



# NVIDIA Linux OpenCL CUDA RTX SUPER Compute

NVIDIA GeForce RTX SUPER Linux OpenCL/CUDA GPU compute benchmarks by Michael Larabel for a future article.

## Automated Executive Summary

*TITAN RTX had the most wins, coming in first place for 89% of the tests.*

*Based on the geometric mean of all complete results, the fastest (TITAN RTX) was 3.565x the speed of the slowest (GTX 970).*

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

*LeelaChessZero (Backend: OpenCL) at 15.043x*

*clpeak (OpenCL Test: Integer Compute INT) at 14.582x*

*LeelaChessZero (Backend: OpenCL) at 12.548x*

*clpeak (OpenCL Test: Integer Compute INT) at 12.138x*

*SHOC Scalable HeterOgeneous Computing (Target: OpenCL - Benchmark: MD5 Hash) at 9.25x*

*NAMD CUDA (System Power Consumption Monitor) at 8.463x*

*SHOC Scalable HeterOgeneous Computing (System Power Consumption Monitor) at 8.247x*

*LeelaChessZero (System Power Consumption Monitor) at 8.163x*

*LuxMark (System Power Consumption Monitor) at 8.118x*

*PlaidML (System Power Consumption Monitor) at 8.02x.*

## Test Systems:

### GTX 970

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: eVGA NVIDIA GeForce GTX 970 4GB (1163/3505MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 1664

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swaps barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of Full generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

### GTX 980

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 980 4GB (1126/3505MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 2048

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swaps barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of Full generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

### GTX 980 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 980 Ti 6GB (999/3505MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.3.0-999-generic (x86\_64) 20190914, Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 2816

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swaps barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of Full generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## GTX 1060

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 1060 6GB (1506/4006MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 1280

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swaps barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of Full generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## GTX 1070

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 1070 8GB (1506/4006MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 1920

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swaps barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of Full generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## GTX 1070 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: Zotac NVIDIA GeForce GTX 1070 Ti 8GB (1607/4006MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 2432

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## GTX 1080

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 1080 8GB (1607/5005MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 2560

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## GTX 1080 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 1080 Ti 11GB (1480/5508MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 3584

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## GTX 1660

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: ASUS NVIDIA GeForce GTX 1660 6GB (1530/4001MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 1408

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## GTX 1660 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: eVGA NVIDIA GeForce GTX 1660 Ti 6GB (1500/6000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 1536

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## RTX 2060

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce RTX 2060 6GB (1365/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 1920

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## RTX 2060 SUPER

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce RTX 2060 SUPER 8GB (1470/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.3.0-999-generic (x86\_64) 20190914, Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 2176

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## RTX 2070

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: ASUS NVIDIA GeForce RTX 2070 8GB (1410/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 2304

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## RTX 2070 SUPER

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce RTX 2070 SUPER 8GB (1605/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.3.0-999-generic (x86\_64) 20190914, Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 2560

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

## RTX 2080

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS), Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: Zotac NVIDIA GeForce RTX 2080 8GB (795/810MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes

--enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v  
Processor Notes: Scaling Governor: intel\_pstate performance  
OpenCL Notes: GPU Compute Cores: 2944  
Python Notes: Python 2.7.16 + Python 3.7.3  
Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW  
STIBP: conditional RSB filling

## RTX 2080 SUPER

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS),  
Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA  
GeForce RTX 2080 SUPER 8GB (1650/7750MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel  
I219-V

OS: Ubuntu 19.04, Kernel: 5.3.0-999-generic (x86\_64) 20190914, Desktop: GNOME Shell 3.32.2, Display Server: X  
Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen  
Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes  
--enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v  
Processor Notes: Scaling Governor: intel\_pstate performance  
OpenCL Notes: GPU Compute Cores: 3072  
Python Notes: Python 2.7.16 + Python 3.7.3  
Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW  
STIBP: conditional RSB filling

## RTX 2080 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS),  
Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA  
GeForce RTX 2080 Ti 11GB (1350/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.0.0-29-generic (x86\_64), Desktop: GNOME Shell 3.32.2, Display Server: X Server 1.20.4,  
Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes  
--enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v  
Processor Notes: Scaling Governor: intel\_pstate performance  
OpenCL Notes: GPU Compute Cores: 4352  
Python Notes: Python 2.7.16 + Python 3.7.3  
Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + 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 generic retpoline IBPB: conditional IBRS\_FW  
STIBP: conditional RSB filling

## TITAN RTX

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0802 BIOS),  
Chipset: Intel Cannon Lake PCH, Memory: 16384MB, Disk: Samsung SSD 970 EVO 250GB, Graphics: NVIDIA TITAN  
RTX 24GB (1350/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel I219-V

OS: Ubuntu 19.04, Kernel: 5.3.0-999-generic (x86\_64) 20190914, Desktop: GNOME Shell 3.32.2, Display Server: X  
Server 1.20.4, Display Driver: NVIDIA 435.21, OpenGL: 4.6.0, Compiler: GCC 8.3.0, File-System: ext4, Screen

Resolution: 3840x2160

Compiler Notes: --build=x86\_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel\_pstate performance

OpenCL Notes: GPU Compute Cores: 4608

Python Notes: Python 2.7.16 + Python 3.7.3

Security Notes: I1tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swaps barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of Full generic retpoline IBPB: conditional IBRS\_FW STIBP: conditional RSB filling

	GTX 970	GTX 980	GTX 980 Ti	GTX 1060	GTX 1070	GTX 1070 Ti	GTX 1080	GTX 1080 Ti	GTX 1660	GTX 1660 Ti	RTX 2060	RTX 2060 SUP ER	RTX 2070	RTX 2070 SUP ER	RTX 2080	RTX 2080 SUP ER	RTX 2080 Ti	TITA N RTX
<b>SHOC</b>	0.05	0.08	0.04	0.11			0.15		0.26	0.26	0.12		0.15	0.30	0.21	0.32		0.37
<b>Scalable Heterogeneous Computing - OpenCL - MD5 Hash (GHash/s/Watt)</b>																		
<b>Normalized</b>	13.51	21.62	10.81	29.73			40.54		70.27	70.27	32.43		40.54	81.08	56.76	86.49		100%
<b>clpeak - S.P.F (GFLOPS/Watt)</b>	39.36	36.49	42.00		70.63	84.46	89.11			71.35	109.2	111.7	74.32	121.6	79.25	84.49	222.09	103.61
<b>Normalized</b>	17.72	16.43	18.91		31.8%	38.03	40.12			32.13	49.17	50.31	33.46	54.76	35.68	38.04	100%	46.65
<b>clpeak - G.M.B (GBPS/Watt)</b>	1.52	1.65	2.02	1.84	2.63	1.92	2.26	2.73	3.16	2.24		4.88	3.33	5.41	5.09	4.86	9.15	6.18
<b>Normalized</b>	16.61	18.03	22.08	20.11	28.74	20.98	24.7%	29.84	34.54	24.48		53.33	36.39	59.13	55.63	53.11	100%	67.54
	%	%	%	%	%	%		%	%	%		%	%	%	%	%		%



SHOC Scalable Heterogeneous Computing - OpenCL - MD5 Hash (GHash/s)	6.50	7.49	9.22	7.26	10.61	11.70	14.25	19.72	11.87	12.98	16.04	18.22	18.77	21.78	24.07	26.56	34.91	36.39
Normalized	17.86	20.58	25.34	19.95	29.16	32.15	39.16	54.19	32.62	35.67	44.08	50.07	51.58	59.85	66.14	72.99	95.93	100%
Standard	0%	0%	0%	0%	0.3%	0%	0.3%	0.4%	0%	0.1%	0.3%	0.2%	0.3%	0.4%	0.1%	0.1%	0.6%	0.3%
Deviation																		
SHOC Scalable Heterogeneous Computing - OpenCL - FFT SP (GFLOPS)	411.2	458.3	725.9	326.6	478.5	525.8	614.6	982.2	452.8	665.6	818.4	971.2	976.7	1088	1100	1204	1486	1578
Normalized	9	5	8	8	4	9	9	3	9	6	4	4	8					
Standard	0.2%	0.4%	0.7%	2.8%	2.7%	0.1%	1%	0.2%	0.1%	0.4%	0.1%	0.6%	0.5%	0%	0.1%	0.1%	0.3%	0.2%
Deviation																		
PlaidML - No Inference - ResNet 50 - OpenCL (FPS/Watt)	1.05	1.15	1.06	1.48	1.80	2.10	2.25	2.66	1.90	1.84	2.16	2.59	3.09	2.89	3.21	3.18	3.70	5.03
Normalized	20.87	22.86	21.07	29.42	35.79	41.75	44.73	52.88	37.77	36.58	42.94	51.49	61.43	57.46	63.82	63.22	73.56	100%
Standard	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
PlaidML - No Inference - Mobilenet - OpenCL (FPS/Watt)	5.20	5.53	6.64	7.59	9.65	11.30	9.35	10.22	9.56	11.20	11.67	12.98	14.84	19.70	14.16	16.60	22.63	23.23
Normalized	22.38	23.81	28.58	32.67	41.54	48.64	40.25	43.99	41.15	48.21	50.24	55.88	63.88	84.8%	60.96	71.46	97.42	100%
Standard	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%

