



Tigerlake Xe Gen12 to Gen9

Tests for a future article

Test Systems:

Gen9 - i7-8565U

Processor: Intel Core i7-8565U @ 4.60GHz (4 Cores / 8 Threads), Motherboard: Dell 0K7W76 (1.0.0 BIOS), Chipset: Intel Cannon Point-LP, Memory: 16GB, Disk: SK hynix PC401 NVMe 256GB, Graphics: Intel UHD 620 3GB (1150MHz), Audio: Realtek ALC3271, Network: Qualcomm Atheros QCA6174 802.11ac

OS: Ubuntu 20.10, Kernel: 5.9.1-050901-generic (x86_64), Desktop: GNOME Shell 3.38.1, Display Server: X Server 1.20.9, Display Driver: modesetting 1.20.9, OpenGL: 4.6 Mesa 20.3.0-devel (git-fdbc45d 2020-10-21 groovy-oibaf-ppa), OpenCL: OpenCL 3.0, Vulkan: 1.2.145, Compiler: GCC 10.2.0, File-System: ext4, Screen Resolution: 1920x1080

```
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++,m2 --enable-libphobos-checking=release --enable-libstdcxx-debug
--enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto
--enable-offload-targets=nvptx-none=/build/gcc-10-JvwpWM/gcc-10-10.2.0/debian/tmp-nvptx/usr,amdgc-n-amdhsa=/build/gcc-10-JvwpWM/gcc-10-10.2.0/debian/tmp-gcn/us
r,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: 0xd6 - Thermald 2.3
Python Notes: Python 3.8.6
Security Notes: itlb_multihit: KVM: Mitigation of VMX disabled + 11tf: Not affected + mds: Mitigation of Clear buffers; SMT vulnerable + meltdown: Not affected +
spec_store_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre_v1: Mitigation of usercopy/swaps barriers and __user pointer sanitization + spectre_v2:
Mitigation of Full generic retpoline IBPB: conditional IBRS_FW STIBP: conditional RSB filling + srbds: Mitigation of Microcode + tsx_async_abort: Not affected
```

Gen11 - i7-1065G7

Processor: Intel Core i7-1065G7 @ 3.90GHz (4 Cores / 8 Threads), Motherboard: Dell 06CDVY (1.0.9 BIOS), Chipset: Intel Device 34ef, Memory: 16GB, Disk: KBG40ZPZ512G NVMe TOSHIBA 512GB, Graphics: Intel Iris Plus G7 3GB (1100MHz), Audio: Realtek ALC289, Network: Intel Killer Wi-Fi 6 AX1650i 160MHz

OS: Ubuntu 20.10, Kernel: 5.9.1-050901-generic (x86_64), Desktop: GNOME Shell 3.38.1, Display Server: X Server 1.20.9, Display Driver: modesetting 1.20.9, OpenGL: 4.6 Mesa 20.3.0-devel (git-4408131 2020-10-20 groovy-oibaf-ppa), OpenCL: OpenCL 3.0, Vulkan: 1.2.145, Compiler: GCC 10.2.0, 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++,m2 --enable-libphobos-checking=release --enable-libstdcxx-debug
--enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto
--enable-offload-targets=nvptx-none=/build/gcc-10-JvwpWM/gcc-10-10.2.0/debian/tmp-nvptx/usr,amdgc-n-amdhsa=/build/gcc-10-JvwpWM/gcc-10-10.2.0/debian/tmp-gcn/us
r,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: 0x78 - Thermald 2.3
Python Notes: Python 3.8.6
Security Notes: itlb_multihit: KVM: Mitigation of VMX disabled + 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
```

Gen12 Xe - i7-1165G7

Processor: Intel Core i7-1165G7 @ 4.70GHz (4 Cores / 8 Threads), Motherboard: Dell 0GG9PT (1.0.3 BIOS), Chipset: Intel Tiger Lake-LP, Memory: 16GB, Disk: KIOXIA KBG40ZNS256G NVMe 256GB, Graphics: Intel UHD 3GB (1300MHz), Audio: Realtek ALC289, Network: Intel Wi-Fi 6 AX201

OS: Ubuntu 20.10, Kernel: 5.9.1-050901-generic (x86_64), Desktop: GNOME Shell 3.38.1, Display Server: X Server 1.20.9, Display Driver: modesetting 1.20.9, OpenGL: 4.6 Mesa 20.3.0-devel (git-fdbc45d 2020-10-21 groovy-oibaf-ppa), OpenCL: OpenCL 3.0, Vulkan: 1.2.145, Compiler: GCC 10.2.0, 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++,m2 --enable-libphobos-checking=release --enable-libstdcxx-debug
--enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto
--enable-offload-targets=nvptx-none=/build/gcc-10-JvwpWM/gcc-10-10.2.0/debian/tmp-nvptx/usr,amdgc-n-amdhsa=/build/gcc-10-JvwpWM/gcc-10-10.2.0/debian/tmp-gcn/us
r,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 2.3
Python Notes: Python 3.8.6
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
```

Gen9 - i7-8565U

Gen11 - i7-1065G7

Gen12 Xe - i7-1165G7

PlaidML - No - Inference - NASNet Large - OpenCL (FPS)	1.84	3.31	5.14
Normalized	35.8%	64.4%	100%
Standard Deviation	0.5%	0.3%	0%
RealSR-NCNN - 4x - Yes (sec)	1615	1120	538.065
Normalized	33.32%	48.04%	100%
Standard Deviation	0.3%	0.5%	0.4%
oneAPI Level Zero Tests - P.I.C (GFLOPS)	138.516	153.147	439.002
Normalized	31.55%	34.89%	100%
Standard Deviation	3%	2.4%	0.5%
Unigine Heaven - 1920 x 1200 - Fullscreen - OpenGL (FPS)	10.3701	16.1363	28.6453
Normalized	36.2%	56.33%	100%
Standard Deviation	0.3%	4.5%	0.4%
PlaidML - No - Inference - DenseNet 201 - OpenCL (FPS)	5.81	11.30	22.44
Normalized	25.89%	50.36%	100%
Standard Deviation	0.4%	1.2%	0.7%
PlaidML - No - Inference - VGG16 - OpenCL (FPS)	8.18	19.93	22.18
Normalized	36.88%	89.86%	100%
Standard Deviation	1.6%	5.9%	0.3%
LeelaChessZero - OpenCL (Nodes/s)	389	579	1535
Normalized	25.34%	37.72%	100%
Standard Deviation	2.3%	2.4%	0.9%
PlaidML - No - Inference - IMDB LSTM - OpenCL (FPS)	7.84	20.10	30.52
Normalized	25.69%	65.86%	100%
Standard Deviation	0.4%	1%	2.3%
PlaidML - No - Inference - VGG19 - OpenCL (FPS)	6.42	16.03	17.49
Normalized	36.71%	91.65%	100%
Standard Deviation	0.2%	2.2%	0.2%
oneAPI Level Zero Tests - P.S.P.C (GB/s)	418.862	920.207	1220
Normalized	34.35%	75.46%	100%
Standard Deviation	0.7%	0%	0%
Unigine Superposition - 1920 x 1200 - Fullscreen - Medium - OpenGL (FPS)	5.8	9.0	16.4
Normalized	35.37%	54.88%	100%
Standard Deviation	1%	0.6%	0%
Unigine Valley - 1920 x 1200 - Fullscreen - OpenGL (FPS)	11.4067	17.7436	30.6054
Normalized	37.27%	57.98%	100%
Standard Deviation	1%	1.1%	0.5%
Unigine Superposition - 1920 x 1200 - Fullscreen - Low - OpenGL (FPS)	12.2	17.0	29.6
Normalized	41.22%	57.43%	100%
Standard Deviation	0.5%	0.9%	0.8%
PlaidML - No - Inference - Inception V3 - OpenCL (FPS)	15.90	34.82	46.37
Normalized	34.29%	75.09%	100%
Standard Deviation	1.7%	0.4%	0.1%
ET: Legacy - Renderer2 - 1920 x 1200 (FPS)	72.1	86.4	122.3
Normalized	58.95%	70.65%	100%

	Standard Deviation	0.1%	8.2%	2.9%
oneAPI Level Zero Tests - P.S.M.C.t.S.M		10.7883	13.8085	15.0859
	Normalized	71.51%	91.53%	100%
	Standard Deviation	2.9%	0.3%	0.5%
RealSR-NCNN - 4x - No (sec)		202.601	143.988	68.311
	Normalized	33.72%	47.44%	100%
	Standard Deviation	0.5%	0.4%	1.1%
SHOC Scalable Heterogeneous Computing - OpenCL - T.R.B (GB/s)		54.2625	174.162	189.757
	Normalized	28.6%	91.78%	100%
	Standard Deviation	0%	1.7%	0%
Xonotic - 1920 x 1200 - Ultimate (FPS)		71.6067804	99.2365089	121.9041377
	Normalized	58.74%	81.41%	100%
	Standard Deviation	1.1%	2.7%	1.5%
cl-mem - Write (GB/s)		24.0	34.4	47.9
	Normalized	50.1%	71.82%	100%
	Standard Deviation	2.7%	5%	0.8%
Xonotic - 1920 x 1200 - Low (FPS)		183.6211229	209.1609960	307.6345501
	Normalized	59.69%	67.99%	100%
	Standard Deviation	0.1%	2.9%	2.9%
Xonotic - 1920 x 1200 - Ultra (FPS)		93.2094800	123.2383929	155.3708840
	Normalized	59.99%	79.32%	100%
	Standard Deviation	1%	2.6%	1.5%
PlaidML - No - Inference - ResNet 50 - OpenCL (FPS)		34.52	75.57	87.29
	Normalized	39.55%	86.57%	100%
	Standard Deviation	1.8%	2.4%	0.1%
NCNN - Vulkan GPU - yolov4-tiny (ms)		62.81	38.95	15.26
	Normalized	24.3%	39.18%	100%
	Standard Deviation	0.8%	0.7%	0.2%
NCNN - Vulkan GPU - resnet50 (ms)		51.46	35.91	15.49
	Normalized	30.1%	43.14%	100%
	Standard Deviation	0.1%	0.2%	0.2%
NCNN - Vulkan GPU - alexnet (ms)		25.77	16.63	9.47
	Normalized	36.75%	56.95%	100%
	Standard Deviation	0.1%	0.3%	0.1%
NCNN - Vulkan GPU - resnet18 (ms)		23.89	17.12	7.18
	Normalized	30.05%	41.94%	100%
	Standard Deviation	0%	0.2%	0.6%
NCNN - Vulkan GPU - vgg16 (ms)		152.79	96.32	38.44
	Normalized	25.16%	39.91%	100%
	Standard Deviation	0.1%	0.1%	0.1%
NCNN - Vulkan GPU - googlenet (ms)		23.72	18.62	10.51
	Normalized	44.31%	56.44%	100%
	Standard Deviation	0.5%	0.1%	0.5%
NCNN - Vulkan GPU - blazeface (ms)		1.81	1.81	0.98
	Normalized	54.14%	54.14%	100%
	Standard Deviation	3.9%	1.8%	3.5%
NCNN - Vulkan GPU - efficientnet-b0 (ms)		16.52	16.72	11.05
	Normalized	66.89%	66.09%	100%
	Standard Deviation	0.1%	0.1%	0.4%
NCNN - Vulkan GPU - mnasnet (ms)		8.36	7.65	5.14
	Normalized	61.48%	67.19%	100%
	Standard Deviation	0.3%	1.1%	9.7%
NCNN - Vulkan GPU - shufflenet-v2 (ms)		6.75	5.71	3.05

	Normalized	45.19%	53.42%	100%
	Standard Deviation	0.5%	1%	0.9%
NCNN - Vulkan GPU-v3-v3 - mobilenet-v3		9.19	8.65	5.47
	Normalized	59.52%	63.24%	100%
	Standard Deviation	0.3%	0.2%	1.4%
NCNN - Vulkan GPU-v2-v2 - mobilenet-v2		8.27	7.53	4.62
	Normalized	55.86%	61.35%	100%
	Standard Deviation	0.8%	1.3%	2.3%
NCNN - Vulkan GPU - mobilenet (ms)		25.74	22.36	9.74
	Normalized	37.84%	43.56%	100%
	Standard Deviation	2.6%	1.6%	1.2%
NCNN - Vulkan GPU - squeezenet (ms)		28.27	24.65	11.51
	Normalized	40.71%	46.69%	100%
	Standard Deviation	0.4%	2.2%	0.3%
Xonotic - 1920 x 1200 - High (FPS)		111.6753218	144.1207832	183.8028138
	Normalized	60.76%	78.41%	100%
	Standard Deviation	1.3%	2.8%	2.4%
GpuTest - Furmark - 1920 x 1200 - Fullscreen (Points)		672	1062	1782
	Normalized	37.71%	59.6%	100%
	Standard Deviation	2.2%	2.5%	2.6%
GpuTest - GiMark - 1920 x 1200 - Fullscreen (Points)		1314	1724	2017
	Normalized	65.15%	85.47%	100%
	Standard Deviation	0.8%	1%	0.1%
Waifu2x-NCNN Vulkan - 2x - 3 - Yes (sec)		75.252	53.103	27.216
	Normalized	36.17%	51.25%	100%
	Standard Deviation	0.1%	3%	0.2%
GpuTest - Pixmark Piano - 1920 x 1200 - Fullscreen (Points)		149	288	584
	Normalized	25.51%	49.32%	100%
	Standard Deviation		1%	1.9%
GpuTest - TessMark - 1920 x 1200 - Fullscreen (Points)		2840	3602	5080
	Normalized	55.91%	70.91%	100%
	Standard Deviation	1.5%	1.3%	0.9%
GpuTest - Pixmark Volplosion - 1920 x 1200 - Fullscreen (Points)		390	761	1595
	Normalized	24.45%	47.71%	100%
	Standard Deviation		1.1%	1.6%
oneAPI Level Zero Tests - P.H.P.C (GFLOPS)		751.623	1352	3037
	Normalized	24.75%	44.53%	100%
	Standard Deviation	0.1%	2.5%	0.1%
MandelGPU - GPU (Samples/sec)		10702353	23689429	44609195
	Normalized	23.99%	53.1%	100%
	Standard Deviation	0%	1.7%	1.2%
cl-mem - Read (GB/s)		22.7	41.3	57.1
	Normalized	39.75%	72.33%	100%
	Standard Deviation	1.8%	0.7%	2.3%
cl-mem - Copy (GB/s)		21.6	35.9	48.0
	Normalized	45%	74.79%	100%
	Standard Deviation	1.9%	0.5%	0.1%

PlaidML - No - Inference - Mobilenet - OpenCL (FPS)	116.08	242.83	299.60
Normalized	38.74%	81.05%	100%
Standard Deviation	0%	2.4%	0.5%
SHOC Scalable Heterogeneous Computing - OpenCL - Max SP Flops (GFLOPS)	1681	5323	7821
Normalized	21.5%	68.06%	100%
Standard Deviation	0.1%	0.1%	0%
oneAPI Level Zero Tests - P.F.G.M.B (GB/s)	24.5276	34.7522	56.9391
Normalized	43.08%	61.03%	100%
Standard Deviation	2.1%	0.1%	0.5%
SHOC Scalable Heterogeneous Computing - OpenCL - MD5 Hash (GHash/s)	0.3691	0.6778	1.6775
Normalized	22%	40.41%	100%
Standard Deviation	0%	0.1%	0%
Tesseract - 1920 x 1200 (FPS)	65.22251	109.4669	140.3486
Normalized	46.47%	78%	100%
Standard Deviation	1.2%	2.3%	1.1%
FinanceBench - Monte-Carlo OpenCL (ms)	2585	579.523733	433.370858
Normalized	16.77%	74.78%	100%
Standard Deviation	8.3%	7.7%	0.9%
oneAPI Level Zero Tests - H.T.D.T.H.I.C	4.83779	10.5736	21.6345
Normalized	22.36%	48.87%	100%
Standard Deviation	0.6%	2.9%	0.9%
JuliaGPU - GPU (Samples/sec)	47930813	85184313	132508055
Normalized	36.17%	64.29%	100%
Standard Deviation	0%	3%	2.6%
clpeak - T.B.e (GBPS)	31.64	39.48	52.15
Normalized	60.67%	75.7%	100%
Standard Deviation	0.1%	6.4%	1%
clpeak - S.P.F (GFLOPS)	438.06	1039	1775
Normalized	24.68%	58.52%	100%
Standard Deviation	0%	0.1%	0%
clpeak - G.M.B (GBPS)	26.73	45.94	57.02
Normalized	46.88%	80.57%	100%
Standard Deviation	0.6%	0.1%	0.1%
oneAPI Level Zero Tests - D.T.H.B (usec)	19941	12233	10057
Normalized	50.44%	82.22%	100%
Standard Deviation	1%	0.1%	0.2%
oneAPI Level Zero Tests - D.T.H.B (GB/s)	13.462183	21.943939	26.690187
Normalized	50.44%	82.22%	100%
Standard Deviation	1%	0.1%	0.2%
oneAPI Level Zero Tests - H.T.D.B (usec)	19750	12241	10049
Normalized	50.88%	82.09%	100%
Standard Deviation	0.9%	0.2%	0.3%
oneAPI Level Zero Tests - H.T.D.B (GB/s)	13.592119	21.928612	26.714162
Normalized	50.88%	82.09%	100%
Standard Deviation	0.9%	0.2%	0.3%
Waifu2x-NCNN Vulkan - 2x - 3 - No (sec)	10.288	7.062	4.095
Normalized	39.8%	57.99%	100%
Standard Deviation	0.6%	0.9%	0.9%
clpeak - Kernel Latency (us)	34.12	36.89	29.53
Normalized	86.55%	80.05%	100%
Standard Deviation	2.3%	0.7%	0.5%

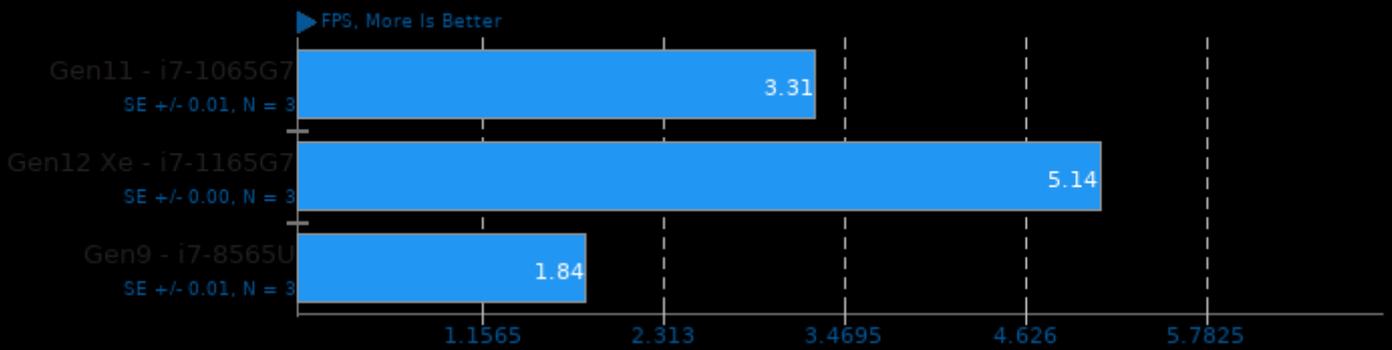
SHOC Scalable HeterOgeneous Computing - OpenCL - Triad (GB/s)	10.9236	14.0210	21.3040
Normalized	51.27%	65.81%	100%
Standard Deviation	6.8%	7.4%	1.9%
SHOC Scalable HeterOgeneous Computing - OpenCL - FFT SP (GFLOPS)	38.9719	140.162	156.162
Normalized	24.96%	89.75%	100%
Standard Deviation	0.3%	0.5%	0.8%
ViennaCL - O.L.F (GFLOPS)	24.6746	48.9117	73.8162
Normalized	33.43%	66.26%	100%
Standard Deviation	0.3%	1.3%	0.4%
SHOC Scalable HeterOgeneous Computing - OpenCL - Bus Speed Readback (GB/s)	23.9319	41.7406	60.4048
Normalized	39.62%	69.1%	100%
Standard Deviation	9.3%	3.9%	1%
oneAPI Level Zero Tests - P.K.L.L (us)	31.4224	28.9568	22.3130
Normalized	71.01%	77.06%	100%
Standard Deviation	1.6%	11%	9.6%
SHOC Scalable HeterOgeneous Computing - OpenCL - Bus Speed Download (GB/s)	23.3048	40.1640	58.4143
Normalized	39.9%	68.76%	100%
Standard Deviation	8.6%	1%	2.9%
FinanceBench - B.S.O (ms)	25.826923	6.343	4.179
Normalized	16.18%	65.88%	100%
Standard Deviation	0.6%	3.2%	0.5%
GpuTest - Pixmark Volplosion - 1920 x 1200 - Fullscreen (Points/Watt)	16.06	45.24	75.39
Normalized	21.3%	60.01%	100%
GpuTest - Pixmark Piano - 1920 x 1200 - Fullscreen (Points/Watt)	6.01	17.13	27.39
Normalized	21.94%	62.54%	100%
GpuTest - TessMark - 1920 x 1200 - Fullscreen (Points/Watt)	100.32	214.06	242.99
Normalized	41.29%	88.09%	100%
GpuTest - Furmark - 1920 x 1200 - Fullscreen (Points/Watt)	25.26	63.19	81.15
Normalized	31.13%	77.87%	100%
GpuTest - GiMark - 1920 x 1200 - Fullscreen (Points/Watt)	47.23	104.65	111.04
Normalized	42.53%	94.25%	100%
Xonotic - 1920 x 1200 - Ultimate (FPS/Watt)	2.65	6.10	5.56
Normalized	43.44%	100%	91.15%
Xonotic - 1920 x 1200 - Ultra (FPS/Watt)	3.44	7.51	7.06
Normalized	45.81%	100%	94.01%
Xonotic - 1920 x 1200 - High (FPS/Watt)	4.25	8.81	8.28
Normalized	48.24%	100%	93.98%
Xonotic - 1920 x 1200 - Low (FPS/Watt)	8.37	12.89	14.14
Normalized	59.19%	91.16%	100%
Unigine Valley - 1920 x 1200 - Fullscreen - OpenGL (FPS/Watt)	0.43	1.13	1.40
Normalized	30.71%	80.71%	100%

Unigine Superposition - 1920 x 1200 - Fullscreen - Medium - OpenGL (FPS/Watt)	0.23	0.50	0.74
Normalized	31.08%	67.57%	100%
Unigine Superposition - 1920 x 1200 - Fullscreen - Low - OpenGL (FPS/Watt)	0.48	0.89	1.34
Normalized	35.82%	66.42%	100%
Unigine Heaven - 1920 x 1200 - Fullscreen - OpenGL (FPS/Watt)	0.38	0.81	1.31
Normalized	29.01%	61.83%	100%
Tesseract - 1920 x 1200 (FPS/Watt)	2.61	5.00	6.28
Normalized	41.56%	79.62%	100%
ET: Legacy - Renderer2 - 1920 x 1200 (FPS/Watt)	3.41	5.40	5.59
Normalized	61%	96.6%	100%
clpeak - T.B.e (GBPS/Watt)	2.60	3.45	3.44
Normalized	75.36%	100%	99.71%
clpeak - G.M.B (GBPS/Watt)	0.92	2.71	2.54
Normalized	33.95%	100%	93.73%
clpeak - S.P.F (GFLOPS/Watt)	20.31	58.93	80.98
Normalized	25.08%	72.77%	100%
MandelGPU - GPU (Samples/sec/Watt)	468417	1544709	1974932
Normalized	23.72%	78.22%	100%
JuliaGPU - GPU (Samples/sec/Watt)	1343813	5650341	5935646
Normalized	22.64%	95.19%	100%
PlaidML - No - Inference - NASNet Large - OpenCL (FPS/Watt)	0.06	0.22	0.24
Normalized	25%	91.67%	100%
PlaidML - No - Inference - Inception V3 - OpenCL (FPS/Watt)	0.54	2.31	2.14
Normalized	23.38%	100%	92.64%
PlaidML - No - Inference - DenseNet 201 - OpenCL (FPS/Watt)	0.22	0.75	1.02
Normalized	21.57%	73.53%	100%
PlaidML - No - Inference - ResNet 50 - OpenCL (FPS/Watt)	1.10	4.79	4.04
Normalized	22.96%	100%	84.34%
PlaidML - No - Inference - Mobilenet - OpenCL (FPS/Watt)	3.63	15.04	13.59
Normalized	24.14%	100%	90.36%
PlaidML - No - Inference - IMDB LSTM - OpenCL (FPS/Watt)	0.24	1.24	1.40
Normalized	17.14%	88.57%	100%
PlaidML - No - Inference - VGG19 - OpenCL (FPS/Watt)	0.21	0.98	0.86
Normalized	21.43%	100%	87.76%
PlaidML - No - Inference - VGG16 - OpenCL (FPS/Watt)	0.26	1.21	1.07
Normalized	21.49%	100%	88.43%
LeelaChessZero - OpenCL (Nodes/s/Watt)	14.05	37.81	70.05
Normalized	20.06%	53.98%	100%
cl-mem - Write (GB/s/Watt)	1.21	2.31	2.08
Normalized	52.38%	100%	90.04%

cl-mem - Read (GB/s/Watt)	1.15	2.70	2.45
Normalized	42.59%	100%	90.74%
cl-mem - Copy (GB/s/Watt)	1.09	1.95	2.02
Normalized	53.96%	96.53%	100%
ViennaCL - O.L.F (GFLOPS/Watt)	1.86	5.86	6.64
Normalized	28.01%	88.25%	100%
SHOC Scalable Heterogeneous Computing - OpenCL - T.R.B (GB/s/Watt)	2.59	10.35	8.35
Normalized	25.02%	100%	80.68%
SHOC Scalable Heterogeneous Computing - OpenCL - Bus Speed Readback (GB/s/Watt)	2.69	6.82	7.39
Normalized	36.4%	92.29%	100%
SHOC Scalable Heterogeneous Computing - OpenCL - Bus Speed Download (GB/s/Watt)	2.65	6.84	7.04
Normalized	37.64%	97.16%	100%
SHOC Scalable Heterogeneous Computing - OpenCL - Max SP Flops (GFLOPS/Watt)	72.86	265.36	343.74
Normalized	21.2%	77.2%	100%
SHOC Scalable Heterogeneous Computing - OpenCL - MD5 Hash (GHash/s/Watt)	0.02	0.03	0.07
Normalized	28.57%	42.86%	100%
SHOC Scalable Heterogeneous Computing - OpenCL - FFT SP (GFLOPS/Watt)	2.21	17.77	17.81
Normalized	12.41%	99.78%	100%
SHOC Scalable Heterogeneous Computing - OpenCL - Triad (GB/s/Watt)	0.92	1.68	1.92
Normalized	47.92%	87.5%	100%
oneAPI Level Zero Tests - P.S.M.C.t.S.M (GB/s/Watt)	0.36	0.90	0.67
Normalized	40%	100%	74.44%
oneAPI Level Zero Tests - P.F.G.M.B (GB/s/Watt)	0.85	2.24	2.49
Normalized	34.14%	89.96%	100%
oneAPI Level Zero Tests - H.T.D.T.H.I.C (GB/s/Watt)	0.14	0.68	0.94
Normalized	14.89%	72.34%	100%
oneAPI Level Zero Tests - P.S.P.C	13.93	59.52	54.99
Normalized	23.4%	100%	92.39%
oneAPI Level Zero Tests - P.H.P.C (GFLOPS/Watt)	24.17	73.98	135.36
Normalized	17.86%	54.65%	100%
oneAPI Level Zero Tests - P.I.C (GFLOPS/Watt)	4.31	9.15	20.00
Normalized	21.55%	45.75%	100%

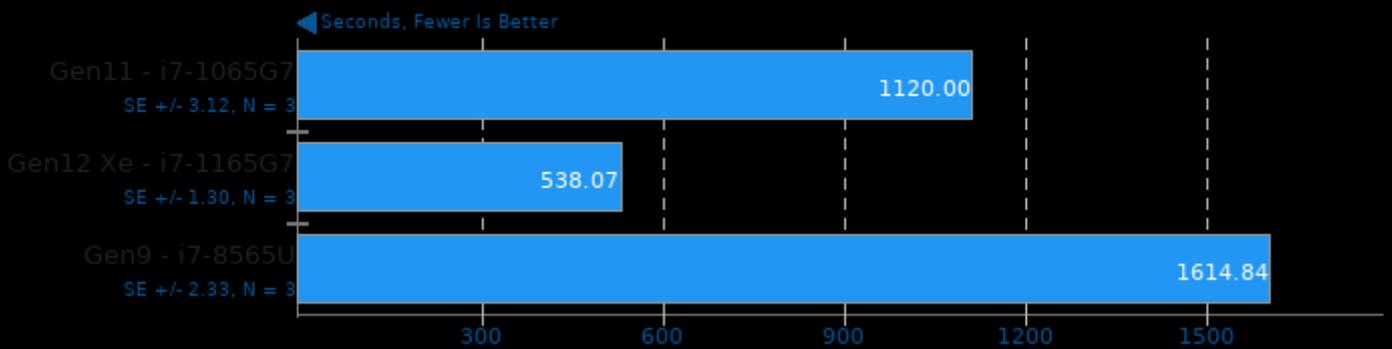
PlaidML

FP16: No - Mode: Inference - Network: NASNet Large - Device: OpenCL



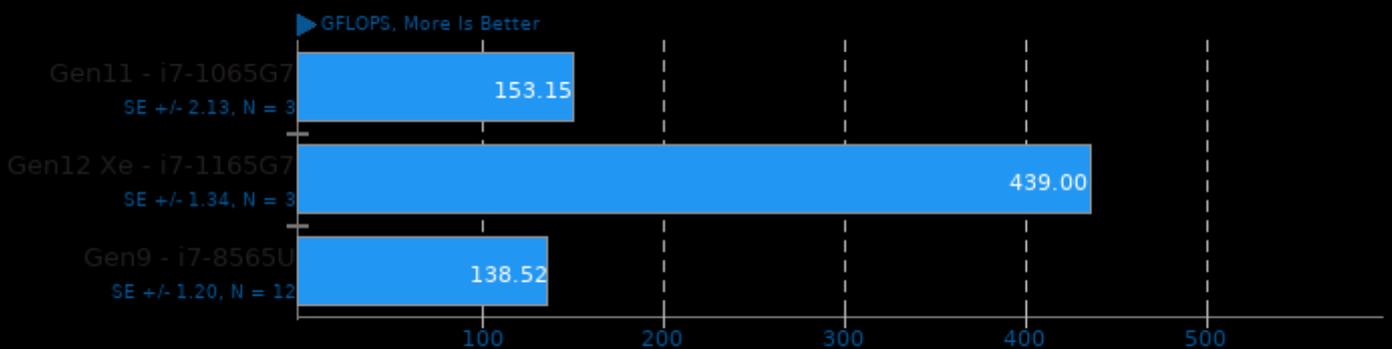
RealSR-NCNN 20200818

Scale: 4x - TAA: Yes



oneAPI Level Zero Tests

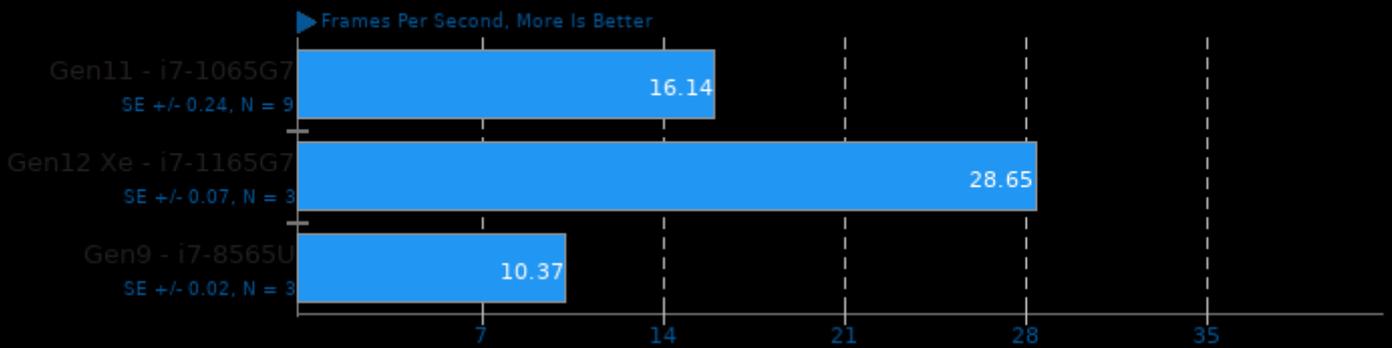
Test: Peak Integer Compute



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

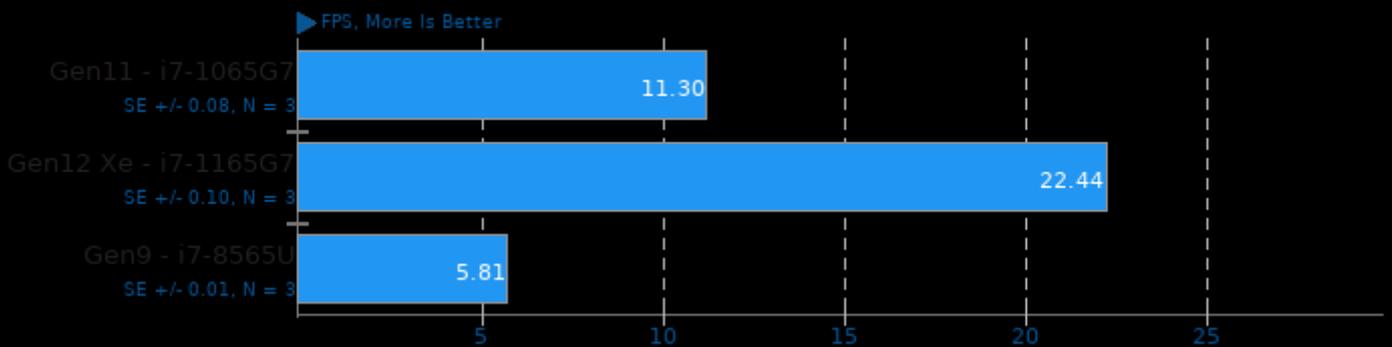
Unigine Heaven 4.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Renderer: OpenGL



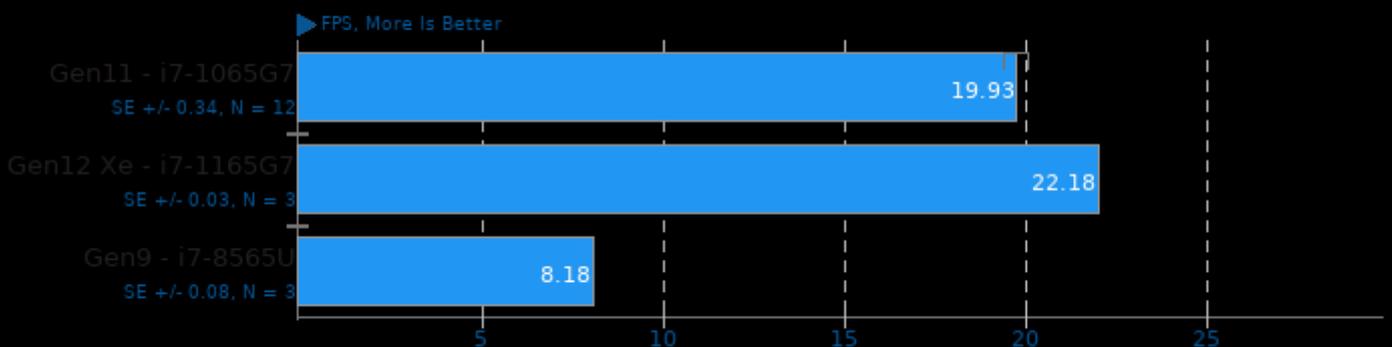
PlaidML

FP16: No - Mode: Inference - Network: DenseNet 201 - Device: OpenCL



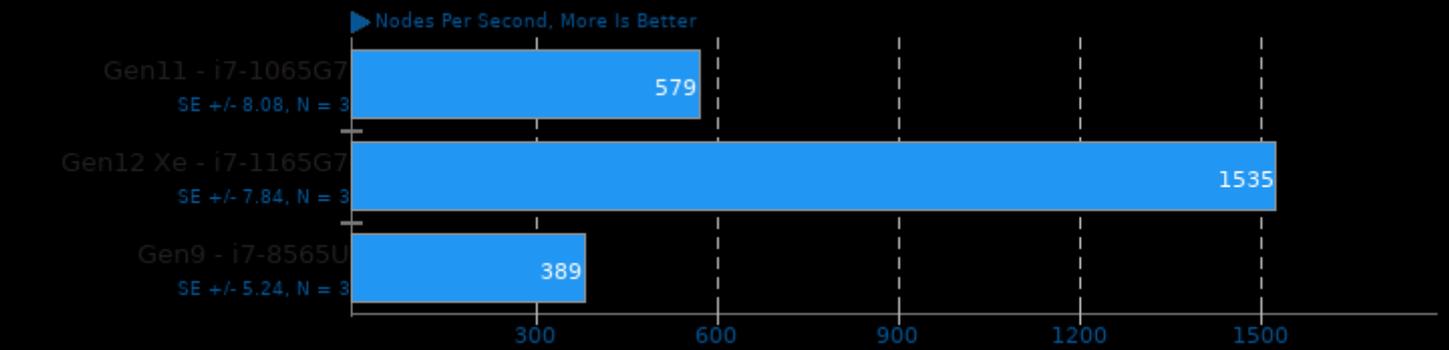
PlaidML

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



LeelaChessZero 0.26

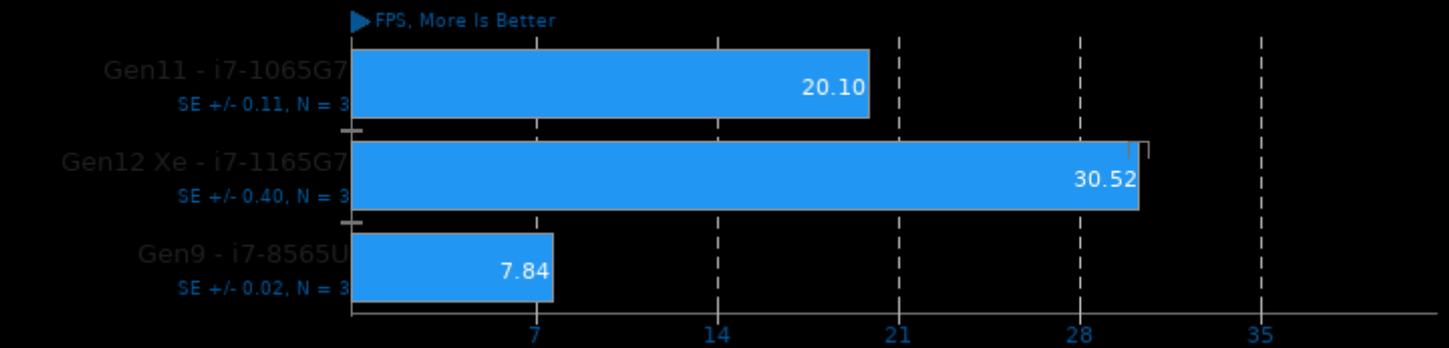
Backend: OpenCL



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

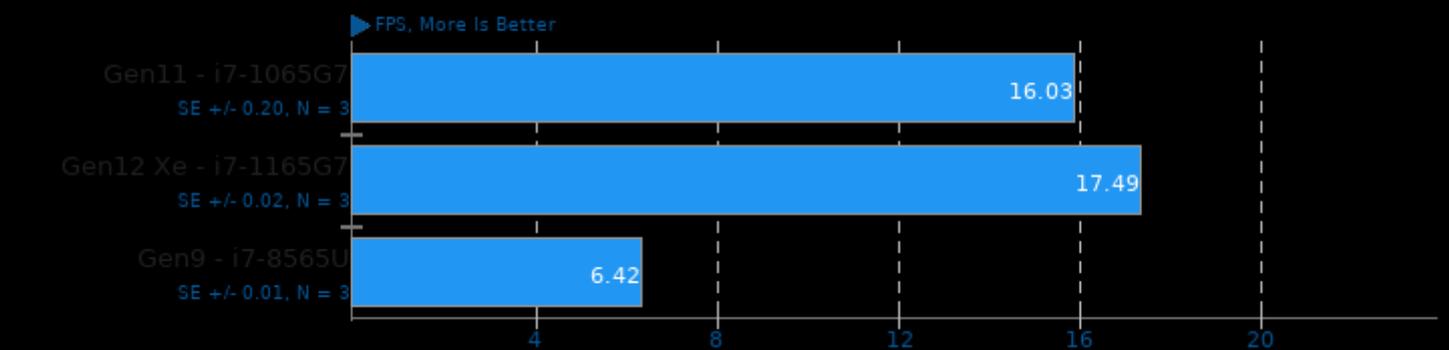
PlaidML

FP16: No - Mode: Inference - Network: IMDB LSTM - Device: OpenCL



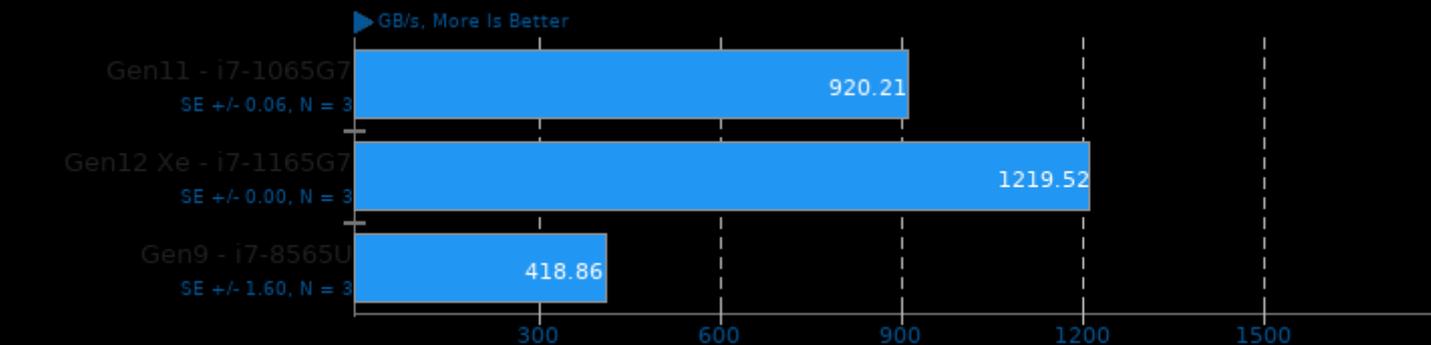
PlaidML

FP16: No - Mode: Inference - Network: VGG19 - Device: OpenCL



oneAPI Level Zero Tests

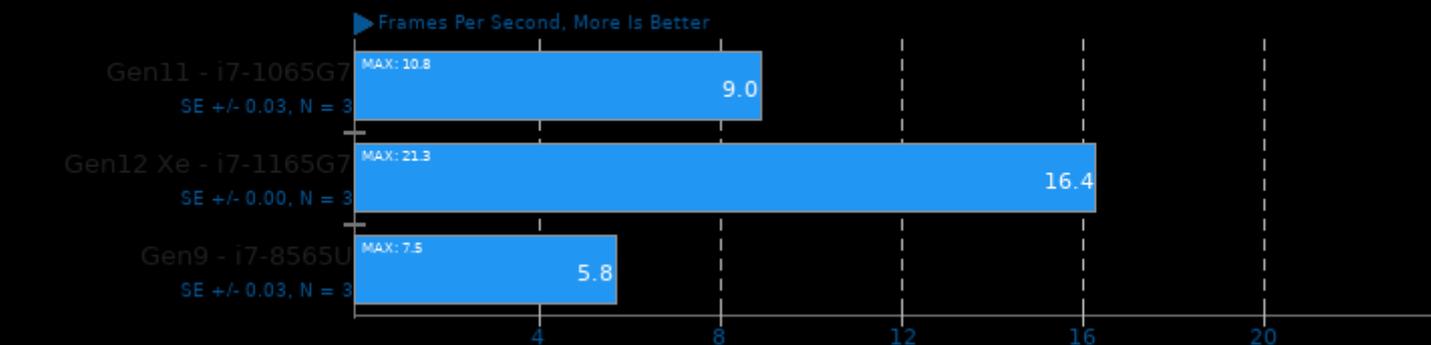
Test: Peak Single-Precision Compute



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

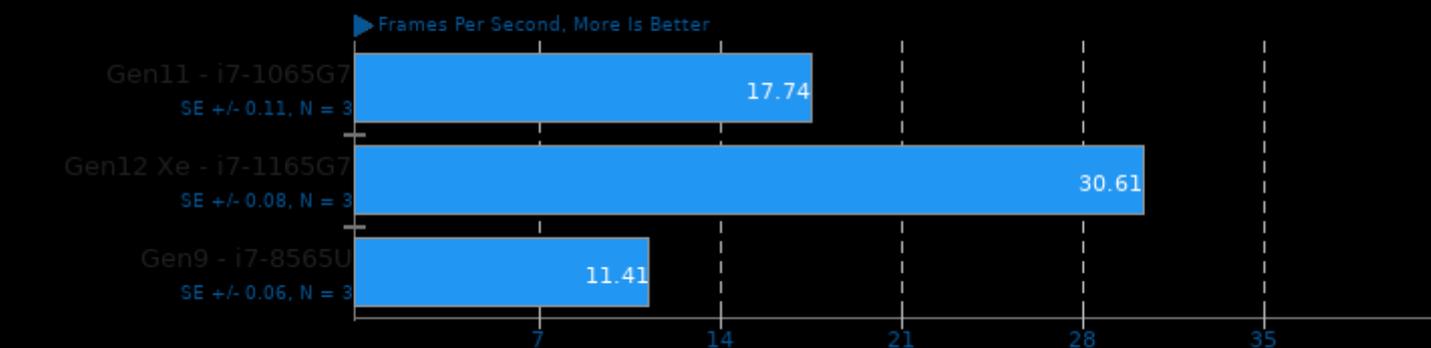
Unigine Superposition 1.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Quality: Medium - Renderer: OpenGL



