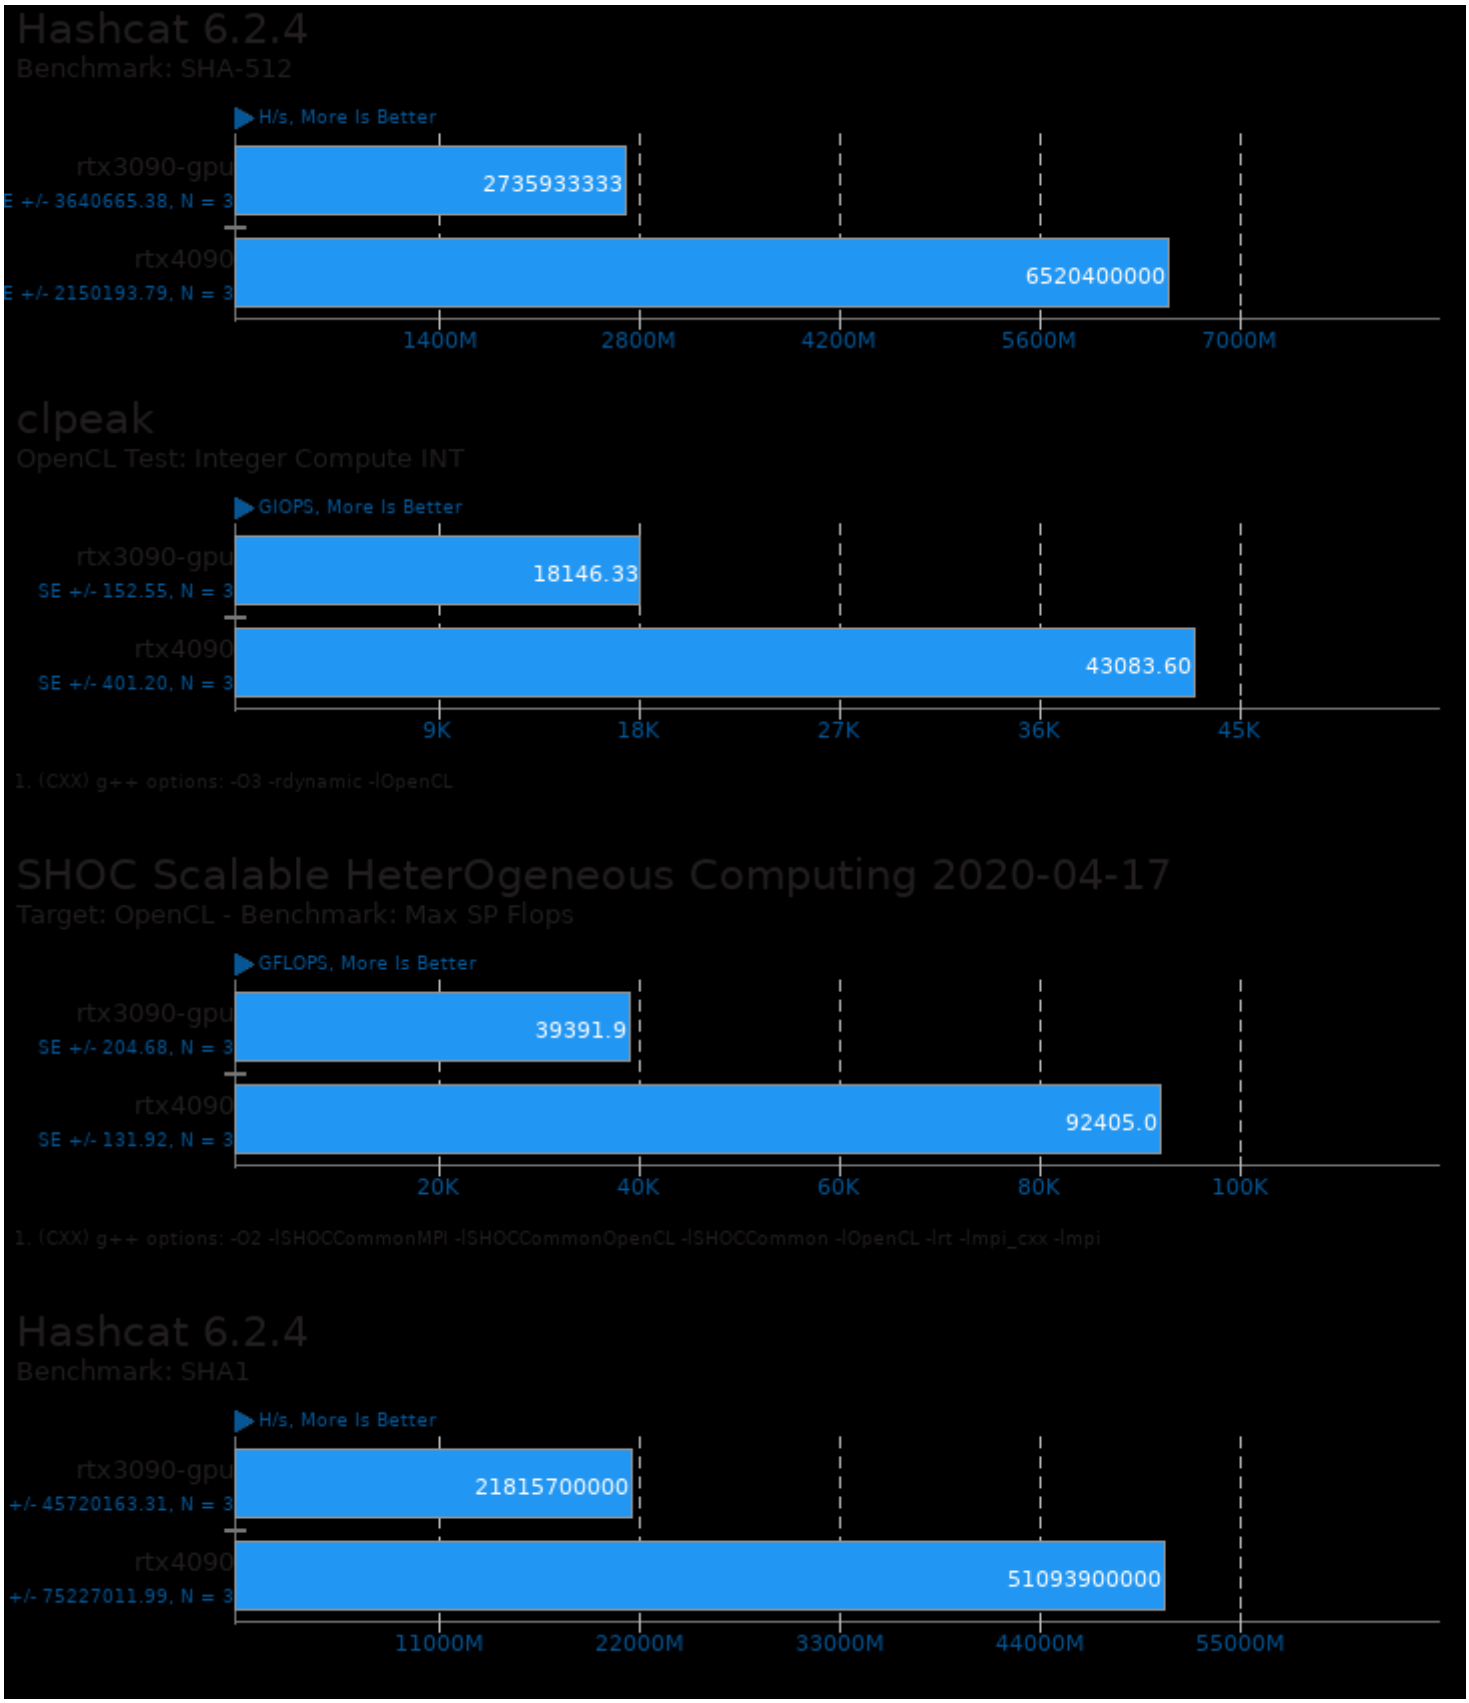
Hashcat 6.2.4
Benchmark: SHA-512
H/s, More Is Better
rtx3090-gpu
E +/- 3640665.38, N = 3
2735933333
rtx4090
E +/- 2150193.79, N = 3
6520400000
1400M
2800M
4200M
5600M
7000M
clpeak
OpenCL Test: Integer Compute INT
GIOPS, More Is Better
rtx3090-gpu
SE +/- 152.55, N = 3
18146.33
rtx4090
SE +/- 401.20, N = 3
43083.60
9K
18K
27K
36K
45K
1. (CXX) g++ options: -O3 -rdynamic -lOpenCL
SHOC Scalable HeterOgeneous Computing 2020-04-17
Target: OpenCL - Benchmark: Max SP Flops
GFLOPS, More Is Better
rtx3090-gpu
SE +/- 204.68, N = 3
39391.9
rtx4090
SE +/- 131.92, N = 3
92405.0
20K
40K
60K
80K
100K
1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi
Hashcat 6.2.4
Benchmark: SHA1
H/s, More Is Better
rtx3090-gpu
+/- 45720163.31, N = 3
21815700000
rtx4090
+/- 75227011.99, N = 3
51093900000
11000M
22000M
33000M
44000M
55000M

# rtx3090

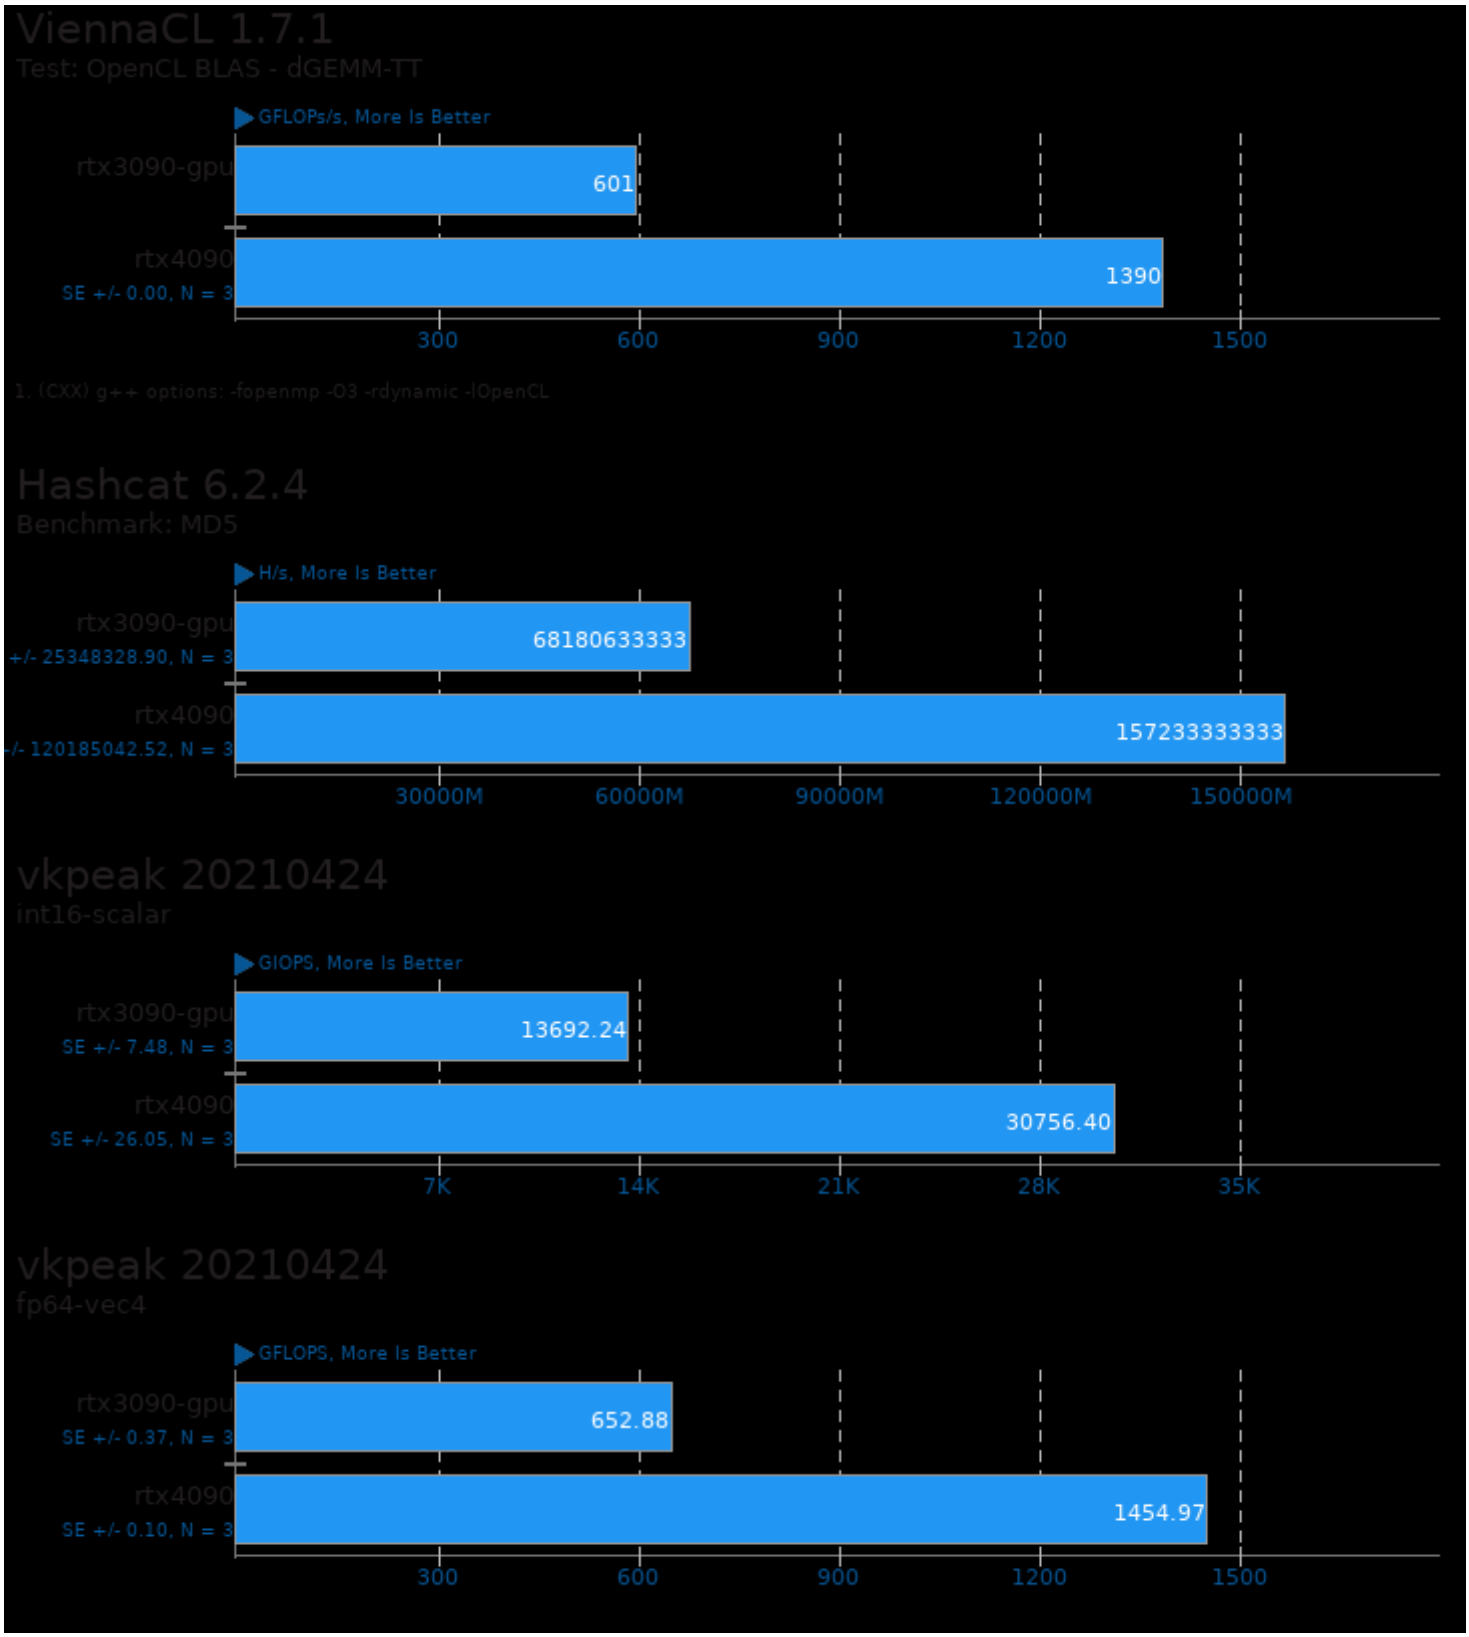
ViennaCL 1.7.1
Test: OpenCL BLAS - dGEMM-TT
GFLOPs/s, More Is Better
rtx3090-gpu
601
rtx4090
SE +/- 0.00, N = 3
1390
300
600
900
1200
1500
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
Hashcat 6.2.4
Benchmark: MD5
H/s, More Is Better
rtx3090-gpu
+/- 25348328.90, N = 3
68180633333
rtx4090
+/- 120185042.52, N = 3
157233333333
30000M
60000M
90000M
120000M
150000M
vkpeak 20210424
int16-scalar
GIOPS, More Is Better
rtx3090-gpu
SE +/- 7.48, N = 3
13692.24
rtx4090
SE +/- 26.05, N = 3
30756.40
7K
14K
21K
28K
35K
vkpeak 20210424
fp64-vec4
GFLOPS, More Is Better
rtx3090-gpu
SE +/- 0.37, N = 3
652.88
rtx4090
SE +/- 0.10, N = 3
1454.97
300
600
900
1200
1500

# rtx3090

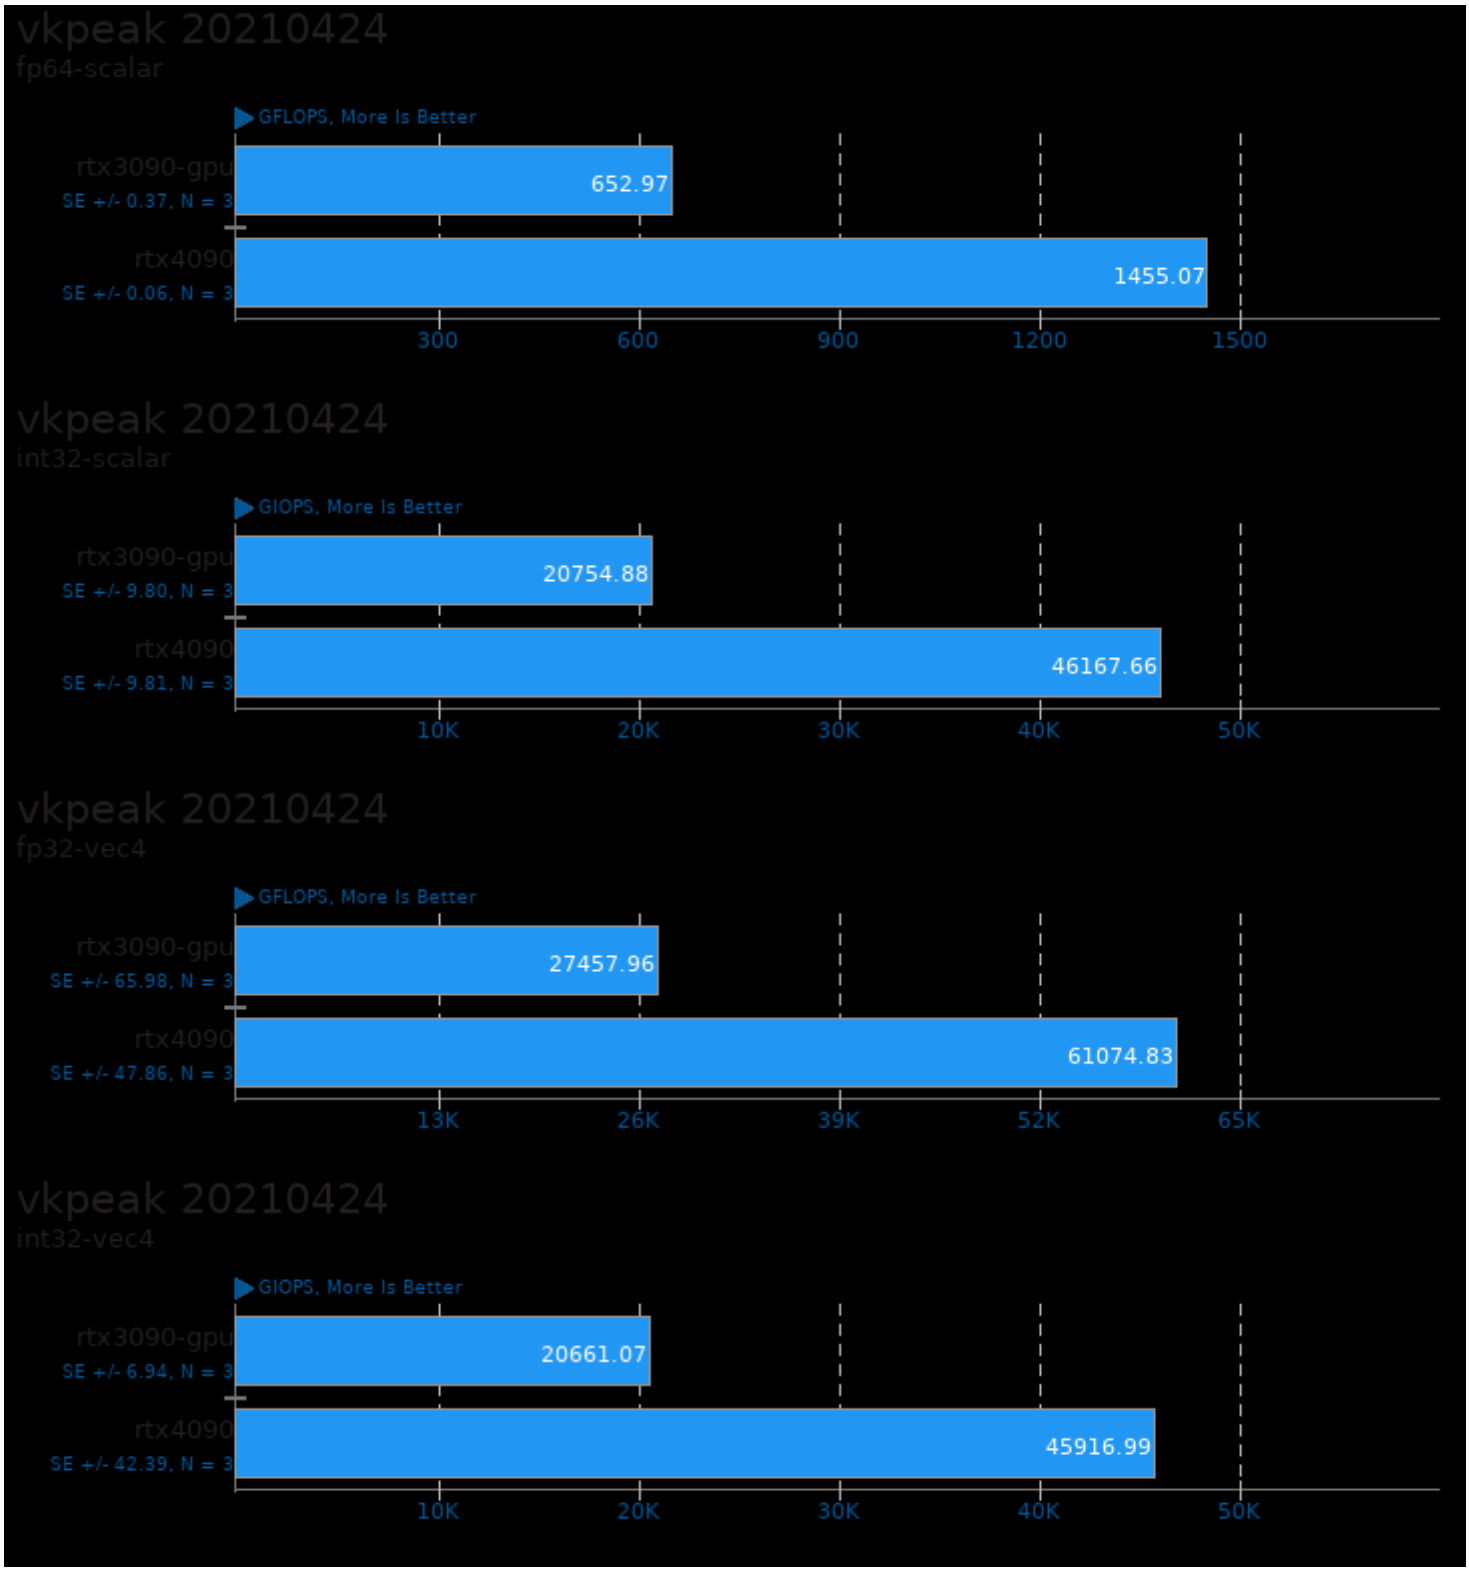
vkpeak 20210424
fp64-scalar
GFLOPS, More Is Better
rtx3090-gpu
SE +/- 0.37, N = 3
652.97
rtx4090
SE +/- 0.06, N = 3
1455.07
300
600
900
1200
1500
vkpeak 20210424
int32-scalar
GIOPS, More Is Better
rtx3090-gpu
SE +/- 9.80, N = 3
20754.88
rtx4090
SE +/- 9.81, N = 3
46167.66
10K
20K
30K
40K
50K
vkpeak 20210424
fp32-vec4
GFLOPS, More Is Better
rtx3090-gpu
SE +/- 65.98, N = 3
27457.96
rtx4090
SE +/- 47.86, N = 3
61074.83
13K
26K
39K
52K
65K
vkpeak 20210424
int32-vec4
GIOPS, More Is Better
rtx3090-gpu
SE +/- 6.94, N = 3
20661.07
rtx4090
SE +/- 42.39, N = 3
45916.99
10K
20K
30K
40K
50K