SHOC	5.12	5.45	6.92	4.00	6.23	6.36	7.87	7.97	5.70	12.25	13.23	17.59	8.67	15.58	15.46		11.54
Scalable Heterogeneous Computing - OpenCL - FFT SP (GFLOPS/Watt)																	
Normalized	29.11	30.98	39.34	22.74	35.42	36.16	44.74	45.31	32.4%	69.64	75.21	100%	49.29	88.57	87.89		65.61
LuxMark - GPU -	7985	8914	1109	7086	9983	1025	8716	1372	9794	1100	1375	1850	1844	2049	1991	2016	2863
Microphone (Score)			6			9		1		3	4	9	1	9	6	4	1
Normalized	25.96	28.98	36.08	23.04	32.46	33.36	28.34	44.61	31.84	35.78	44.72	60.18	59.96	66.65	64.75	65.56	93.09
Standard	0%	0.4%	0.8%	0%	0.1%	0.3%	0%	0.2%	0%	0%	0%	0.4%	0.2%	0.1%	0.1%	0.9%	0.1%
Deviation																	
Darktable - Server Room - OpenCL (sec)	3.16	3.10	1.51	1.28	1.11	1.17	1.09	1.05	1.18	1.15	0.83	0.75	0.75	0.80	0.80	0.79	0.75
Normalized	23.42	23.87	49.01	57.81	66.67	63.25	67.89	70.48	62.71	64.35	89.16	98.67	98.67	92.5%	92.5%	93.67	98.67
Standard	0.8%	0.2%	0.4%	0.9%	1.3%	0.7%	1.2%	2.6%	0.1%	0.2%	0.6%	0.5%	0.3%	0.6%	0.5%	0.5%	0.3%
Deviation																	
PlaidML - Yes - Inference - Inception V3 - OpenCL	88.16	101.9	116.3	97.83	141.5	147.0	173.7	228.6	137.6	161.3	193.7	219.7	220.9	239.6	268.3	301.1	363.1
		2	1		5	4	9	8	6	3	5	6	9	1	3	1	6
Normalized	23.59	27.27	31.12	26.18	37.88	39.35	46.5%	61.19	36.84	43.17	51.85	58.8%	59.13	64.12	71.8%	80.57	97.18
Standard	0%	0.5%	0.4%	0.3%	0.2%	0.1%	0.3%	0.2%	0.3%	0.1%	0.3%	0.3%	0.3%	0.1%	0.4%	0.1%	0.2%
Deviation																	
PlaidML - Yes - Inference - VGG19 - OpenCL (FPS)	39.34	43.63	56.13	44.37	63.56	73.12	83.83	115.0	52.52	58.03	71.07	80.82	83.56	97.82	105.5	115.4	154.2
								3							4	4	0
Normalized	24.28	26.92	34.64	27.38	39.22	45.12	51.73	70.98	32.41	35.81	43.86	49.87	51.56	60.36	65.13	71.24	95.16
Standard	0.8%	2.9%	2.4%	0.5%	0.5%	0.1%	0.6%	0.7%	0.3%	0.3%	0.5%	0.5%	0.7%	0.3%	0.7%	0.3%	0.4%
Deviation																	

PlaidML -	47.70	53.49	68.47	54.78	76.89	90.34	101.0	138.6	67.24	73.35	88.04	101.1	103.8	122.3	129.5	143.0	186.1	196.4
No -							0	8				7	9	2	8	2	9	4
Inference -																		
VGG19 -																		
OpenCL																		
(FPS)																		
Normalized	24.28	27.23	34.86	27.89	39.14	45.99	51.42	70.6%	34.23	37.34	44.82	51.5%	52.89	62.27	65.96	72.81	94.78	100%
Standard	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%
Deviation	0%	1.5%	1.6%	0.3%	0.2%	0.1%	0.3%	0.4%	0.4%	0.3%	0.3%	0.3%	0.4%	0.3%	0.6%	0.3%	0.4%	0.5%
PlaidML -	49.76	55.70	71.26	56.07	80.27	92.03	106.0	145.4	66.24	73.18	89.55	102.5	105.6	123.2	133.2	145.6	194.8	204.6
Yes -							9	2				8	0	0	0	5	3	7
Inference -																		
VGG16 -																		
OpenCL																		
(FPS)																		
Normalized	24.31	27.21	34.82	27.4%	39.22	44.97	51.83	71.05	32.36	35.76	43.75	50.12	51.6%	60.19	65.08	71.16	95.19	100%
Standard	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%
Deviation	0.7%	3%	2.7%	0.6%	0.6%	0.2%	0.7%	0.7%	0.5%	0.3%	0.4%	0.6%	0.6%	0.4%	0.7%	0.3%	0.5%	0.5%
PlaidML -	60.32	68.95	87.39	69.31	97.32	114.0	128.0	175.3	84.83	92.45	110.9	128.2	131.5	154.6	164.0	180.4	234.6	247.8
No -							3	5	7			3	2	4	4	5	5	3
Inference -																		
VGG16 -																		
OpenCL																		
(FPS)																		
Normalized	24.34	27.82	35.26	27.97	39.27	46.01	51.67	70.76	34.23	37.3%	44.76	51.74	53.08	62.4%	66.19	72.81	94.68	100%
Standard	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%
Deviation	0.1%	2%	2.7%	0.3%	0.3%	0.1%	0.5%	0.6%	0.4%	0.5%	0.5%	0.5%	0.6%	0.3%	0.8%	0.3%	0.4%	0.3%
PlaidML -	21.42	24.24	30.64	23.16	33.07	35.41	40.38	56.64	28.99	34.12	41.11	48.96	49.94	55.24	58.47	64.80	84.97	87.17
Yes -																		
Inference -																		
NASNet																		
Large -																		
OpenCL																		
(FPS)																		
Normalized	24.57	27.81	35.15	26.57	37.94	40.62	46.32	64.98	33.26	39.14	47.16	56.17	57.29	63.37	67.08	74.34	97.48	100%
Standard	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%
Deviation	0.2%	1.1%	0.6%	0.3%	0.3%	0.1%	0.3%	0.4%	0.2%	0.2%	0.3%	0.5%	0.4%	0.2%	0.4%	0.2%	0.4%	0.5%
cl-mem -	142.1	164.5	265.3	153.5	205.2	205.2	228.8	337.5	162.9	250.2	296.2	395.6	395.7	395.6	395.7	436.2	544.3	566.7
Read	0	0	3	0	7	7	0	3	0	3	0	7	0	7	0	7	3	0
Normalized	25.07	29.03	46.82	27.09	36.22	36.22	40.37	59.56	28.75	44.16	52.27	69.82	69.83	69.82	69.83	76.98	96.05	100%
Standard	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%
Deviation	3%	0.1%	0.1%	0%	0%	0.1%	0.1%	0.1%	0%	0%	0.1%	0.1%	0.1%	0.1%	0.1%	0.1%	0.1%	0.1%
clpeak -	137.4	159.6	196.6	149.5	223.9	242.6	297.9	414.9	168.4	185.1	231.4	263.1	268.3	310.2	345.8	378.9	521.9	547.9
D.P.D	7	8	0	8	1	3	4	6	4	2	4	1	2	7	1	8	0	7
(GFLOPS)																		

Normalized	25.09	29.14	35.88	27.3%	40.86	44.28	54.37	75.73	30.74	33.78	42.24	48.02	48.97	56.62	63.11	69.16	95.24	100%
Standard	0%	0.1%	0.3%	0.2%	0.5%	0%	0.4%	0.5%	0%	0%	0%	0.5%	0.5%	0.4%	0.5%	0%	0.5%	0.4%
Deviation																		
PlaidML -	1.28	1.43	1.41	1.83	2.37	2.59	2.81	3.00	2.98	3.14	3.10	3.23	3.68	3.67	3.95	4.52	4.56	5.09
No -																		
Inference -																		
IMDB LSTM																		
- OpenCL																		
(FPS/Watt)																		
Normalized	25.15	28.09	27.7%	35.95	46.56	50.88	55.21	58.94	58.55	61.69	60.9%	63.46	72.3%	72.1%	77.6%	88.8%	89.59	100%
SHOC	4362	5053	6214	4789	7115	7715	9408	1323	5330	5893	7321	8379	8529	9869	1095	1201	1665	1733
Scalable								1							3	6	6	4
HeterOgene																		
ous																		
Computing																		
- OpenCL -																		
Max SP																		
Flops																		
(GFLOPS)																		
Normalized	25.17	29.15	35.85	27.63	41.05	44.51	54.27	76.33	30.75	34%	42.23	48.34	49.2%	56.94	63.19	69.32	96.09	100%
Standard	0%	0.1%	0.3%	0.2%	0%	0%	0.1%	0.1%	0.4%	0%	0.5%	0.6%	0.9%	0.9%	0.5%	0.9%	0.9%	0.9%
Deviation																		
LuxMark -	1173	1324	1687	1224	1729	1692	1378	2168	1516	1614	2168	3009	3009	2973	2916	3017	4297	4654
GPU -	2	6	1	5	0	5	8	9	2	7	3	3	7	4	4	9	4	6
Luxball																		
HDR																		
Normalized	25.21	28.46	36.25	26.31	37.15	36.36	29.62	46.6%	32.57	34.69	46.58	64.65	64.66	63.88	62.66	64.84	92.33	100%
Standard	0.1%	0.3%	0.4%	0.1%	0%	0.4%	0.2%	0.1%	0.4%	0.2%	0%	0%	0%	0.6%	0%	0%	0.7%	0.1%
Deviation																		
PlaidML -	204.3	231.6	270.0	224.8	315.3	333.0	399.0	523.9	287.6	315.8	389.4	440.7	446.2	499.6	548.8	589.8	758.2	790.7
No -	2	6	0	9	7	3	0	8	8	3	0	9	3	9	3	6	1	7
Inference -																		
IMDB LSTM																		
- OpenCL																		
(FPS)																		
Normalized	25.84	29.3%	34.14	28.44	39.88	42.11	50.46	66.26	36.38	39.94	49.24	55.74	56.43	63.19	69.4%	74.59	95.88	100%
Standard	0%	0.4%	0.4%	0.1%	0.1%	0.3%	0.1%	0.2%	0%	0.1%	0.2%	0.1%	0.1%	0.2%	0.5%	0.4%	0.2%	0.5%
Deviation																		

