



[www.phoronix-test-suite.com](http://www.phoronix-test-suite.com)

## 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-multilib --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 + l1tf: Not affected + mds: Not affected + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swapgs barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of 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 + l1tf: Not affected + mds: Not affected + meltdown: Not affected + spec\_store\_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre\_v1: Mitigation of usercopy/swapgs barriers and \_\_user pointer sanitization + spectre\_v2: Mitigation of Enhanced IBRS IBPB: conditional RSB filling + srbds: Mitigation of TSX disabled + tsx\_async\_abort: Not affected

## Coder Radio XPS 13 ML Ubuntu Benchmark

	XPS 13 Tiger Lake Ubuntu 20.04	ThinkPad X1 Fedora Comet Lake
<b>Numenta Anomaly Benchmark - EXPoSE (sec)</b>	1139	
Standard Deviation	0.4%	
<b>PlaidML - No - Inference - VGG16 - CPU (FPS)</b>	<b>6.47</b>	<b>3.74</b>
Normalized	100%	57.81%
Standard Deviation	0.6%	3%
<b>PlaidML - No - Inference - ResNet 50 - CPU (FPS)</b>	<b>3.27</b>	<b>2.57</b>
Normalized	100%	78.59%
Standard Deviation	0.8%	0.4%
<b>NCNN - CPUv2-yolov3v2-yolov3 - mobilenetv2-yolov3</b>		45.37
(ms)		
Standard Deviation		7.3%
<b>AI Benchmark Alpha - Device AI Score (Score)</b>	1186	
AI Benchmark Alpha - D.T.S (Score)	630	
AI Benchmark Alpha - D.I.S (Score)	556	
oneDNN - R.N.N.I - f32 - CPU (ms)	<b>5736</b>	<b>8215</b>
Normalized	100%	69.83%
Standard Deviation	73.4%	7.4%
Mobile Neural Network - inception-v3 (ms)	<b>68.523</b>	<b>97.71</b>
Normalized	100%	70.13%
Standard Deviation	1.2%	8.8%
Mobile Neural Network - mobilenet-v1-1.0 (ms)	<b>8.320</b>	<b>13.15</b>
Normalized	100%	63.27%
Standard Deviation	6.8%	13.8%
Mobile Neural Network - MobileNetV2_224 (ms)	<b>6.238</b>	<b>9.62</b>
Normalized	100%	64.84%
Standard Deviation	0.7%	21.1%
Mobile Neural Network - resnet-v2-50 (ms)	<b>54.691</b>	<b>88.35</b>
Normalized	100%	61.9%
Standard Deviation	1%	3.3%
Mobile Neural Network - SqueezeNetV1.0 (ms)	<b>11.256</b>	<b>18.04</b>
Normalized	100%	62.39%
Standard Deviation	7%	0.9%
oneDNN - R.N.N.T - bf16bf16bf16 - CPU (ms)	<b>8859</b>	<b>15837</b>
Normalized	100%	55.94%
Standard Deviation	0%	7%
oneDNN - R.N.N.T - u8s8f32 - CPU (ms)	<b>8875</b>	<b>15829</b>
Normalized	100%	56.07%
Standard Deviation	0.4%	6.4%
oneDNN - R.N.N.T - f32 - CPU (ms)	<b>8863</b>	<b>15730</b>
Normalized	100%	56.34%
Standard Deviation	0.1%	6.3%
<b>Numenta Anomaly Benchmark - Earthgecko Skyline</b>	333.639	
Standard Deviation	0.4%	
<b>NCNN - CPU - regnety_400m (ms)</b>	<b>21.11</b>	<b>22.34</b>
Normalized	100%	94.49%
Standard Deviation	1.4%	9.7%
<b>NCNN - CPU - squeezenet_ss (ms)</b>	<b>39.61</b>	<b>45.61</b>
Normalized	100%	86.84%
Standard Deviation	1%	7.6%
<b>NCNN - CPU - yolov4-tiny (ms)</b>	<b>44.65</b>	<b>57.10</b>
Normalized	100%	78.2%
Standard Deviation	2.1%	7.1%

## Coder Radio XPS 13 ML Ubuntu Benchmark

<b>NCNN - CPU - resnet50 (ms)</b>	<b>51.01</b>	<b>67.20</b>
Normalized	100%	75.91%
Standard Deviation	0.4%	11%
<b>NCNN - CPU - alexnet (ms)</b>	<b>19.06</b>	<b>26.55</b>
Normalized	100%	71.79%
Standard Deviation	0.1%	5.7%
<b>NCNN - CPU - resnet18 (ms)</b>	<b>22.07</b>	<b>30.23</b>
Normalized	100%	73.01%
Standard Deviation	0.2%	9.1%
<b>NCNN - CPU - vgg16 (ms)</b>	<b>68.50</b>	<b>108.71</b>
Normalized	100%	63.01%
Standard Deviation	0.6%	4.3%
<b>NCNN - CPU - googlenet (ms)</b>	<b>25.01</b>	<b>31.39</b>
Normalized	100%	79.68%
Standard Deviation	0.8%	9.8%
<b>NCNN - CPU - blazeface (ms)</b>	<b>2.85</b>	<b>3.80</b>
Normalized	100%	75%
Standard Deviation	1.4%	10.3%
<b>NCNN - CPU - efficientnet-b0 (ms)</b>	<b>12.53</b>	<b>15.22</b>
Normalized	100%	82.33%
Standard Deviation	0.5%	11.7%
<b>NCNN - CPU - mnasnet (ms)</b>	<b>8.10</b>	<b>9.76</b>
Normalized	100%	82.99%
Standard Deviation	10.8%	13.9%
<b>NCNN - CPU - shufflenet-v2 (ms)</b>	<b>9.99</b>	<b>14.58</b>
Normalized	100%	68.52%
Standard Deviation	10.2%	15.8%
<b>NCNN - CPU-v3-v3 - mobilenet-v3 (ms)</b>	<b>6.71</b>	<b>9.10</b>
Normalized	100%	73.74%
Standard Deviation	0.5%	13%
<b>NCNN - CPU-v2-v2 - mobilenet-v2 (ms)</b>	<b>7.78</b>	<b>10.67</b>
Normalized	100%	72.91%
Standard Deviation	1.2%	18.8%
<b>NCNN - CPU - mobilenet (ms)</b>	<b>35.03</b>	<b>45.37</b>
Normalized	100%	77.21%
Standard Deviation	2.8%	7.3%
<b>oneDNN - R.N.N.I - u8s8f32 - CPU (ms)</b>	<b>4528</b>	<b>8004</b>
Normalized	100%	56.57%
Standard Deviation	0.4%	9.7%
<b>oneDNN - R.N.N.I - bf16bf16bf16 - CPU (ms)</b>	<b>4517</b>	<b>7969</b>
Normalized	100%	56.68%
Standard Deviation	0.2%	9.8%
<b>Numpy Benchmark (Score)</b>	<b>293.18</b>	<b>299.99</b>
Normalized	97.73%	100%
Standard Deviation	0.3%	1.8%
<b>Milpack Benchmark - scikit_qda (sec)</b>	138.24	
Standard Deviation	0.1%	
<b>TensorFlow Lite - SqueezeNet (us)</b>	<b>627487</b>	<b>769601</b>
Normalized	100%	81.53%
Standard Deviation	2.8%	7.9%
<b>TensorFlow Lite - NASNet Mobile (us)</b>	<b>455525</b>	<b>597569</b>
Normalized	100%	76.23%
Standard Deviation	1.7%	6.7%
<b>TensorFlow Lite - Mobilenet Quant (us)</b>	<b>419506</b>	<b>535897</b>
Normalized	100%	78.28%

TensorFlow Lite - Mobilenet Float (us)	424309	525312
Normalized	100%	80.77%
Standard Deviation	2%	7.6%
TensorFlow Lite - Inception V4 (us)	9220230	11134467
Normalized	100%	82.81%
Standard Deviation	1.2%	7.5%
TensorFlow Lite - I.R.V (us)	8329360	10177933
Normalized	100%	81.84%
Standard Deviation	0.6%	0.3%
NCNN - V.G.y.y - mobilenetv2-yolov3 (ms)		46.41
Standard Deviation	0.7%	1.1%
Mipack Benchmark - scikit_ica (sec)	123.23	
Standard Deviation	2.7%	0.5%
NCNN - Vulkan GPU - regnety_400m (ms)	21.01	20.14
Normalized	95.86%	100%
Standard Deviation	7.4%	18.7%
NCNN - Vulkan GPU - squeezenet_ssd (ms)	39.53	44.71
Normalized	100%	88.41%
Standard Deviation	1%	9.6%
NCNN - Vulkan GPU - yolov4-tiny (ms)	44.68	58.16
Normalized	100%	76.82%
Standard Deviation	2.2%	4.7%
NCNN - Vulkan GPU - resnet50 (ms)	51.13	65.59
Normalized	100%	77.95%
Standard Deviation	0.4%	15.4%
NCNN - Vulkan GPU - alexnet (ms)	19.23	25.75
Normalized	100%	74.68%
Standard Deviation	1.8%	6.2%
NCNN - Vulkan GPU - resnet18 (ms)	22.25	28.33
Normalized	100%	78.54%
Standard Deviation	1.3%	10.9%
NCNN - Vulkan GPU - vgg16 (ms)	69.16	107.70
Normalized	100%	64.22%
Standard Deviation	1.3%	2.8%
NCNN - Vulkan GPU - googlenet (ms)	24.58	28.85
Normalized	100%	85.2%
Standard Deviation	4.2%	15.1%
NCNN - Vulkan GPU - blazeface (ms)	2.80	3.24
Normalized	100%	86.42%
Standard Deviation	3.8%	27.7%
NCNN - Vulkan GPU - efficientnet-b0 (ms)	11.83	13.74
Normalized	100%	86.1%
Standard Deviation	9.7%	19.8%
NCNN - Vulkan GPU - mnasnet (ms)	8.13	8.77
Normalized	100%	92.7%
Standard Deviation	11.4%	22.8%
NCNN - Vulkan GPU - shufflenet-v2 (ms)	10.30	12.80
Normalized	100%	80.47%
Standard Deviation	11.2%	28.3%
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-v2-v2 - mobilenet-v2 (ms)	7.71	9.64
Normalized	100%	79.98%

	Standard Deviation	0.6%	24.2%
NCNN - Vulkan GPU - mobilenet (ms)	35.11	46.41	
	Normalized	100%	75.65%
	Standard Deviation	2.4%	0.5%
<b>Numenta Anomaly Benchmark - B.C (sec)</b>	81.062		
	Standard Deviation	2.6%	
oneDNN - D.B.s - f32 - CPU (ms)	14.6326	20.71	
	Normalized	100%	70.65%
	Standard Deviation	2.9%	13.8%
<b>Numenta Anomaly Benchmark - Relative Entropy</b>	48.609		
	Standard Deviation	3%	
<b>Mipack Benchmark - scikit_linearridge_regression</b>	13.50		
	Standard Deviation	0.3%	
oneDNN - D.B.s - u8s8f32 - CPU (ms)	2.94262	21.88	
	Normalized	100%	13.45%
	Standard Deviation	3.4%	0.7%
oneDNN - M.M.B.S.T - bf16bf16bf16 - CPU (ms)	11.8072		
	Standard Deviation	3.6%	
oneDNN - IP Shapes 1D - f32 - CPU (ms)	10.2773	15.77	
	Normalized	100%	65.17%
	Standard Deviation	1%	15.2%
oneDNN - D.B.s - bf16bf16bf16 - CPU (ms)	57.0546		
	Standard Deviation	2.8%	
oneDNN - IP Shapes 1D - u8s8f32 - CPU (ms)	2.43051	5.88	
	Normalized	100%	41.34%
	Standard Deviation	3.6%	0.8%
<b>Mipack Benchmark - scikit_svm (sec)</b>	34.66		
	Standard Deviation	0.2%	
oneDNN - M.M.B.S.T - f32 - CPU (ms)	3.93298	6.73	
	Normalized	100%	58.44%
	Standard Deviation	0.7%	9.6%
oneDNN - M.M.B.S.T - u8s8f32 - CPU (ms)	2.12307	11.82	
	Normalized	100%	17.96%
	Standard Deviation	0.6%	12.1%
<b>RNNNoise (sec)</b>	31.947		
	Standard Deviation	0%	
<b>Numenta Anomaly Benchmark - Windowed Gaussian (sec)</b>	27.702		
	Standard Deviation	2.6%	
oneDNN - C.B.S.A - f32 - CPU (ms)	11.5632	22.02	
	Normalized	100%	52.51%
	Standard Deviation	4.3%	0.1%
OpenCV - DNN - D.N.N (ms)	5351	7968	
	Normalized	100%	67.16%
	Standard Deviation	2.8%	42.3%
<b>Scikit-Learn (sec)</b>	17.900		
	Standard Deviation	1.4%	
oneDNN - IP Shapes 1D - bf16bf16bf16 - CPU (ms)	25.8282		
	Standard Deviation	0.7%	
oneDNN - IP Shapes 3D - u8s8f32 - CPU (ms)	2.62981	3.95	
	Normalized	100%	66.58%
	Standard Deviation	0.1%	2.5%
oneDNN - C.B.S.A - bf16bf16bf16 - CPU (ms)	52.4284		
	Standard Deviation	0.1%	