## ViennaCL 1.7.1

Test: OpenCL BLAS - dGEMM-TN

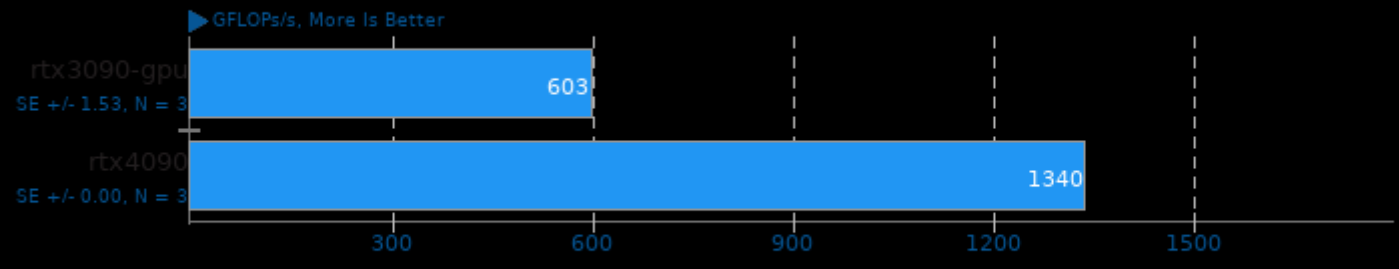


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## vkpeak 20210424

fp16-scalar

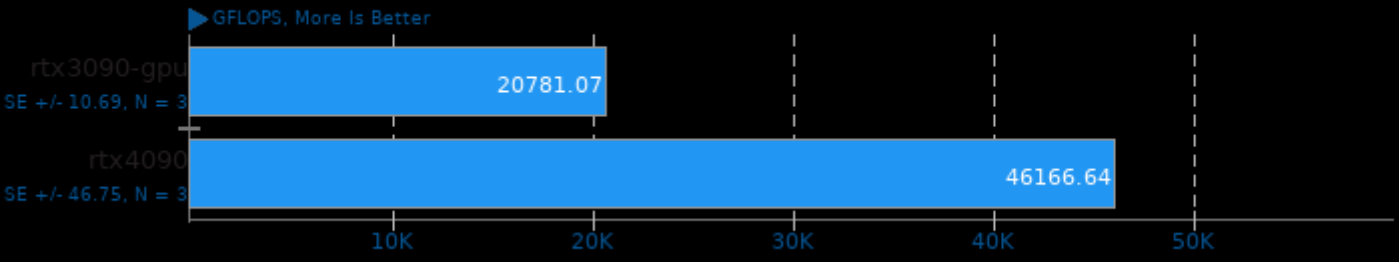


## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: MD5 Hash

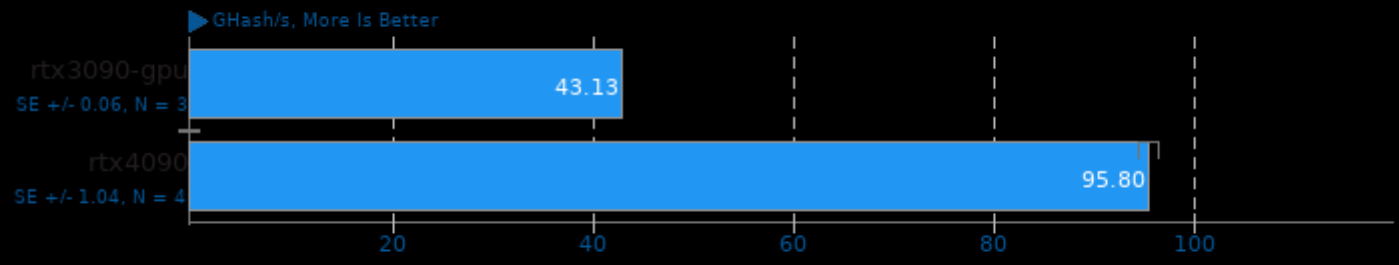


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

## vkpeak 20210424

fp16-vec4

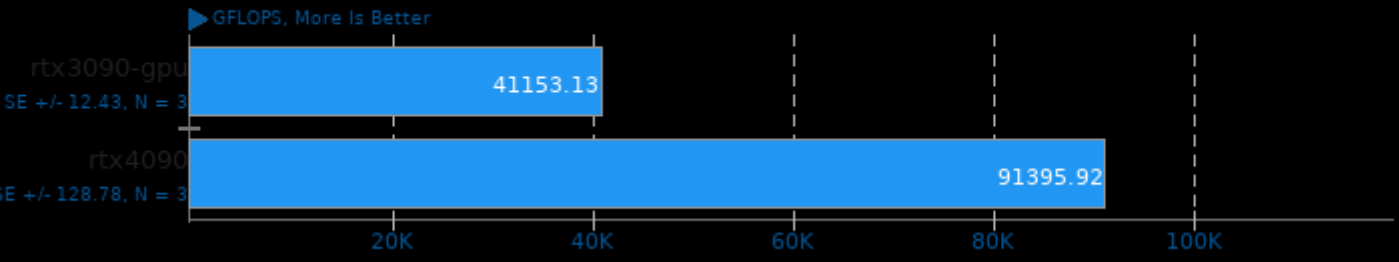

## clpeak

OpenCL Test: Double-Precision Double

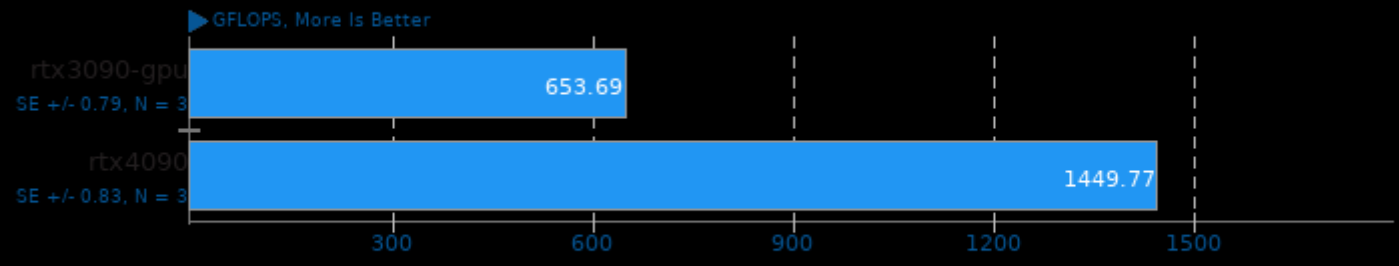


1. (CXX) g++ options: -O3 -rdynamic -lOpenCL

## vkpeak 20210424

fp32-scalar

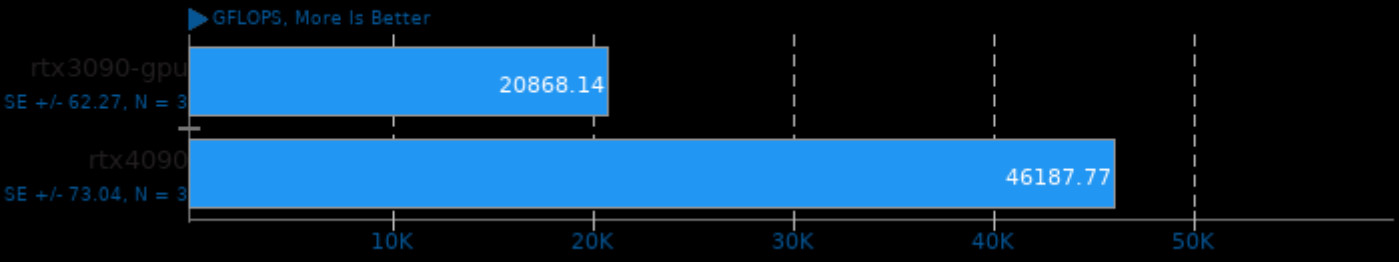


## VkResample 1.0

Upscale: 2x - Precision: Double

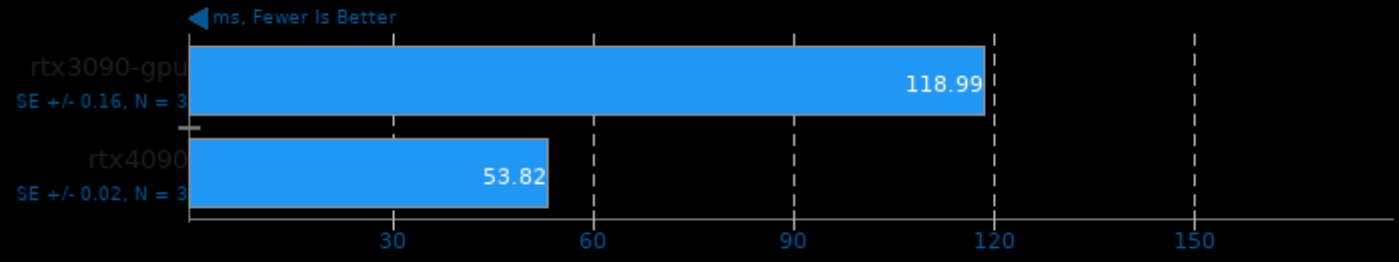


1. (CXX) g++ options: -O3

## ViennaCL 1.7.1

Test: OpenCL BLAS - dGEMM-NT

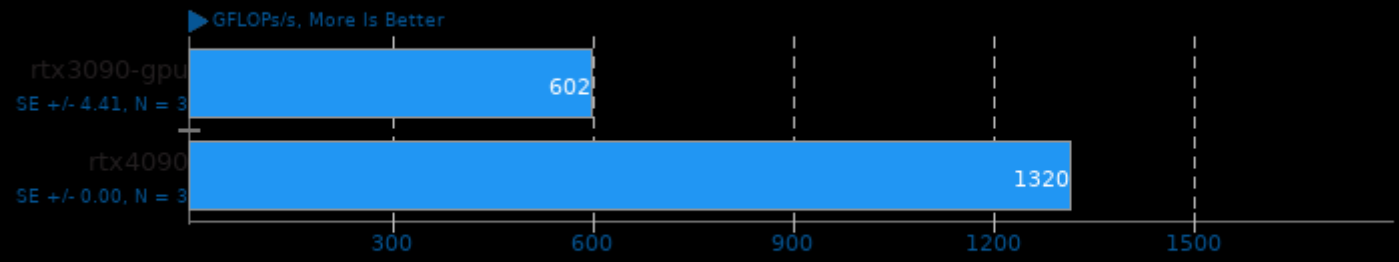


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## FinanceBench 2016-07-25

Benchmark: Black-Scholes OpenCL

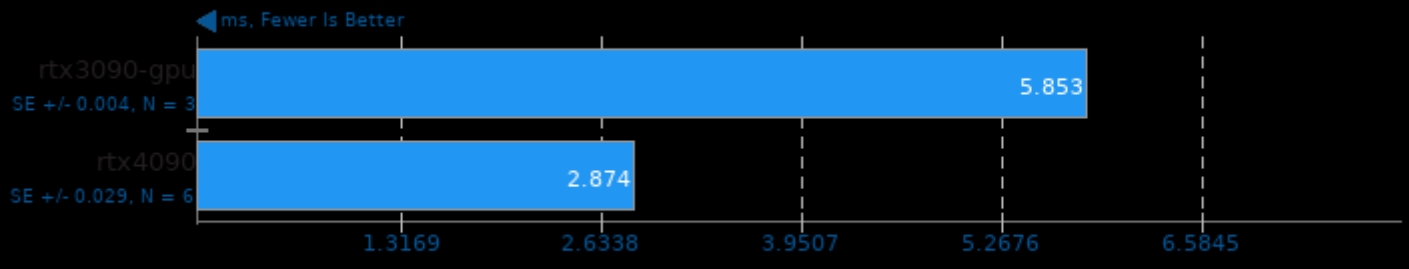


1. (CXX) g++ options: -O3 -march=native -fopenmp

## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Bus Speed Download

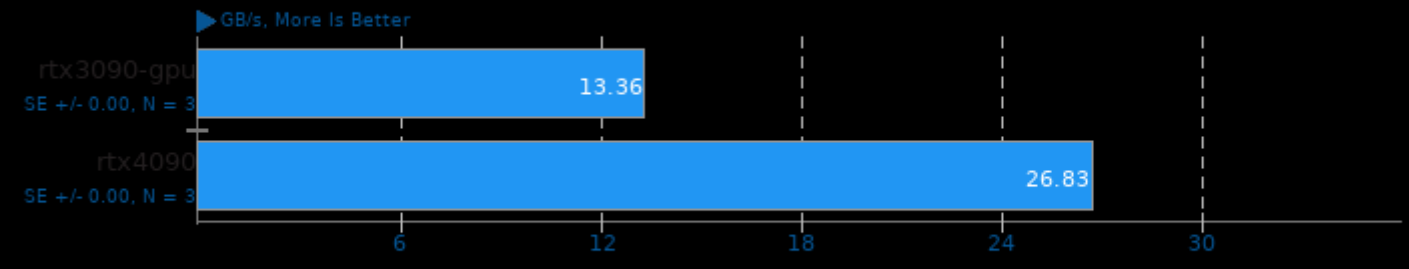


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Triad

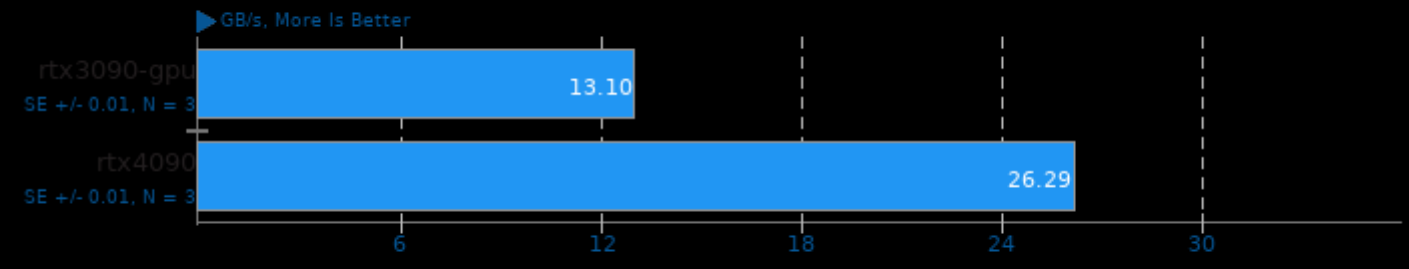


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Bus Speed Readback

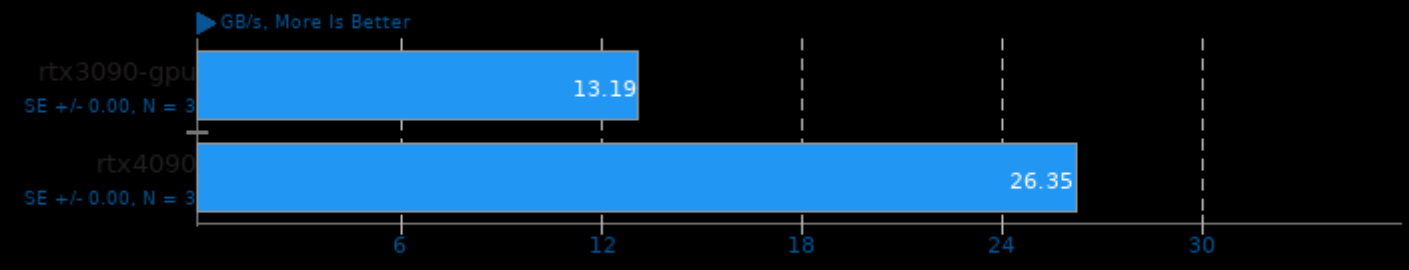


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

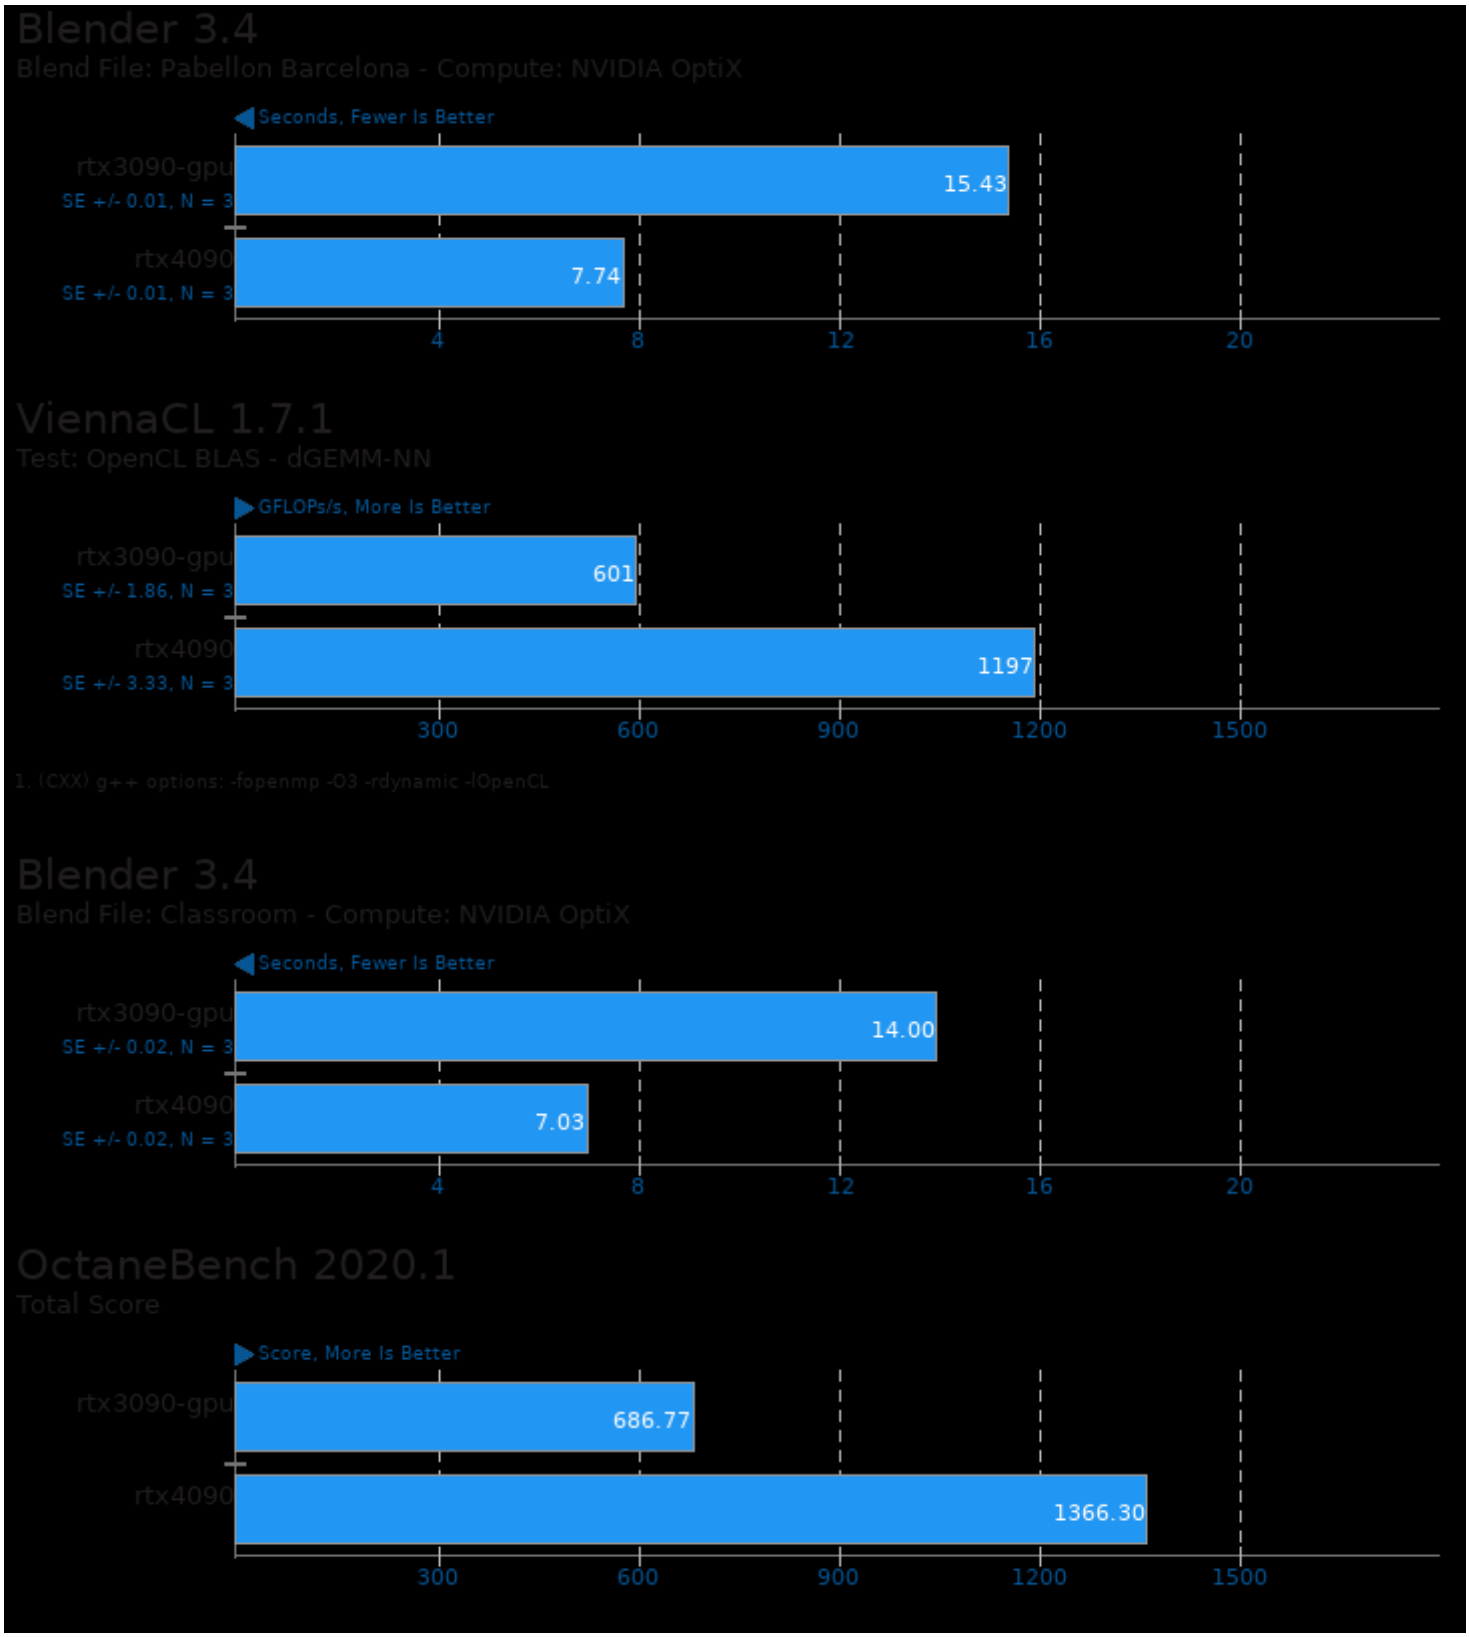