PlaidML - Yes - Inference - VGG16 - OpenCL (FPS/Watt)	0.27	0.28	0.29	0.37	0.49	0.57	0.56	0.62	0.48	0.48	0.55	0.55	0.61	0.66	0.77	0.80	1.04	0.94
Normalized	25.96	26.92	27.88	35.58	47.12	54.81	53.85	59.62	46.15	46.15	52.88	52.88	58.65	63.46	74.04	76.92	100%	90.38
PlaidML - Yes - Inference - Mobilenet - OpenCL (FPS/Watt)	6.98	7.17	7.67	9.57	12.35	14.16	14.02	15.21	14.79	16.06	18.20	20.68	21.35	24.75	26.21	25.06	26.71	22.19
Normalized	26.13	26.84	28.72	35.83	46.24	53.01	52.49	56.94	55.37	60.13	68.14	77.42	79.93	92.66	98.13	93.82	100%	83.08
PlaidML - No - Inference - Inception V3 - OpenCL	0.52	0.56	0.61	0.70	0.99	1.08	0.94	1.34	0.95	1.02	1.11	1.23	1.18	1.34	1.36	1.52	1.66	1.95
Normalized	26.67	28.72	31.28	35.9%	50.77	55.38	48.21	68.72	48.72	52.31	56.92	63.08	60.51	68.72	69.74	77.95	85.13	100%
PlaidML - Yes - Inference - ResNet 50 - OpenCL (FPS)	116.0	142.1	162.6	130.6	196.0	205.8	240.6	333.5	156.5	172.6	220.4	250.9	251.0	283.9	305.7	327.8	424.0	433.8
Normalized	26.75	32.77	37.49	30.11	45.18	47.45	55.47	76.89	36.08	39.81	50.81	57.85	57.88	65.46	70.48	75.56	97.75	100%
Standard	0%	0.3%	0.1%	0.2%	0.3%	0.1%	0.1%	0.2%	0.3%	0.1%	0.1%	0.3%	0.4%	0.4%	0.3%	0.1%	0.4%	0.3%
Deviation																		
cl-mem - Write	132.6	151.5	241.5	144.5	190.9	190.0	214.6	340.2	148.8	208.6	241.2	333.2	322.6	319.5	331.5	345.7	441.2	495.3
Normalized	26.77	30.58	48.75	29.18	38.54	38.36	43.34	68.69	30.05	42.12	48.69	67.26	65.14	64.51	66.92	69.8%	89.07	100%
Standard	0%	0%	0%	0%	0%	0%	0%	0.1%	0%	0.1%	0.3%	0.4%	0.5%	0.2%	0.1%	0.3%	0.1%	0.5%
Deviation																		
LuxMark - GPU - Hotel (Score)	2730	3048	3885	2648	3870	4170	3807	5649	3733	3805	4833	6148	6150	6565	6590	6839	9215	9836
Normalized	27.76	30.99	39.5%	26.92	39.35	42.4%	38.7%	57.43	37.95	38.68	49.14	62.51	62.53	66.74	67%	69.53	93.69	100%
Standard	0.1%	0.2%	0.4%	0.2%	0.1%	0.1%	0%	0.2%	0%		0%	0%	0.1%		0%	0%	0.1%	0%
Deviation																		



<b>clpeak -</b>	<b>143.4</b>	164.2	263.2	146.6	196.4	197.2	222.0	328.9	157.6	234.6	275.8	367.9	368.9	369.1	368.4	405.0	506.3	<b>528.8</b>
<b>G.M.B</b>	<b>8</b>	5	2	2	7	6	9	0	0	7	6	2	6	1	7	0	3	<b>9</b>
<b>(GBPS)</b>																		
<b>Normalized</b>	27.13	31.06	49.77	27.72	37.15	37.3%	41.99	62.19	29.8%	44.37	52.16	69.56	69.76	69.79	69.67	76.58	95.73	100%
<b>Standard</b>	0.1%	0%	0.3%	0.1%	0.2%	0%	0.1%	0.4%	0%	0.1%	0.1%	0.1%	0%	0.2%	0.1%	0%	0.1%	0.2%
<b>Deviation</b>																		
<b>PlaidML -</b>	<b>671.4</b>	759.4	990.1	726.3	987.6	1009	1124	1673	814.6	1002	1232	1531	1517	1649	1664	1762	2335	<b>2467</b>
<b>No -</b>	<b>9</b>	3	2	0	9				5									
<b>Inference -</b>																		
<b>MobileNet -</b>																		
<b>OpenCL</b>																		
<b>(FPS)</b>																		
<b>Normalized</b>	27.22	30.79	40.14	29.44	40.04	40.92	45.56	67.81	33.03	40.62	49.97	62.07	61.5%	66.86	67.45	71.43	94.68	100%
<b>Standard</b>	0.2%	0.1%	0%	0.1%	0.3%	0.1%	0.2%	0.5%	0.3%	0.1%	0.3%	0.2%	0.2%	0.1%	0.4%	0.3%	0.5%	1%
<b>Deviation</b>																		
<b>PlaidML -</b>	<b>94.72</b>	105.6	135.1	104.4	147.1	153.3	184.8	257.5	116.9	140.0	167.8	199.4	198.0	215.8	238.7	259.2	343.4	<b>347.9</b>
<b>No -</b>		4	4	4	7	4	0	5	0	5	7	3	3	5	3	6	9	<b>3</b>
<b>Inference -</b>																		
<b>Inception</b>																		
<b>V3 -</b>																		
<b>OpenCL</b>																		
<b>Normalized</b>	27.22	30.36	38.84	30.02	42.3%	44.07	53.11	74.02	33.6%	40.25	48.25	57.32	56.92	62.04	68.61	74.51	98.72	100%
<b>Standard</b>	0.3%	0.3%	0.1%	0.2%	0.1%	0.1%	0.4%	0.1%	0.2%	0.3%	0.2%	0.2%	0.3%	0.1%	0.4%	0.1%	0.2%	0.2%
<b>Deviation</b>																		
<b>cl-mem -</b>	<b>1.05</b>	1.19	1.76	1.34	1.68	1.86	1.90	2.10	1.69	2.67	2.76	3.09	3.28	3.47	3.26	<b>3.85</b>	3.38	3.37
<b>Read</b>																		
<b>(GB/s/Watt)</b>																		
<b>Normalized</b>	27.27	30.91	45.71	34.81	43.64	48.31	49.35	54.55	43.9%	69.35	71.69	80.26	85.19	90.13	84.68	100%	87.79	87.53
<b>Standard</b>	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
<b>clpeak -</b>	<b>3888</b>	4477	5539	4200	6269	6771	8314	1172	4610	4794	5329	7090	7199	8503	8884	1033	1353	<b>1420</b>
<b>S.P.F</b>								1								6	3	<b>4</b>
<b>(GFLOPS)</b>																		
<b>Normalized</b>	27.37	31.52	38.99	29.57	44.14	47.67	58.54	82.52	32.46	33.75	37.52	49.91	50.68	59.86	62.55	72.77	95.28	100%
<b>Standard</b>	4%	2.9%	3.1%	0.1%	0.1%	0.2%	0.1%	0.1%	0.3%	5%	2.5%	4.6%	4.4%	3.5%	0.1%	2.1%	2.7%	6.2%
<b>Deviation</b>																		
<b>PlaidML -</b>	<b>0.20</b>	0.22	0.23	0.28	0.37	0.46	0.43	0.49	0.38	0.38	0.40	0.39	0.44	0.49	0.58	0.59	<b>0.73</b>	0.71
<b>Yes -</b>																		
<b>Inference -</b>																		
<b>VGG19 -</b>																		
<b>OpenCL</b>																		
<b>(FPS/Watt)</b>																		
<b>Normalized</b>	27.4%	30.14	31.51	38.36	50.68	63.01	58.9%	67.12	52.05	52.05	54.79	53.42	60.27	67.12	79.45	80.82	100%	97.26
<b>Standard</b>		%	%	%	%	%		%	%	%	%	%	%	%	%	%		%

PlaidML - Yes - Inference - Inception V3 - OpenCL	0.56	0.61	0.60	0.78	0.95	1.12	1.10	1.13	1.22	1.32	1.34	1.49	1.57	1.72	1.75	1.99	2.04	2.04
Normalized	27.45	29.9%	29.41	38.24	46.57	54.9%	53.92	55.39	59.8%	64.71	65.69	73.04	76.96	84.31	85.78	97.55	100%	100%
PlaidML - Yes - Inference - ResNet 50 - OpenCL (FPS/Watt)	0.71	0.78	0.79	1.01	1.24	1.57	1.51	1.76	1.42	1.47	1.55	1.74	1.69	2.10	2.19	2.22	2.53	2.56
Normalized	27.73	30.47	30.86	39.45	48.44	61.33	58.98	68.75	55.47	57.42	60.55	67.97	66.02	82.03	85.55	86.72	98.83	100%
cl-mem - Write (GB/s/Watt)	1.07	1.03	1.48	1.44	1.54	1.76	1.64	2.02	1.61	2.33	2.24	2.61	3.71	2.60	2.74	2.85	3.46	3.13
Normalized	28.84	27.76	39.89	38.81	41.51	47.44	44.2%	54.45	43.4%	62.8%	60.38	70.35	100%	70.08	73.85	76.82	93.26	84.37
PlaidML - No - Inference - ResNet 50 - OpenCL (FPS)	181.8	203.9	247.4	196.0	276.8	294.3	337.4	496.2	228.1	259.4	319.2	374.7	376.7	421.5	441.2	468.7	634.0	648.3
Normalized	28.05	31.45	38.17	30.23	42.71	45.4%	52.05	76.54	35.19	40.01	49.24	57.79	58.11	65.02	68.05	72.29	97.8%	100%
Standard	0.1%	0%	0.4%	0.1%	0.1%	0.1%	0.2%	0.3%	0.1%	0.3%	0.2%	0.2%	0.1%	0.1%	0.1%	0.2%	0.2%	0.2%
Deviation	0.1%	0%	0.4%	0.1%	0.1%	0.1%	0.2%	0.3%	0.1%	0.3%	0.2%	0.2%	0.1%	0.1%	0.1%	0.2%	0.2%	0.2%
OctaneBenchmark - Total Score (Score)	95.60	109.7	142.2	91.44	132.8	141.3	147.9	212.0	118.4	132.2	164.5	205.2	206.8	220.7	222.8	233.5	309.0	322.7
Normalized	29.62	33.99	44.06	28.33	41.16	43.79	45.85	65.7%	36.71	40.98	50.97	63.59	64.1%	68.4%	69.05	72.36	95.75	100%
PlaidML - Yes - Inference - Mobilenet - OpenCL (FPS)	841.4	973.3	1203	908.4	1289	1371	1545	2163	1161	1394	1694	1984	1970	2265	2292	2437	2751	2842
Normalized	29.6%	34.25	42.32	31.96	45.35	48.24	54.34	76.08	40.83	49.05	59.6%	69.79	69.29	79.69	80.65	85.75	96.77	100%
Standard	0.3%	0.4%	0.2%	0.5%	0.6%	0.1%	0.5%	0.6%	0.2%	0.3%	0.4%	1%	0.5%	0.4%	0.4%	0.5%	0.9%	0.2%
Deviation	0.3%	0.4%	0.2%	0.5%	0.6%	0.1%	0.5%	0.6%	0.2%	0.3%	0.4%	1%	0.5%	0.4%	0.4%	0.5%	0.9%	0.2%


