



Coder Radio XPS 13 ML Ubuntu Benchmark

Intel Core i7-10510U testing with a LENOVO 20U9CTO1WW (N2WET24W 1.14 BIOS) and Intel UHD 3GB on Fedora 33 via the Phoronix Test Suite.

Automated Executive Summary

XPS 13 Tiger Lake Ubuntu 20.04 had the most wins, coming in first place for 96% of the tests.

Based on the geometric mean of all complete results, the fastest (XPS 13 Tiger Lake Ubuntu 20.04) was 1.488x the speed of the slowest (ThinkPad X1 Fedora Comet Lake).

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

oneDNN (Harness: Deconvolution Batch shapes_1d - Data Type: u8s8f32 - Engine: CPU) at 7.436x

oneDNN (Harness: IP Shapes 3D - Data Type: f32 - Engine: CPU) at 2.722x

oneDNN (Harness: Convolution Batch Shapes Auto - Data Type: u8s8f32 - Engine: CPU) at 2.664x

oneDNN (Harness: Deconvolution Batch shapes_3d - Data Type: u8s8f32 - Engine: CPU) at 2.608x

oneDNN (Harness: IP Shapes 1D - Data Type: u8s8f32 - Engine: CPU) at 2.419x

oneDNN (Harness: Convolution Batch Shapes Auto - Data Type: f32 - Engine: CPU) at 1.904x

PlaidML (FP16: No - Mode: Inference - Network: VGG16 - Device: CPU) at 1.73x

Mobile Neural Network (Model: resnet-v2-50) at 1.615x

NCNN (Target: CPU - Model: vgg16) at 1.587x

NCNN (Target: Vulkan GPU - Model: vgg16) at 1.557x.

Test Systems:

XPS 13 Tiger Lake Ubuntu 20.04

Processor: Intel Core i5-1135G7 @ 4.20GHz (4 Cores / 8 Threads), Motherboard: Dell 0THX8P (1.1.1 BIOS), Chipset: Intel Device a0ef, Memory: 16GB, Disk: Micron 2300 NVMe 512GB, Graphics: Intel Xe 3GB (1300MHz), Audio: Realtek ALC289, Network: Intel Device a0f0

OS: Ubuntu 20.04, Kernel: 5.6.0-1036-oem (x86_64), Desktop: GNOME Shell 3.36.4, Display Server: X Server 1.20.8, Display Driver: modesetting 1.20.8, OpenGL: 4.6 Mesa 20.0.8, Vulkan: 1.2.131, Compiler: GCC 9.3.0, File-System: ext4, Screen Resolution: 1920x1200

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --enable-checking=release --enable-clocale=gnu --enable-default-pie --enable-gnu-unique-object --enable-languages=c,ada,c++,go,brig,d,fortran,objc,obj-c++,gm2 --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none=/build/gcc-9-HskZEa/gcc-9-9.3.0/debian/tmp-nvptx/usr,hsa --enable-plugin --enable-shared --enable-threads=posix --host=x86_64-linux-gnu --program-prefix=x86_64-linux-gnu- --target=x86_64-linux-gnu --with-abi=m64 --with-arch-32=i686 --with-default-libstdcxx-abi=new --with-gcc-major-version-only --with-multilib-list=m32,m64,mx32 --with-target-system-zlib=auto --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate powersave - CPU Microcode: 0x60 - ThermalD 1.9.1
Python Notes: Python 3.8.5
Security Notes: itlb_multihit: Not affected + 11tf: Not affected + mds: Not affected + 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 Enhanced IBRS IBPB: conditional RSB filling + srbds: Not affected + tsx_async_abort: Not affected

ThinkPad X1 Fedora Comet Lake

Processor: Intel Core i7-10510U @ 4.90GHz (4 Cores / 8 Threads), Motherboard: LENOVO 20U9CTO1WW (N2WET24W 1.14 BIOS), Chipset: Intel Comet Lake PCH-LP, Memory: 2 x 8 GB LPDDR3-2133MT/s Samsung, Disk: 256GB Western Digital PC SN730 SDBQNTY-256G-1001, Graphics: Intel UHD 3GB (1150MHz), Audio: Realtek ALC285, Network: Intel + Intel Comet Lake PCH-LP CNVi WiFi

OS: Fedora 33, Kernel: 5.9.16-200.fc33.x86_64 (x86_64), Desktop: KDE Plasma 5.20.4, Display Server: X Server 1.20.10, Display Driver: modesetting 1.20.10, OpenGL: 4.6 Mesa 20.2.6, Compiler: GCC 10.2.1 20201125 + Clang 11.0.0, File-System: btrfs, Screen Resolution: 2560x1440

Compiler Notes: --build=x86_64-redhat-linux --disable-libunwind-exceptions --enable-__cxa_atexit --enable-bootstrap --enable-cet --enable-checking=release --enable-gnu-indirect-function --enable-gnu-unique-object --enable-initfini-array --enable-languages=c,c++,fortran,objc,obj-c++,ada,go,d,lto --enable-multilib --enable-offload-targets=nvptx-none --enable-plugin --enable-shared --enable-threads=posix --mandir=/usr/share/man --with-arch_32=i686 --with-gcc-major-version-only --with-isl --with-linker-hash-style=gnu --with-tune=generic --without-cuda-driver
Processor Notes: Scaling Governor: intel_pstate powersave
Python Notes: Python 3.9.1
Security Notes: SELinux + itlb_multihit: KVM: Mitigation of VMX unsupported + 11tf: Not affected + mds: Not affected + 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 Enhanced IBRS IBPB: conditional RSB filling + srbds: Mitigation of TSX disabled + tsx_async_abort: Not affected

	XPS 13 Tiger Lake Ubuntu 20.04	ThinkPad X1 Fedora Comet Lake
AI Benchmark Alpha - D.I.S (Score)	556	
AI Benchmark Alpha - D.T.S (Score)	630	
AI Benchmark Alpha - Device AI Score (Score)	1186	
Mlpack Benchmark - scikit_ica (sec)	123.23	
Standard Deviation	2.7%	
Mlpack Benchmark - scikit_qda (sec)	138.24	
Standard Deviation	0.1%	
Mlpack Benchmark - scikit_svm (sec)	34.66	
Standard Deviation	0.2%	
Mlpack Benchmark - scikit_linearridgeregression	13.50	
Standard Deviation	0.3%	
Mobile Neural Network - SqueezeNetV1.0 (ms)	11.256	18.04
Normalized	100%	62.39%
Standard Deviation	7%	0.9%
Mobile Neural Network - resnet-v2-50 (ms)	54.691	88.35
Normalized	100%	61.9%
Standard Deviation	1%	3.3%
Mobile Neural Network - MobileNetV2_224 (ms)	6.238	9.62
Normalized	100%	64.84%
Standard Deviation	0.7%	21.1%
Mobile Neural Network - mobilenet-v1-1.0 (ms)	8.320	13.15
Normalized	100%	63.27%
Standard Deviation	6.8%	13.8%
Mobile Neural Network - inception-v3 (ms)	68.523	97.71
Normalized	100%	70.13%
Standard Deviation	1.2%	8.8%
NCNN - CPU - mobilenet (ms)	35.03	45.37
Normalized	100%	77.21%
Standard Deviation	2.8%	7.3%
NCNN - CPU-v2-v2 - mobilenet-v2 (ms)	7.78	10.67
Normalized	100%	72.91%
Standard Deviation	1.2%	18.8%
NCNN - CPU-v3-v3 - mobilenet-v3 (ms)	6.71	9.10
Normalized	100%	73.74%
Standard Deviation	0.5%	13%
NCNN - CPU - shufflenet-v2 (ms)	9.99	14.58
Normalized	100%	68.52%
Standard Deviation	10.2%	15.8%
NCNN - CPU - mnasnet (ms)	8.10	9.76
Normalized	100%	82.99%
Standard Deviation	10.8%	13.9%
NCNN - CPU - efficientnet-b0 (ms)	12.53	15.22
Normalized	100%	82.33%
Standard Deviation	0.5%	11.7%
NCNN - CPU - blazeface (ms)	2.85	3.80
Normalized	100%	75%
Standard Deviation	1.4%	10.3%
NCNN - CPU - googlenet (ms)	25.01	31.39
Normalized	100%	79.68%
Standard Deviation	0.8%	9.8%
NCNN - CPU - vgg16 (ms)	68.50	108.71
Normalized	100%	63.01%

	Standard Deviation	0.6%	4.3%
NCNN - CPU - resnet18 (ms)		22.07	30.23
	Normalized	100%	73.01%
	Standard Deviation	0.2%	9.1%
NCNN - CPU - alexnet (ms)		19.06	26.55
	Normalized	100%	71.79%
	Standard Deviation	0.1%	5.7%
NCNN - CPU - resnet50 (ms)		51.01	67.20
	Normalized	100%	75.91%
	Standard Deviation	0.4%	11%
NCNN - CPU - yolov4-tiny (ms)		44.65	57.10
	Normalized	100%	78.2%
	Standard Deviation	2.1%	7.1%
NCNN - CPU - squeezenet_ssd (ms)		39.61	45.61
	Normalized	100%	86.84%
	Standard Deviation	1%	7.6%
NCNN - CPU - regnety_400m (ms)		21.11	22.34
	Normalized	100%	94.49%
	Standard Deviation	1.4%	9.7%
NCNN - Vulkan GPU - mobilenet (ms)		35.11	46.41
	Normalized	100%	75.65%
	Standard Deviation	2.4%	0.5%
NCNN - Vulkan GPU-v2-v2 - mobilenet-v2 (ms)		7.71	9.64
	Normalized	100%	79.98%
	Standard Deviation	0.6%	24.2%
NCNN - Vulkan GPU-v3-v3 - mobilenet-v3 (ms)		6.70	8.28
	Normalized	100%	80.92%
	Standard Deviation	0.2%	20.2%
NCNN - Vulkan GPU - shufflenet-v2 (ms)		10.30	12.80
	Normalized	100%	80.47%
	Standard Deviation	11.2%	28.3%
NCNN - Vulkan GPU - mnasnet (ms)		8.13	8.77
	Normalized	100%	92.7%
	Standard Deviation	11.4%	22.8%
NCNN - Vulkan GPU - efficientnet-b0 (ms)		11.83	13.74
	Normalized	100%	86.1%
	Standard Deviation	9.7%	19.8%
NCNN - Vulkan GPU - blazeface (ms)		2.80	3.24
	Normalized	100%	86.42%
	Standard Deviation	3.8%	27.7%
NCNN - Vulkan GPU - googlenet (ms)		24.58	28.85
	Normalized	100%	85.2%
	Standard Deviation	4.2%	15.1%
NCNN - Vulkan GPU - vgg16 (ms)		69.16	107.70
	Normalized	100%	64.22%
	Standard Deviation	1.3%	2.8%
NCNN - Vulkan GPU - resnet18 (ms)		22.25	28.33
	Normalized	100%	78.54%
	Standard Deviation	1.3%	10.9%
NCNN - Vulkan GPU - alexnet (ms)		19.23	25.75
	Normalized	100%	74.68%
	Standard Deviation	1.8%	6.2%
NCNN - Vulkan GPU - resnet50 (ms)		51.13	65.59
	Normalized	100%	77.95%
	Standard Deviation	0.4%	15.4%

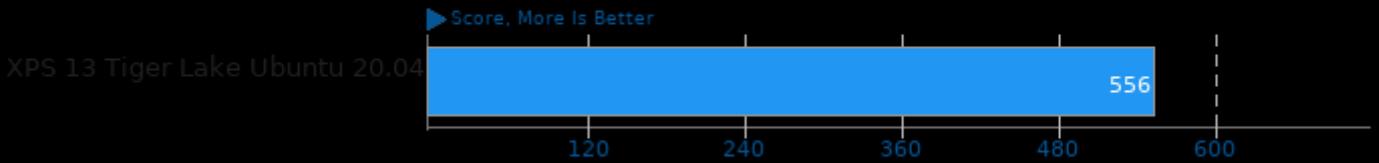
NCNN - Vulkan GPU - yolov4-tiny (ms)	44.68	58.16
Normalized	100%	76.82%
Standard Deviation	2.2%	4.7%
NCNN - Vulkan GPU - squeezeenet_ssd (ms)	39.53	44.71
Normalized	100%	88.41%
Standard Deviation	1%	9.6%
NCNN - Vulkan GPU - regnety_400m (ms)	21.01	20.14
Normalized	95.86%	100%
Standard Deviation	7.4%	18.7%
NCNN - CPUv2-yolov3v2-yolov3 - mobilenetv2-yolov3 (ms)		45.37
Standard Deviation		7.3%
NCNN - V.G.y.y - mobilenetv2-yolov3 (ms)		46.41
Standard Deviation		0.5%
Numenta Anomaly Benchmark - EXPoSE (sec)	1139	
Standard Deviation	0.4%	
Numenta Anomaly Benchmark - Relative Entropy	48.609	
Standard Deviation	3%	
Numenta Anomaly Benchmark - Windowed Gaussian (sec)	27.702	
Standard Deviation	2.6%	
Numenta Anomaly Benchmark - Earthgecko Skyline	333.639	
Standard Deviation	0.4%	
Numenta Anomaly Benchmark - B.C (sec)	81.062	
Standard Deviation	2.6%	
Numpy Benchmark (Score)	293.18	299.99
Normalized	97.73%	100%
Standard Deviation	0.3%	1.8%
oneDNN - IP Shapes 1D - f32 - CPU (ms)	10.2773	15.77
Normalized	100%	65.17%
Standard Deviation	1%	15.2%
oneDNN - IP Shapes 3D - f32 - CPU (ms)	6.59350	17.95
Normalized	100%	36.73%
Standard Deviation	0.4%	0.1%
oneDNN - IP Shapes 1D - u8s8f32 - CPU (ms)	2.43051	5.88
Normalized	100%	41.34%
Standard Deviation	3.6%	0.8%
oneDNN - IP Shapes 3D - u8s8f32 - CPU (ms)	2.62981	3.95
Normalized	100%	66.58%
Standard Deviation	0.1%	2.5%
oneDNN - IP Shapes 1D - bf16bf16bf16 - CPU (ms)	25.8282	
Standard Deviation	0.7%	
oneDNN - IP Shapes 3D - bf16bf16bf16 - CPU (ms)	6.39979	
Standard Deviation	1.7%	
oneDNN - C.B.S.A - f32 - CPU (ms)	11.5632	22.02
Normalized	100%	52.51%
Standard Deviation	4.3%	0.1%
oneDNN - D.B.s - f32 - CPU (ms)	14.6326	20.71
Normalized	100%	70.65%
Standard Deviation	2.9%	13.8%
oneDNN - D.B.s - f32 - CPU (ms)	13.4125	15.85
Normalized	100%	84.62%
Standard Deviation	0.7%	2.9%
oneDNN - C.B.S.A - u8s8f32 - CPU (ms)	7.93450	21.14