Blender 3.4
Blend File: Pabellon Barcelona - Compute: NVIDIA OptiX
Seconds, Fewer Is Better
rtx3090-gpu
SE +/- 0.01, N = 3
15.43
rtx4090
SE +/- 0.01, N = 3
7.74
4
8
12
16
20
ViennaCL 1.7.1
Test: OpenCL BLAS - dGEMM-NN
GFLOPs/s, More Is Better
rtx3090-gpu
SE +/- 1.86, N = 3
601
rtx4090
SE +/- 3.33, N = 3
1197
300
600
900
1200
1500
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
Blender 3.4
Blend File: Classroom - Compute: NVIDIA OptiX
Seconds, Fewer Is Better
rtx3090-gpu
SE +/- 0.02, N = 3
14.00
rtx4090
SE +/- 0.02, N = 3
7.03
4
8
12
16
20
OctaneBench 2020.1
Total Score
Score, More Is Better
rtx3090-gpu
686.77
rtx4090
1366.30
300
600
900
1200
1500

# rtx3090

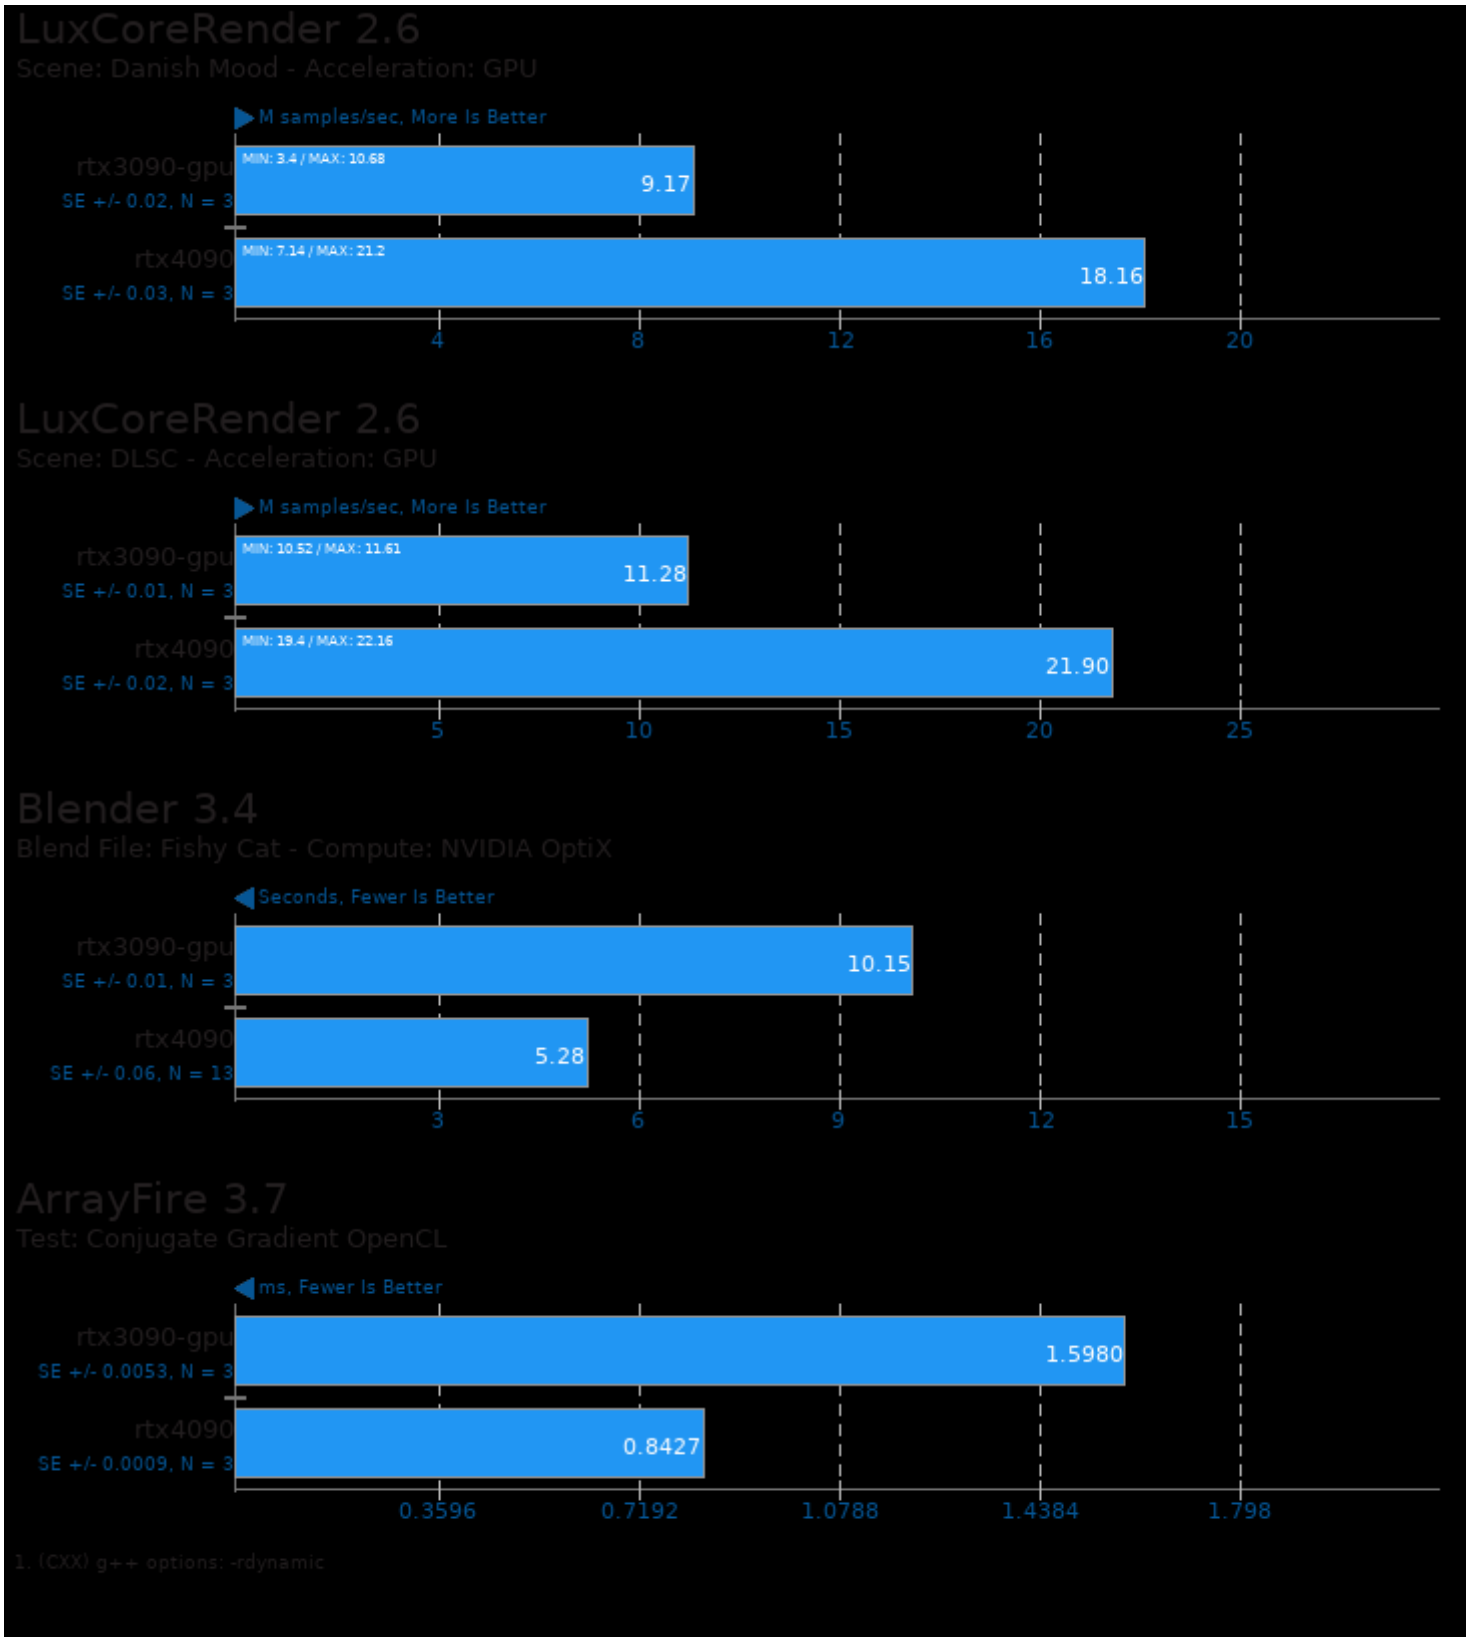
LuxCoreRender 2.6
Scene: Danish Mood - Acceleration: GPU
M samples/sec, More Is Better
rtx3090-gpu
SE +/- 0.02, N = 3
MIN: 3.4 / MAX: 10.68
9.17
rtx4090
SE +/- 0.03, N = 3
MIN: 7.14 / MAX: 21.2
18.16
4
8
12
16
20
LuxCoreRender 2.6
Scene: DLSC - Acceleration: GPU
M samples/sec, More Is Better
rtx3090-gpu
SE +/- 0.01, N = 3
MIN: 10.52 / MAX: 11.61
11.28
rtx4090
SE +/- 0.02, N = 3
MIN: 19.4 / MAX: 22.16
21.90
5
10
15
20
25
Blender 3.4
Blend File: Fishy Cat - Compute: NVIDIA OptiX
Seconds, Fewer Is Better
rtx3090-gpu
SE +/- 0.01, N = 3
10.15
rtx4090
SE +/- 0.06, N = 13
5.28
3
6
9
12
15
ArrayFire 3.7
Test: Conjugate Gradient OpenCL
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.0053, N = 3
1.5980
rtx4090
SE +/- 0.0009, N = 3
0.8427
0.3596
0.7192
1.0788
1.4384
1.798
1. (CXX) g++ options: -rdynamic