**oneDNN - IP Shapes 3D - bf16bf16bf16 - CPU (ms)** 6.39979

Standard Deviation 1.7%

**oneDNN - IP Shapes 3D - f32 - CPU (ms)** **6.59350****17.95**

Normalized 100%

36.73%

Standard Deviation 0.4%

0.1%

**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 - bf16bf16bf16 - CPU (ms)** 52.6824

Standard Deviation 0.2%

**oneDNN - D.B.s - u8s8f32 - CPU (ms)** **3.13684****8.18**

Normalized 100%

38.35%

Standard Deviation 0.2%

1.9%

**oneDNN - D.B.s - f32 - CPU (ms)** **13.4125****15.85**

Normalized 100%

84.62%

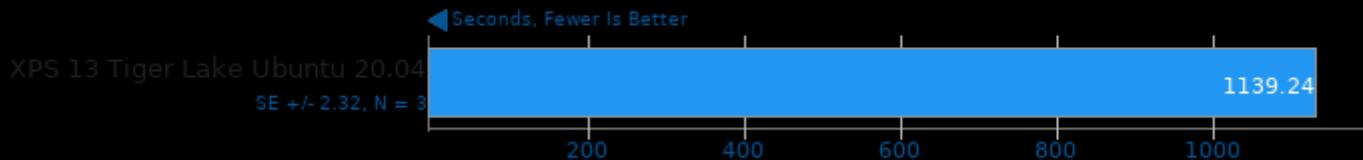
Standard Deviation 0.7%

2.9%

## Coder Radio XPS 13 ML Ubuntu Benchmark

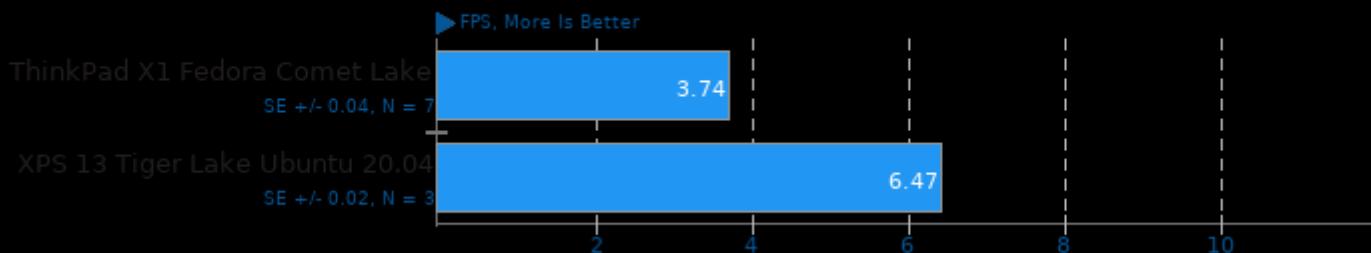
### Numenta Anomaly Benchmark 1.1

Detector: EXPoSE



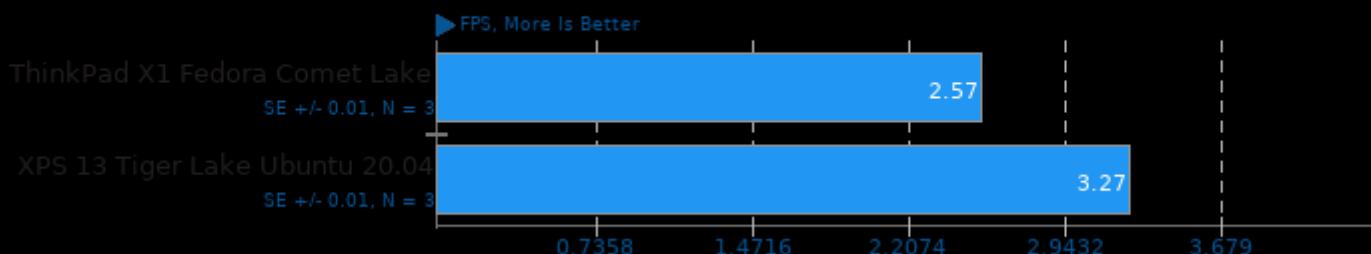
### PlaidML

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



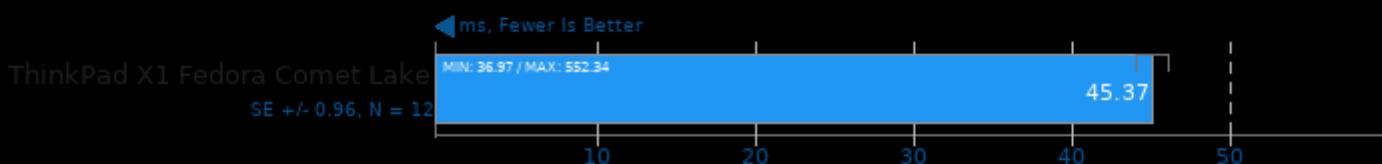
### PlaidML

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



### NCNN 20201218

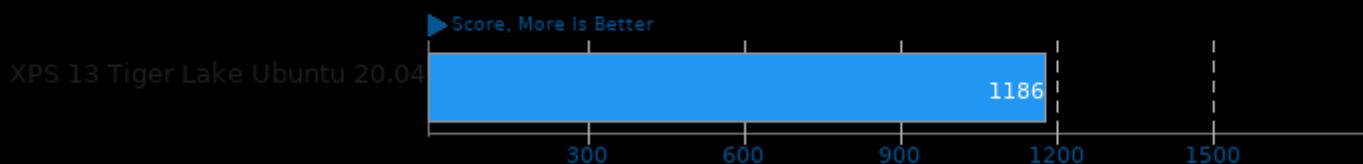
Target: CPUv2-yolov3v2-yolov3 - Model: mobilenetv2-yolov3



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

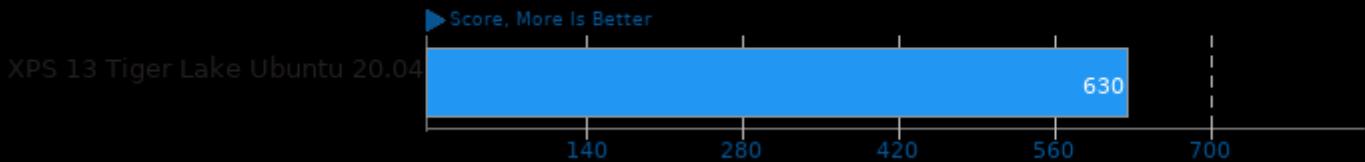
### AI Benchmark Alpha 0.1.2

Device AI Score



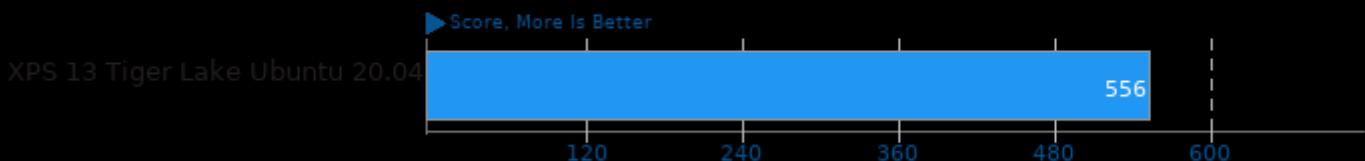
### AI Benchmark Alpha 0.1.2

Device Training Score



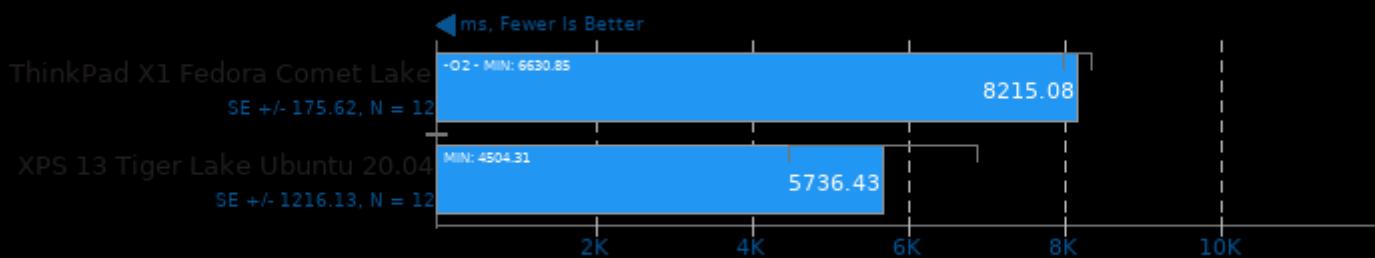
### AI Benchmark Alpha 0.1.2

Device Inference Score



### oneDNN 2.0

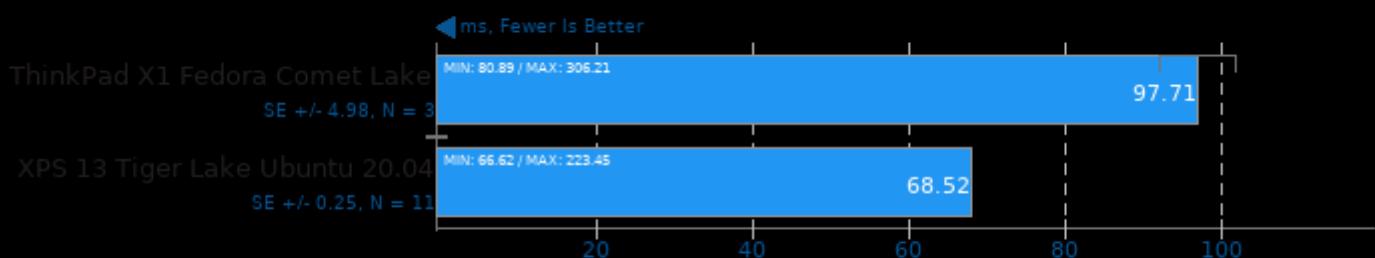
Harness: Recurrent Neural Network Inference - Data Type: f32 - Engine: CPU



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

### Mobile Neural Network 2020-09-17

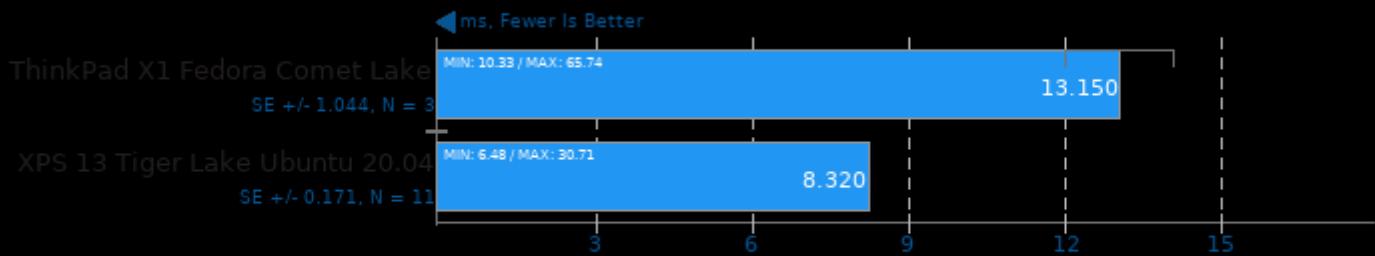
Model: inception-v3



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

## Mobile Neural Network 2020-09-17

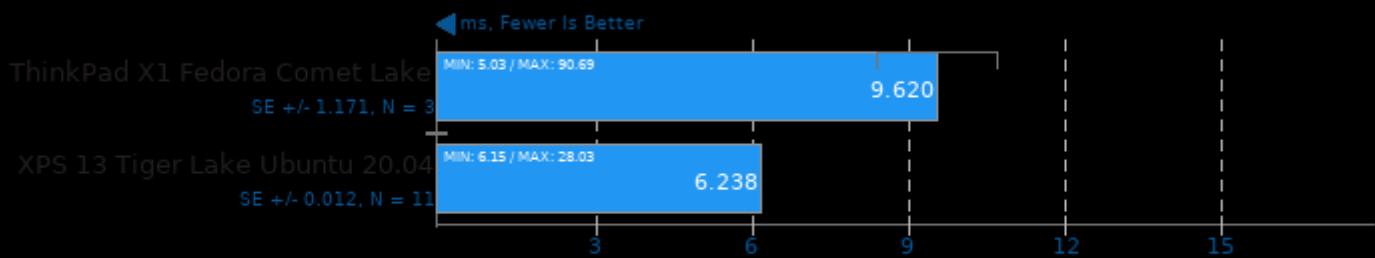
Model: mobilenet-v1-1.0



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

## Mobile Neural Network 2020-09-17

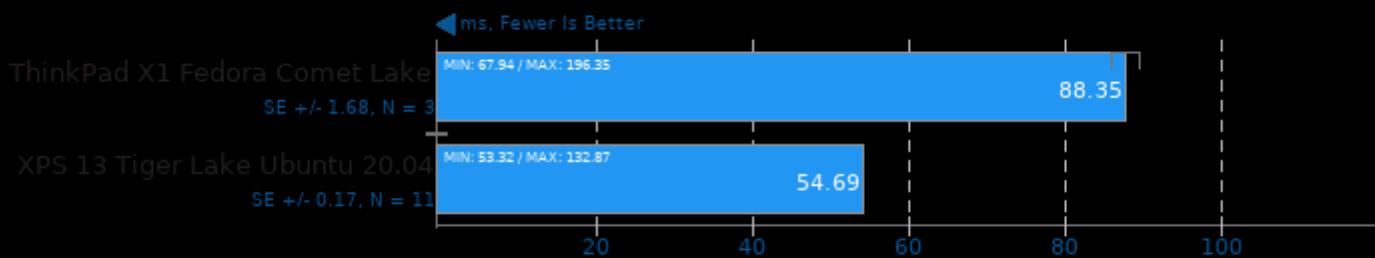
Model: MobileNetV2\_224



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

## Mobile Neural Network 2020-09-17

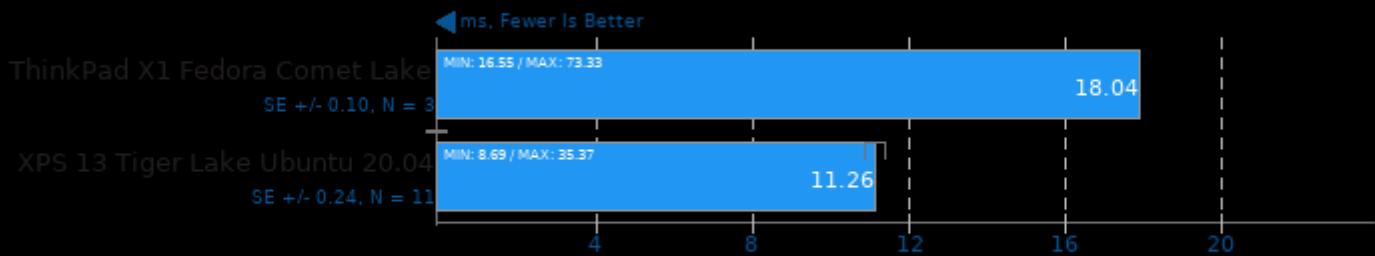
Model: resnet-v2-50



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

## Mobile Neural Network 2020-09-17

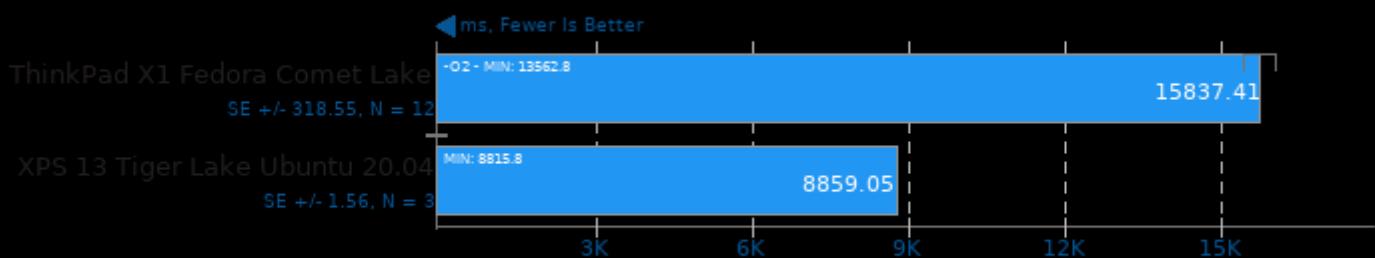
Model: SqueezeNetV1.0



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

## oneDNN 2.0

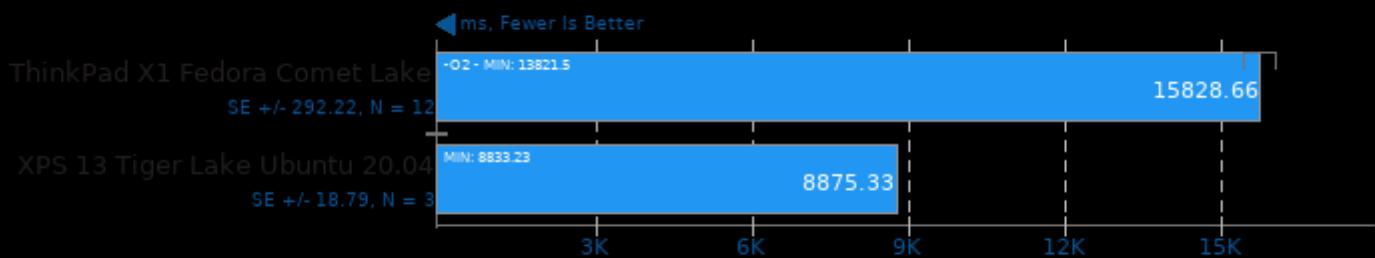
Harness: Recurrent Neural Network Training - Data Type: bf16bf16bf16 - Engine: CPU