	Normalized	100%	37.53%
	Standard Deviation	0.4%	1.3%
oneDNN - D.B.s - u8s8f32 - CPU (ms)		2.94262	21.88
	Normalized	100%	13.45%
	Standard Deviation	3.4%	0.7%
oneDNN - D.B.s - u8s8f32 - CPU (ms)		3.13684	8.18
	Normalized	100%	38.35%
	Standard Deviation	0.2%	1.9%
oneDNN - R.N.N.T - f32 - CPU (ms)		8863	15730
	Normalized	100%	56.34%
	Standard Deviation	0.1%	6.3%
oneDNN - R.N.N.I - f32 - CPU (ms)		5736	8215
	Normalized	100%	69.83%
	Standard Deviation	73.4%	7.4%
oneDNN - R.N.N.T - u8s8f32 - CPU (ms)		8875	15829
	Normalized	100%	56.07%
	Standard Deviation	0.4%	6.4%
oneDNN - C.B.S.A - bf16bf16bf16 - CPU (ms)		52.4284	
	Standard Deviation	0.1%	
oneDNN - D.B.s - bf16bf16bf16 - CPU (ms)		57.0546	
	Standard Deviation	2.8%	
oneDNN - D.B.s - bf16bf16bf16 - CPU (ms)		52.6824	
	Standard Deviation	0.2%	
oneDNN - R.N.N.I - u8s8f32 - CPU (ms)		4528	8004
	Normalized	100%	56.57%
	Standard Deviation	0.4%	9.7%
oneDNN - M.M.B.S.T - f32 - CPU (ms)		3.93298	6.73
	Normalized	100%	58.44%
	Standard Deviation	0.7%	9.6%
oneDNN - R.N.N.T - bf16bf16bf16 - CPU (ms)		8859	15837
	Normalized	100%	55.94%
	Standard Deviation	0%	7%
oneDNN - R.N.N.I - bf16bf16bf16 - CPU (ms)		4517	7969
	Normalized	100%	56.68%
	Standard Deviation	0.2%	9.8%
oneDNN - M.M.B.S.T - u8s8f32 - CPU (ms)		2.12307	11.82
	Normalized	100%	17.96%
	Standard Deviation	0.6%	12.1%
oneDNN - M.M.B.S.T - bf16bf16bf16 - CPU (ms)		11.8072	
	Standard Deviation	3.6%	
OpenCV - DNN - D.N.N (ms)		5351	7968
	Normalized	100%	67.16%
	Standard Deviation	2.8%	42.3%
PlaidML - No - Inference - VGG16 - CPU (FPS)		6.47	3.74
	Normalized	100%	57.81%
	Standard Deviation	0.6%	3%
PlaidML - No - Inference - ResNet 50 - CPU (FPS)		3.27	2.57
	Normalized	100%	78.59%
	Standard Deviation	0.8%	0.4%
RNNoise (sec)		31.947	
	Standard Deviation	0%	
Scikit-Learn (sec)		17.900	
	Standard Deviation	1.4%	
TensorFlow Lite - SqueezeNet (us)		627487	769601
	Normalized	100%	81.53%

	Standard Deviation	2.8%	7.9%
TensorFlow Lite - Inception V4 (us)		9220230	11134467
	Normalized	100%	82.81%
	Standard Deviation	0.6%	0.3%
TensorFlow Lite - NASNet Mobile (us)		455525	597569
	Normalized	100%	76.23%
	Standard Deviation	1.7%	6.7%
TensorFlow Lite - Mobilenet Float (us)		424309	525312
	Normalized	100%	80.77%
	Standard Deviation	1.2%	7.5%
TensorFlow Lite - Mobilenet Quant (us)		419506	535897
	Normalized	100%	78.28%
	Standard Deviation	2%	7.6%
TensorFlow Lite - I.R.V (us)		8329360	10177933
	Normalized	100%	81.84%
	Standard Deviation	0.7%	1.1%

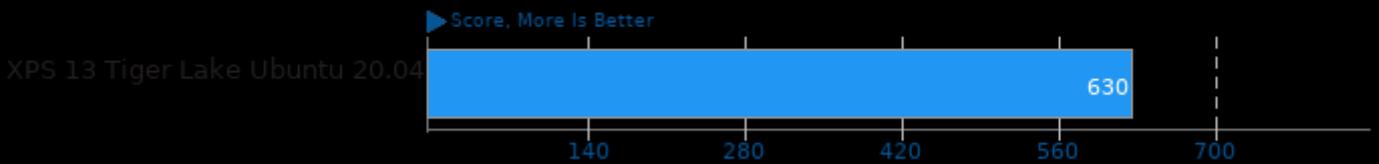
AI Benchmark Alpha 0.1.2

Device Inference Score



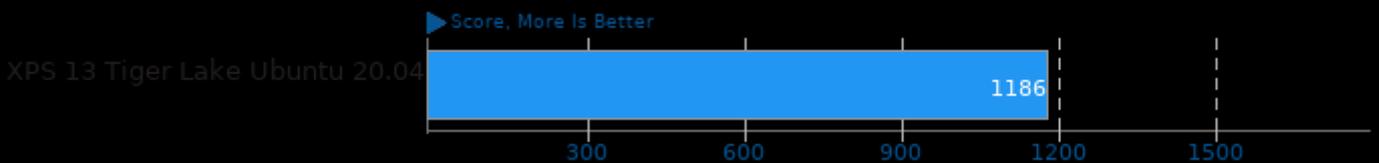
AI Benchmark Alpha 0.1.2

Device Training Score



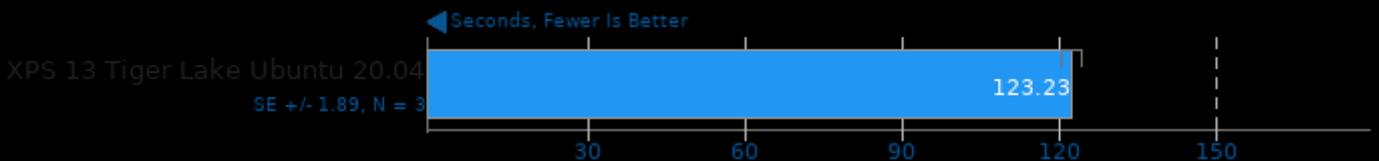
AI Benchmark Alpha 0.1.2

Device AI Score



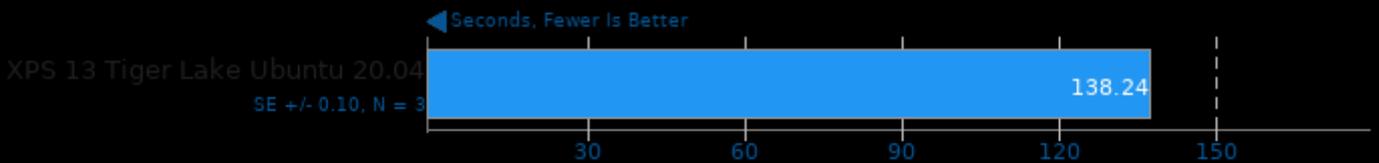
Mlpack Benchmark

Benchmark: scikit_ica



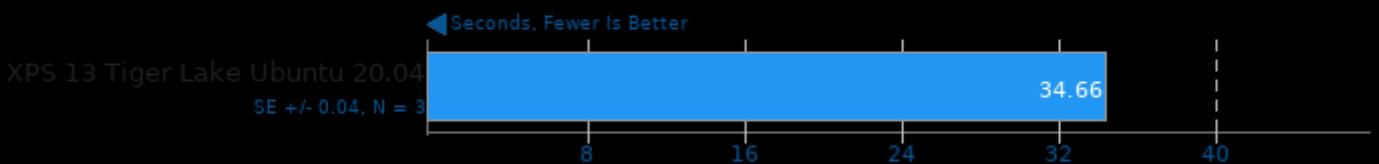
Mlpack Benchmark

Benchmark: scikit_qda



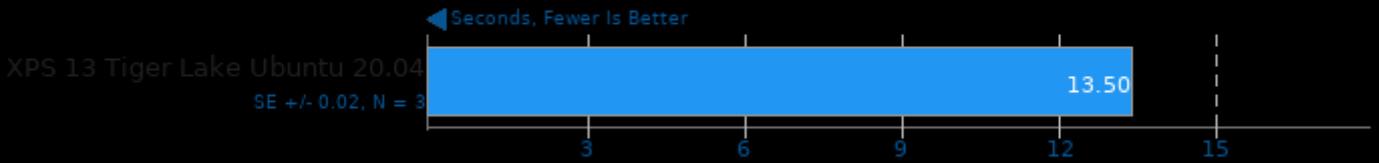
Mlpack Benchmark

Benchmark: scikit_svm



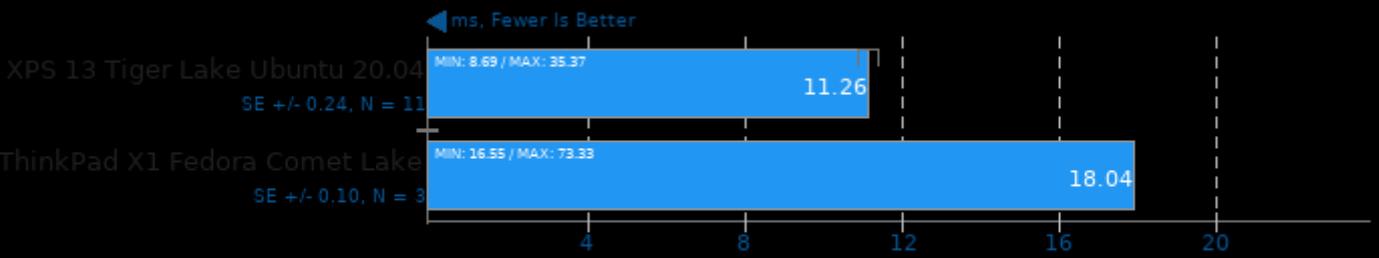
Mlpack Benchmark

Benchmark: scikit_linearidgegression



Mobile Neural Network 2020-09-17

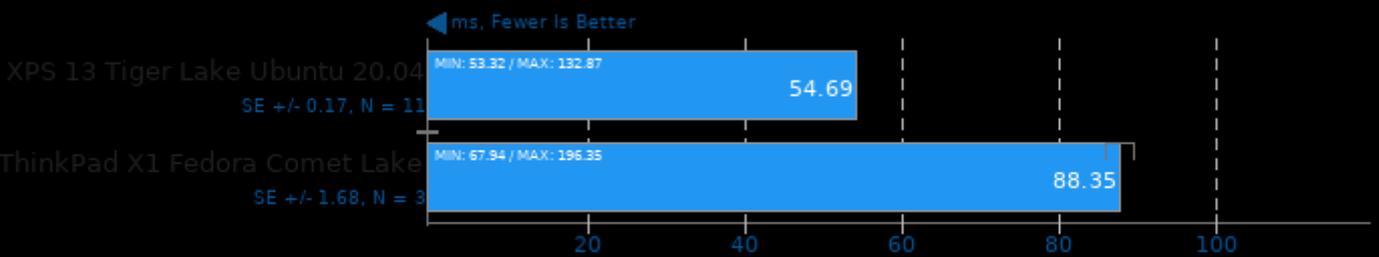
Model: SqueezeNetV1.0



1. (CXX) g++ options: -std=c++11 -O3 -fvisibility=hidden -fomit-frame-pointer -fstrict-aliasing -ffunction-sections -fdata-sections -ffast-math -fno-rtti -fno-

Mobile Neural Network 2020-09-17

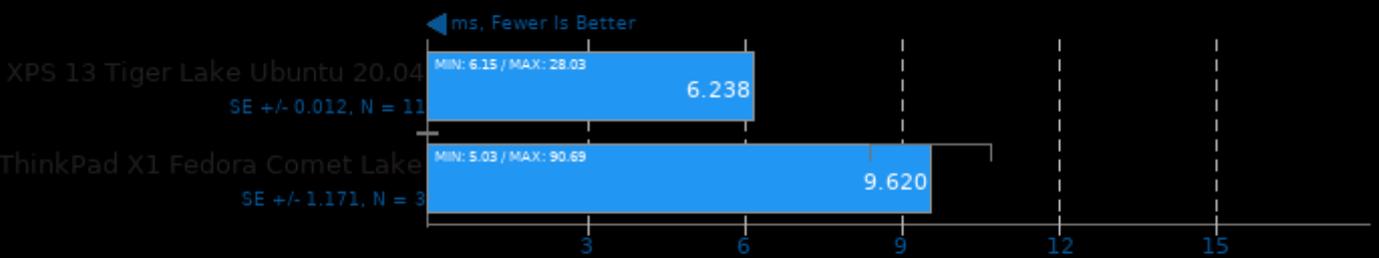
Model: resnet-v2-50



1. (CXX) g++ options: -std=c++11 -O3 -fvisibility=hidden -fomit-frame-pointer -fstrict-aliasing -ffunction-sections -fdata-sections -ffast-math -fno-rtti -fno-

Mobile Neural Network 2020-09-17

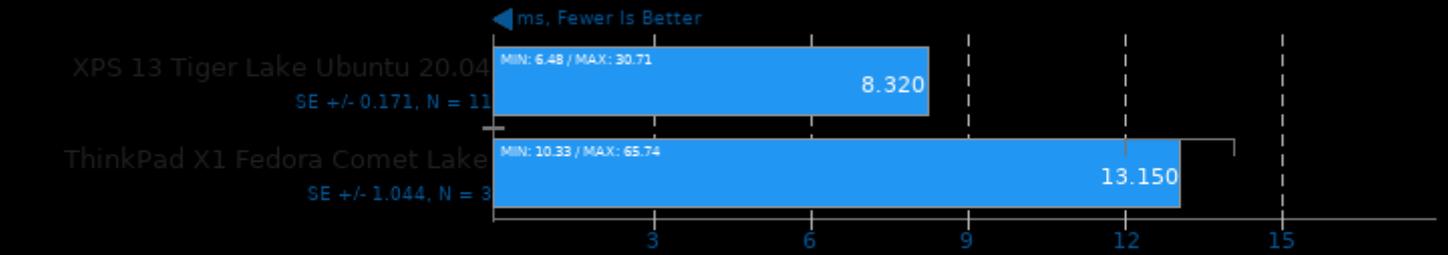
Model: MobileNetV2_224



1. (CXX) g++ options: -std=c++11 -O3 -fvisibility=hidden -fomit-frame-pointer -fstrict-aliasing -ffunction-sections -fdata-sections -ffast-math -fno-rtti -fno-

Mobile Neural Network 2020-09-17

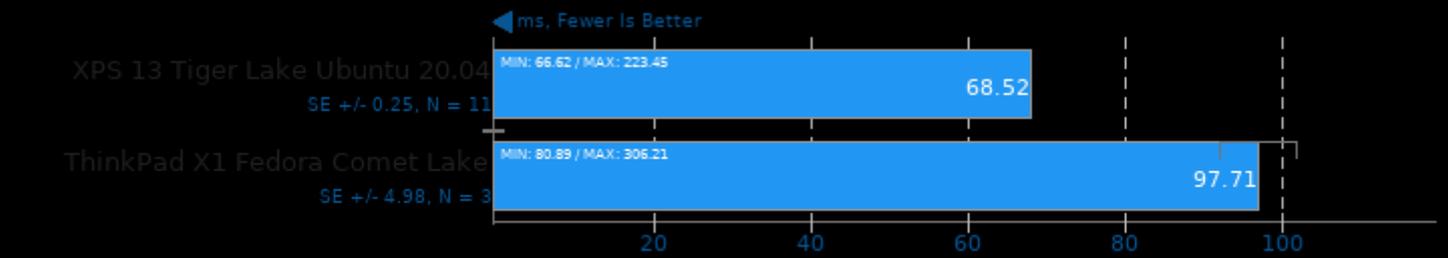
Model: mobilenet-v1-1.0



1. (CXX) g++ options: -std=c++11 -O3 -fvisibility=hidden -fomit-frame-pointer -fstrict-aliasing -ffunction-sections -fdata-sections -ffast-math -fno-rtti -fno-

Mobile Neural Network 2020-09-17

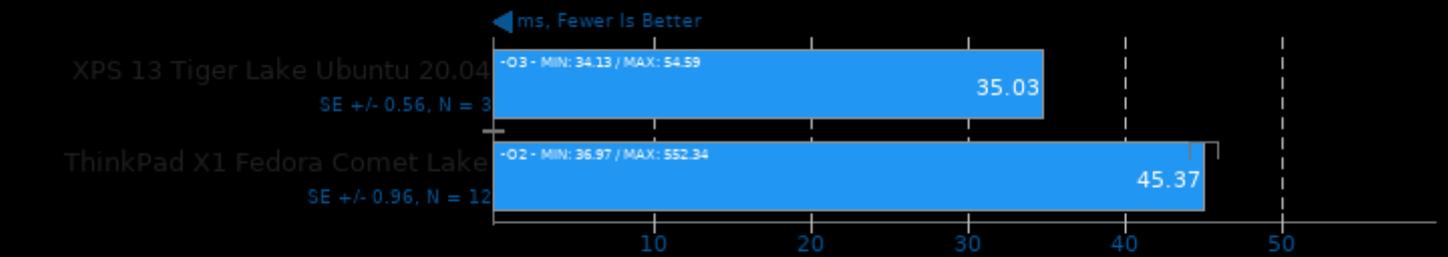
Model: inception-v3



1. (CXX) g++ options: -std=c++11 -O3 -fvisibility=hidden -fomit-frame-pointer -fstrict-aliasing -ffunction-sections -fdata-sections -ffast-math -fno-rtti -fno-

