



## 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-multilib --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-nvptx=/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-nvptx=/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
<b>FinanceBench - B.S.O (ms)</b>	<b>25.826923</b>	6.343	<b>4.179</b>
<b>Normalized</b>	16.18%	65.88%	100%
<b>Standard Deviation</b>	0.6%	3.2%	0.5%

<b>SHOC Scalable HeterOgeneous Computing - OpenCL - Max SP Flops (GFLOPS)</b>	<b>1681</b>	5323	<b>7821</b>
Normalized	21.5%	68.06%	100%
Standard Deviation	0.1%	0.1%	0%
<b>SHOC Scalable HeterOgeneous Computing - OpenCL - MD5 Hash (GHash/s)</b>	<b>0.3691</b>	0.6778	<b>1.6775</b>
Normalized	22%	40.41%	100%
Standard Deviation	0%	0.1%	0%
<b>oneAPI Level Zero Tests - H.T.D.T.H.I.C</b>	<b>4.83779</b>	10.5736	<b>21.6345</b>
Normalized	22.36%	48.87%	100%
Standard Deviation	0.6%	2.9%	0.9%
<b>MandelGPU - GPU (Samples/sec)</b>	<b>10702353</b>	23689429	<b>44609195</b>
Normalized	23.99%	53.1%	100%
Standard Deviation	0%	1.7%	1.2%
<b>NCNN - Vulkan GPU - yolov4-tiny (ms)</b>	<b>62.81</b>	38.95	<b>15.26</b>
Normalized	24.3%	39.18%	100%
Standard Deviation	0.8%	0.7%	0.2%
<b>GpuTest - Pixmark Volplosion - 1920 x 1200 - Fullscreen (Points)</b>	<b>390</b>	761	<b>1595</b>
Normalized	24.45%	47.71%	100%
Standard Deviation		1.1%	1.6%
<b>clpeak - S.P.F (GFLOPS)</b>	<b>438.06</b>	1039	<b>1775</b>
Normalized	24.68%	58.52%	100%
Standard Deviation	0%	0.1%	0%
<b>oneAPI Level Zero Tests - P.H.P.C (GFLOPS)</b>	<b>751.623</b>	1352	<b>3037</b>
Normalized	24.75%	44.53%	100%
Standard Deviation	0.1%	2.5%	0.1%
<b>SHOC Scalable HeterOgeneous Computing - OpenCL - FFT SP (GFLOPS)</b>	<b>38.9719</b>	140.162	<b>156.162</b>
Normalized	24.96%	89.75%	100%
Standard Deviation	0.3%	0.5%	0.8%
<b>NCNN - Vulkan GPU - vgg16 (ms)</b>	<b>152.79</b>	96.32	<b>38.44</b>
Normalized	25.16%	39.91%	100%
Standard Deviation	0.1%	0.1%	0.1%
<b>LeelaChessZero - OpenCL (Nodes/s)</b>	<b>389</b>	579	<b>1535</b>
Normalized	25.34%	37.72%	100%
Standard Deviation	2.3%	2.4%	0.9%
<b>GpuTest - Pixmark Piano - 1920 x 1200 - Fullscreen (Points)</b>	<b>149</b>	288	<b>584</b>
Normalized	25.51%	49.32%	100%
Standard Deviation		1%	1.9%
<b>PlaidML - No - Inference - IMDB LSTM - OpenCL (FPS)</b>	<b>7.84</b>	20.10	<b>30.52</b>
Normalized	25.69%	65.86%	100%
Standard Deviation	0.4%	1%	2.3%
<b>PlaidML - No - Inference - DenseNet 201 - OpenCL (FPS)</b>	<b>5.81</b>	11.30	<b>22.44</b>
Normalized	25.89%	50.36%	100%
Standard Deviation	0.4%	1.2%	0.7%
<b>SHOC Scalable HeterOgeneous Computing - OpenCL - T.R.B (GB/s)</b>	<b>54.2625</b>	174.162	<b>189.757</b>
Normalized	28.6%	91.78%	100%
Standard Deviation	0%	1.7%	0%
<b>NCNN - Vulkan GPU - resnet18 (ms)</b>	<b>23.89</b>	17.12	<b>7.18</b>