cl-mem - Copy (GB/s/Watt)	0.96	1.04	1.45	1.25	1.35	1.53	1.65	1.62	1.58	1.96	2.66	2.43	2.44	2.82	2.28	2.83	3.23	2.17
Normalized	29.72	32.2%	44.89	38.7%	41.8%	47.37	51.08	50.15	48.92	60.68	82.35	75.23	75.54	87.31	70.59	87.62	100%	67.18
PlaidML - Yes - Inference - NASNet Large - OpenCL (FPS/Watt)	0.11	0.12	0.13	0.15	0.18	0.23	0.20	0.24	0.21	0.22	0.23	0.25	0.26	0.30	0.29	0.32	0.37	0.37
Normalized	29.73	32.43	35.14	40.54	48.65	62.16	54.05	64.86	56.76	59.46	62.16	67.57	70.27	81.08	78.38	86.49	100%	100%
PlaidML - No - Inference - VGG16 - OpenCL (FPS/Watt)	0.31	0.32	0.32	0.42	0.56	0.67	0.61	0.75	0.58	0.56	0.62	0.71	0.68	0.81	0.82	0.82	0.98	1.03
Normalized	30.1%	31.07	31.07	40.78	54.37	65.05	59.22	72.82	56.31	54.37	60.19	68.93	66.02	78.64	79.61	79.61	95.15	100%
PlaidML - Yes - Inference - DenseNet 201 - OpenCL (FPS)	79.71	92.92	111.4	91.07	127.0	130.1	143.5	191.3	102.8	114.5	154.1	179.5	176.4	190.0	189.2	200.3	254.4	263.1
Normalized	30.29	35.31	42.36	34.61	48.27	49.46	54.55	72.71	39.08	43.55	58.6%	68.22	67.07	72.21	71.93	76.15	96.7%	100%
Standard	0%	0.2%	0.3%	0.3%	0.1%	0%	0%	0%	0.2%	0%	0.1%	0%	0.2%	0.1%	0.2%	0.1%	0.1%	0.1%
Deviation FAHBench (Ns/Day)	91.53	102.8	116.0	102.6	140.2	138.2	155.2	198.2	126.6	139.8	183.4	205.2	204.7	229.1	242.7	256.8	301.8	301.9
Normalized	30.31	34.07	38.43	33.98	46.46	45.78	51.41	65.66	41.94	46.31	60.75	67.96	67.8%	75.89	80.4%	85.06	99.99	100%
Standard	0.1%	0.1%	0.1%	0.3%	0.5%	0.1%	0.6%	0.4%	0.4%	0.2%	0.6%	0.3%	0.5%	0.5%	0.7%	0.5%	0.5%	0.4%
Deviation PlaidML - No - Inference - VGG19 - OpenCL (FPS/Watt)	0.24	0.26	0.25	0.32	0.43	0.50	0.48	0.56	0.42	0.44	0.46	0.52	0.53	0.61	0.56	0.64	0.74	0.75
Normalized	32%	34.67	33.33	42.67	57.33	66.67	64%	74.67	56%	58.67	61.33	69.33	70.67	81.33	74.67	85.33	98.67	100%
Rodinia - O.P.F (sec)	12.99	11.69	9.95	11.93	8.26	7.88	6.49	4.96	11.87	10.82	8.90	7.87	7.75	6.82	6.22	5.74	4.39	4.24
Normalized	32.64	36.27	42.61	35.54	51.33	53.81	65.33	85.48	35.72	39.19	47.64	53.88	54.71	62.17	68.17	73.87	96.58	100%
	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	%	

Standard	0.3%	1.1%	0.6%	0.3%	0.1%	1.5%	0.2%	0.8%	0.2%	0.3%	0.3%	0%	0.1%	1.8%	0.4%	0.3%	0.6%	0.4%
Deviation																		
PlaidML -	0.54	0.59	0.58	0.72	0.94	1.01	0.95	0.94	0.96	0.99	1.15	1.26	1.27	1.53	1.29	1.41	1.47	1.52
Yes -																		
Inference -																		
DenseNet																		
201 -																		
OpenCL																		
(FPS/Watt)																		
Normalized	35.29	38.56	37.91	47.06	61.44	66.01	62.09	61.44	62.75	64.71	75.16	82.35	83.01	100%	84.31	92.16	96.08	99.35
%																		
clpeak -	0.98	1.09	1.09	1.46	1.90	2.14	2.20	2.49	1.73	1.72	1.84	1.85	1.95	2.25	2.25	2.50	2.65	2.70
D.P.D																		
(GFLOPS/W																		
att)																		
Normalized	36.3%	40.37	40.37	54.07	70.37	79.26	81.48	92.22	64.07	63.7%	68.15	68.52	72.22	83.33	83.33	92.59	98.15	100%
%																		
Darktable -	4.42	3.92	3.27	3.65	2.89	2.93	2.72	2.29	3.35	2.92	2.22	1.95	1.93	1.93	1.90	1.82	1.64	1.64
Boat -																		
OpenCL																		
(sec)																		
Normalized	37.1%	41.84	50.15	44.93	56.75	55.97	60.29	71.62	48.96	56.16	73.87	84.1%	84.97	84.97	86.32	90.11	100%	100%
%																		
Standard	0.2%	0.2%	0.3%	0.2%	0.2%	0.1%	0.3%	0.6%	0.5%	0.1%	0.4%	0.5%	0.3%	0.3%	0.2%	0.2%	0.2%	1.7%
Deviation																		
SHOC	29.29	31.40	32.05	40.06	50.30	61.43	59.87	66.26	50.17	49.63	50.36	56.38	57.13	65.76	63.93	73.23	76.69	77.72
Scalable																		
HeterOgene																		
ous																		
Computing																		
- OpenCL -																		
Max SP																		
Flops																		
(GFLOPS/W																		
att)																		
Normalized	37.69	40.4%	41.24	51.54	64.72	79.04	77.03	85.25	64.55	63.86	64.8%	72.54	73.51	84.61	82.26	94.22	98.67	100%
%																		
JuliaGPU -	1275	1325	1099	1658	1809	2260	1906	1682	2790	2880	2469	2670	2271	2632	2691	2379	1910	1923
GPU	494	998	178	142	790	230	370	593	570	366	451	027	142	330	799	445	019	117
(Samples/s																		
ec/Watt)																		
Normalized	44.28	46.04	38.16	57.57	62.83	78.47	66.18	58.42	96.88	100%	85.73	92.7%	78.85	91.39	93.45	82.61	66.31	66.77
%																		
cl-mem -	124.7	143.4	216.8	137.2	182.0	182.7	205.9	283.2	145.3	208.5	237.7	287.8	284.4	291.8	289.2	301.2	324.4	322.2
Copy	7	0	3	0	0	0	3	0	7	0	7	0	3	7	0	3	0	0
Normalized	38.46	44.2%	66.84	42.29	56.1%	56.32	63.48	87.3%	44.81	64.27	73.3%	88.72	87.68	89.97	89.15	92.86	100%	99.32
%																		
Standard	0%	0%	0.1%	0.1%	0.1%	0%	0%	0.2%	0%	0.1%	0.1%	0.2%	0.1%	0.1%	0.1%	0.1%	0.2%	0.4%
Deviation																		