NCNN 20201218

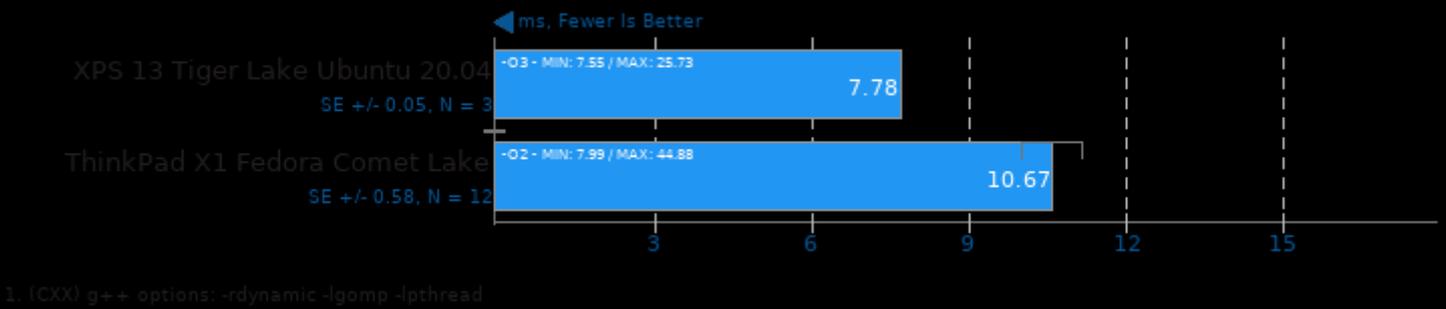
Target: CPU - Model: mobilenet



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

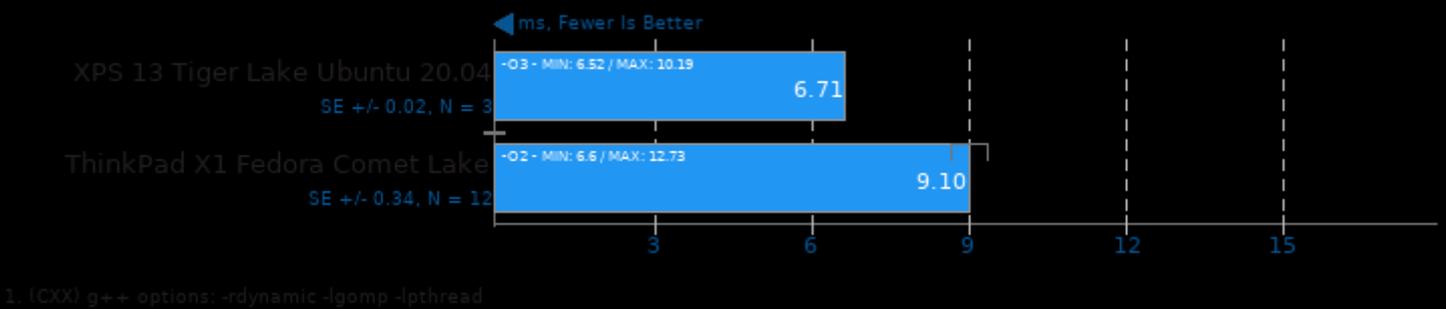
NCNN 20201218

Target: CPU-v2-v2 - Model: mobilenet-v2



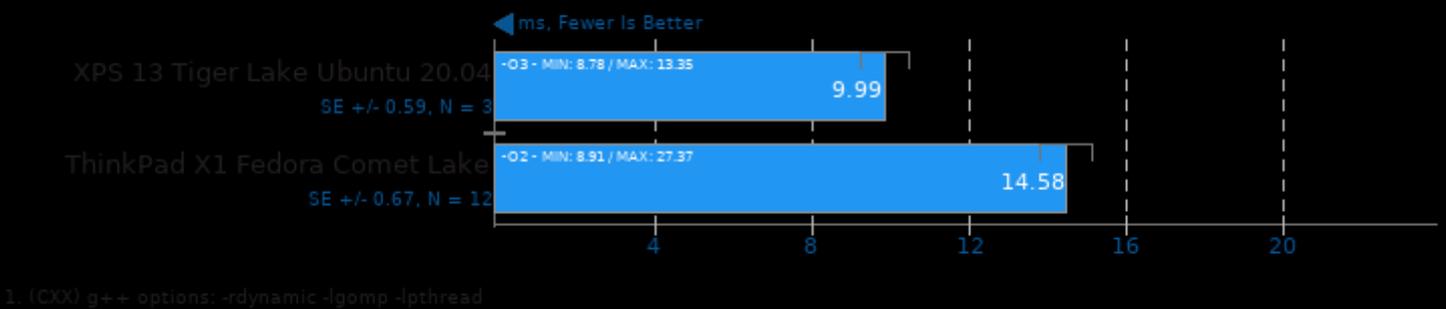
NCNN 20201218

Target: CPU-v3-v3 - Model: mobilenet-v3



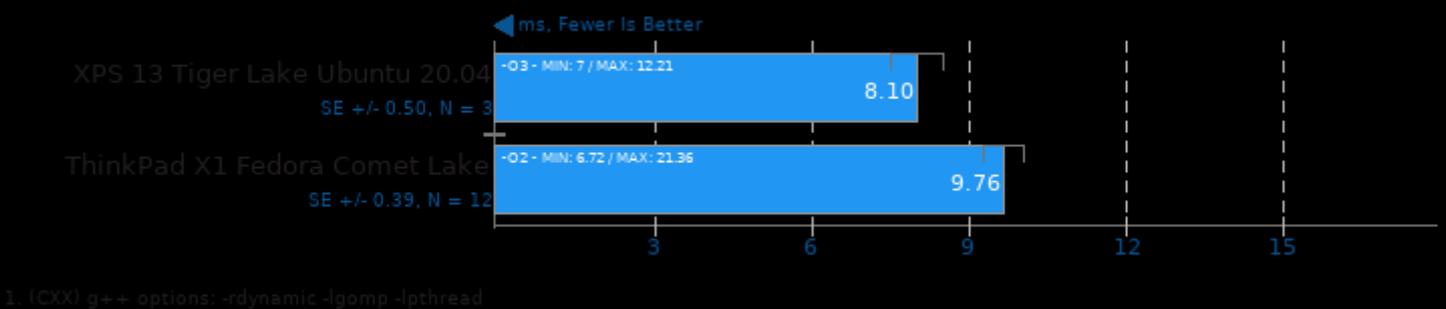
NCNN 20201218

Target: CPU - Model: shufflenet-v2



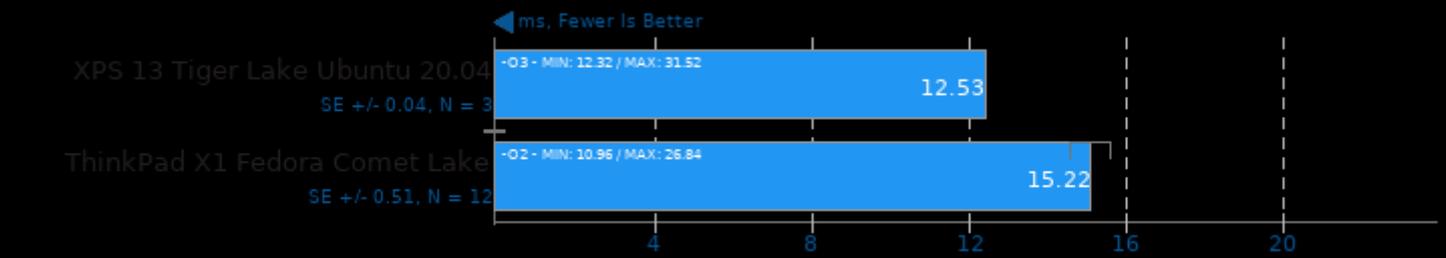
NCNN 20201218

Target: CPU - Model: mnasnet



NCNN 20201218

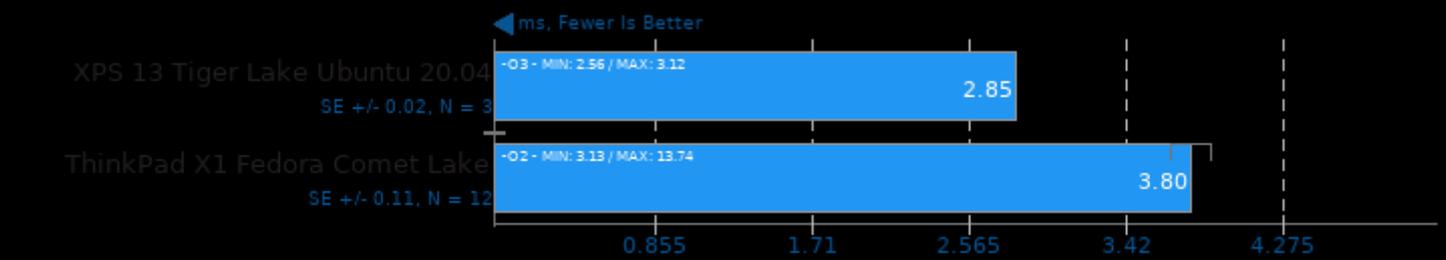
Target: CPU - Model: efficientnet-b0



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

NCNN 20201218

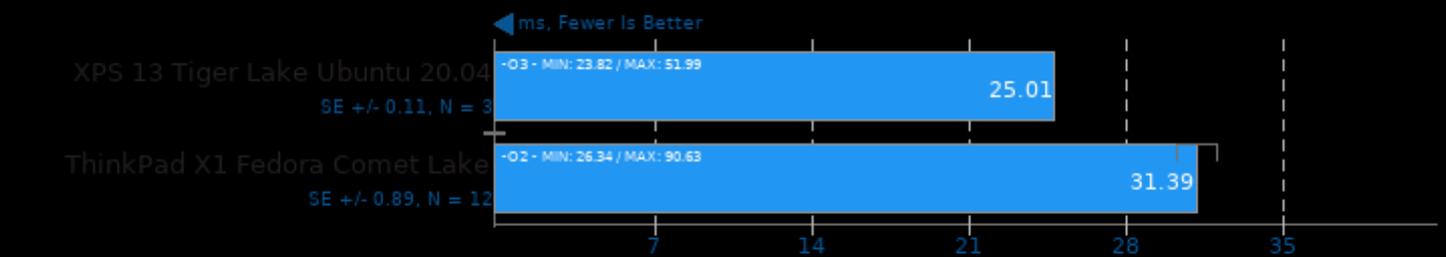
Target: CPU - Model: blazeface



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

NCNN 20201218

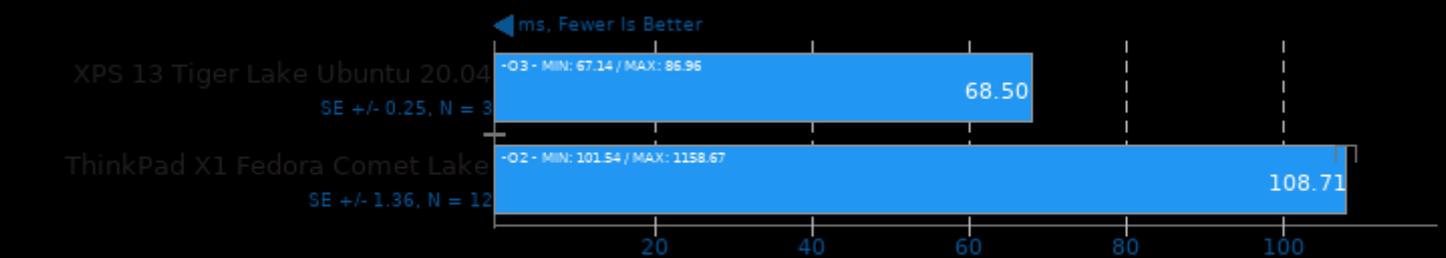
Target: CPU - Model: googlenet



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

NCNN 20201218

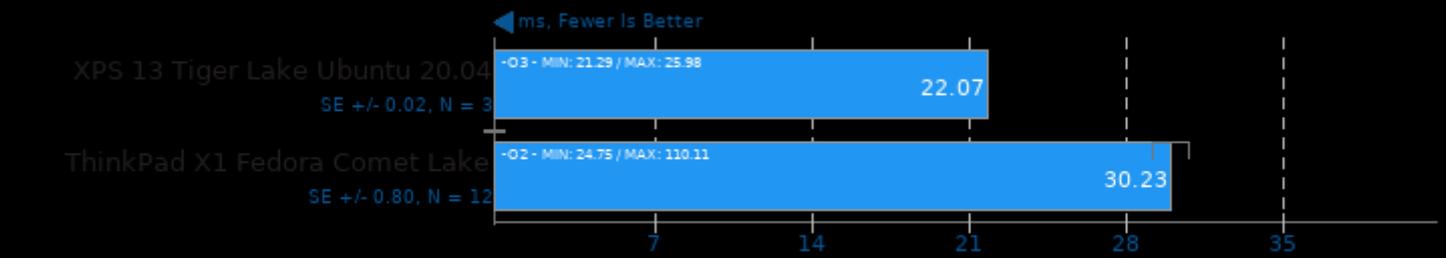
Target: CPU - Model: vgg16



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

NCNN 20201218

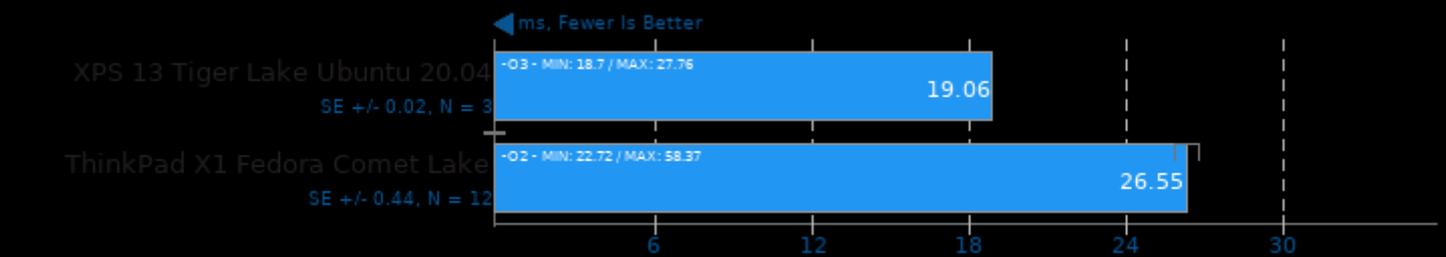
Target: CPU - Model: resnet18



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

NCNN 20201218

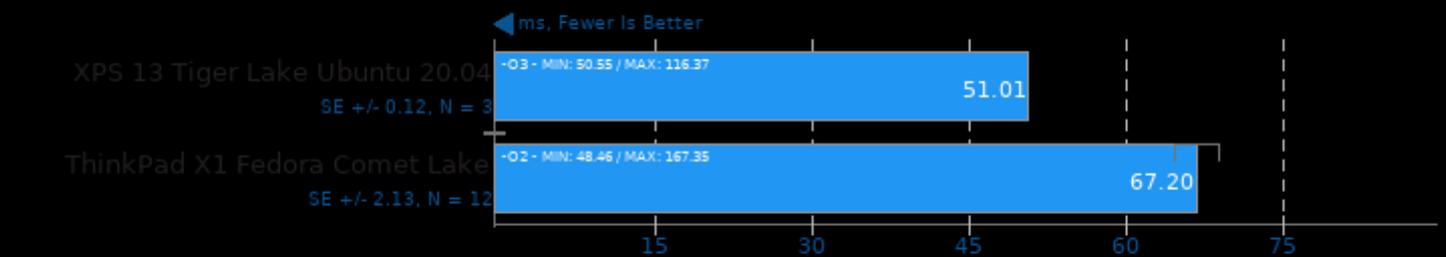
Target: CPU - Model: alexnet



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

NCNN 20201218

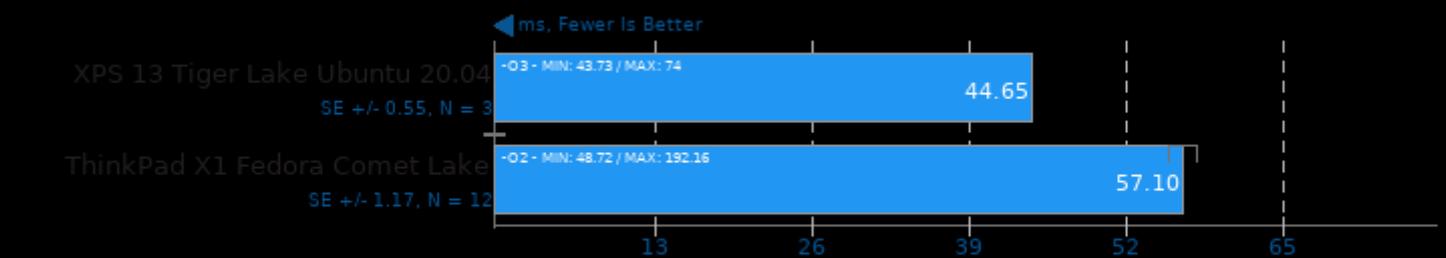