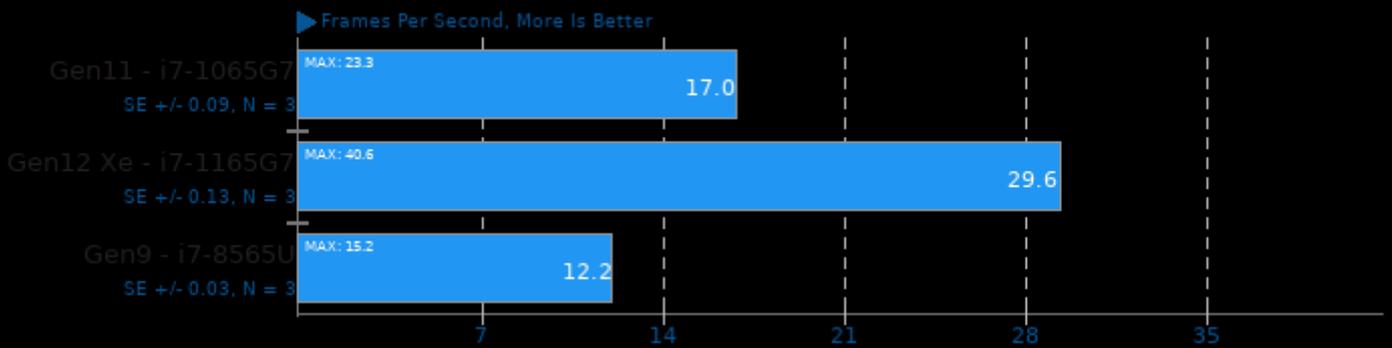
Unigine Valley 1.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Renderer: OpenGL



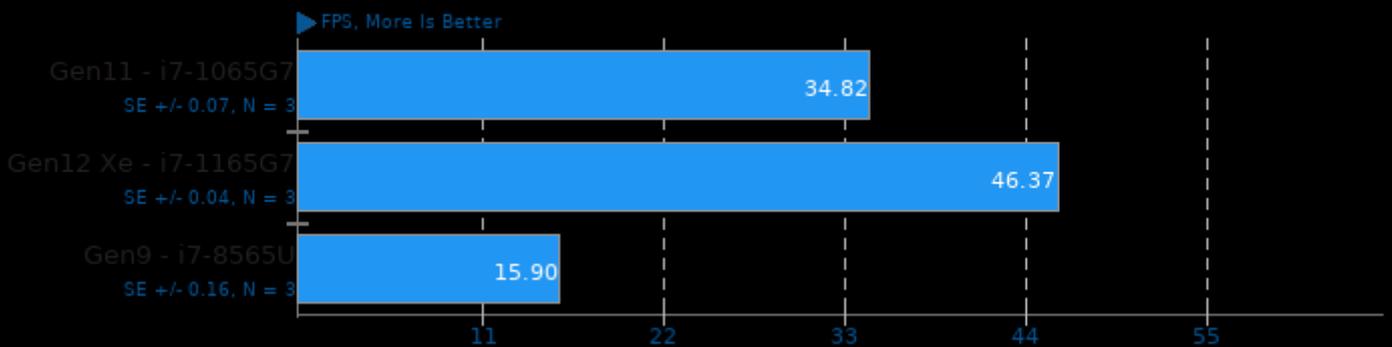
Unigine Superposition 1.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Quality: Low - Renderer: OpenGL



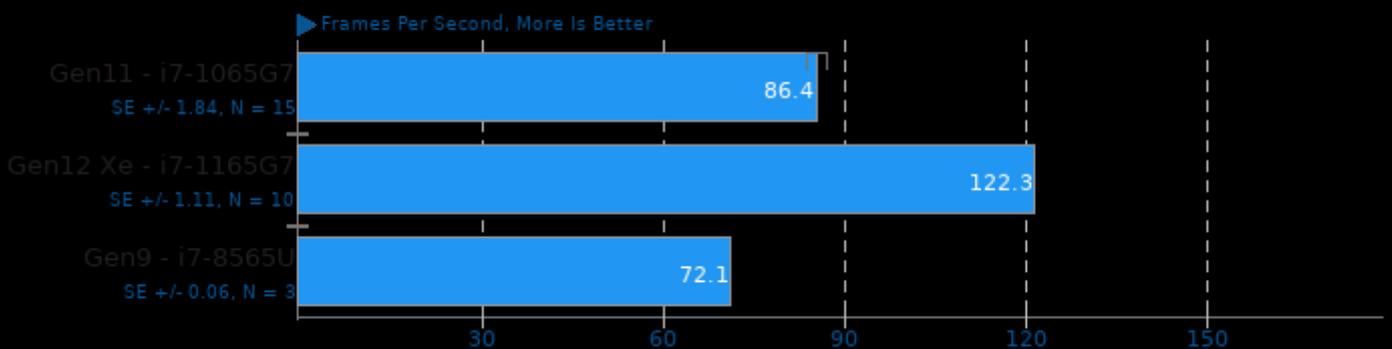
PlaidML

FP16: No - Mode: Inference - Network: Inception V3 - Device: OpenCL



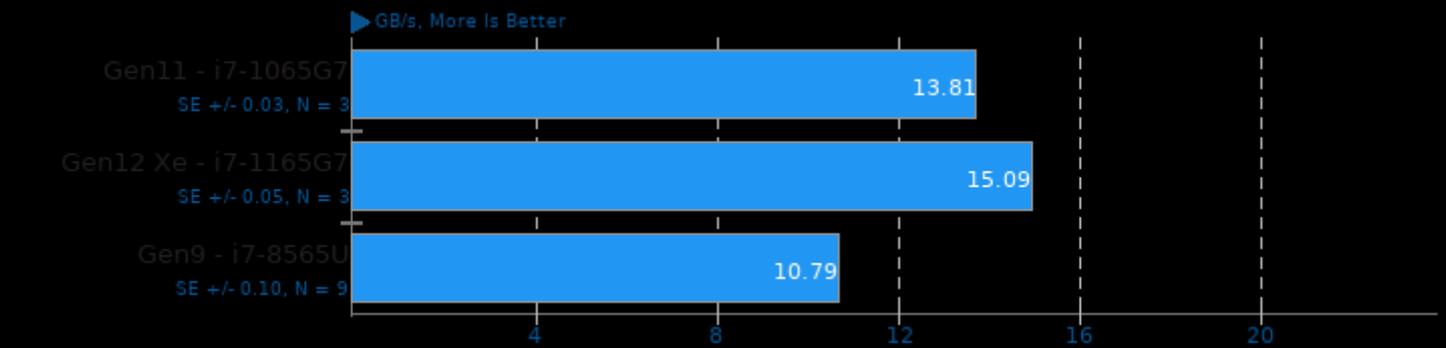
ET: Legacy 2.75

Renderer: Renderer2 - Resolution: 1920 x 1200



oneAPI Level Zero Tests

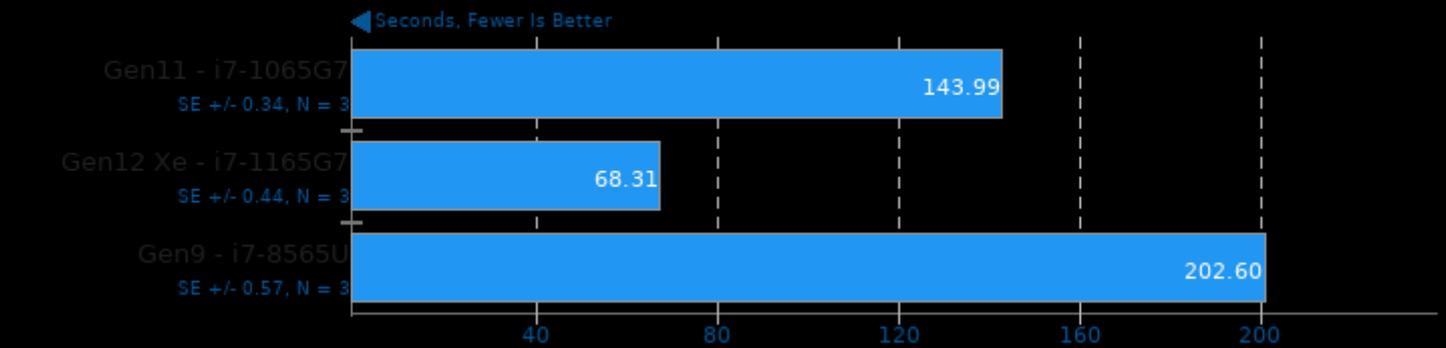
Test: Peak System Memory Copy to Shared Memory



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

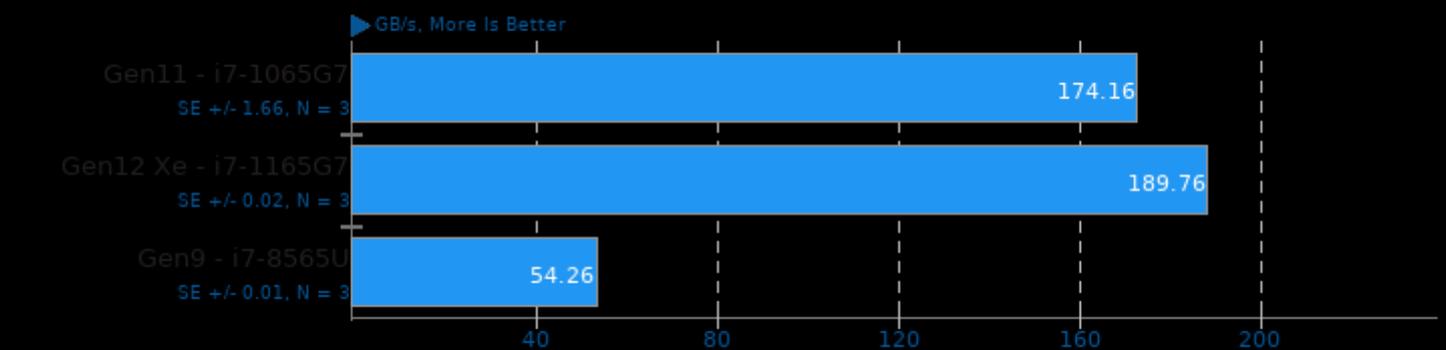
RealSR-NCNN 20200818

Scale: 4x - TAA: No



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Texture Read Bandwidth

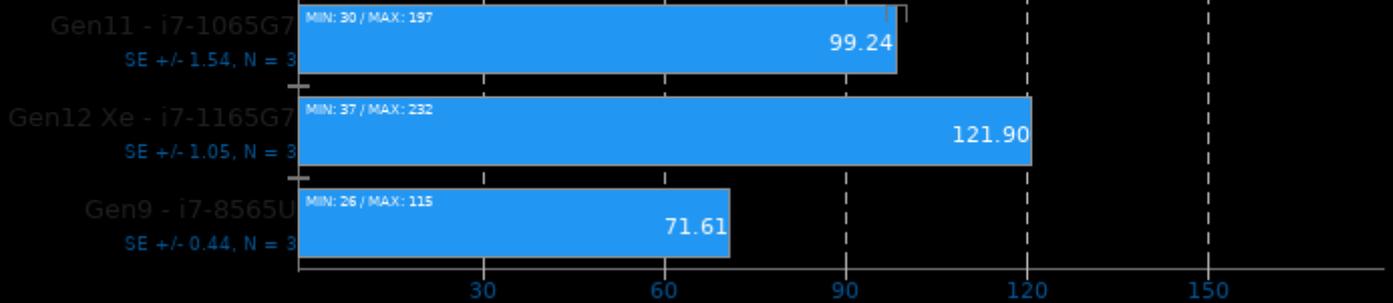


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Ultimate

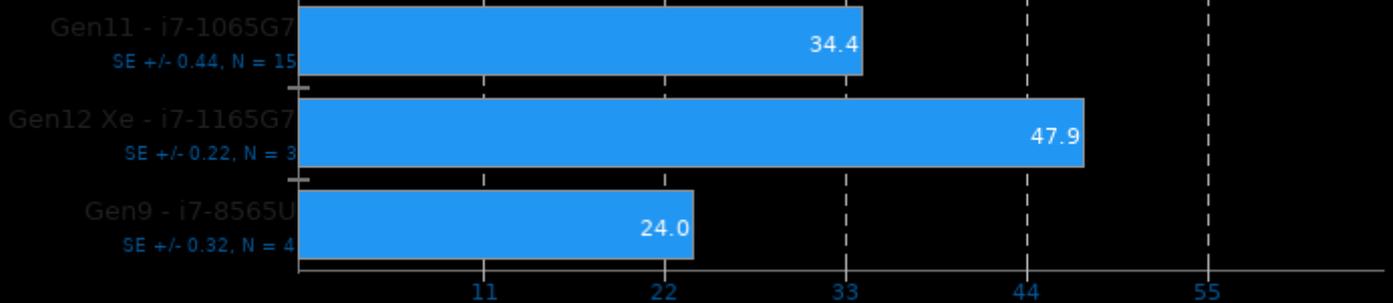
► Frames Per Second, More Is Better



cl-mem 2017-01-13

Benchmark: Write

► GB/s, More Is Better

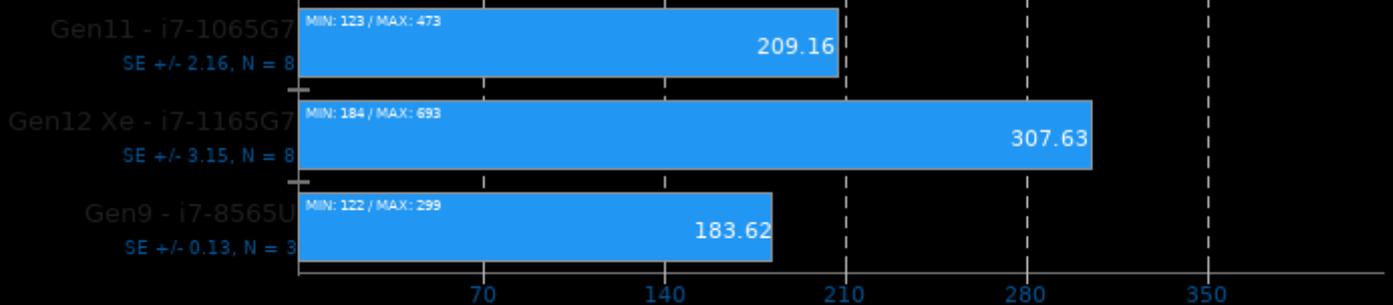


1. (CC) gcc options: -O2 -fno -fOpenCL

Xonotic 0.8.2

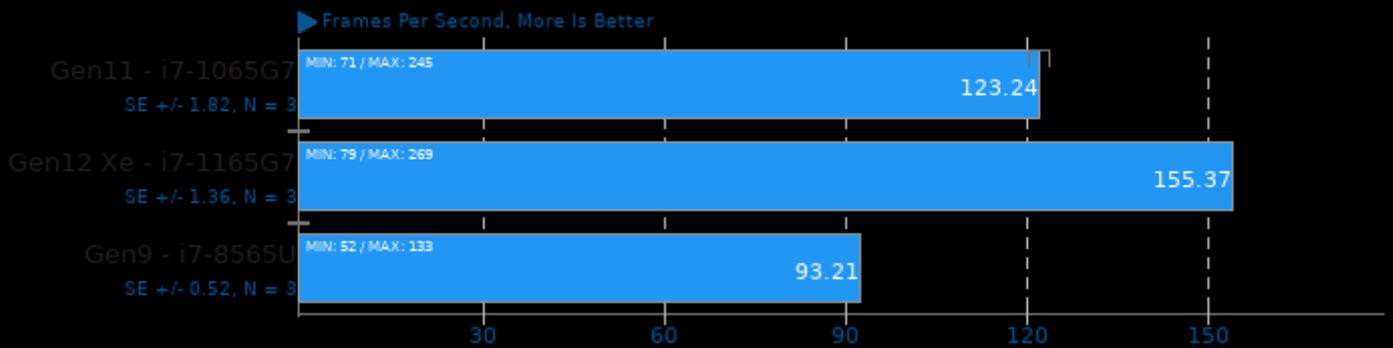
Resolution: 1920 x 1200 - Effects Quality: Low

► Frames Per Second, More Is Better



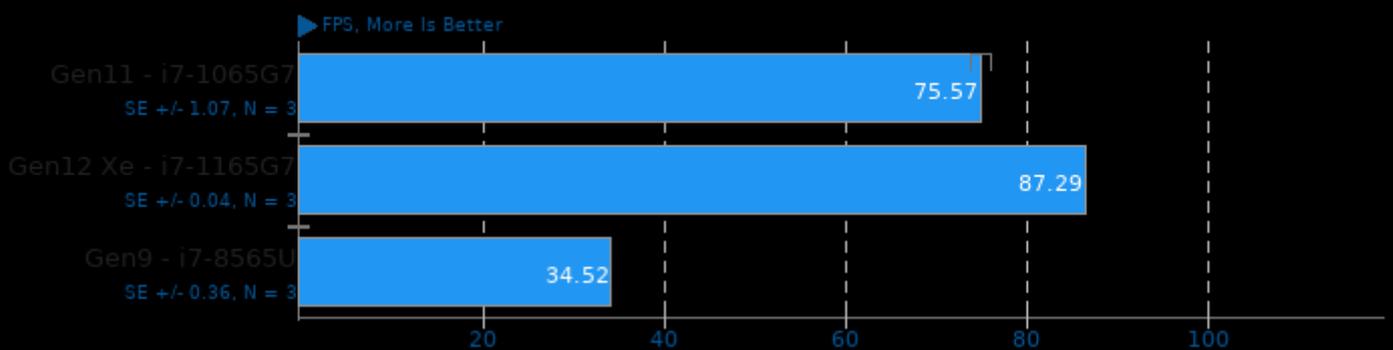
Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Ultra



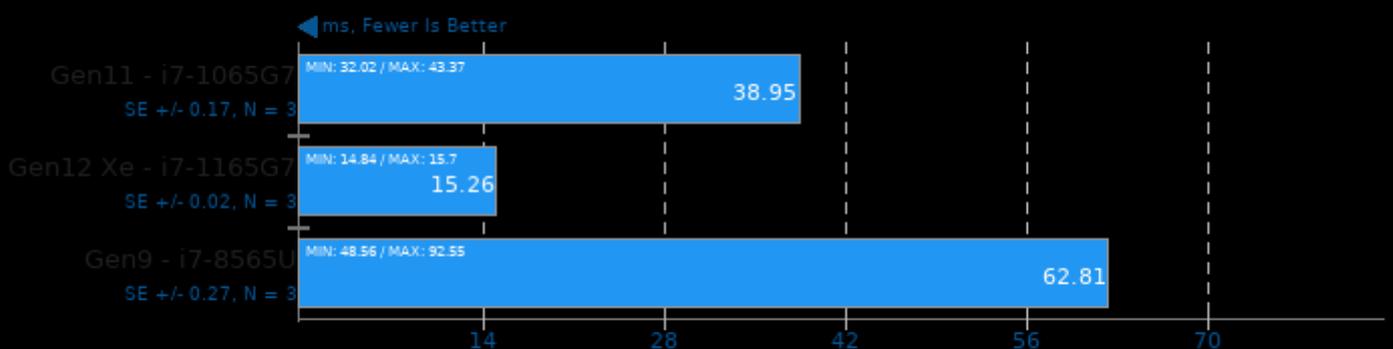
PlaidML

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



NCNN 20200916

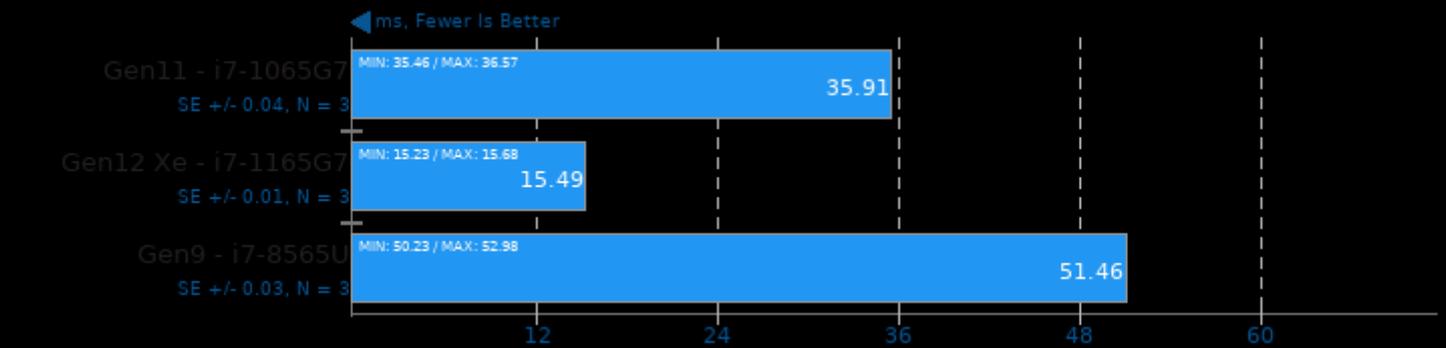
Target: Vulkan GPU - Model: yolov4-tiny



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

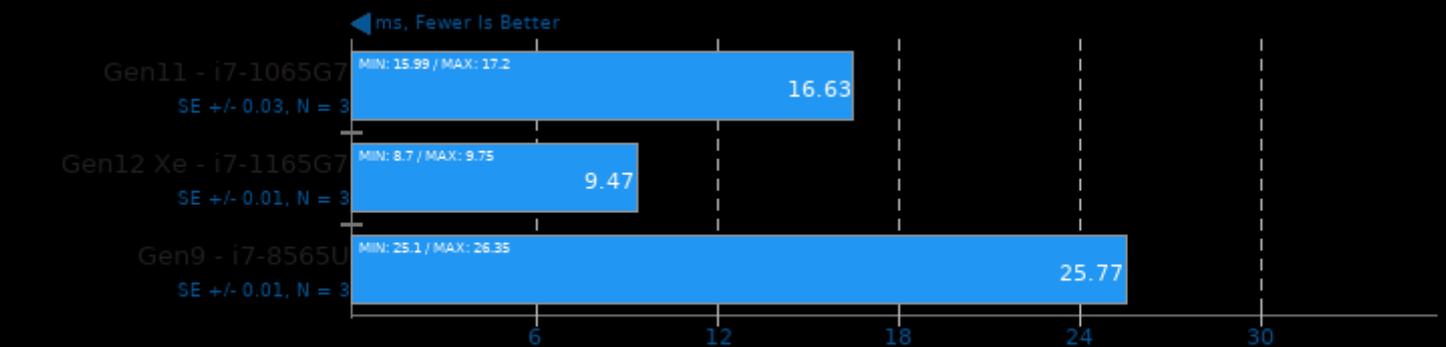
NCNN 20200916

Target: Vulkan GPU - Model: resnet50



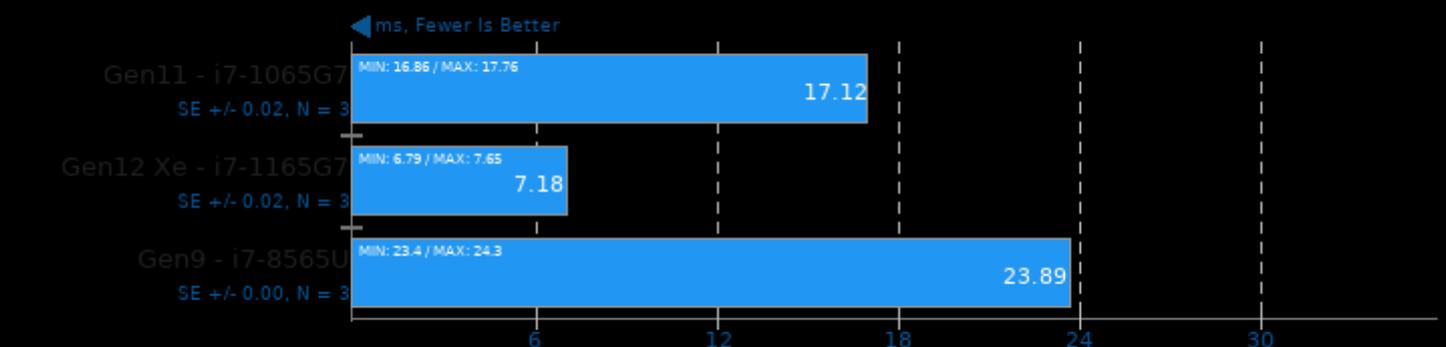
NCNN 20200916

Target: Vulkan GPU - Model: alexnet



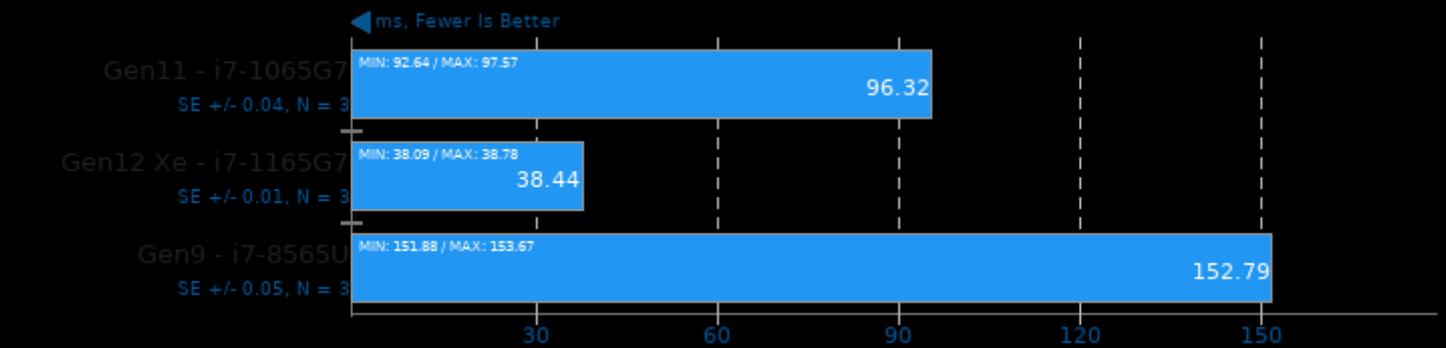
NCNN 20200916

Target: Vulkan GPU - Model: resnet18



NCNN 20200916

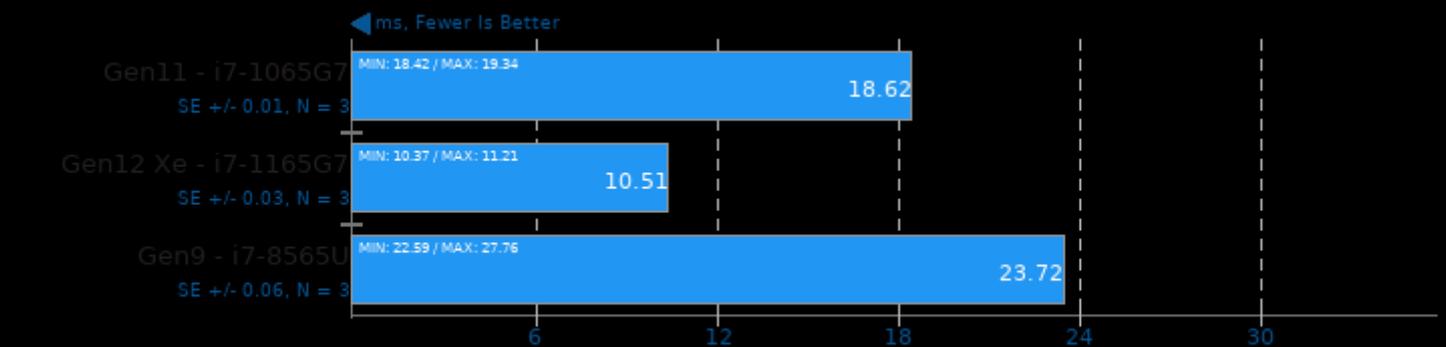
Target: Vulkan GPU - Model: vgg16



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

NCNN 20200916

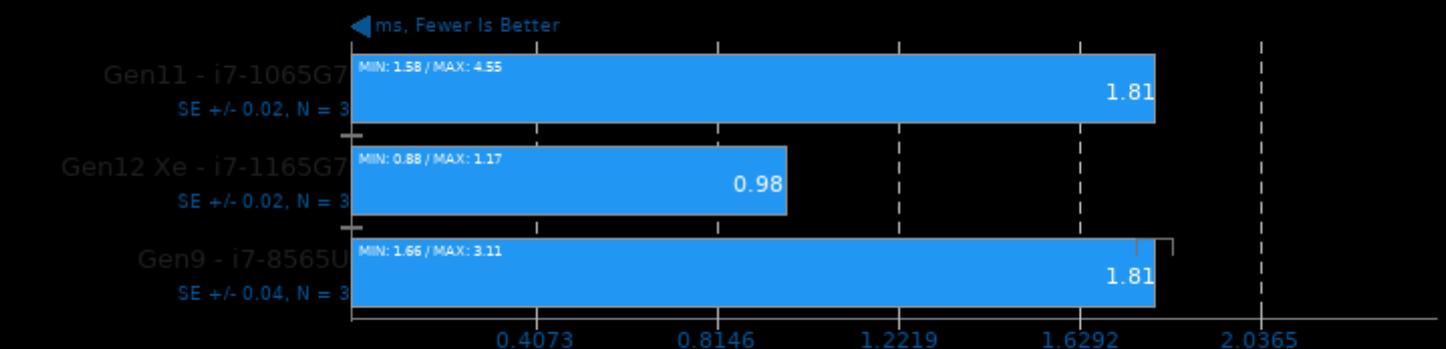
Target: Vulkan GPU - Model: googlenet



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

NCNN 20200916

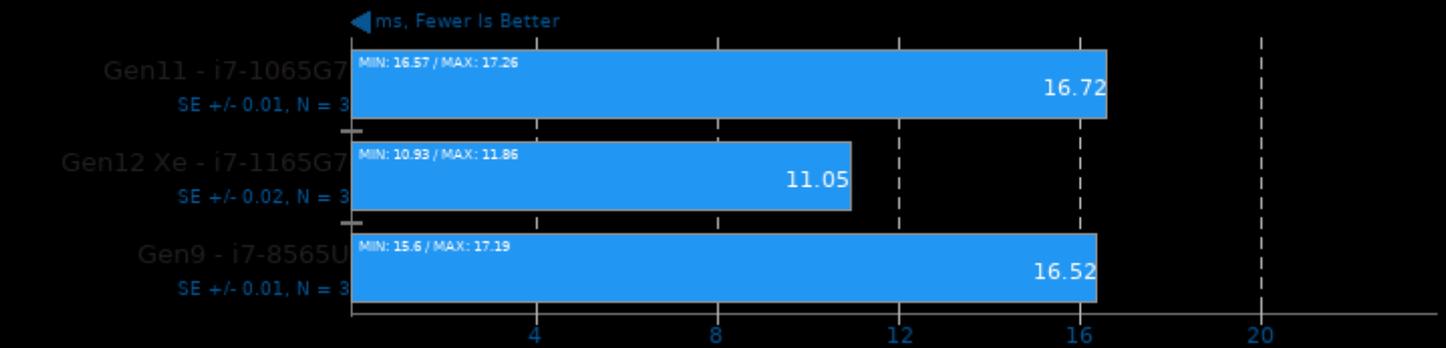
Target: Vulkan GPU - Model: blazeface



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

NCNN 20200916

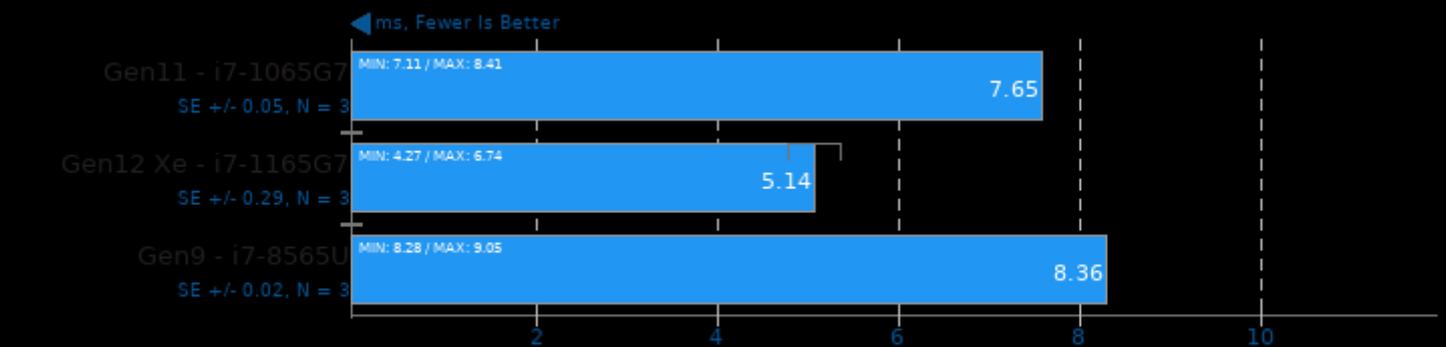
Target: Vulkan GPU - Model: efficientnet-b0



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

NCNN 20200916

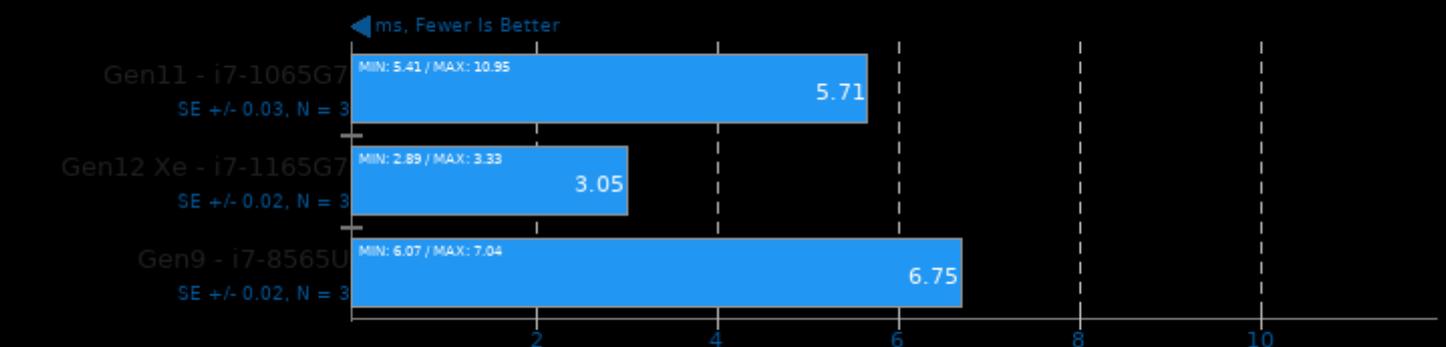
Target: Vulkan GPU - Model: mnasnet



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

NCNN 20200916

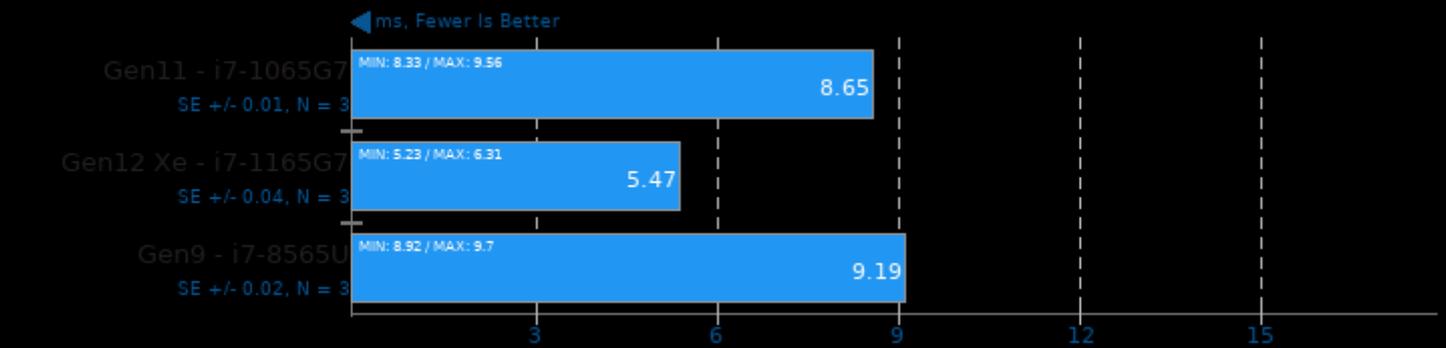
Target: Vulkan GPU - Model: shufflenet-v2



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

NCNN 20200916

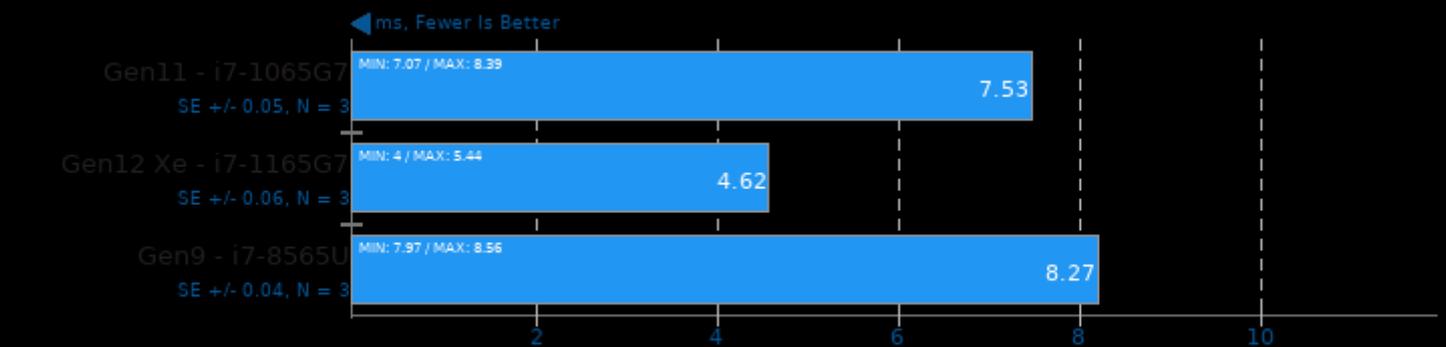
Target: Vulkan GPU-v3-v3 - Model: mobilenet-v3



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

NCNN 20200916

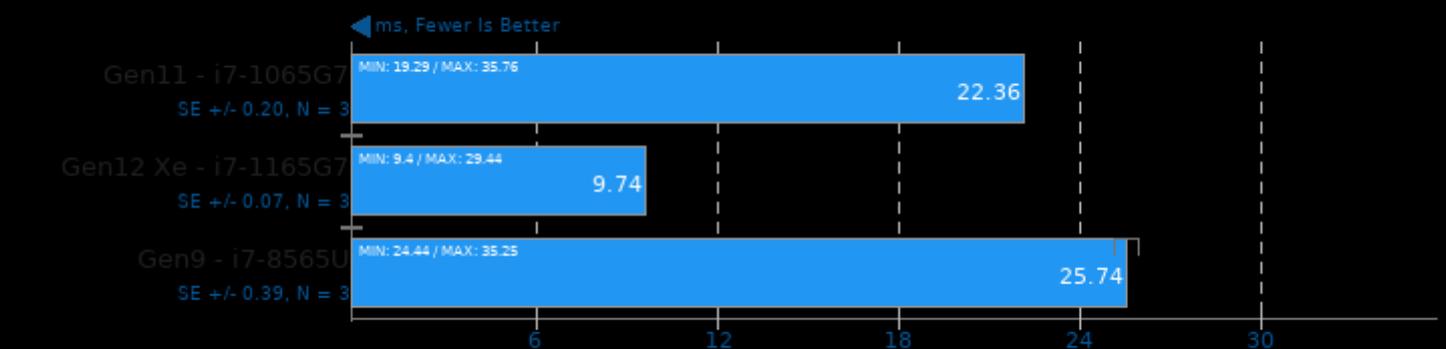
Target: Vulkan GPU-v2-v2 - Model: mobilenet-v2