	Normalized	30.05%	41.94%	100%
	Standard Deviation	0%	0.2%	0.6%
<b>NCNN - Vulkan GPU - resnet50 (ms)</b>		<b>51.46</b>	35.91	<b>15.49</b>
	Normalized	30.1%	43.14%	100%
	Standard Deviation	0.1%	0.2%	0.2%
<b>oneAPI Level Zero Tests - P.I.C (GFLOPS)</b>		<b>138.516</b>	153.147	<b>439.002</b>
	Normalized	31.55%	34.89%	100%
	Standard Deviation	3%	2.4%	0.5%
<b>RealSR-NCNN - 4x - Yes (sec)</b>		<b>1615</b>	1120	<b>538.065</b>
	Normalized	33.32%	48.04%	100%
	Standard Deviation	0.3%	0.5%	0.4%
<b>ViennaCL - O.L.F (GFLOPS)</b>		<b>24.6746</b>	48.9117	<b>73.8162</b>
	Normalized	33.43%	66.26%	100%
	Standard Deviation	0.3%	1.3%	0.4%
<b>RealSR-NCNN - 4x - No (sec)</b>		<b>202.601</b>	143.988	<b>68.311</b>
	Normalized	33.72%	47.44%	100%
	Standard Deviation	0.5%	0.4%	1.1%
<b>PlaidML - No - Inference - Inception V3 - OpenCL (FPS)</b>		<b>15.90</b>	34.82	<b>46.37</b>
	Normalized	34.29%	75.09%	100%
	Standard Deviation	1.7%	0.4%	0.1%
<b>oneAPI Level Zero Tests - P.S.P.C (GB/s)</b>		<b>418.862</b>	920.207	<b>1220</b>
	Normalized	34.35%	75.46%	100%
	Standard Deviation	0.7%	0%	0%
<b>Unigine Superposition - 1920 x 1200 - Fullscreen - Medium - OpenGL (FPS)</b>		<b>5.8</b>	9.0	<b>16.4</b>
	Normalized	35.37%	54.88%	100%
	Standard Deviation	1%	0.6%	0%
<b>PlaidML - No - Inference - NASNet Large - OpenCL (FPS)</b>		<b>1.84</b>	3.31	<b>5.14</b>
	Normalized	35.8%	64.4%	100%
	Standard Deviation	0.5%	0.3%	0%
<b>Waifu2x-NCNN Vulkan - 2x - 3 - Yes (sec)</b>		<b>75.252</b>	53.103	<b>27.216</b>
	Normalized	36.17%	51.25%	100%
	Standard Deviation	0.1%	3%	0.2%
<b>JuliaGPU - GPU (Samples/sec)</b>		<b>47930813</b>	85184313	<b>132508055</b>
	Normalized	36.17%	64.29%	100%
	Standard Deviation	0%	3%	2.6%
<b>Unigine Heaven - 1920 x 1200 - Fullscreen - OpenGL (FPS)</b>		<b>10.3701</b>	16.1363	<b>28.6453</b>
	Normalized	36.2%	56.33%	100%
	Standard Deviation	0.3%	4.5%	0.4%
<b>PlaidML - No - Inference - VGG19 - OpenCL (FPS)</b>		<b>6.42</b>	16.03	<b>17.49</b>
	Normalized	36.71%	91.65%	100%
	Standard Deviation	0.2%	2.2%	0.2%
<b>NCNN - Vulkan GPU - alexnet (ms)</b>		<b>25.77</b>	16.63	<b>9.47</b>
	Normalized	36.75%	56.95%	100%
	Standard Deviation	0.1%	0.3%	0.1%
<b>PlaidML - No - Inference - VGG16 - OpenCL (FPS)</b>		<b>8.18</b>	19.93	<b>22.18</b>
	Normalized	36.88%	89.86%	100%
	Standard Deviation	1.6%	5.9%	0.3%

