




# Accelerated Neural Networks on OpenCL Devices Using SYCL-DNN

John Lawson  
Project lead SYCL-DNN

IWOCL — May 2019

# ImageNet classification error




ImageNet Large Scale Visual Recognition Challenge [image-net.org/challenges/LSVRC](http://image-net.org/challenges/LSVRC).



# 2012: Alexnet 60 million parameters



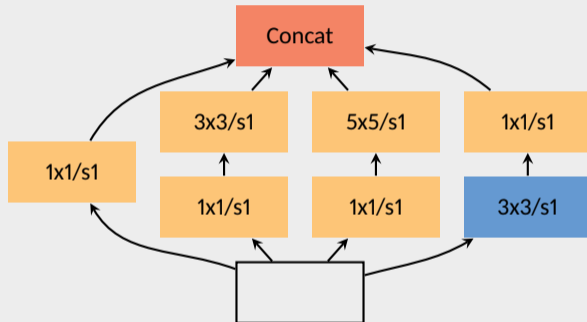
# 2014: VGG 138 million parameters





# 2014: GoogLeNet 5 million parameters




# Inception block



# 2014: GoogLeNet 5 million parameters

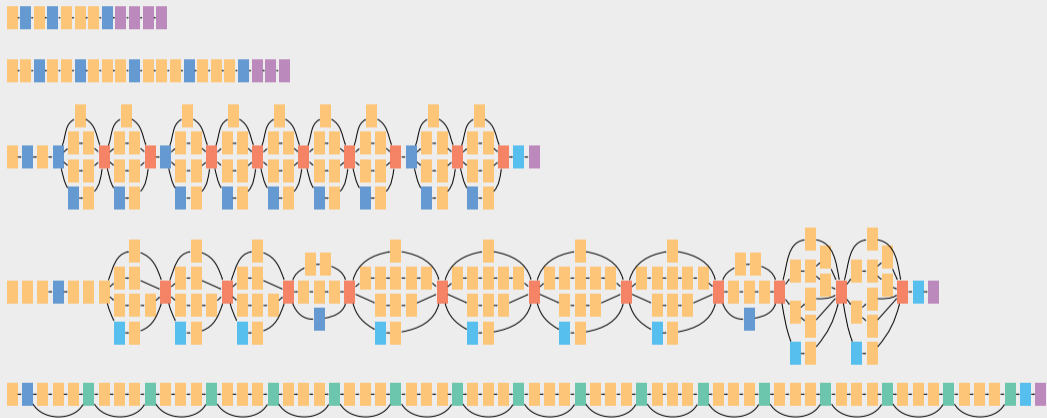


# 2015: InceptionV2 23 million parameters





# 2015: ResNet-50 25 million parameters



# Resnet block






# 2015: ResNet-152 60 million parameters




# Accelerator adoption TOP500 supercomputers




**arm** COMPUTE LIBRARY

AMD MiOpen

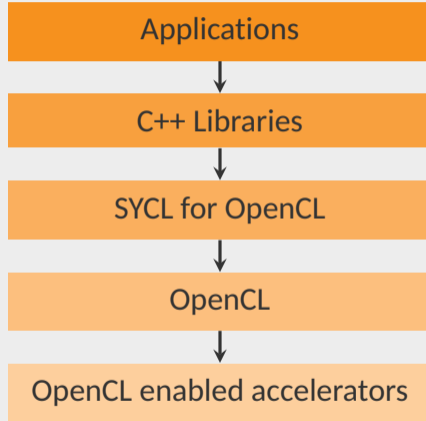
Intel MKL-DNN





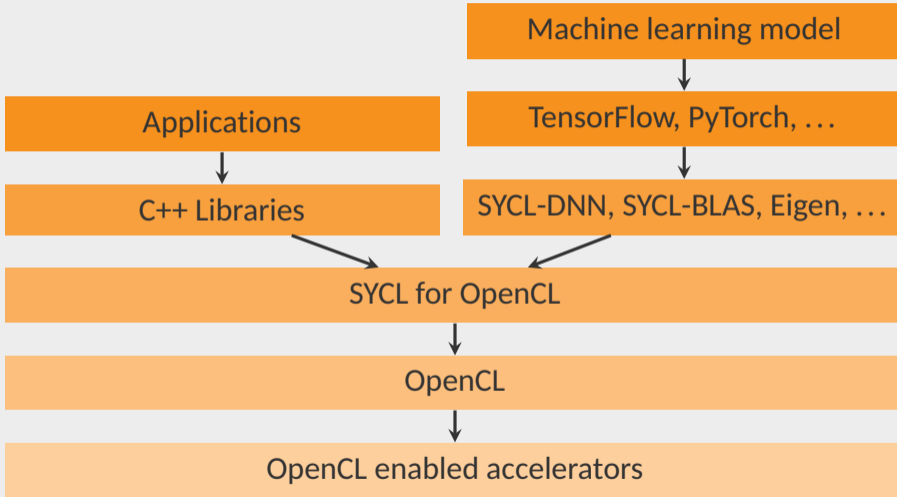