# rtx3090

## LeelaChessZero 0.28

Backend: OpenCL

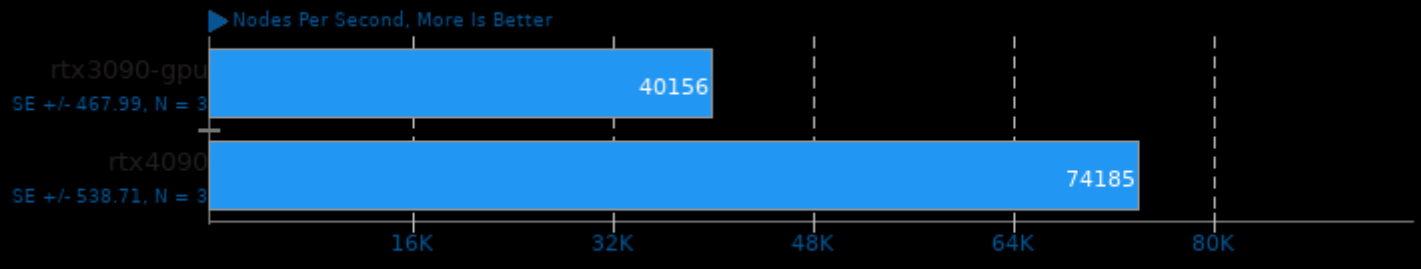


1. (CXX) g++ options: -flto -pthread

## IndigoBench 4.4

Acceleration: OpenCL GPU - Scene: Bedroom

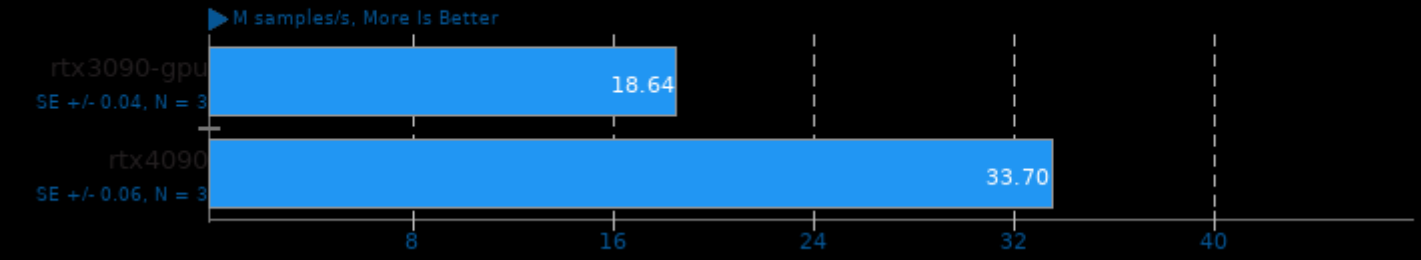


## Rodinia 3.1

Test: OpenCL Particle Filter

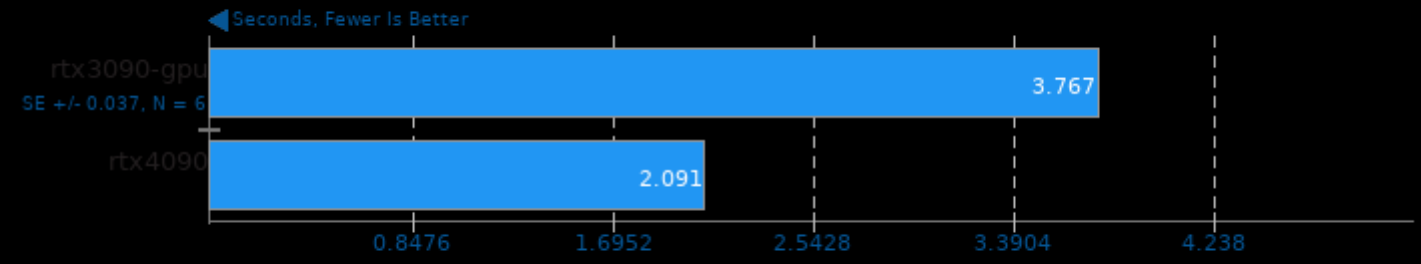


1. (CXX) g++ options: -O2 -lOpenCL

## Blender 3.4

Blend File: Barbershop - Compute: NVIDIA OptiX

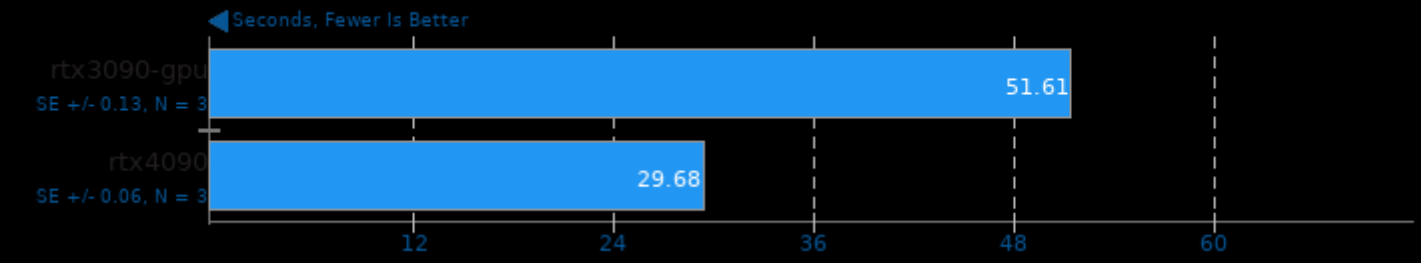

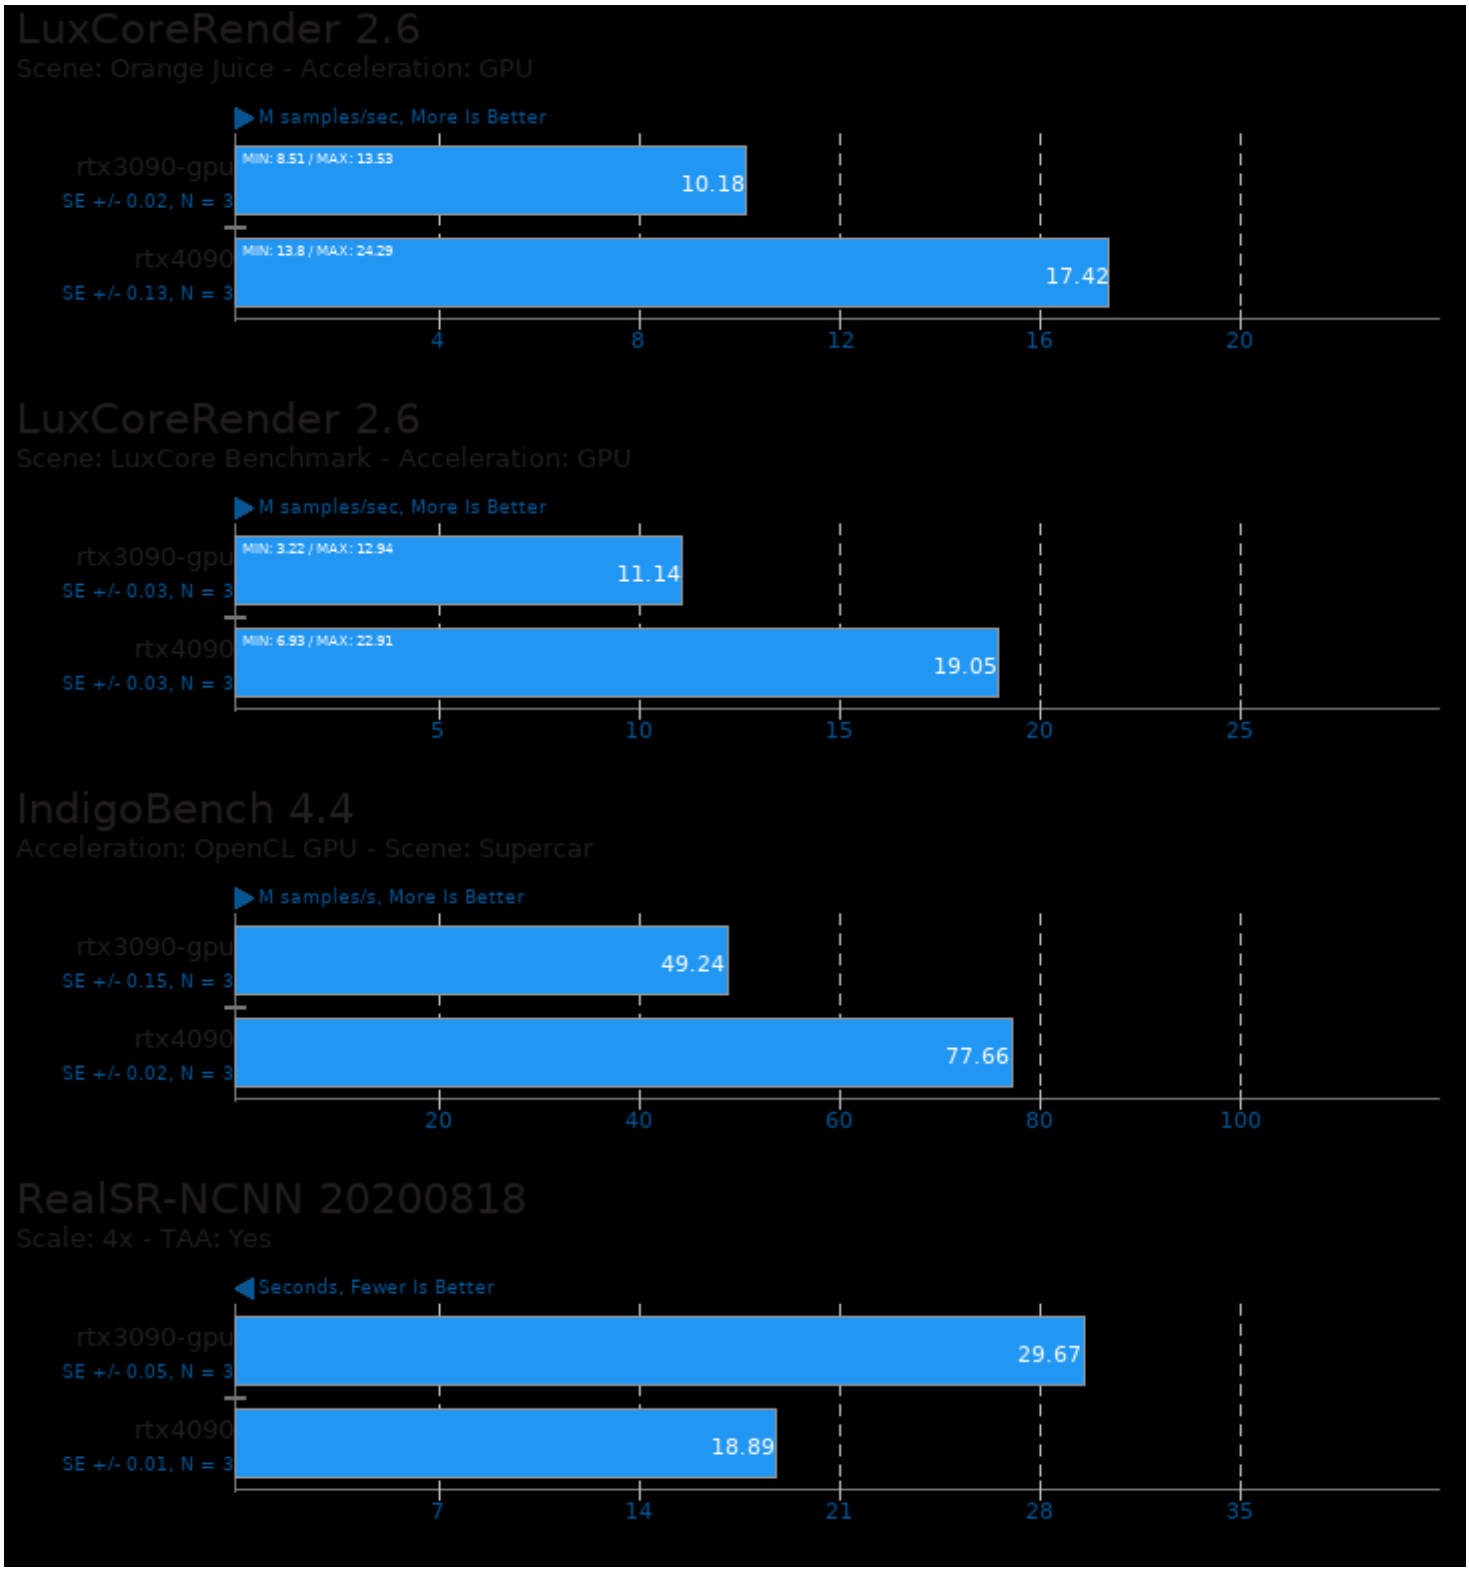
LuxCoreRender 2.6
Scene: Orange Juice - Acceleration: GPU
M samples/sec, More Is Better
rtx3090-gpu
SE +/- 0.02, N = 3
MIN: 8.51 / MAX: 13.53
10.18
rtx4090
SE +/- 0.13, N = 3
MIN: 13.8 / MAX: 24.29
17.42
4
8
12
16
20
LuxCoreRender 2.6
Scene: LuxCore Benchmark - Acceleration: GPU
M samples/sec, More Is Better
rtx3090-gpu
SE +/- 0.03, N = 3
MIN: 3.22 / MAX: 12.94
11.14
rtx4090
SE +/- 0.03, N = 3
MIN: 6.93 / MAX: 22.91
19.05
5
10
15
20
25
IndigoBench 4.4
Acceleration: OpenCL GPU - Scene: Supercar
M samples/s, More Is Better
rtx3090-gpu
SE +/- 0.15, N = 3
49.24
rtx4090
SE +/- 0.02, N = 3
77.66
20
40
60
80
100
RealSR-NCNN 20200818
Scale: 4x - TAA: Yes
Seconds, Fewer Is Better
rtx3090-gpu
SE +/- 0.05, N = 3
29.67
rtx4090
SE +/- 0.01, N = 3
18.89
7
14
21
28
35

## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: S3D

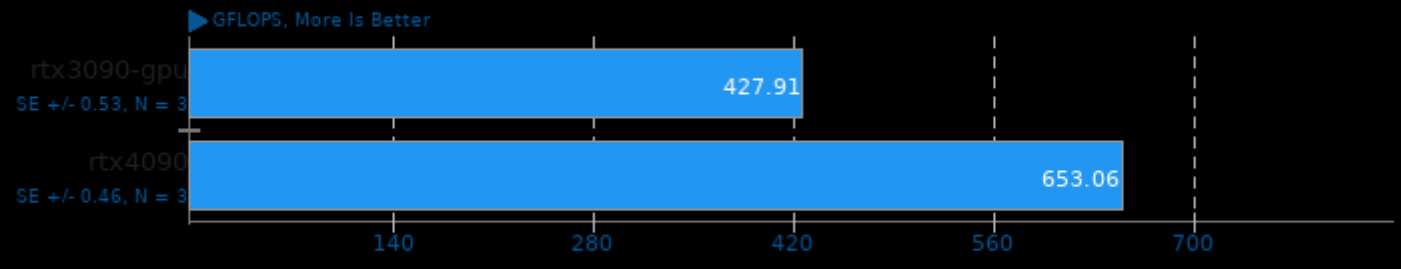


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

## VkFFT 1.1.1

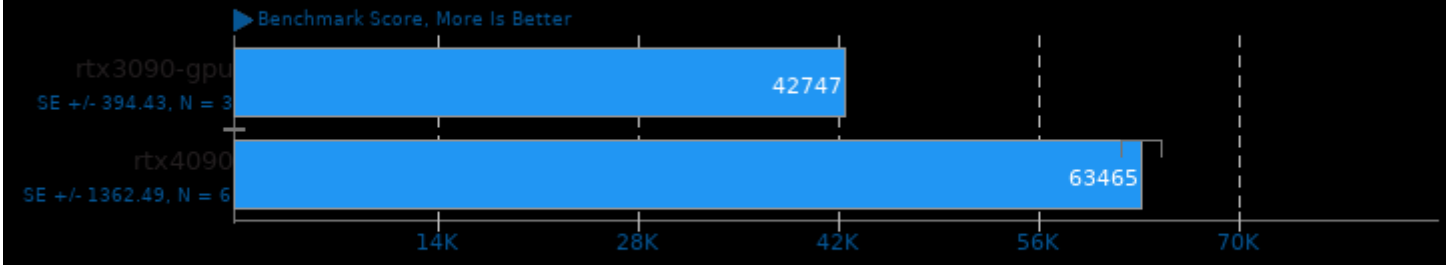


1. (CXX) g++ options: -O3

## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: Texture Read Bandwidth

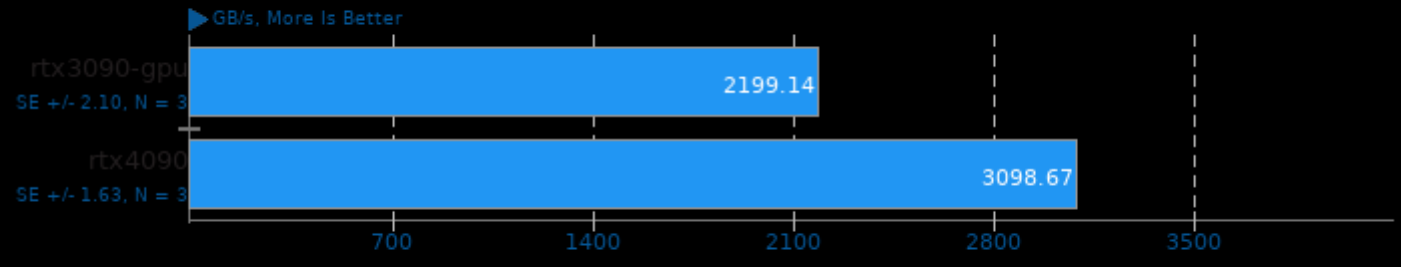


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

## FAHBench 2.3.2

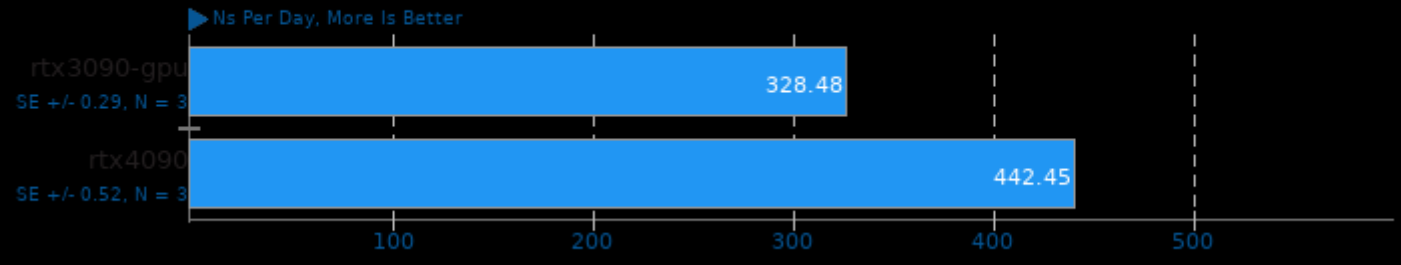

## NeatBench 5

Acceleration: GPU

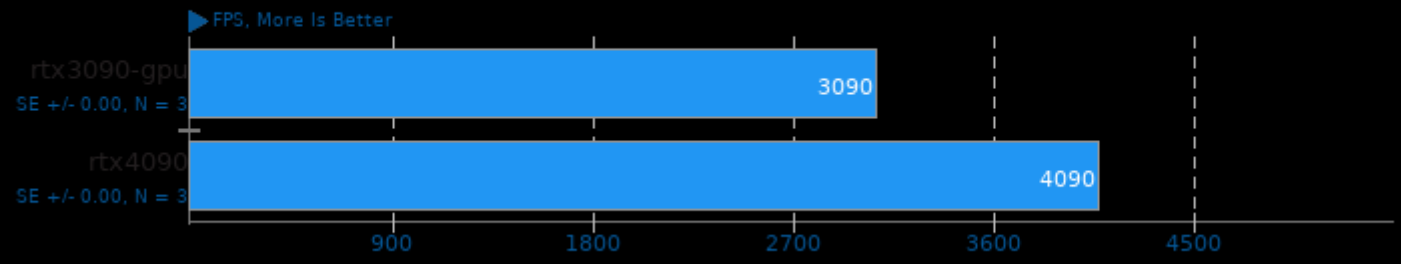


## SHOC Scalable HeterOgeneous Computing 2020-04-17

Target: OpenCL - Benchmark: FFT SP

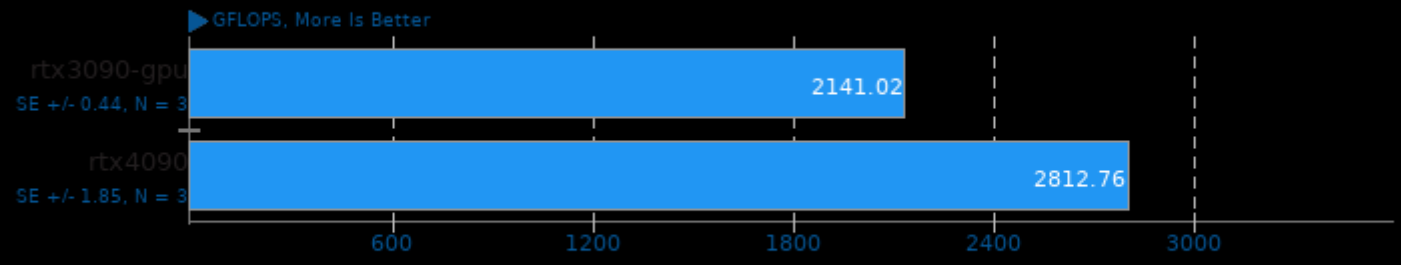


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -lmpi_cxx -lmpi

## ViennaCL 1.7.1

Test: OpenCL BLAS - sCOPY

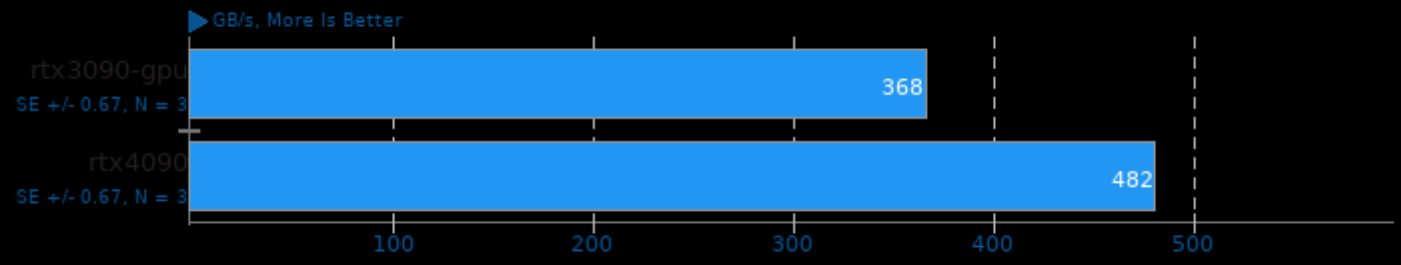


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## Waifu2x-NCNN Vulkan 20200818