Target: CPU - Model: resnet50



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

NCNN 20201218

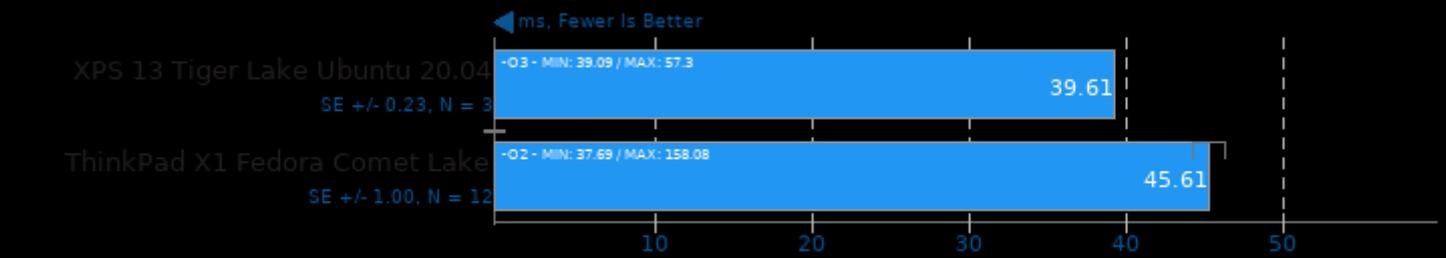
Target: CPU - Model: yolov4-tiny



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

NCNN 20201218

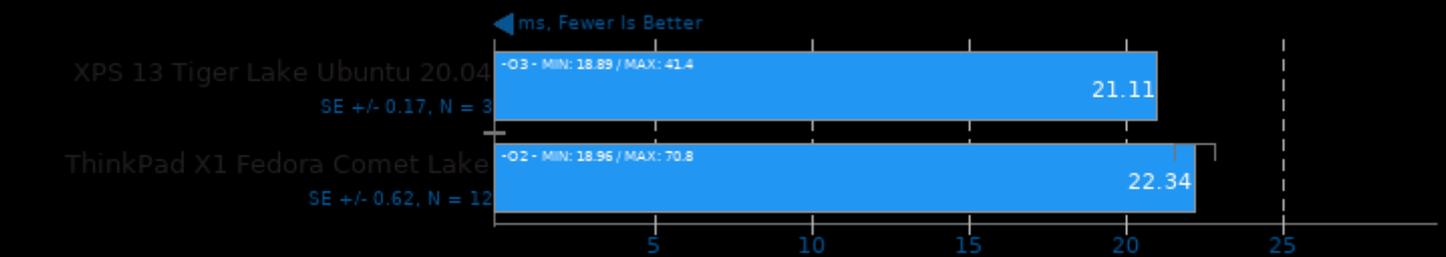
Target: CPU - Model: squeezenet_ssd



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

NCNN 20201218

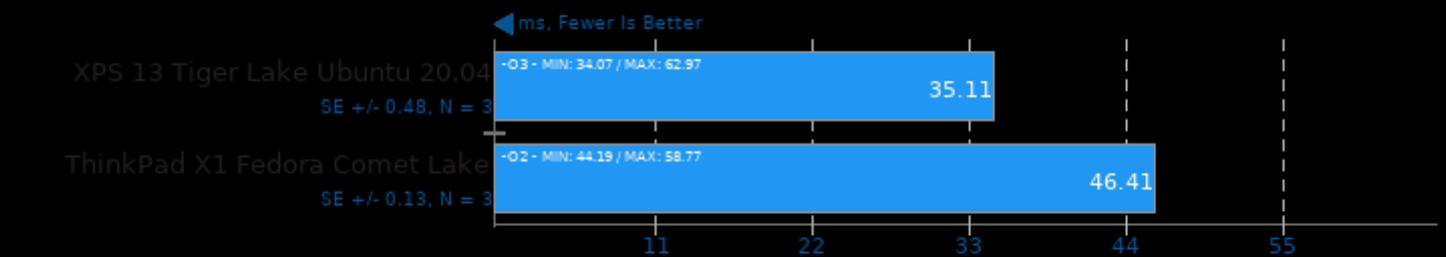
Target: CPU - Model: regnety_400m



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

NCNN 20201218

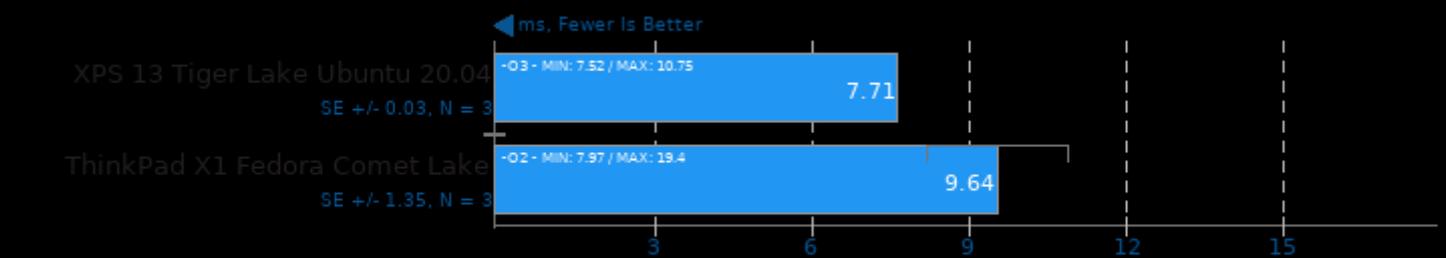
Target: Vulkan GPU - Model: mobilenet



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

NCNN 20201218

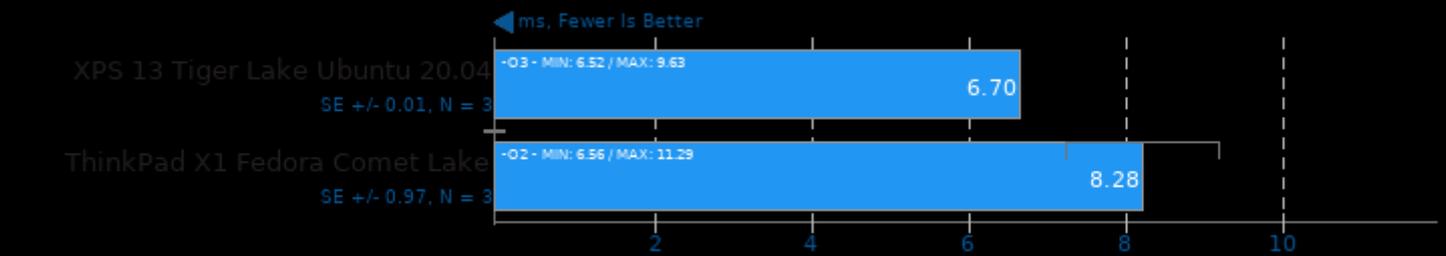
Target: Vulkan GPU-v2-v2 - Model: mobilenet-v2



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

NCNN 20201218

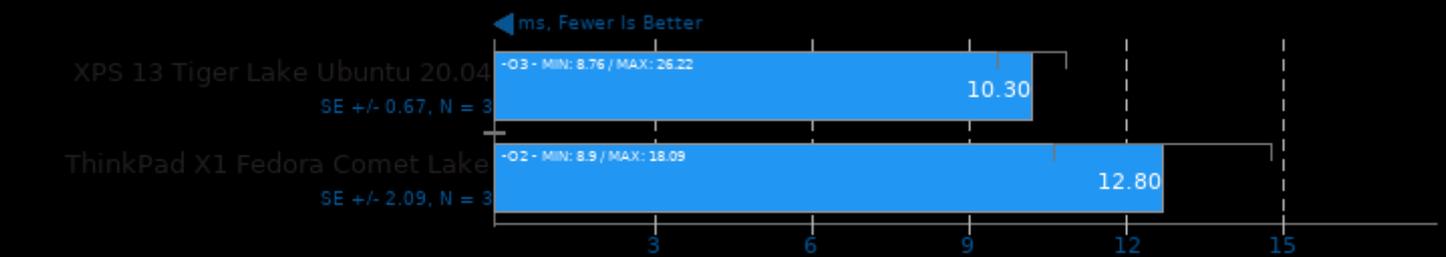
Target: Vulkan GPU-v3-v3 - Model: mobilenet-v3



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

NCNN 20201218

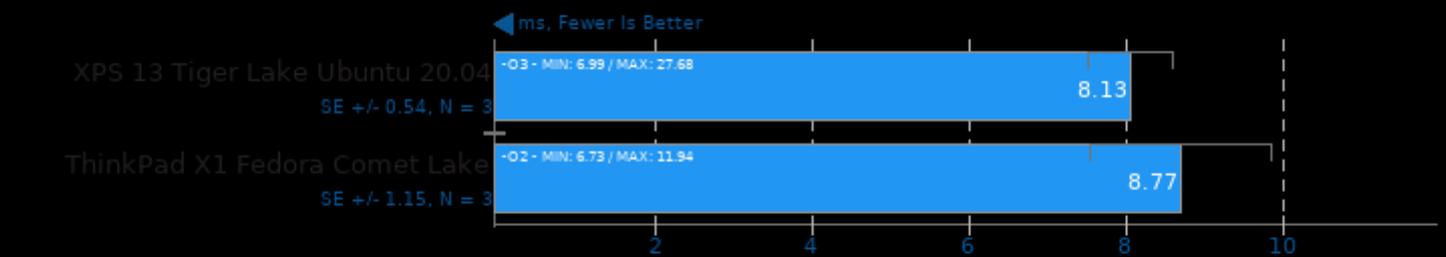
Target: Vulkan GPU - Model: shufflenet-v2



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

NCNN 20201218

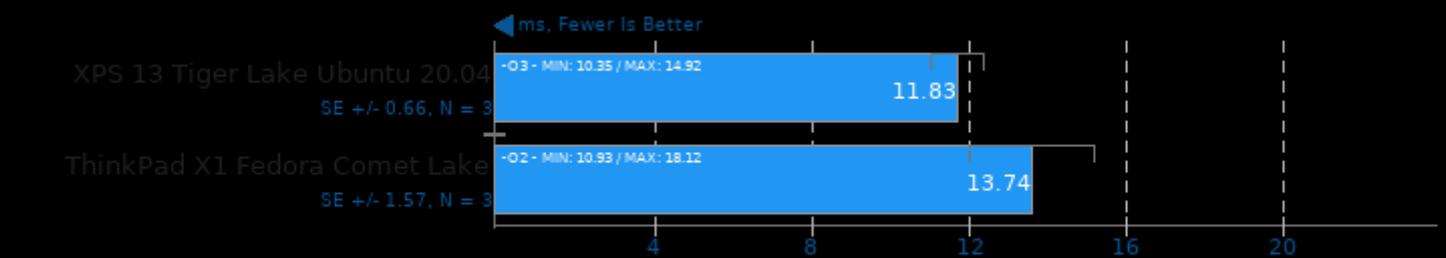
Target: Vulkan GPU - Model: mnasnet



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

NCNN 20201218

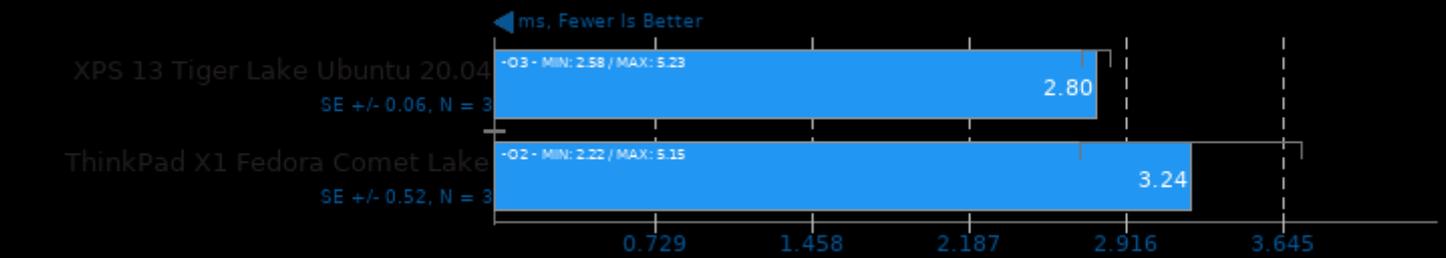
Target: Vulkan GPU - Model: efficientnet-b0