<b>Unigine Valley - 1920 x 1200 - Fullscreen - OpenGL (FPS)</b>	<b>11.4067</b>	17.7436	<b>30.6054</b>
Normalized	37.27%	57.98%	100%
Standard Deviation	1%	1.1%	0.5%
<b>GpuTest - Furmark - 1920 x 1200 - Fullscreen (Points)</b>	<b>672</b>	1062	<b>1782</b>
Normalized	37.71%	59.6%	100%
Standard Deviation	2.2%	2.5%	2.6%
<b>NCNN - Vulkan GPU - mobilenet (ms)</b>	<b>25.74</b>	22.36	<b>9.74</b>
Normalized	37.84%	43.56%	100%
Standard Deviation	2.6%	1.6%	1.2%
<b>PlaidML - No - Inference - Mobilenet - OpenCL (FPS)</b>	<b>116.08</b>	242.83	<b>299.60</b>
Normalized	38.74%	81.05%	100%
Standard Deviation	0%	2.4%	0.5%
<b>PlaidML - No - Inference - ResNet 50 - OpenCL (FPS)</b>	<b>34.52</b>	75.57	<b>87.29</b>
Normalized	39.55%	86.57%	100%
Standard Deviation	1.8%	2.4%	0.1%
<b>cl-mem - Read (GB/s)</b>	<b>22.7</b>	41.3	<b>57.1</b>
Normalized	39.75%	72.33%	100%
Standard Deviation	1.8%	0.7%	2.3%
<b>Waifu2x-NCNN Vulkan - 2x - 3 - No (sec)</b>	<b>10.288</b>	7.062	<b>4.095</b>
Normalized	39.8%	57.99%	100%
Standard Deviation	0.6%	0.9%	0.9%
<b>NCNN - Vulkan GPU - squeezenet (ms)</b>	<b>28.27</b>	24.65	<b>11.51</b>
Normalized	40.71%	46.69%	100%
Standard Deviation	0.4%	2.2%	0.3%
<b>Unigine Superposition - 1920 x 1200 - Fullscreen - Low - OpenGL (FPS)</b>	<b>12.2</b>	17.0	<b>29.6</b>
Normalized	41.22%	57.43%	100%
Standard Deviation	0.5%	0.9%	0.8%
<b>oneAPI Level Zero Tests - P.F.G.M.B (GB/s)</b>	<b>24.5276</b>	34.7522	<b>56.9391</b>
Normalized	43.08%	61.03%	100%
Standard Deviation	2.1%	0.1%	0.5%
<b>NCNN - Vulkan GPU - googlenet (ms)</b>	<b>23.72</b>	18.62	<b>10.51</b>
Normalized	44.31%	56.44%	100%
Standard Deviation	0.5%	0.1%	0.5%
<b>cl-mem - Copy (GB/s)</b>	<b>21.6</b>	35.9	<b>48.0</b>
Normalized	45%	74.79%	100%
Standard Deviation	1.9%	0.5%	0.1%
<b>NCNN - Vulkan GPU - shufflenet-v2 (ms)</b>	<b>6.75</b>	5.71	<b>3.05</b>
Normalized	45.19%	53.42%	100%
Standard Deviation	0.5%	1%	0.9%
<b>Tesseract - 1920 x 1200 (FPS)</b>	<b>65.22251</b>	109.4669	<b>140.3486</b>
Normalized	46.47%	78%	100%
Standard Deviation	1.2%	2.3%	1.1%
<b>clpeak - G.M.B (GBPS)</b>	<b>26.73</b>	45.94	<b>57.02</b>
Normalized	46.88%	80.57%	100%
Standard Deviation	0.6%	0.1%	0.1%
<b>cl-mem - Write (GB/s)</b>	<b>24.0</b>	34.4	<b>47.9</b>
Normalized	50.1%	71.82%	100%
Standard Deviation	2.7%	5%	0.8%
<b>oneAPI Level Zero Tests - D.T.H.B (usec)</b>	<b>19941</b>	12233	<b>10057</b>