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

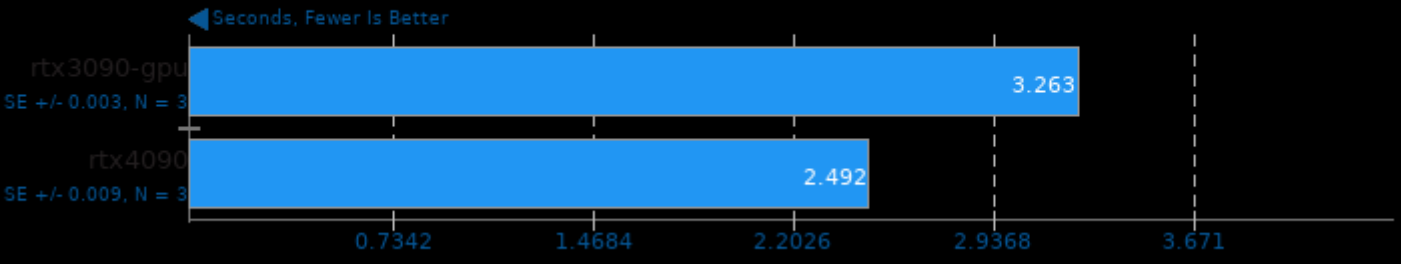

## RealSR-NCNN 20200818

Scale: 4x - TAA: No

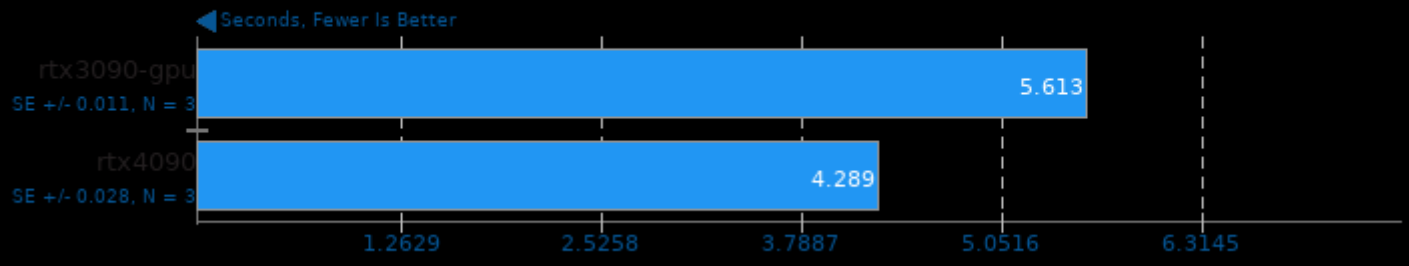


## LuxCoreRender 2.6

Scene: Rainbow Colors and Prism - Acceleration: GPU

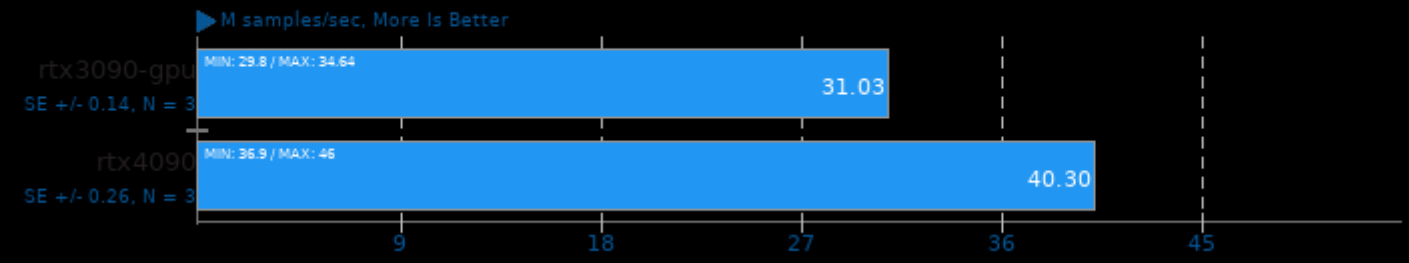


## ViennaCL 1.7.1

Test: OpenCL BLAS - sAXPY

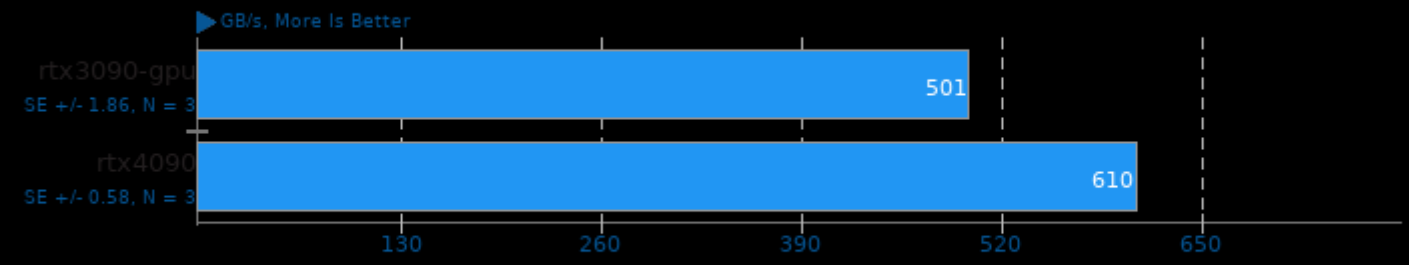


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## ViennaCL 1.7.1

Test: OpenCL BLAS - sDOT

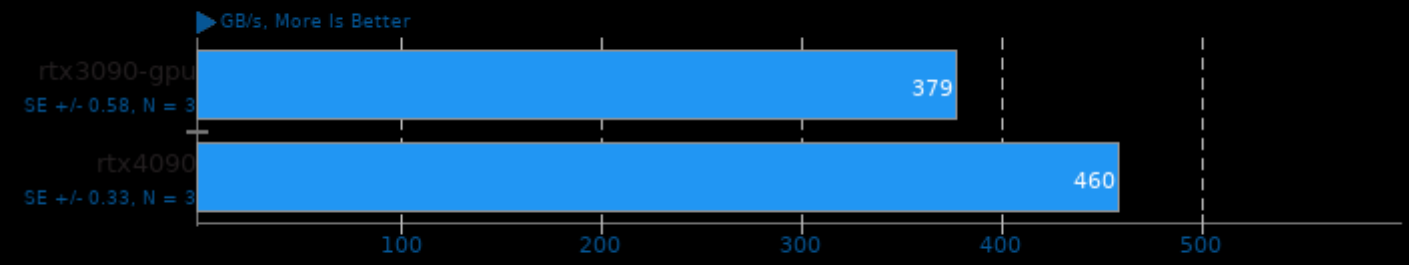


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## VkResample 1.0

Upscale: 2x - Precision: Single

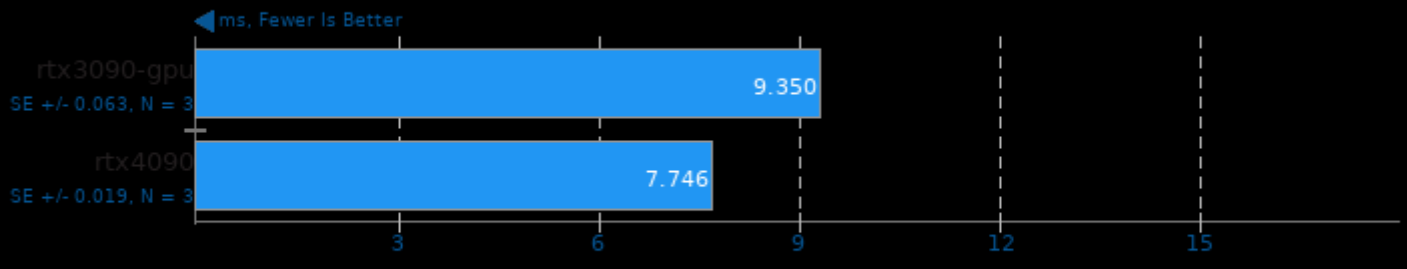


1. (CXX) g++ options: -O3

## ViennaCL 1.7.1

Test: OpenCL BLAS - dGEMV-T

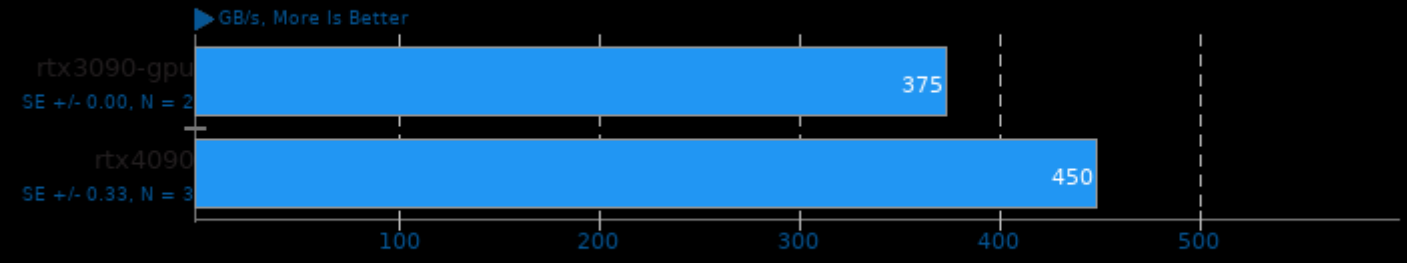


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## NAMD CUDA 2.14

ATPase Simulation - 327,506 Atoms

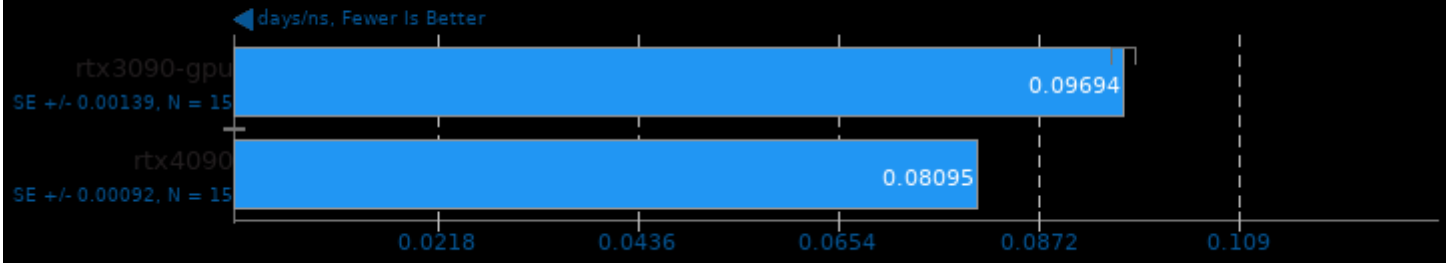


## ViennaCL 1.7.1

Test: OpenCL BLAS - dGEMV-N

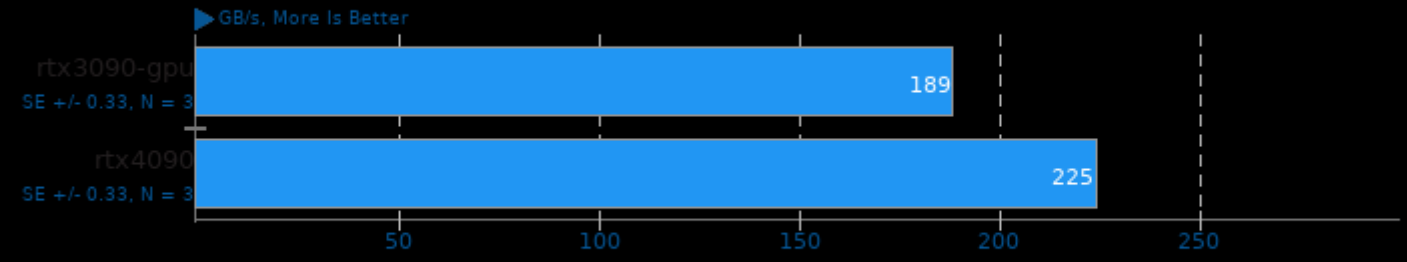


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## cl-mem 2017-01-13

Benchmark: Copy

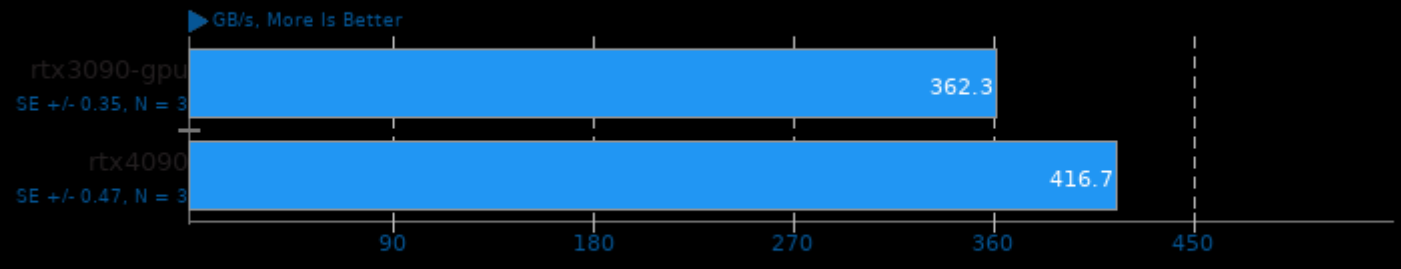


1. (CC) gcc options: -O2 -flto -lOpenCL

## ViennaCL 1.7.1

Test: CPU BLAS - sDOT

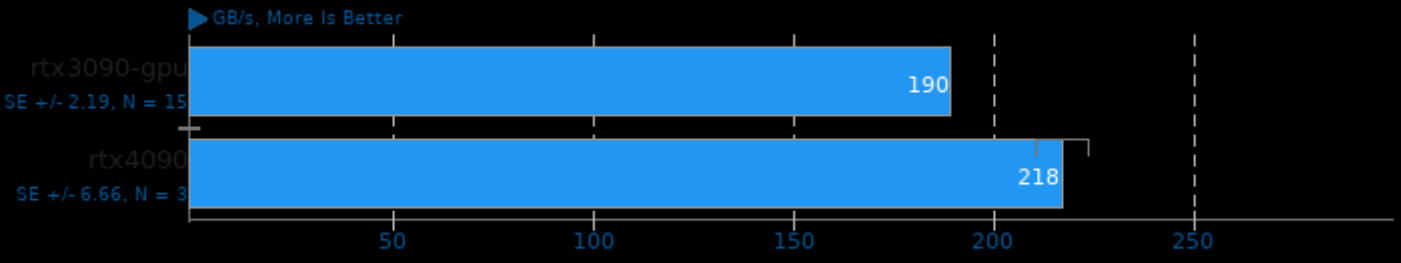


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## ViennaCL 1.7.1

Test: OpenCL BLAS - dDOT

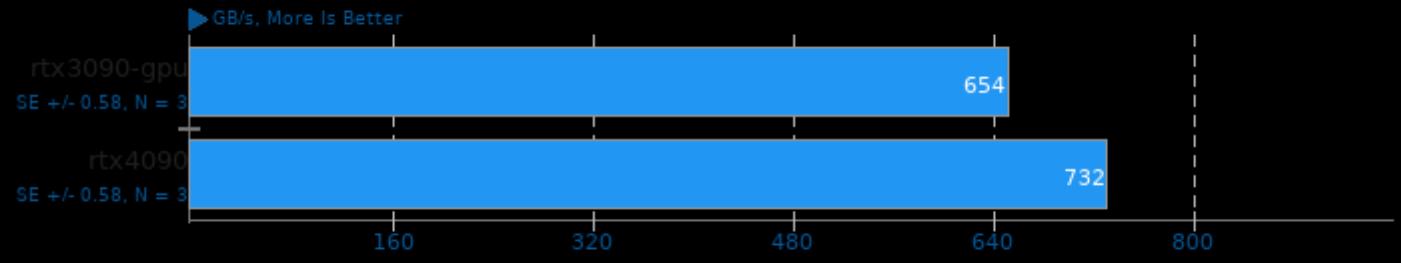


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## ViennaCL 1.7.1

Test: CPU BLAS - sAXPY

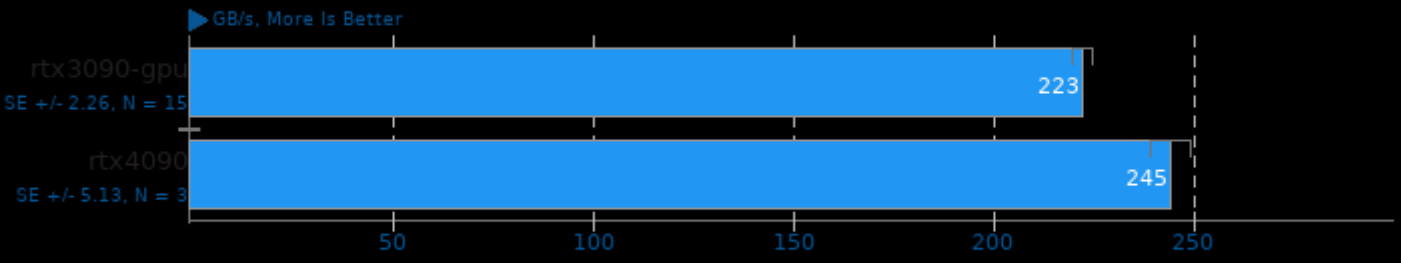


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

# rtx3090

## ViennaCL 1.7.1

Test: OpenCL BLAS - dCOPY

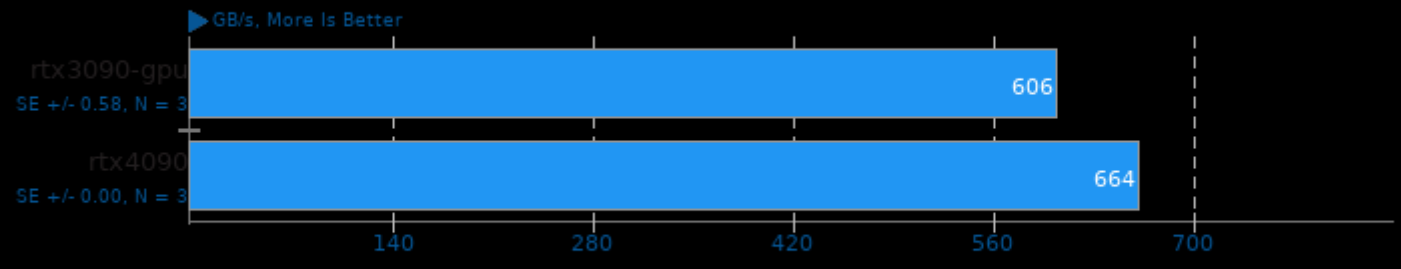


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## cl-mem 2017-01-13

Benchmark: Write

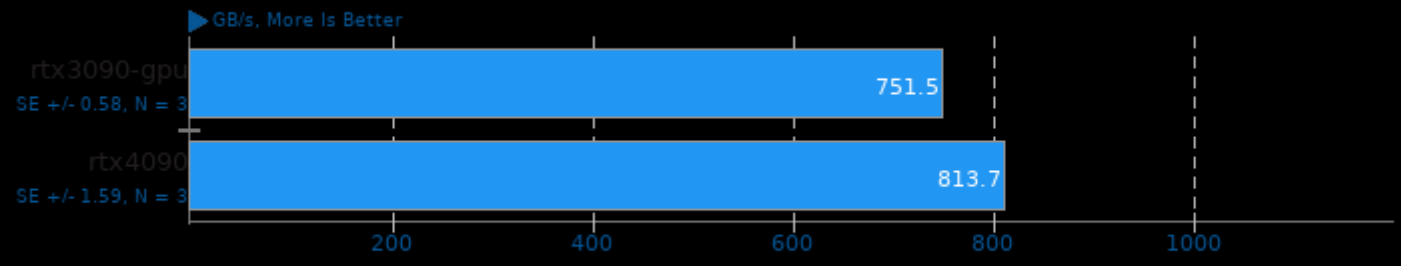


1. (CC) gcc options: -O2 -flto -lOpenCL

## cl-mem 2017-01-13

Benchmark: Read

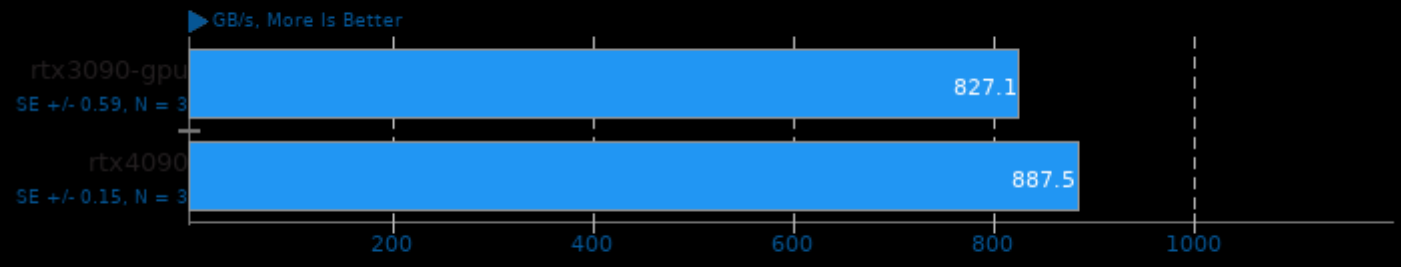


1. (CC) gcc options: -O2 -flto -lOpenCL

## ViennaCL 1.7.1

Test: OpenCL BLAS - dAXPY

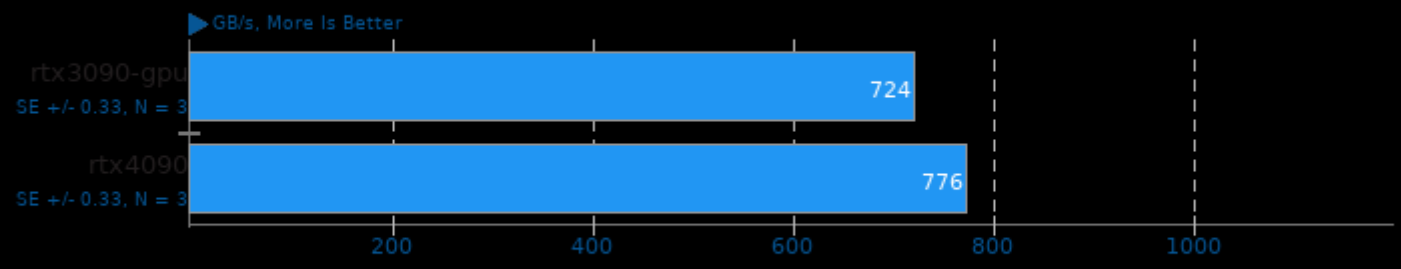


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