ViennaCL - O.L.F (GFLOPS/Watt)	0.66	0.79	0.60	1.00	0.80	0.71	0.70	0.59	1.08	1.10	0.73	1.40	1.35	0.90	0.71	1.18	1.31	1.35
Normalized	47.14	56.43	42.86	71.43	57.14	50.71	50%	42.14	77.14	78.57	52.14	100%	96.43	64.29	50.71	84.29	93.57	96.43
FAHBench (Ns/Day/Watt)	0.62	0.65	0.62	0.83	0.96	1.10	1.00	1.02	1.11	1.10	1.21	1.27	1.29	1.42	1.38	1.45	1.38	1.33
Normalized	42.76	44.83	42.76	57.24	66.21	75.86	68.97	70.34	76.55	75.86	83.45	87.59	88.97	97.93	95.17	100%	95.17	91.72
LuxMark - GPU - Luxball HDR	65.07	67.87	70.64	83.34	95.03	111.5	76.55	88.26	108.7	102.7	112.4	131.3	130.3	136.8	119.0	130.9	131.6	137.3
Normalized	47.38	49.42	51.43	60.68	69.19	81.21	55.74	64.26	79.15	74.81	81.86	95.62	94.87	99.64	86.71	95.33	95.86	100%
LuxMark - GPU - Microphone (Score/Watt)	44.91	46.08	47.25	51.14	58.63	70.32	50.29	57.68	71.22	70.33	72.11	83.09	83.92	92.78	82.60	89.58	89.77	91.00
Normalized	48.4%	49.67	50.93	55.12	63.19	75.79	54.2%	62.17	76.76	75.8%	77.72	89.56	90.45	100%	89.03	96.55	96.76	98.08
LuxMark - GPU - Hotel (Score/Watt)	13.72	14.16	15.21	17.00	19.59	25.64	18.51	20.02	24.22	22.50	22.92	26.38	25.52	27.18	24.50	27.41	26.99	27.46
Normalized	49.96	51.57	55.39	61.91	71.34	93.37	67.41	72.91	88.2%	81.94	83.47	96.07	92.94	98.98	89.22	99.82	98.29	100%
LeelaChess Zero - OpenCL (Nodes/s)	207.8	280.9	466.8	262.4	642.1	753.6	1147	1886	335.8	420.0	704.6	998.1	999.7	1604	1880	2105	3032	3126
Normalized	6.65%	8.99%	14.93	8.39%	20.54	24.11	36.68	60.33	10.74	13.43	22.54	31.93	31.98	51.31	60.13	67.34	96.98	100%
Standard Deviation	3.6%	1.7%	4.7%	2.8%	7.4%	2.8%	2.2%	3%	4.7%	5.4%	0.5%	7%	6.6%	2.9%	2%	2.5%	0.5%	1.6%
LeelaChess Zero - OpenCL (Nodes/s/Watt)	8.44	9.00	10.83	23.89	25.59	26.82	16.52	79.94	61.38	63.25	80.74	52.78	60.39	123.07	72.16	54.76	66.58	
Normalized	7.47%	6.86%	7.31%	8.8%	19.41	20.79	21.79	13.42	64.95	49.87	51.39	65.6%	42.89	49.07	100%	58.63	44.5%	54.1%
LeelaChess Zero - OpenCL (Nodes/s/Watt)	1.26	1.83	2.16	1.90	4.21	5.86	7.10	8.75	2.84	3.13	4.79	6.57	5.78	8.98	10.73	12.32	14.00	15.81
Normalized	7.97%	11.57	13.66	12.02	26.63	37.07	44.91	55.34	17.96	19.8%	30.3%	41.56	36.56	56.8%	67.87	77.93	88.55	100%
LeelaChess Zero - OpenCL (Nodes/s/Watt)	1140	1312	1610	1268	1685	2077	2431	3302	4658	4775	5269	7023	7161	8571	9661	1040	1360	1384
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)																		
Normalized																		
LeelaChess Zero - OpenCL (Nodes/s/Watt)		</																




Normalized	8.24%	9.48%	11.63	9.16%	12.17	15.01	17.56	23.86	33.65	34.5%	38.07	50.74	51.74	61.92	69.8%	75.17	98.32	100%
Standard	3%	0%	2.5%	4.1%	2.5%	0.1%	1.9%	3.5%	1.2%	5.2%	2.1%	4.2%	4.6%	3.7%	0.1%	3.1%	0.7%	4.4%
Deviation																		
NAMD	0.363	0.331	0.286	0.317	0.239	0.222	0.209	0.196	0.250	0.231	0.202	0.193	0.194	0.195	0.193	0.193	0.189	0.188
CUDA - ATPase Simulation - 327,506 Atoms (days/ns)	68	10	75	89	57	69	78	69	31	85	50	36	40	01	07	85	09	67
Normalized	51.88	56.98	65.8%	59.35	78.75	84.72	89.94	95.92	75.37	81.38	93.17	97.57	97.05	96.75	97.72	97.33	99.78	100%
Standard	0.1%	0.8%	0.6%	1.7%	1.4%	0.8%	0.9%	0.4%	2%	1.4%	1.5%	2.2%	2.2%	2.2%	1.8%	2.1%	2.2%	2.7%
Deviation																		
JuliaGPU - GPU (Samples/sec)	1692	1818	1956	1821	2169	2204	2379	2622	2309	2393	2519	2641	2670	2758	2801	2868	2998	3031
GPU	5804	8271	5372	8831	4214	4775	1497	4889	1964	0083	2517	7244	8627	6815	6247	1833	3472	9857
(Samples/sec)	5	0	8	9	5	6	4	9	2	8	7	1	5	1	0	1	9	4
Normalized	55.82	59.99	64.53	60.09	71.55	72.71	78.47	86.49	76.16	78.93	83.09	87.13	88.09	90.99	92.4%	94.6%	98.89	100%
Standard	0.4%	0.1%	0.1%	0.1%	0.1%	0.1%	0.2%	0.4%	0.1%	0.4%	0.1%	0.2%	0.4%	0.2%	0.1%	0.7%	0%	0.3%
Deviation																		
OctaneBenchmark - Total Score (Score/Watt)	0.60	0.61	0.64	0.76	0.84	1.07	0.88	0.89	0.93	0.91	0.87	0.97	0.96	1.04	0.93	1.04	0.96	0.99
Normalized	56.07	57.01	59.81	71.03	78.5%	100%	82.24	83.18	86.92	85.05	81.31	90.65	89.72	97.2%	86.92	97.2%	89.72	92.52
Standard	45.48	47.38	52.13	34.69	35.10	38.70	34.87	35.91	30.97	30.67	31.11	30.42	31.64	30.76	31.65	30.63	31.05	31.69
Deviation																		
Rodinia - OpenCL Myocyte (sec)	52.13	47.38	52.13	34.69	35.10	38.70	34.87	35.91	30.97	30.67	31.11	30.42	31.64	30.76	31.65	30.63	31.05	31.69
Normalized	66.89	64.2%	58.35	87.69	86.67	78.6%	87.24	84.71	98.22	99.18	97.78	100%	96.14	98.89	96.11	99.31	97.97	95.99
Standard	0.2%	0.2%	0.6%	0.5%	0.4%	0.6%	0.2%	0.1%	0.5%	0.3%	0.3%	0.2%	0.2%	0.8%	0.2%	0.2%	0.4%	0.2%
Deviation																		
Darktable - Masskrug - OpenCL (sec)	5.78	5.74	4.12	4.07	3.92	3.97	3.94	3.92	4.01	4.02	3.73	3.68	3.67	3.71	3.71	3.69	3.66	3.66
Normalized	63.32	63.76	88.83	89.93	93.37	92.19	92.89	93.37	91.27	91.04	98.12	99.46	99.73	98.65	98.65	99.19	100%	100%
Standard	0.8%	0.7%	0.3%	0.7%	0.4%	0.3%	0.6%	0.5%	0.4%	0.1%	0.2%	0.7%	0.6%	0.7%	0.5%	0.4%	0.5%	0.7%
Deviation																		
ViennaCL - O.L.F (GFLOPS)	56.95	59.21	62.01	58.67	63.86	64.68	66.67	69.09	65.10	66.06	67.92	69.17	69.41	70.45	70.96	71.47	72.71	72.96
Normalized	78.06	81.15	84.99	80.41	87.53	88.65	91.38	94.7%	89.23	90.54	93.09	94.81	95.13	96.56	97.26	97.96	99.66	100%
Standard	0.3%	0.1%	0.1%	0.1%	0.1%	0%	0.5%	0.8%	0.1%	0.1%	0.5%	0.1%	0.1%	0.1%	0.2%	0.6%	1.7%	1.4%
Deviation																		