	Normalized	50.44%	82.22%	100%
	Standard Deviation	1%	0.1%	0.2%
<b>oneAPI Level Zero Tests - D.T.H.B (GB/s)</b>		<b>13.462183</b>	21.943939	<b>26.690187</b>
	Normalized	50.44%	82.22%	100%
	Standard Deviation	1%	0.1%	0.2%
<b>oneAPI Level Zero Tests - H.T.D.B (usec)</b>		<b>19750</b>	12241	<b>10049</b>
	Normalized	50.88%	82.09%	100%
	Standard Deviation	0.9%	0.2%	0.3%
<b>oneAPI Level Zero Tests - H.T.D.B (GB/s)</b>		<b>13.592119</b>	21.928612	<b>26.714162</b>
	Normalized	50.88%	82.09%	100%
	Standard Deviation	0.9%	0.2%	0.3%
<b>NCNN - Vulkan GPU - blazeface (ms)</b>		<b>1.81</b>	<b>1.81</b>	<b>0.98</b>
	Normalized	54.14%	54.14%	100%
	Standard Deviation	3.9%	1.8%	3.5%
<b>NCNN - Vulkan GPU-v2-v2 - mobilenet-v2</b>		<b>8.27</b>	7.53	<b>4.62</b>
	Normalized	55.86%	61.35%	100%
	Standard Deviation	0.8%	1.3%	2.3%
<b>GpuTest - TessMark - 1920 x 1200 - Fullscreen (Points)</b>		<b>2840</b>	3602	<b>5080</b>
	Normalized	55.91%	70.91%	100%
	Standard Deviation	1.5%	1.3%	0.9%
<b>Xonotic - 1920 x 1200 - Ultimate (FPS)</b>		<b>71.6067804</b>	99.2365089	<b>121.9041377</b>
	Normalized	58.74%	81.41%	100%
	Standard Deviation	1.1%	2.7%	1.5%
<b>NCNN - Vulkan GPU-v3-v3 - mobilenet-v3</b>		<b>9.19</b>	8.65	<b>5.47</b>
	Normalized	59.52%	63.24%	100%
	Standard Deviation	0.3%	0.2%	1.4%
<b>Xonotic - 1920 x 1200 - Low (FPS)</b>		<b>183.6211229</b>	209.1609960	<b>307.6345501</b>
	Normalized	59.69%	67.99%	100%
	Standard Deviation	0.1%	2.9%	2.9%
<b>Xonotic - 1920 x 1200 - Ultra (FPS)</b>		<b>93.2094800</b>	123.2383929	<b>155.3708840</b>
	Normalized	59.99%	79.32%	100%
	Standard Deviation	1%	2.6%	1.5%
<b>Xonotic - 1920 x 1200 - High (FPS)</b>		<b>111.6753218</b>	144.1207832	<b>183.8028138</b>
	Normalized	60.76%	78.41%	100%
	Standard Deviation	1.3%	2.8%	2.4%
<b>GpuTest - GiMark - 1920 x 1200 - Fullscreen (Points)</b>		<b>1314</b>	1724	<b>2017</b>
	Normalized	65.15%	85.47%	100%
	Standard Deviation	0.8%	1%	0.1%
<b>NCNN - Vulkan GPU - efficientnet-b0 (ms)</b>		16.52	<b>16.72</b>	<b>11.05</b>
	Normalized	66.89%	66.09%	100%
	Standard Deviation	0.1%	0.1%	0.4%
<b>oneAPI Level Zero Tests - P.S.M.C.t.S.M</b>		<b>10.7883</b>	13.8085	<b>15.0859</b>
	Normalized	71.51%	91.53%	100%
	Standard Deviation	2.9%	0.3%	0.5%
<b>clpeak - Kernel Latency (us)</b>		34.12	<b>36.89</b>	<b>29.53</b>
	Normalized	86.55%	80.05%	100%
	Standard Deviation	2.3%	0.7%	0.5%
<b>GpuTest - Pixmark Volplosion - 1920 x 1200 - Fullscreen (Points/Watt)</b>		<b>16.06</b>	45.24	<b>75.39</b>
	Normalized	21.3%	60.01%	100%