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

## oneDNN 2.0

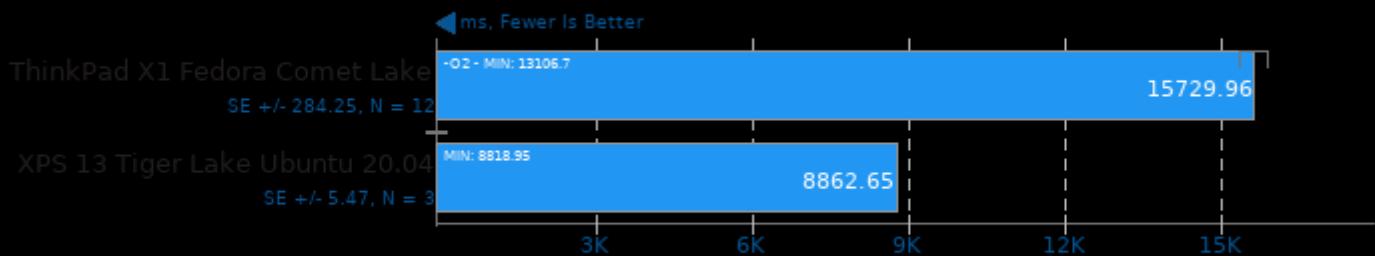
Harness: Recurrent Neural Network Training - Data Type: u8s8f32 - Engine: CPU



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

## oneDNN 2.0

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

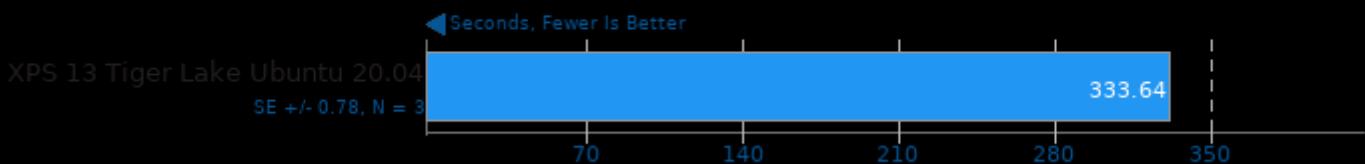


1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fPIC -pie -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

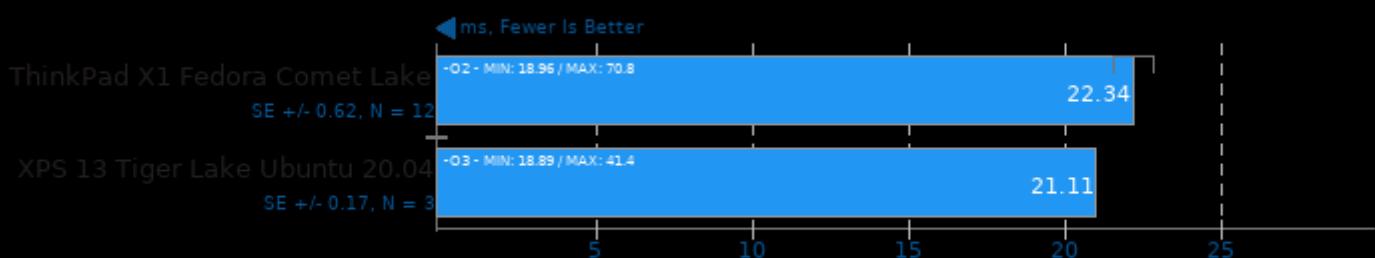
### Numenta Anomaly Benchmark 1.1

Detector: Earthgecko Skyline



### NCNN 20201218

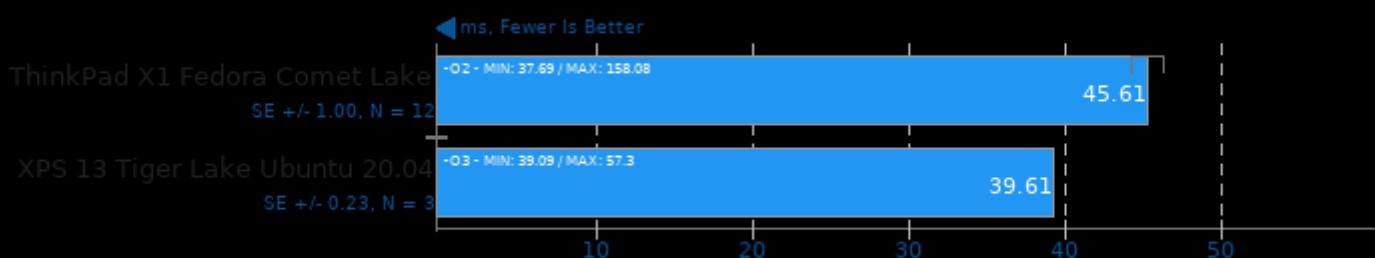
Target: CPU - Model: regnety\_400m



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: yolov4-tiny

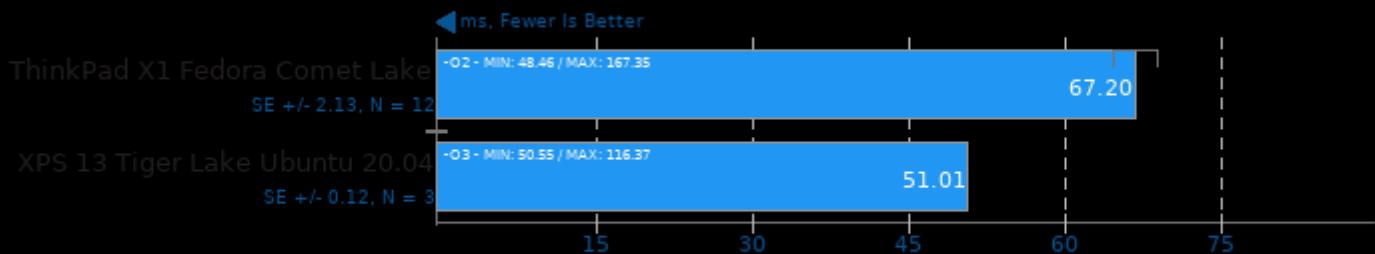


1. (CXX) g++ options: -rdynamic -lgomp -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

### NCNN 20201218

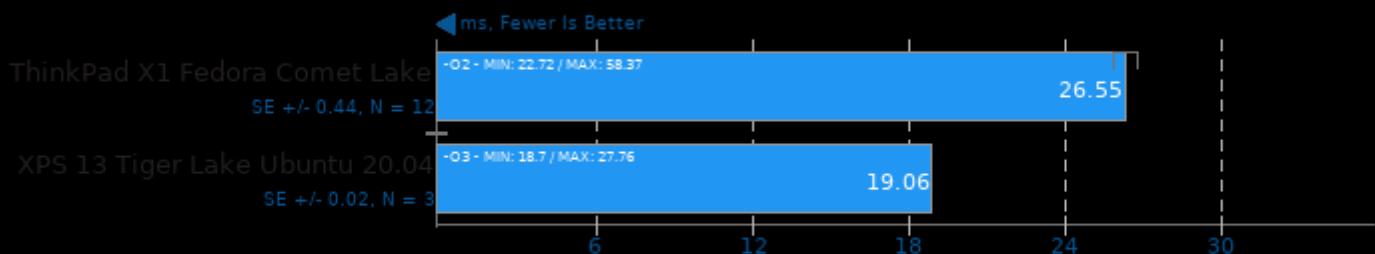
Target: CPU - Model: resnet50



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

### NCNN 20201218

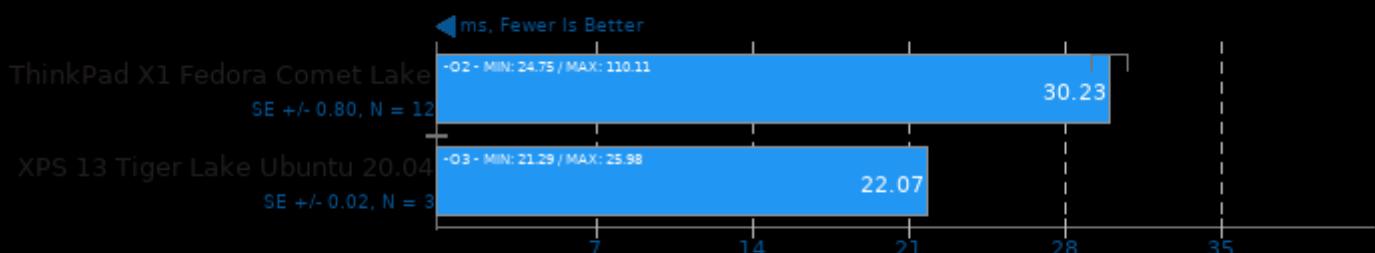
Target: CPU - Model: alexnet



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

### NCNN 20201218

Target: CPU - Model: resnet18



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

### NCNN 20201218

Target: CPU - Model: vgg16



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

## Coder Radio XPS 13 ML Ubuntu Benchmark

### NCNN 20201218

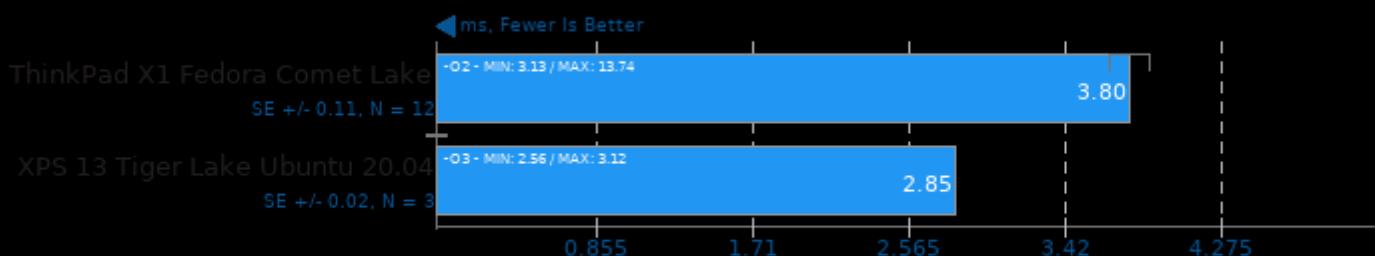
Target: CPU - Model: googlenet



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

### NCNN 20201218

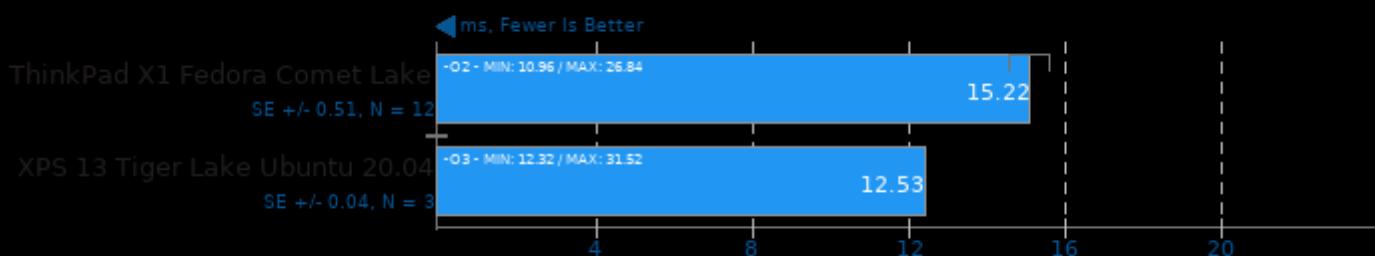
Target: CPU - Model: blazeface



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

### NCNN 20201218

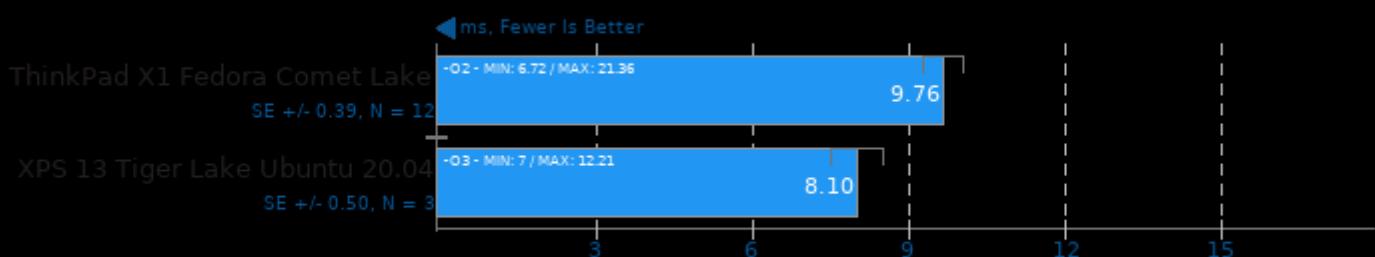
Target: CPU - Model: efficientnet-b0



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

### NCNN 20201218

Target: CPU - Model: mnasnet

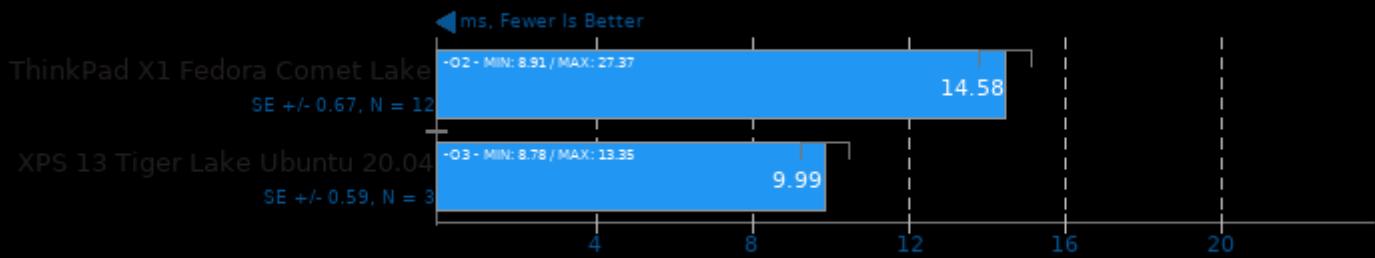


1. (CXX) g++ options: -rdynamic -lgomp -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