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

NCNN 20200916

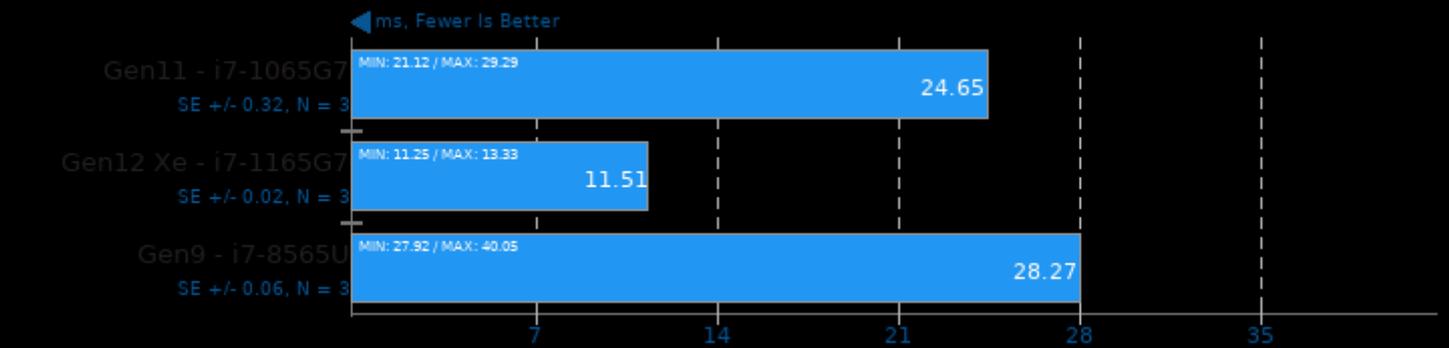
Target: Vulkan GPU - Model: mobilenet



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

NCNN 20200916

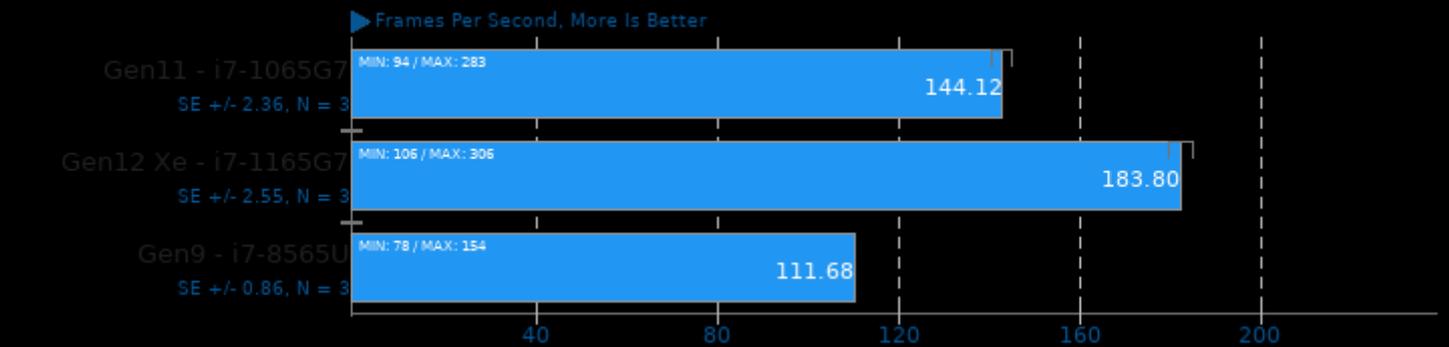
Target: Vulkan GPU - Model: squeezenet



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

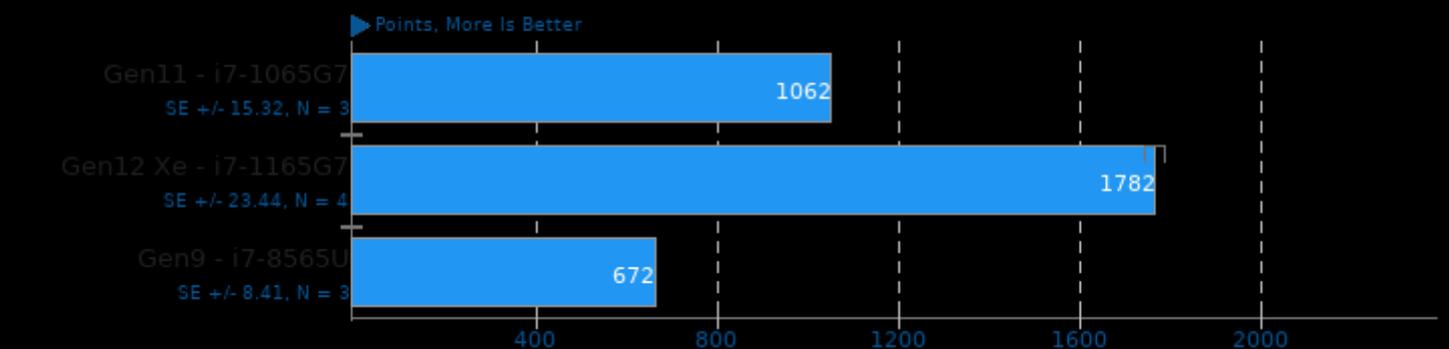
Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: High



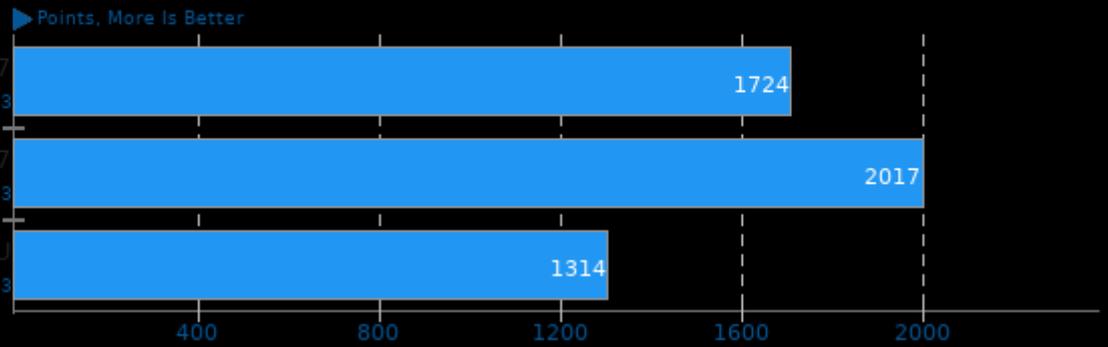
GpuTest 0.7.0

Test: Furmark - Resolution: 1920 x 1200 - Mode: Fullscreen



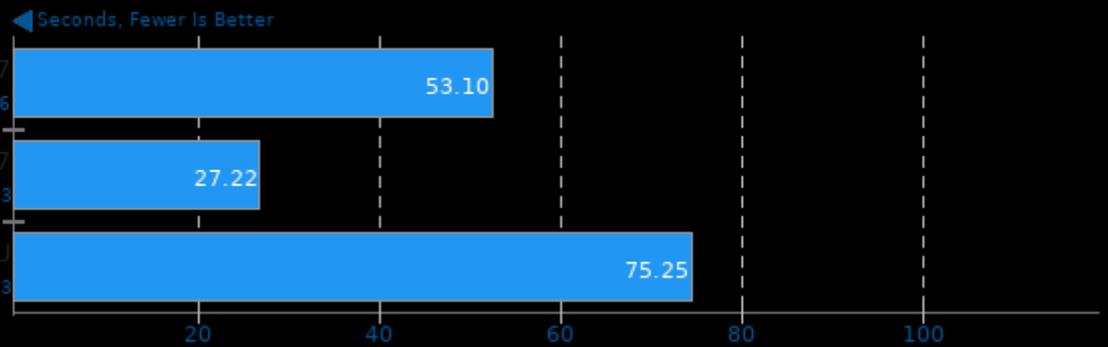
GpuTest 0.7.0

Test: GiMark - Resolution: 1920 x 1200 - Mode: Fullscreen



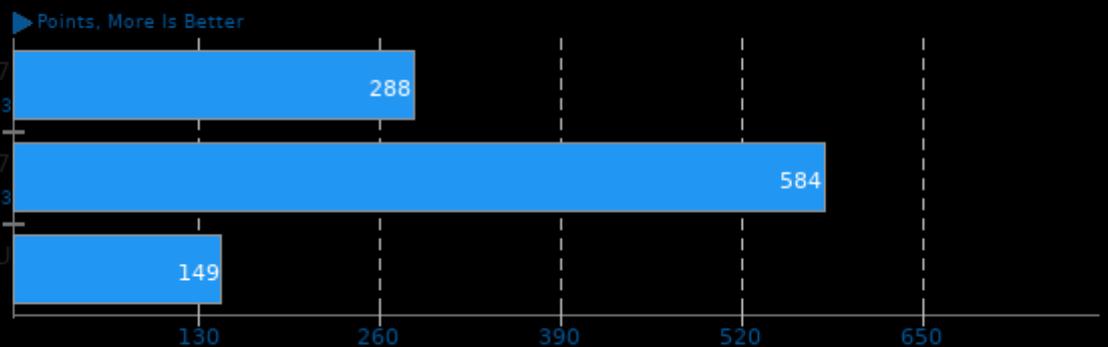
Waifu2x-NCNN Vulkan 20200818

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



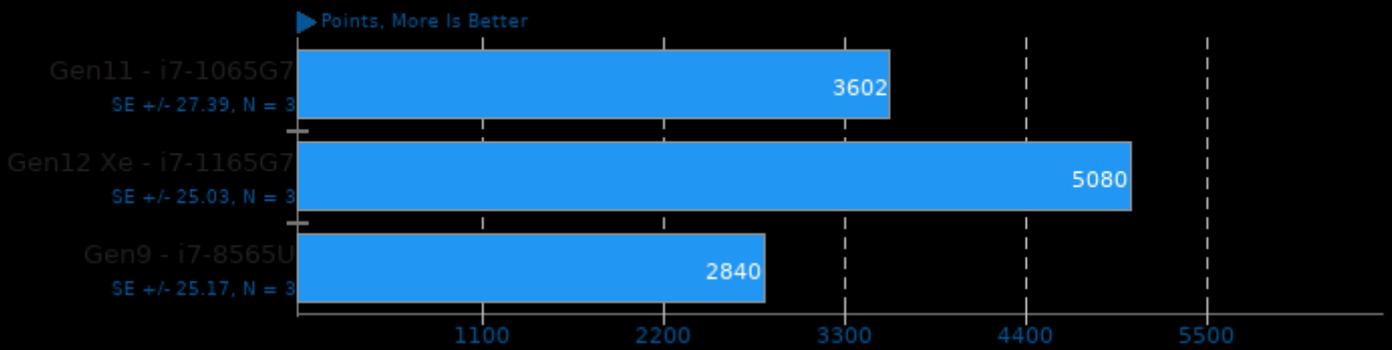
GpuTest 0.7.0

Test: Pixmark Piano - Resolution: 1920 x 1200 - Mode: Fullscreen



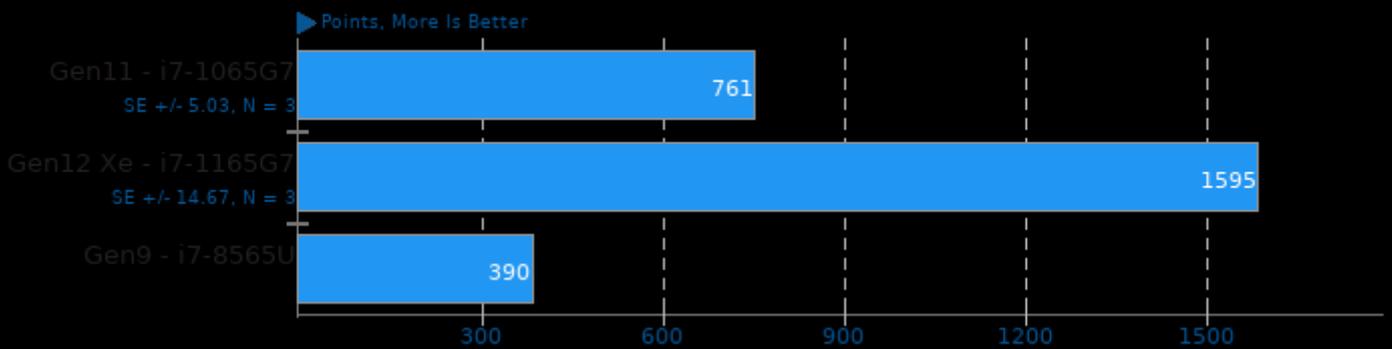
GpuTest 0.7.0

Test: TessMark - Resolution: 1920 x 1200 - Mode: Fullscreen



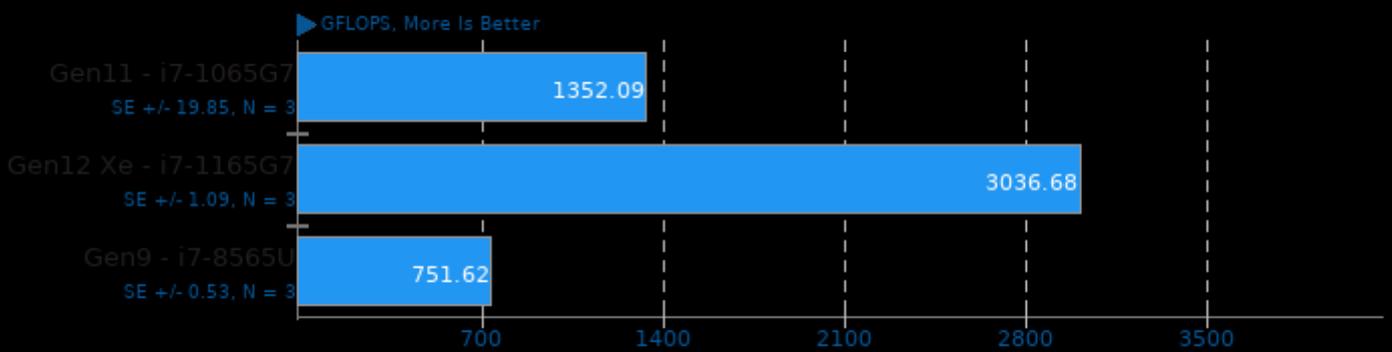
GpuTest 0.7.0

Test: Pixmark Volplosion - Resolution: 1920 x 1200 - Mode: Fullscreen



oneAPI Level Zero Tests

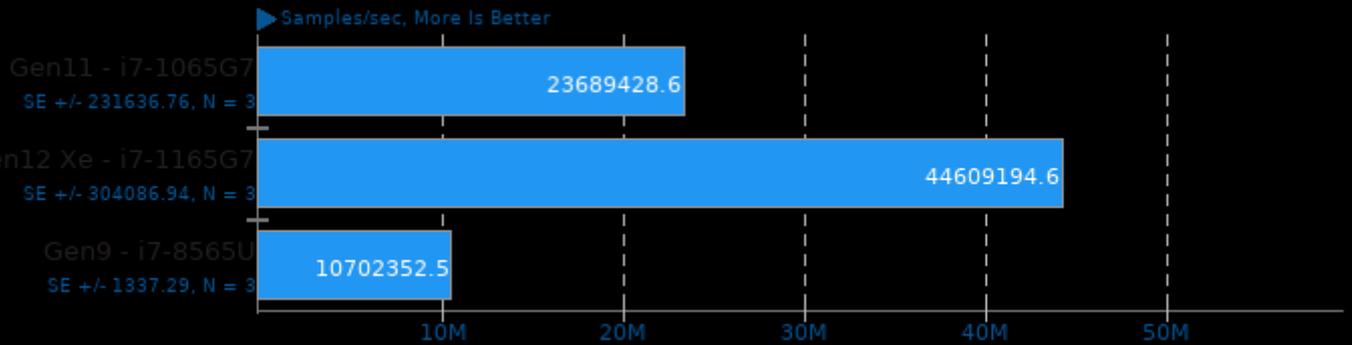
Test: Peak Half-Precision Compute



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

MandelGPU 1.3pts1

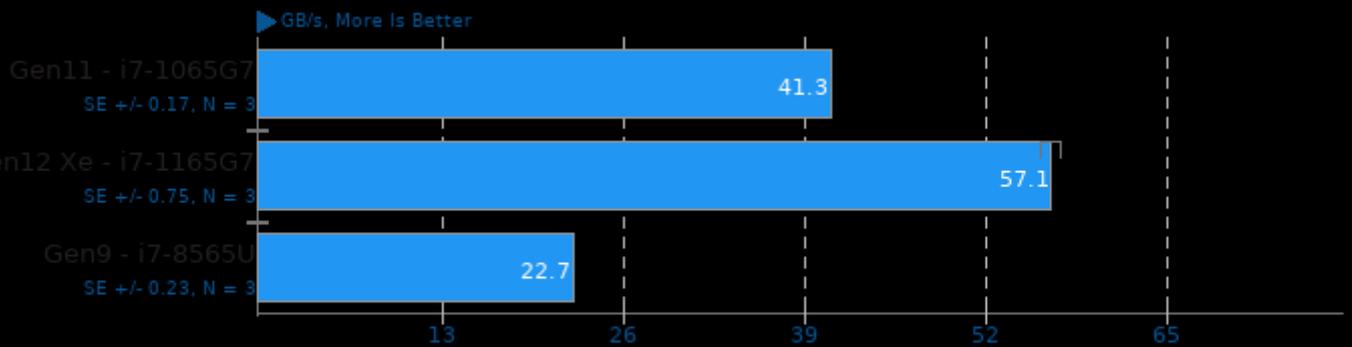
OpenCL Device: GPU



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

cl-mem 2017-01-13

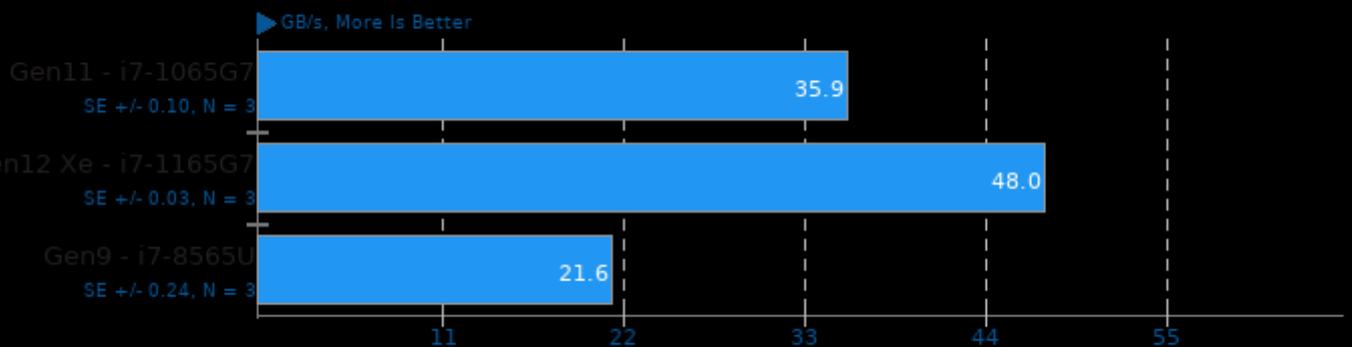
Benchmark: Read



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

cl-mem 2017-01-13

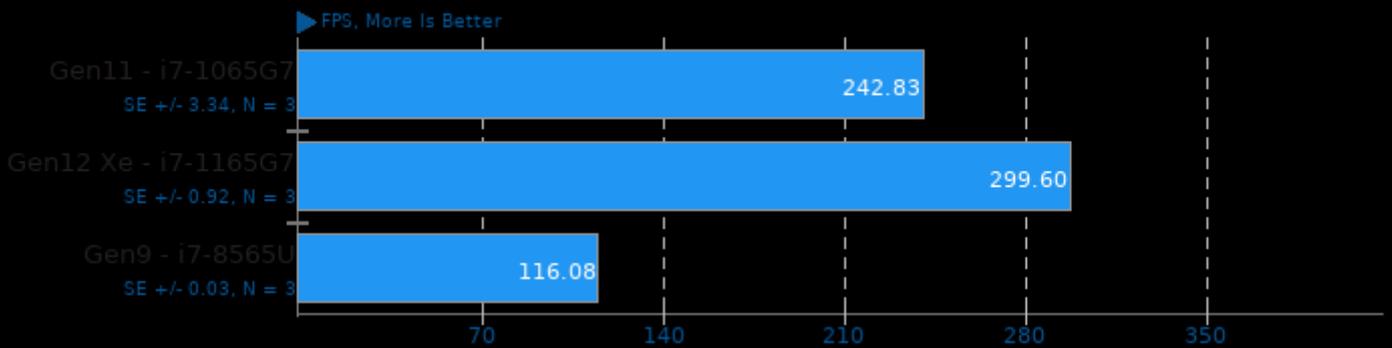
Benchmark: Copy



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

PlaidML

FP16: No - Mode: Inference - Network: Mobilenet - Device: OpenCL



SHOC Scalable Heterogeneous Computing 2015-11-10

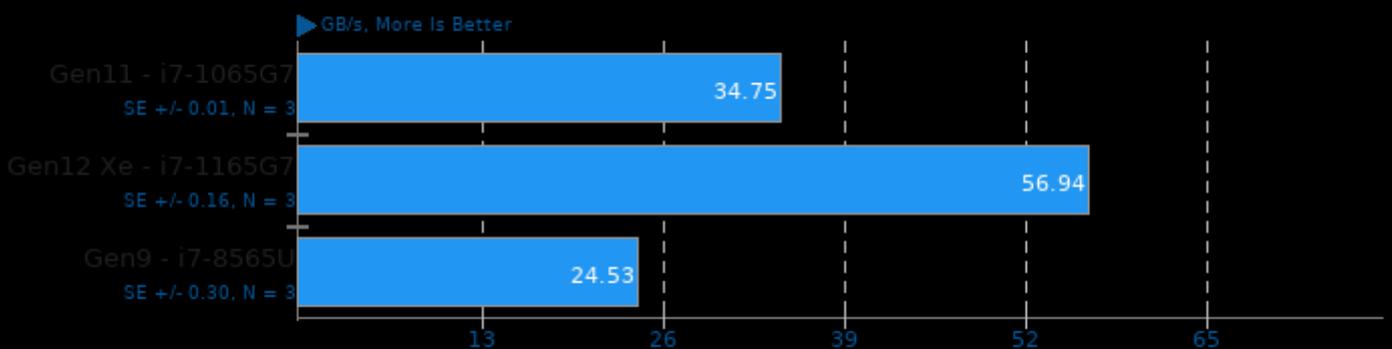
Target: OpenCL - Benchmark: Max SP Flops



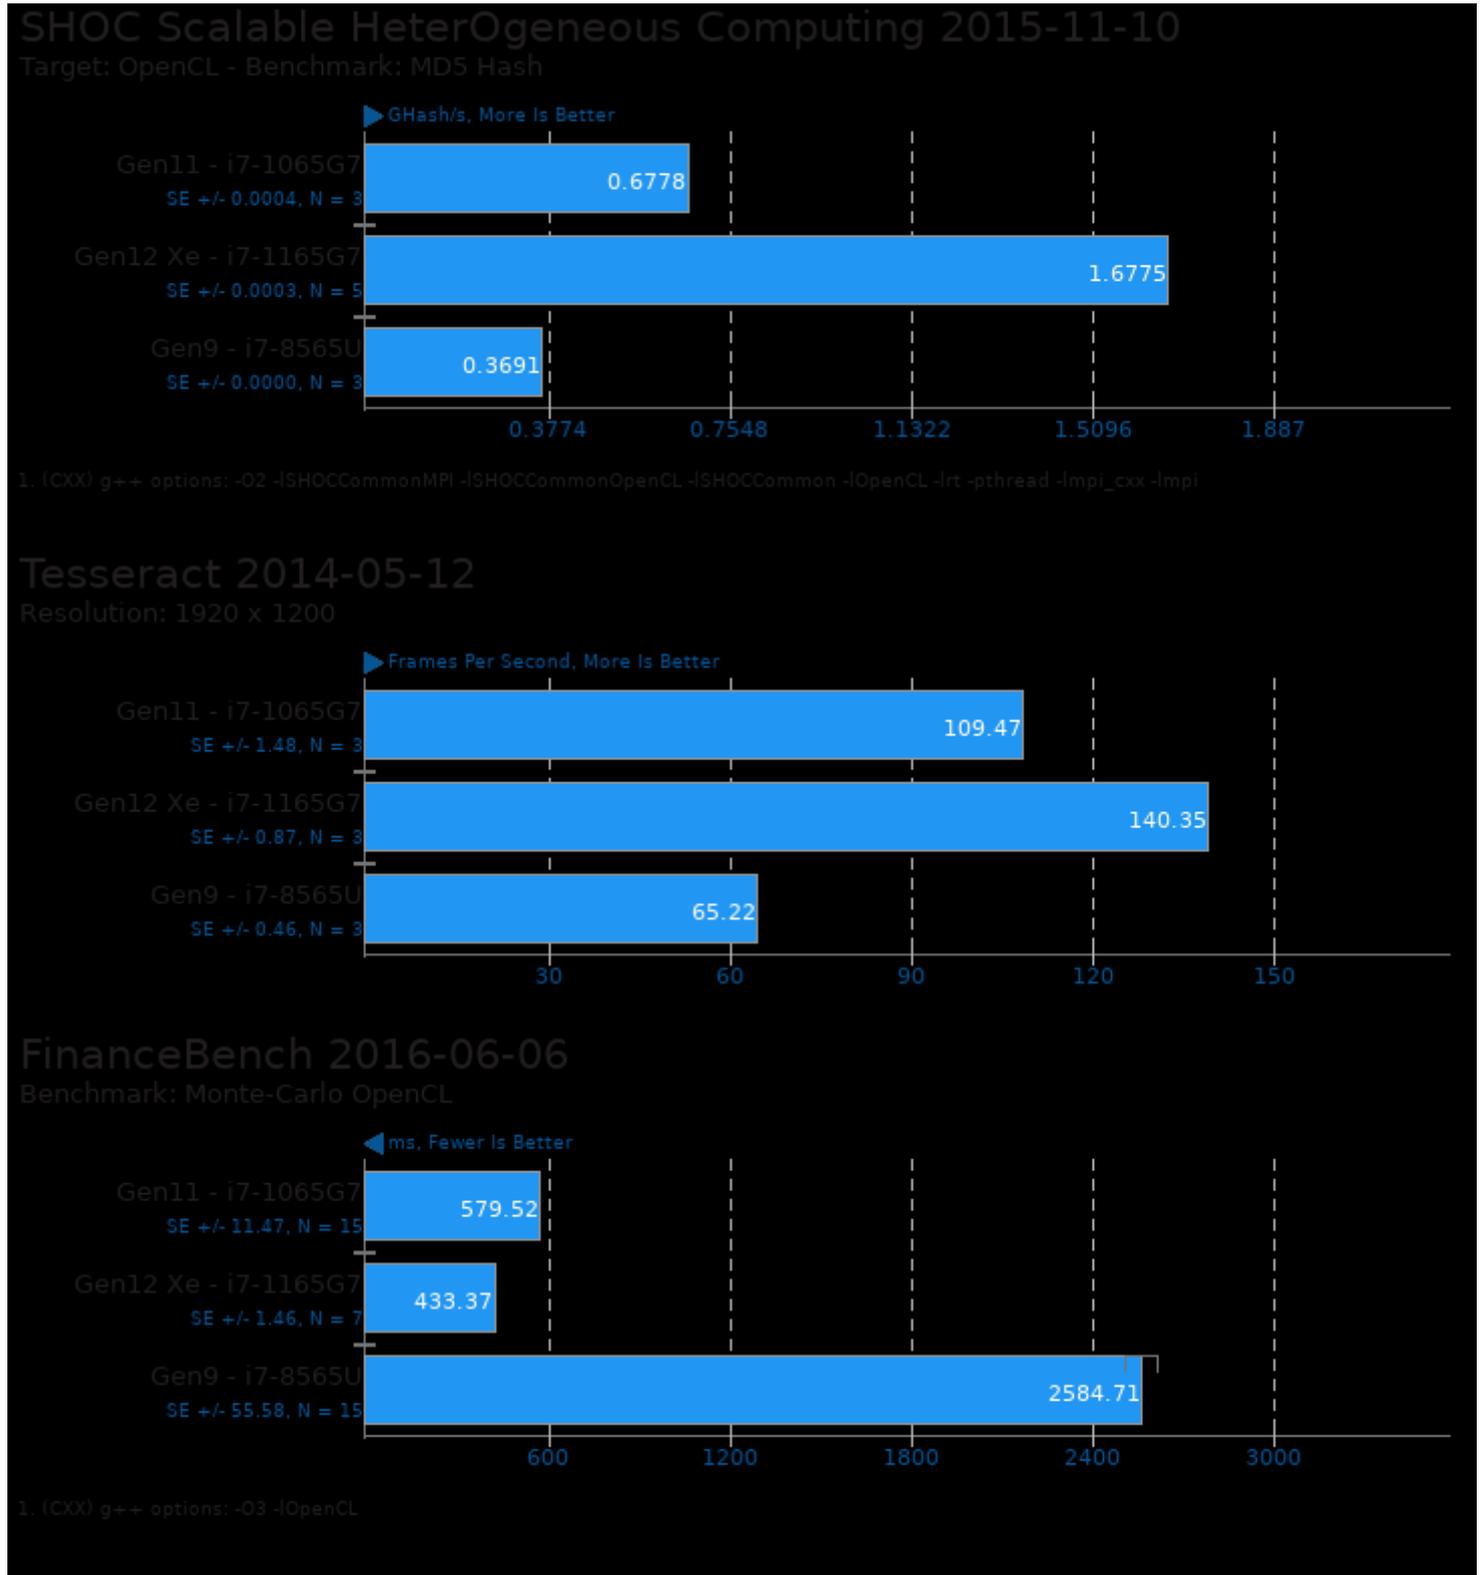
1. (CXX) g++ options: -O2 -ISHOCCommonMPI -ISHOCCommonOpenCL -ISHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

oneAPI Level Zero Tests

Test: Peak Float16 Global Memory Bandwidth

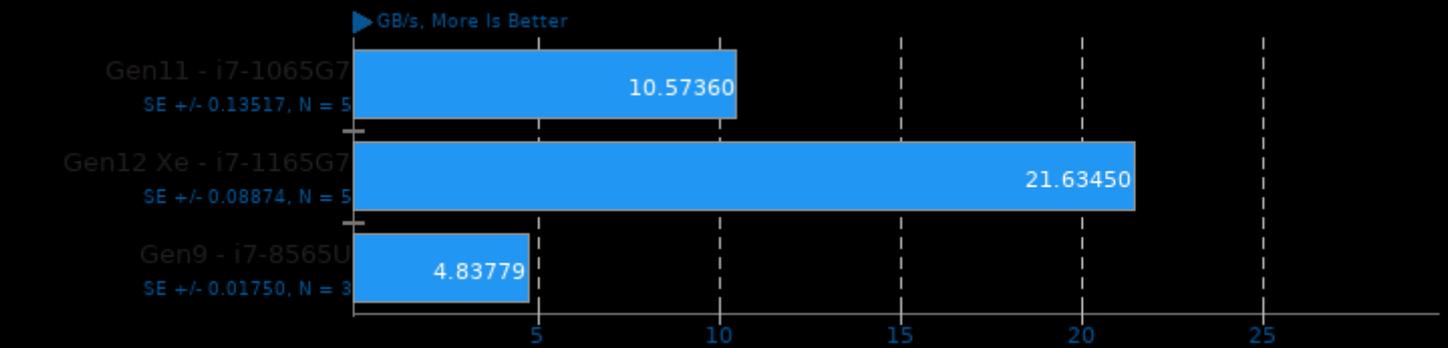


1. (CXX) g++ options: -ldl -pthread



oneAPI Level Zero Tests

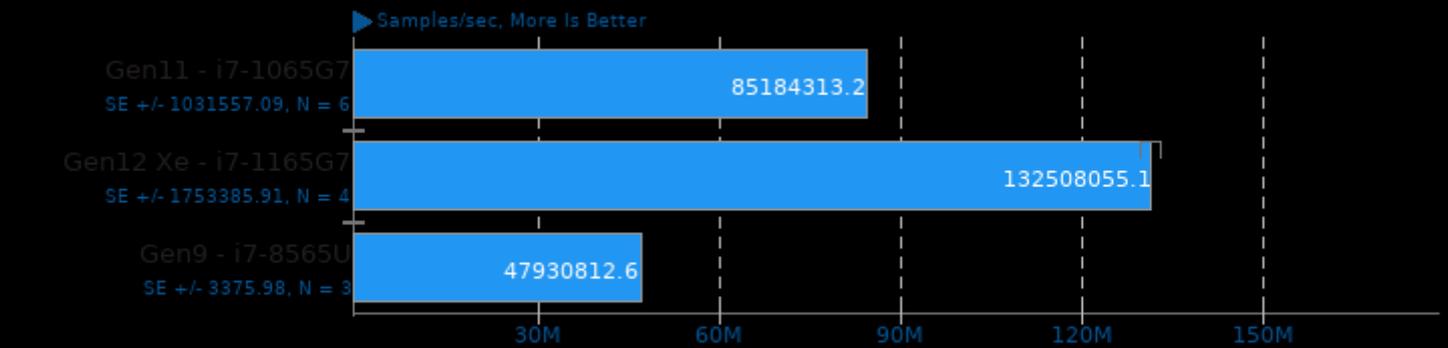
Test: Host-To-Device-To-Host Image Copy



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

JuliaGPU 1.2pts1

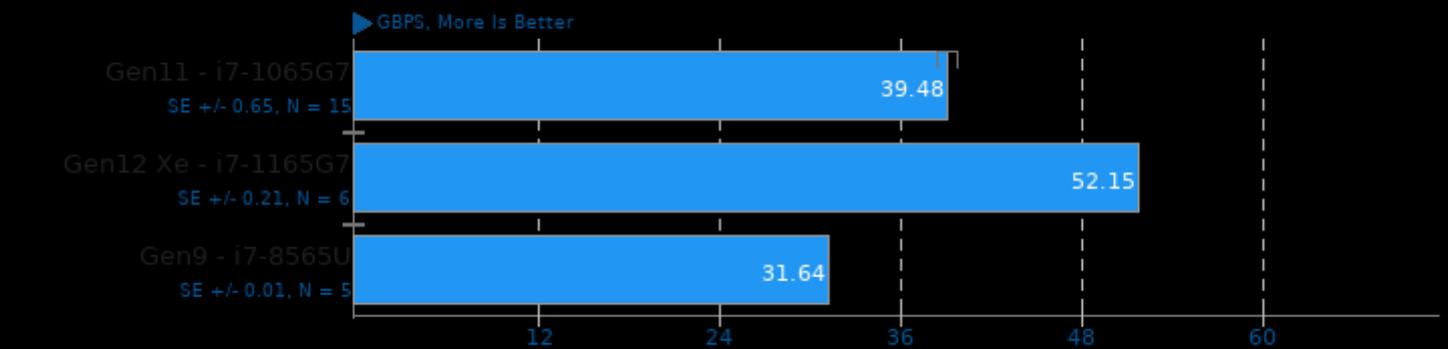
OpenCL Device: GPU



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

clpeak

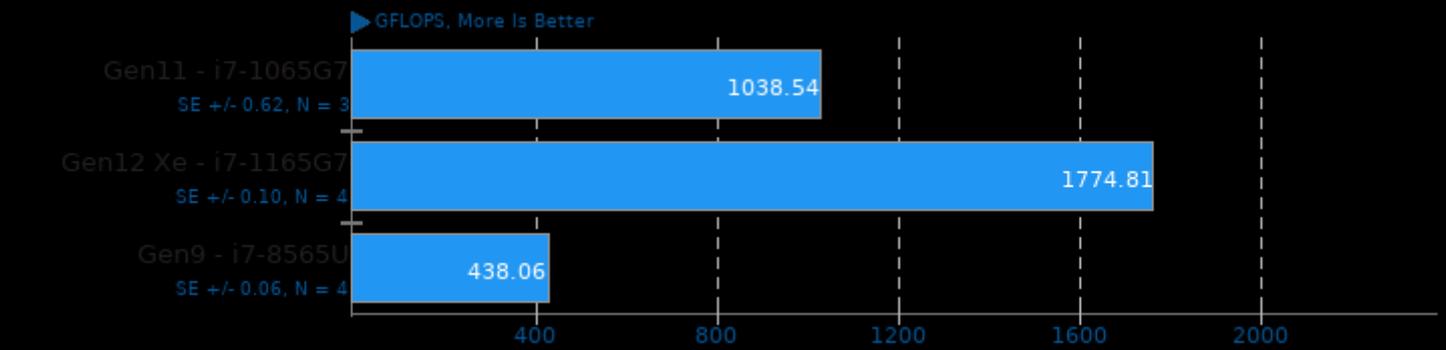
OpenCL Test: Transfer Bandwidth enqueueWriteBuffer