GpuTest - Pixmark Piano - 1920 x 1200 - Fullscreen (Points/Watt)	<b>6.01</b>	17.13	<b>27.39</b>
Normalized	21.94%	62.54%	100%
GpuTest - TessMark - 1920 x 1200 - Fullscreen (Points/Watt)	<b>100.32</b>	214.06	<b>242.99</b>
Normalized	41.29%	88.09%	100%
GpuTest - Furmark - 1920 x 1200 - Fullscreen (Points/Watt)	<b>25.26</b>	63.19	<b>81.15</b>
Normalized	31.13%	77.87%	100%
GpuTest - GiMark - 1920 x 1200 - Fullscreen (Points/Watt)	<b>47.23</b>	104.65	<b>111.04</b>
Normalized	42.53%	94.25%	100%
Xonotic - 1920 x 1200 - Ultimate (FPS/Watt)	<b>2.65</b>	<b>6.10</b>	5.56
Normalized	43.44%	100%	91.15%
Xonotic - 1920 x 1200 - Ultra (FPS/Watt)	<b>3.44</b>	<b>7.51</b>	7.06
Normalized	45.81%	100%	94.01%
Xonotic - 1920 x 1200 - High (FPS/Watt)	<b>4.25</b>	<b>8.81</b>	8.28
Normalized	48.24%	100%	93.98%
Xonotic - 1920 x 1200 - Low (FPS/Watt)	<b>8.37</b>	12.89	<b>14.14</b>
Normalized	59.19%	91.16%	100%
Unigine Valley - 1920 x 1200 - Fullscreen - OpenGL (FPS/Watt)	<b>0.43</b>	1.13	<b>1.40</b>
Normalized	30.71%	80.71%	100%
Unigine Superposition - 1920 x 1200 - Fullscreen - Medium - OpenGL (FPS/Watt)	<b>0.23</b>	0.50	<b>0.74</b>
Normalized	31.08%	67.57%	100%
Unigine Superposition - 1920 x 1200 - Fullscreen - Low - OpenGL (FPS/Watt)	<b>0.48</b>	0.89	<b>1.34</b>
Normalized	35.82%	66.42%	100%
Unigine Heaven - 1920 x 1200 - Fullscreen - OpenGL (FPS/Watt)	<b>0.38</b>	0.81	<b>1.31</b>
Normalized	29.01%	61.83%	100%
Tesseract - 1920 x 1200 (FPS/Watt)	<b>2.61</b>	5.00	<b>6.28</b>
Normalized	41.56%	79.62%	100%
ET: Legacy - Renderer2 - 1920 x 1200 (FPS/Watt)	<b>3.41</b>	5.40	<b>5.59</b>
Normalized	61%	96.6%	100%
clpeak - T.B.e (GBPS/Watt)	<b>2.60</b>	<b>3.45</b>	3.44
Normalized	75.36%	100%	99.71%
clpeak - G.M.B (GBPS/Watt)	<b>0.92</b>	<b>2.71</b>	2.54
Normalized	33.95%	100%	93.73%
clpeak - S.P.F (GFLOPS/Watt)	<b>20.31</b>	58.93	<b>80.98</b>
Normalized	25.08%	72.77%	100%
MandelGPU - GPU (Samples/sec/Watt)	<b>468417</b>	1544709	<b>1974932</b>
Normalized	23.72%	78.22%	100%
JuliaGPU - GPU (Samples/sec/Watt)	<b>1343813</b>	5650341	<b>5935646</b>
Normalized	22.64%	95.19%	100%
PlaidML - No - Inference - NASner Large - OpenCL (FPS/Watt)	<b>0.06</b>	0.22	<b>0.24</b>
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)	<b>0.92</b>	1.68	<b>1.92</b>
Normalized	47.92%	87.5%	100%
oneAPI Level Zero Tests - P.S.M.C.t.S.M (GB/s/Watt)	<b>0.36</b>	<b>0.90</b>	0.67
Normalized	40%	100%	74.44%
oneAPI Level Zero Tests - P.F.G.M.B (GB/s/Watt)	<b>0.85</b>	2.24	<b>2.49</b>
Normalized	34.14%	89.96%	100%
oneAPI Level Zero Tests - H.T.D.T.H.I.C (GB/s/Watt)	<b>0.14</b>	0.68	<b>0.94</b>
Normalized	14.89%	72.34%	100%
oneAPI Level Zero Tests - P.S.P.C	<b>13.93</b>	<b>59.52</b>	54.99
Normalized	23.4%	100%	92.39%
oneAPI Level Zero Tests - P.H.P.C (GFLOPS/Watt)	<b>24.17</b>	73.98	<b>135.36</b>
Normalized	17.86%	54.65%	100%
oneAPI Level Zero Tests - P.I.C (GFLOPS/Watt)	<b>4.31</b>	9.15	<b>20.00</b>
Normalized	21.55%	45.75%	100%
ET: Legacy - Renderer2 - 1920 x 1200 (FPS)	<b>72.1</b>	86.4	<b>122.3</b>
Normalized	58.95%	70.65%	100%
Standard Deviation	0.1%	8.2%	2.9%
clpeak - T.B.e (GBPS)	<b>31.64</b>	39.48	<b>52.15</b>
Normalized	60.67%	75.7%	100%
Standard Deviation	0.1%	6.4%	1%
NCNN - Vulkan GPU - mnasnet (ms)	<b>8.36</b>	7.65	<b>5.14</b>
Normalized	61.48%	67.19%	100%
Standard Deviation	0.3%	1.1%	9.7%
SHOC Scalable Heterogeneous Computing - OpenCL - Bus Speed Readback (GB/s)	<b>23.9319</b>	41.7406	<b>60.4048</b>
Normalized	39.62%	69.1%	100%
Standard Deviation	9.3%	3.9%	1%
SHOC Scalable Heterogeneous Computing - OpenCL - Bus Speed Download (GB/s)	<b>23.3048</b>	40.1640	<b>58.4143</b>
Normalized	39.9%	68.76%	100%
Standard Deviation	8.6%	1%	2.9%
SHOC Scalable Heterogeneous Computing - OpenCL - Triad (GB/s)	<b>10.9236</b>	14.0210	<b>21.3040</b>
Normalized	51.27%	65.81%	100%
Standard Deviation	6.8%	7.4%	1.9%
FinanceBench - Monte-Carlo OpenCL (ms)	<b>2585</b>	579.523733	<b>433.370858</b>
Normalized	16.77%	74.78%	100%
Standard Deviation	8.3%	7.7%	0.9%
oneAPI Level Zero Tests - P.K.L.L (us)	<b>31.4224</b>	28.9568	<b>22.3130</b>
Normalized	71.01%	77.06%	100%
Standard Deviation	1.6%	11%	9.6%