# rtx3090

## NCNN 20220729

Target: Vulkan GPU - Model: blazeface

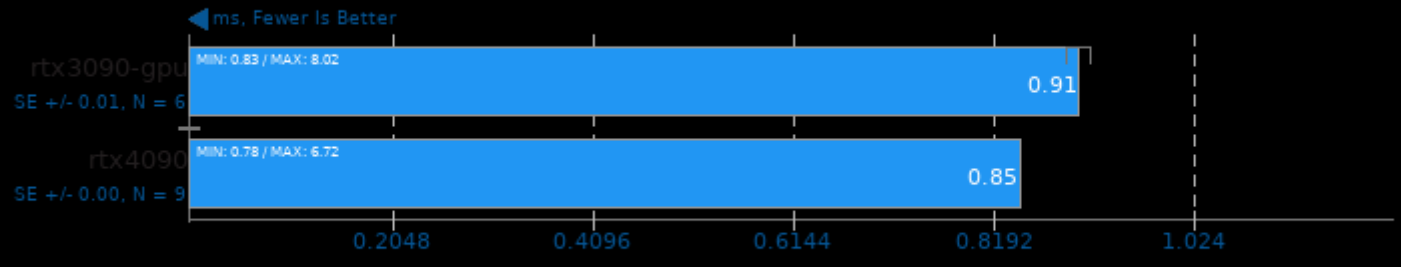


1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

## clpeak

OpenCL Test: Global Memory Bandwidth

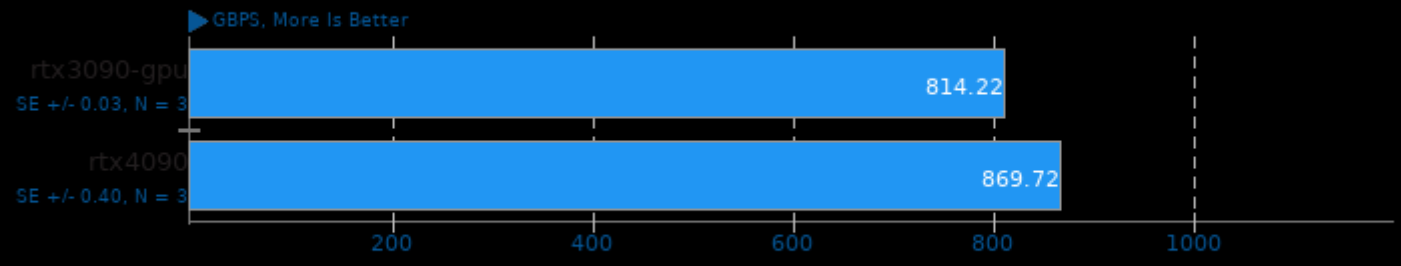


1. (CXX) g++ options: -O3 -rdynamic -lOpenCL

## ViennaCL 1.7.1

Test: CPU BLAS - sCOPY

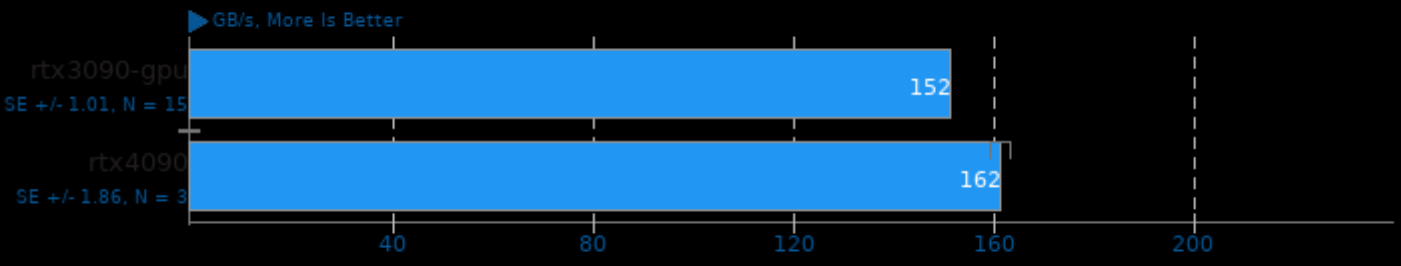


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## ViennaCL 1.7.1

Test: CPU BLAS - dDOT

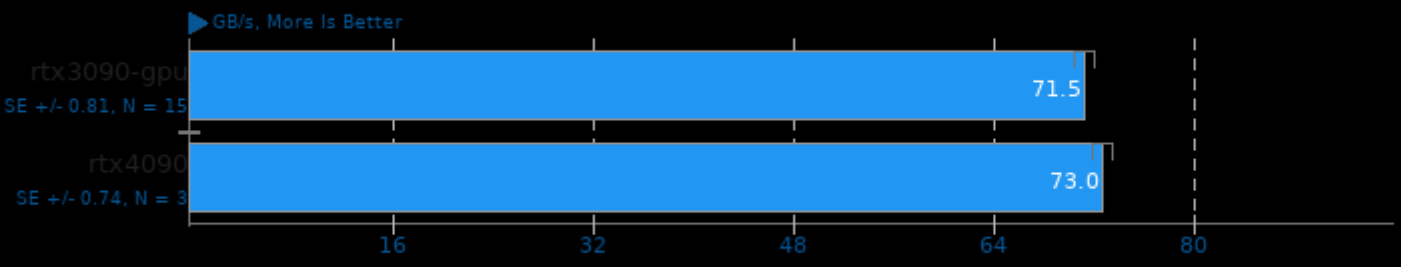


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

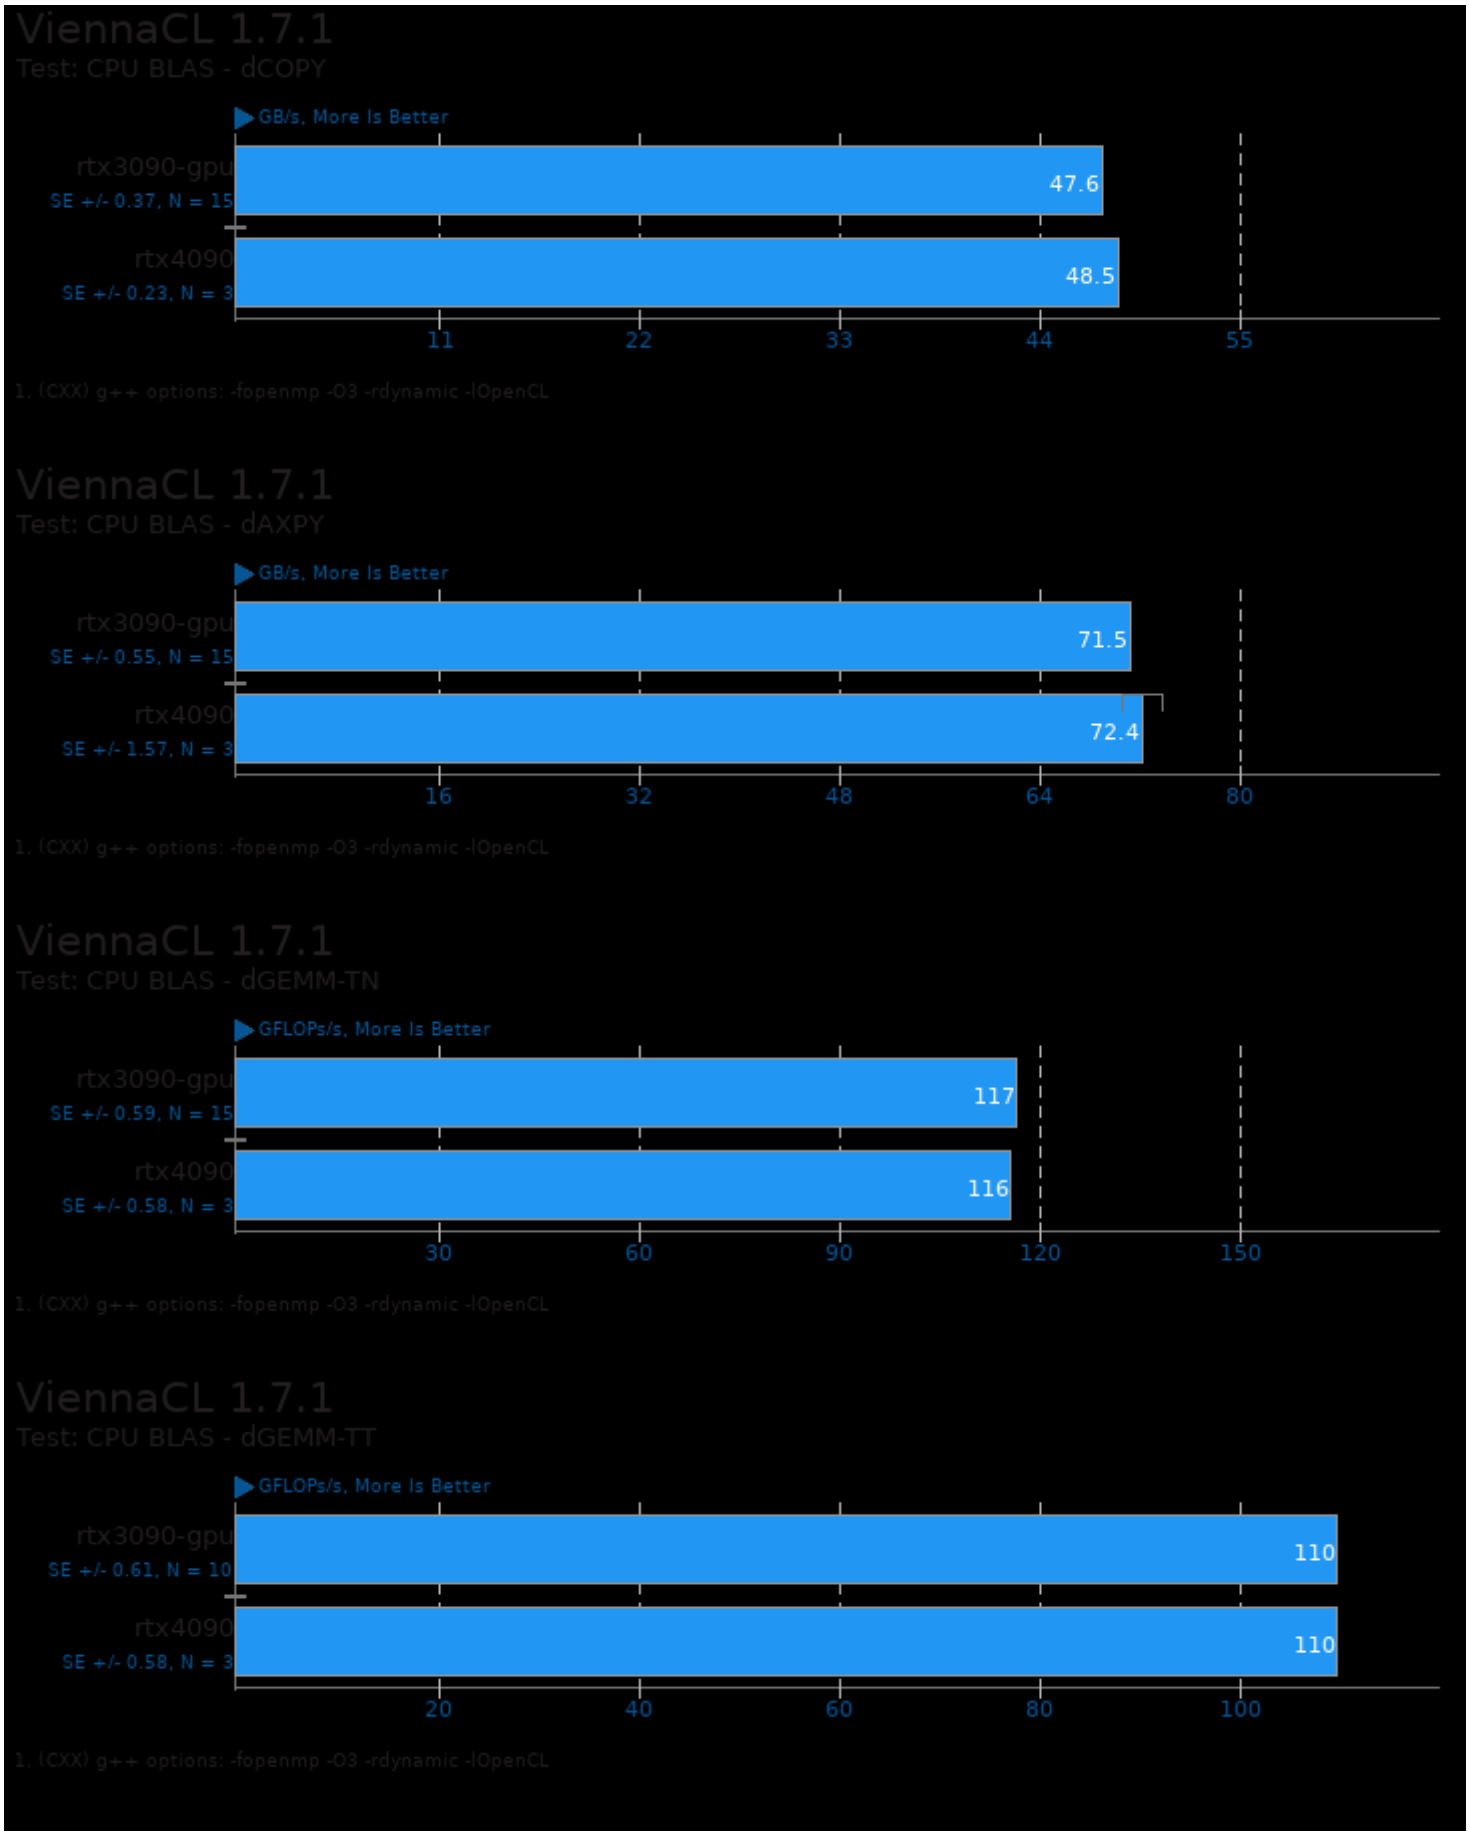
ViennaCL 1.7.1
Test: CPU BLAS - dCOPY
GB/s, More Is Better
rtx3090-gpu
SE +/- 0.37, N = 15
47.6
rtx4090
SE +/- 0.23, N = 3
48.5
11
22
33
44
55
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dAXPY
GB/s, More Is Better
rtx3090-gpu
SE +/- 0.55, N = 15
71.5
rtx4090
SE +/- 1.57, N = 3
72.4
16
32
48
64
80
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dGEMM-TN
GFLOPs/s, More Is Better
rtx3090-gpu
SE +/- 0.59, N = 15
117
rtx4090
SE +/- 0.58, N = 3
116
30
60
90
120
150
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dGEMM-TT
GFLOPs/s, More Is Better
rtx3090-gpu
SE +/- 0.61, N = 10
110
rtx4090
SE +/- 0.58, N = 3
110
20
40
60
80
100
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

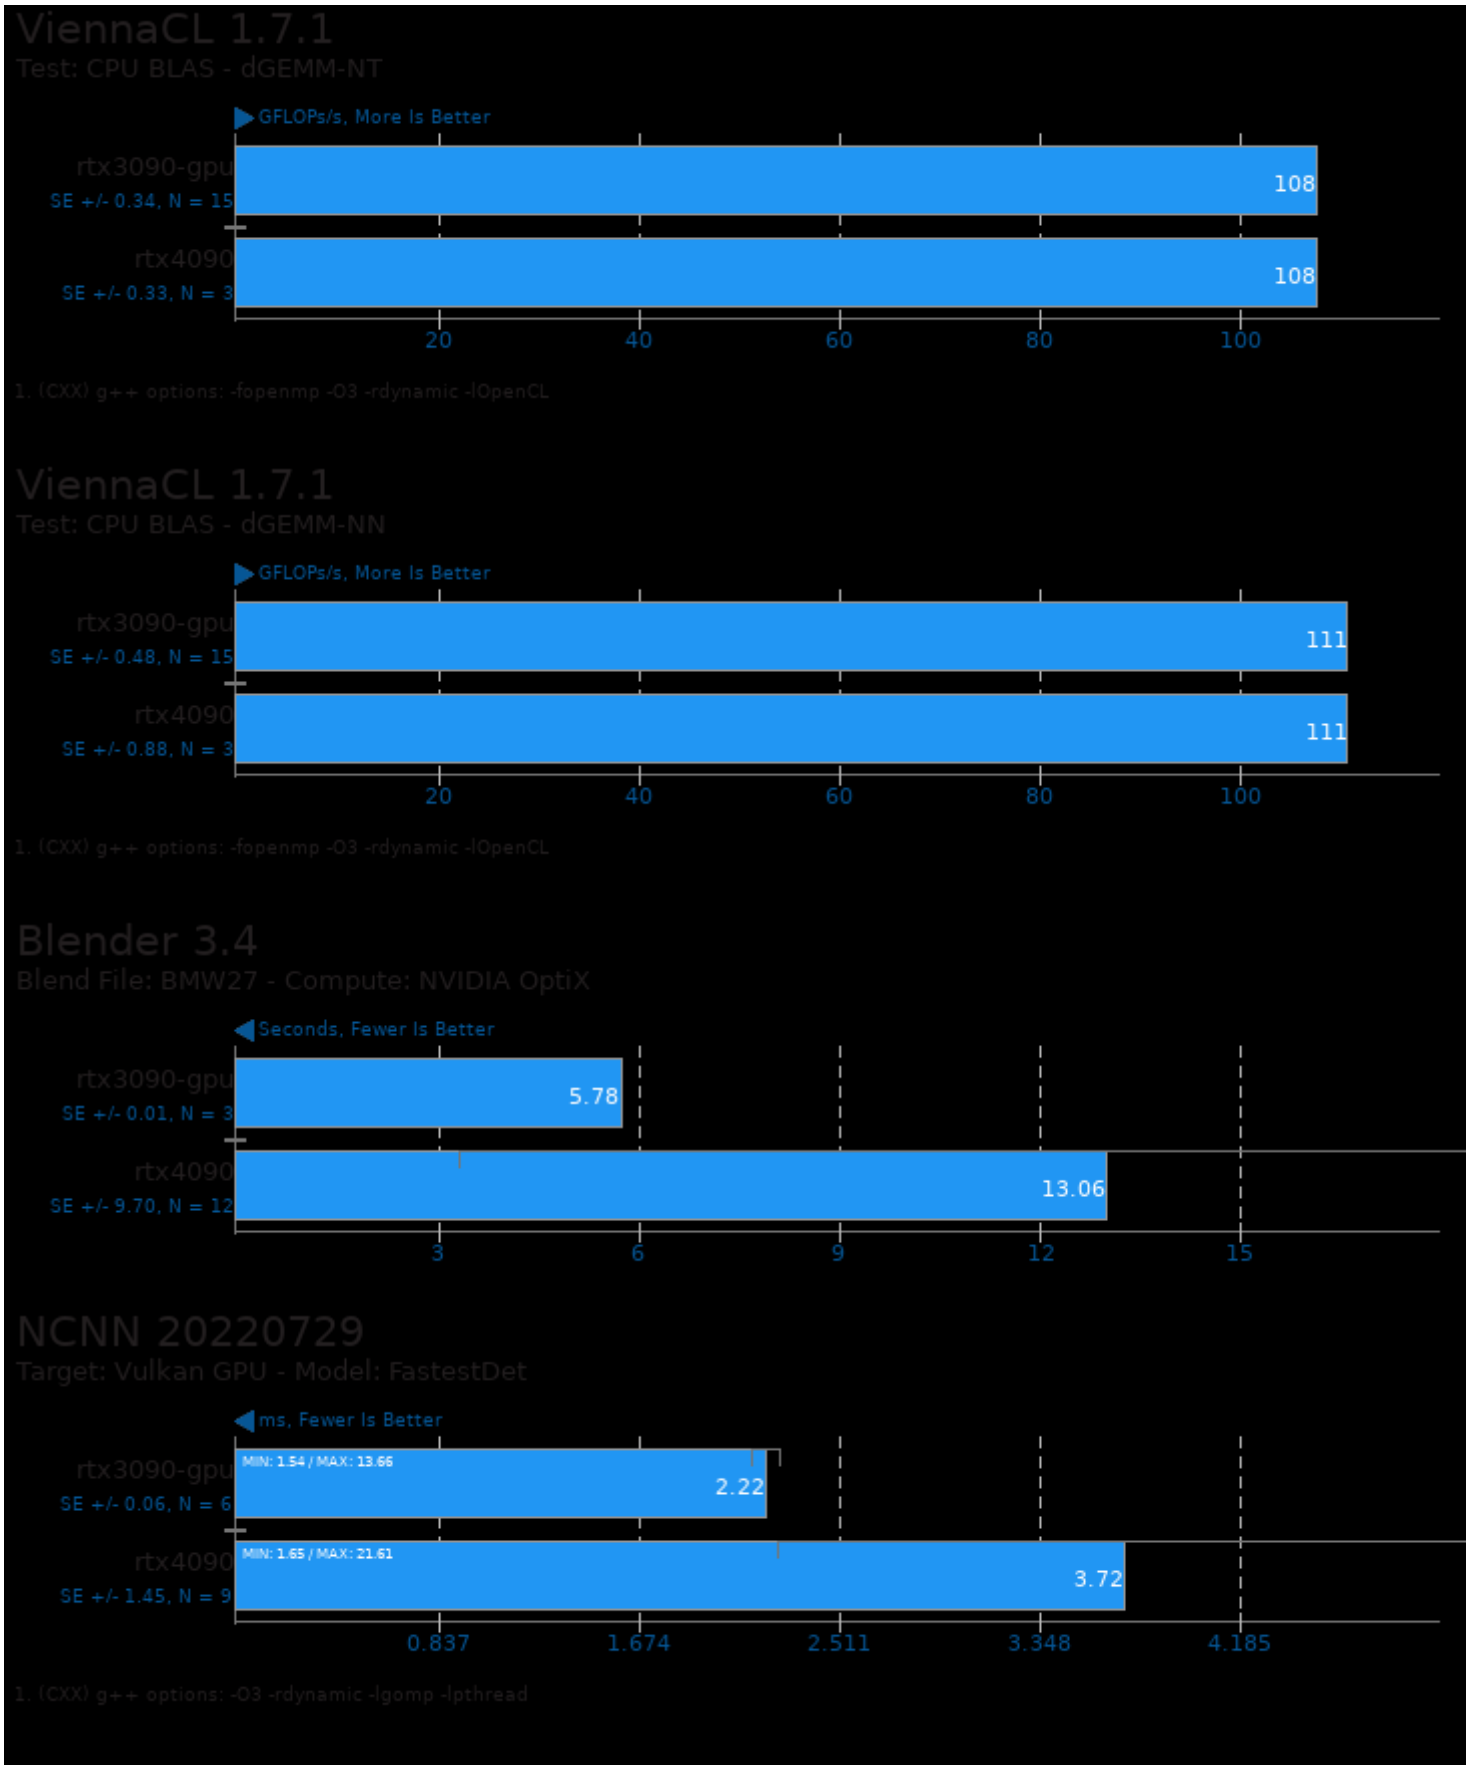
ViennaCL 1.7.1
Test: CPU BLAS - dGEMM-NT
GFLOPs/s, More Is Better
rtx3090-gpu
SE +/- 0.34, N = 15
108
rtx4090
SE +/- 0.33, N = 3
108
20
40
60
80
100
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
ViennaCL 1.7.1
Test: CPU BLAS - dGEMM-NN
GFLOPs/s, More Is Better
rtx3090-gpu
SE +/- 0.48, N = 15
111
rtx4090
SE +/- 0.88, N = 3
111
20
40
60
80
100
1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL
Blender 3.4
Blend File: BMW27 - Compute: NVIDIA OptiX
Seconds, Fewer Is Better
rtx3090-gpu
SE +/- 0.01, N = 3
5.78
rtx4090
SE +/- 9.70, N = 12
13.06
3
6
9
12
15
NCNN 20220729
Target: Vulkan GPU - Model: FastestDet
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.06, N = 6
MIN: 1.54 / MAX: 13.66
2.22
rtx4090
SE +/- 1.45, N = 9
MIN: 1.65 / MAX: 21.61
3.72
0.837
1.674
2.511
3.348
4.185
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