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

NCNN 20201218

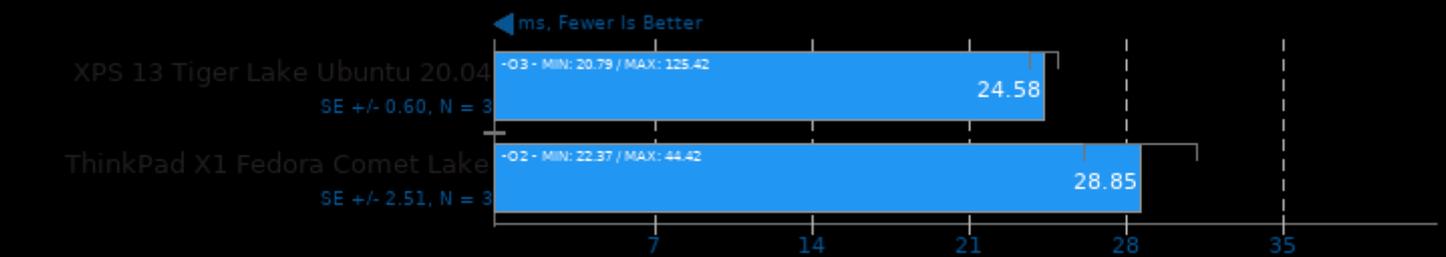
Target: Vulkan GPU - Model: blazeface



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

NCNN 20201218

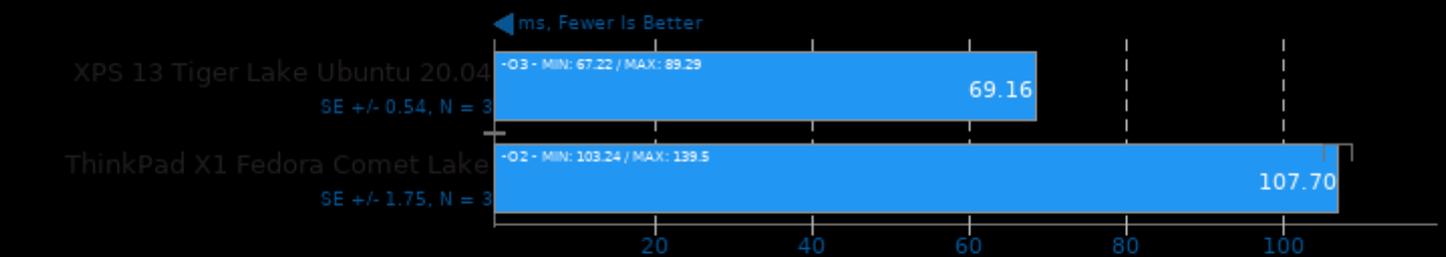
Target: Vulkan GPU - Model: googlenet



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

NCNN 20201218

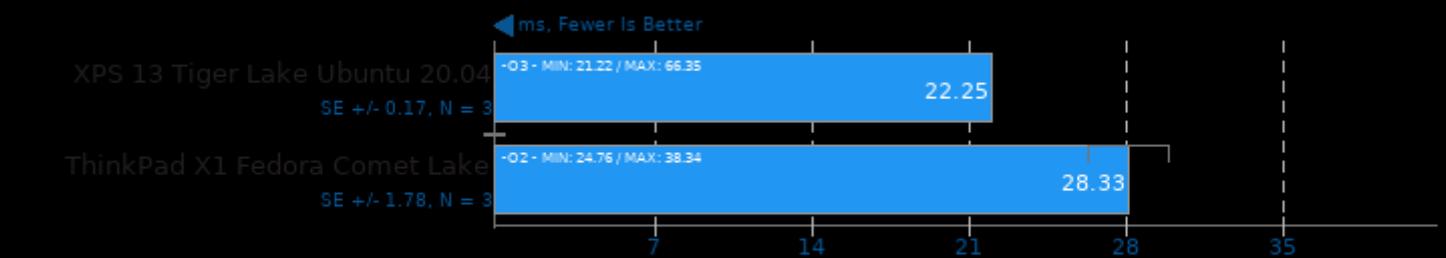
Target: Vulkan GPU - Model: vgg16



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

NCNN 20201218

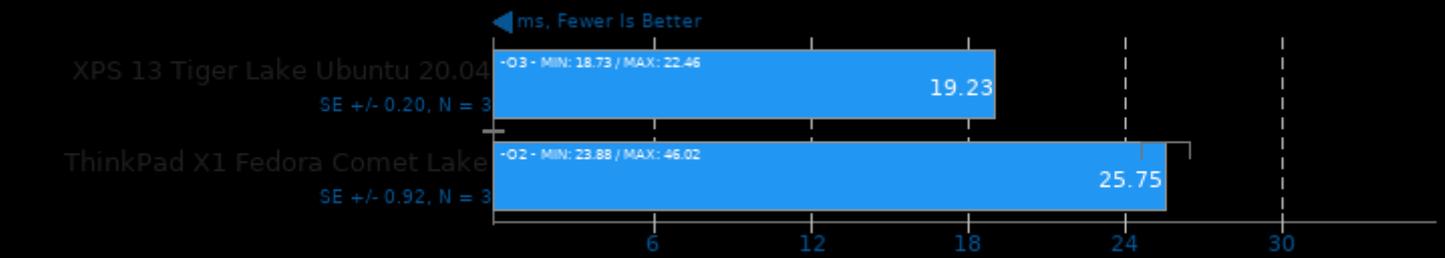
Target: Vulkan GPU - Model: resnet18



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

NCNN 20201218

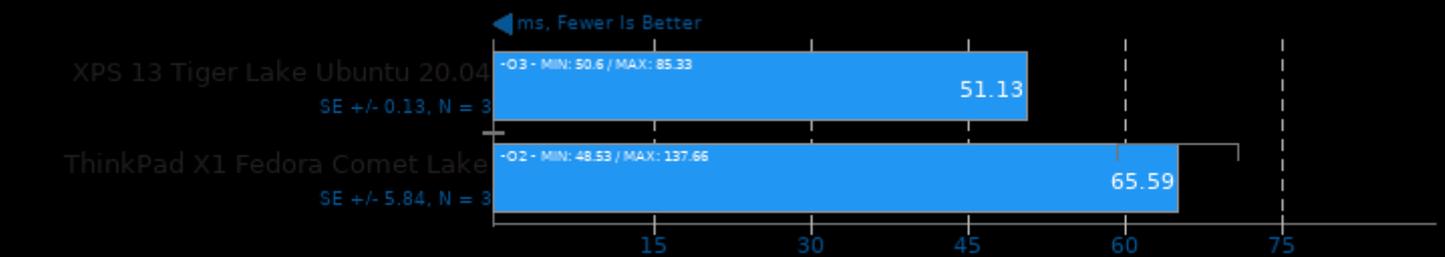
Target: Vulkan GPU - Model: alexnet



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

NCNN 20201218

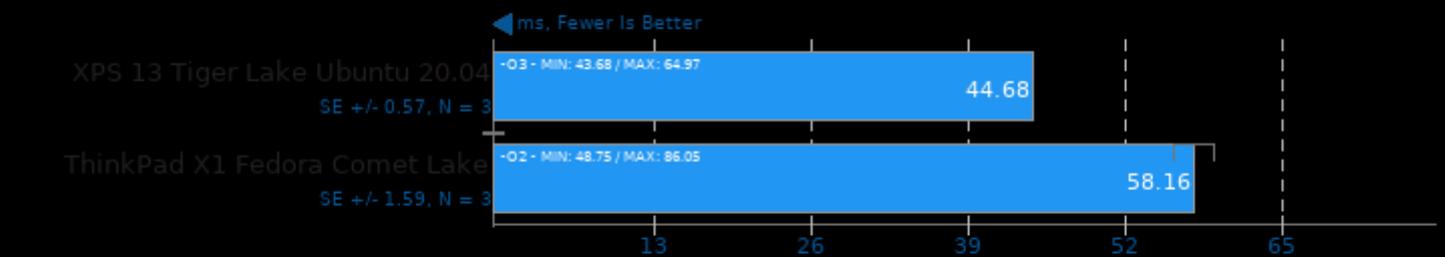
Target: Vulkan GPU - Model: resnet50



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

NCNN 20201218

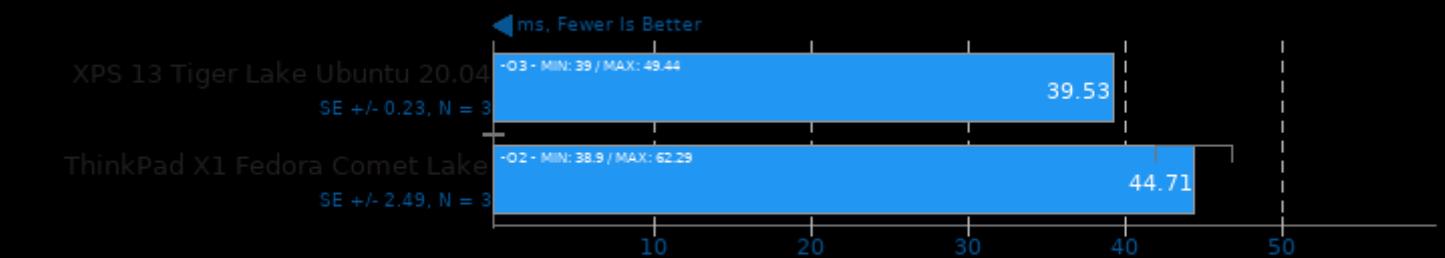
Target: Vulkan GPU - Model: yolov4-tiny



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

NCNN 20201218

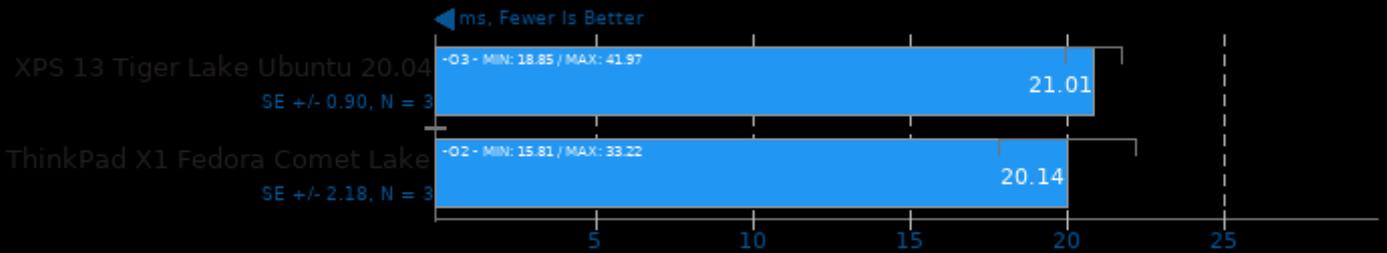
Target: Vulkan GPU - Model: squeezenet_ssd



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

NCNN 20201218

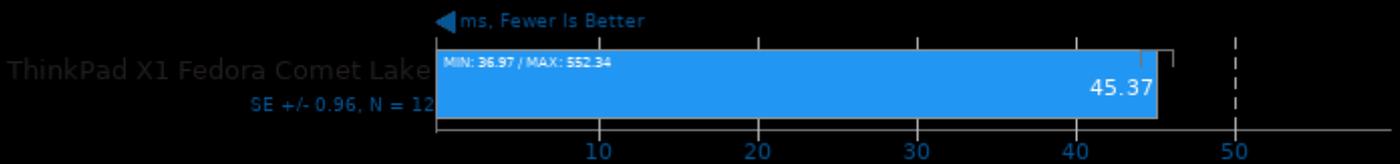
Target: Vulkan GPU - Model: regnety_400m



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

NCNN 20201218

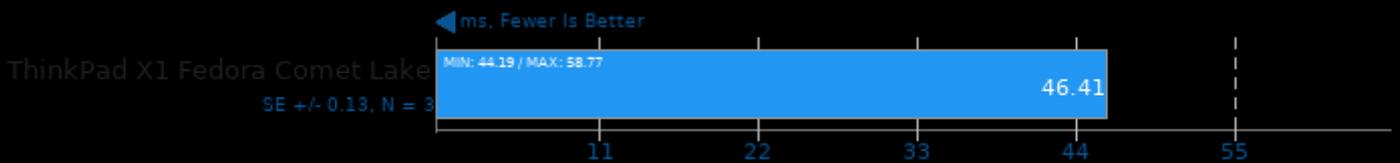
Target: CPUv2-yolov3v2-yolov3 - Model: mobilenetv2-yolov3



1. (CXX) g++ options: -O2 -rdynamic -lgomp -lpthread

NCNN 20201218

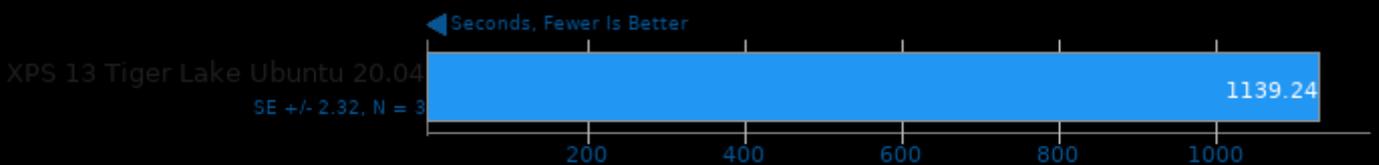
Target: Vulkan GPUv2-yolov3v2-yolov3 - Model: mobilenetv2-yolov3



1. (CXX) g++ options: -O2 -rdynamic -lgomp -lpthread

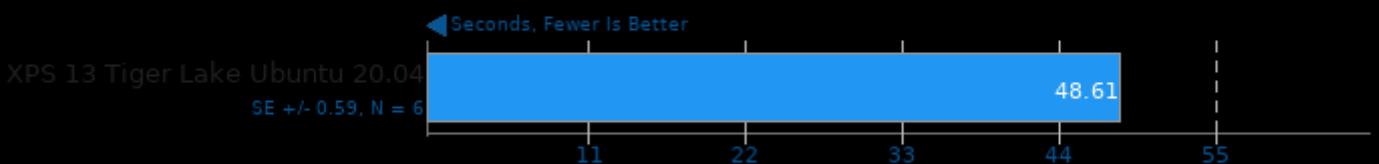
Numenta Anomaly Benchmark 1.1

Detector: EXPoSE



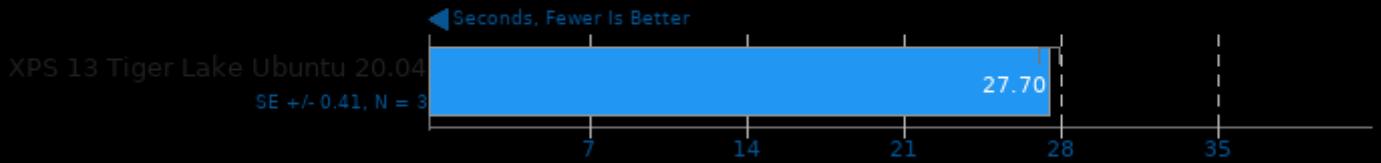
Numenta Anomaly Benchmark 1.1

Detector: Relative Entropy



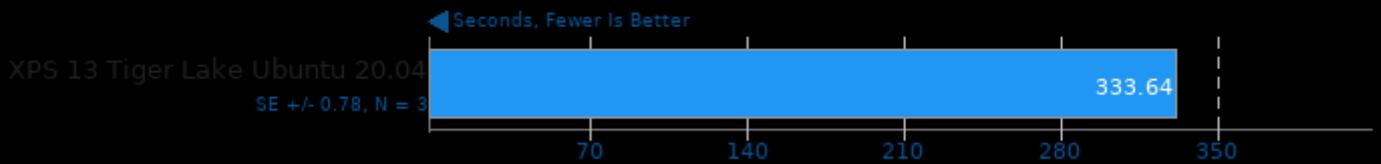
Numenta Anomaly Benchmark 1.1

Detector: Windowed Gaussian



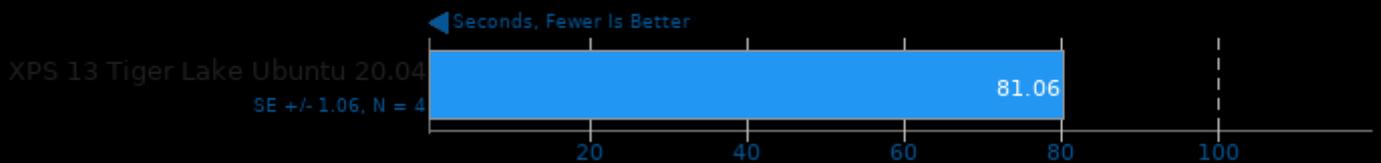
Numenta Anomaly Benchmark 1.1

Detector: Earthgecko Skyline

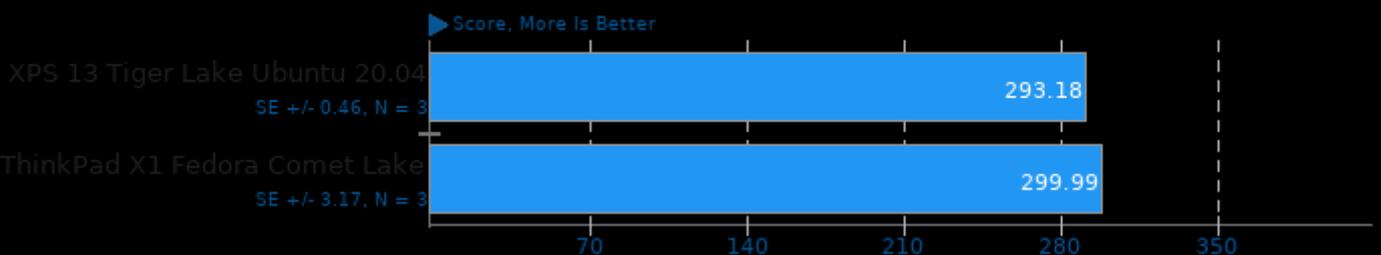


Numenta Anomaly Benchmark 1.1

Detector: Bayesian Changepoint

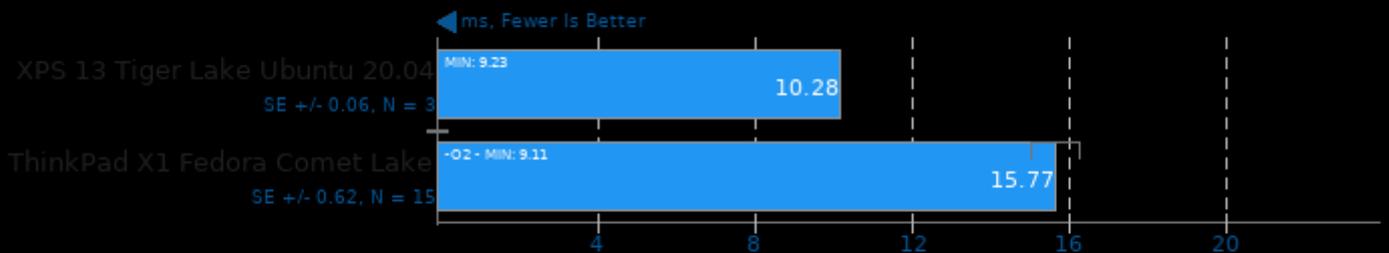


Numpy Benchmark



oneDNN 2.0

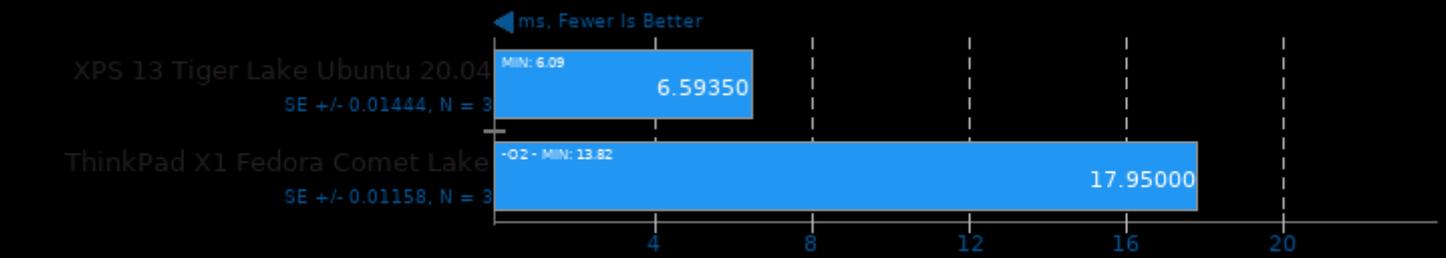
Harness: IP Shapes 1D - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -mssse4.1 -fPIC -pie -pthread

oneDNN 2.0

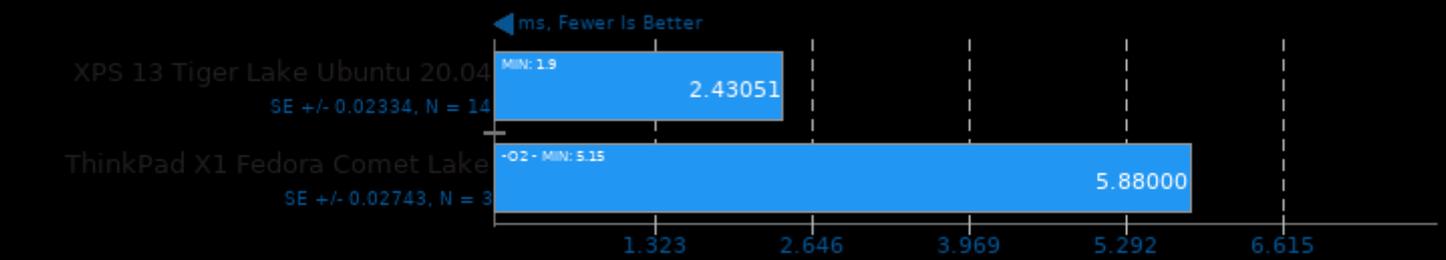
Harness: IP Shapes 3D - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