### FinanceBench 2016-06-06


Benchmark: Black-Scholes OpenCL



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

### SHOC Scalable Heterogeneous Computing 2015-11-10


Target: OpenCL - Benchmark: Max SP Flops



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

### SHOC Scalable Heterogeneous Computing 2015-11-10


Target: OpenCL - Benchmark: MD5 Hash



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

## oneAPI Level Zero Tests


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



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

## MandelGPU 1.3pts1


OpenCL Device: GPU



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

## NCNN 20200916


Target: Vulkan GPU - Model: yolov4-tiny



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


## GpuTest 0.7.0

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



## clpeak


OpenCL Test: Single-Precision Float



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

## oneAPI Level Zero Tests


Test: Peak Half-Precision Compute



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

## 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

## NCNN 20200916


Target: Vulkan GPU - Model: vgg16



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

## LeelaChessZero 0.26


Backend: OpenCL



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


## GpuTest 0.7.0

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




## PlaidML

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




## PlaidML

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



## SHOC Scalable Heterogeneous Computing 2015-11-10


Target: OpenCL - Benchmark: Texture Read Bandwidth



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

## NCNN 20200916


Target: Vulkan GPU - Model: resnet18



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

## NCNN 20200916


Target: Vulkan GPU - Model: resnet50



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


## oneAPI Level Zero Tests

Test: Peak Integer Compute




## RealSR-NCNN 20200818

Scale: 4x - TAA: Yes




## ViennaCL 1.4.2

OpenCL LU Factorization




## RealSR-NCNN 20200818

Scale: 4x - TAA: No




## PlaidML

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




## PlaidML

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




## Waifu2x-NCNN Vulkan 20200818

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



## JuliaGPU 1.2pts1


OpenCL Device: GPU



1. (CC) gcc options: -O3 -march=native -ftriple-precision-float -funroll-loops -lglut -lOpenCL -lGL -lm