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

clpeak

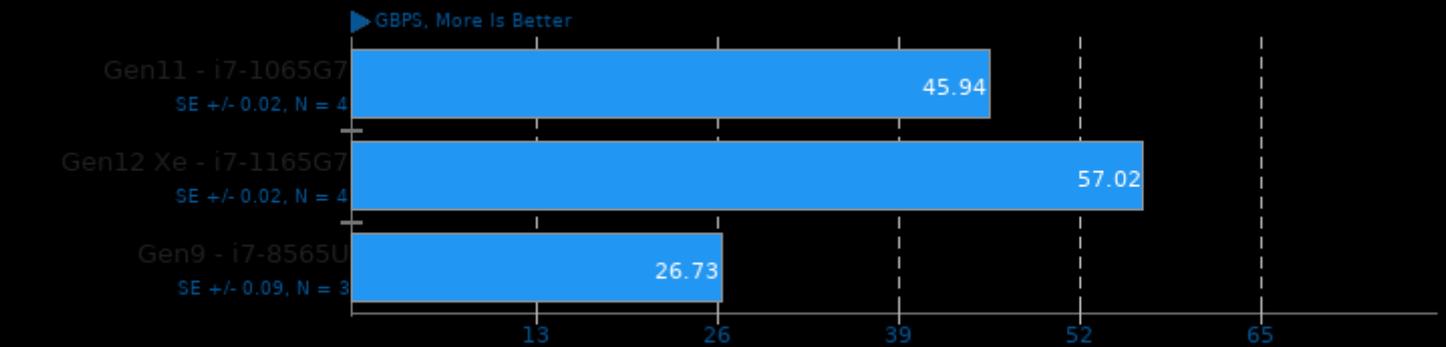
OpenCL Test: Single-Precision Float



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

clpeak

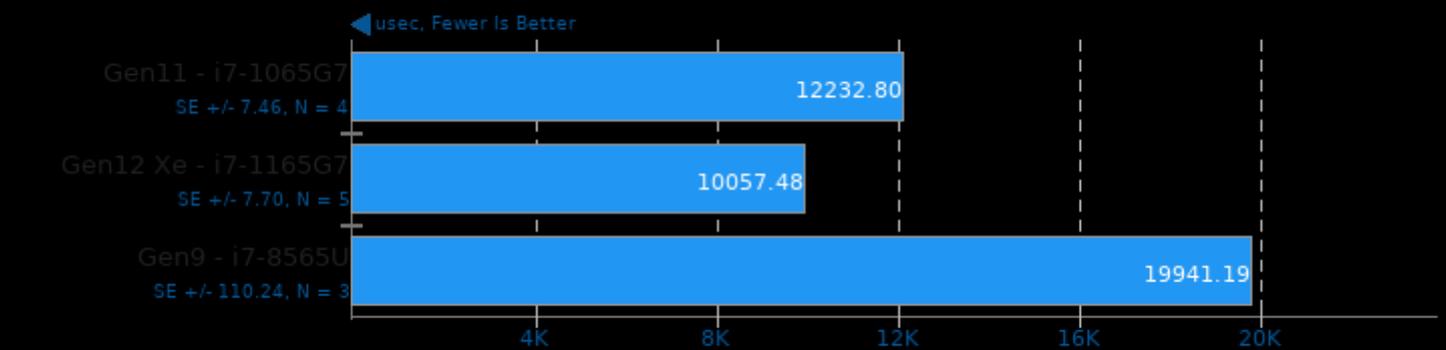
OpenCL Test: Global Memory Bandwidth



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

oneAPI Level Zero Tests

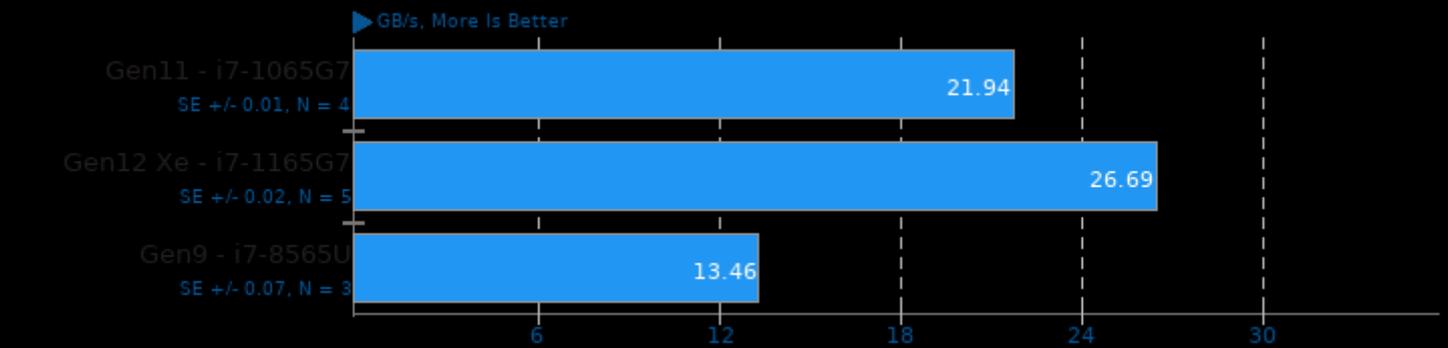
Test: Device-To-Host Bandwidth



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

oneAPI Level Zero Tests

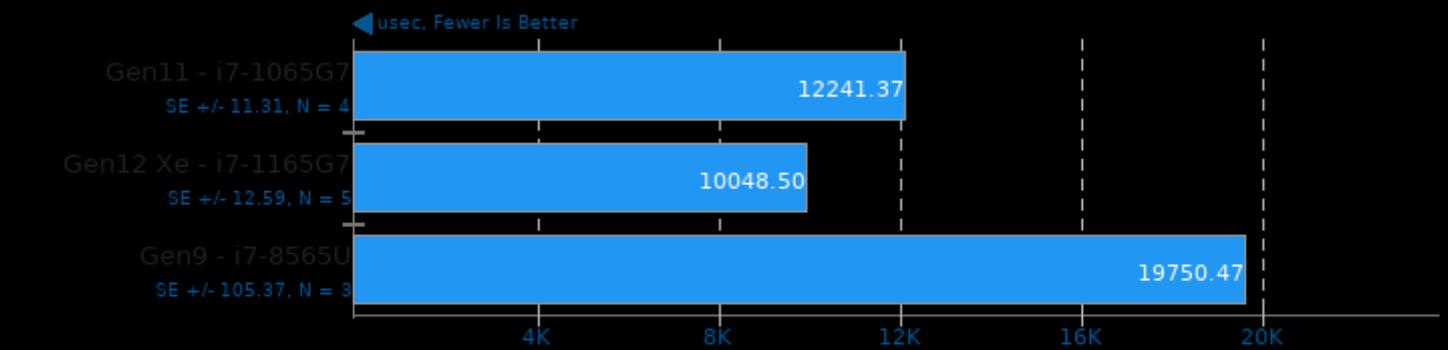
Test: Device-To-Host Bandwidth



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

oneAPI Level Zero Tests

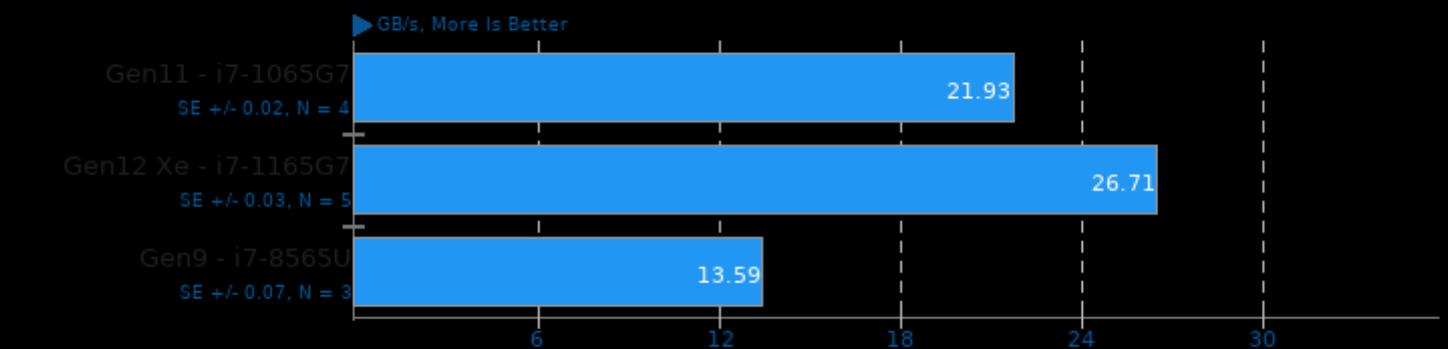
Test: Host-To-Device Bandwidth



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

oneAPI Level Zero Tests

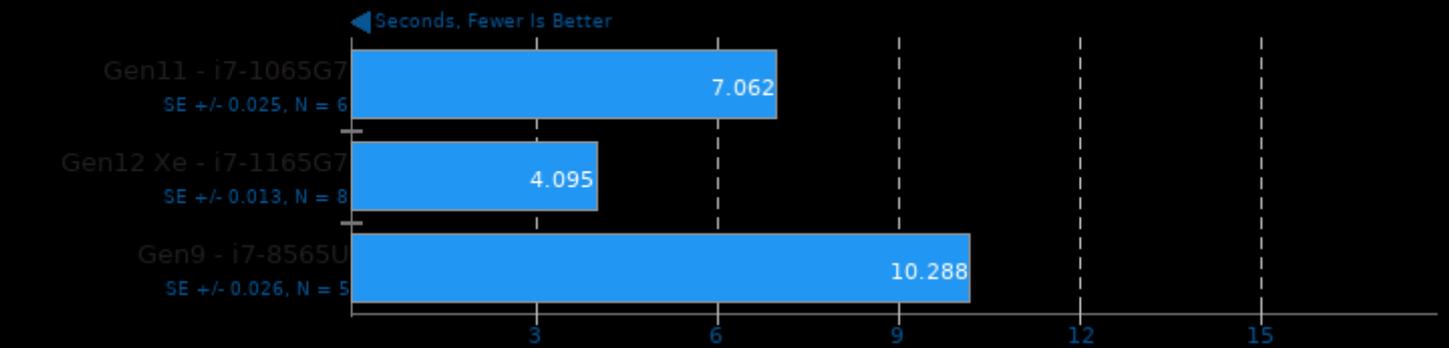
Test: Host-To-Device Bandwidth



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

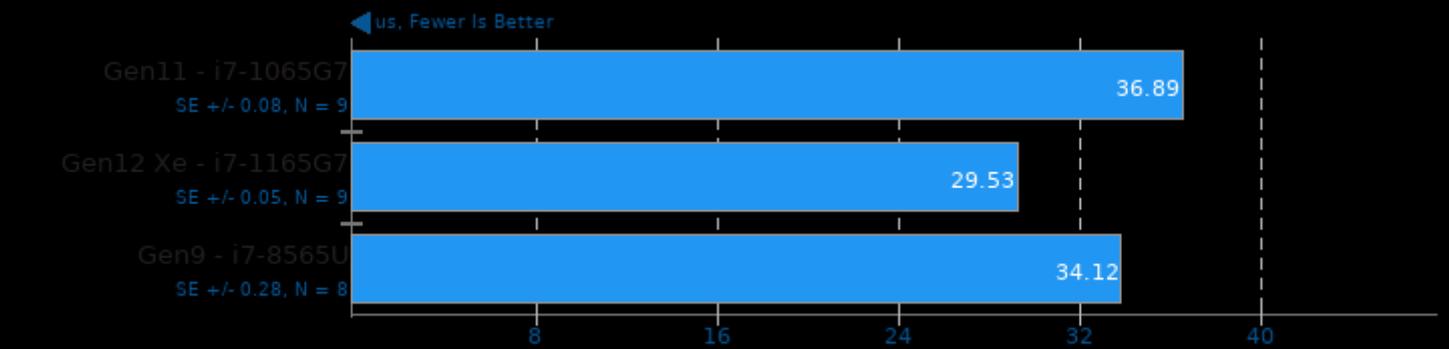
Waifu2x-NCNN Vulkan 20200818

Scale: 2x - Denoise: 3 - TAA: No



clpeak

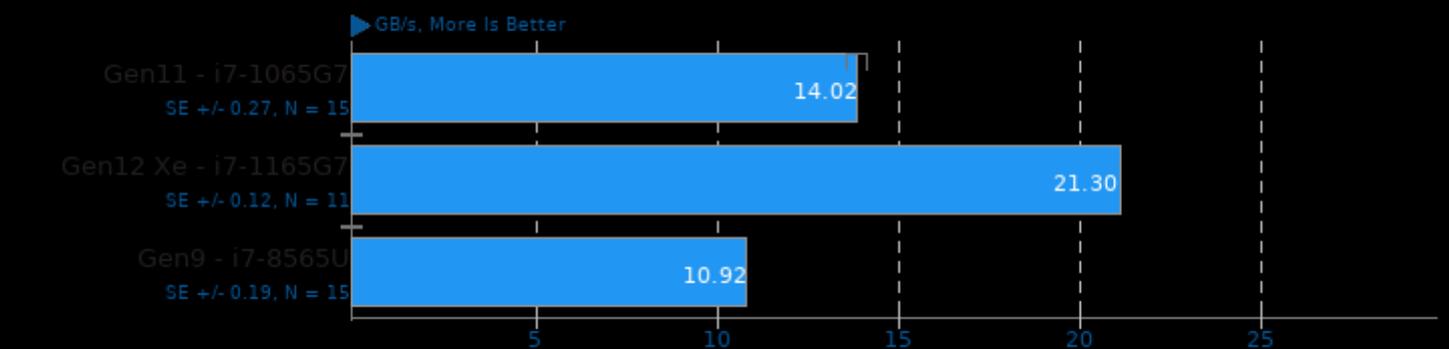
OpenCL Test: Kernel Latency



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

SHOC Scalable Heterogeneous Computing 2015-11-10

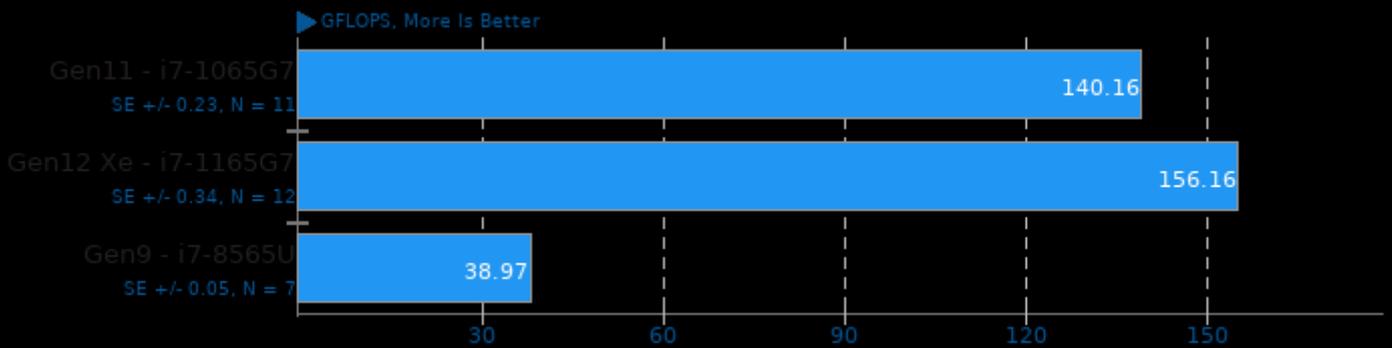
Target: OpenCL - Benchmark: Triad



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

SHOC Scalable Heterogeneous Computing 2015-11-10

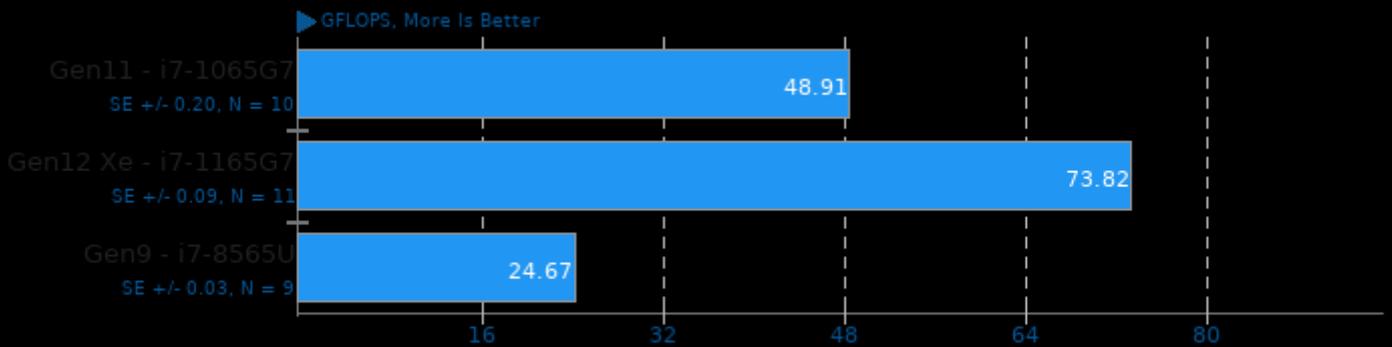
Target: OpenCL - Benchmark: FFT SP



1. (CXX) g++ options: -O2 -ISHOCCommonMPI -ISHOCCommonOpenCL -ISHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

ViennaCL 1.4.2

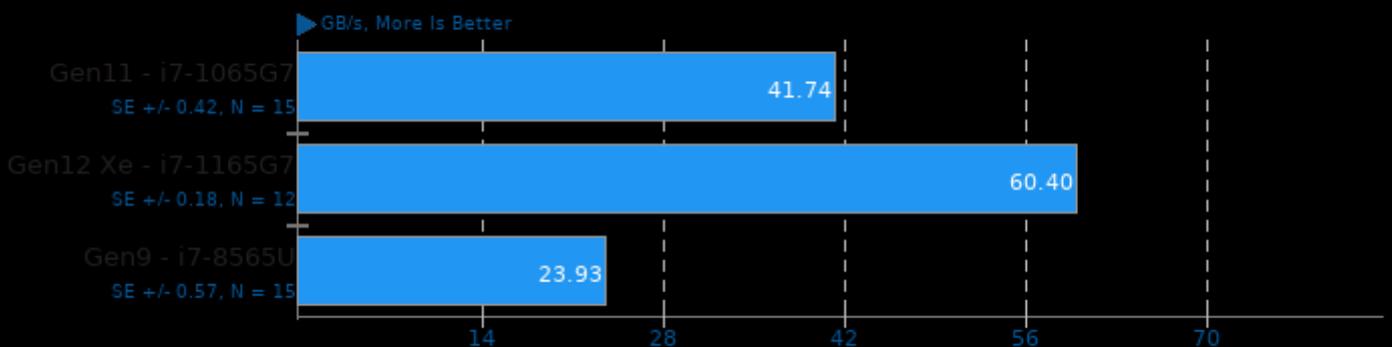
OpenCL LU Factorization



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

SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Bus Speed Readback

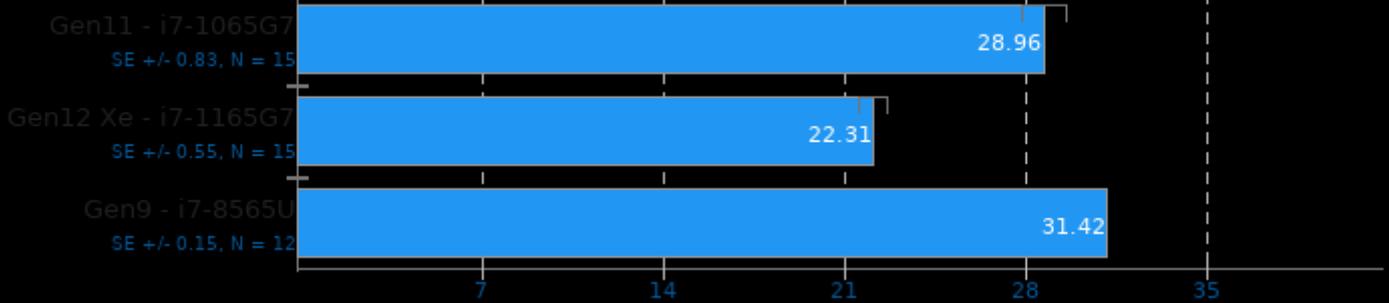


1. (CXX) g++ options: -O2 -ISHOCCommonMPI -ISHOCCommonOpenCL -ISHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

oneAPI Level Zero Tests

Test: Peak Kernel Launch Latency

us, Fewer Is Better

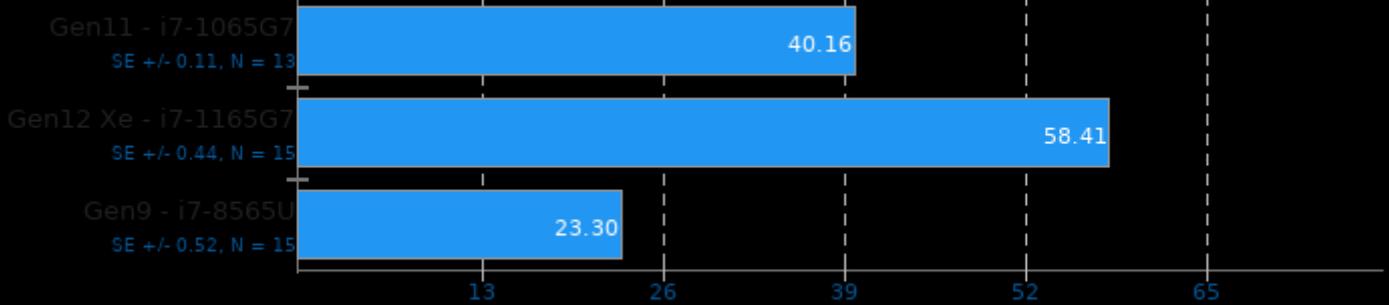


1. (CXX) g++ options: -ldl -pthread

SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Bus Speed Download

GB/s, More Is Better

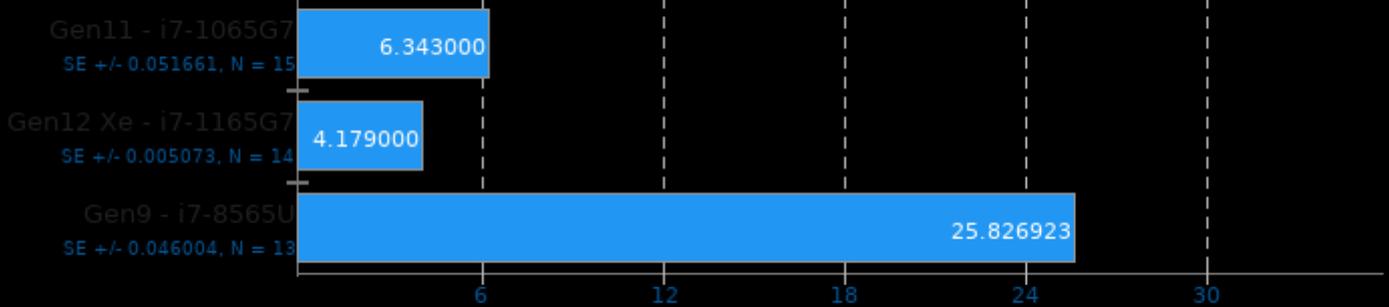


1. (CXX) g++ options: -O2 -lSHOCCommonMPI -lSHOCCommonOpenCL -lSHOCCommon -lOpenCL -lrt -pthread -lmpi_cxx -lmpi

FinanceBench 2016-06-06

Benchmark: Black-Scholes OpenCL

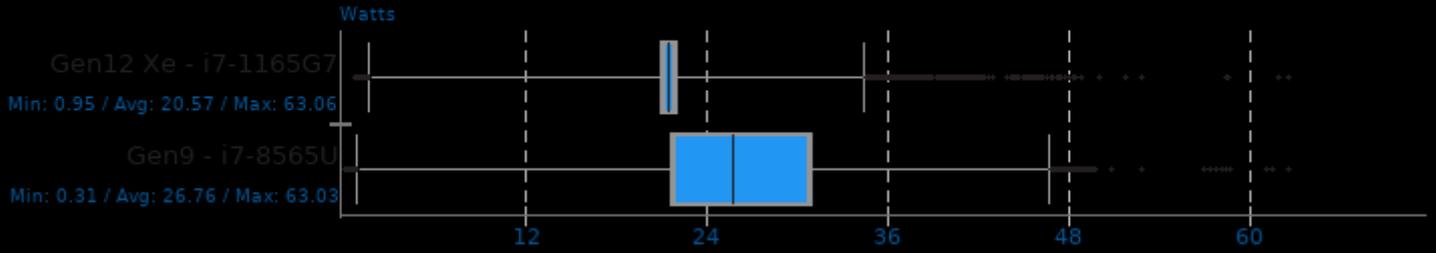
ms, Fewer Is Better



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

CPU Power Consumption Monitor

Phoronix Test Suite System Monitoring



Waifu2x-NCNN Vulkan 20200818

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	3.8	17.3	22.6
Gen12 Xe - i7-1165G7	1.6	19.8	23.3
Gen9 - i7-8565U	1.0	24.8	27.0

▼ Watts, Fewer Is Better

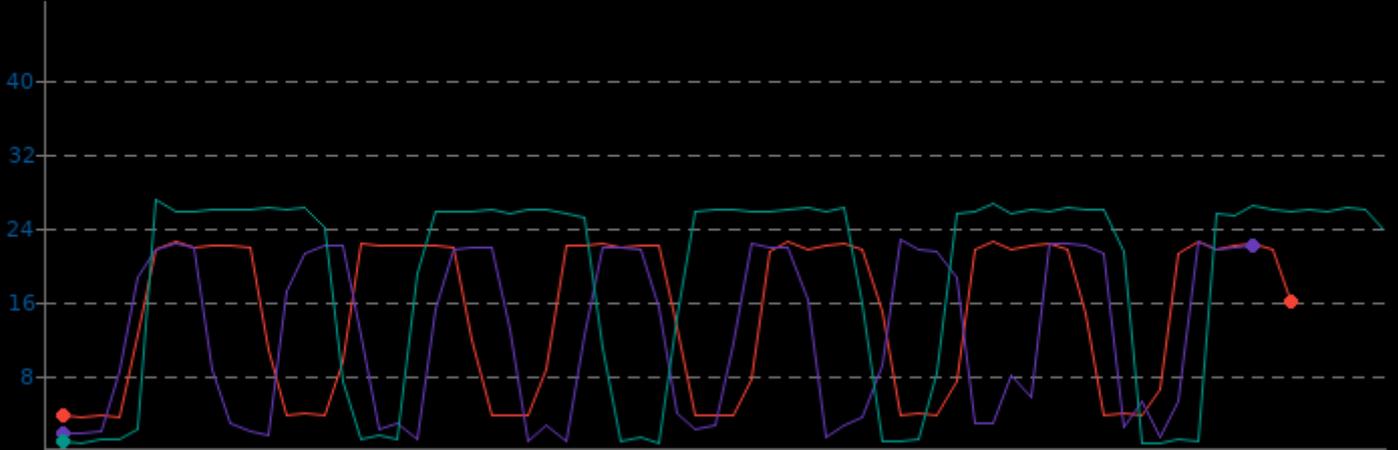


Waifu2x-NCNN Vulkan 20200818

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	3.7	14.9	22.6
Gen12 Xe - i7-1165G7	1.1	12.7	22.7
Gen9 - i7-8565U	0.8	18.5	27.0

▼ Watts, Fewer Is Better

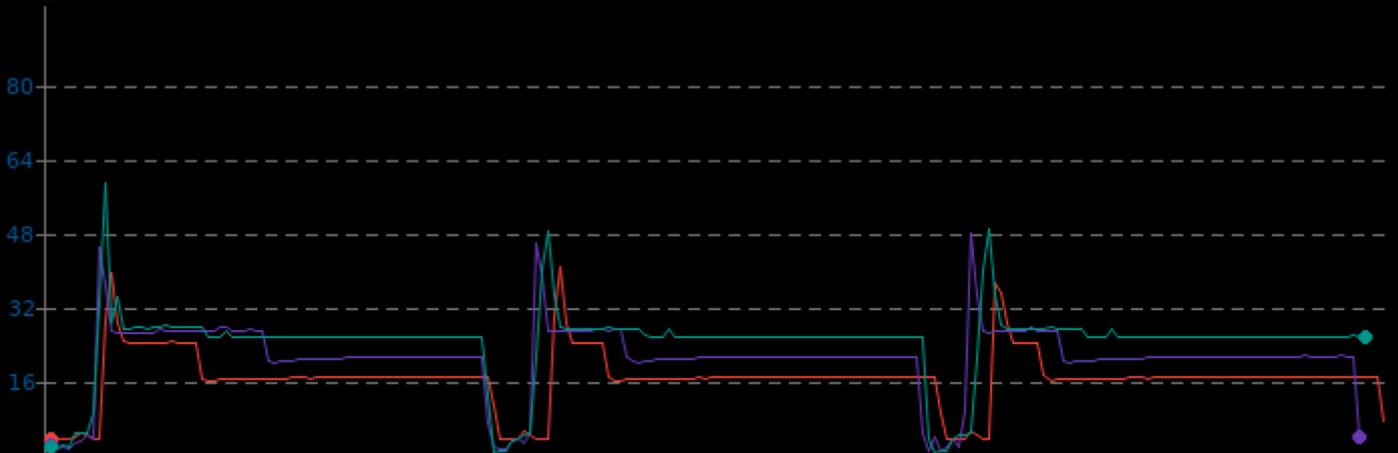


GpuTest 0.7.0

CPU Power Consumption Monitor

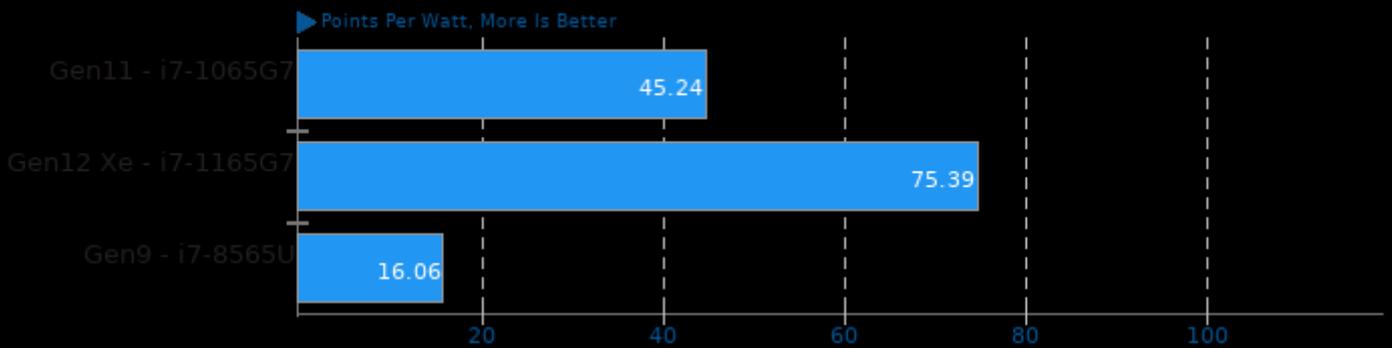
	Min	Avg	Max
Gen11 - i7-1065G7	3.8	16.8	40.9
Gen12 Xe - i7-1165G7	1.2	21.2	47.8
Gen9 - i7-8565U	0.9	24.3	58.9

▼ Watts, Fewer Is Better



GpuTest 0.7.0

Test: Pixmark Volplosion - Resolution: 1920 x 1200 - Mode: Fullscreen

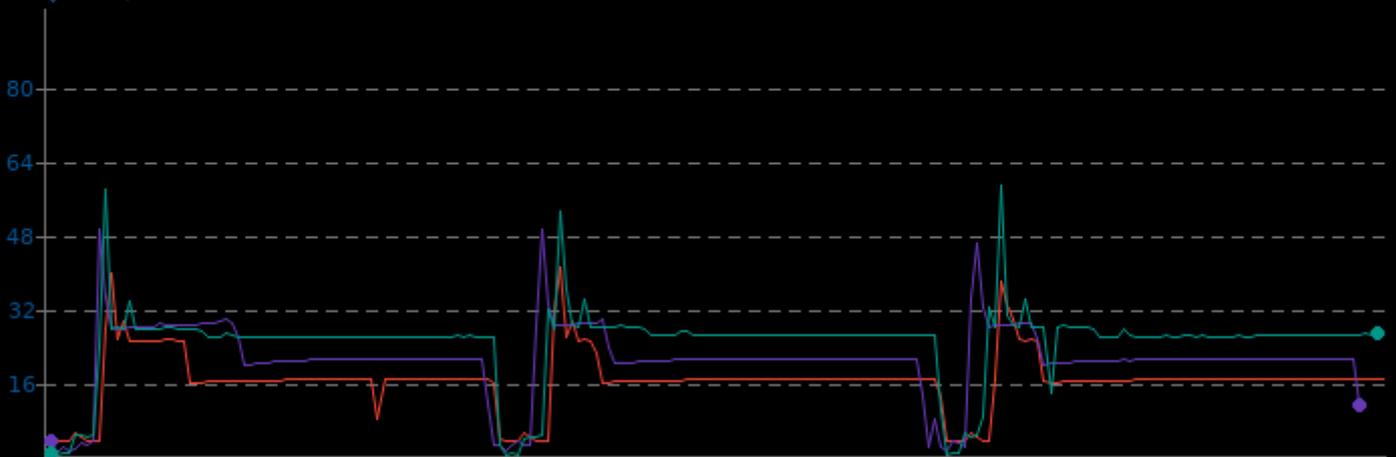


GpuTest 0.7.0

CPU Power Consumption Monitor

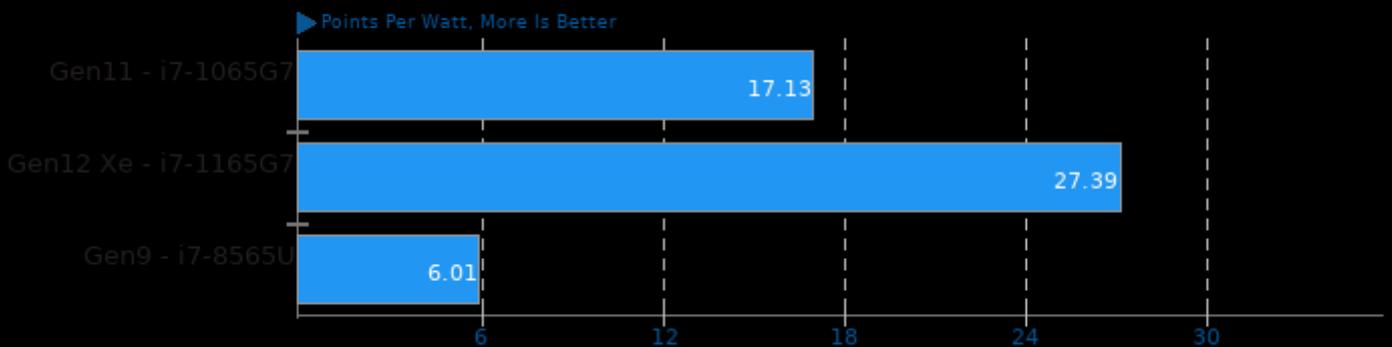
	Min	Avg	Max
Gen11 - i7-1065G7	3.8	16.8	41.2
Gen12 Xe - i7-1165G7	1.3	21.3	49.3
Gen9 - i7-8565U	0.9	24.8	58.7

▼ Watts, Fewer Is Better



GpuTest 0.7.0

Test: Pixmark Piano - Resolution: 1920 x 1200 - Mode: Fullscreen

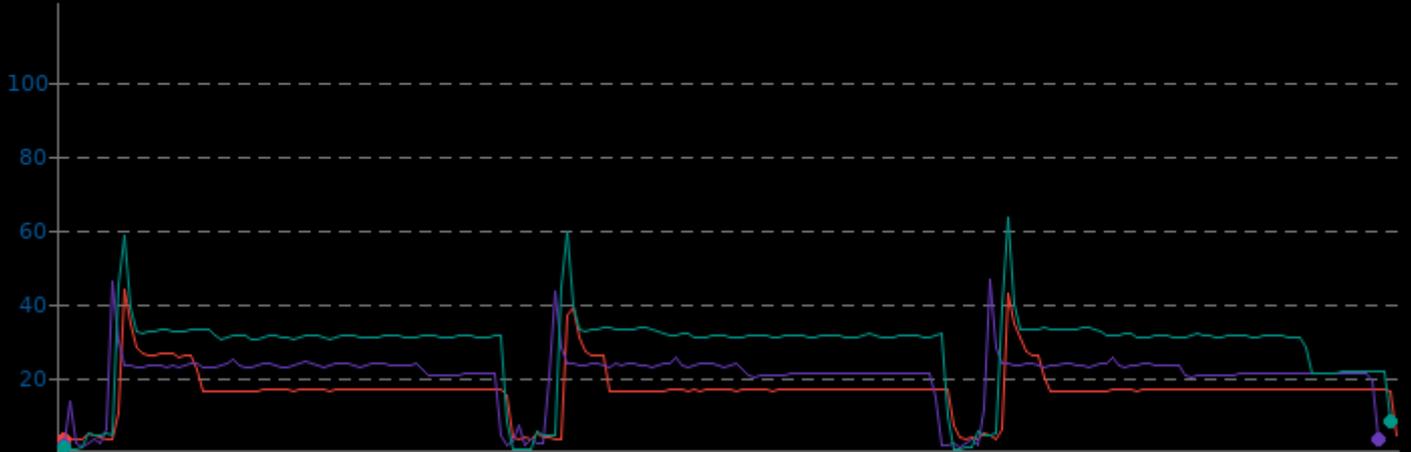


GpuTest 0.7.0

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	3.8	16.8	43.8
Gen12 Xe - i7-1165G7	1.5	20.9	46.8
Gen9 - i7-8565U	0.9	28.3	63.0

▼ Watts, Fewer Is Better



GpuTest 0.7.0

Test: TessMark - Resolution: 1920 x 1200 - Mode: Fullscreen

► Points Per Watt, More Is Better

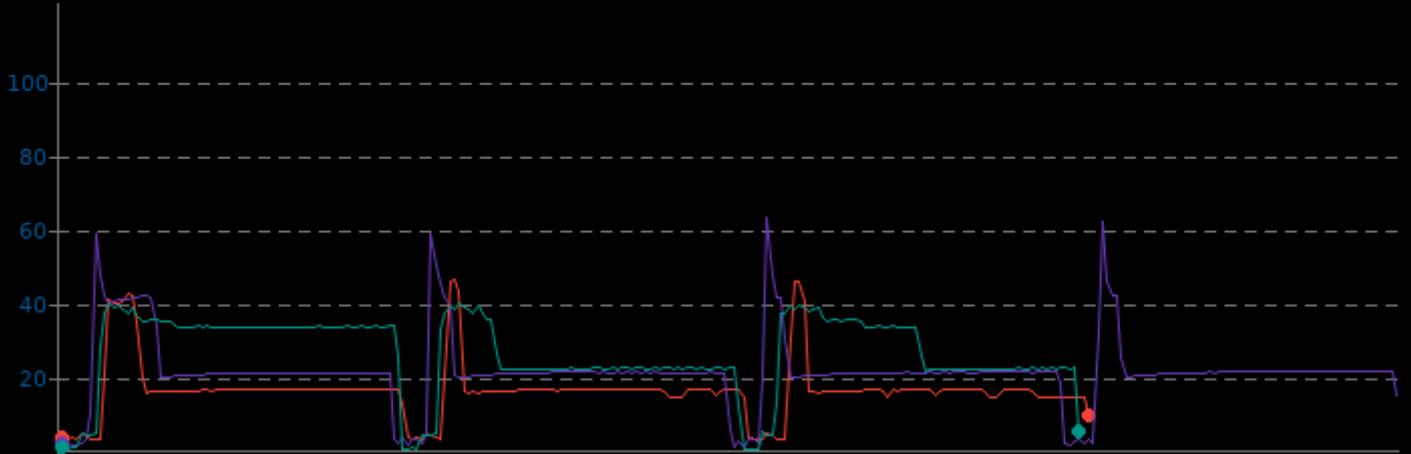


GpuTest 0.7.0

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	3.8	16.8	46.5
Gen12 Xe - i7-1165G7	1.6	22.0	63.1
Gen9 - i7-8565U	1.0	26.6	40.3

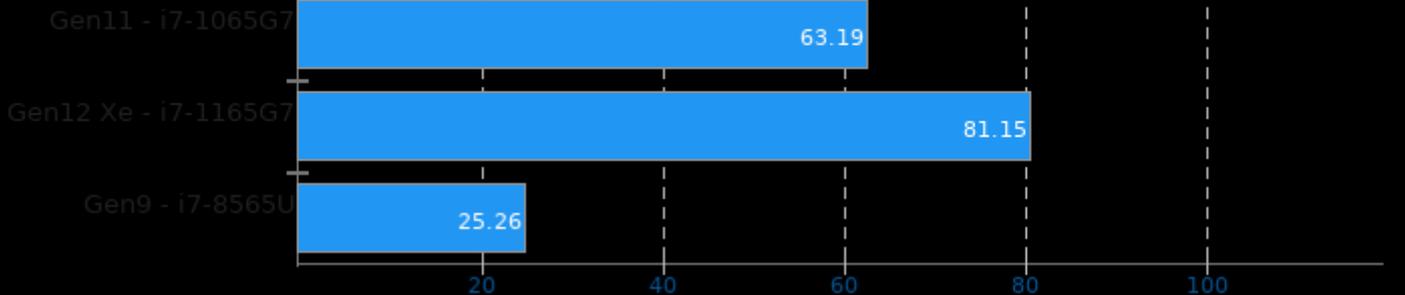
▼ Watts, Fewer Is Better



GpuTest 0.7.0

Test: Furmark - Resolution: 1920 x 1200 - Mode: Fullscreen

► Points Per Watt, More Is Better

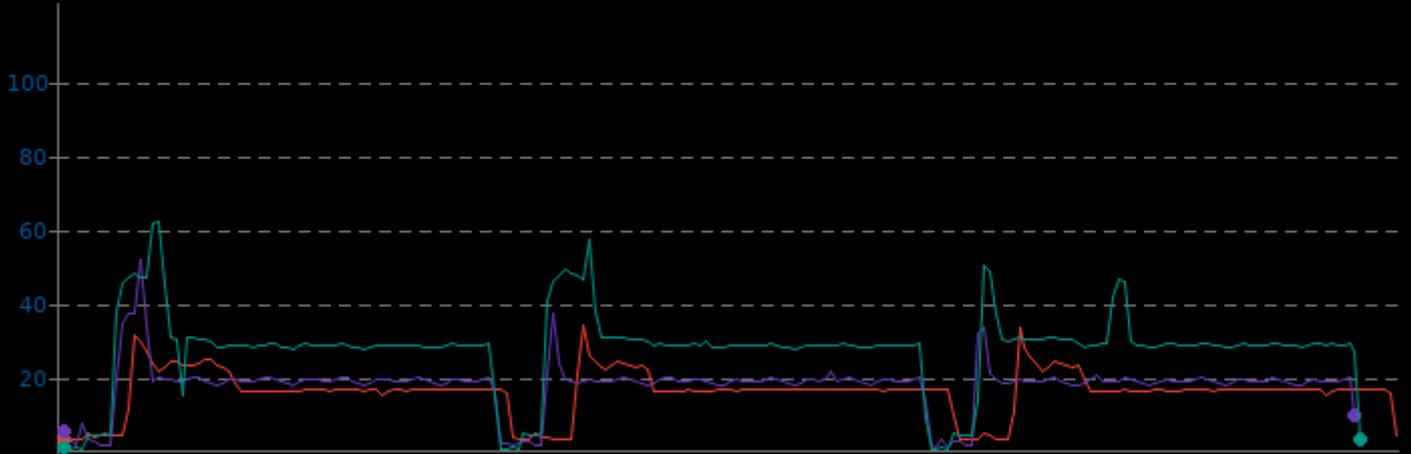


GpuTest 0.7.0

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	3.7	16.5	34.4
Gen12 Xe - i7-1165G7	1.4	18.2	52.2
Gen9 - i7-8565U	0.8	27.8	62.0

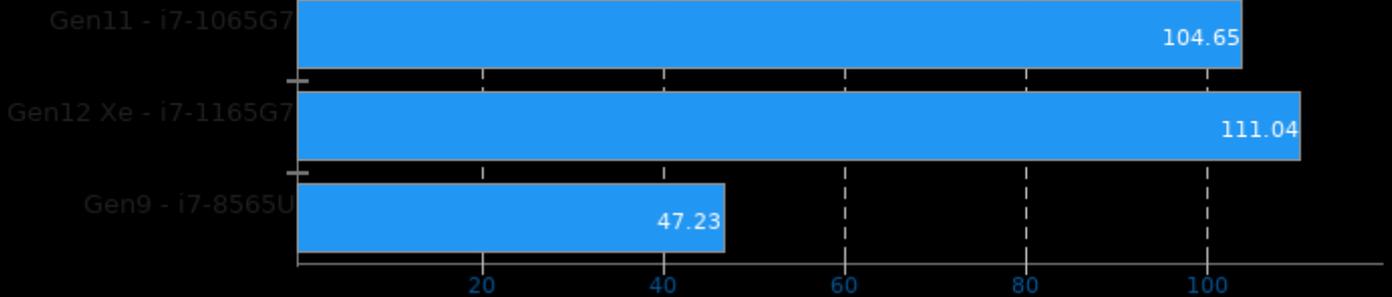
▼ Watts, Fewer Is Better

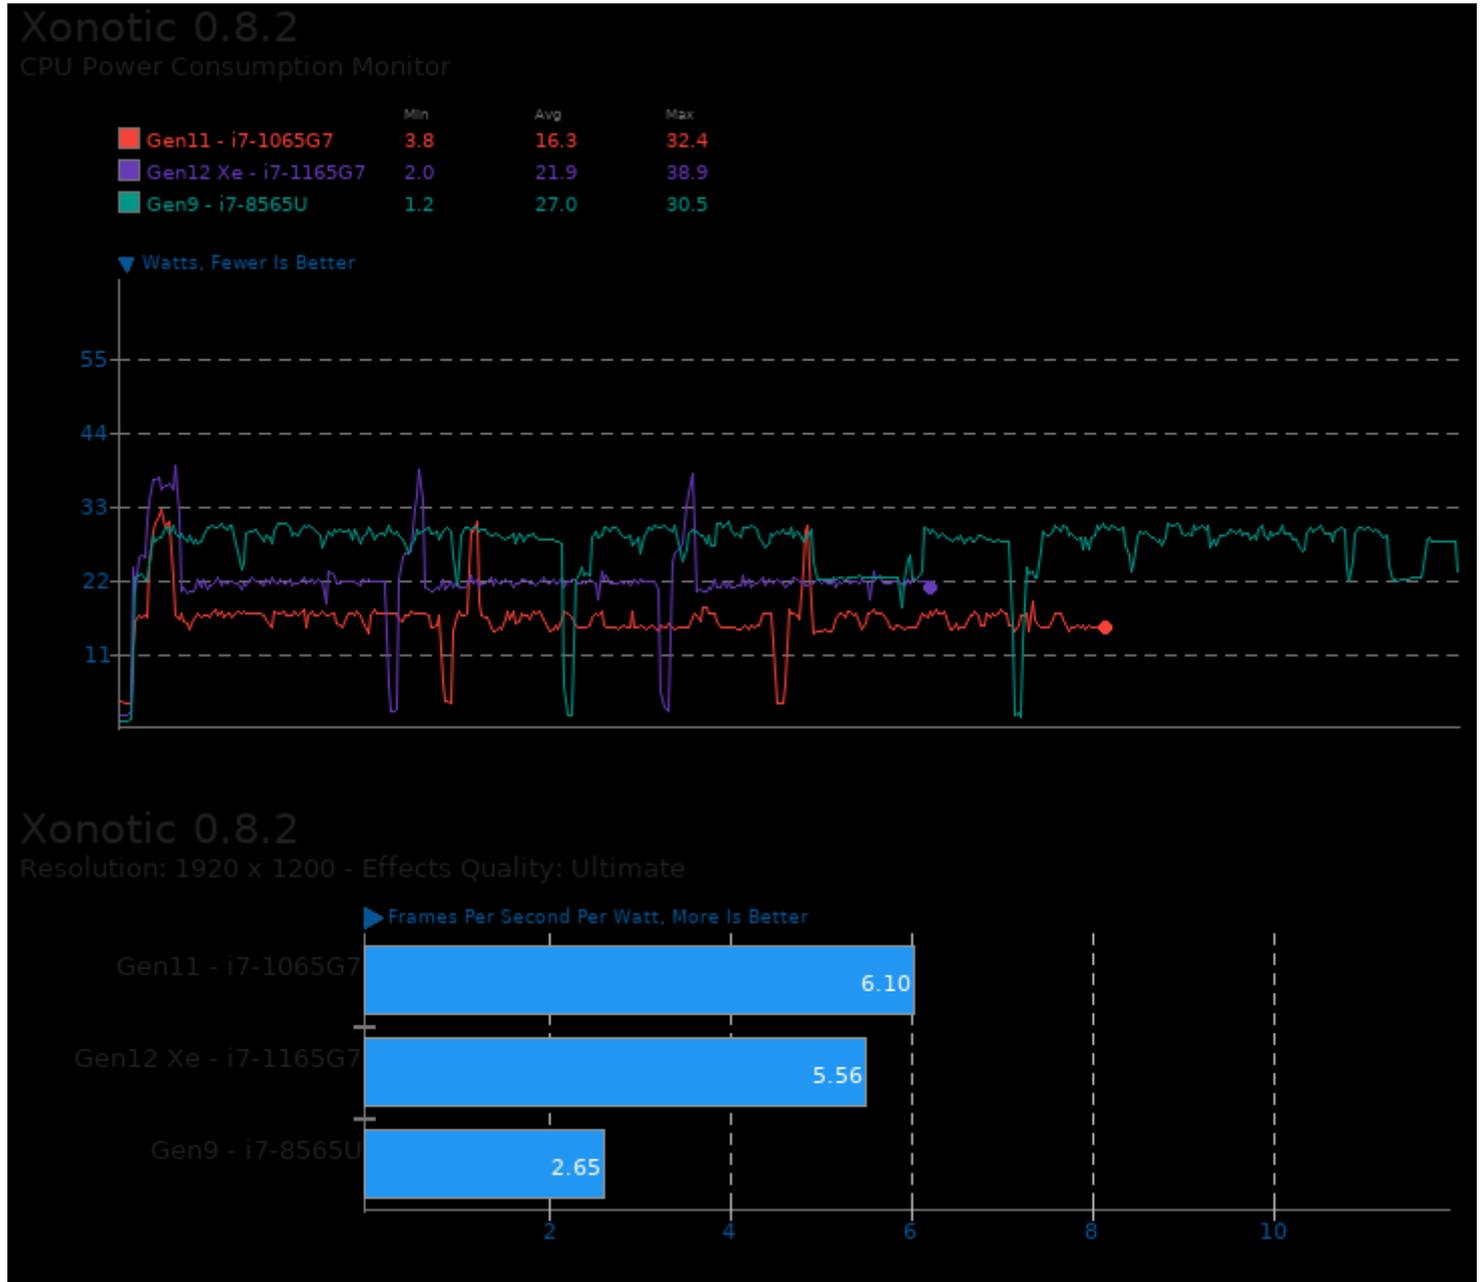


GpuTest 0.7.0

Test: GiMark - Resolution: 1920 x 1200 - Mode: Fullscreen

► Points Per Watt, More Is Better



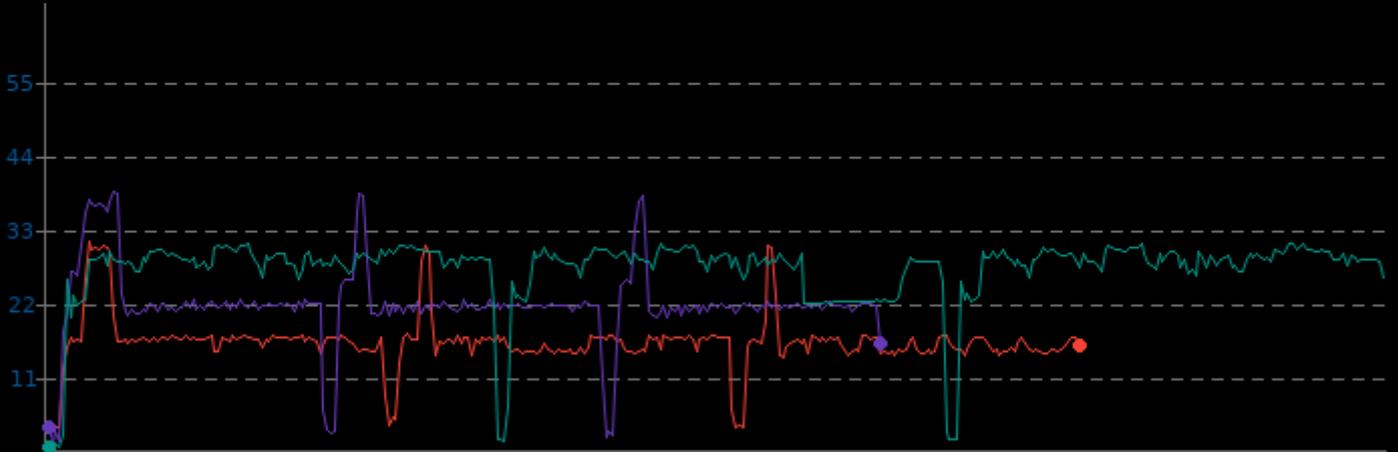


Xonotic 0.8.2

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	3.7	16.4	31.2
Gen12 Xe - i7-1165G7	1.5	22.0	38.7
Gen9 - i7-8565U	0.9	27.1	30.9