### NCNN 20201218

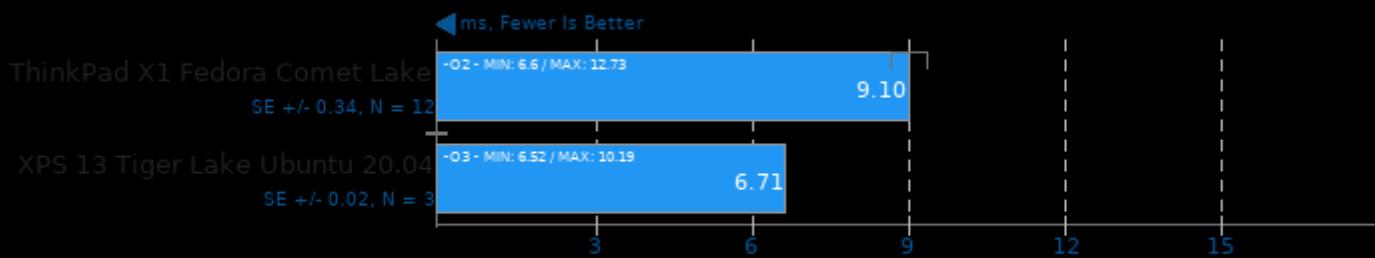
Target: CPU - Model: shufflenet-v2



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

### NCNN 20201218

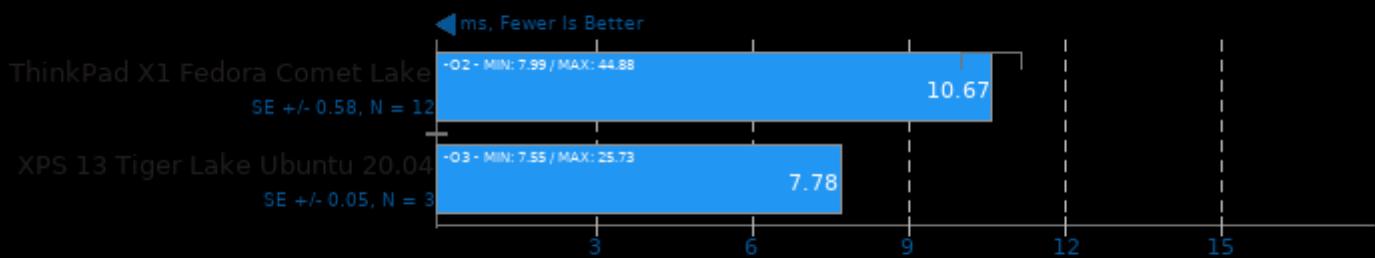
Target: CPU-v3-v3 - Model: mobilenet-v3



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

### NCNN 20201218

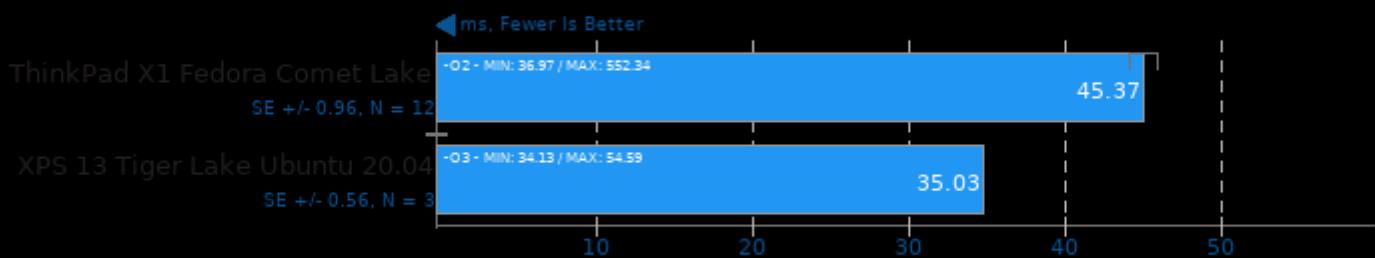
Target: CPU-v2-v2 - Model: mobilenet-v2



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

### NCNN 20201218

Target: CPU - Model: mobilenet

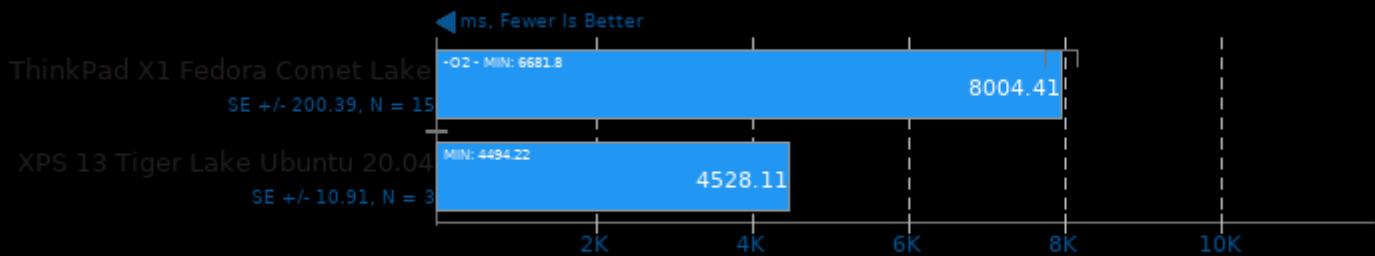


1. (CXX) g++ options: -rdynamic -lgomp -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

### oneDNN 2.0

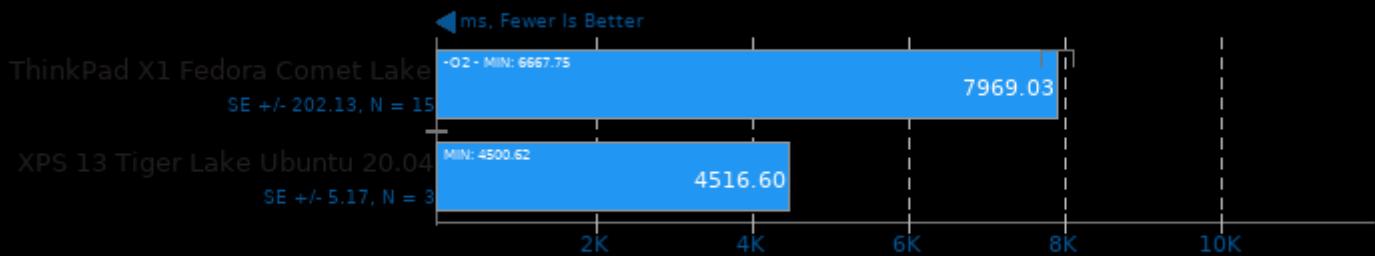
Harness: Recurrent Neural Network Inference - Data Type: u8s8f32 - Engine: CPU



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

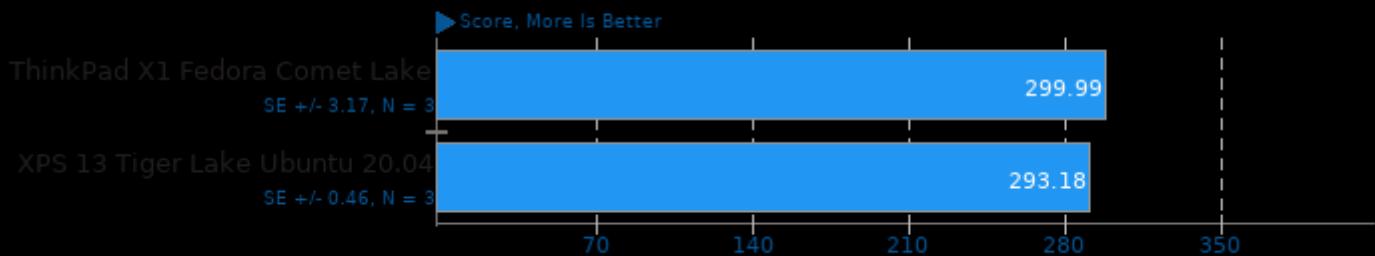
### oneDNN 2.0

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



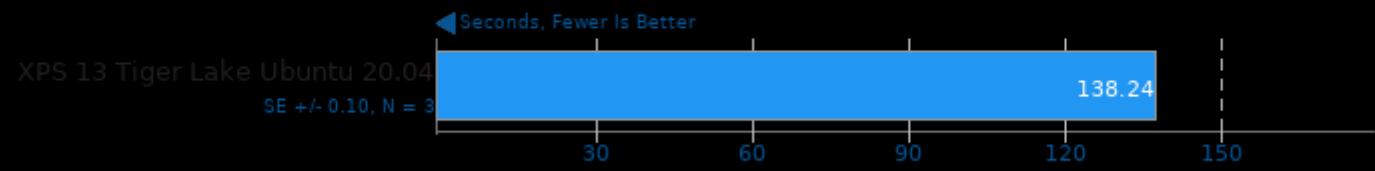
1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fPIC -pie -lpthread