## Unigine Heaven 4.0

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




## PlaidML

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



## NCNN 20200916


Target: Vulkan GPU - Model: alexnet



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


## PlaidML

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




## Unigine Valley 1.0

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




## GpuTest 0.7.0

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



## NCNN 20200916


Target: Vulkan GPU - Model: mobilenet



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


## PlaidML

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




## PlaidML

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



## cl-mem 2017-01-13


Benchmark: Read



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


## Waifu2x-NCNN Vulkan 20200818

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



## NCNN 20200916


Target: Vulkan GPU - Model: squeezenet



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


## Unigine Superposition 1.0

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



## oneAPI Level Zero Tests


Test: Peak Float16 Global Memory Bandwidth



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

### NCNN 20200916


Target: Vulkan GPU - Model: googlenet



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

### cl-mem 2017-01-13


Benchmark: Copy



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

### NCNN 20200916


Target: Vulkan GPU - Model: shufflenet-v2



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


## Tesseract 2014-05-12

Resolution: 1920 x 1200



## clpeak


OpenCL Test: Global Memory Bandwidth



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

## cl-mem 2017-01-13

Benchmark: Write




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

## oneAPI Level Zero Tests

Test: Device-To-Host Bandwidth

usec, Fewer Is Better




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

## oneAPI Level Zero Tests

Test: Device-To-Host Bandwidth

GB/s, More Is Better




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

## oneAPI Level Zero Tests

Test: Host-To-Device Bandwidth


usec, Fewer Is Better



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

### oneAPI Level Zero Tests


Test: Host-To-Device Bandwidth



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

### NCNN 20200916


Target: Vulkan GPU - Model: blazeface



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

### NCNN 20200916


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



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


## GpuTest 0.7.0

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




## Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Ultimate



## NCNN 20200916


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



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


## Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Low




## Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Ultra




## Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: High




## GpuTest 0.7.0

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



## NCNN 20200916


Target: Vulkan GPU - Model: efficientnet-b0



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

## oneAPI Level Zero Tests


Test: Peak System Memory Copy to Shared Memory



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

## clpeak


OpenCL Test: Kernel Latency



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

## CPU Power Consumption Monitor

Phoronix Test Suite System Monitoring




## Waifu2x-NCNN Vulkan 20200818

CPU Power Consumption Monitor

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

Watts, Fewer Is Better




## Waifu2x-NCNN Vulkan 20200818

CPU Power Consumption Monitor

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

▼ Watts, Fewer Is Better




## GpuTest 0.7.0

CPU Power Consumption Monitor


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

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

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


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


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


▼ 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
Gen9 - i7-8565U	1.2	27.0	30.5
Gen11 - i7-1065G7	3.8	16.3	32.4
Gen12 Xe - i7-1165G7	2.0	21.9	38.9


▼ Watts, Fewer Is Better




## Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Ultimate

► Frames Per Second Per Watt, More Is Better






## Xonotic 0.8.2

CPU Power Consumption Monitor

	Min	Avg	Max
Gen9 - i7-8565U	1.0	26.2	30.3
Gen11 - i7-1065G7	3.8	16.4	29.9
Gen12 Xe - i7-1165G7	2.2	22.2	37.0


▼ Watts, Fewer Is Better



## Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: High

► Frames Per Second Per Watt, More Is Better




## Xonotic 0.8.2

CPU Power Consumption Monitor

	Min	Avg	Max
Gen9 - i7-8565U	1.0	21.9	26.0
Gen11 - i7-1065G7	3.6	16.2	24.7
Gen12 Xe - i7-1165G7	1.8	21.8	37.2


▼ Watts, Fewer Is Better



## Xonotic 0.8.2

Resolution: 1920 x 1200 - Effects Quality: Low

► Frames Per Second Per Watt, More Is Better




## Unigine Valley 1.0

CPU Power Consumption Monitor

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


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


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


▼ 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

	Min	Avg	Max
Gen9 - i7-8565U	0.8	27.1	30.7
Gen11 - i7-1065G7	1.5	20.0	35.9
Gen12 Xe - i7-1165G7	1.5	21.9	33.9