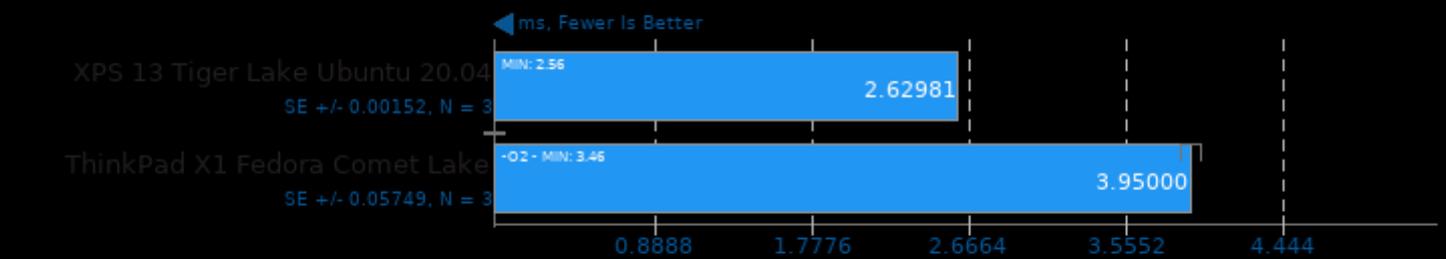
Harness: IP Shapes 1D - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

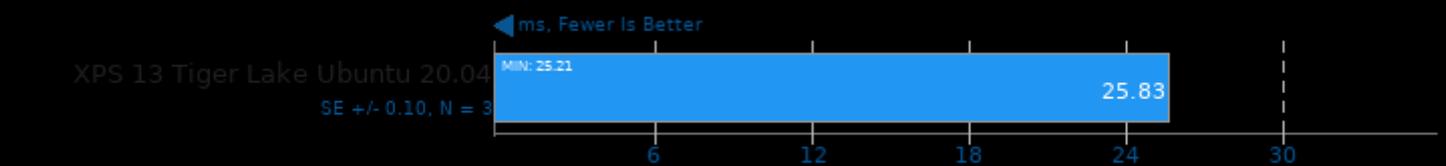
Harness: IP Shapes 3D - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

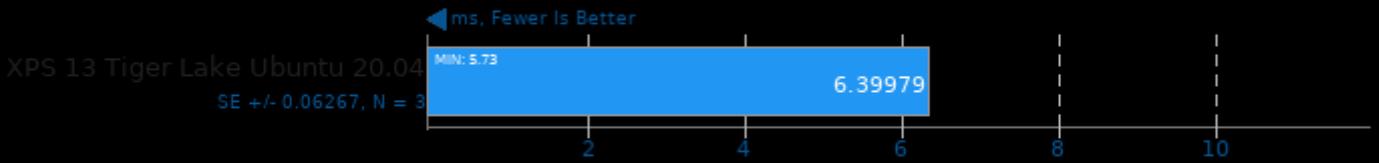
Harness: IP Shapes 1D - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

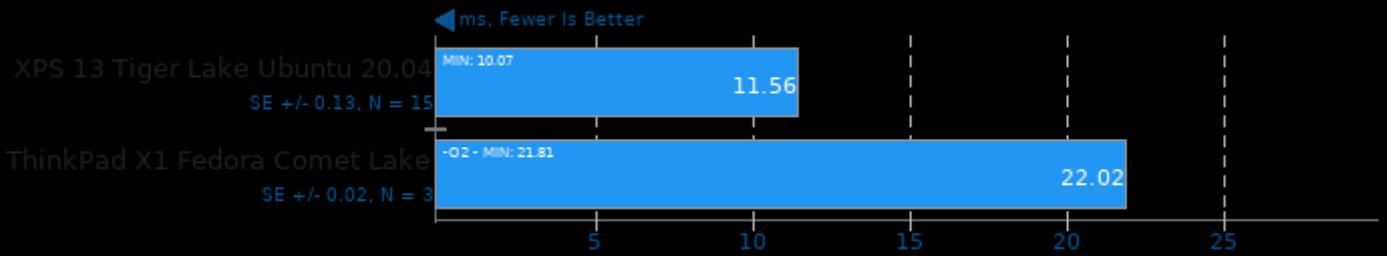
Harness: IP Shapes 3D - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

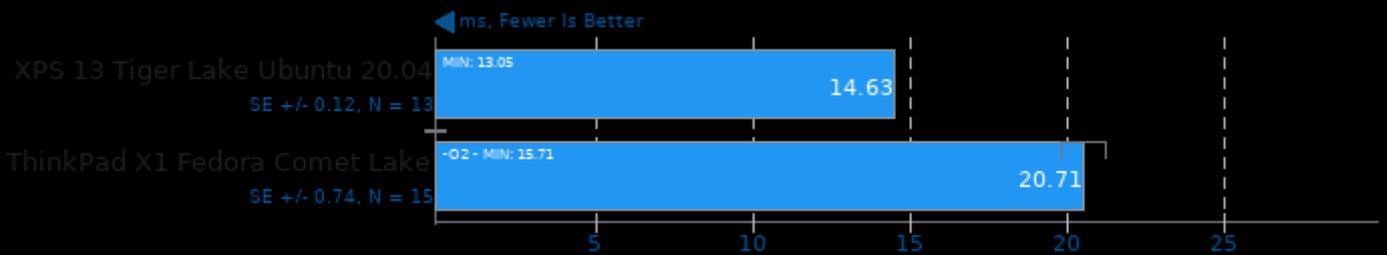
Harness: Convolution Batch Shapes Auto - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

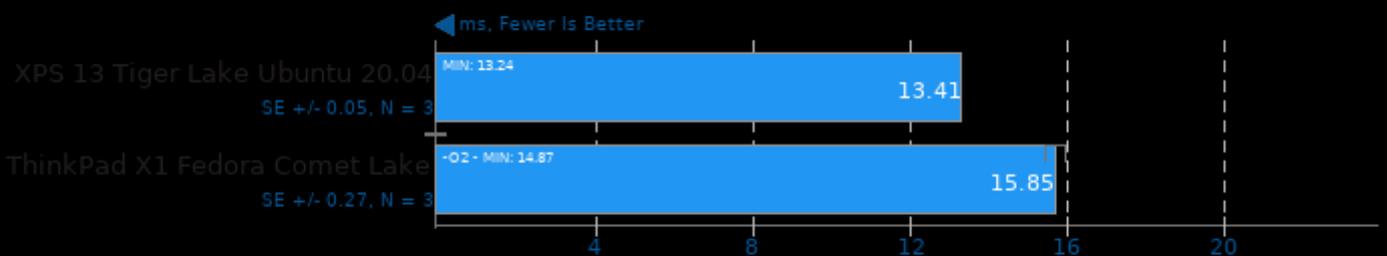
Harness: Deconvolution Batch shapes_1d - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

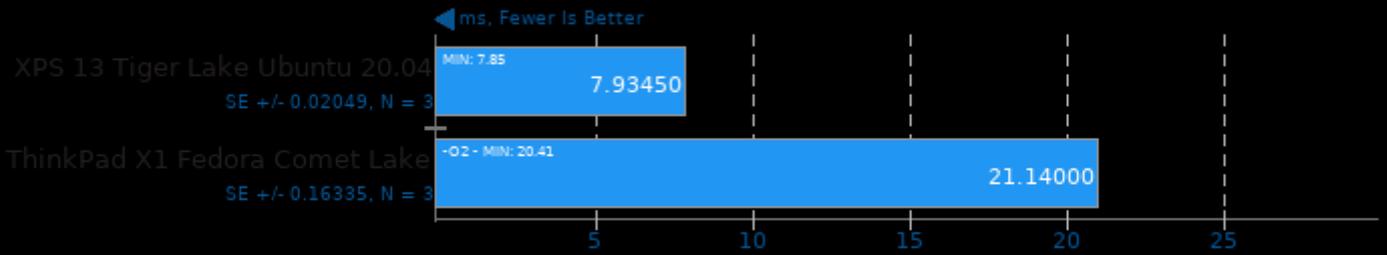
Harness: Deconvolution Batch shapes_3d - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

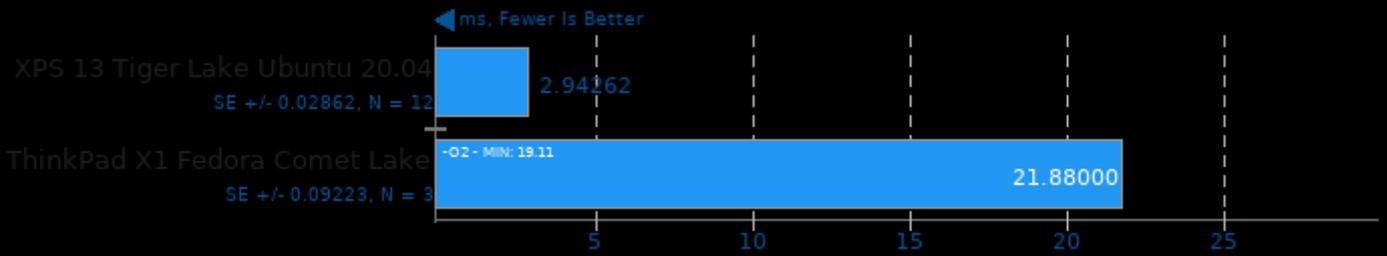
Harness: Convolution Batch Shapes Auto - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

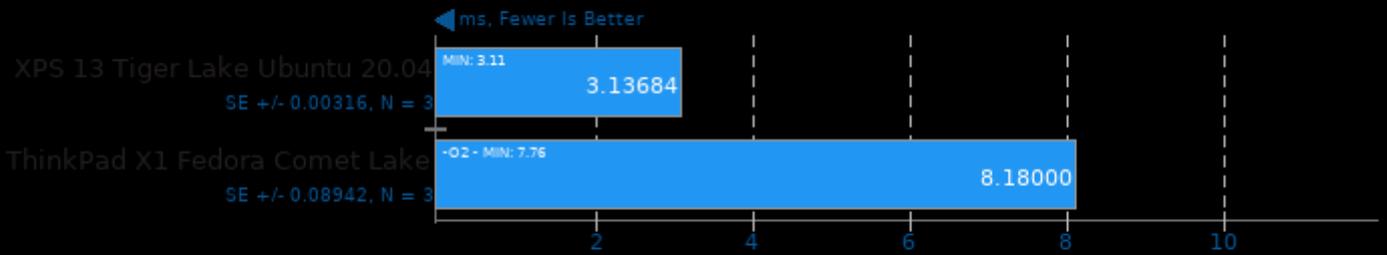
Harness: Deconvolution Batch shapes_1d - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

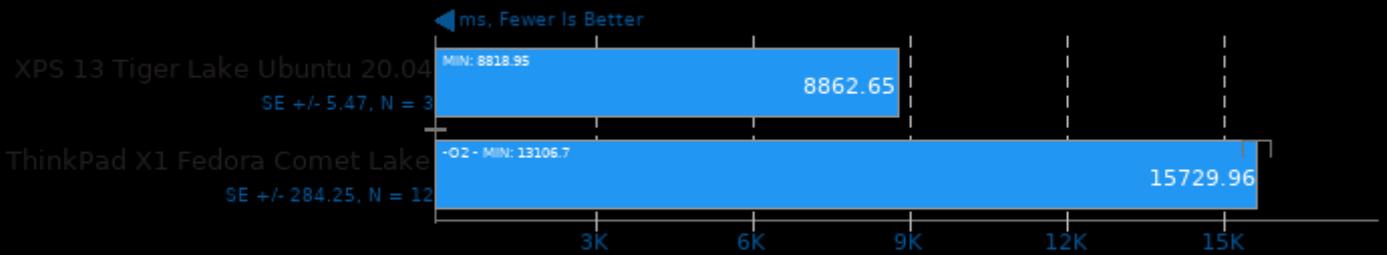
Harness: Deconvolution Batch shapes_3d - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

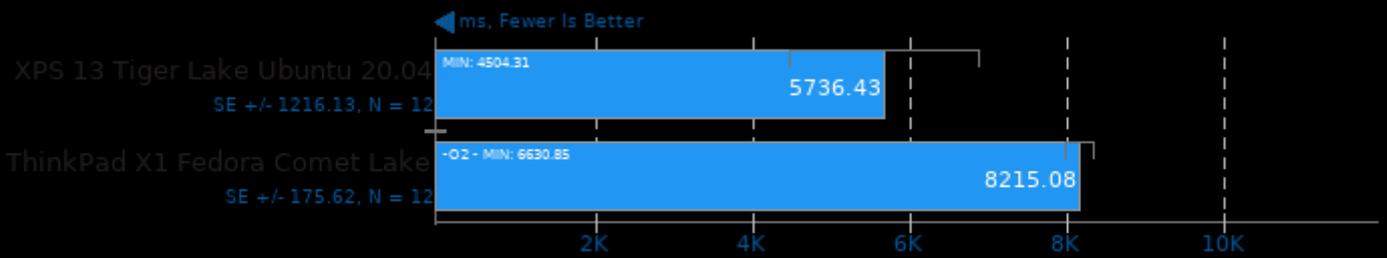
Harness: Recurrent Neural Network Training - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

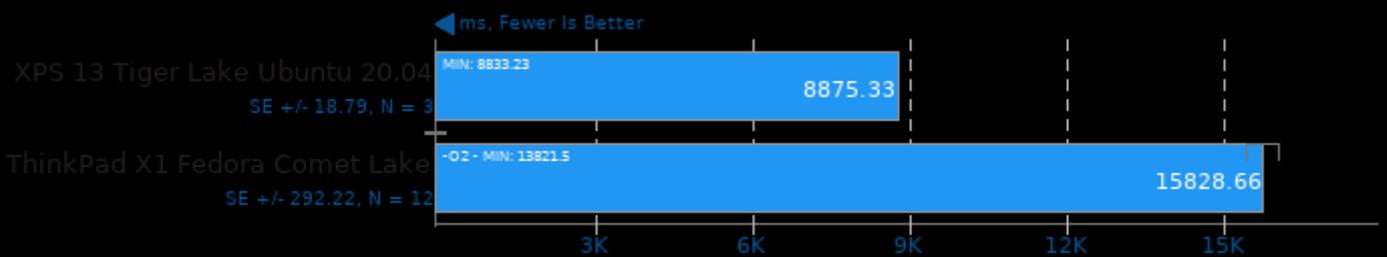
Harness: Recurrent Neural Network Inference - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