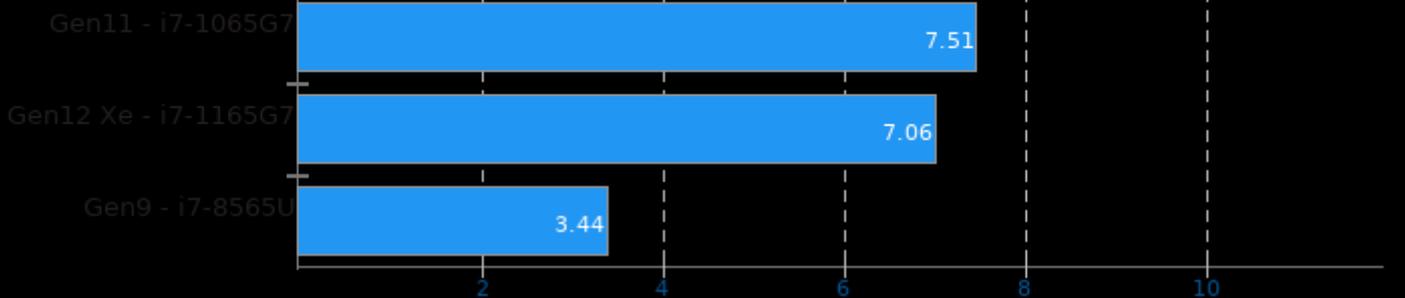
▼ Watts, Fewer Is Better

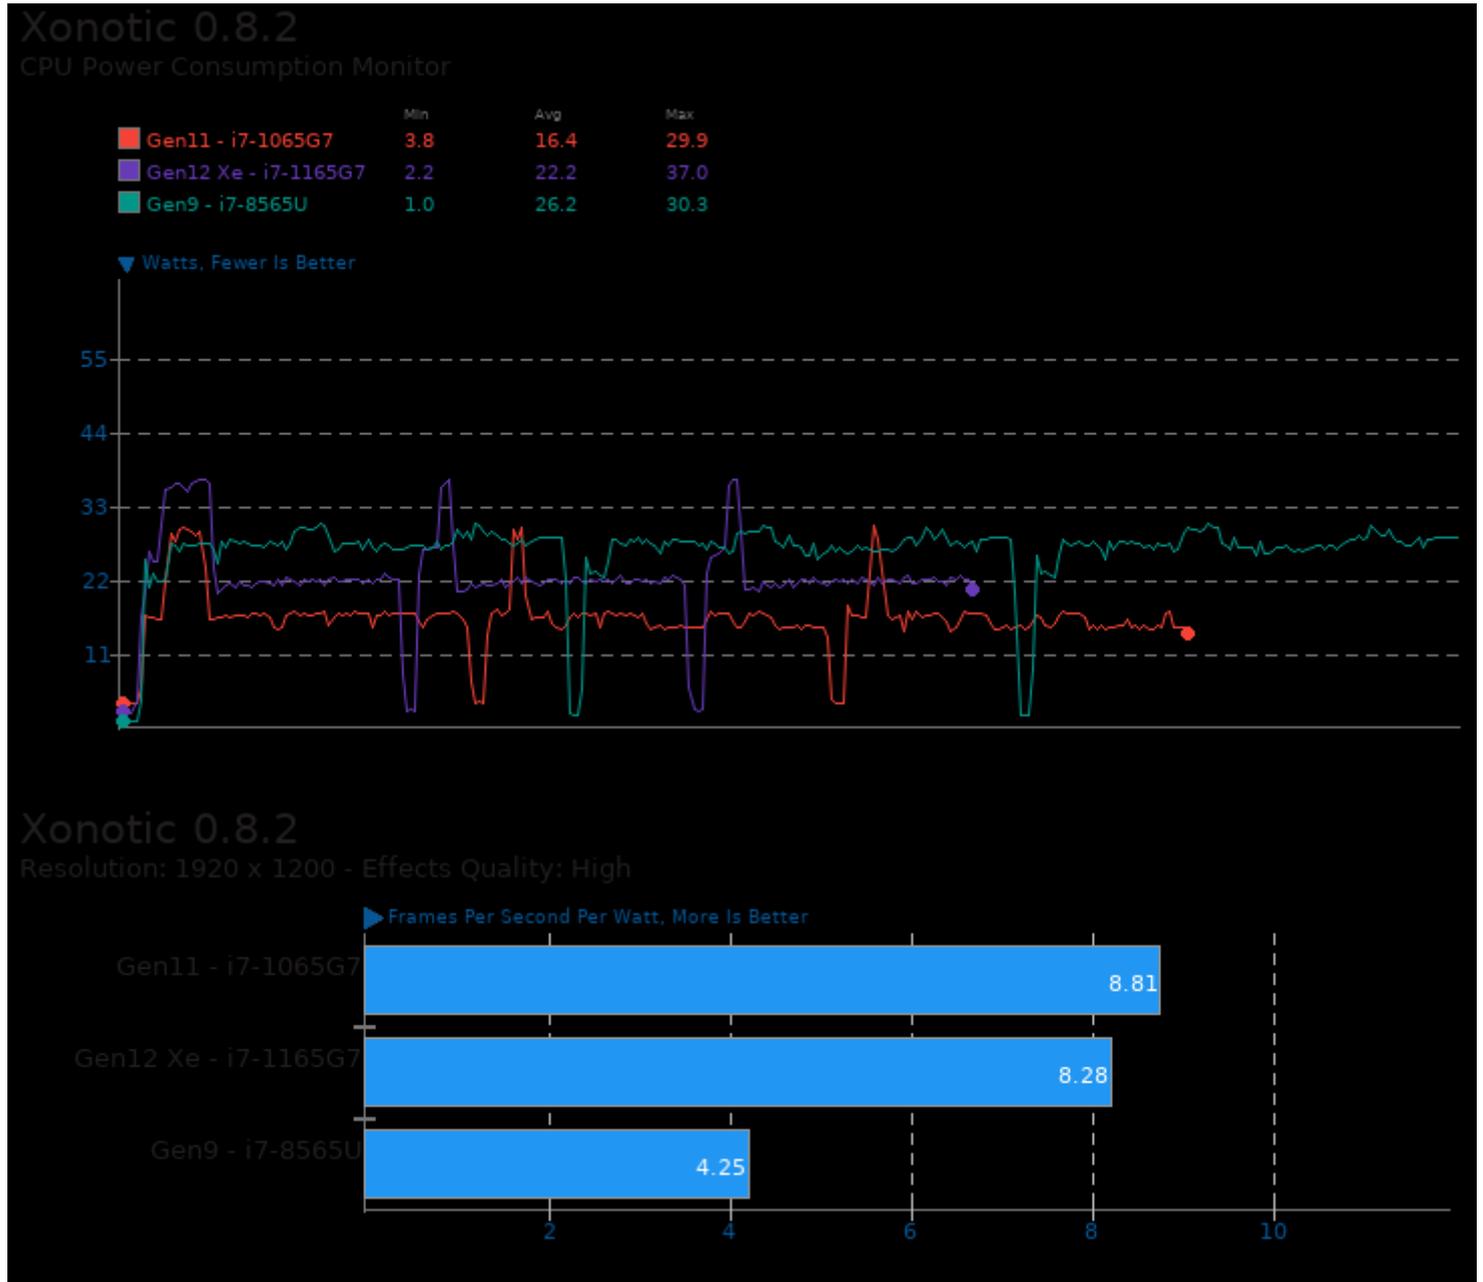


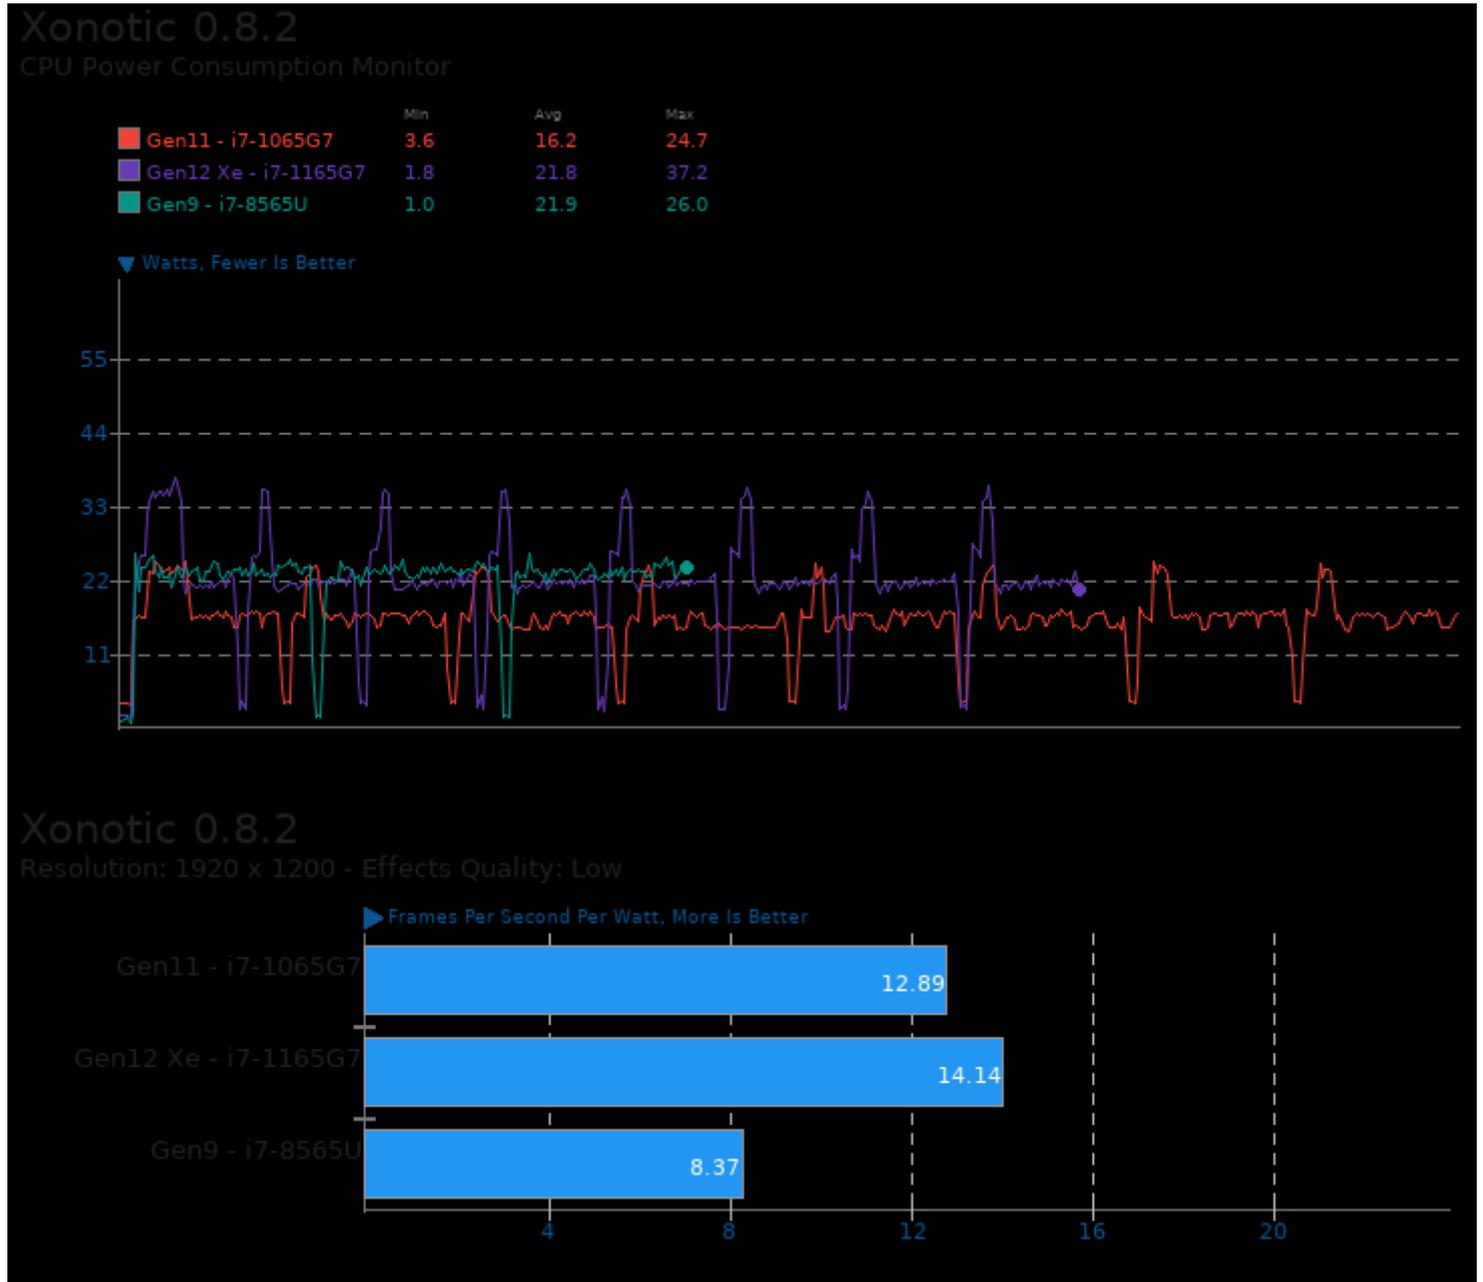
Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Ultra

► Frames Per Second Per Watt, More Is Better







Unigine Valley 1.0

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.9	15.7	29.9
Gen12 Xe - i7-1165G7	1.8	21.9	36.2
Gen9 - i7-8565U	1.0	26.7	36.6

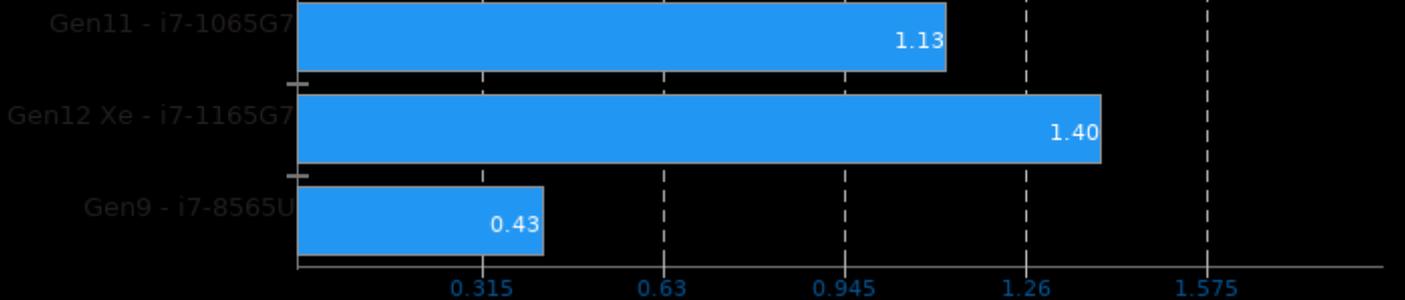
▼ Watts, Fewer Is Better



Unigine Valley 1.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Renderer: OpenGL

► Frames Per Second Per Watt, More Is Better

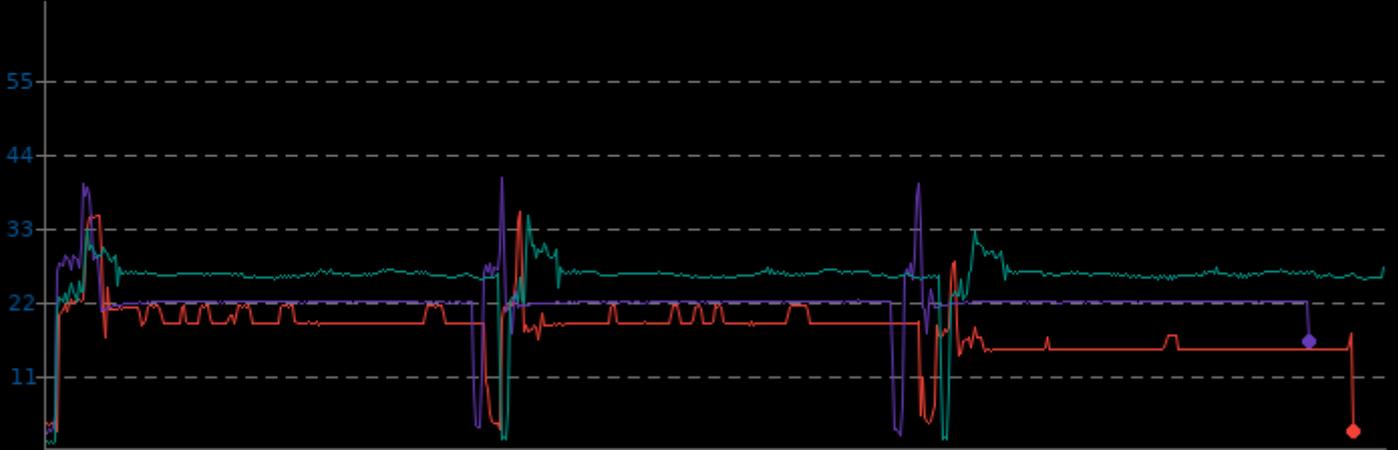


Unigine Superposition 1.0

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	2.9	18.0	35.3
Gen12 Xe - i7-1165G7	2.4	22.0	40.4
Gen9 - i7-8565U	1.1	25.7	34.7

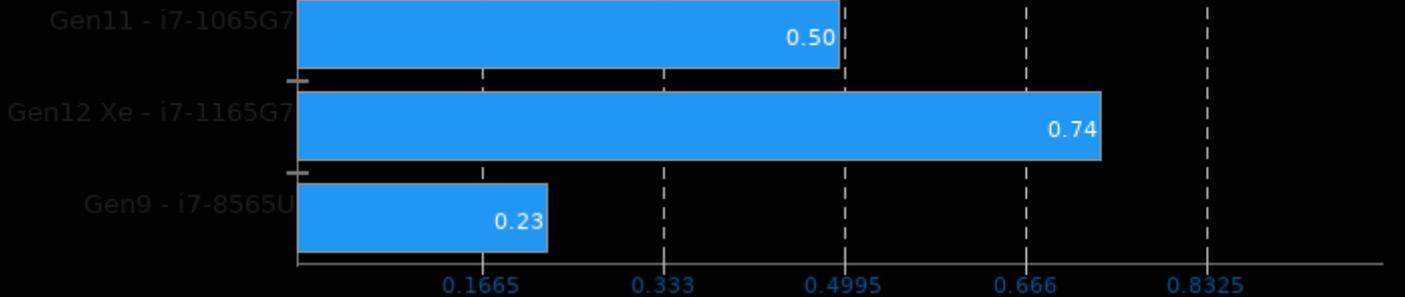
▼ Watts, Fewer Is Better



Unigine Superposition 1.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Quality: Medium - Renderer: OpenGL

► Frames Per Second Per Watt, More Is Better

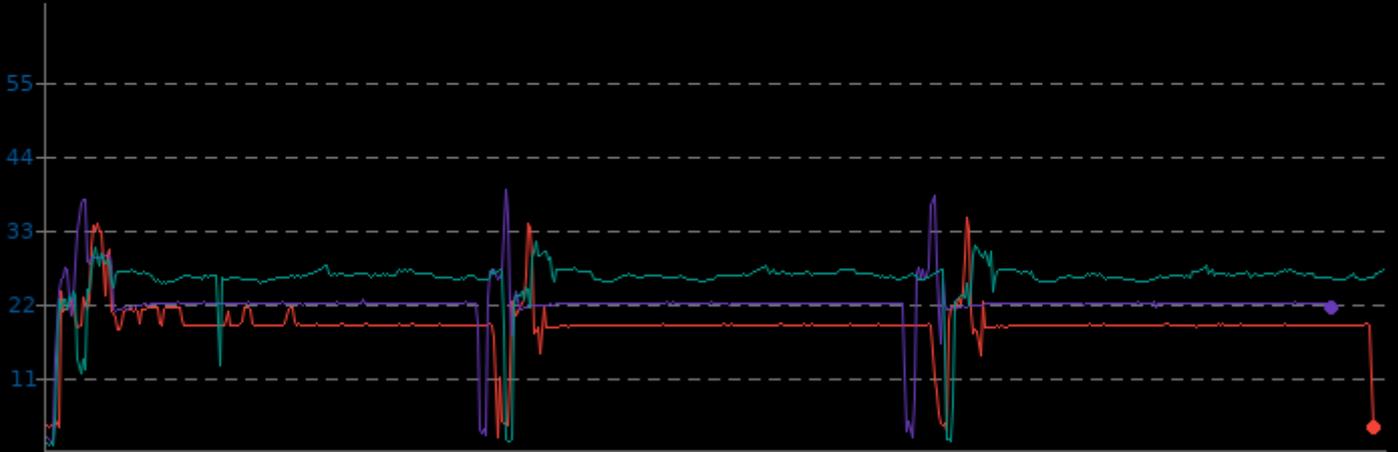


Unigine Superposition 1.0

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	2.3	19.0	34.8
Gen12 Xe - i7-1165G7	1.7	22.1	38.8
Gen9 - i7-8565U	1.1	25.6	31.3

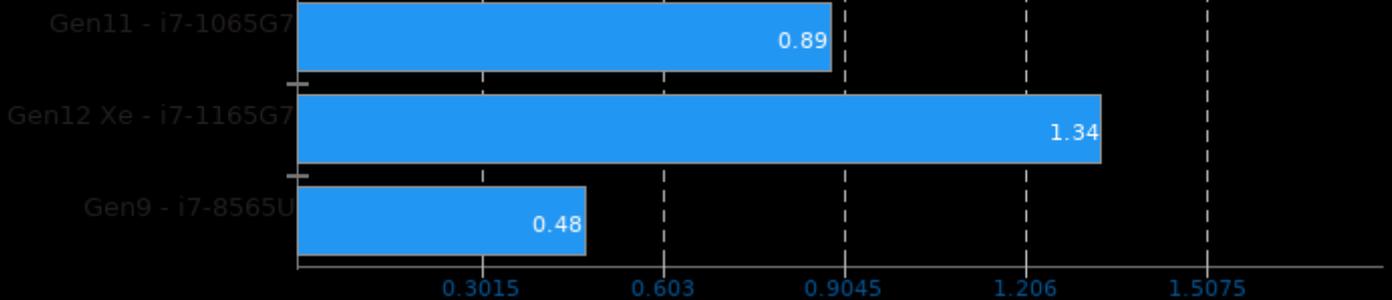
▼ Watts, Fewer Is Better



Unigine Superposition 1.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Quality: Low - Renderer: OpenGL

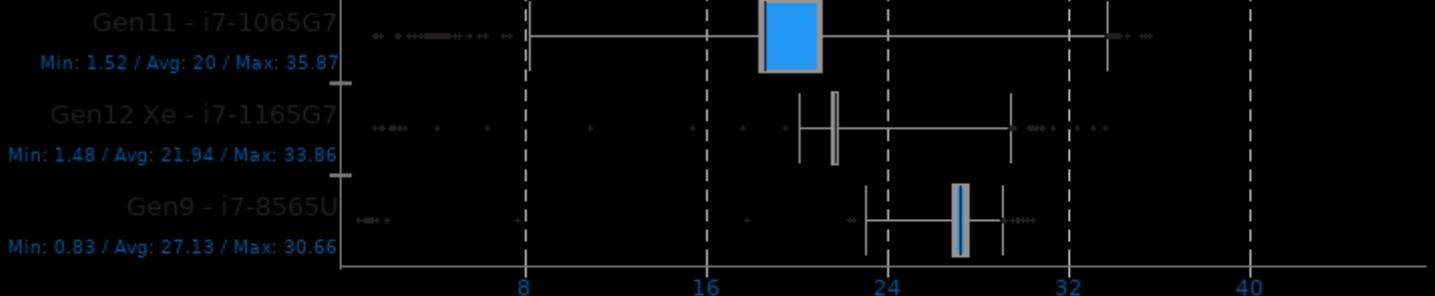
► Frames Per Second Per Watt, More Is Better



Unigine Heaven 4.0

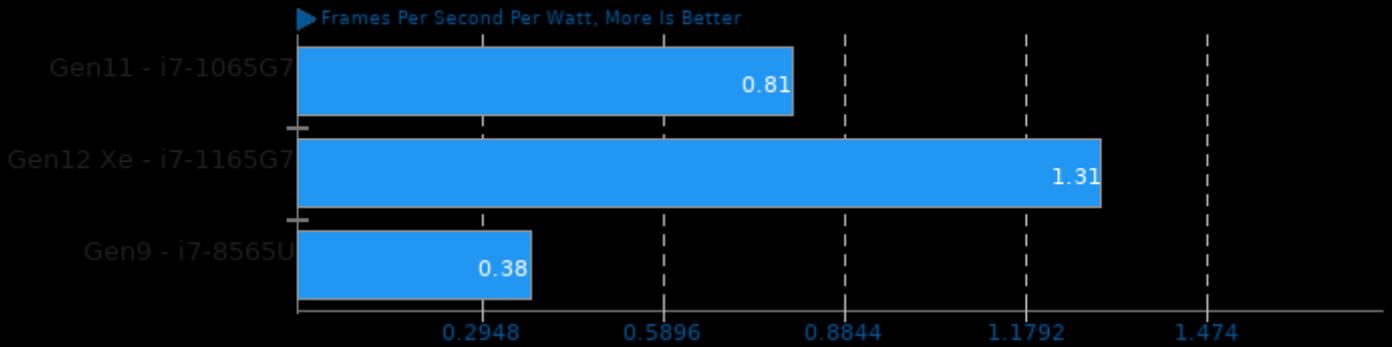
CPU Power Consumption Monitor

◀ Watts, Fewer Is Better



Unigine Heaven 4.0

Resolution: 1920 x 1200 - Mode: Fullscreen - Renderer: OpenGL

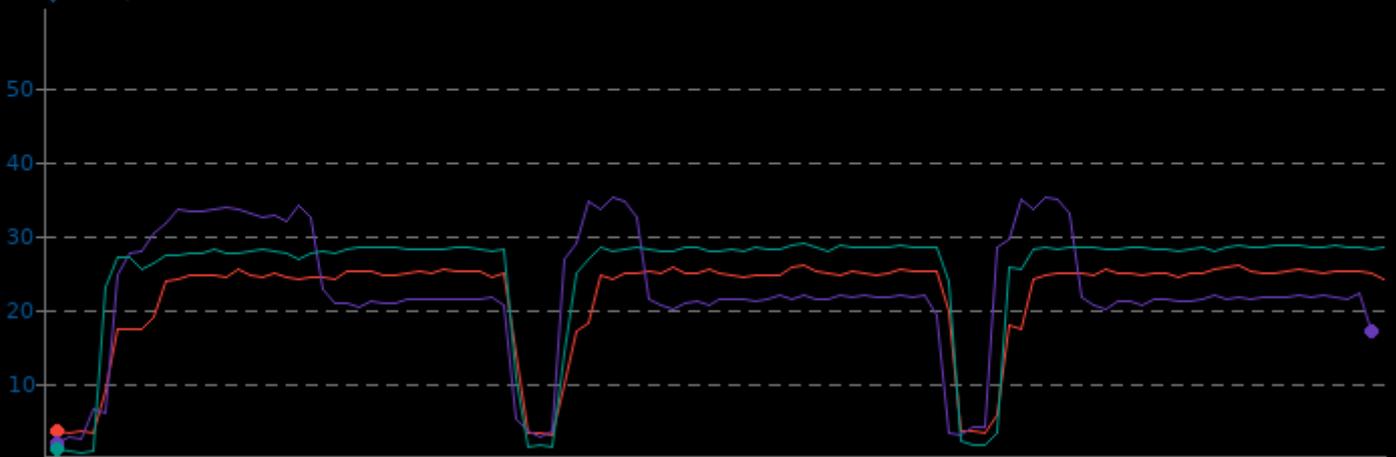


Tesseract 2014-05-12

CPU Power Consumption Monitor

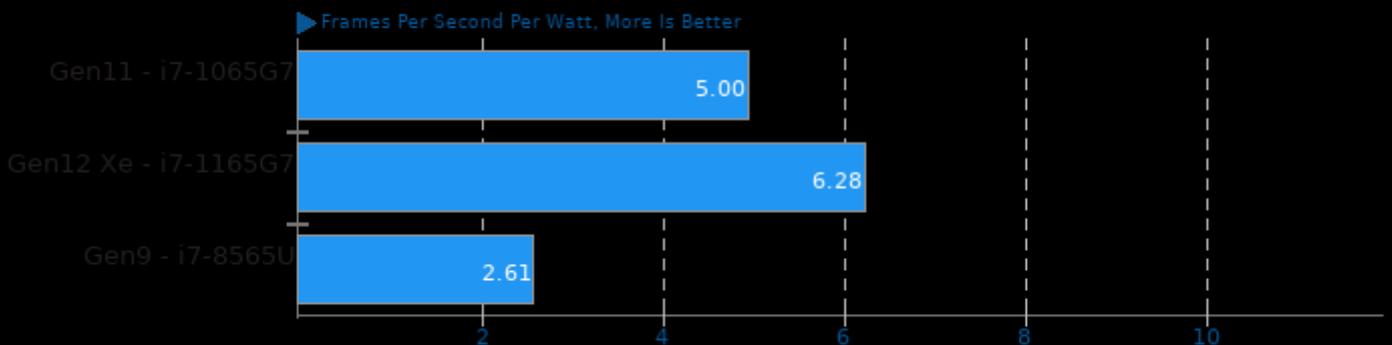
	Min	Avg	Max
Gen11 - i7-1065G7	3.3	21.9	26.0
Gen12 Xe - i7-1165G7	2.1	22.3	35.1
Gen9 - i7-8565U	0.8	25.0	28.9

▼ Watts, Fewer Is Better



Tesseract 2014-05-12

Resolution: 1920 x 1200



ET: Legacy 2.75

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.7	16.0	30.9
Gen12 Xe - i7-1165G7	1.5	21.9	33.4
Gen9 - i7-8565U	1.0	21.1	25.6

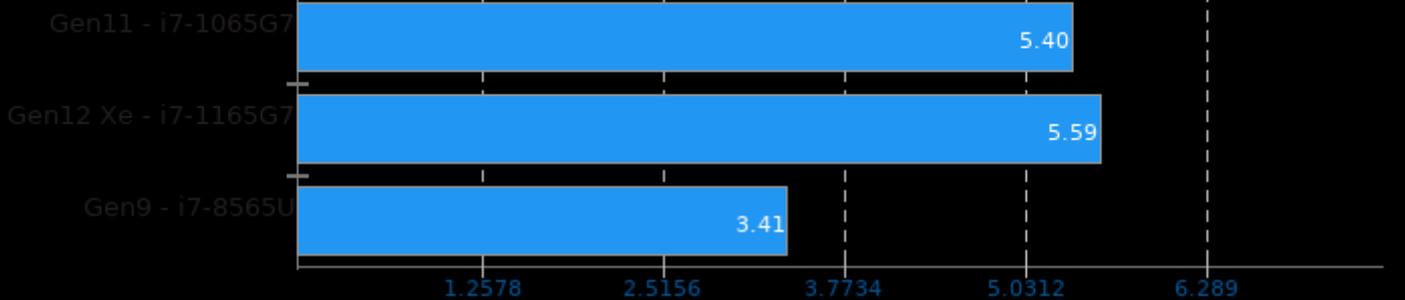
▼ Watts, Fewer Is Better



ET: Legacy 2.75

Renderer: Renderer2 - Resolution: 1920 x 1200

► Frames Per Second Per Watt, More Is Better

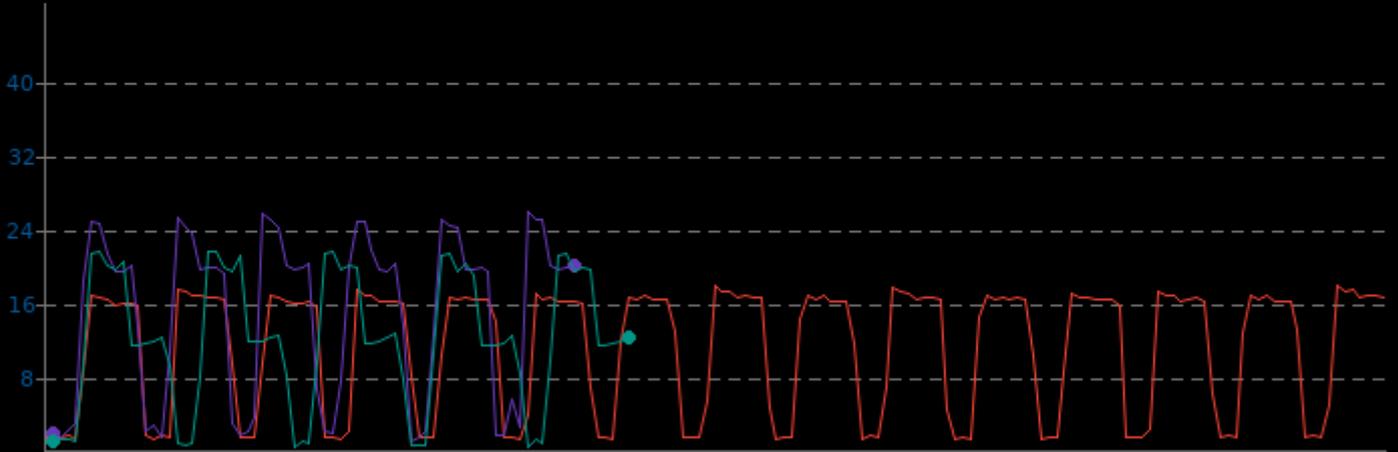


clpeak

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.5	11.4	18.0
Gen12 Xe - i7-1165G7	1.3	15.2	26.0
Gen9 - i7-8565U	0.7	12.2	21.7

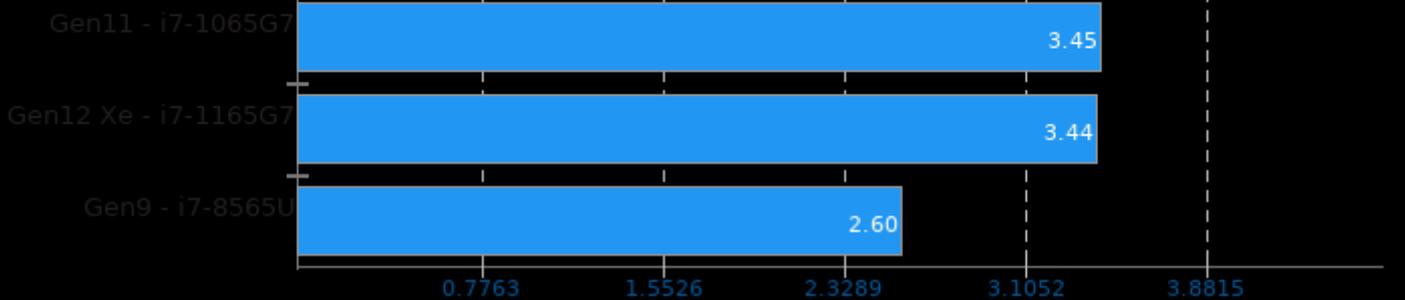
▼ Watts, Fewer Is Better



clpeak

OpenCL Test: Transfer Bandwidth enqueueWriteBuffer

▶ GBPS Per Watt, More Is Better

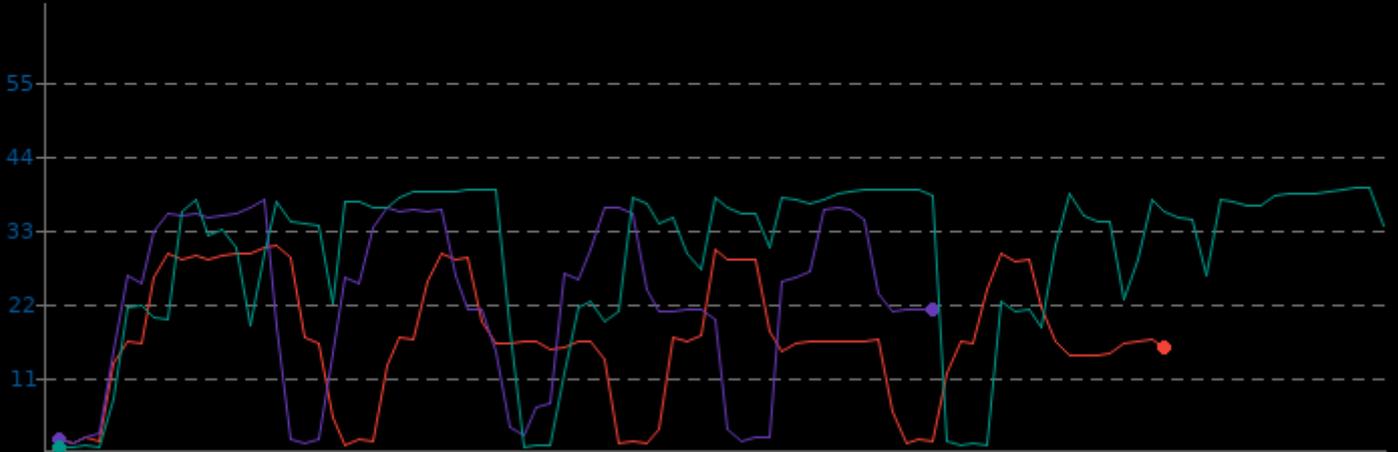


clpeak

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.3	16.9	30.6
Gen12 Xe - i7-1165G7	1.5	22.4	37.3
Gen9 - i7-8565U	0.8	29.1	39.1

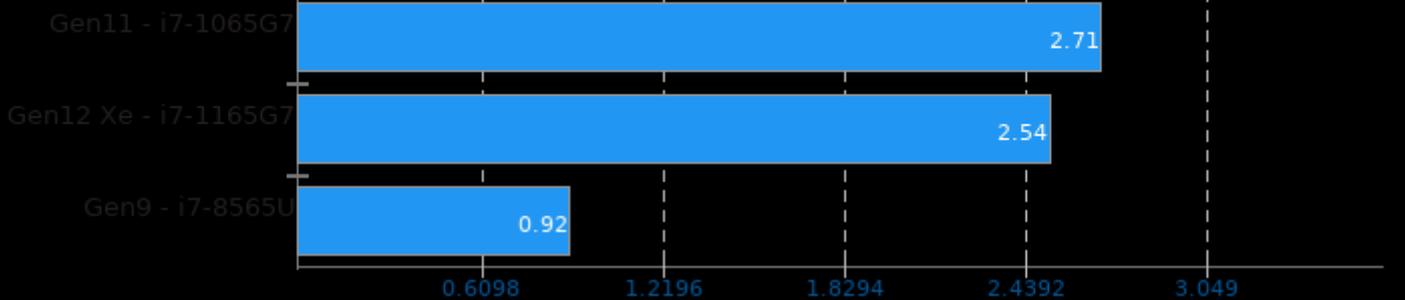
▼ Watts, Fewer Is Better



clpeak

OpenCL Test: Global Memory Bandwidth

▶ GBPS Per Watt, More Is Better

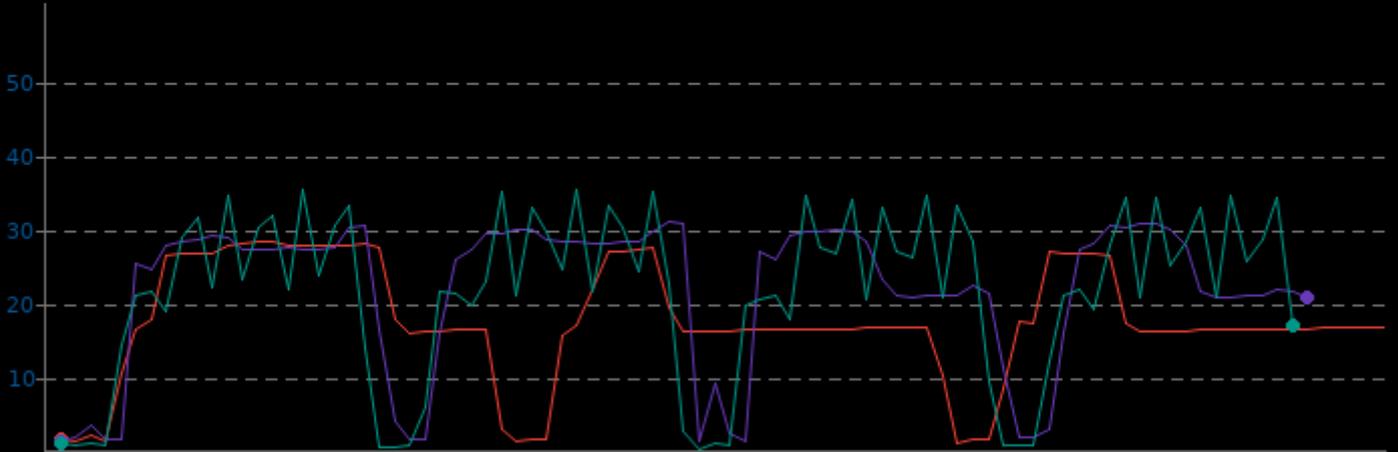


clpeak

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.4	17.6	28.3
Gen12 Xe - i7-1165G7	1.5	21.9	31.2
Gen9 - i7-8565U	0.6	21.6	35.5