### Numpy Benchmark



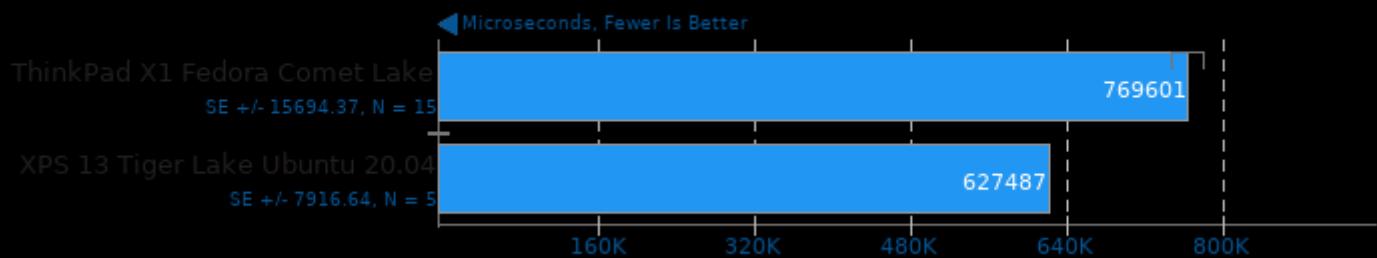
### Milpack Benchmark

Benchmark: scikit\_qda



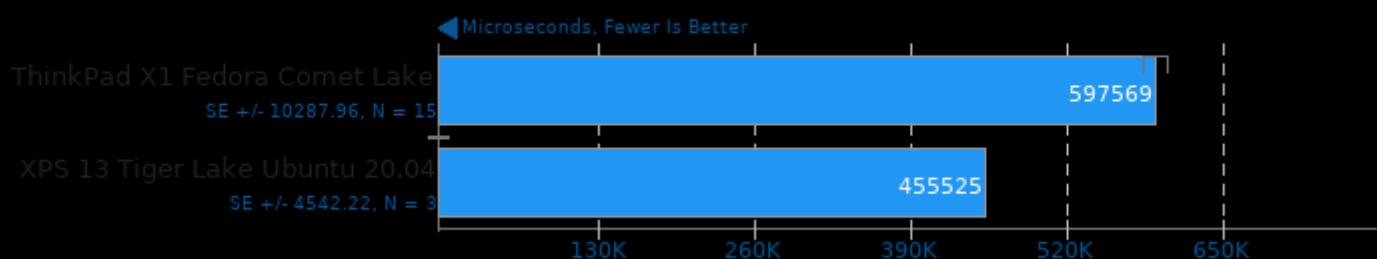
### TensorFlow Lite 2020-08-23

Model: SqueezeNet



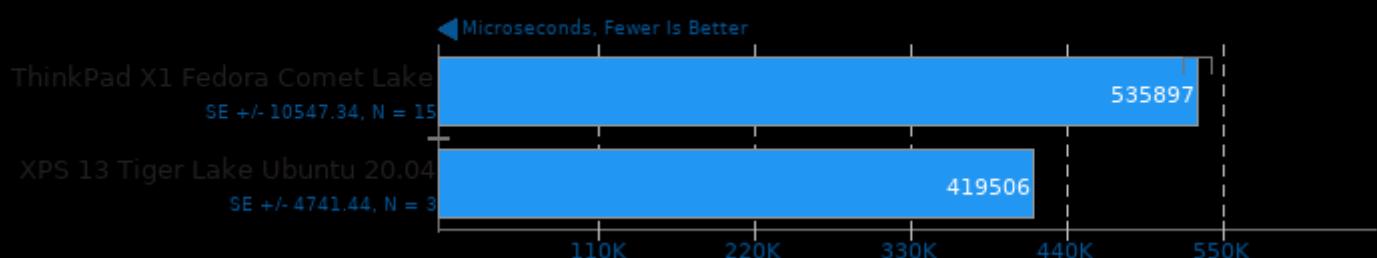
### TensorFlow Lite 2020-08-23

Model: NASNet Mobile



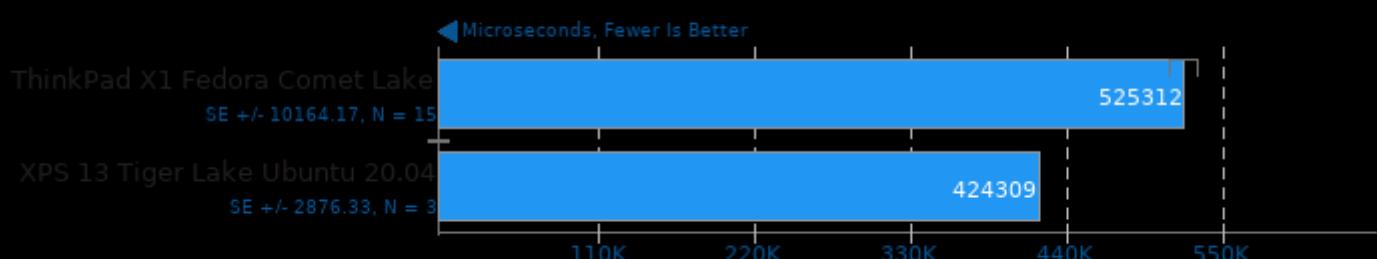
### TensorFlow Lite 2020-08-23

Model: Mobilenet Quant



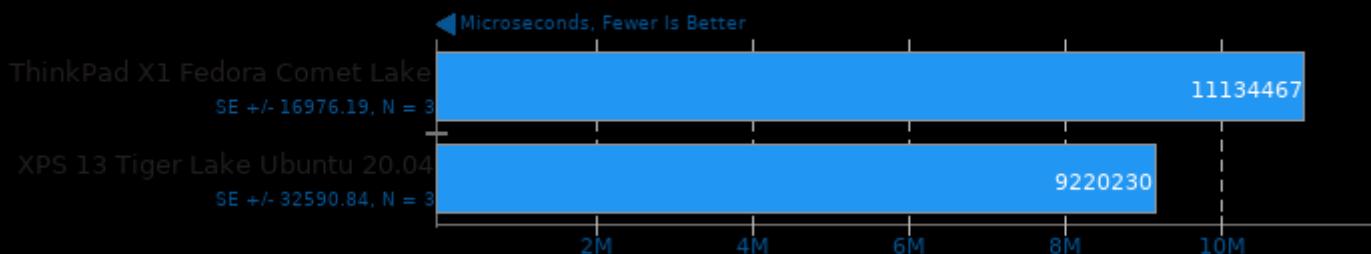
### TensorFlow Lite 2020-08-23

Model: Mobilenet Float



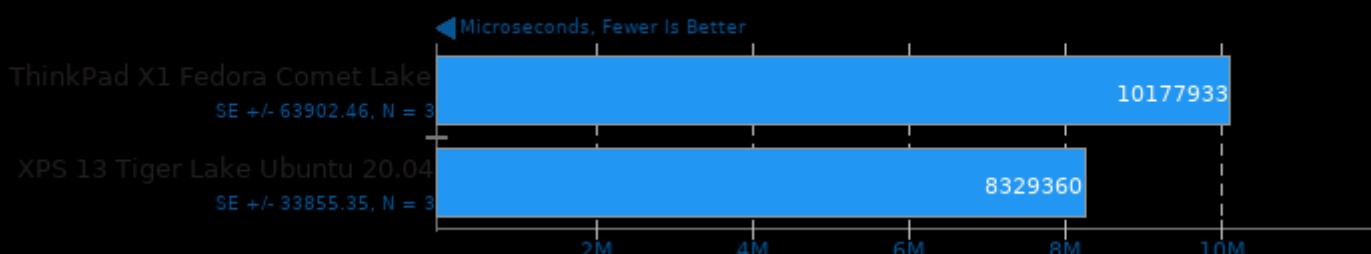
### TensorFlow Lite 2020-08-23

Model: Inception V4



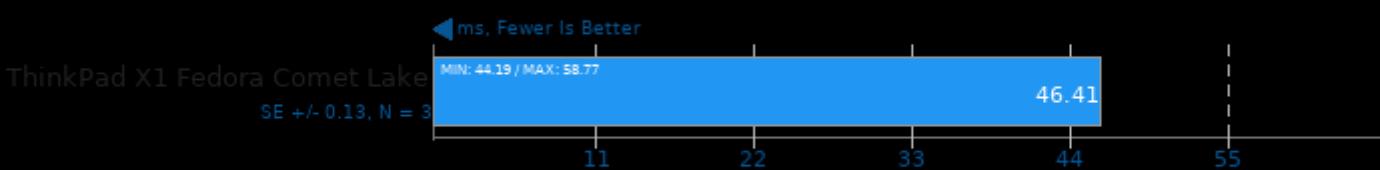
### TensorFlow Lite 2020-08-23

Model: Inception ResNet V2



### NCNN 20201218

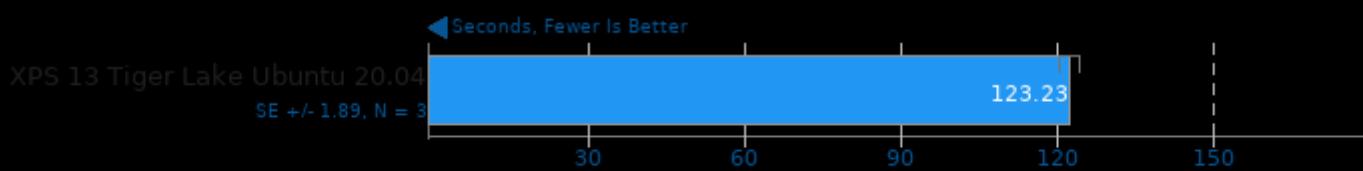
Target: Vulkan GPUv2-yolov3v2-yolov3 - Model: mobilenetv2-yolov3



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

### Milpck Benchmark

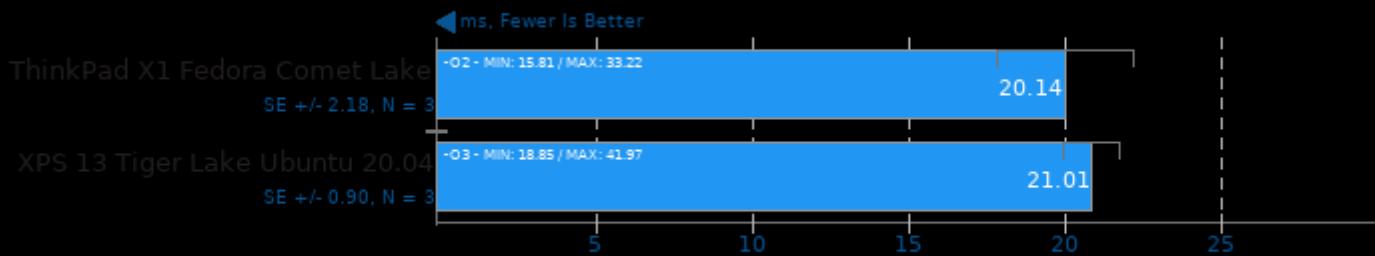
Benchmark: scikit\_ica



## Coder Radio XPS 13 ML Ubuntu Benchmark

### NCNN 20201218

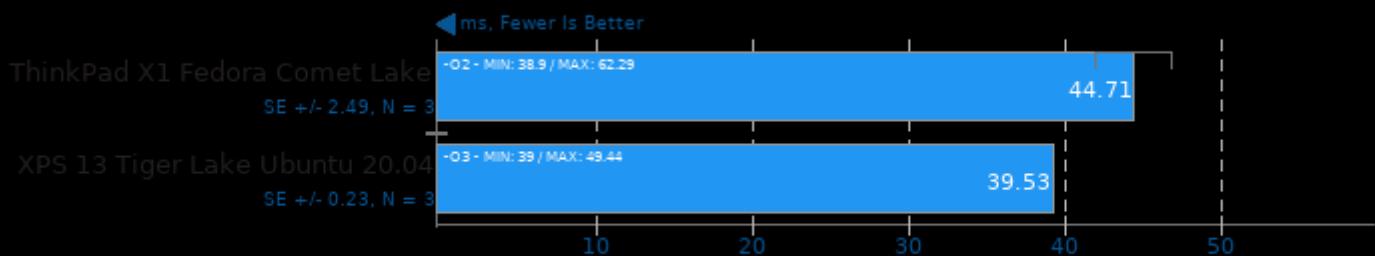
Target: Vulkan GPU - Model: regnety\_400m



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

### NCNN 20201218

Target: Vulkan GPU - Model: squeezenet\_ssdl



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

### NCNN 20201218

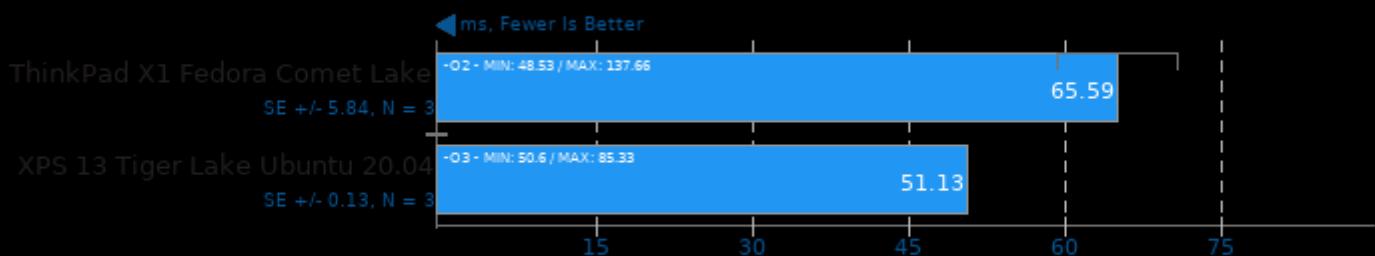
Target: Vulkan GPU - Model: yolov4-tiny



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

### NCNN 20201218

Target: Vulkan GPU - Model: resnet50

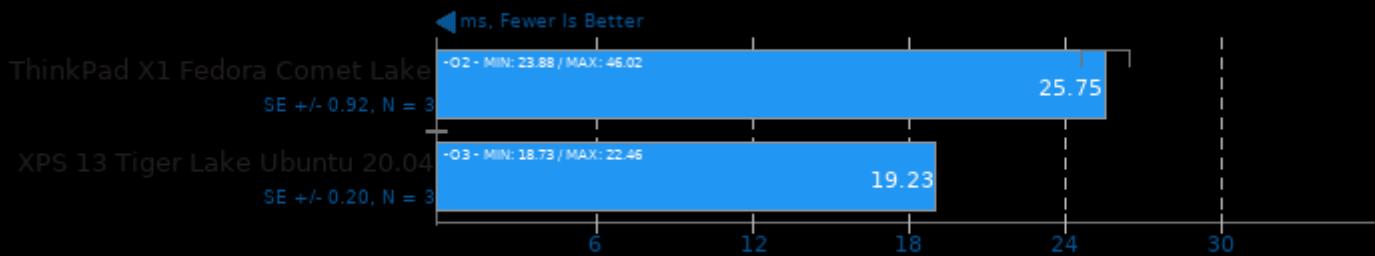


1. (CXX) g++ options: -rdynamic -lgomp -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

### NCNN 20201218

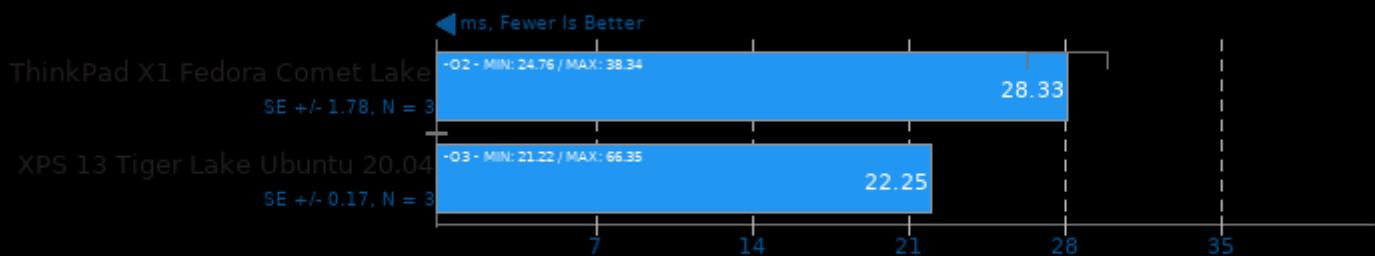
Target: Vulkan GPU - Model: alexnet



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

### NCNN 20201218

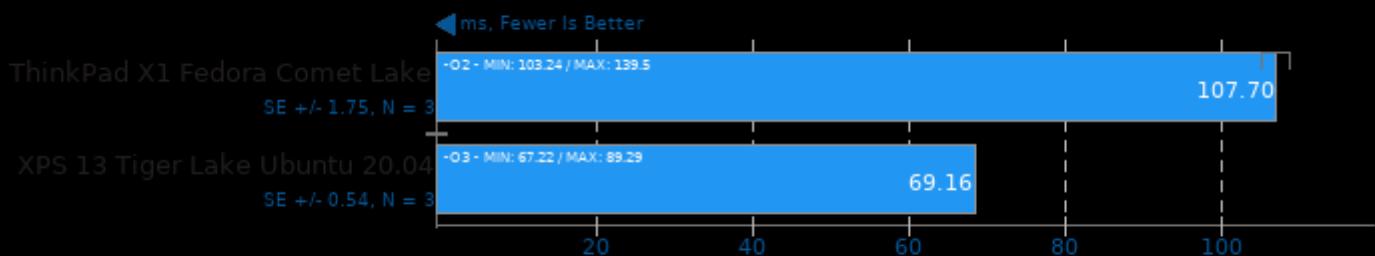
Target: Vulkan GPU - Model: resnet18



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

### NCNN 20201218

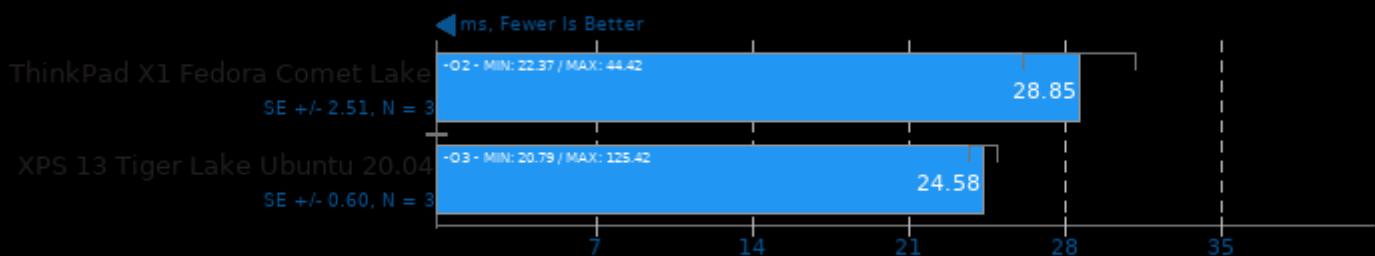
Target: Vulkan GPU - Model: vgg16



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

### NCNN 20201218

Target: Vulkan GPU - Model: googlenet

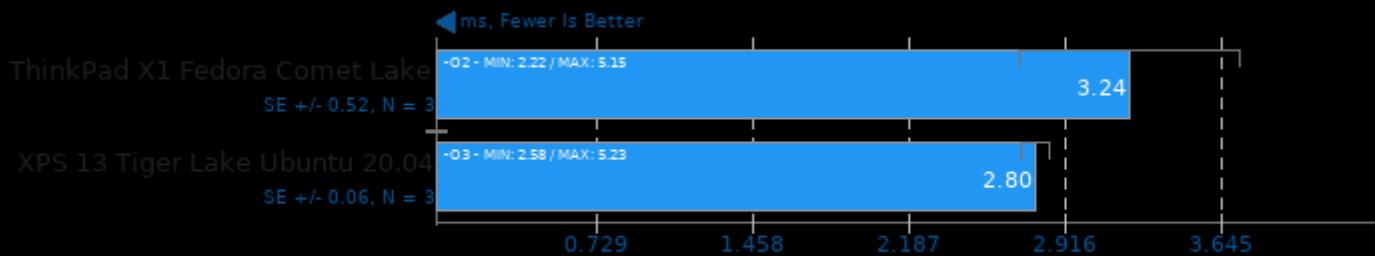


1. (CXX) g++ options: -rdynamic -lgomp -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

### NCNN 20201218

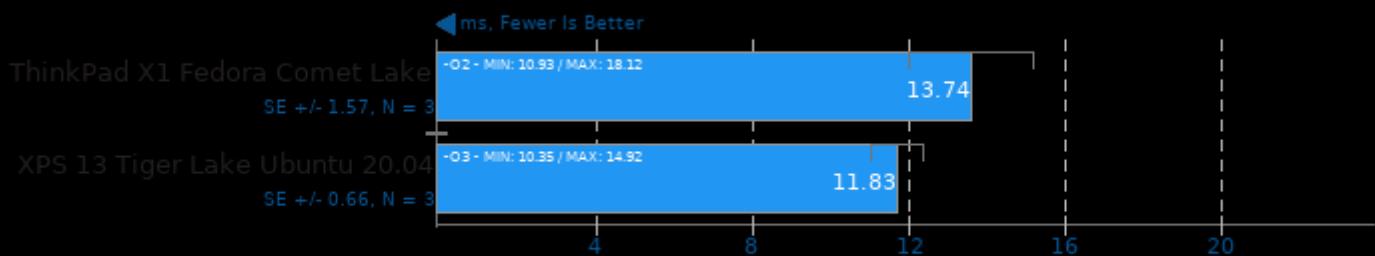
Target: Vulkan GPU - Model: blazeface



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

### NCNN 20201218

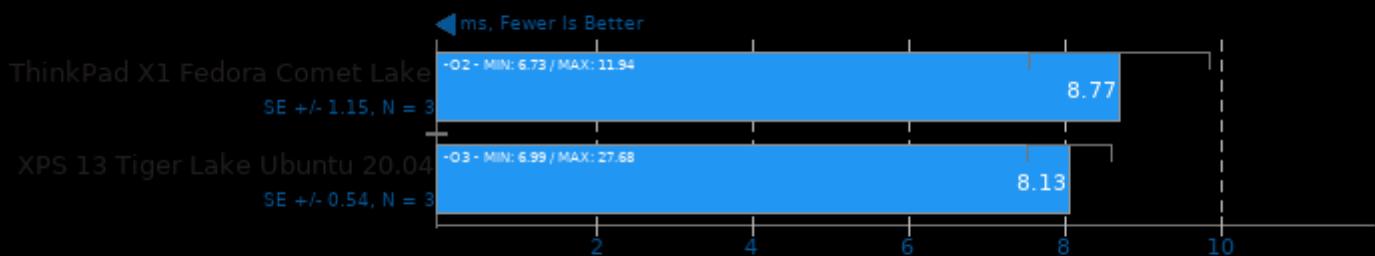
Target: Vulkan GPU - Model: efficientnet-b0