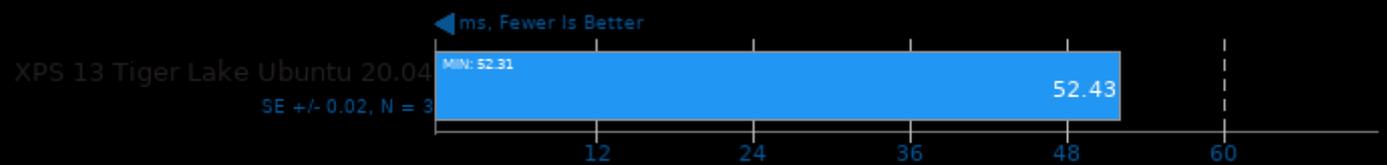
Harness: Recurrent Neural Network Training - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

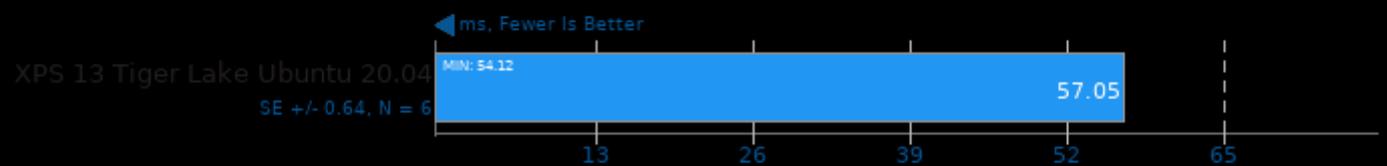
Harness: Convolution Batch Shapes Auto - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

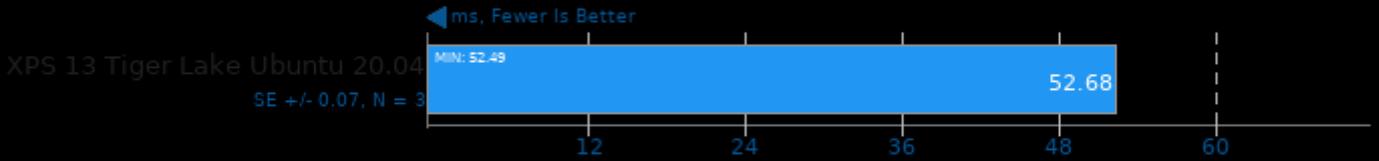
Harness: Deconvolution Batch shapes_1d - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

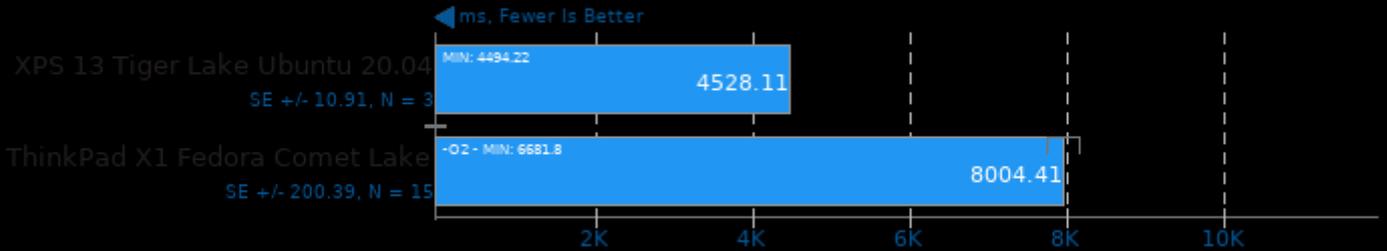
Harness: Deconvolution Batch shapes_3d - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -mssse4.1 -fpic -pie -lpthread

oneDNN 2.0

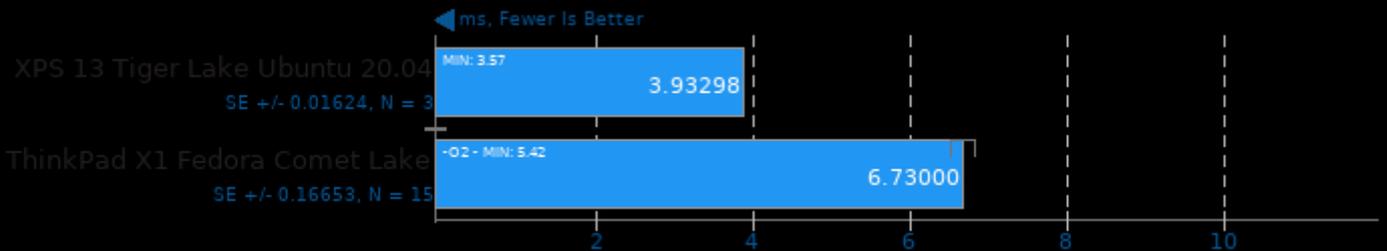
Harness: Recurrent Neural Network Inference - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -mssse4.1 -fpic -pie -lpthread

oneDNN 2.0

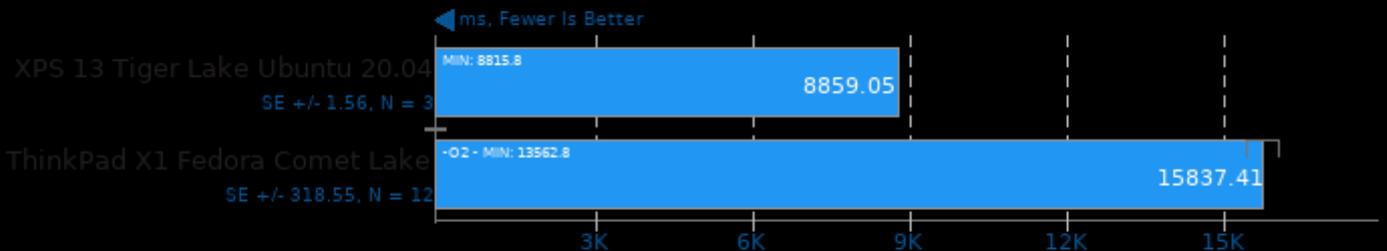
Harness: Matrix Multiply Batch Shapes Transformer - Data Type: f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -mssse4.1 -fpic -pie -lpthread

oneDNN 2.0

Harness: Recurrent Neural Network Training - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -mssse4.1 -fpic -pie -lpthread

oneDNN 2.0

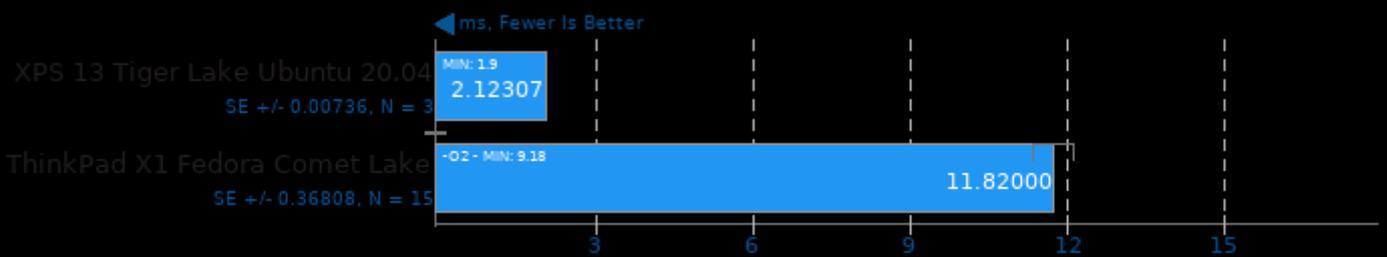
Harness: Recurrent Neural Network Inference - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

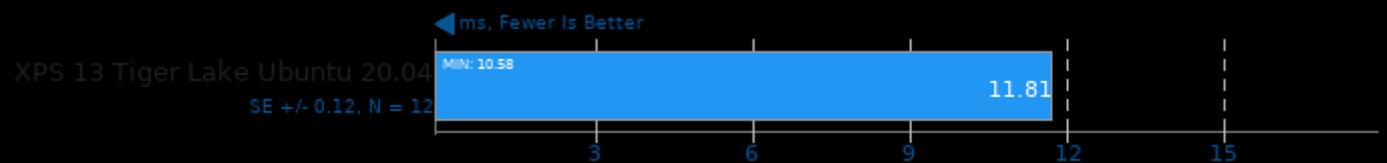
Harness: Matrix Multiply Batch Shapes Transformer - Data Type: u8s8f32 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

oneDNN 2.0

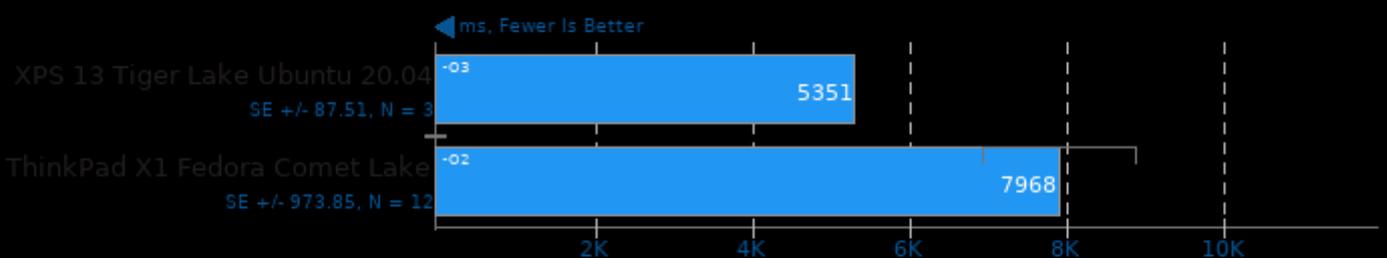
Harness: Matrix Multiply Batch Shapes Transformer - Data Type: bf16bf16bf16 - Engine: CPU



1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fpic -pie -pthread

OpenCV 4.4

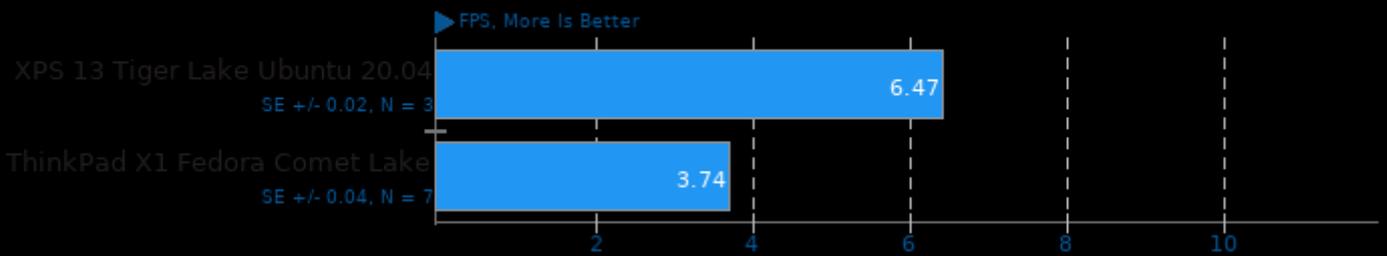
Test: DNN - Deep Neural Network



1. (CXX) g++ options: -fsigned-char -pthread -fomit-frame-pointer -ffunction-sections -fdata-sections -msse -msse2 -msse3 -fvisibility=hidden -ld -lm -lp