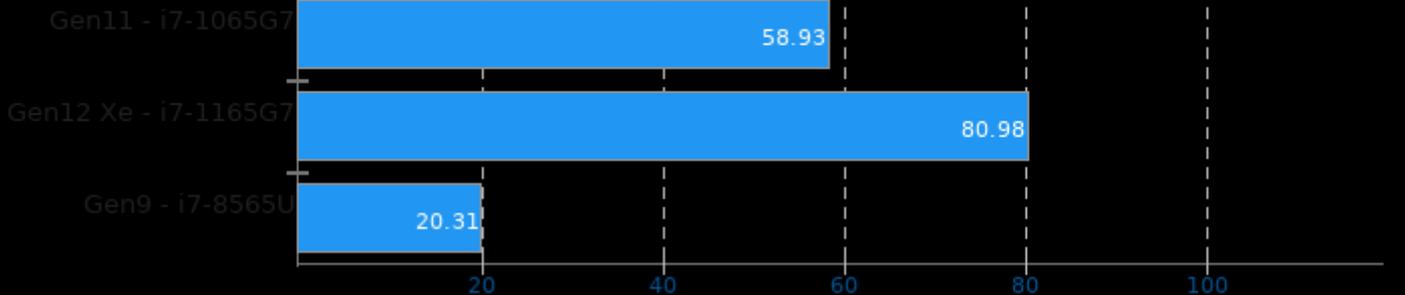
▼ Watts, Fewer Is Better



clpeak

OpenCL Test: Single-Precision Float

▶ GFLOPS Per Watt, More Is Better

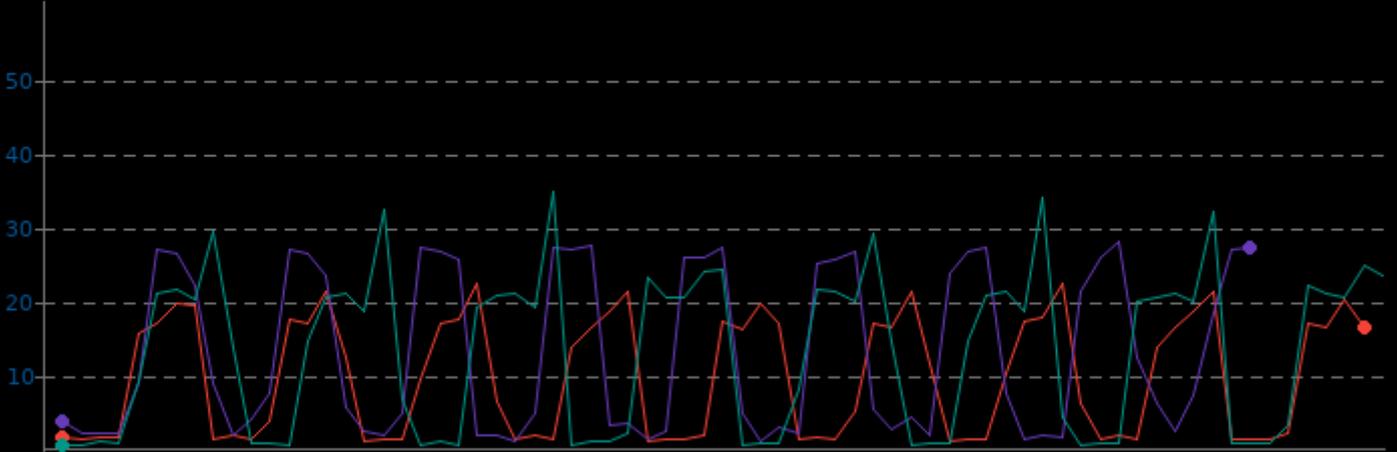


clpeak

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.3	10.0	22.5
Gen12 Xe - i7-1165G7	1.4	13.3	28.1
Gen9 - i7-8565U	0.7	13.4	34.7

▼ Watts, Fewer Is Better

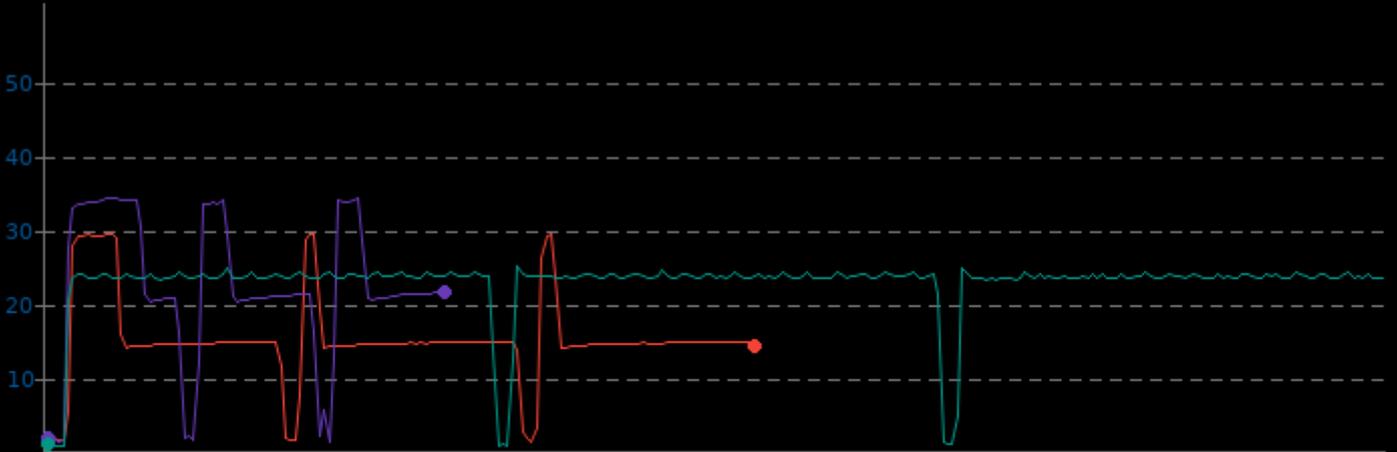


MandelGPU 1.3pts1

CPU Power Consumption Monitor

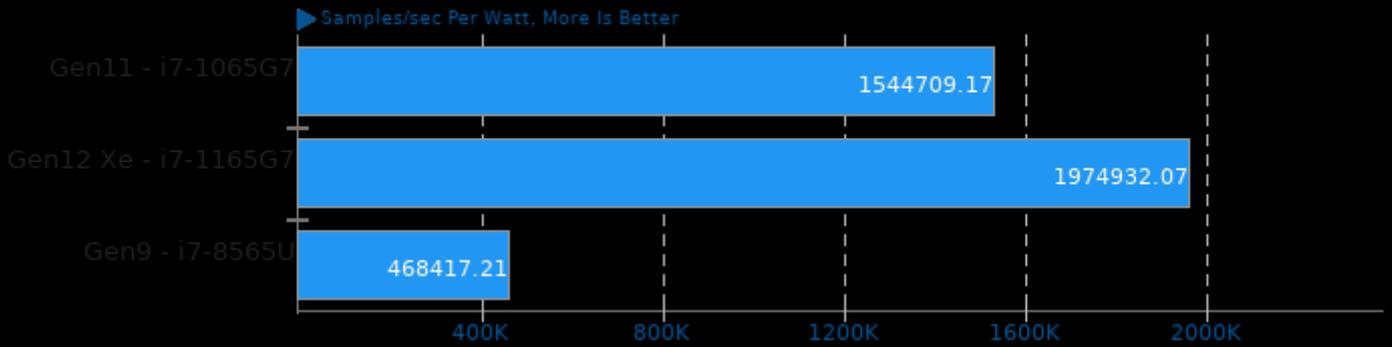
	Min	Avg	Max
Gen11 - i7-1065G7	1.7	15.3	29.6
Gen12 Xe - i7-1165G7	1.6	22.6	34.3
Gen9 - i7-8565U	1.0	22.8	25.1

▼ Watts, Fewer Is Better



MandelGPU 1.3pts1

OpenCL Device: GPU

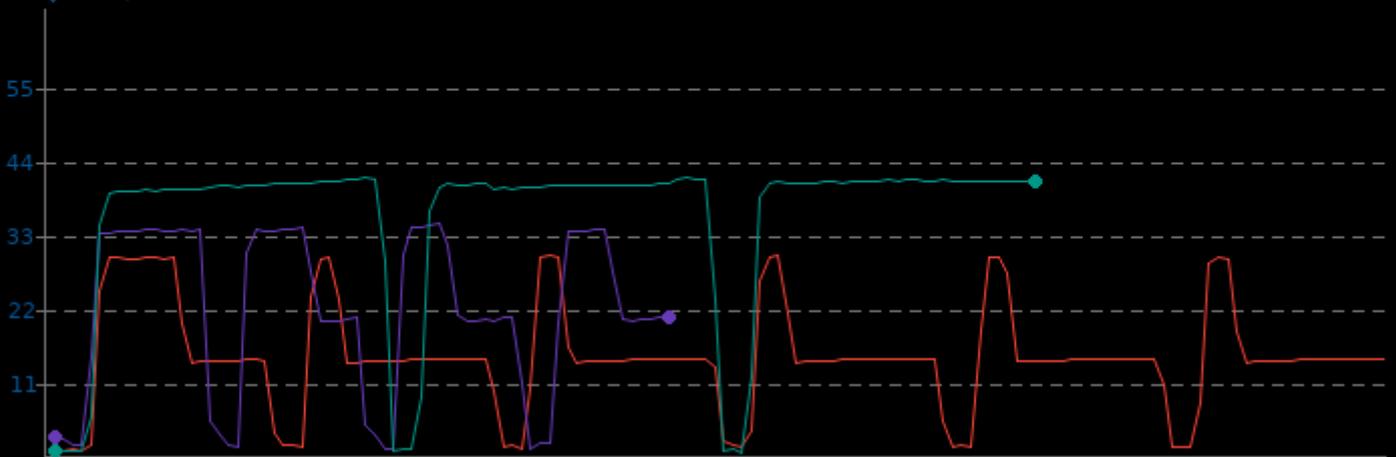


JuliaGPU 1.2pts1

CPU Power Consumption Monitor

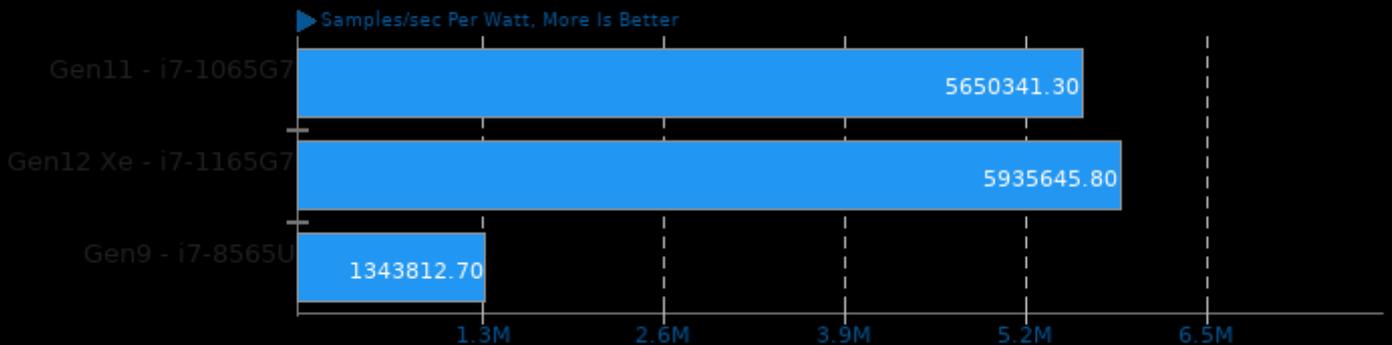
	Min	Avg	Max
Gen11 - i7-1065G7	1.1	15.1	30.0
Gen12 Xe - i7-1165G7	1.5	22.3	34.7
Gen9 - i7-8565U	1.0	35.7	41.4

Watts, Fewer Is Better



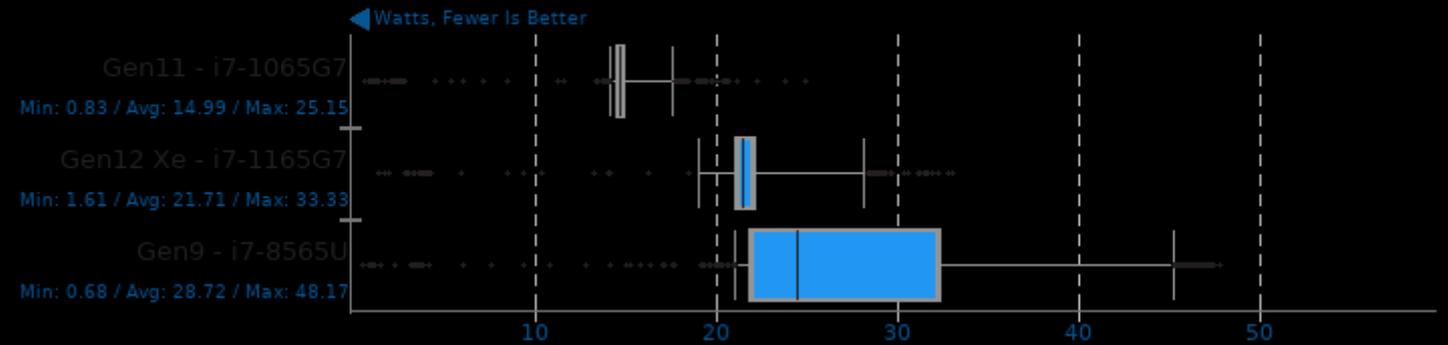
JuliaGPU 1.2pts1

OpenCL Device: GPU



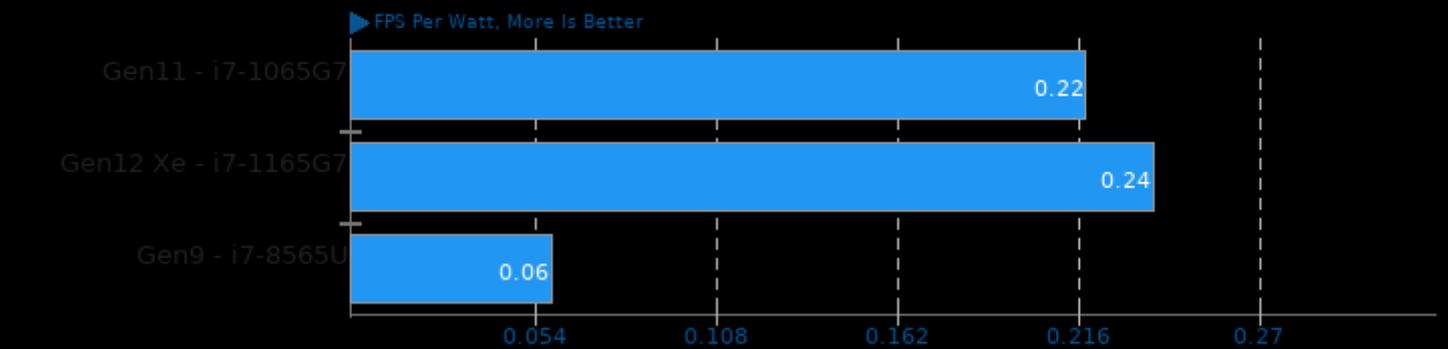
PlaidML

CPU Power Consumption Monitor



PlaidML

FP16: No - Mode: Inference - Network: NASNet Large - Device: OpenCL



PlaidML

CPU Power Consumption Monitor

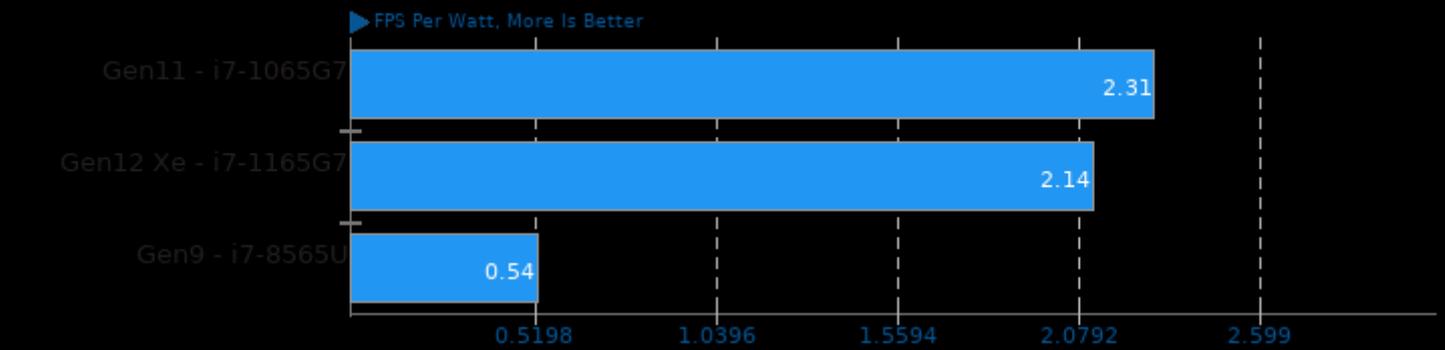
	Min	Avg	Max
Gen11 - i7-1065G7	1.0	15.1	25.7
Gen12 Xe - i7-1165G7	1.9	21.7	33.4
Gen9 - i7-8565U	0.7	29.6	46.6

Watts, Fewer Is Better



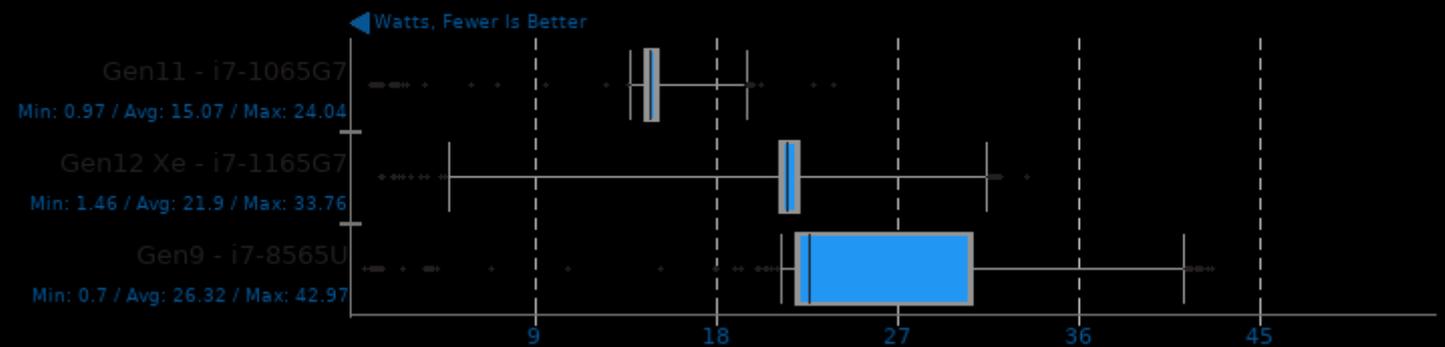
PlaidML

FP16: No - Mode: Inference - Network: Inception V3 - Device: OpenCL



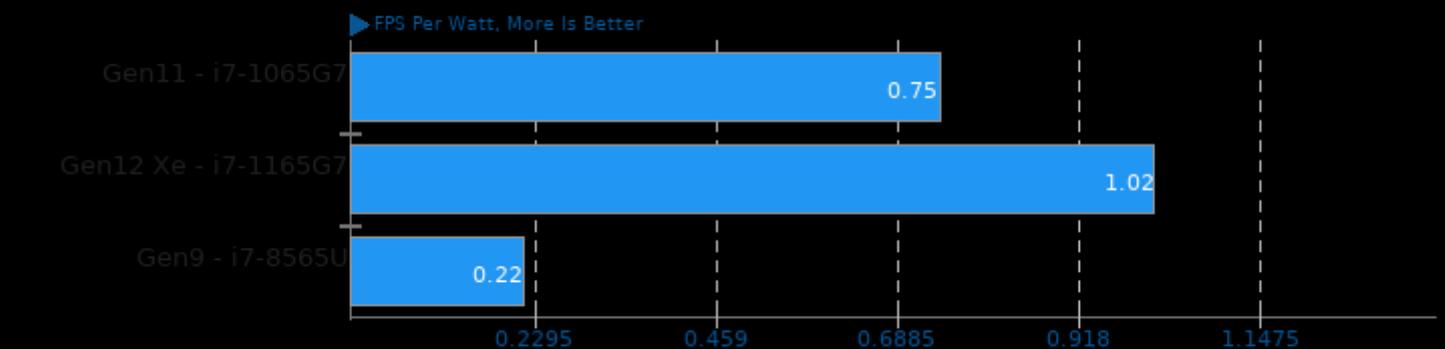
PlaidML

CPU Power Consumption Monitor



PlaidML

FP16: No - Mode: Inference - Network: DenseNet 201 - Device: OpenCL

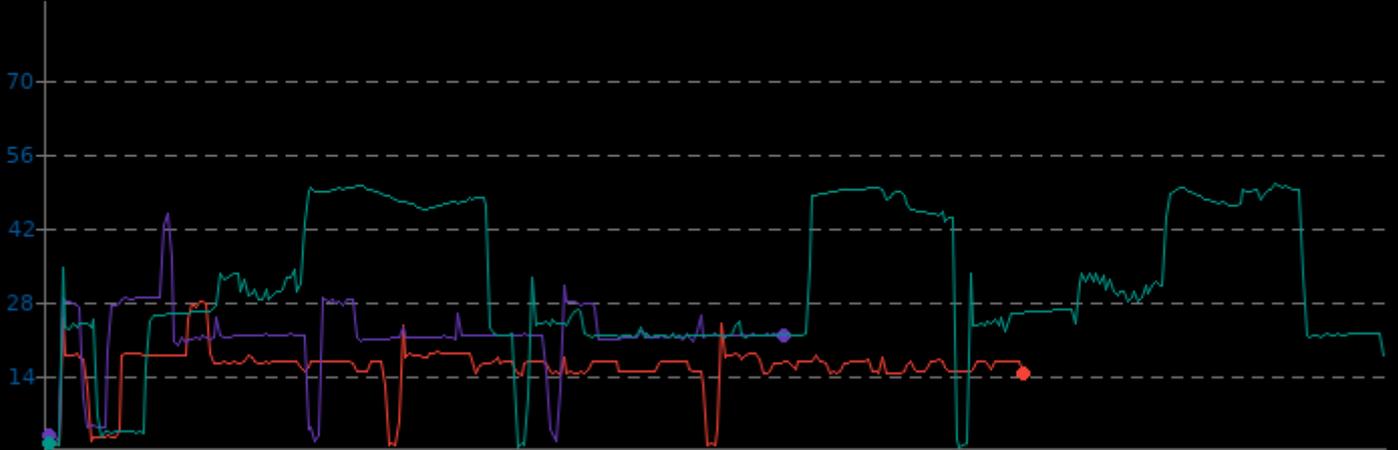


PlaidML

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	15.8	28.3
Gen12 Xe - i7-1165G7	1.5	21.6	44.6
Gen9 - i7-8565U	0.6	31.5	50.1

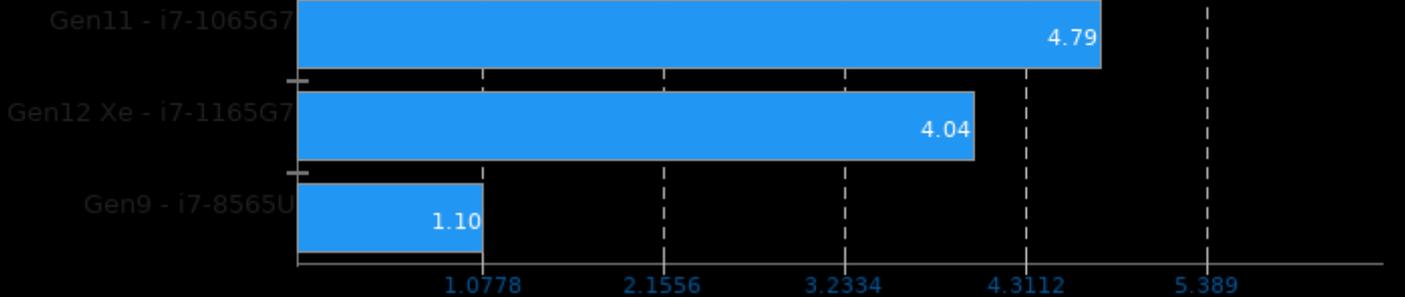
▼ Watts, Fewer Is Better



PlaidML

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

▶ FPS Per Watt, More Is Better

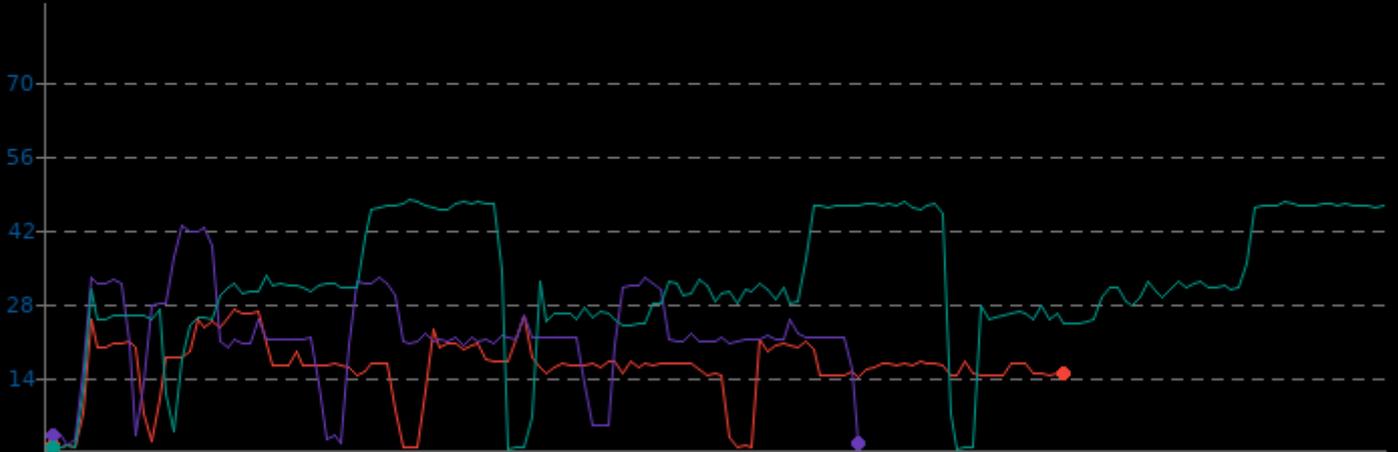


PlaidML

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.1	16.1	27.0
Gen12 Xe - i7-1165G7	1.7	22.0	42.9
Gen9 - i7-8565U	0.7	32.0	47.5

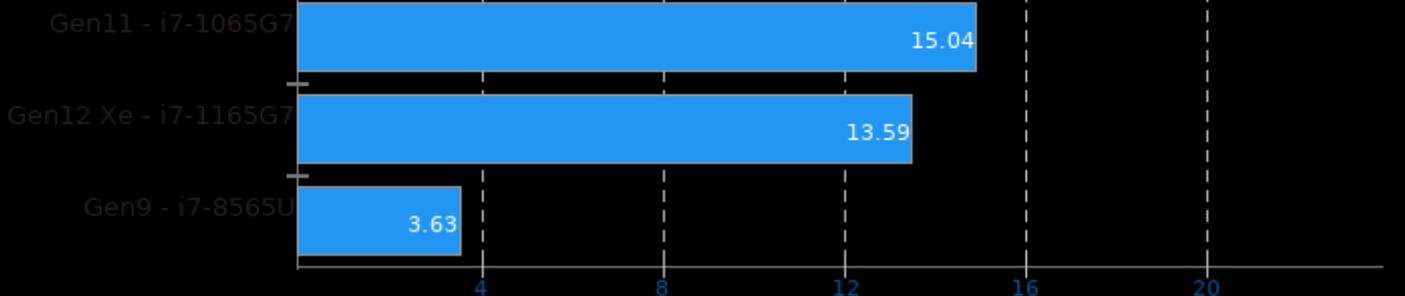
▼ Watts, Fewer Is Better



PlaidML

FP16: No - Mode: Inference - Network: Mobilenet - Device: OpenCL

▶ FPS Per Watt, More Is Better

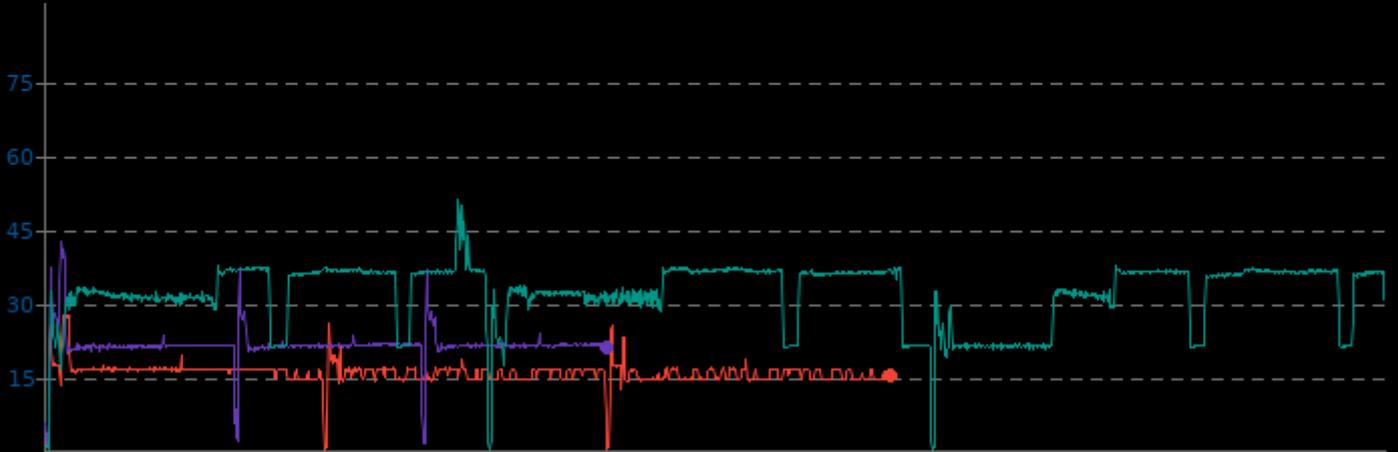


PlaidML

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	16.2	27.9
Gen12 Xe - i7-1165G7	1.7	21.8	42.4
Gen9 - i7-8565U	0.9	32.2	51.2

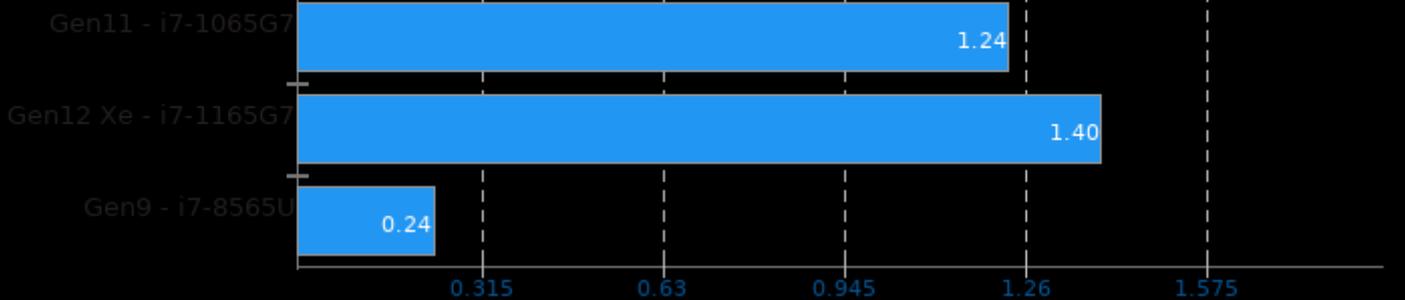
▼ Watts, Fewer Is Better



PlaidML

FP16: No - Mode: Inference - Network: IMDB LSTM - Device: OpenCL

► FPS Per Watt, More Is Better

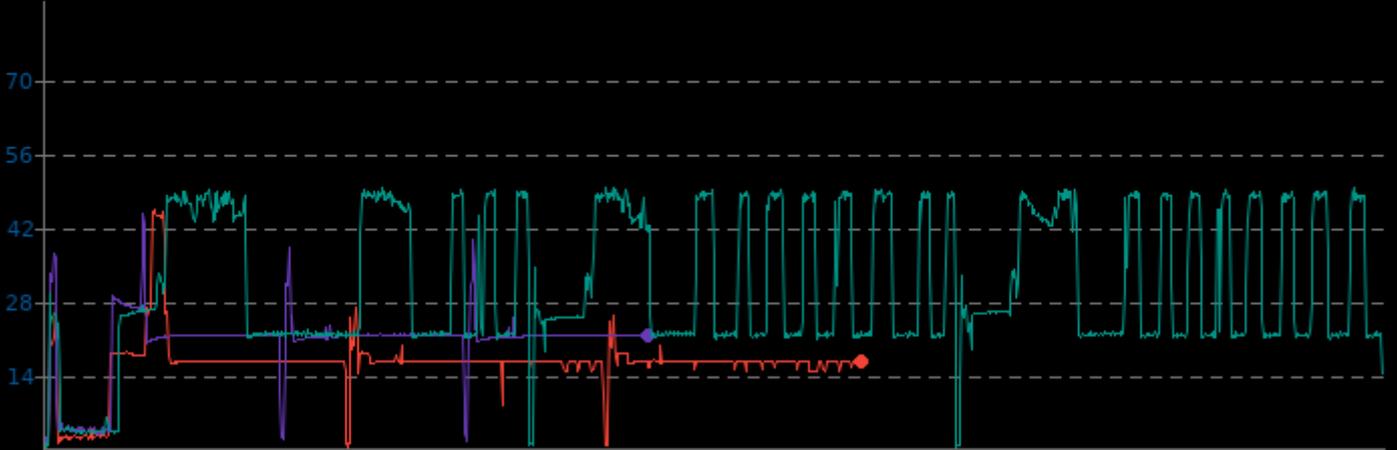


PlaidML

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	0.9	16.4	45.3
Gen12 Xe - i7-1165G7	1.8	20.4	44.6
Gen9 - i7-8565U	0.7	31.2	49.6

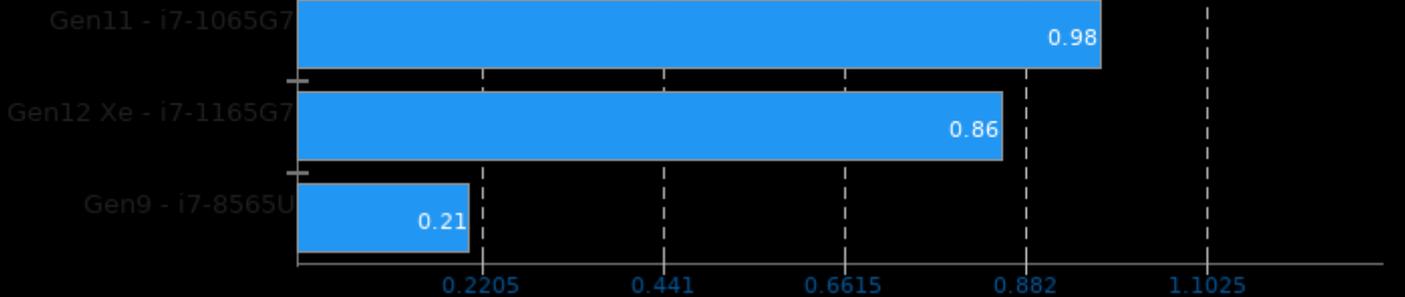
▼ Watts, Fewer Is Better



PlaidML

FP16: No - Mode: Inference - Network: VGG19 - Device: OpenCL

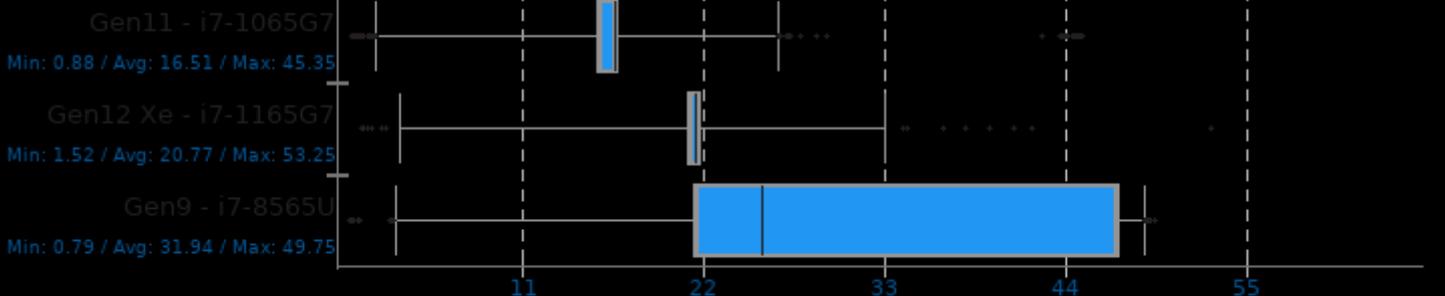
► FPS Per Watt, More Is Better



PlaidML

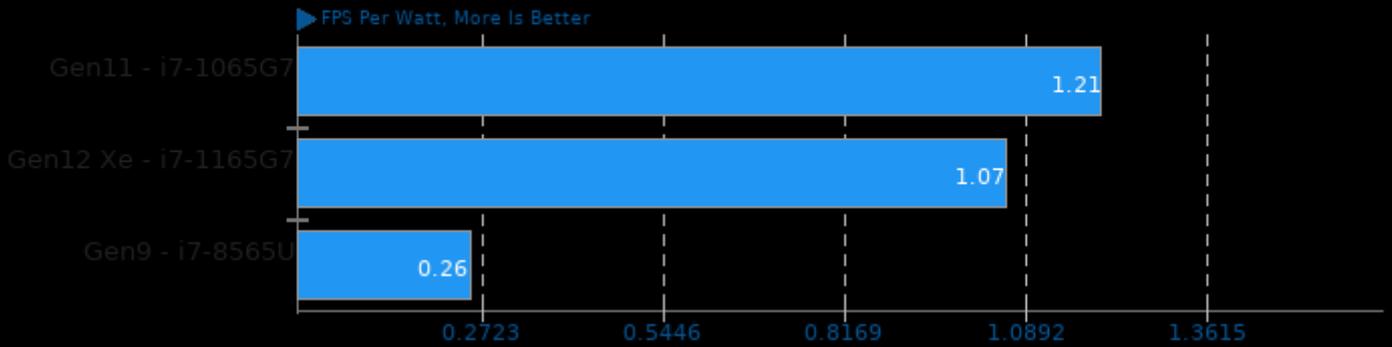
CPU Power Consumption Monitor

◀ Watts, Fewer Is Better



PlaidML

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



NCNN 20200916

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	0.9	15.5	24.0
Gen12 Xe - i7-1165G7	1.1	17.7	26.2
Gen9 - i7-8565U	0.8	20.4	35.3

Watts, Fewer Is Better



LeelaChessZero 0.26 CPU Power Consumption Monitor

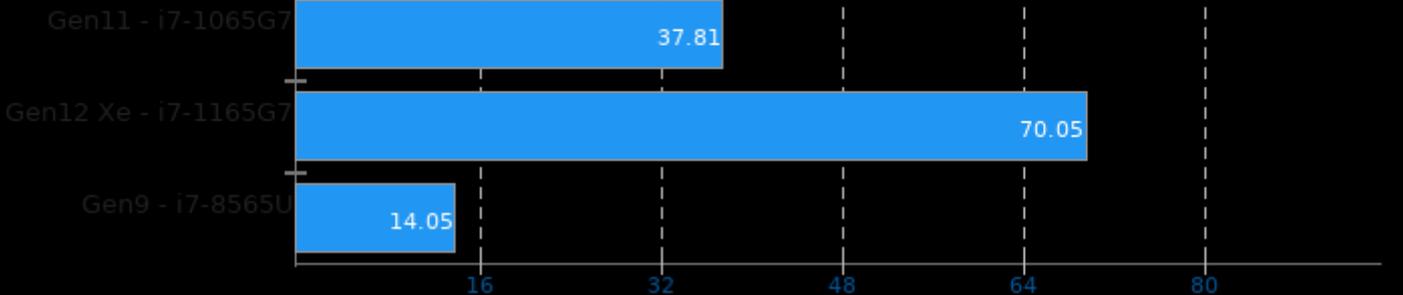
	Min	Avg	Max
Gen11 - i7-1065G7	0.9	15.3	27.3
Gen12 Xe - i7-1165G7	1.8	21.9	35.2
Gen9 - i7-8565U	0.7	27.7	38.7

▼ Watts, Fewer Is Better



LeelaChessZero 0.26 Backend: OpenCL

► Nodes Per Second Per Watt, More Is Better



cl-mem 2017-01-13

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	0.6	14.9	31.6
Gen12 Xe - i7-1165G7	1.5	23.0	38.2
Gen9 - i7-8565U	0.7	19.9	22.9

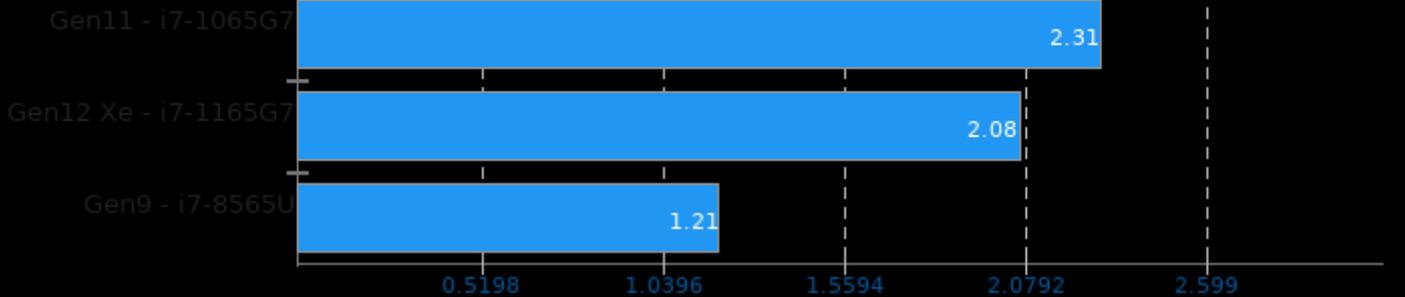
▼ Watts, Fewer Is Better



cl-mem 2017-01-13

Benchmark: Write

▶ GB/s Per Watt, More Is Better



cl-mem 2017-01-13

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	15.3	31.4
Gen12 Xe - i7-1165G7	1.8	23.3	38.0
Gen9 - i7-8565U	0.7	19.8	23.0

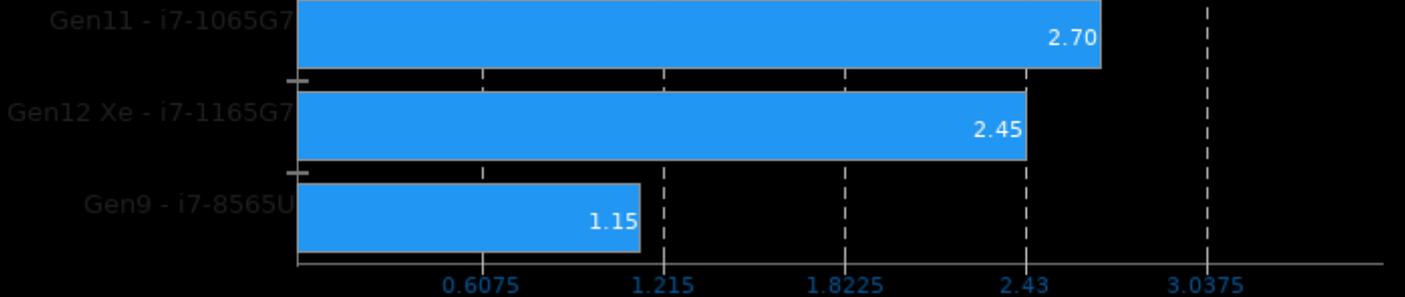
▼ Watts, Fewer Is Better



cl-mem 2017-01-13

Benchmark: Read

► GB/s Per Watt, More Is Better

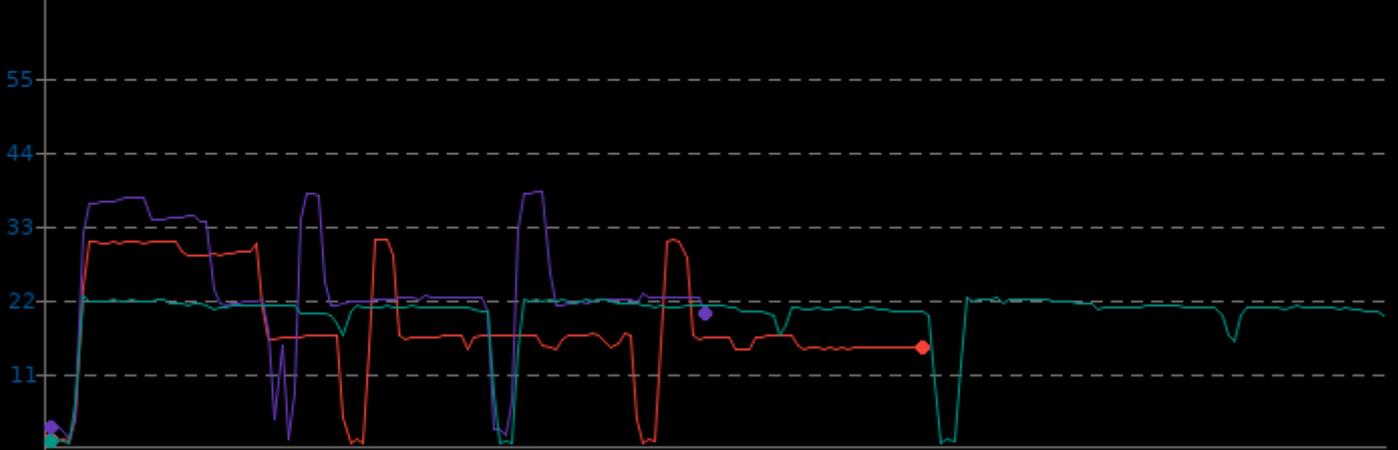


cl-mem 2017-01-13

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	18.4	31.0
Gen12 Xe - i7-1165G7	1.3	23.8	38.0
Gen9 - i7-8565U	0.8	19.8	22.7