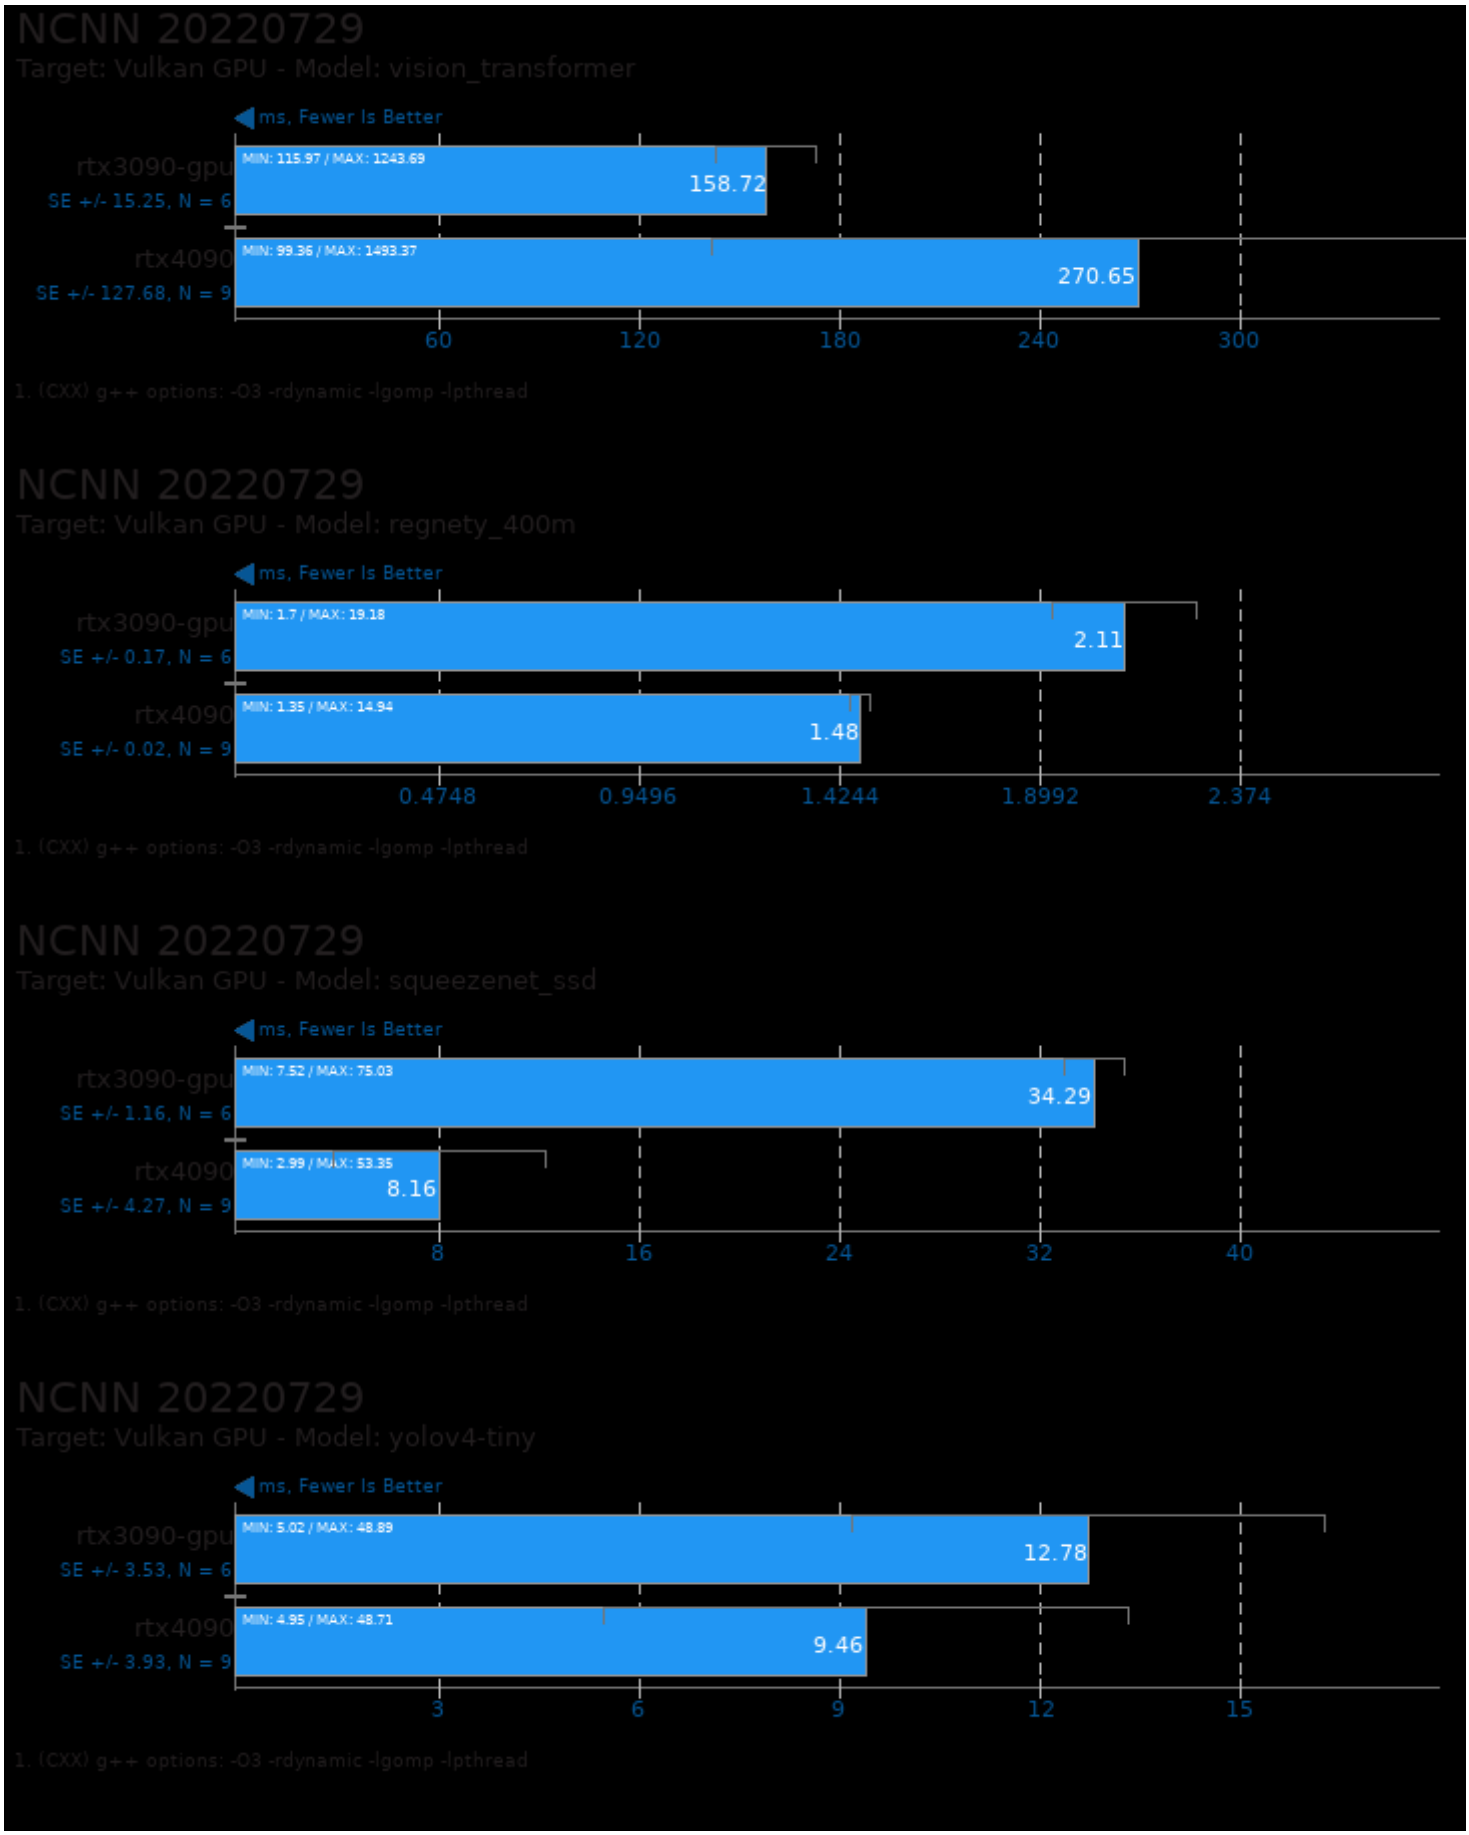
NCNN 20220729
Target: Vulkan GPU - Model: vision_transformer
ms, Fewer Is Better
rtx3090-gpu
SE +/- 15.25, N = 6
MIN: 115.97 / MAX: 1243.69
158.72
rtx4090
SE +/- 127.68, N = 9
MIN: 99.36 / MAX: 1493.37
270.65
60
120
180
240
300
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: regnety_400m
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.17, N = 6
MIN: 1.7 / MAX: 19.18
2.11
rtx4090
SE +/- 0.02, N = 9
MIN: 1.35 / MAX: 14.94
1.48
0.4748
0.9496
1.4244
1.8992
2.374
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: squeezenet_ssd
ms, Fewer Is Better
rtx3090-gpu
SE +/- 1.16, N = 6
MIN: 7.52 / MAX: 75.03
34.29
rtx4090
SE +/- 4.27, N = 9
MIN: 2.99 / MAX: 53.35
8.16
8
16
24
32
40
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: yolov4-tiny
ms, Fewer Is Better
rtx3090-gpu
SE +/- 3.53, N = 6
MIN: 5.02 / MAX: 48.89
12.78
rtx4090
SE +/- 3.93, N = 9
MIN: 4.95 / MAX: 48.71
9.46
3
6
9
12
15
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

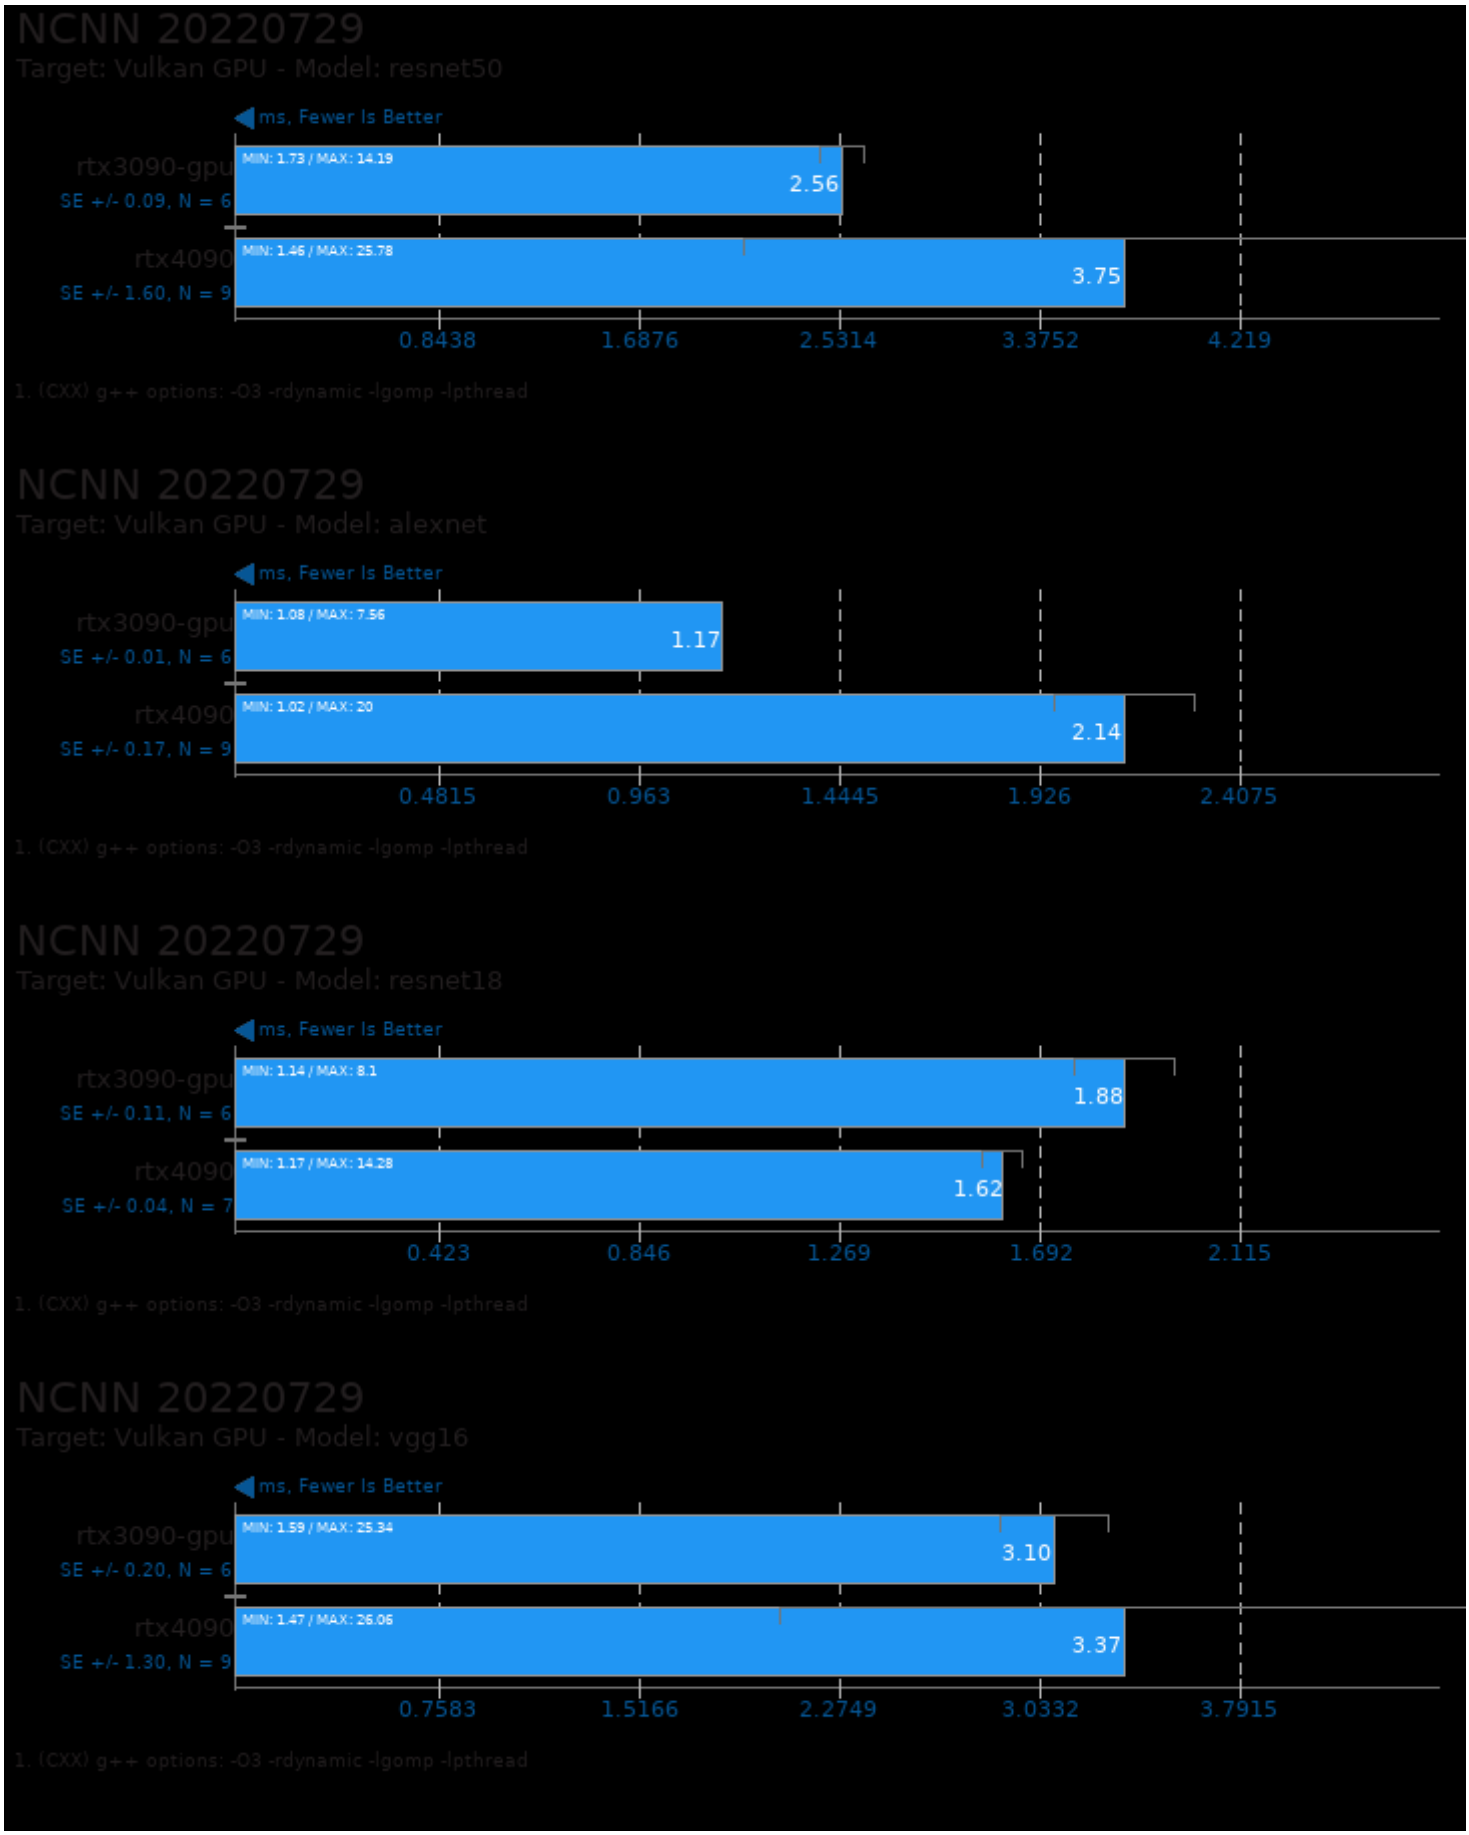
NCNN 20220729
Target: Vulkan GPU - Model: resnet50
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.09, N = 6
MIN: 1.73 / MAX: 14.19
2.56
rtx4090
SE +/- 1.60, N = 9
MIN: 1.46 / MAX: 25.78
3.75
0.8438
1.6876
2.5314
3.3752
4.219
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: alexnet
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.01, N = 6
MIN: 1.08 / MAX: 7.56
1.17
rtx4090
SE +/- 0.17, N = 9
MIN: 1.02 / MAX: 20
2.14
0.4815
0.963
1.4445
1.926
2.4075
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: resnet18
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.11, N = 6
MIN: 1.14 / MAX: 8.1
1.88
rtx4090
SE +/- 0.04, N = 7
MIN: 1.17 / MAX: 14.28
1.62
0.423
0.846
1.269
1.692
2.115
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: vgg16
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.20, N = 6
MIN: 1.59 / MAX: 25.34
3.10
rtx4090
SE +/- 1.30, N = 9
MIN: 1.47 / MAX: 26.06
3.37
0.7583
1.5166
2.2749
3.0332
3.7915
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread


NCNN 20220729
Target: Vulkan GPU - Model: efficientnet-b0
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.96, N = 6
MIN: 2.22 / MAX: 29.26
5.35
rtx4090
SE +/- 1.43, N = 9
MIN: 1.76 / MAX: 23.24
3.78
1.2038
2.4076
3.6114
4.8152
6.019
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: mnasnet
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.01, N = 6
MIN: 1.13 / MAX: 7.82
1.23
rtx4090
SE +/- 1.41, N = 9
MIN: 0.98 / MAX: 22.9
2.50
0.5625
1.125
1.6875
2.25
2.8125
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU - Model: shufflenet-v2
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.05, N = 6
MIN: 1.34 / MAX: 14.72
1.53
rtx4090
SE +/- 1.51, N = 9
MIN: 1.17 / MAX: 21.26
2.89
0.6503
1.3006
1.9509
2.6012
3.2515
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread
NCNN 20220729
Target: Vulkan GPU-v3-v3 - Model: mobilenet-v3
ms, Fewer Is Better
rtx3090-gpu
SE +/- 0.03, N = 6
MIN: 1.39 / MAX: 15.23
1.53
rtx4090
SE +/- 1.19, N = 9
MIN: 1.25 / MAX: 41.37
2.93
0.6593
1.3186
1.9779
2.6372
3.2965
1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