Royalty free open standards from the Khronos Group.

# SYCL Ecosystem





# SYCL Ecosystem



codeplaysoftware / SYCL-DNN

Watch

15

★ Star

16

Fork

2

Code

Issues 0

Pull requests 0

Insights

The SYCL-DNN neural network acceleration library.

sycl

cpp

cpp11

cplusplus

opencl

gpgpu

machine-learning

neural-network

239 commits

1 branch

5 releases

3 contributors

Apache-2.0

# Supported operations

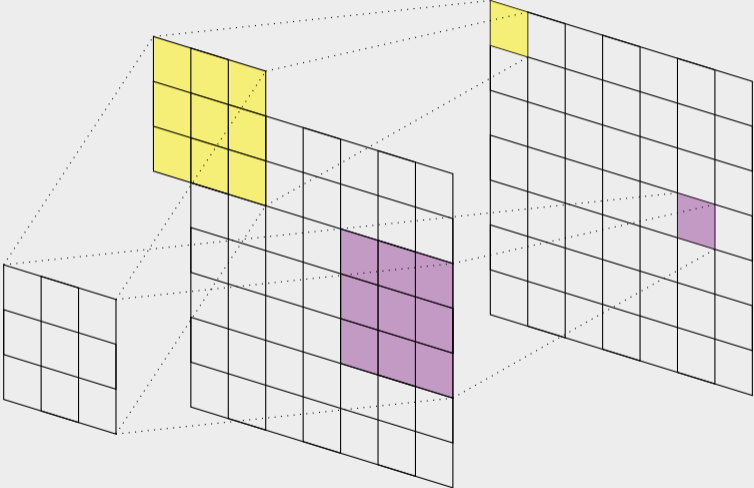
2D convolutions

2D depthwise convolutions

2D max & average pooling

Relu & tanh activations

# 2D Convolution

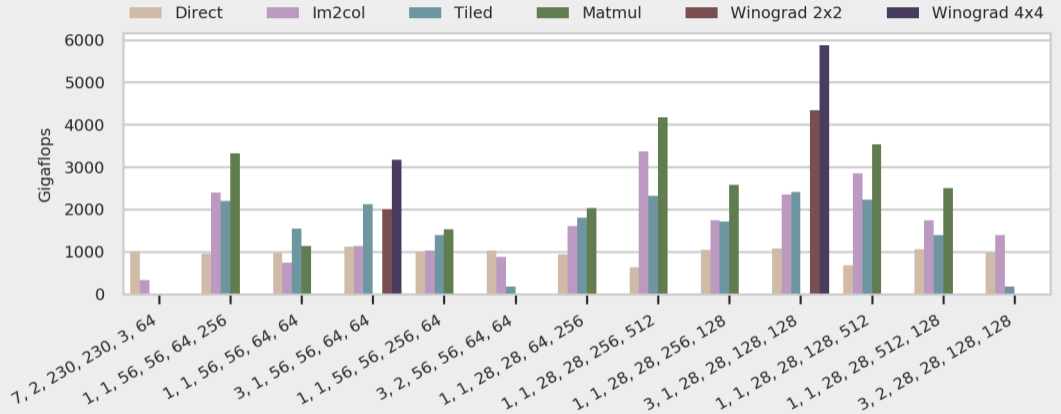


# Highly parameterized kernels

```
template <typename T, typename Index, typename ConvType, int TileRows,  
         int TileCols, int ChannelVectorWidth, int FeatureVectorWidth,  
         bool UseFastDiv, int WindowRows, int WindowCols, int Stride>  
struct TiledConv2D;
```