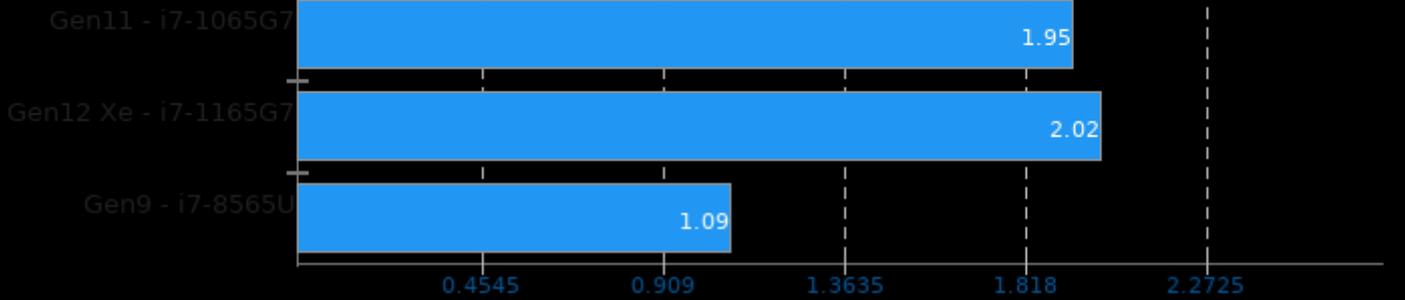
▼ Watts, Fewer Is Better

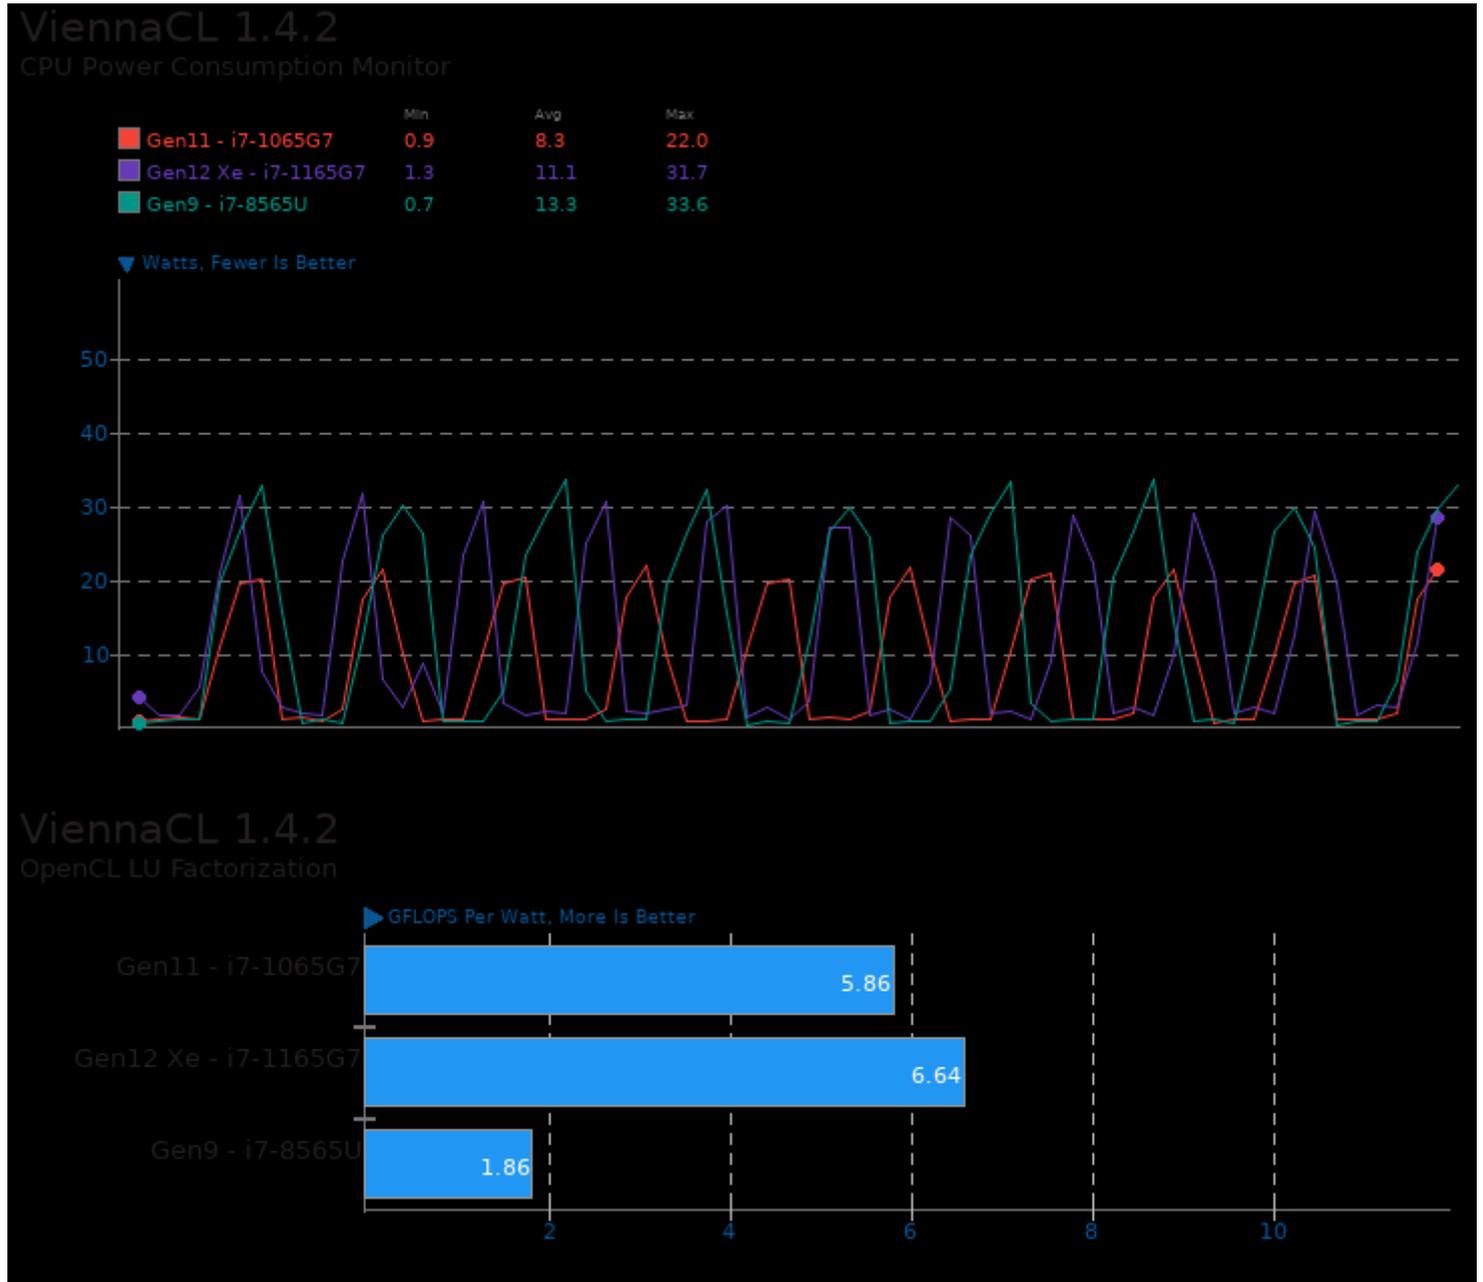


cl-mem 2017-01-13

Benchmark: Copy

▶ GB/s Per Watt, More Is Better



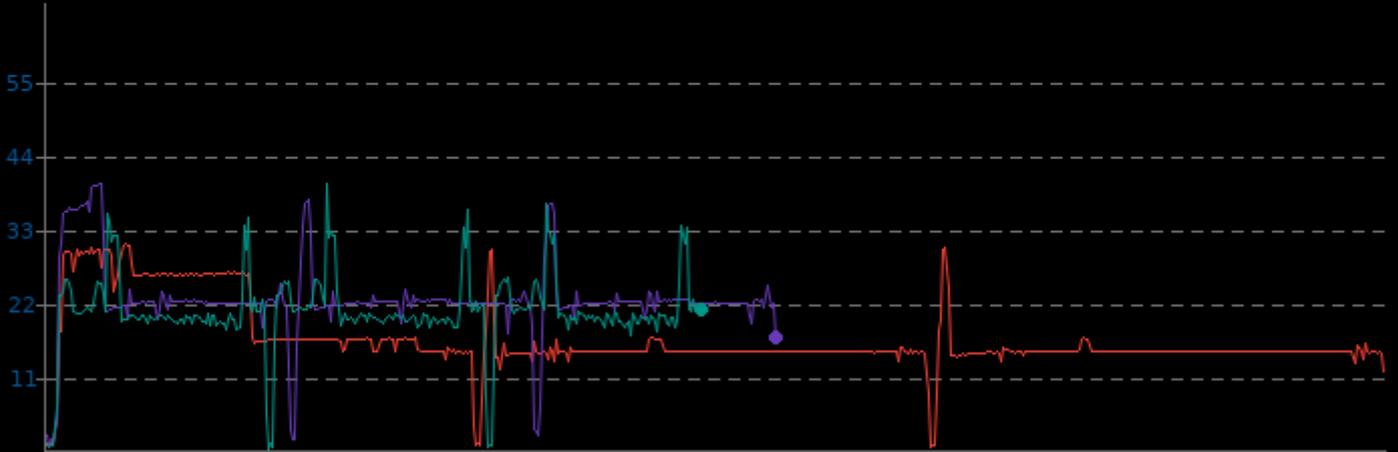


SHOC Scalable Heterogeneous Computing 2015-11-10

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	16.8	30.9
Gen12 Xe - i7-1165G7	1.1	22.7	39.9
Gen9 - i7-8565U	0.7	21.0	39.7

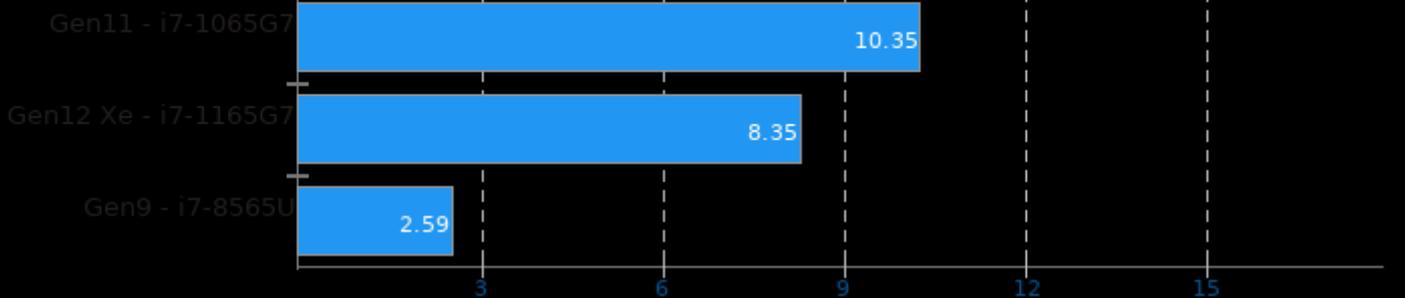
▼ Watts, Fewer Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Texture Read Bandwidth

▶ GB/s Per Watt, More Is Better

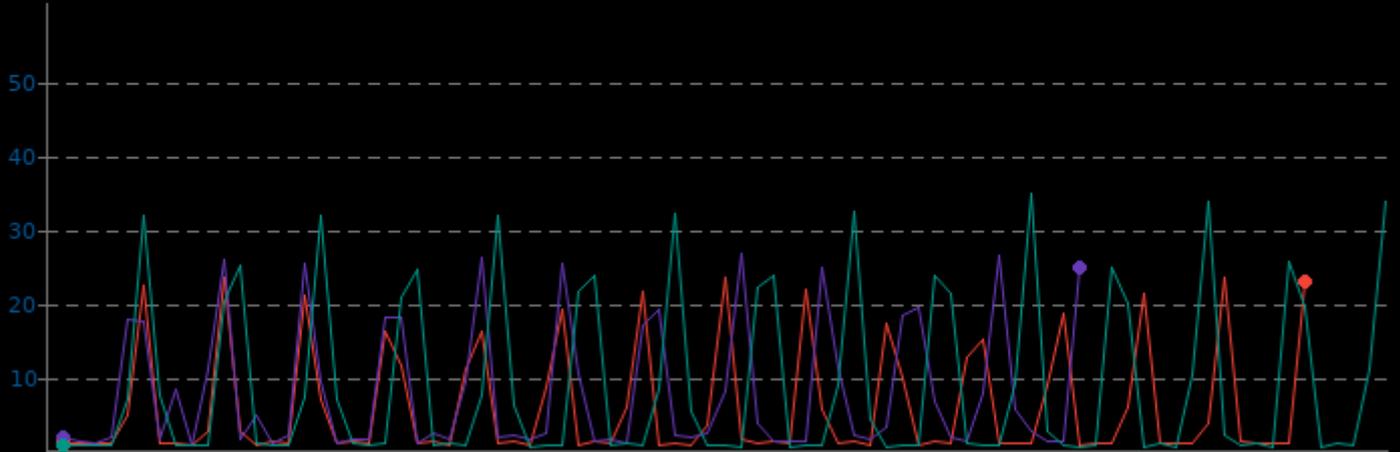


SHOC Scalable Heterogeneous Computing 2015-11-10

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	6.1	23.7
Gen12 Xe - i7-1165G7	1.2	8.2	26.7
Gen9 - i7-8565U	0.7	8.9	34.7

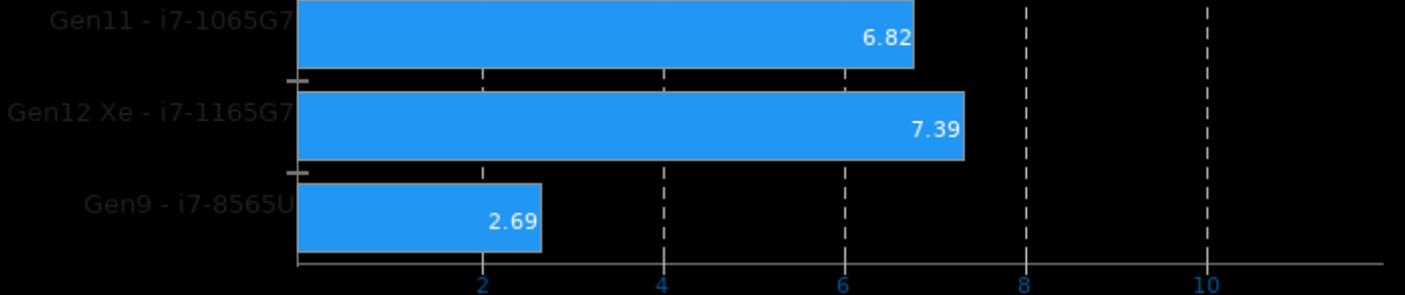
▼ Watts, Fewer Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Bus Speed Readback

▶ GB/s Per Watt, More Is Better

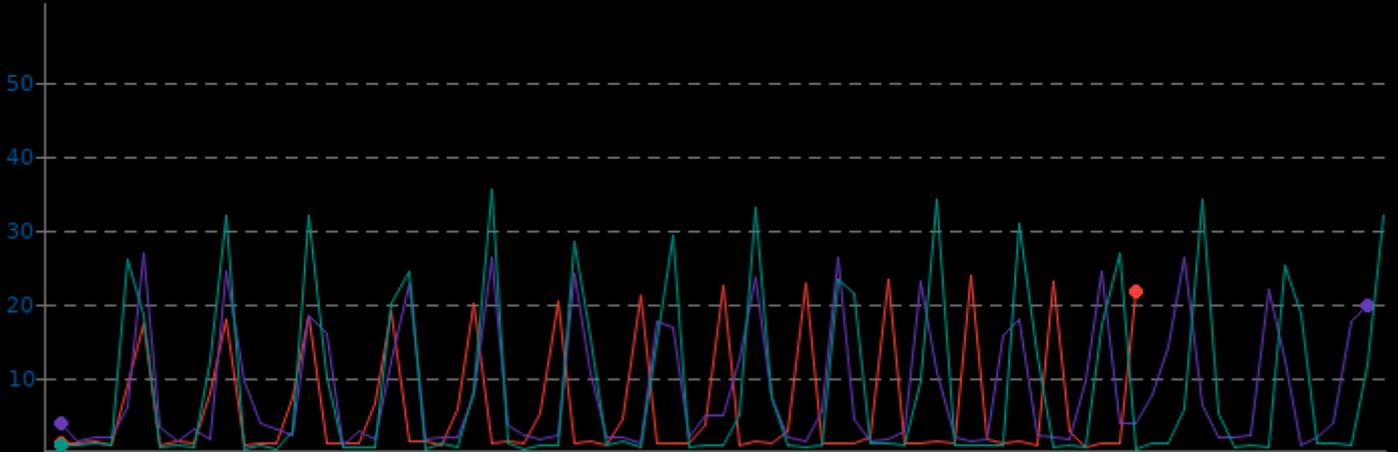


SHOC Scalable Heterogeneous Computing 2015-11-10

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	0.9	5.9	23.7
Gen12 Xe - i7-1165G7	1.0	8.3	26.9
Gen9 - i7-8565U	0.3	8.8	35.3

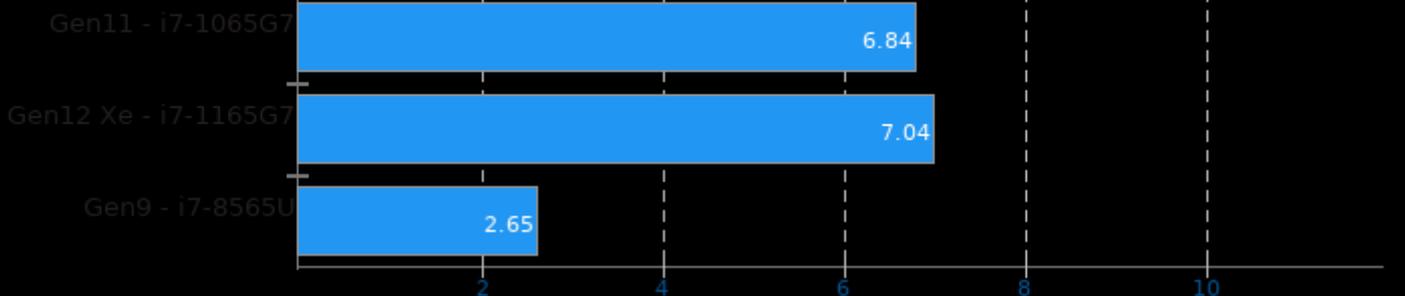
▼ Watts, Fewer Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Bus Speed Download

▶ GB/s Per Watt, More Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.1	20.1	32.3
Gen12 Xe - i7-1165G7	1.7	22.8	37.3
Gen9 - i7-8565U	0.4	23.1	33.5

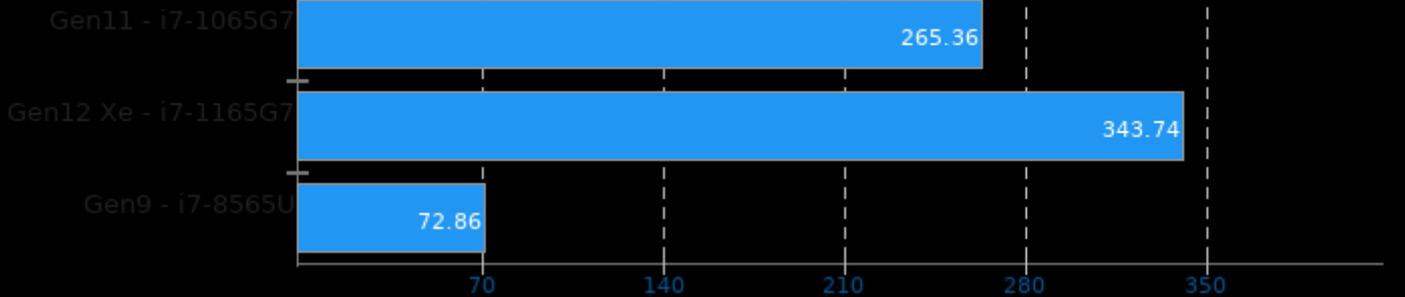
▼ Watts, Fewer Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Max SP Flops

▶ GFLOPS Per Watt, More Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	25.1	28.4
Gen12 Xe - i7-1165G7	1.0	23.8	42.3
Gen9 - i7-8565U	0.9	21.9	24.7

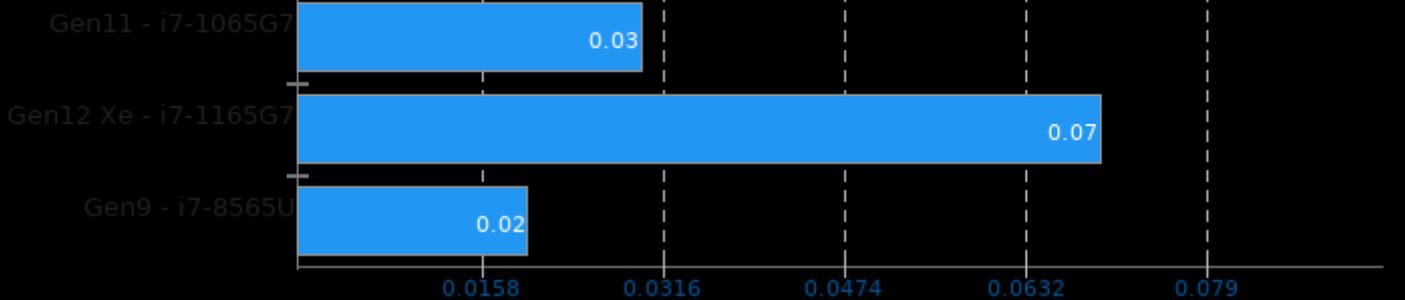
▼ Watts, Fewer Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: MD5 Hash

▶ GHash/s Per Watt, More Is Better

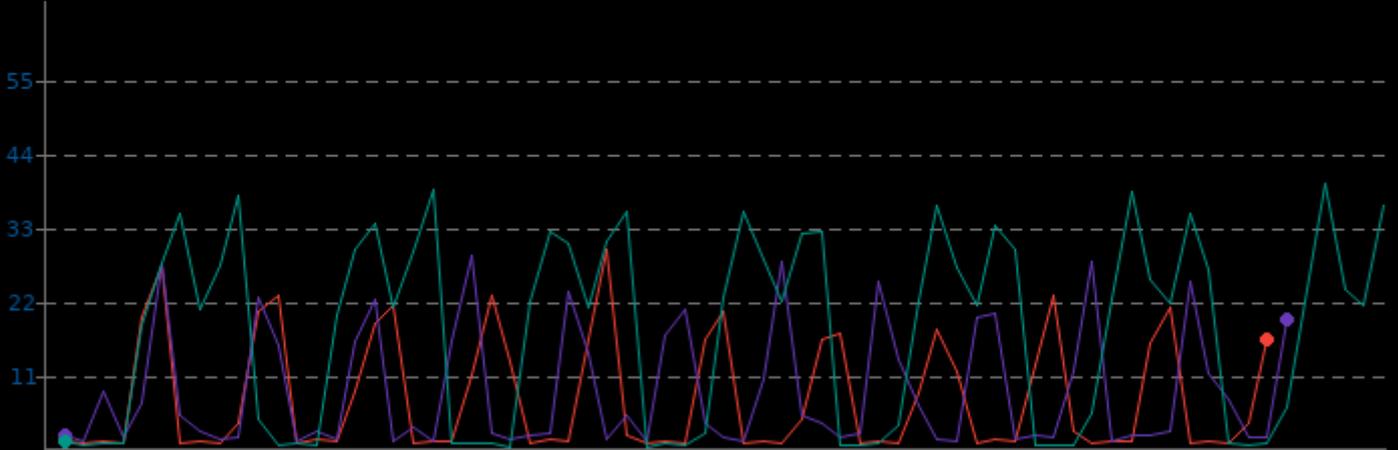


SHOC Scalable Heterogeneous Computing 2015-11-10

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.1	7.9	29.6
Gen12 Xe - i7-1165G7	1.3	8.8	28.8
Gen9 - i7-8565U	0.7	17.6	39.6

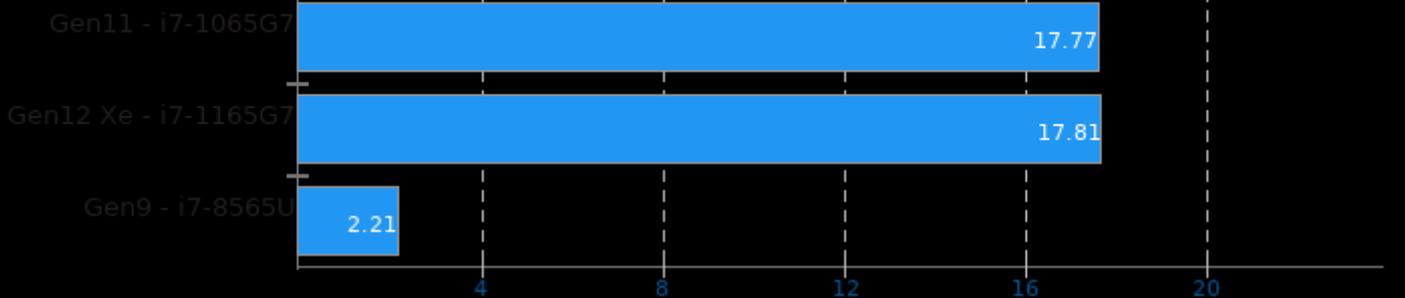
▼ Watts, Fewer Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: FFT SP

▶ GFLOPS Per Watt, More Is Better

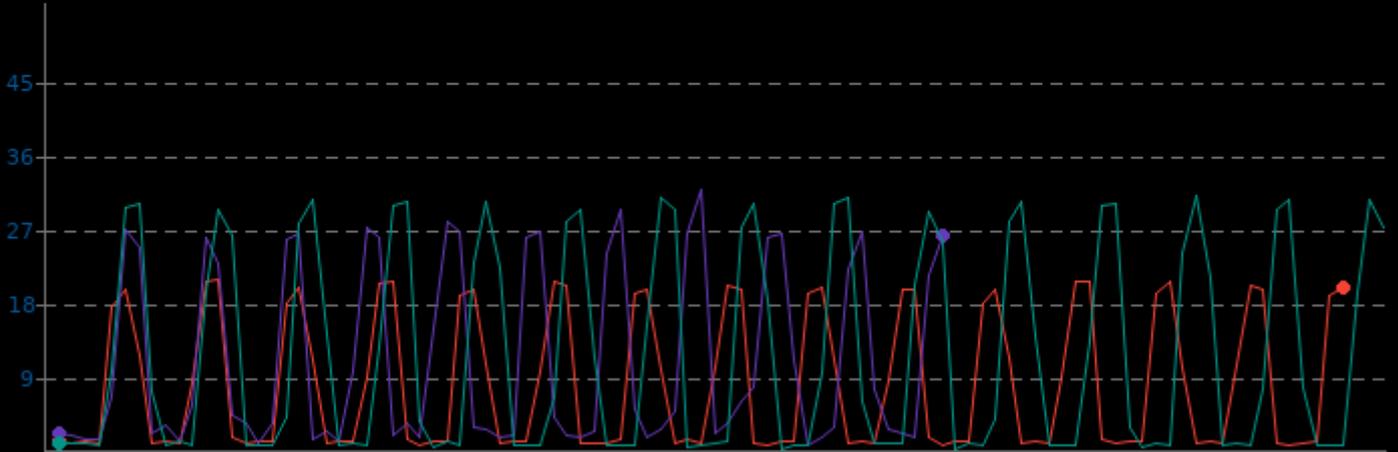


SHOC Scalable Heterogeneous Computing 2015-11-10

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	8.3	20.9
Gen12 Xe - i7-1165G7	1.0	11.1	31.8
Gen9 - i7-8565U	0.6	11.9	31.2

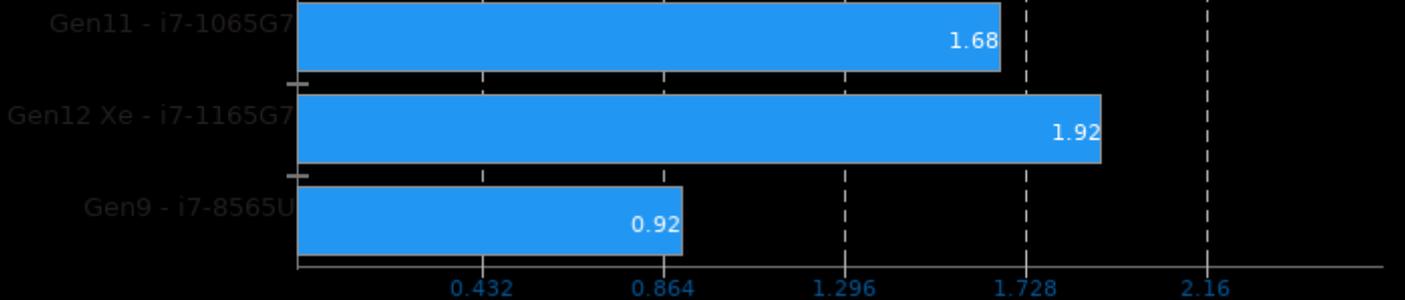
▼ Watts, Fewer Is Better



SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Triad

▶ GB/s Per Watt, More Is Better

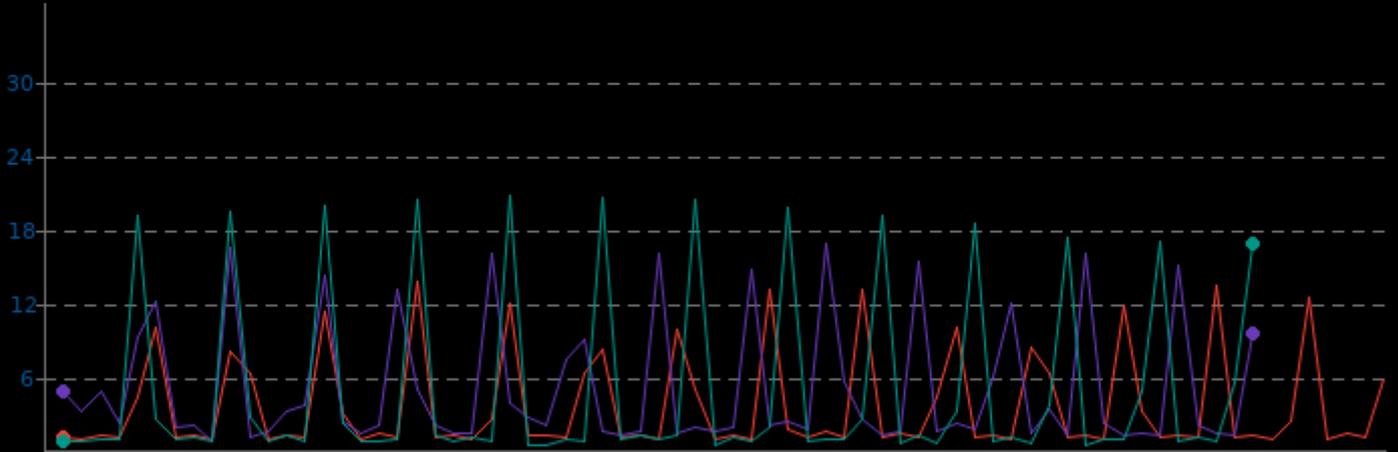


FinanceBench 2016-06-06

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.1	3.8	13.9
Gen12 Xe - i7-1165G7	1.0	5.1	16.9
Gen9 - i7-8565U	0.7	5.0	20.7

▼ Watts, Fewer Is Better

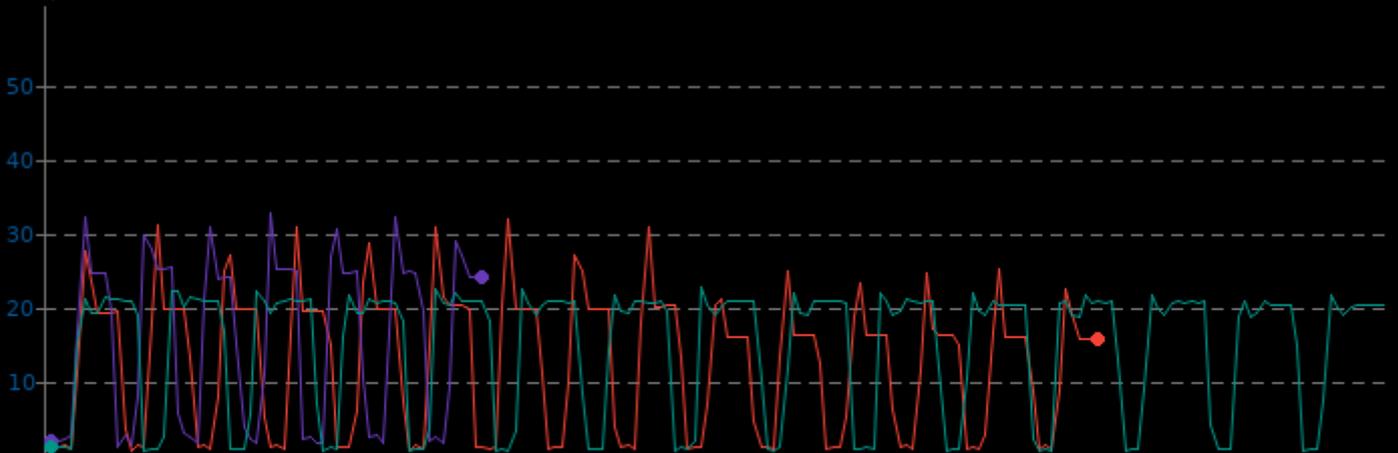


FinanceBench 2016-06-06

CPU Power Consumption Monitor

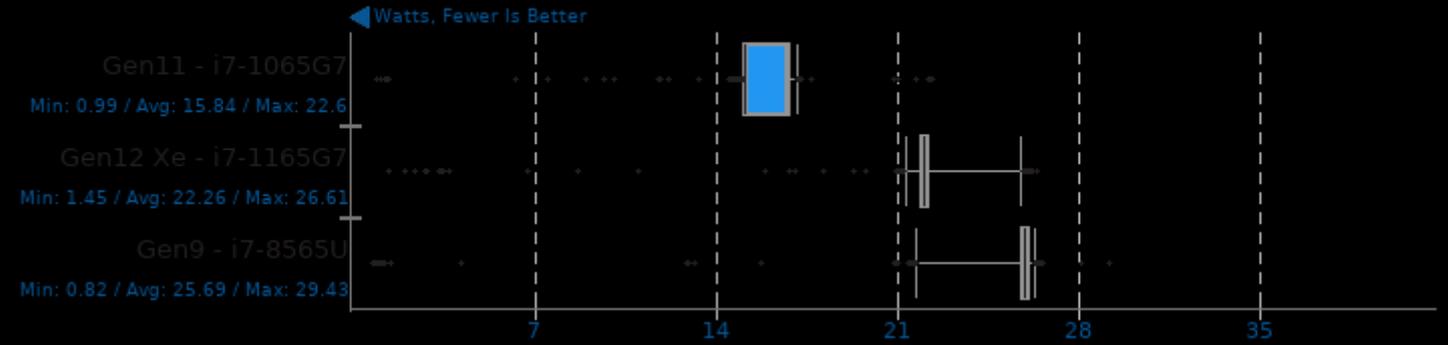
	Min	Avg	Max
Gen11 - i7-1065G7	0.9	12.9	31.9
Gen12 Xe - i7-1165G7	1.3	15.9	32.7
Gen9 - i7-8565U	0.7	14.8	22.7

▼ Watts, Fewer Is Better



RealSR-NCNN 20200818

CPU Power Consumption Monitor



RealSR-NCNN 20200818

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	15.5	22.7
Gen12 Xe - i7-1165G7	1.5	22.1	26.6
Gen9 - i7-8565U	0.7	25.7	30.7

Watts, Fewer Is Better

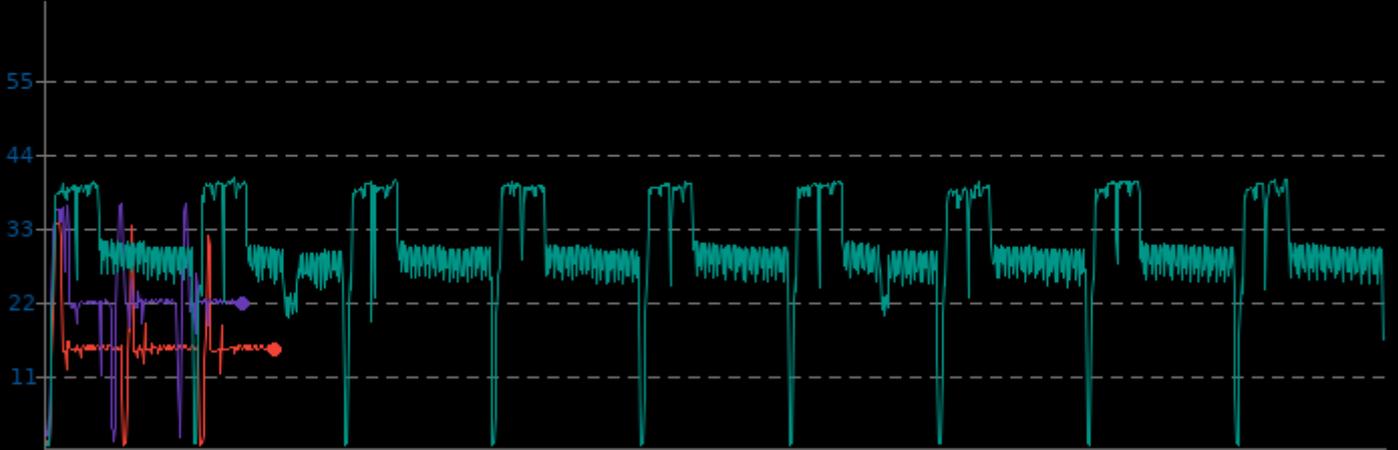


oneAPI Level Zero Tests

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	15.4	33.9
Gen12 Xe - i7-1165G7	1.5	22.4	36.6
Gen9 - i7-8565U	0.8	30.2	40.4

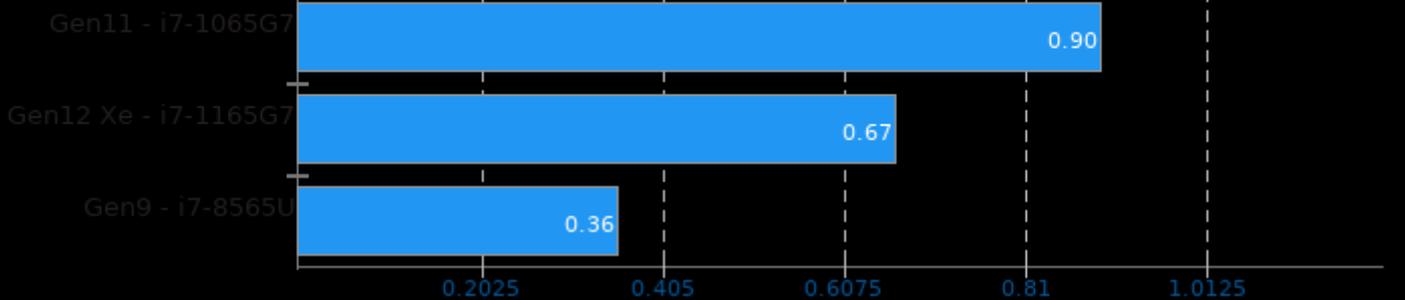
▼ Watts, Fewer Is Better



oneAPI Level Zero Tests

Test: Peak System Memory Copy to Shared Memory

▶ GB/s Per Watt, More Is Better

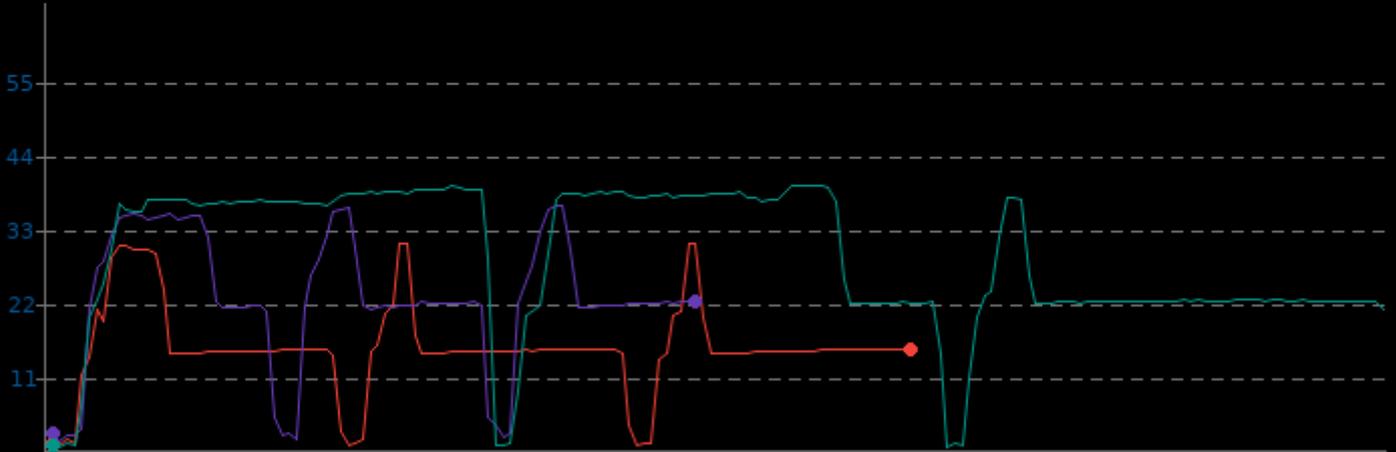


oneAPI Level Zero Tests

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.2	15.5	31.1
Gen12 Xe - i7-1165G7	1.9	22.9	36.7
Gen9 - i7-8565U	0.8	29.0	39.5