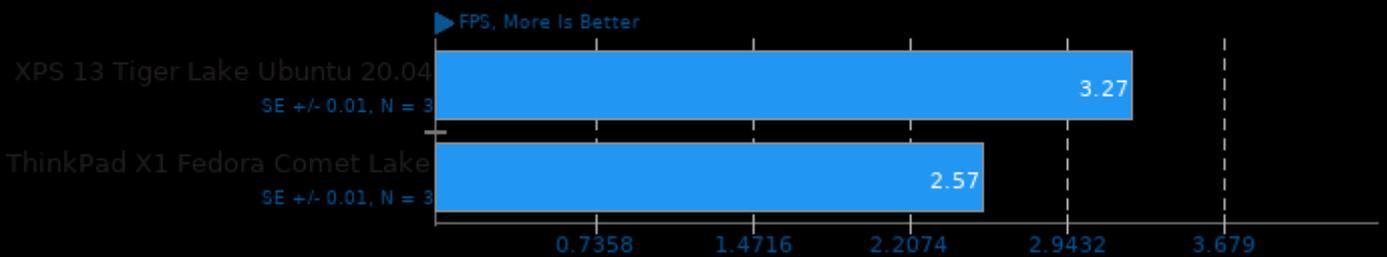
PlaidML

FP16: No - Mode: Inference - Network: VGG16 - Device: CPU



PlaidML

FP16: No - Mode: Inference - Network: ResNet 50 - Device: CPU



RNNoise 2020-06-28

◀ Seconds, Fewer Is Better



1. (CC) gcc options: -O2 -pedantic -fvisibility=hidden

Scikit-Learn 0.22.1

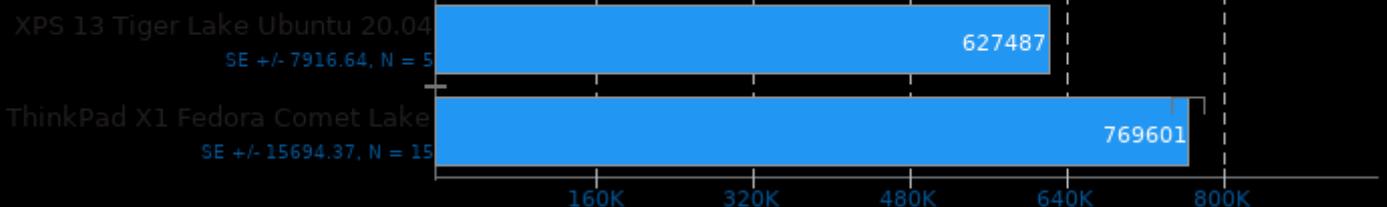
◀ Seconds, Fewer Is Better



TensorFlow Lite 2020-08-23

Model: SqueezeNet

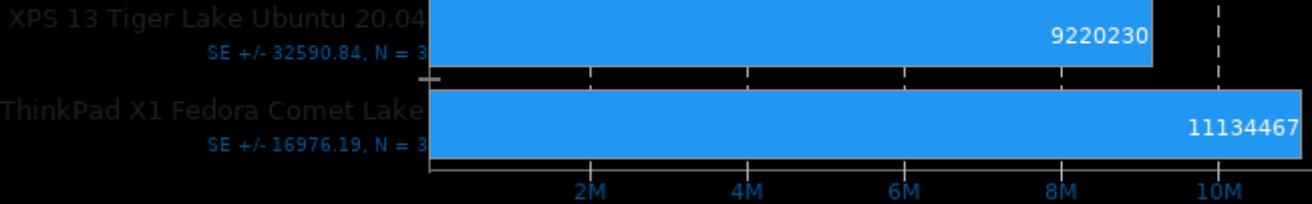
◀ Microseconds, Fewer Is Better



TensorFlow Lite 2020-08-23

Model: Inception V4

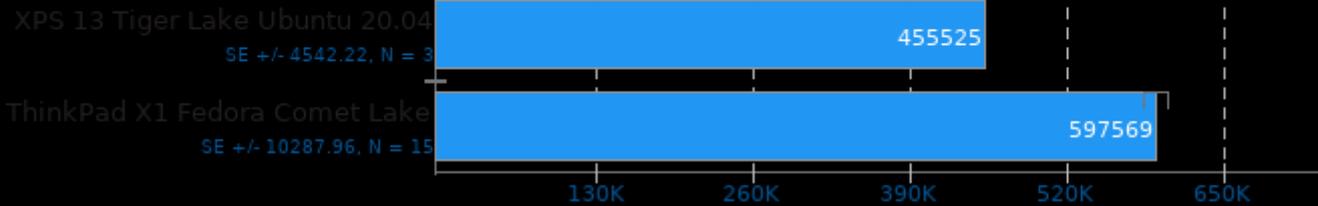
Microseconds, Fewer Is Better



TensorFlow Lite 2020-08-23

Model: NASNet Mobile

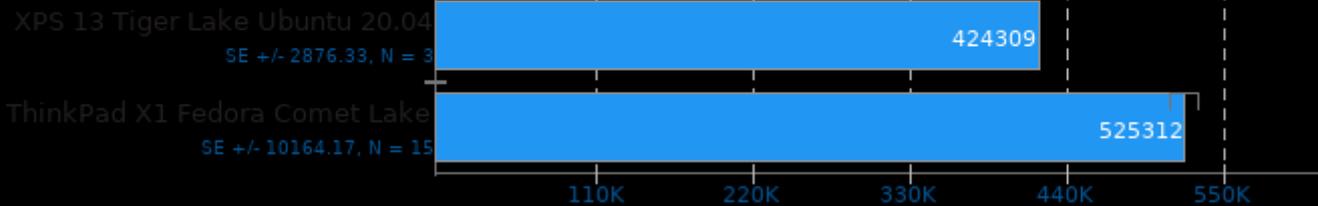
Microseconds, Fewer Is Better



TensorFlow Lite 2020-08-23

Model: Mobilenet Float

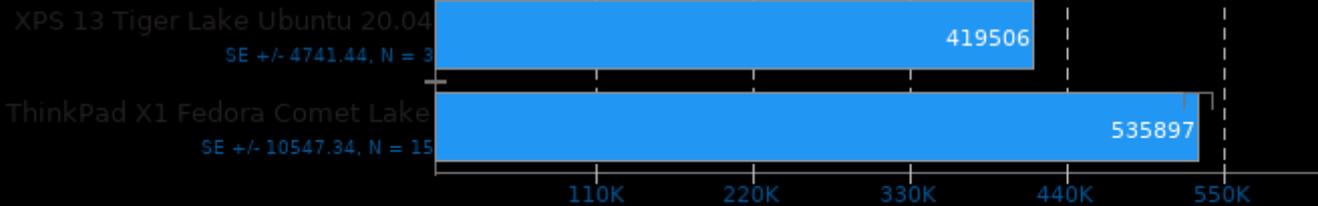
Microseconds, Fewer Is Better

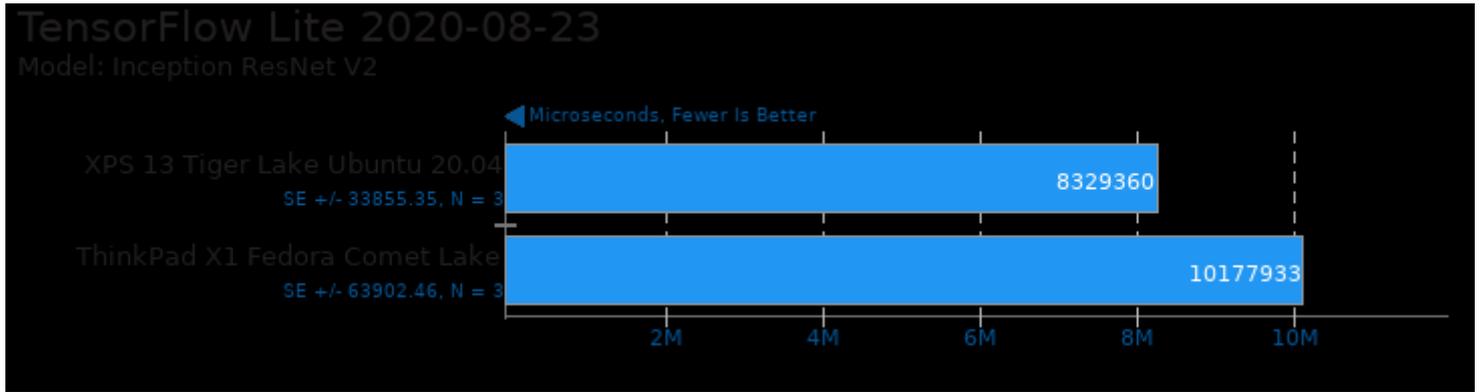


TensorFlow Lite 2020-08-23

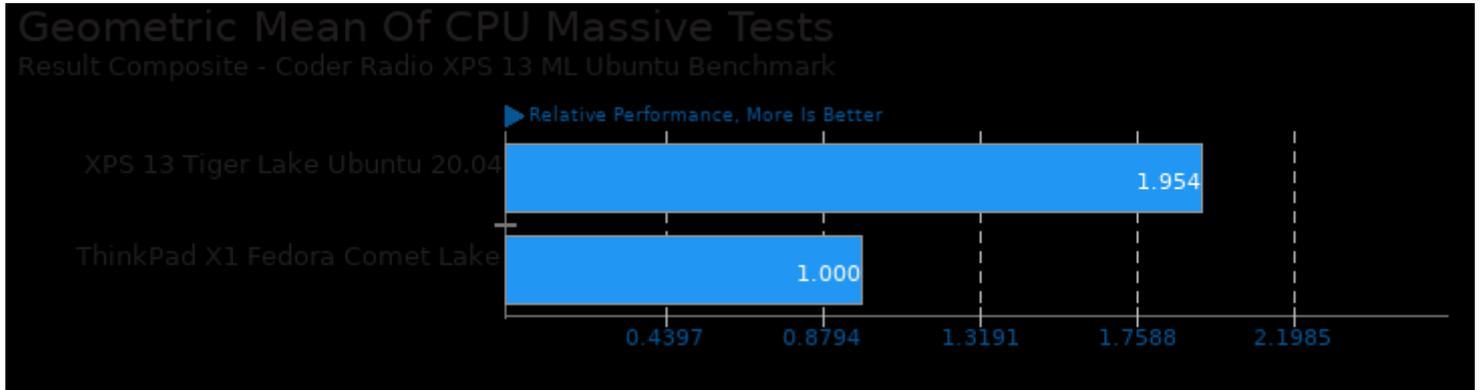
Model: Mobilenet Quant

Microseconds, Fewer Is Better

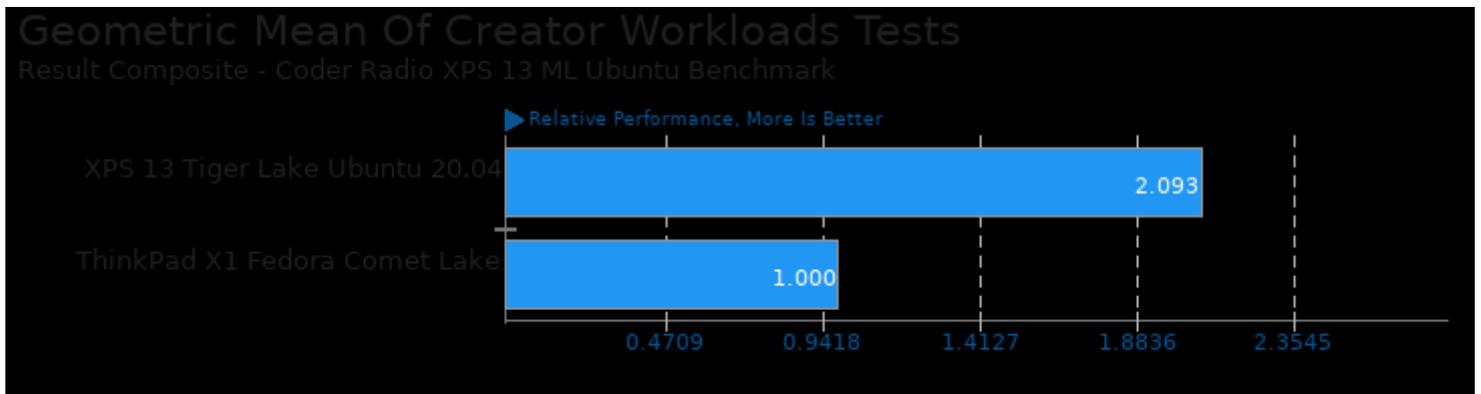




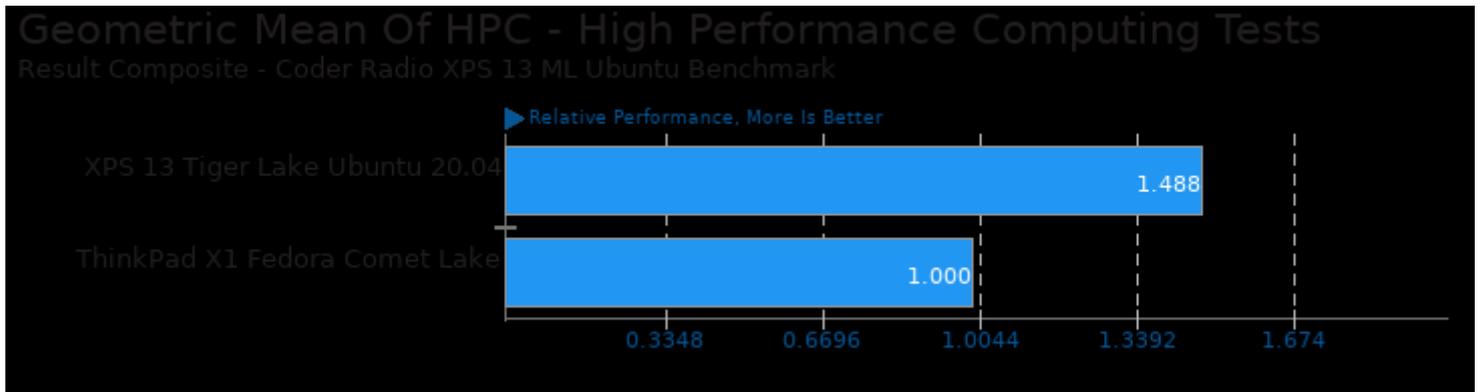
These geometric means are based upon test groupings / test suites for this result file.



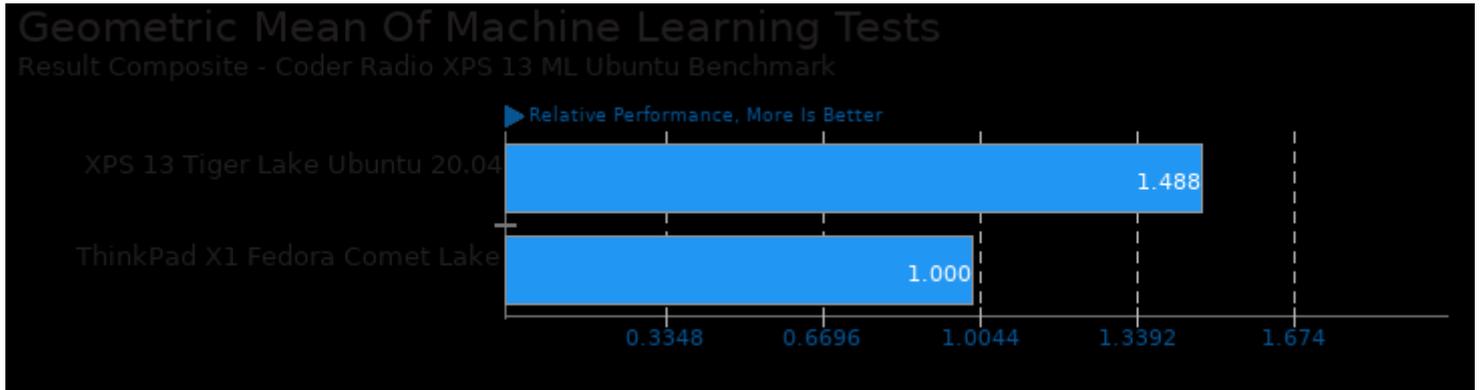
Geometric mean based upon tests: pts/onednn, pts/numenta-nab, pts/numpy, pts/plaidml and pts/scikit-learn



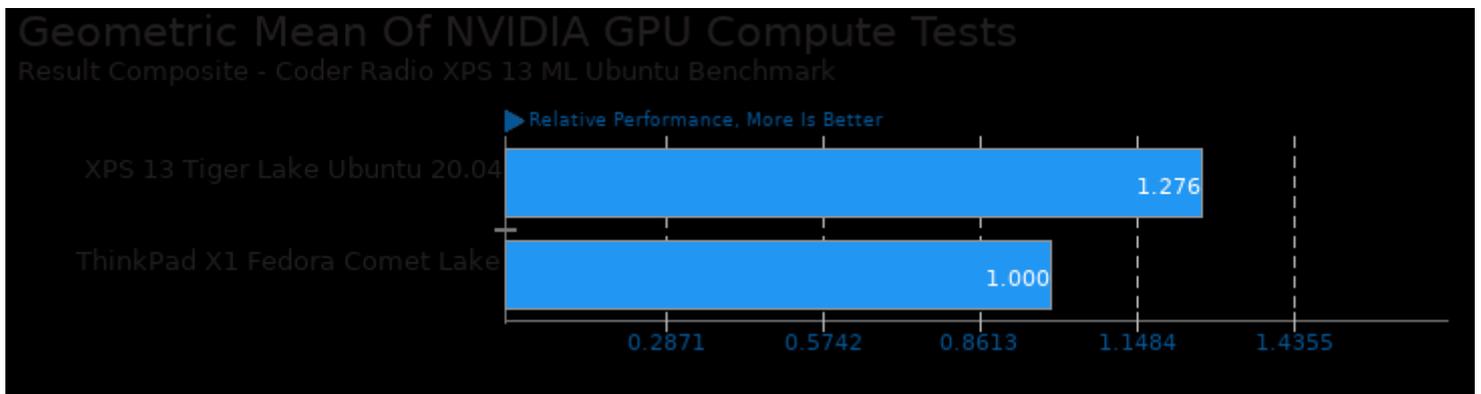
Geometric mean based upon tests: pts/onednn and pts/rnnoise



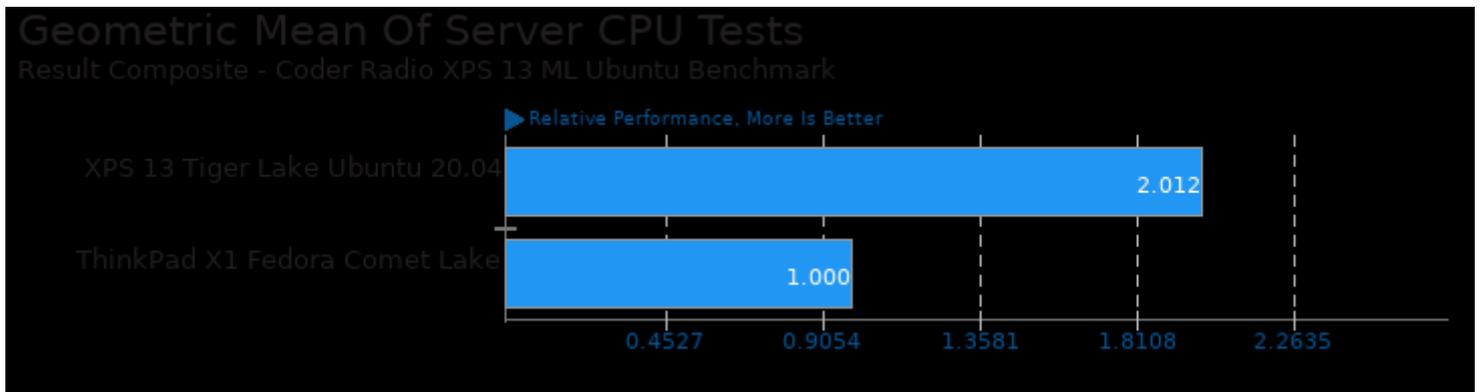
Geometric mean based upon tests: pts/mnn, pts/ncnn, pts/opencv, pts/numpy, pts/ai-benchmark, pts/rnnoise, pts/scikit-learn, pts/mlpack, pts/numenta-nab, pts/tensorflow-lite, pts/onednn and pts/plaidml



Geometric mean based upon tests: pts/mnn, pts/ncnn, pts/opencv, pts/numpy, pts/ai-benchmark, pts/rnnoise, pts/scikit-learn, pts/mlpack, pts/numenta-nab, pts/tensorflow-lite, pts/onednn and pts/plaidml



Geometric mean based upon tests: pts/plaidml and pts/ncnn



Geometric mean based upon tests: pts/onednn, pts/numpy, pts/numenta-nab and pts/scikit-learn

This file was automatically generated via the Phoronix Test Suite benchmarking software on Friday, 29 March 2024 07:17.