## NAMD CUDA 2.13

System Power Consumption Monitor

Watts, Fewer Is Better











## LuxMark 3.1

System Power Consumption Monitor

Watts, Fewer Is Better








# OctaneBench 4.00c

System Power Consumption Monitor

Watts, Fewer Is Better











## LuxMark 3.1

System Power Consumption Monitor

Watts, Fewer Is Better








## LuxMark 3.1


System Power Consumption Monitor


Watts, Fewer Is Better

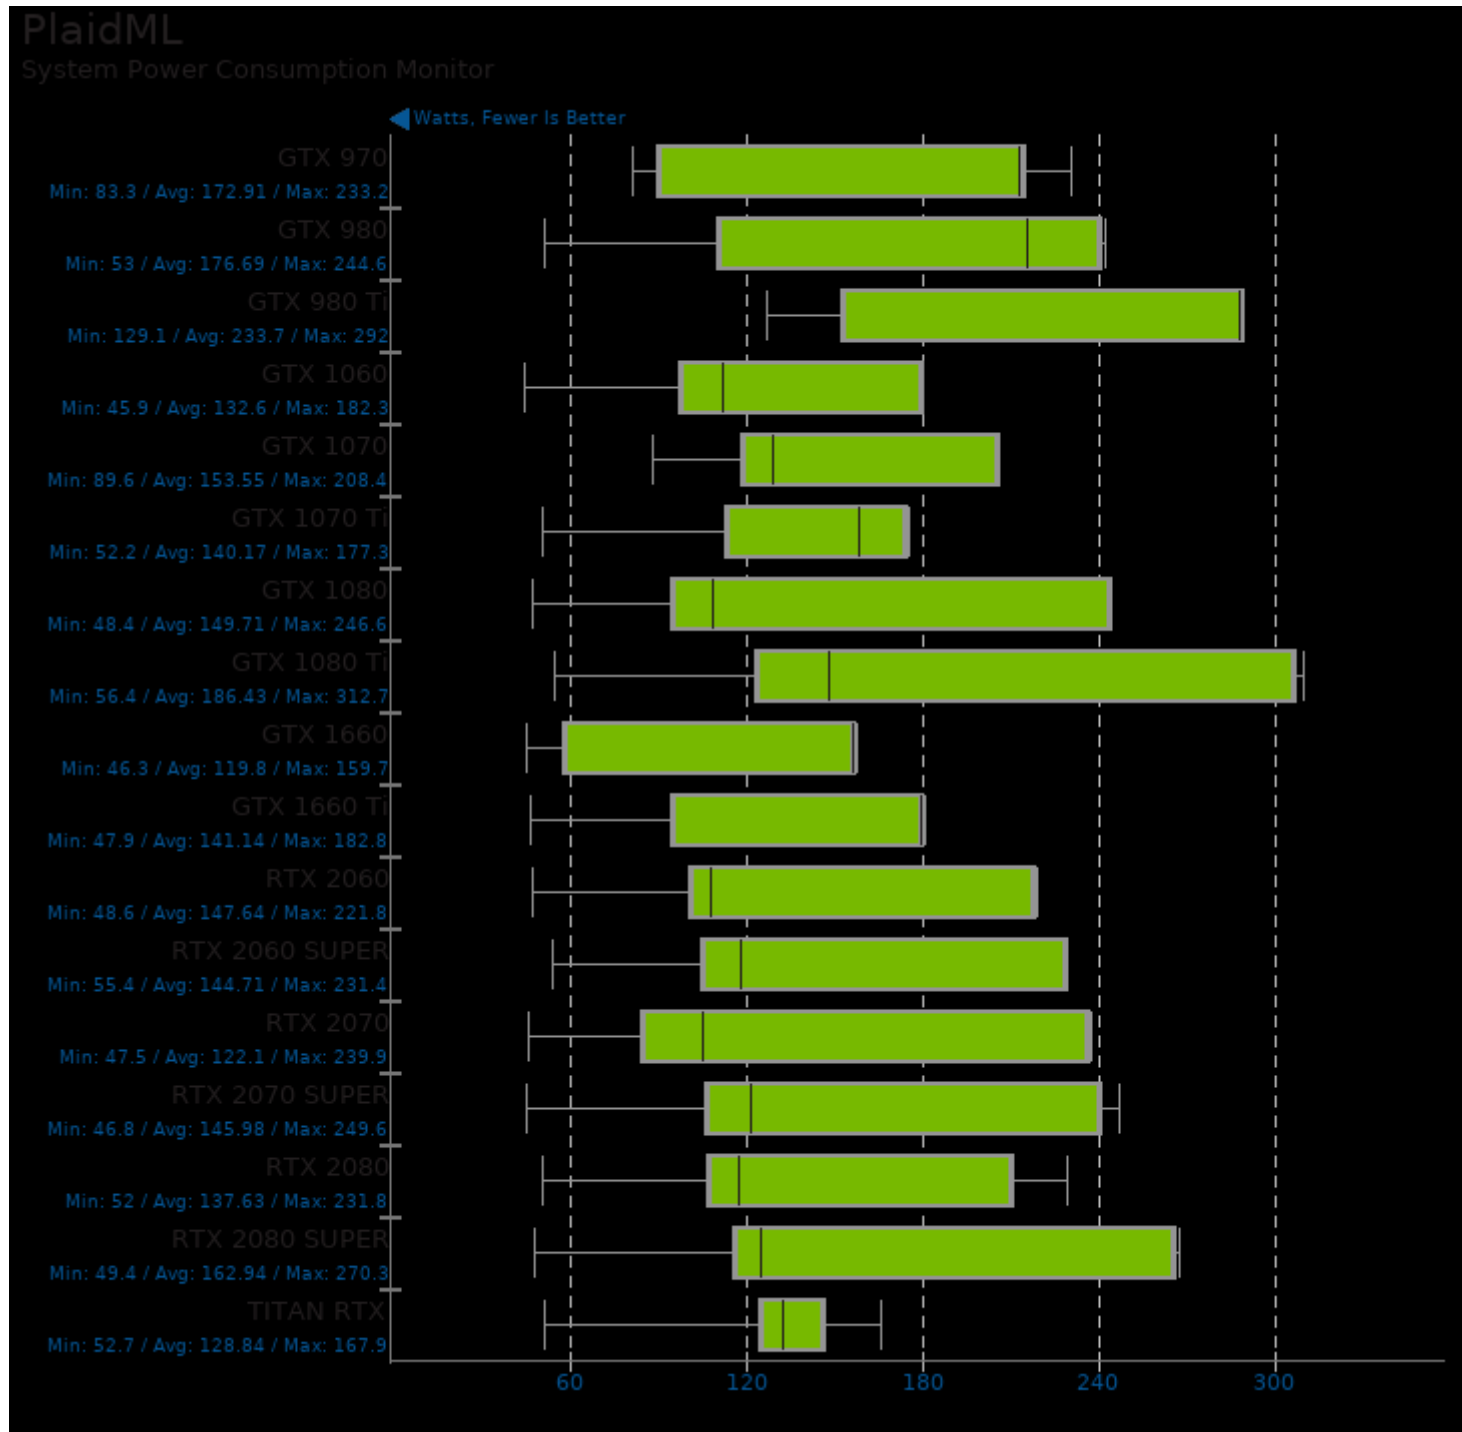







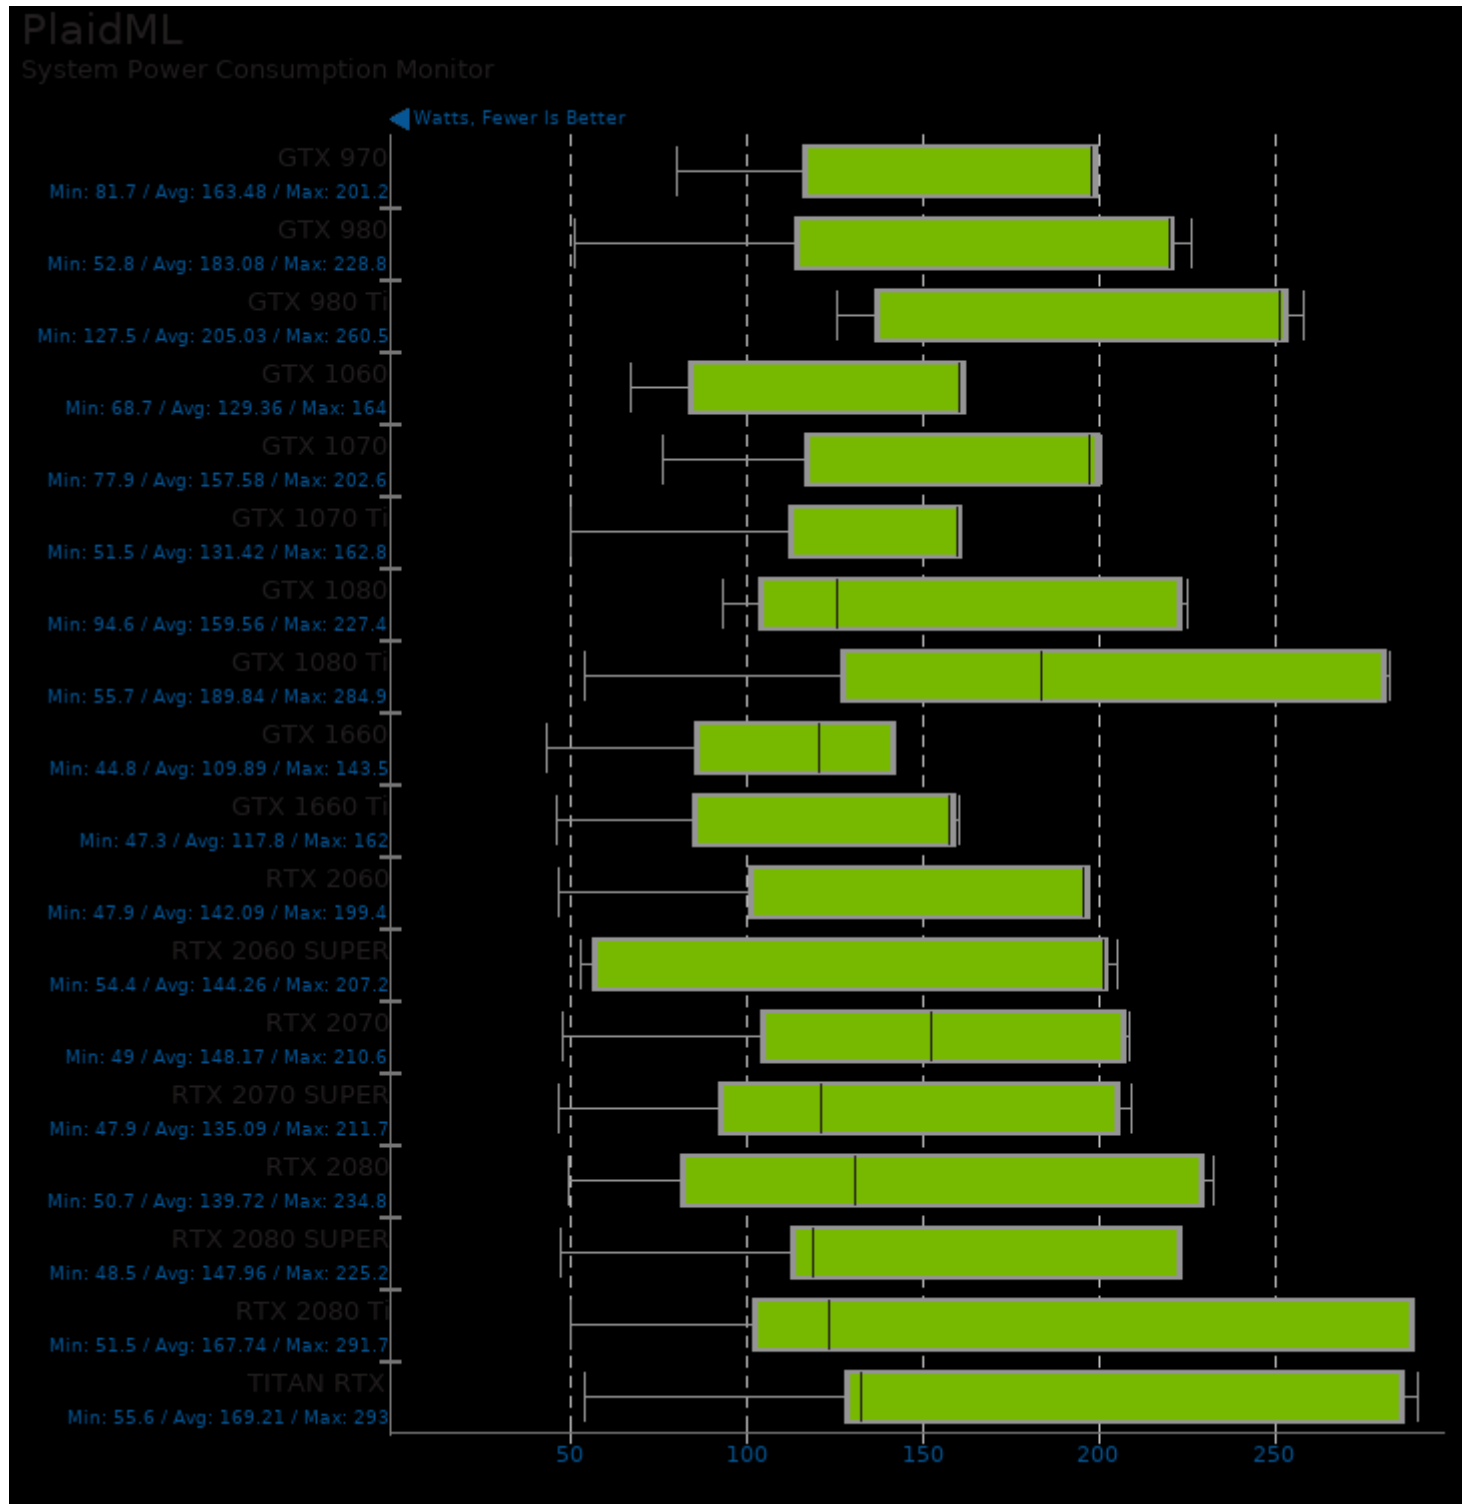





























# SHOC Scalable HeterOgeneous Computing 2015-11-10


Target: OpenCL - Benchmark: MD5 Hash



1. (CXX) g++ options: -O2 -ISHOCCCommonMPI -ISHOCCCommonOpenCL -ISHOCCCommon -lOpenCL -lrt -pthread -lmpi\_cxx -lmpi










## SHOC Scalable HeterOgeneous Computing 2015-11-10

Target: OpenCL - Benchmark: FFT SP



1. (CXX) g++ options: -O2 -ISHOCCCommonMPI -ISHOCCCommonOpenCL -ISHOCCCommon -lOpenCL -lrt -pthread -lmpi\_cxx -lmpi