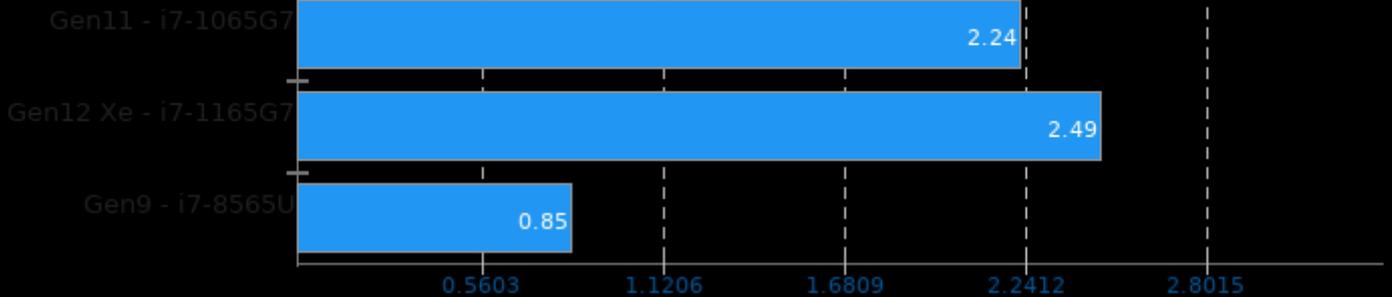
▼ Watts, Fewer Is Better



oneAPI Level Zero Tests

Test: Peak Float16 Global Memory Bandwidth

▶ GB/s Per Watt, More Is Better

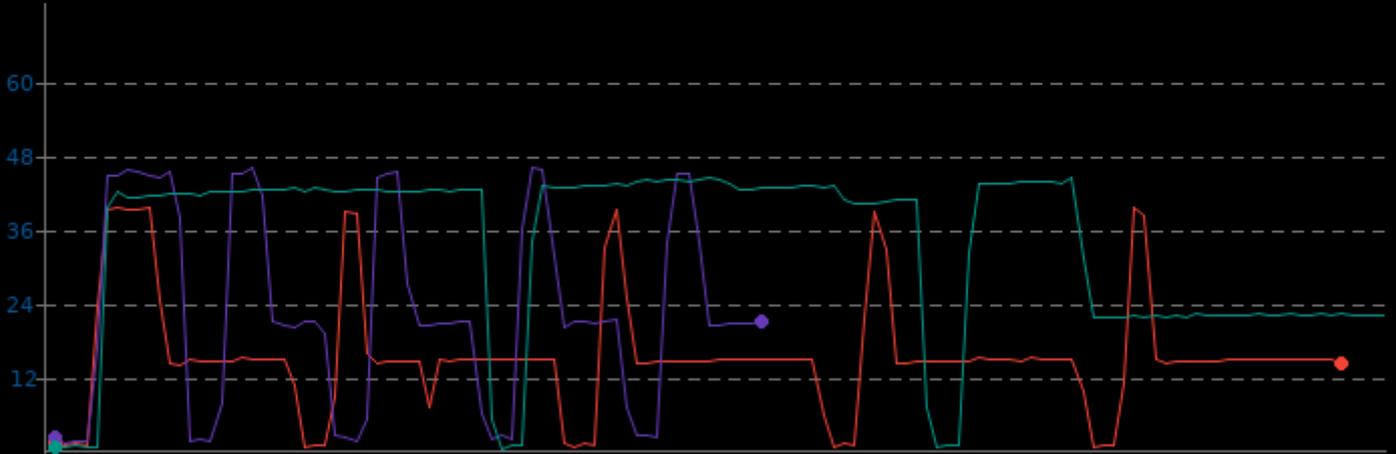


oneAPI Level Zero Tests

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	15.5	39.6
Gen12 Xe - i7-1165G7	1.6	23.1	46.0
Gen9 - i7-8565U	0.7	33.6	44.3

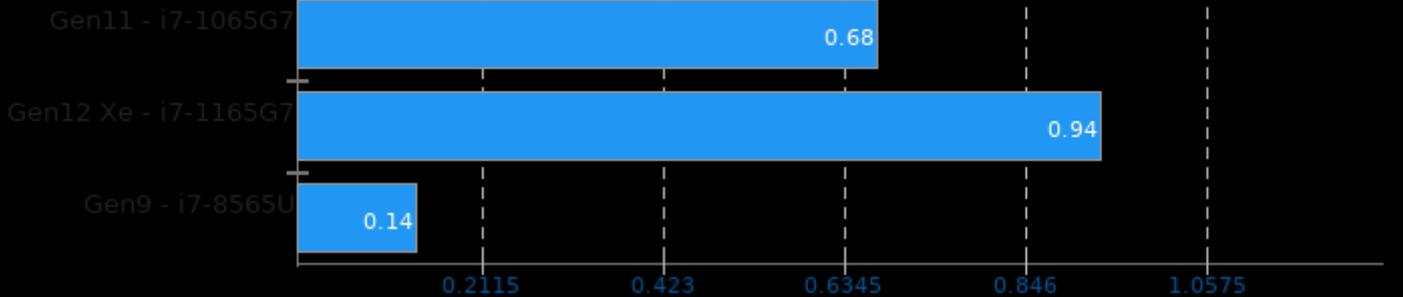
▼ Watts, Fewer Is Better



oneAPI Level Zero Tests

Test: Host-To-Device-To-Host Image Copy

► GB/s Per Watt, More Is Better

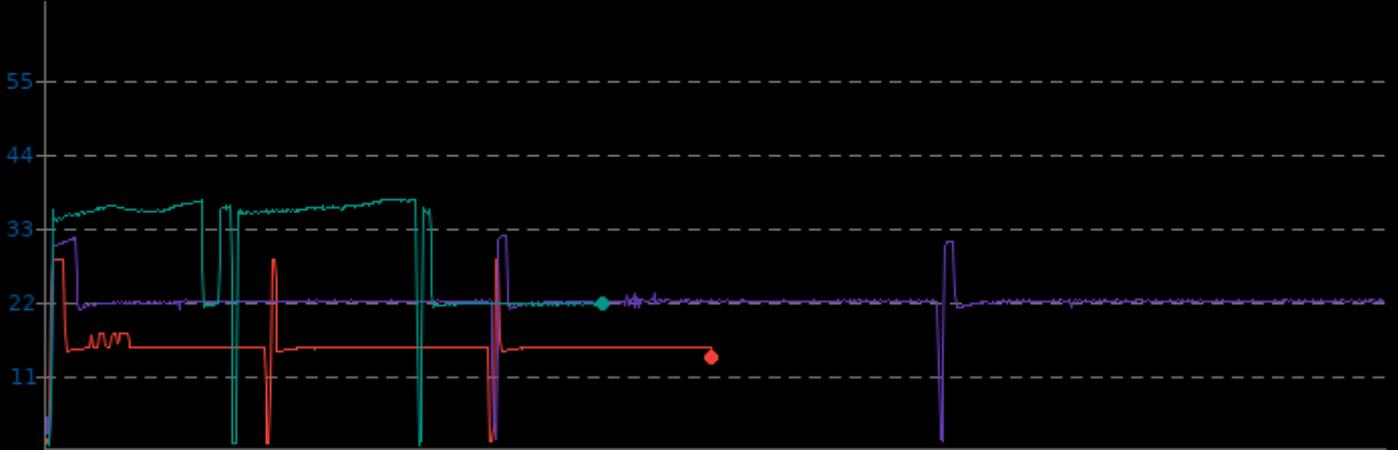


oneAPI Level Zero Tests

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.1	15.5	28.4
Gen12 Xe - i7-1165G7	1.4	22.2	31.9
Gen9 - i7-8565U	0.8	30.1	37.3

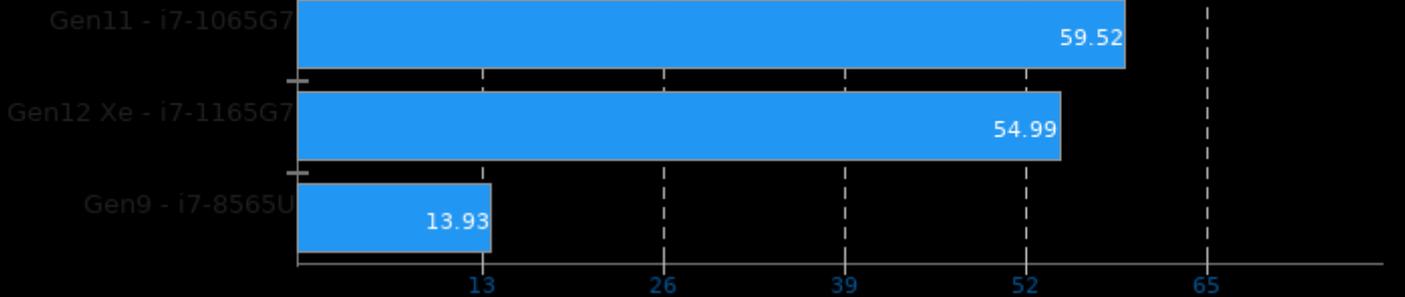
▼ Watts, Fewer Is Better



oneAPI Level Zero Tests

Test: Peak Single-Precision Compute

▶ GB/s Per Watt, More Is Better

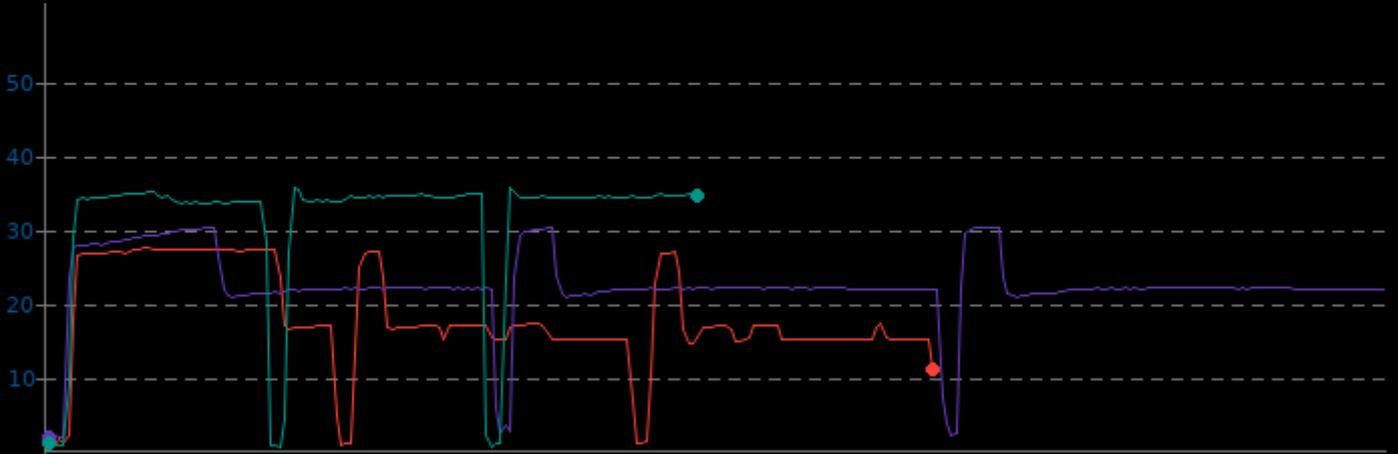


oneAPI Level Zero Tests

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	18.3	27.5
Gen12 Xe - i7-1165G7	2.0	22.4	30.4
Gen9 - i7-8565U	0.8	31.1	35.7

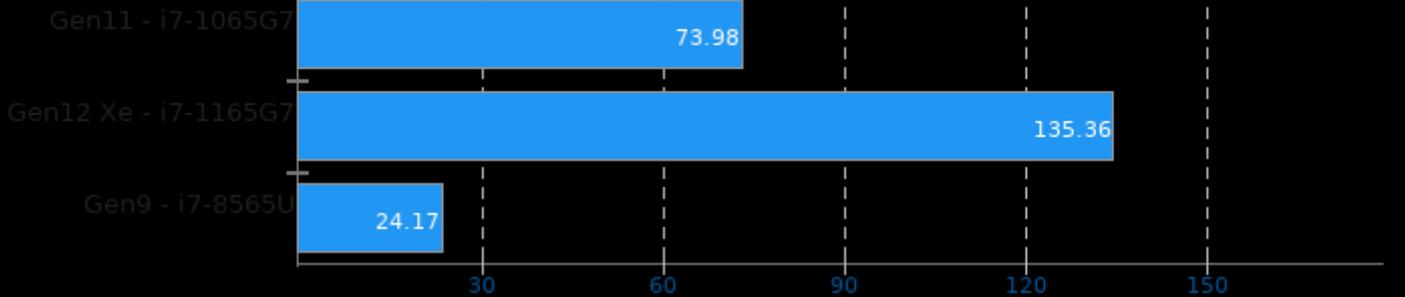
▼ Watts, Fewer Is Better



oneAPI Level Zero Tests

Test: Peak Half-Precision Compute

▶ GFLOPS Per Watt, More Is Better

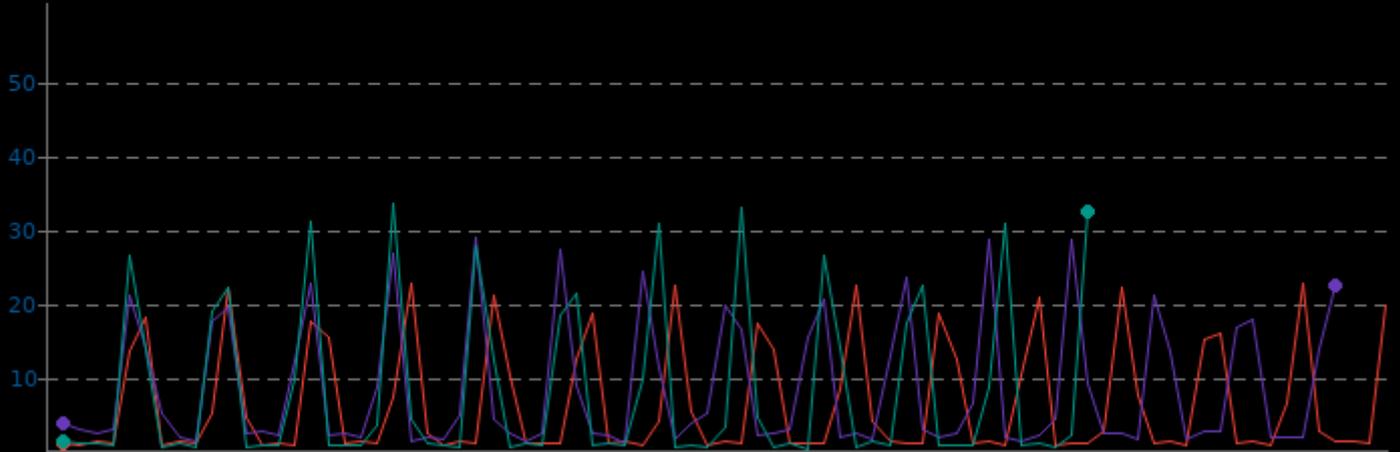


oneAPI Level Zero Tests

CPU Power Consumption Monitor

	Min	Avg	Max
Gen11 - i7-1065G7	1.0	6.6	22.9
Gen12 Xe - i7-1165G7	1.3	8.5	28.9
Gen9 - i7-8565U	0.5	8.3	33.4

▼ Watts, Fewer Is Better

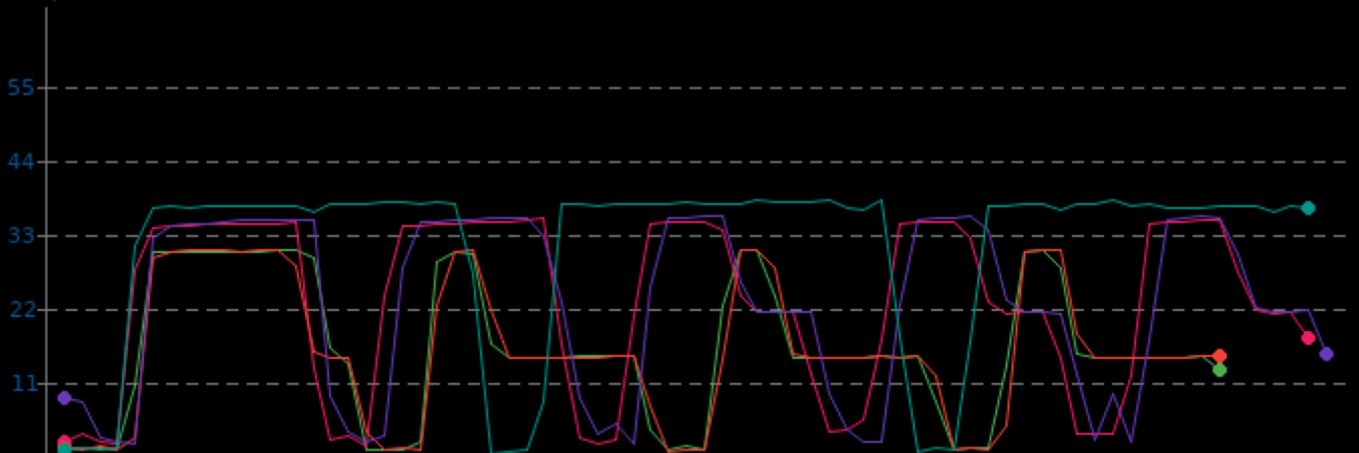


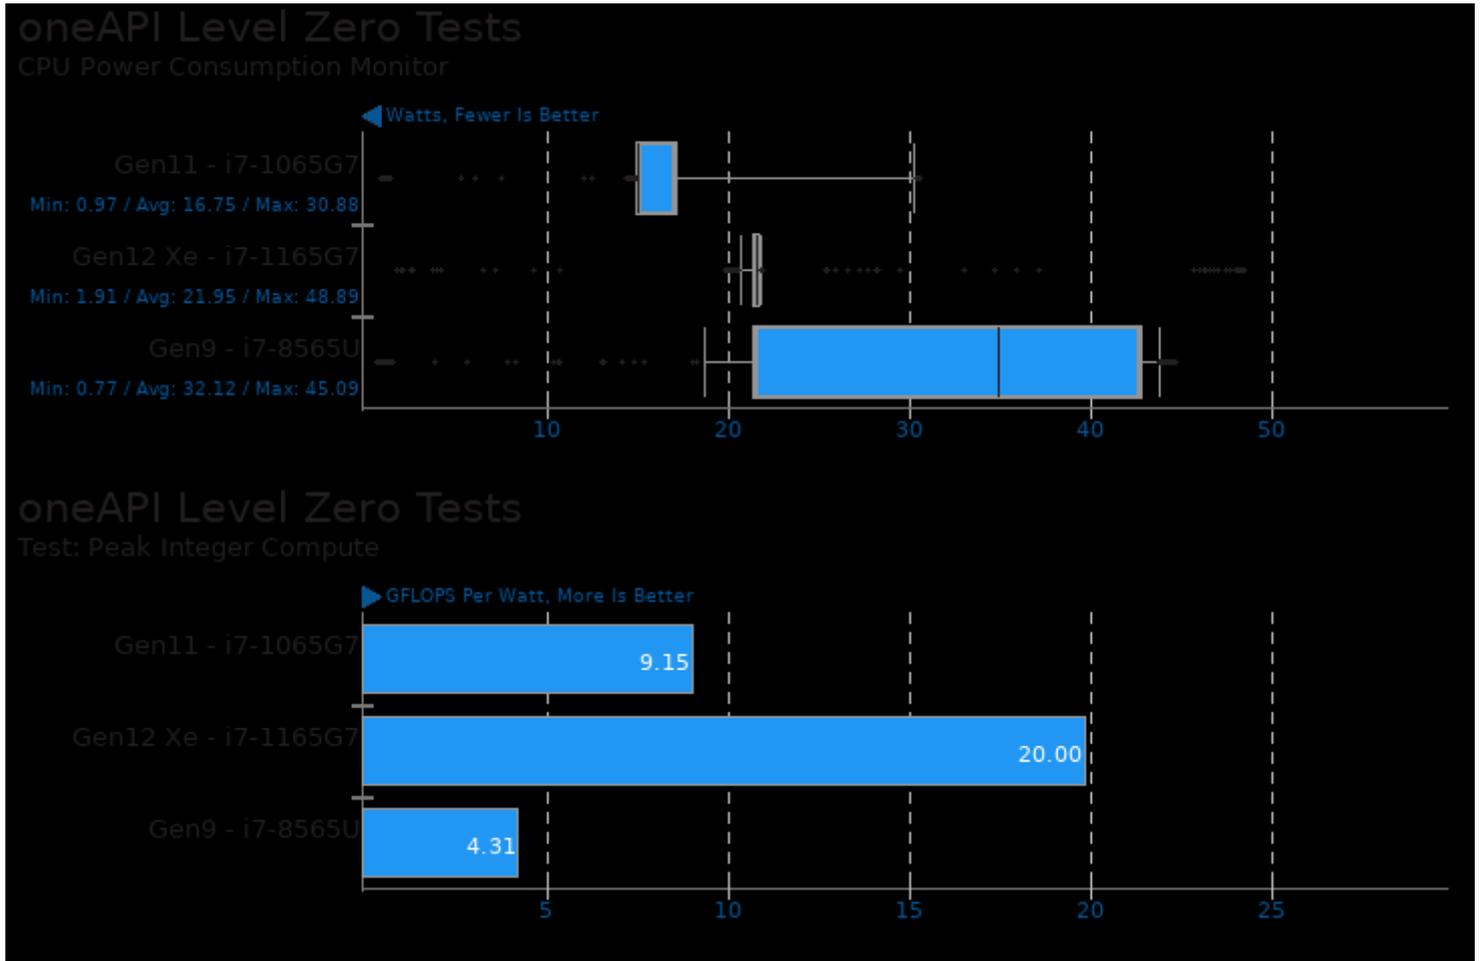
oneAPI Level Zero Tests

CPU Power Consumption Monitor

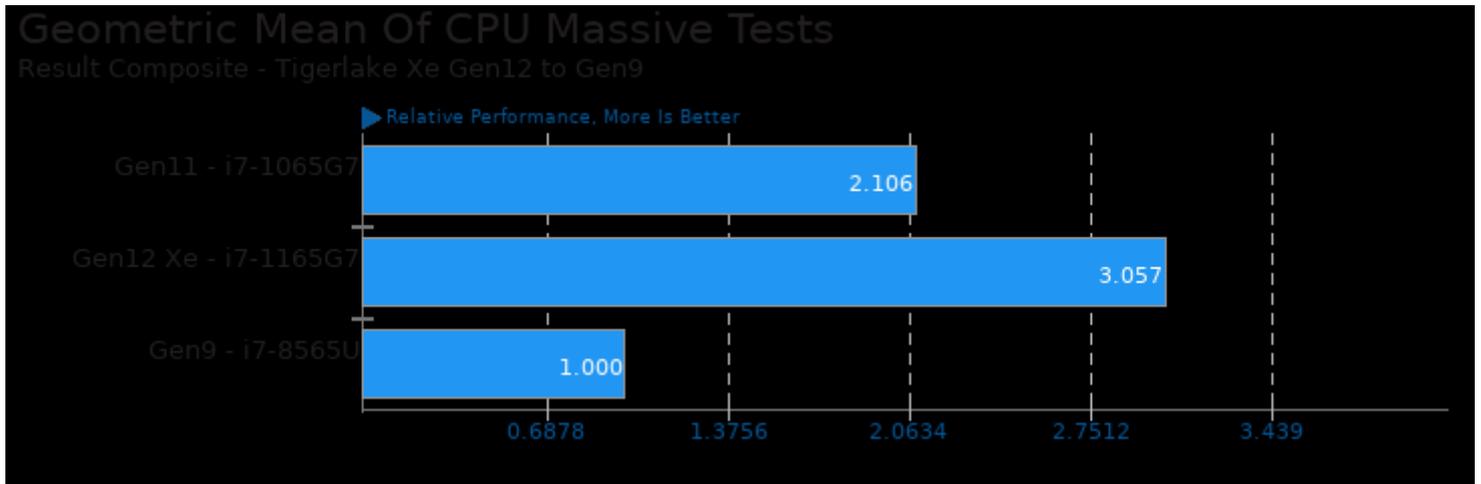
	Min	Avg	Max
11th Gen - i7-1165G7	1.1	16.1	30.7
12th Gen Xe - i7-1165G7	1.7	23.0	35.2
Gen11 - i7-1065G7	1.0	15.7	30.7
Gen12 Xe - i7-1165G7	2.0	22.7	35.7
Gen9 - i7-8565U	0.7	31.0	38.1

▼ Watts, Fewer Is Better

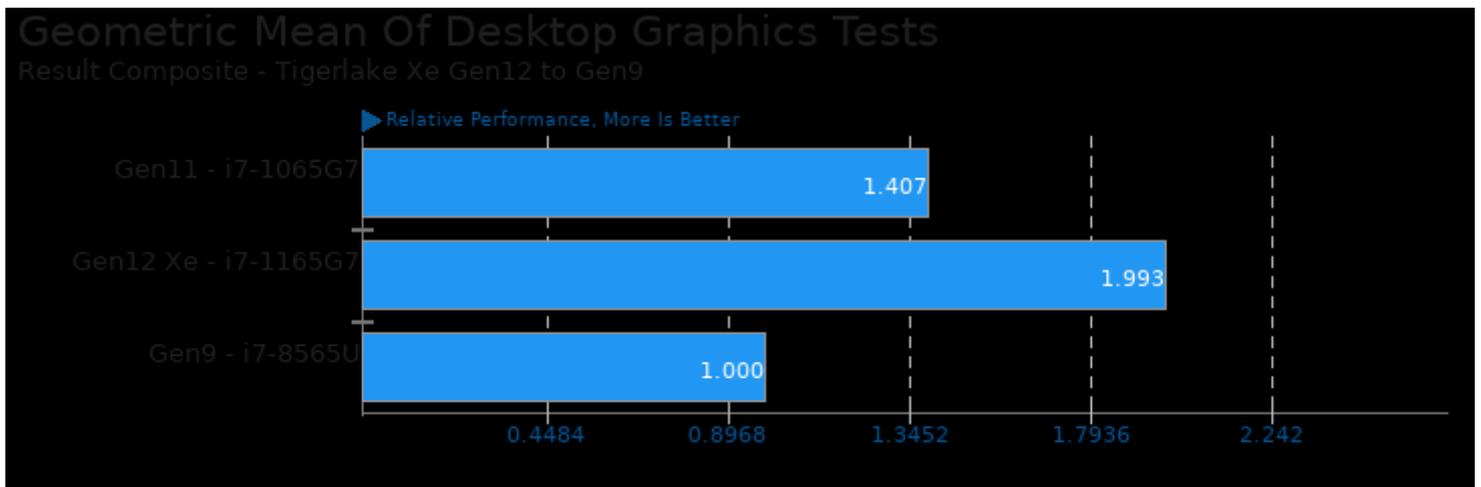




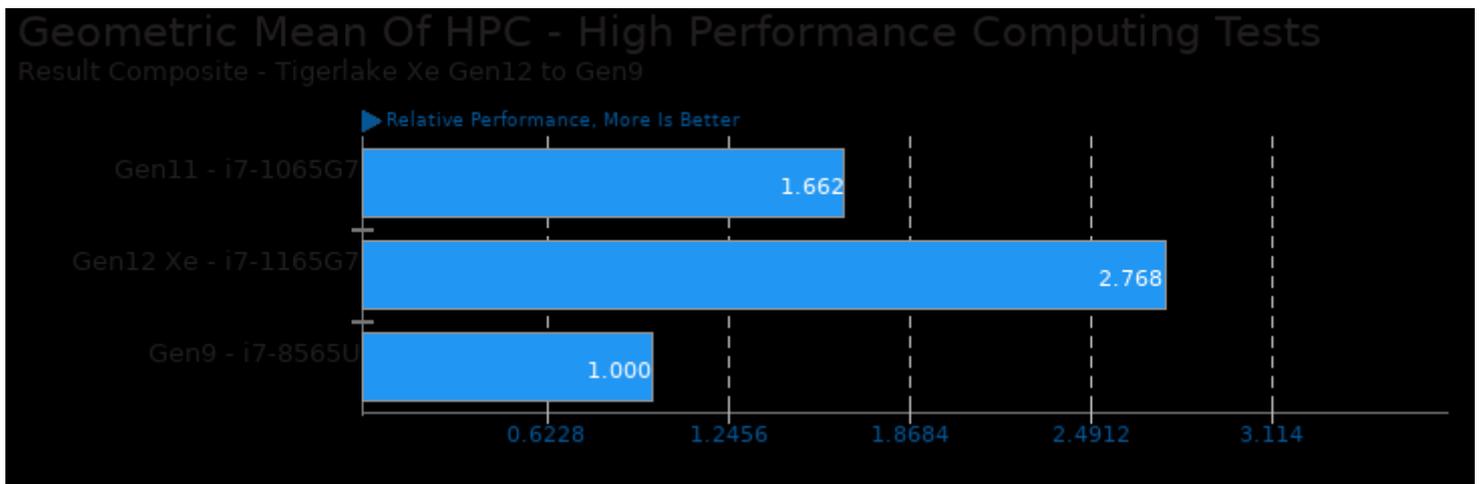
These geometric means are based upon test groupings / test suites for this result file.



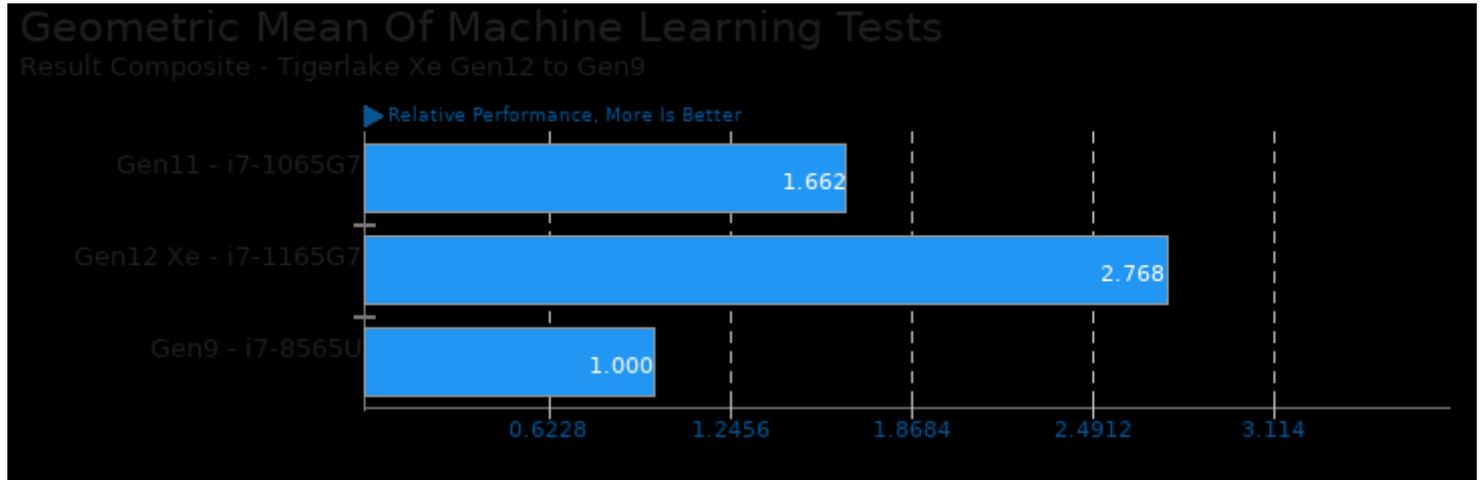
Geometric mean based upon tests: pts/lczero and pts/plaidml



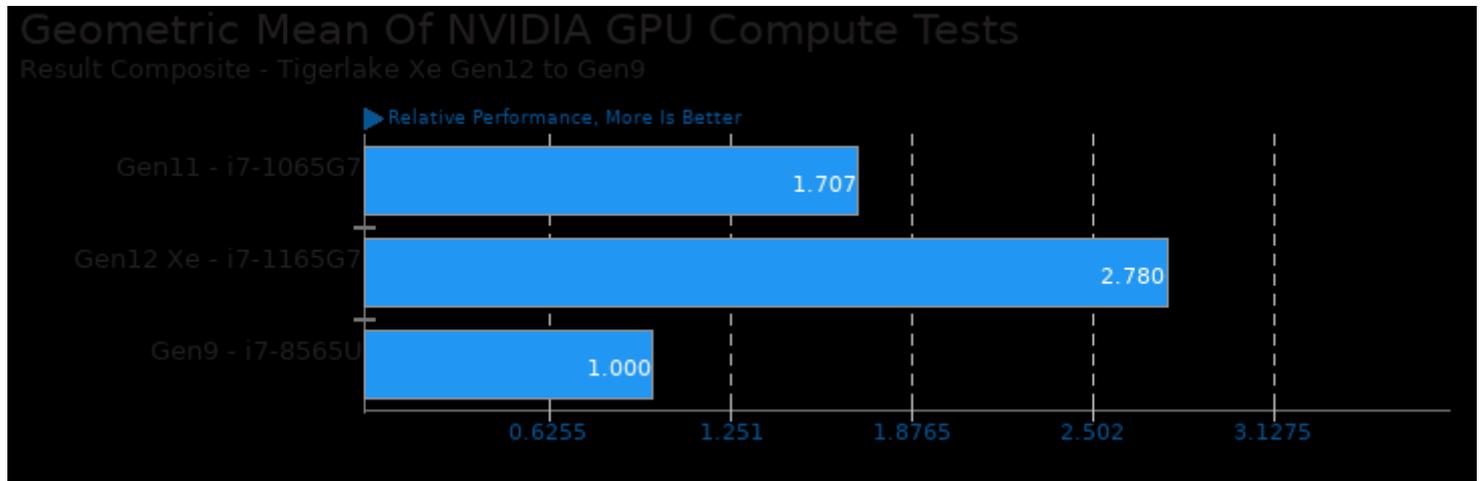
Geometric mean based upon tests: pts/xonotic, pts/tesseract, pts/unigine-valley and pts/unigine-heaven



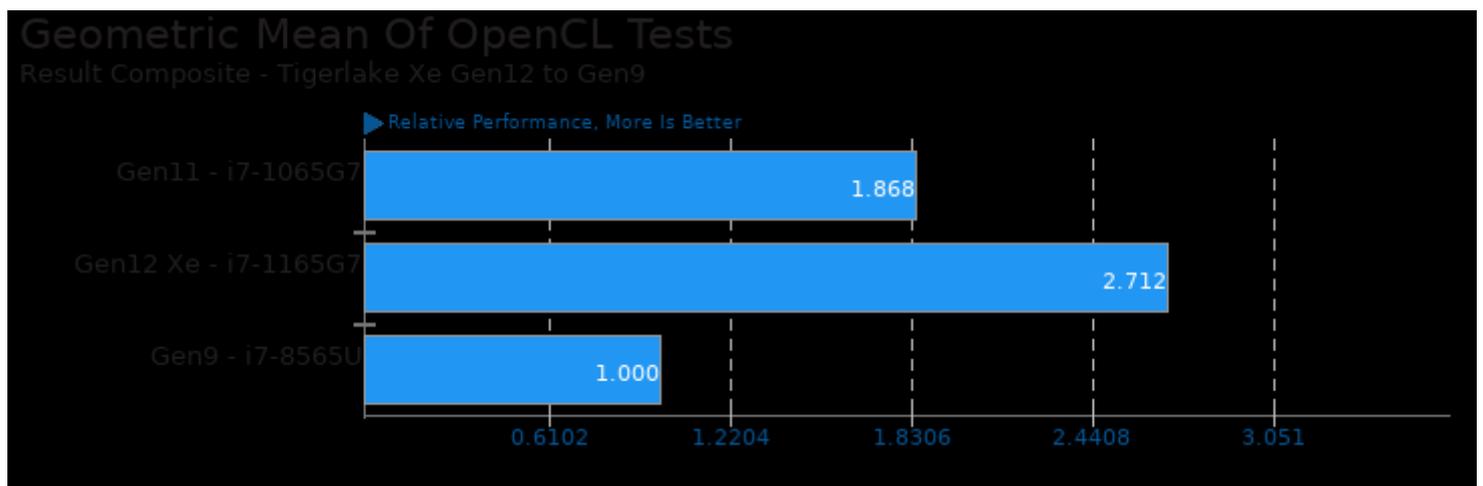
Geometric mean based upon tests: pts/ncnn, pts/shoc, pts/plaidml and pts/lczero



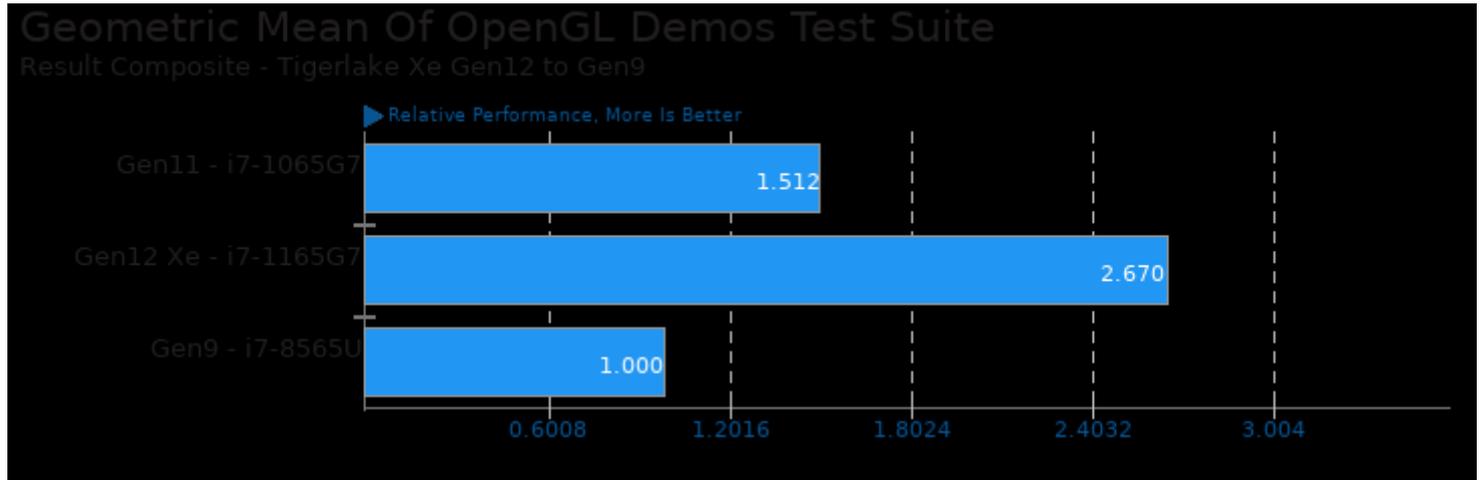
Geometric mean based upon tests: pts/ncnn, pts/shoc, pts/plaidml and pts/lczero



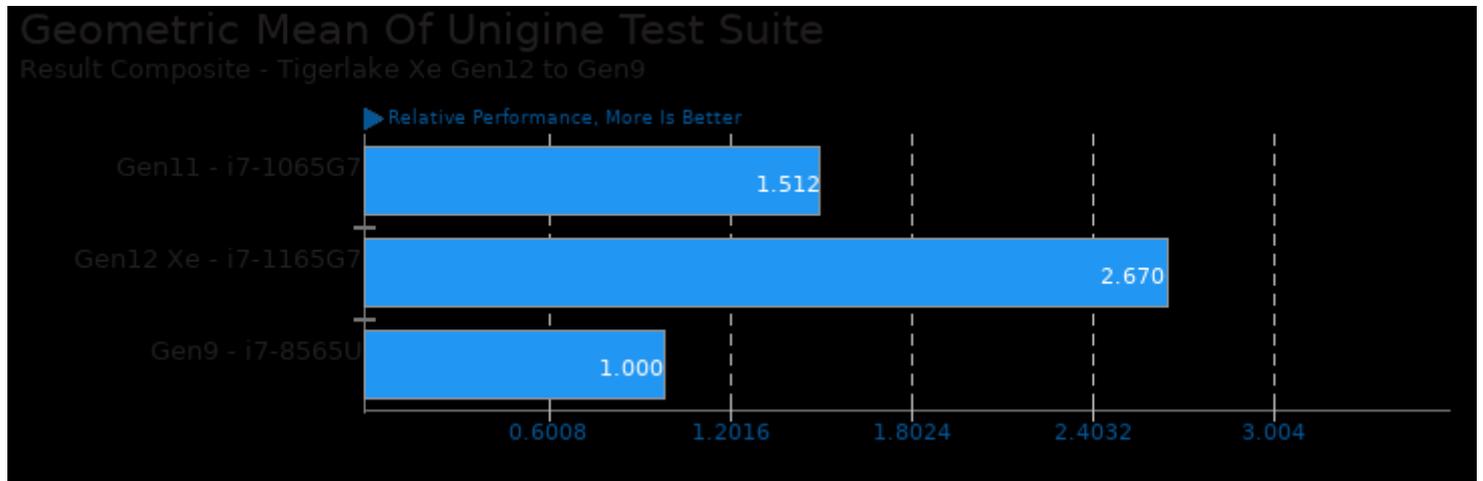
Geometric mean based upon tests: pts/clpeak, pts/financebench, pts/plaidml, pts/lczero, pts/cl-mem, pts/mandelgpu, pts/viennacl, pts/shoc, pts/ncnn, pts/realsr-ncnn and pts/waifu2x-ncnn



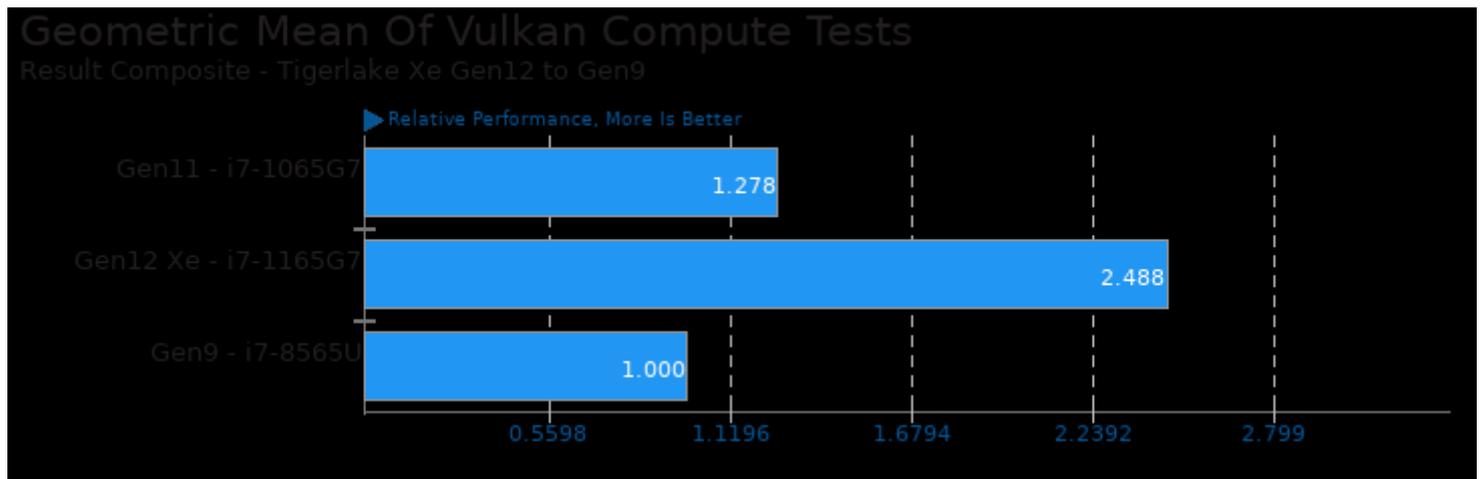
Geometric mean based upon tests: pts/juliagpu, pts/mandelgpu, pts/shoc, pts/cl-mem, pts/clpeak and pts/viennacl



Geometric mean based upon tests: pts/unigine-valley, pts/unigine-heaven and pts/unigine-super



Geometric mean based upon tests: pts/unigine-valley, pts/unigine-heaven and pts/unigine-super



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

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