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

### NCNN 20201218

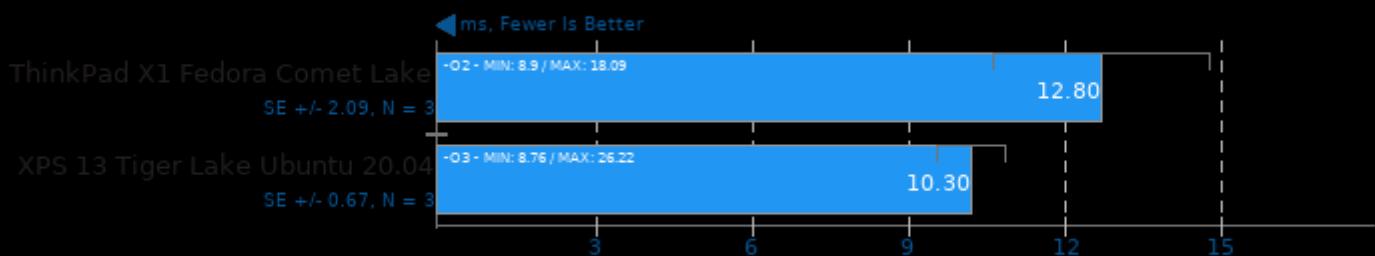
Target: Vulkan GPU - Model: mnasnet



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

### NCNN 20201218

Target: Vulkan GPU - Model: shufflenet-v2

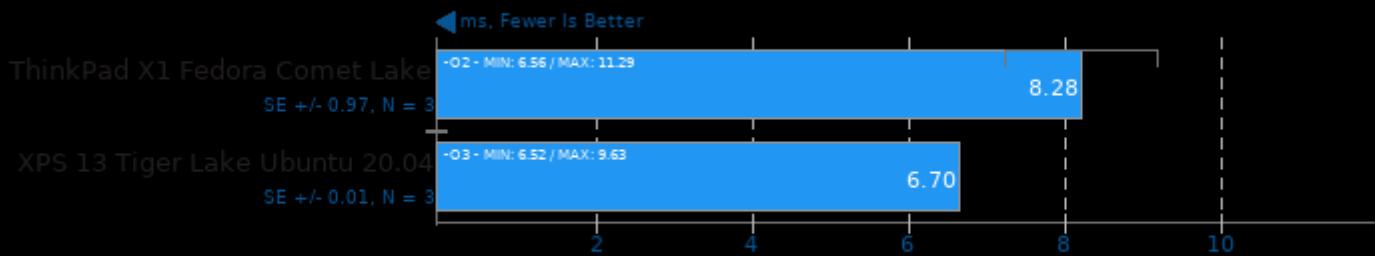


1. (CXX) g++ options: -rdynamic -lgomp -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

### NCNN 20201218

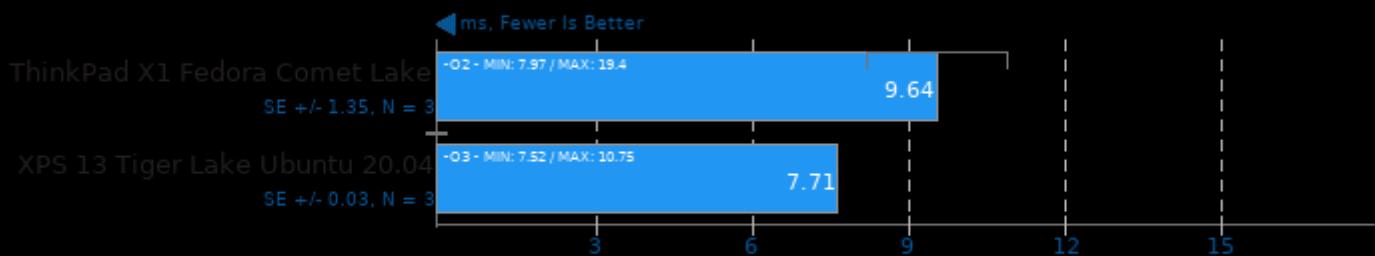
Target: Vulkan GPU-v3-v3 - Model: mobilenet-v3



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

### NCNN 20201218

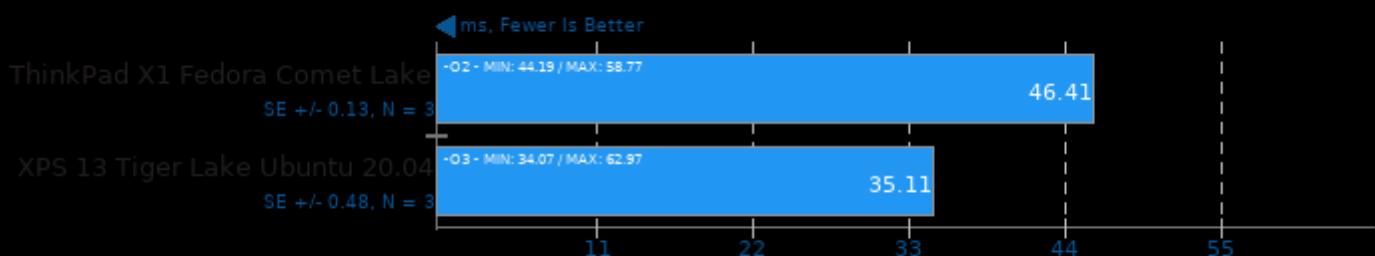
Target: Vulkan GPU-v2-v2 - Model: mobilenet-v2



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

### NCNN 20201218

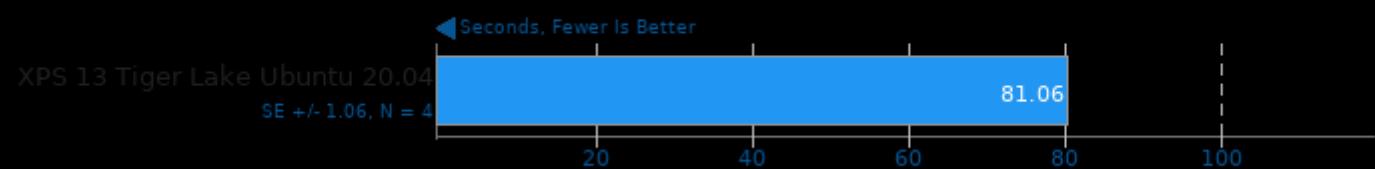
Target: Vulkan GPU - Model: mobilenet



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

### Numenta Anomaly Benchmark 1.1

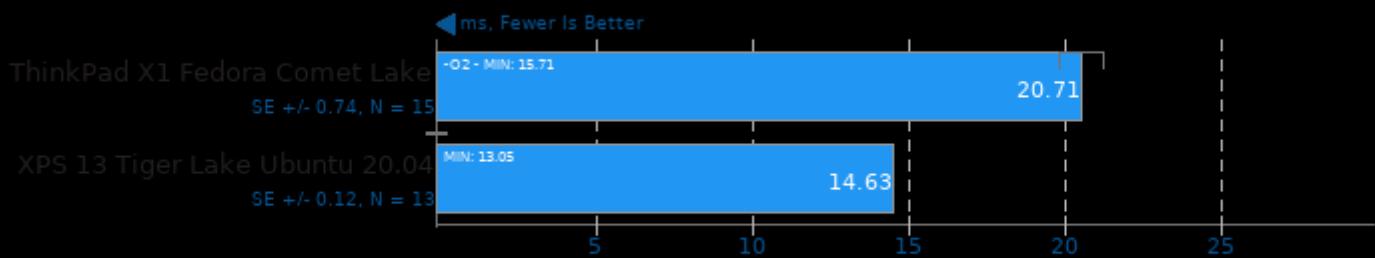
Detector: Bayesian Changepoint



## Coder Radio XPS 13 ML Ubuntu Benchmark

### oneDNN 2.0

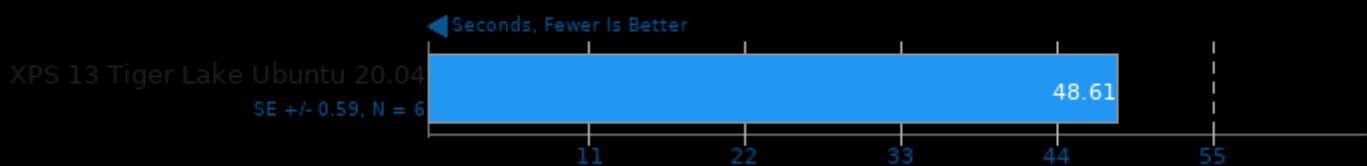
Harness: Deconvolution Batch shapes\_1d - Data Type: f32 - Engine: CPU



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

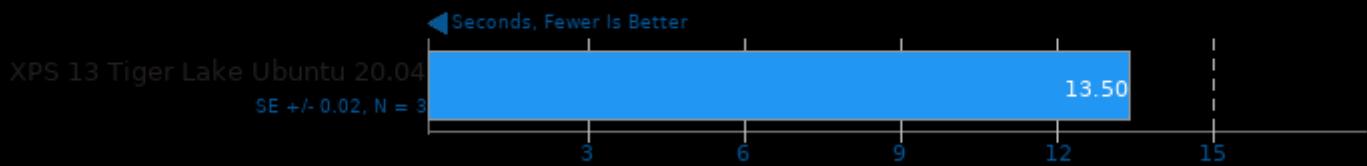
### Numenta Anomaly Benchmark 1.1

Detector: Relative Entropy



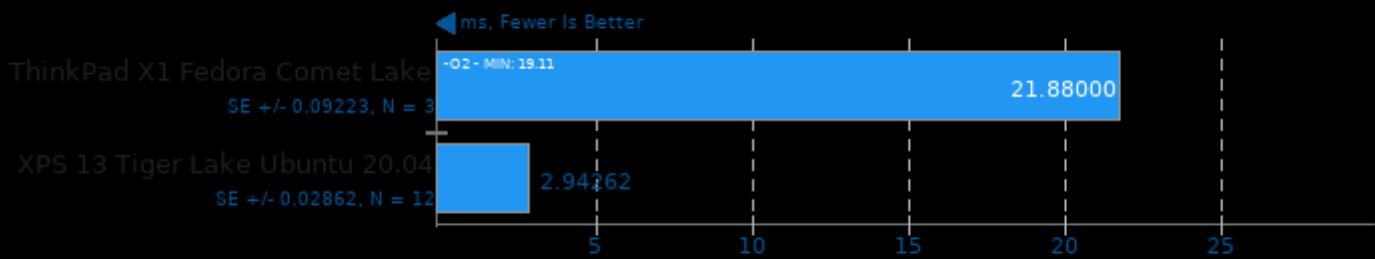
### Milpack Benchmark

Benchmark: scikit\_linearridge\_regression



### oneDNN 2.0

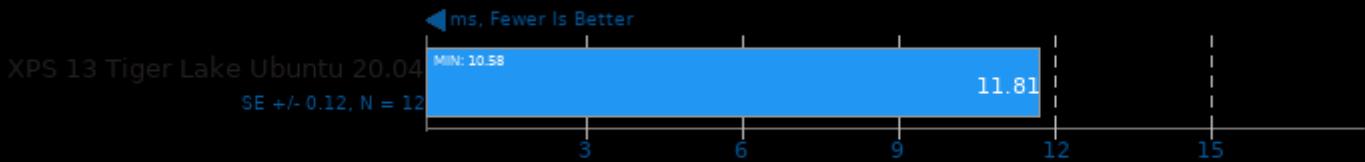
Harness: Deconvolution Batch shapes\_1d - Data Type: u8s8f32 - Engine: CPU



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

### oneDNN 2.0

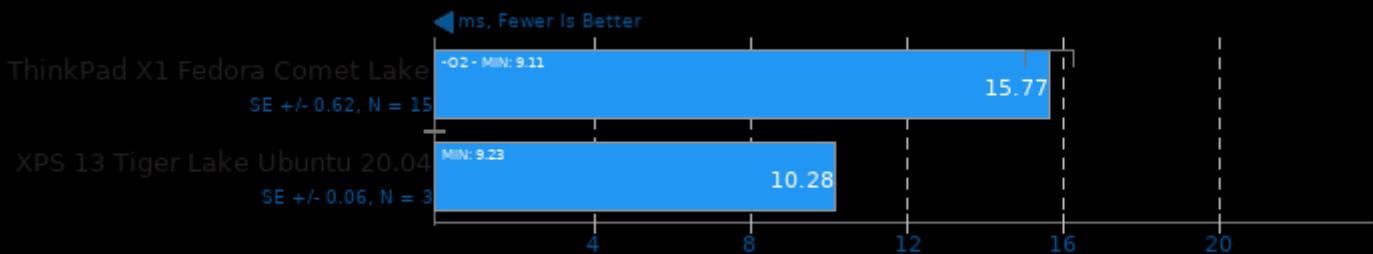
Harness: Matrix Multiply Batch Shapes Transformer - Data Type: bf16bf16bf16 - Engine: CPU



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

### oneDNN 2.0

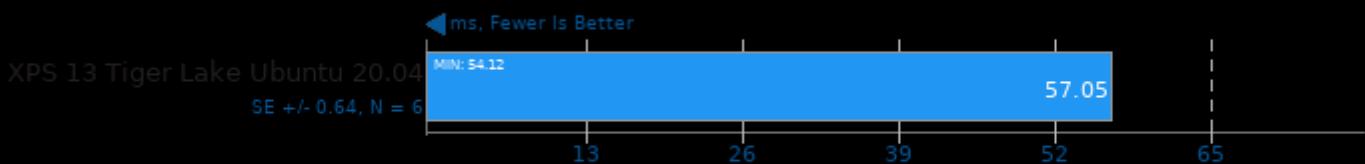
Harness: IP Shapes 1D - Data Type: f32 - Engine: CPU