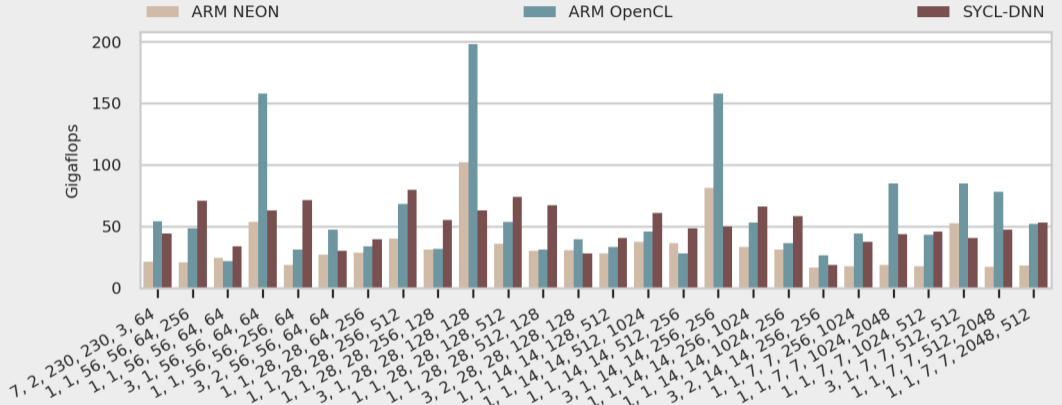
```
template <typename T, typename Index, int ChannelVector, int M, int N,  
         int R, int S, typename ConvType>  
struct WinogradInputTiles;
```

# Different algorithms First layers in ResNet



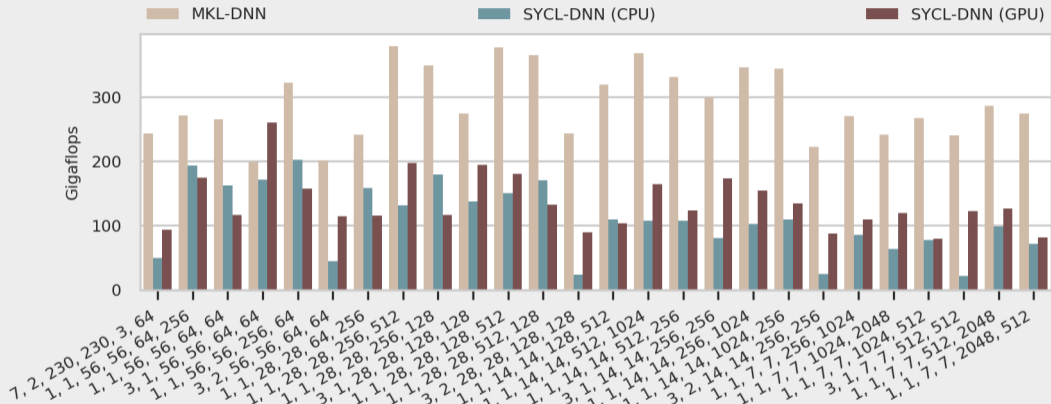
Run on AMD R9 Nano, batch size 32.

# Performance on ARM HiKey 960



ARM NEON and OpenCL from ARM Compute Library v18.11.

# Performance on Intel i7-6700K



MKL-DNN v1.0-pc running on CPU.



# Still to come

Performance improvements


Quantized integer support

Integration of third party libraries

More operations

...



 [codeplay.com](https://codeplay.com)

 [codeplaysoftware/sycl-dnn](https://github.com/codeplaysoftware/sycl-dnn)

 [info@codeplay.com](mailto:info@codeplay.com)

 [@codeplaysoft](https://twitter.com/codeplaysoft)