▼ Watts, Fewer Is Better



## Unigine Heaven 4.0

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

► Frames Per Second Per Watt, More Is Better




## Tesseract 2014-05-12

CPU Power Consumption Monitor

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


▼ Watts, Fewer Is Better



## Tesseract 2014-05-12

Resolution: 1920 x 1200

► Frames Per Second Per Watt, More Is Better



## ET: Legacy 2.75

CPU Power Consumption Monitor

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


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


▼ Watts, Fewer Is Better



## clpeak

OpenCL Test: Transfer Bandwidth enqueueWriteBuffer

▶ GBPS Per Watt, More Is Better




## clpeak

CPU Power Consumption Monitor

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


▼ Watts, Fewer Is Better



## clpeak

OpenCL Test: Global Memory Bandwidth

▶ GBPS Per Watt, More Is Better




## clpeak

CPU Power Consumption Monitor

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


▼ Watts, Fewer Is Better



## clpeak

OpenCL Test: Single-Precision Float

▶ GFLOPS Per Watt, More Is Better




## clpeak

CPU Power Consumption Monitor

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

▼ Watts, Fewer Is Better




## MandelGPU 1.3pts1

CPU Power Consumption Monitor


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

▼ Watts, Fewer Is Better



MandelGPU 1.3pts1

OpenCL Device: GPU




JuliaGPU 1.2pts1

CPU Power Consumption Monitor


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

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

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


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


▼ Watts, Fewer Is Better



## PlaidML

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


► FPS Per Watt, More Is Better



## PlaidML


CPU Power Consumption Monitor

◀ Watts, Fewer Is Better




## PlaidML

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




## PlaidML

CPU Power Consumption Monitor



## PlaidML

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




## PlaidML

CPU Power Consumption Monitor

	Min	Avg	Max
Gen9 - i7-8565U	0.8	31.9	49.8
Gen11 - i7-1065G7	0.9	16.5	45.4
Gen12 Xe - i7-1165G7	1.5	20.8	53.3


▼ Watts, Fewer Is Better



## PlaidML

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

▶ FPS Per Watt, More Is Better



## NCNN 20200916

CPU Power Consumption Monitor

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


▼ Watts, Fewer Is Better



## LeelaChessZero 0.26

CPU Power Consumption Monitor


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


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


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


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


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


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


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


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


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


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


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

▼ Watts, Fewer Is Better




## FinanceBench 2016-06-06

CPU Power Consumption Monitor


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

▼ Watts, Fewer Is Better



RealSR-NCNN 20200818

CPU Power Consumption Monitor




RealSR-NCNN 20200818

CPU Power Consumption Monitor


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

Watts, Fewer Is Better




oneAPI Level Zero Tests

CPU Power Consumption Monitor



oneAPI Level Zero Tests

Test: Peak System Memory Copy to Shared Memory




oneAPI Level Zero Tests

CPU Power Consumption Monitor


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

Watts, Fewer Is Better



oneAPI Level Zero Tests

Test: Peak Float16 Global Memory Bandwidth




## oneAPI Level Zero Tests

CPU Power Consumption Monitor

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


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


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


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

▼ 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
Gen9 - i7-8565U	0.7	31.0	38.1
Gen11 - i7-1065G7	1.0	15.7	30.7
Gen12 Xe - i7-1165G7	2.0	22.7	35.7

▼ Watts, Fewer Is Better




## oneAPI Level Zero Tests

CPU Power Consumption Monitor




## oneAPI Level Zero Tests

Test: Peak Integer Compute



## ET: Legacy 2.75


Renderer: Renderer2 - Resolution: 1920 x 1200



## clpeak

OpenCL Test: Transfer Bandwidth enqueueWriteBuffer

▶ GBPS, More Is Better




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

## NCNN 20200916

Target: Vulkan GPU - Model: mnasnet

◀ ms, Fewer Is Better




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

## SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Bus Speed Readback


▶ GB/s, More Is Better



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

### SHOC Scalable Heterogeneous Computing 2015-11-10


Target: OpenCL - Benchmark: Bus Speed Download



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

### SHOC Scalable Heterogeneous Computing 2015-11-10


Target: OpenCL - Benchmark: Triad




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

### FinanceBench 2016-06-06


Benchmark: Monte-Carlo OpenCL




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




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 17:58.