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

### oneDNN 2.0

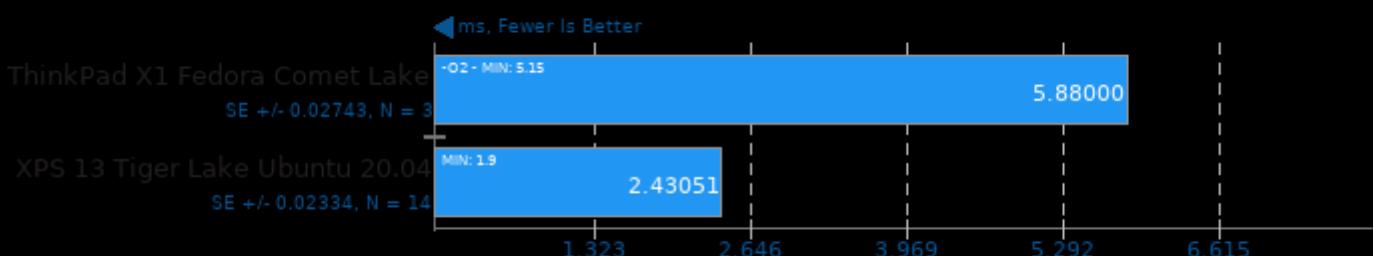
Harness: Deconvolution Batch shapes\_1d - Data Type: bf16bf16bf16 - Engine: CPU



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

### oneDNN 2.0

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

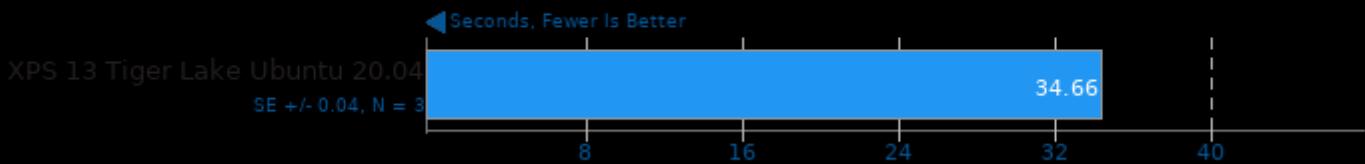


1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fPIC -pie -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

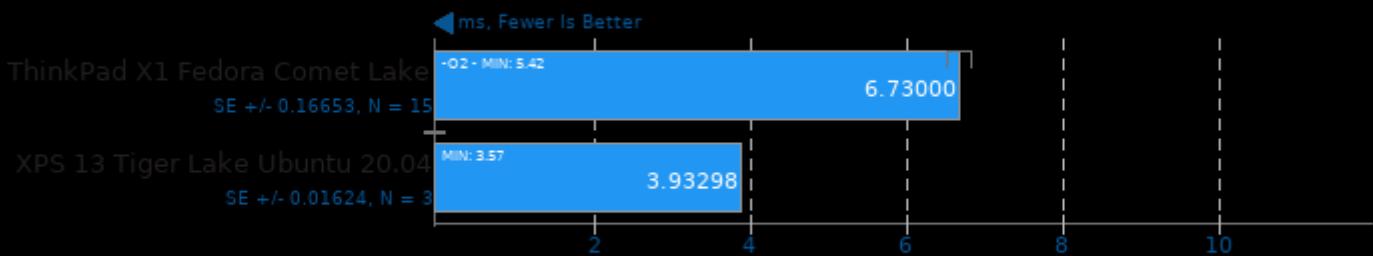
### Milpack Benchmark

Benchmark: scikit\_svm



### oneDNN 2.0

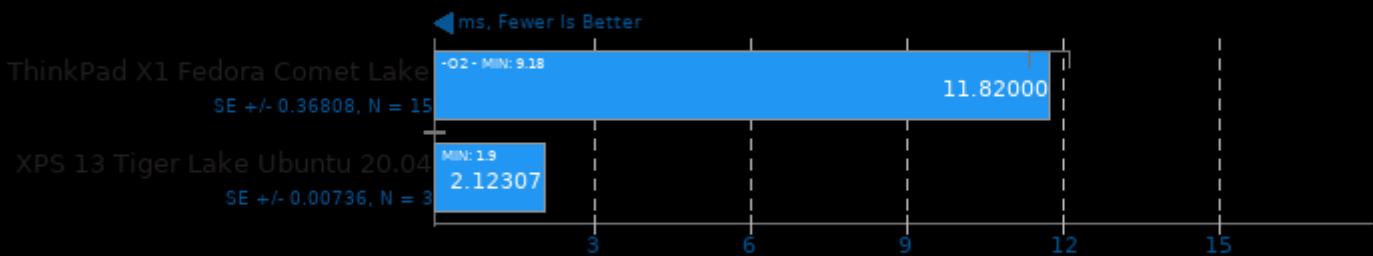
Harness: Matrix Multiply Batch Shapes Transformer - Data Type: f32 - Engine: CPU



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

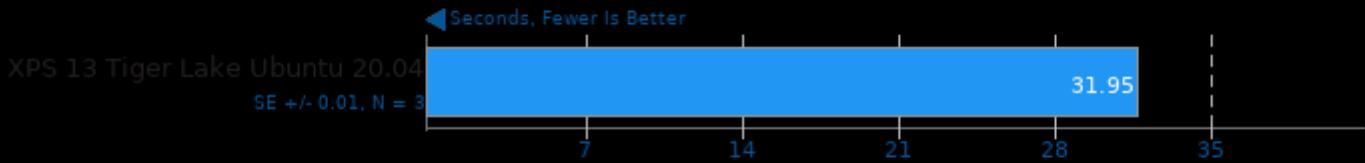
### oneDNN 2.0

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



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

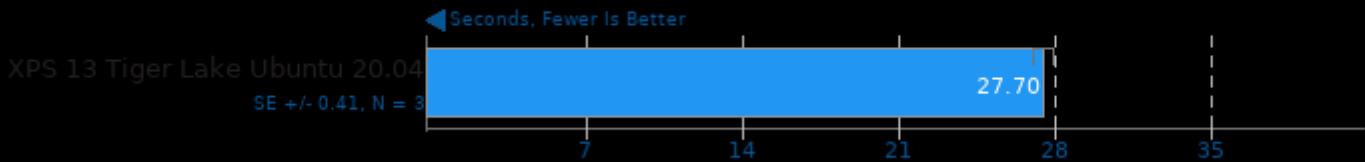
### RNNoise 2020-06-28



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

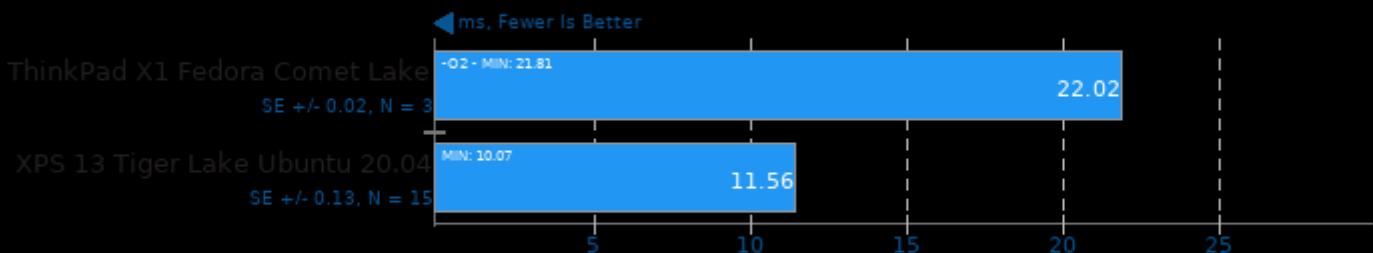
## Numenta Anomaly Benchmark 1.1

Detector: Windowed Gaussian



## oneDNN 2.0

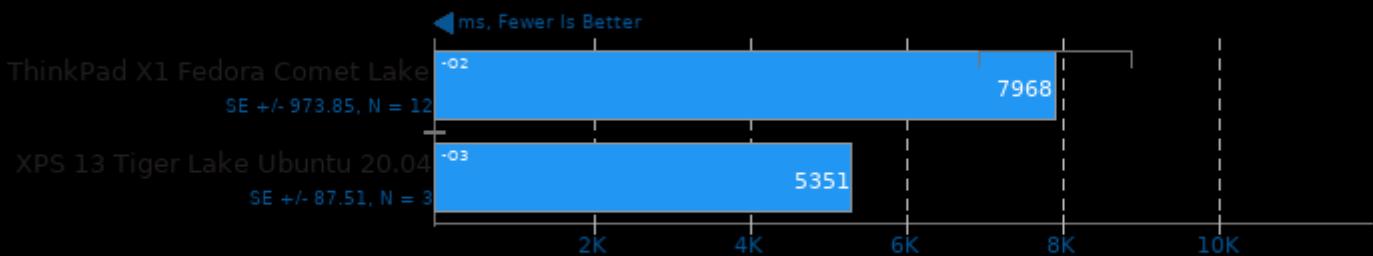
Harness: Convolution Batch Shapes Auto - Data Type: f32 - Engine: CPU



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

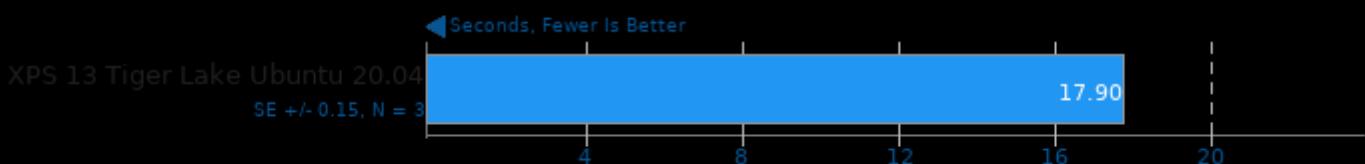
## OpenCV 4.4

Test: DNN - Deep Neural Network



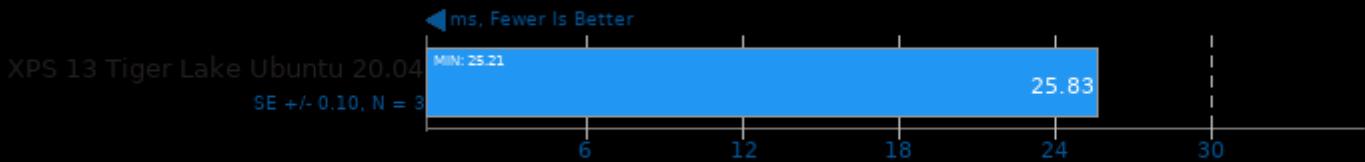
1. (CXX) g++ options: -fsigned-char -pthread -fomit-frame-pointer -ffunction-sections -fdata-sections -msse -msse2 -msse3 -fvisibility=hidden -ldl -lm -lpthread

## Scikit-Learn 0.22.1



## oneDNN 2.0

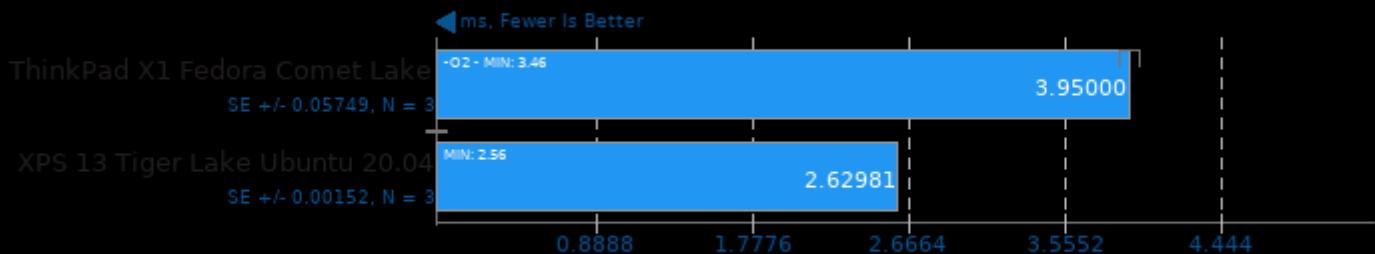
Harness: IP Shapes 1D - Data Type: bf16bf16bf16 - Engine: CPU



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

## oneDNN 2.0

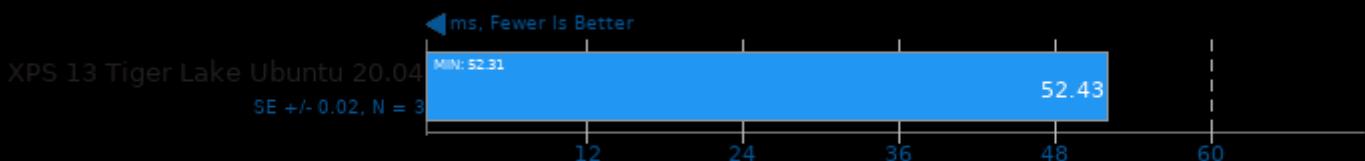
Harness: IP Shapes 3D - Data Type: u8s8f32 - Engine: CPU



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

## oneDNN 2.0

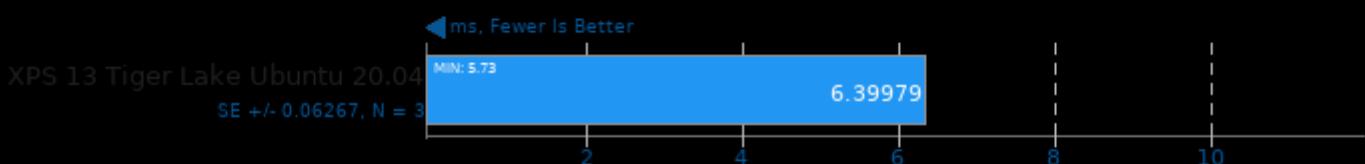
Harness: Convolution Batch Shapes Auto - Data Type: bf16bf16bf16 - Engine: CPU



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

## oneDNN 2.0

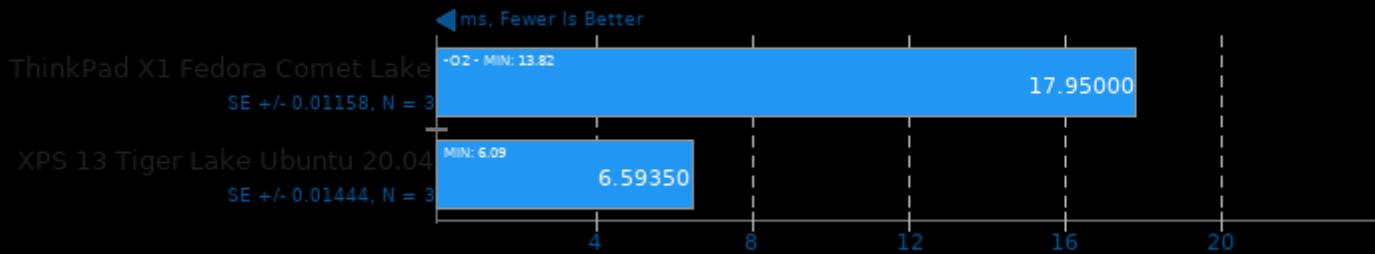
Harness: IP Shapes 3D - Data Type: bf16bf16bf16 - Engine: CPU