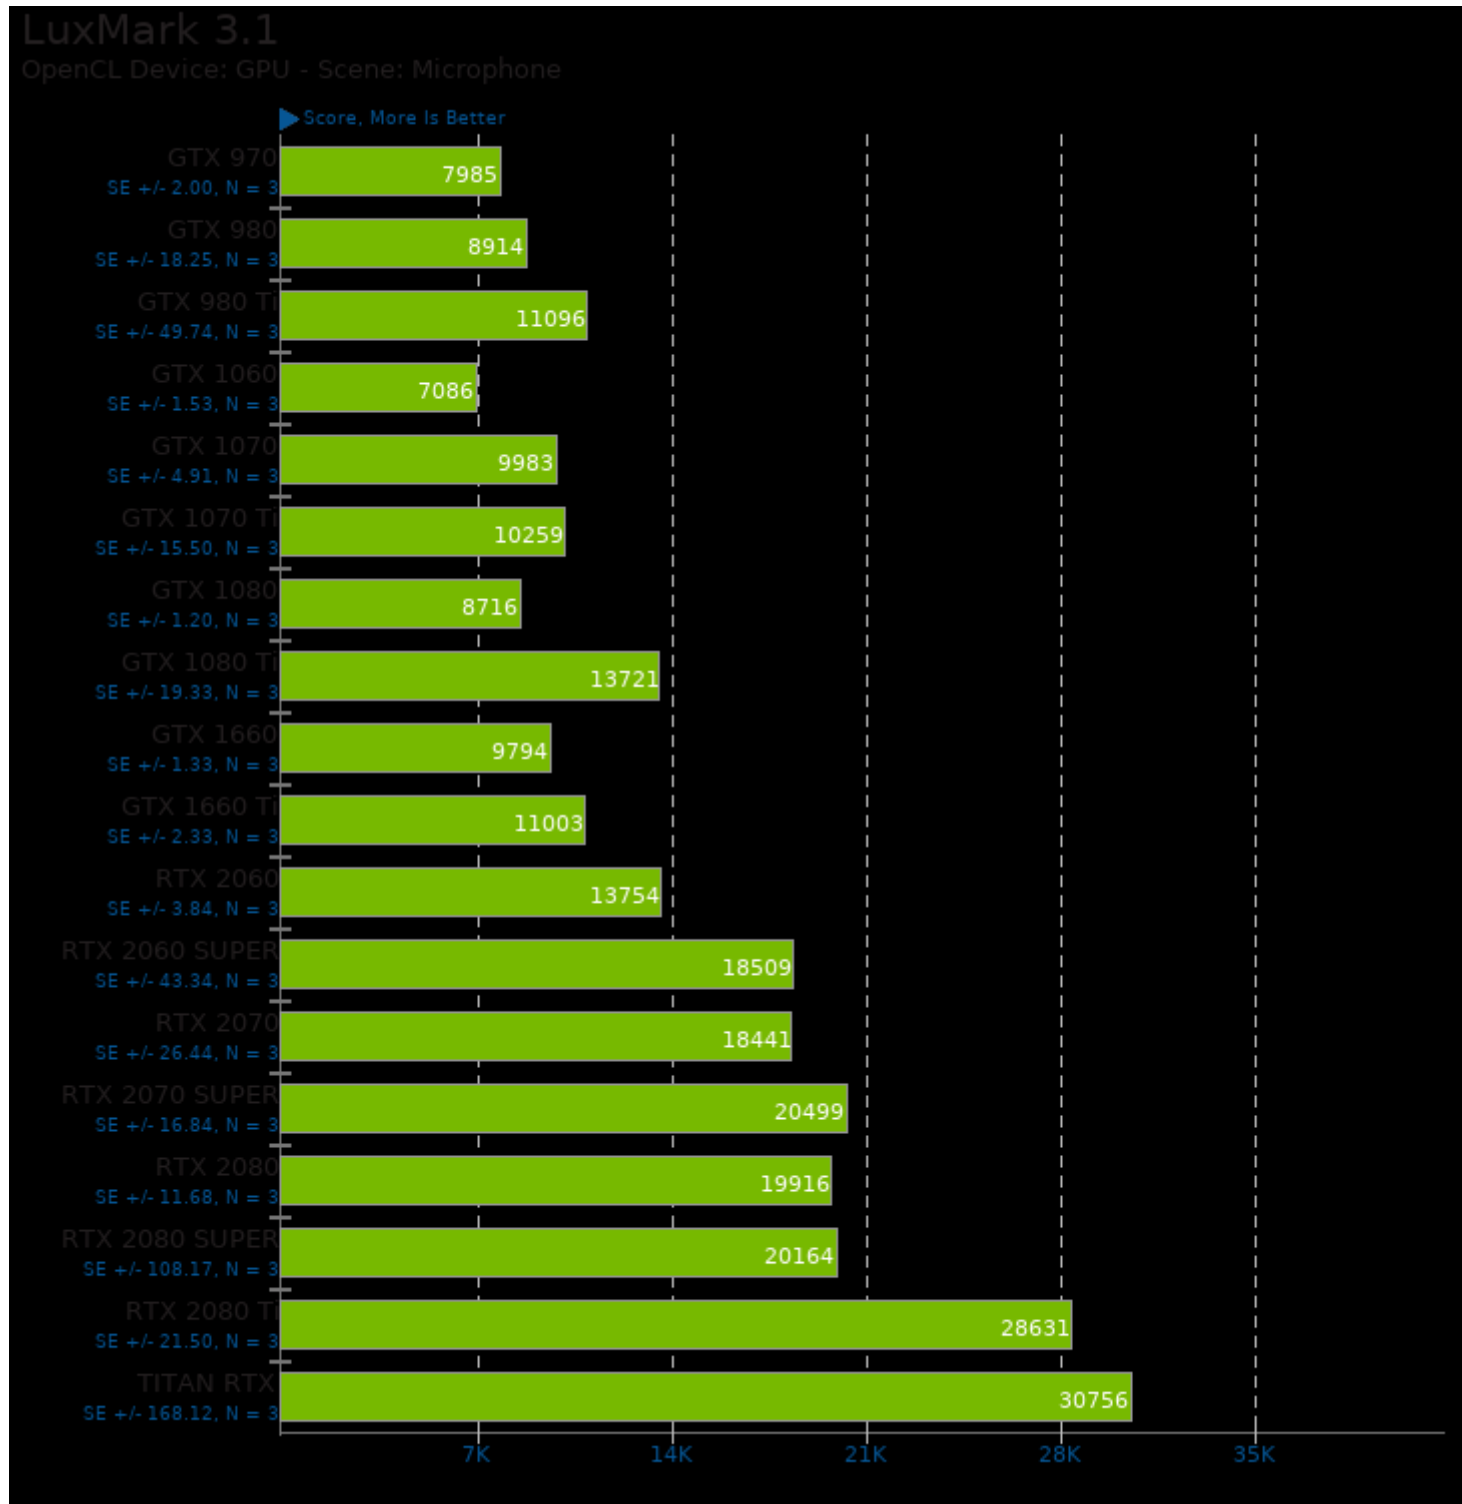


# SHOC Scalable HeterOgeneous Computing 2015-11-10

Target: OpenCL - Benchmark: FFT SP

GFLOPS Per Watt, More Is Better







## Darktable 2.6.0


Test: Server Room - Acceleration: OpenCL


← Seconds, Fewer Is Better



















## PlaidML

FP16: Yes - Mode: Inference - Network: NASNet Large - Device: OpenCL




cl-mem 2017-01-13


Benchmark: Read



1, (CC) gcc options: -O2 -fno -fOpenCL






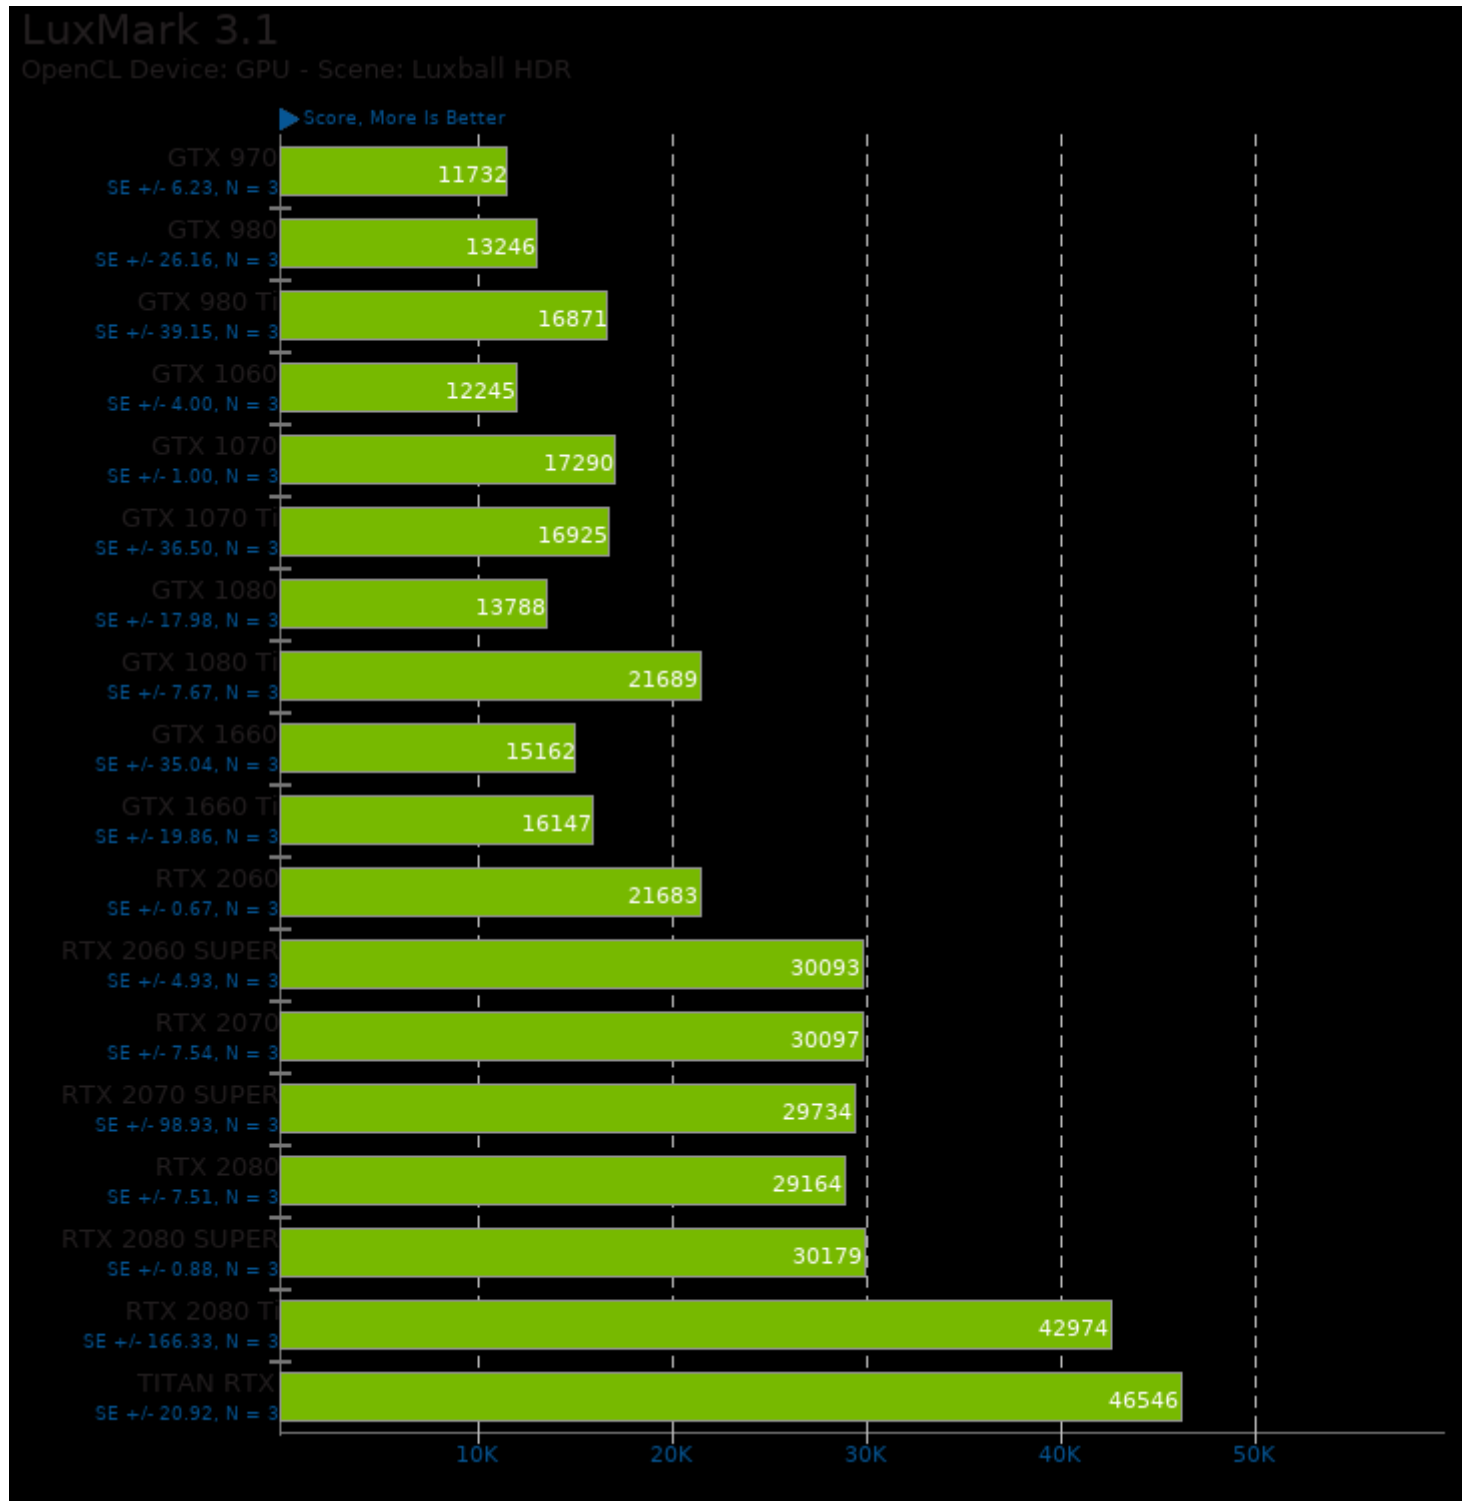



## SHOC Scalable Heterogeneous Computing 2015-11-10


Target: OpenCL - Benchmark: Max SP Flops





1. (CXX) g++ options: -O2 -ISHOCCCommonMPI -ISHOCCCommonOpenCL -ISHOCCCommon -lOpenCL -lrt -pthread -lmpi\_cxx -lmpi

















cl-mem 2017-01-13

Benchmark: Write



1, (CC) gcc options: -O2 -fno -fOpenCL

## Rodinia 2.4


System Power Consumption Monitor


◀ Watts, Fewer Is Better




## LuxMark 3.1

OpenCL Device: GPU - Scene: Hotel











cl-mem 2017-01-13


Benchmark: Read
















cl-mem 2017-01-13

Benchmark: Write

GB/s Per Watt, More Is Better





## OctaneBench 4.00c


Total Score

Score, More Is Better



## PlaidML

FP16: Yes - Mode: Inference - Network: Mobilenet - Device: OpenCL




cl-mem 2017-01-13


Benchmark: Copy


GB/s Per Watt, More Is Better



















## Rodinia 2.4


Test: OpenCL Particle Filter

◀ Seconds, Fewer Is Better





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
















cl-mem 2017-01-13


Benchmark: Copy



1, (CC) gcc options: -O2 -fno -fOpenCL








## LuxMark 3.1


OpenCL Device: GPU - Scene: Microphone


Score Per Watt, More Is Better






















## JuliaGPU 1.2pts1

OpenCL Device: GPU

► Samples/sec, More Is Better



1. (CC) gcc options: -O3 -march=native -ftree-vectorize -funroll-loops -lglt -lOpenCL -lGL -lm







## Rodinia 2.4

Test: OpenCL Myocyte

Seconds, Fewer Is Better




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




## ViennaCL 1.4.2


OpenCL LU Factorization

► GFLOPS, More Is Better




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







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



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




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




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





Geometric mean based upon tests: pts/fahbench, pts/namd-cuda, pts/octanebench, pts/rodinia, pts/clpeak, pts/plaidml, pts/lczero, pts/cl-mem, pts/viennacl and pts/shoc



Geometric mean based upon tests: pts/juliagpu, pts/luxmark, pts/shoc, pts/cl-mem, pts/clpeak, pts/rodinia, system/darktable and pts/viennacl

This file was automatically generated via the Phoronix Test Suite benchmarking software on Sunday, 22 December 2024 11:39.