## NCNN 20220729

Target: Vulkan GPU-v2-v2 - Model: mobilenet-v2

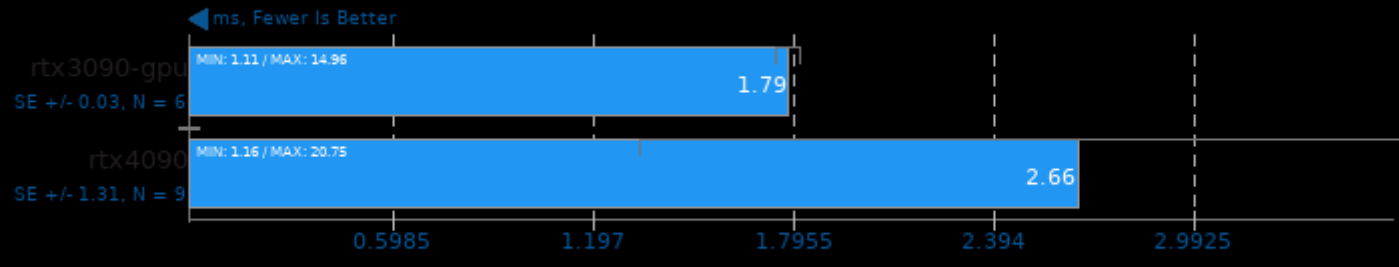


1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

## NCNN 20220729

Target: Vulkan GPU - Model: mobilenet

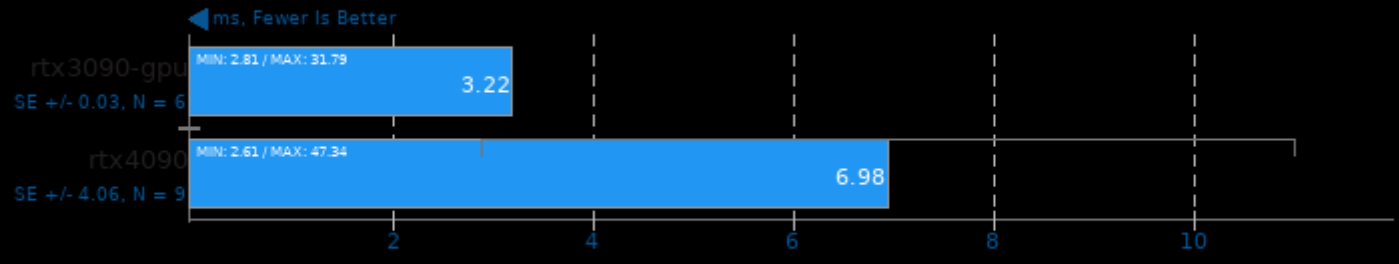


1. (CXX) g++ options: -O3 -rdynamic -lgomp -lpthread

## ViennaCL 1.7.1

Test: CPU BLAS - dGEMV-T

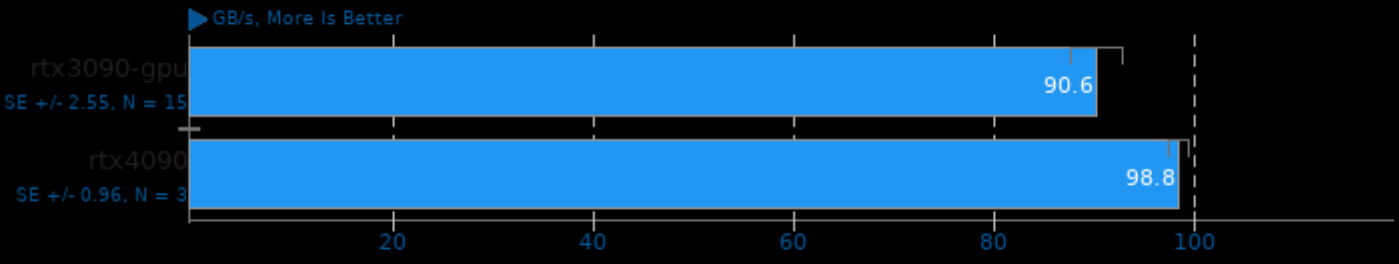


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

## ViennaCL 1.7.1

Test: CPU BLAS - dGEMV-N

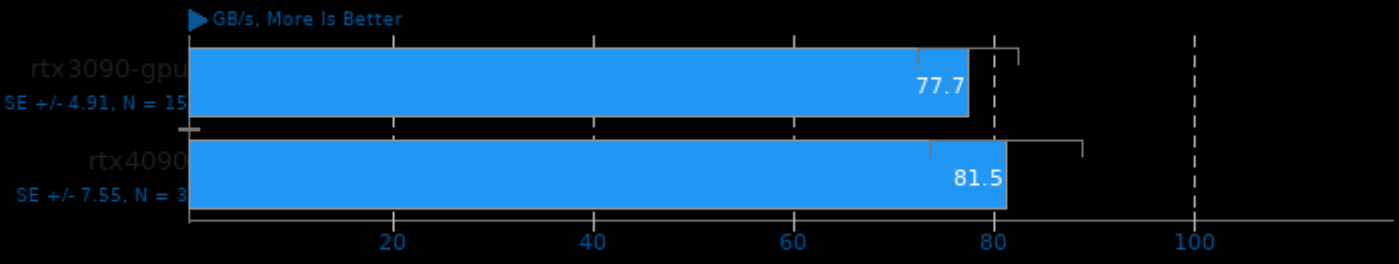


1. (CXX) g++ options: -fopenmp -O3 -rdynamic -lOpenCL

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

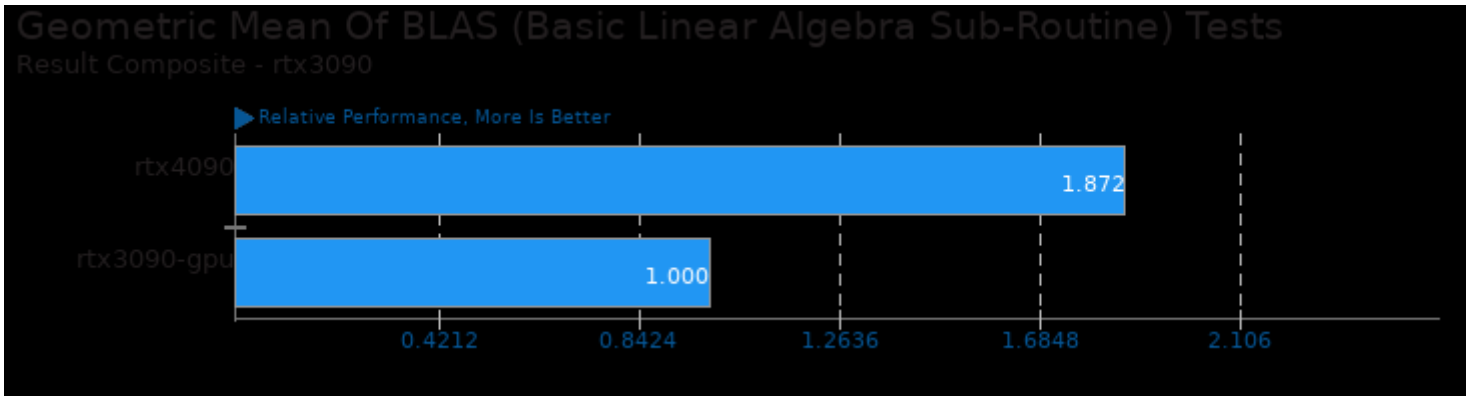


Geometric mean based upon tests: pts/arrayfire, pts/lczero and pts/caffe

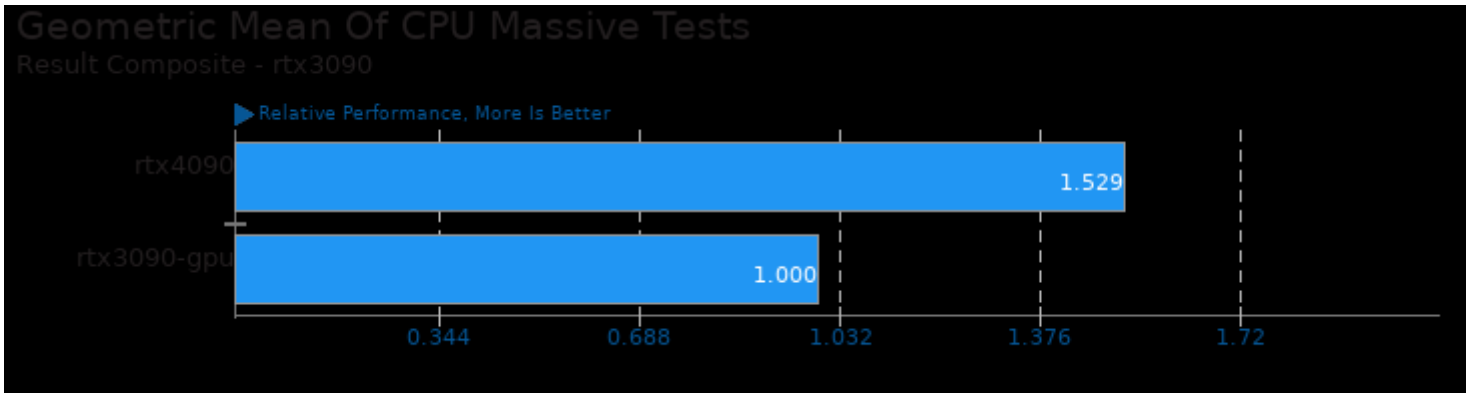


Geometric mean based upon tests: pts/lczero, pts/plaidml, pts/rodinia and pts/blender

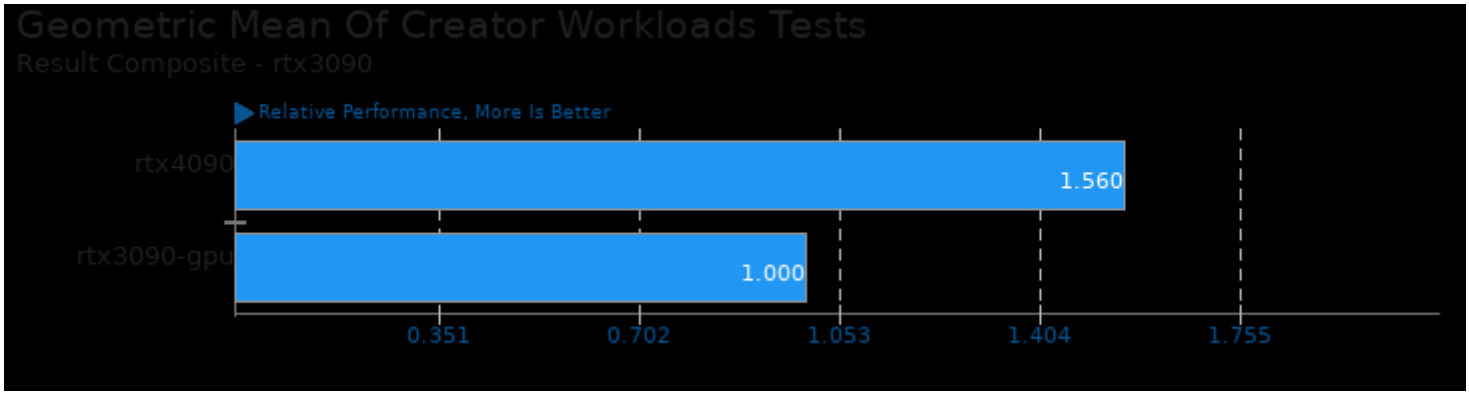


Geometric mean based upon tests: pts/blender, pts/luxcorerender, pts/indigobench, pts/neatbench and pts/betsy

Geometric mean based upon tests: pts/betsy and pts/blender

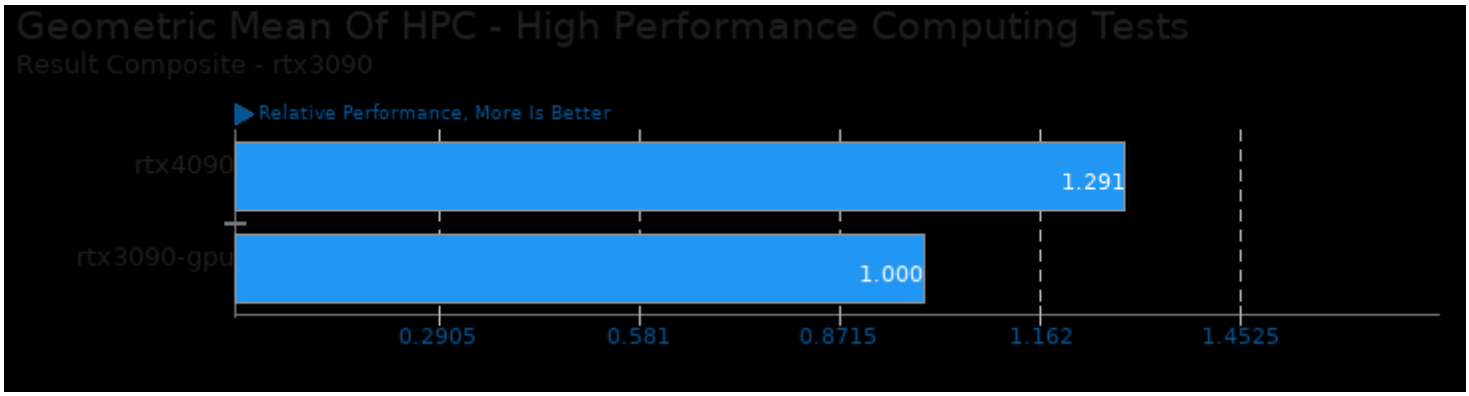


Geometric mean based upon tests: pts/rodinia, pts/arrayfire, pts/gromacs, pts/ncnn, pts/caffe, pts/shoc, pts/plaidml and pts/lczero

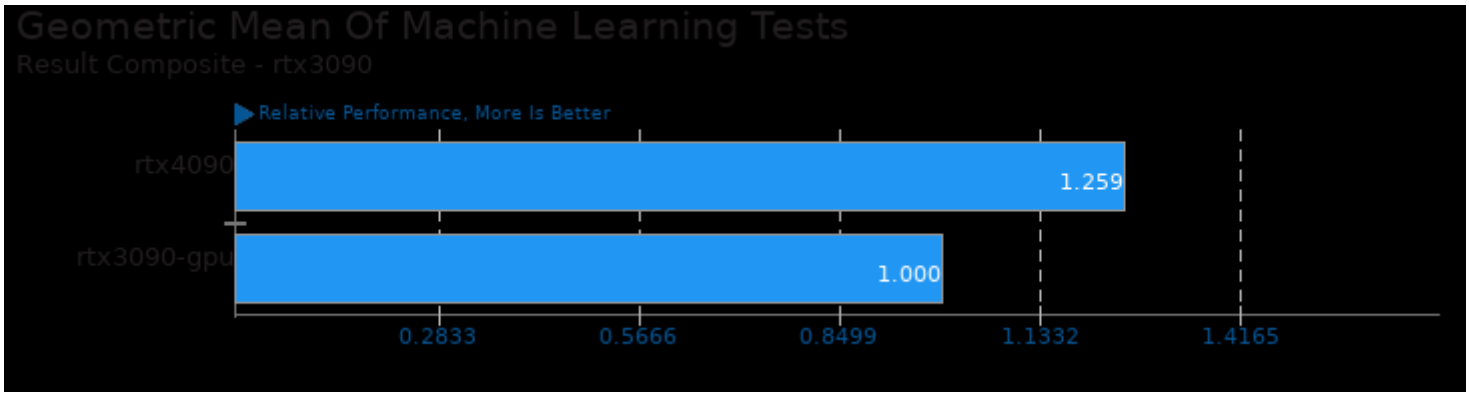


Geometric mean based upon tests: pts/ncnn, pts/caffe, pts/shoc, pts/plaidml and pts/lczero

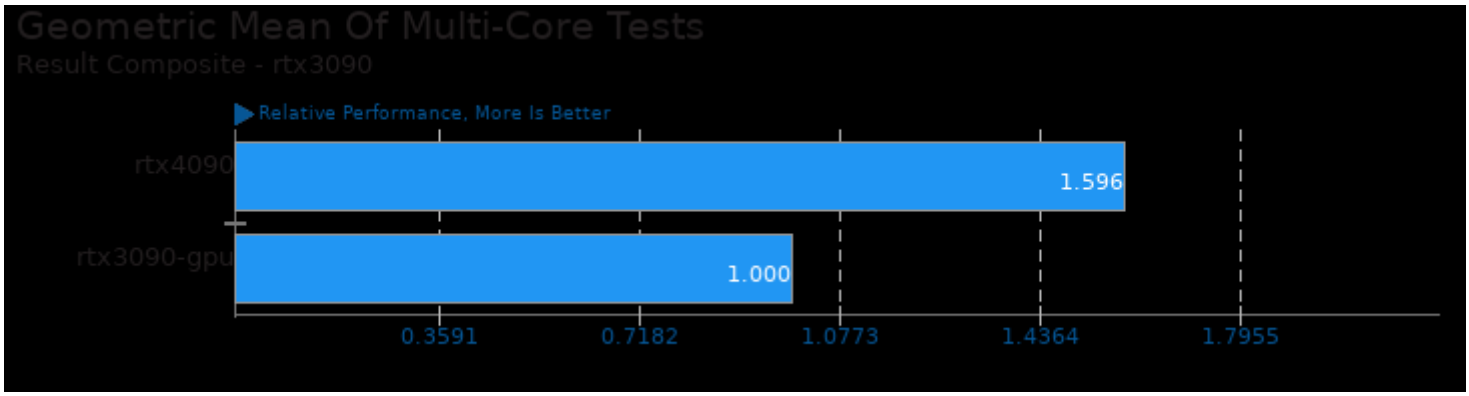


Geometric mean based upon tests: pts/blender, pts/arrayfire, pts/rodinia, pts/gromacs, pts/luxcorerender, pts/indigobench and pts/neatbench

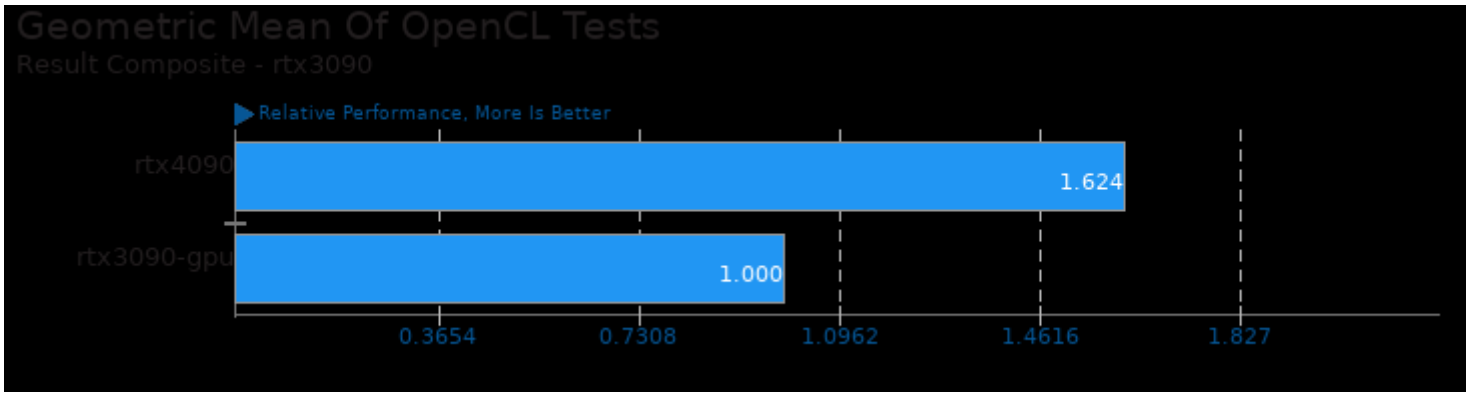


Geometric mean based upon tests: pts/mandelgpu, pts/shoc, pts/cl-mem, pts/clpeak, pts/rodinia and pts/viennacl

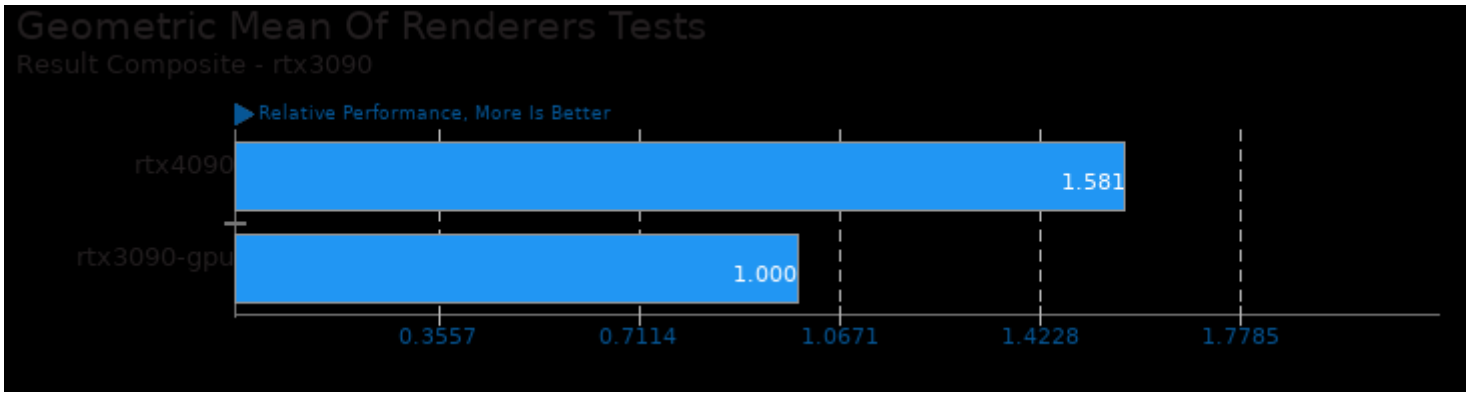


Geometric mean based upon tests: pts/blender, pts/luxcorerender and pts/indigobench

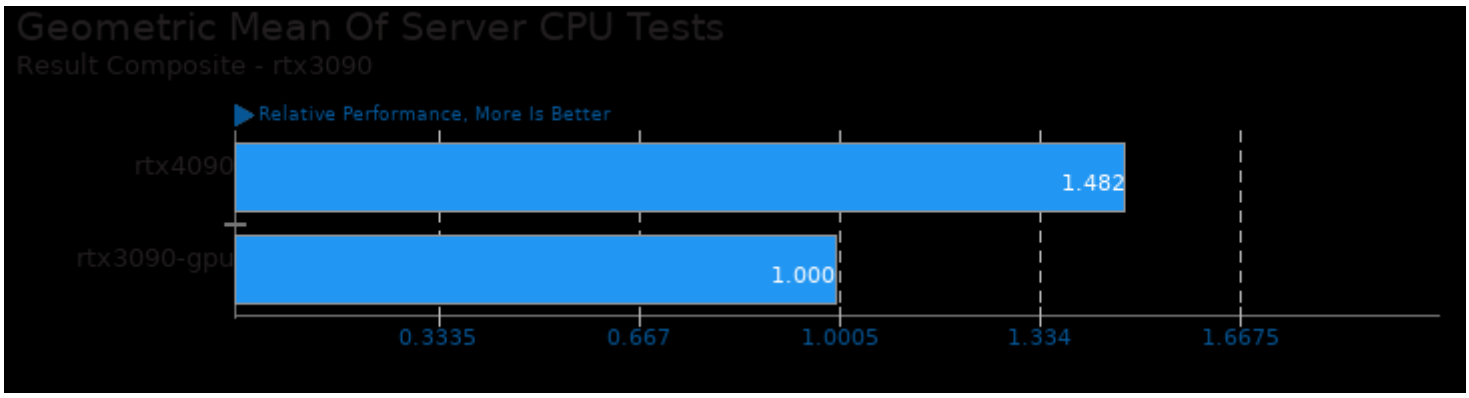


Geometric mean based upon tests: pts/rodinia and pts/blender

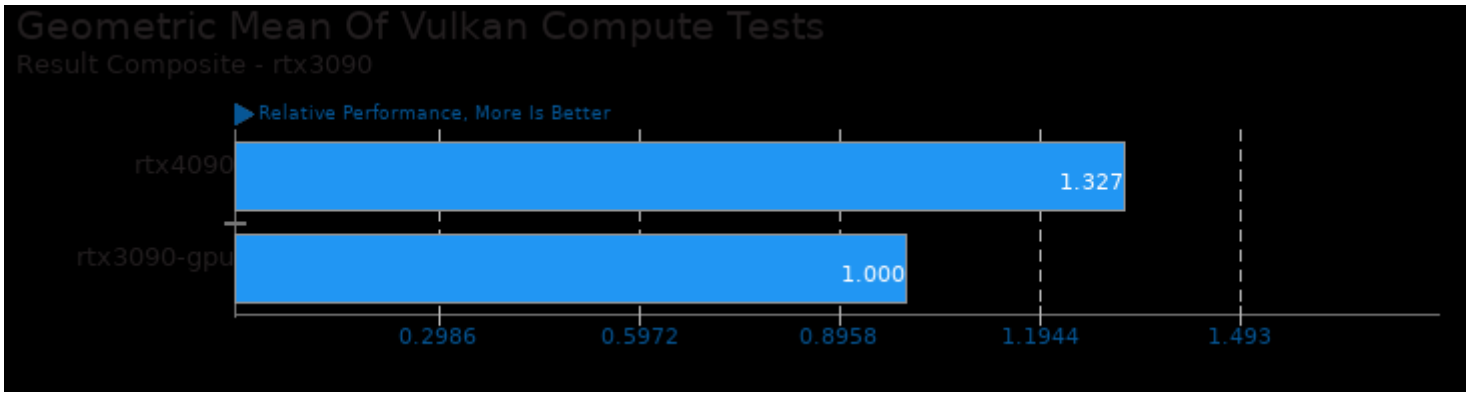


Geometric mean based upon tests: pts/vkfft, pts/vkresample, pts/vkpeak, pts/ncnn, pts/realsr-ncnn, pts/waifu2x-ncnn, pts/betsy and pts/libplacebo

# rtx3090

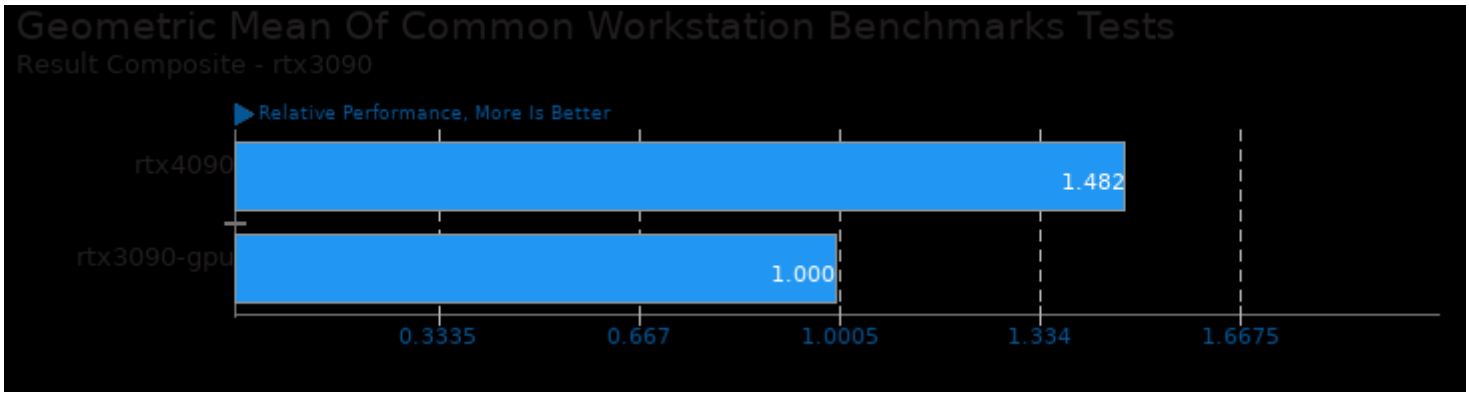


Geometric mean based upon tests: pts/blender and pts/rodinia

*This file was automatically generated via the Phoronix Test Suite benchmarking software on Thursday, 28 March 2024 16:22.*