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

### oneDNN 2.0

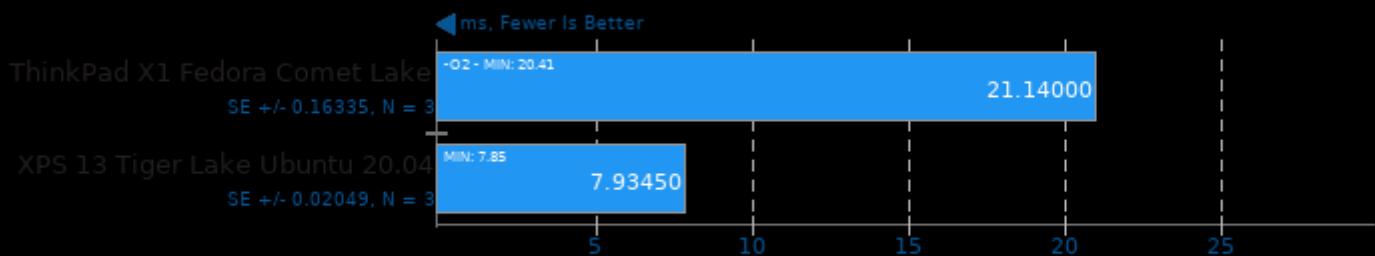
Harness: IP Shapes 3D - Data Type: f32 - Engine: CPU



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

### oneDNN 2.0

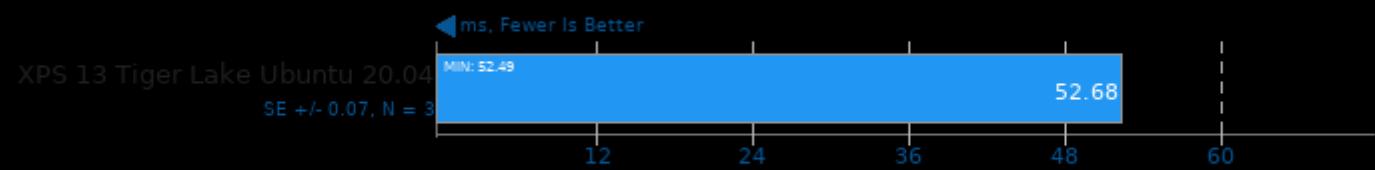
Harness: Convolution Batch Shapes Auto - Data Type: u8s8f32 - Engine: CPU



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

### oneDNN 2.0

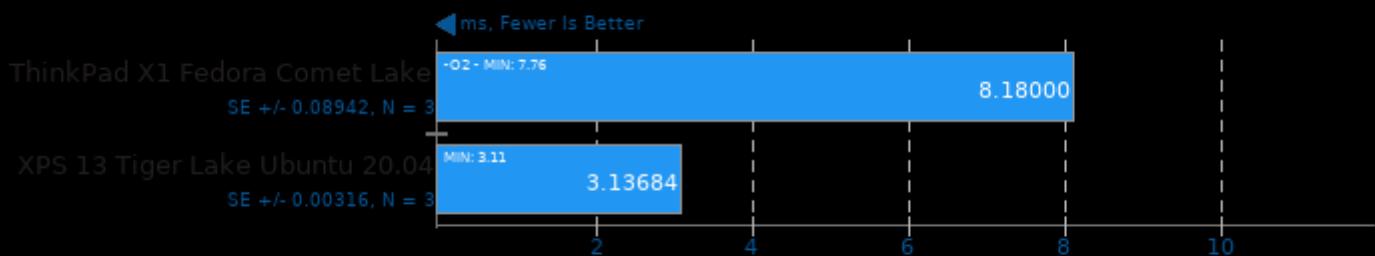
Harness: Deconvolution Batch shapes\_3d - Data Type: bf16bf16bf16 - Engine: CPU



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

### oneDNN 2.0

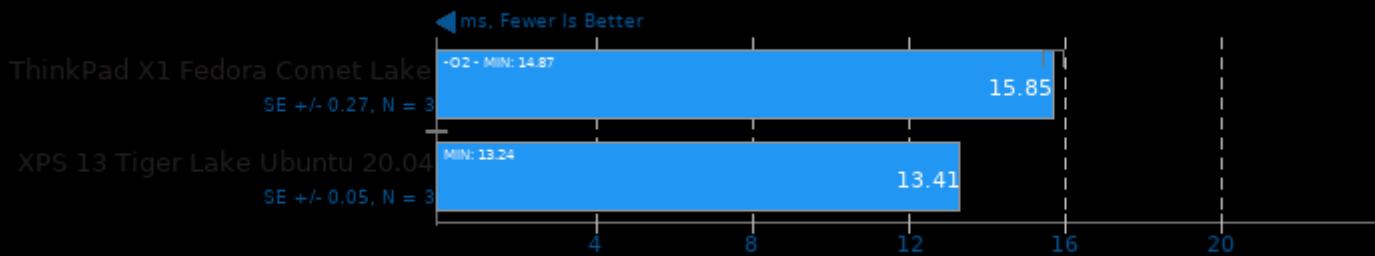
Harness: Deconvolution Batch shapes\_3d - Data Type: u8s8f32 - Engine: CPU



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

## oneDNN 2.0

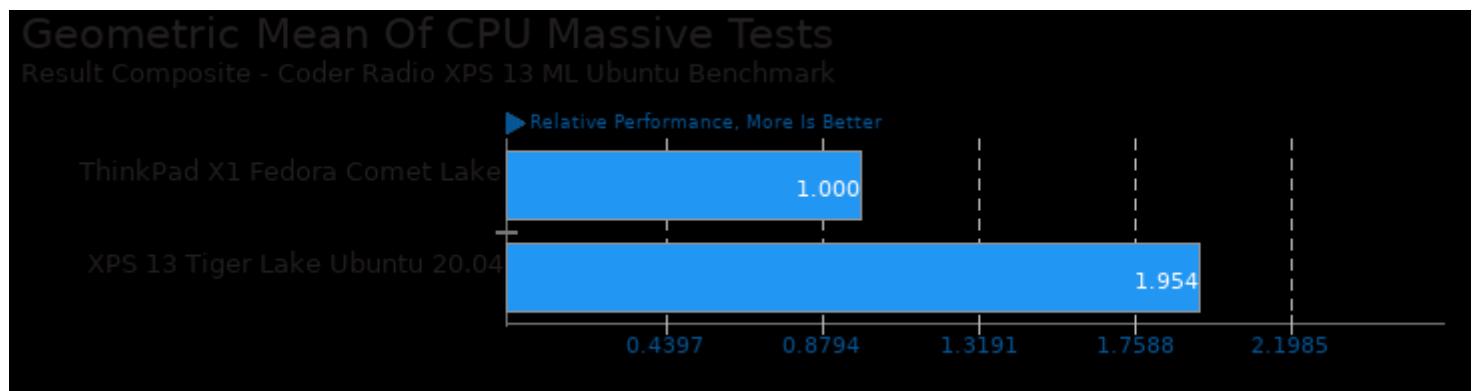
Harness: Deconvolution Batch shapes\_3d - Data Type: f32 - Engine: CPU



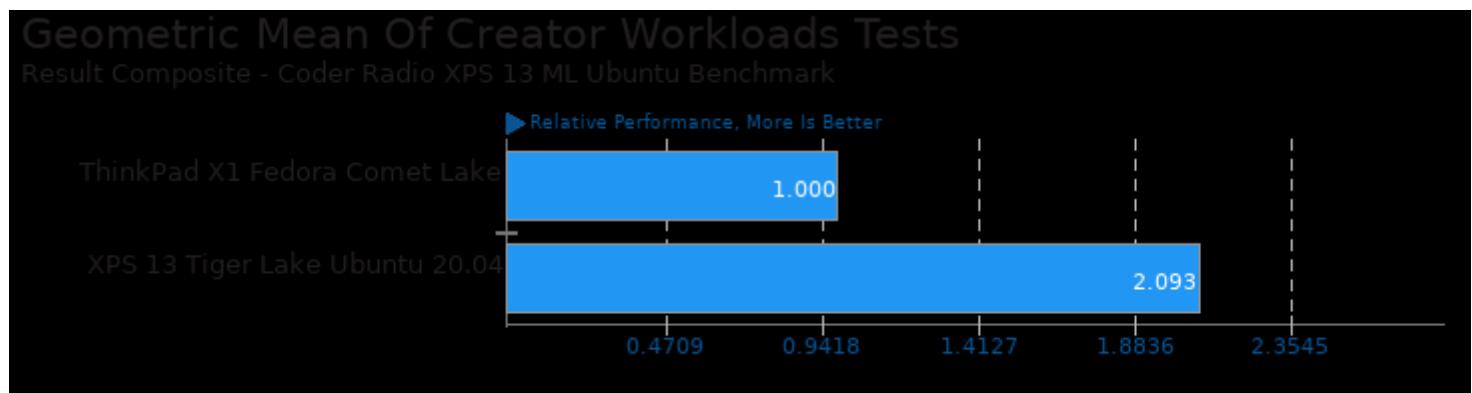
1. (CXX) g++ options: -O3 -std=c++11 -fopenmp -msse4.1 -fPIC -pie -lpthread

## Coder Radio XPS 13 ML Ubuntu Benchmark

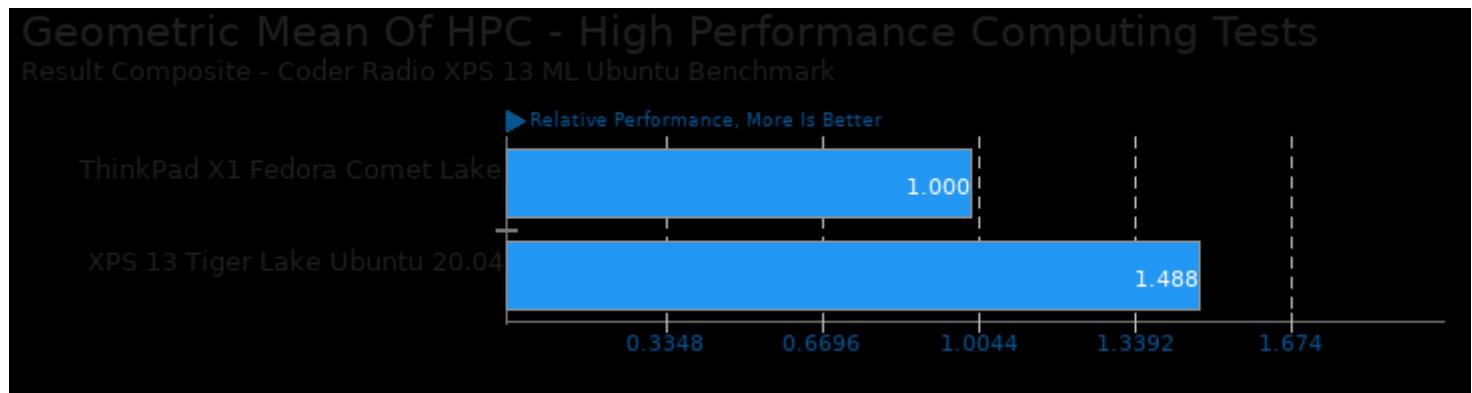
These geometric means are based upon test groupings / test suites for this result file.



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



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

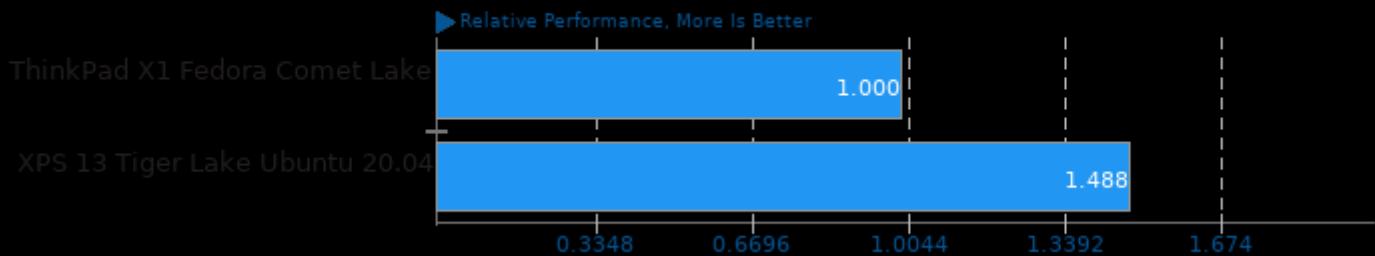


Geometric mean based upon tests: pts/mnn, pts/ncnn, pts/opencv, pts/numpy, pts/ai-benchmark, pts/rnnoise, pts/scikit-learn, pts/mlpack, pts/numenata-nab, pts/tensorflow-lite, pts/onnednn and pts/plaidml

## Coder Radio XPS 13 ML Ubuntu Benchmark

### Geometric Mean Of Machine Learning Tests

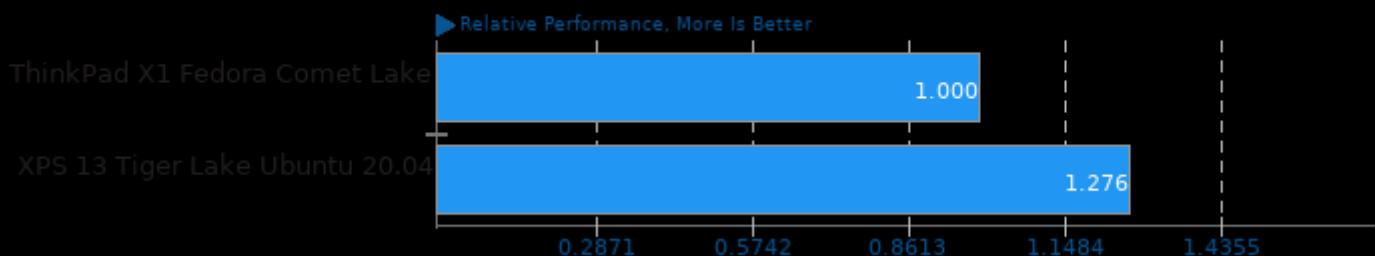
Result Composite - Coder Radio XPS 13 ML Ubuntu Benchmark



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

### Geometric Mean Of NVIDIA GPU Compute Tests

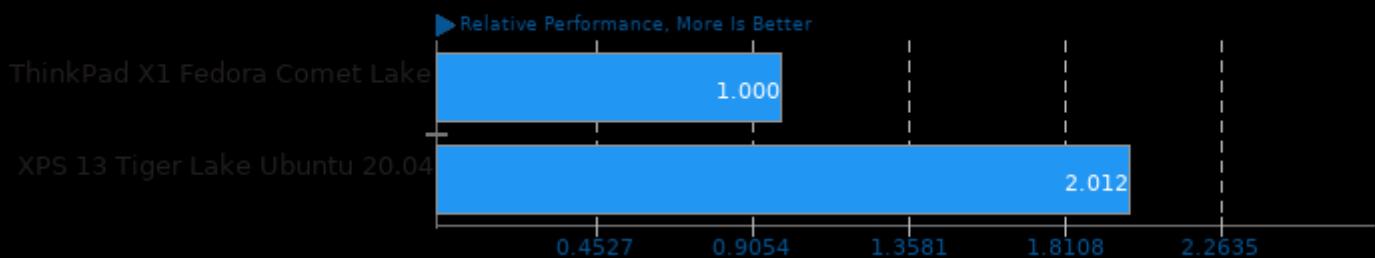
Result Composite - Coder Radio XPS 13 ML Ubuntu Benchmark



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

### Geometric Mean Of Server CPU Tests

Result Composite - Coder Radio XPS 13 ML Ubuntu Benchmark



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

*This file was automatically generated via the Phoronix Test Suite benchmarking software on Friday, 29 March 2024 05:15.*