



# OPEN TECH THOUGHTS

AV DATASETS & MULTIGPU TRAINING

Dr. Adolf Hohl

4.7.2018

# NVIDIA GPU TECHNOLOGY CONFERENCE

**OCTOBER 9 - 11, 2018 ... in MUNICH!!!**

NVIDIA's GPU Technology Conference (GTC) Europe is part of the largest global series of events focused on artificial intelligence and its applications across many important fields.

Join us in Munich and discover the latest breakthroughs in autonomous vehicles, high performance computing, smart cities, healthcare, big data, virtual reality and more.

<https://www.nvidia.com/en-eu/gtc/>

**Enjoy Early Bird Pricing until 10th July and use following code for 25% discount:**

**NVAHOHL**




# END-TO-END SYSTEM FOR AV

## COLLECT DATA




## TRAIN MODELS




## SIMULATE



## RE-SIMULATE



## MAPPING




# RISING DATASETS AND DEMANDS


Estimate (moving):

1. Amount of training data
2. Training data size

Minibatching - memorybound




# RISING DATASETS AND DEMANDS





# RISING DATASETS AND DEMANDS

## NVIDIA Innovations

1. 32 GB Memory
2. TensorCores & Mixed Precision
3. MultiGPU
4. NVLINK/NVSWITCH



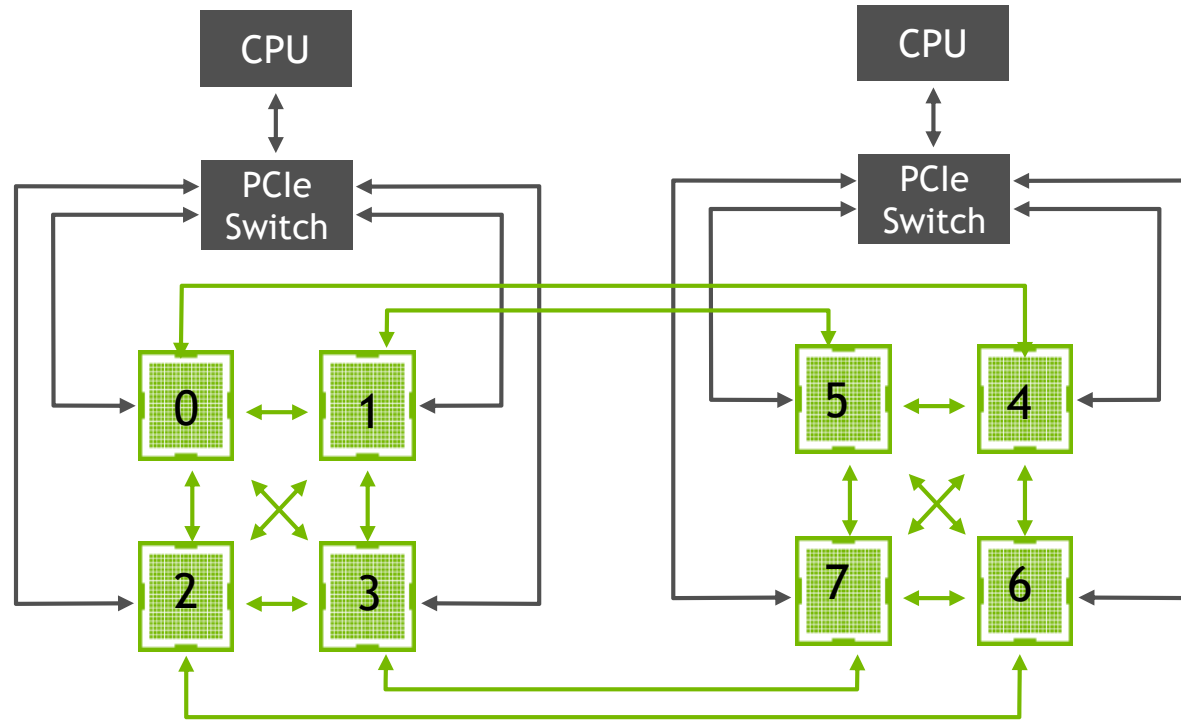
# GPU COMPUTING TO THE RESCUE



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten New plot and data collected for 2010-2015 by K. Rupp

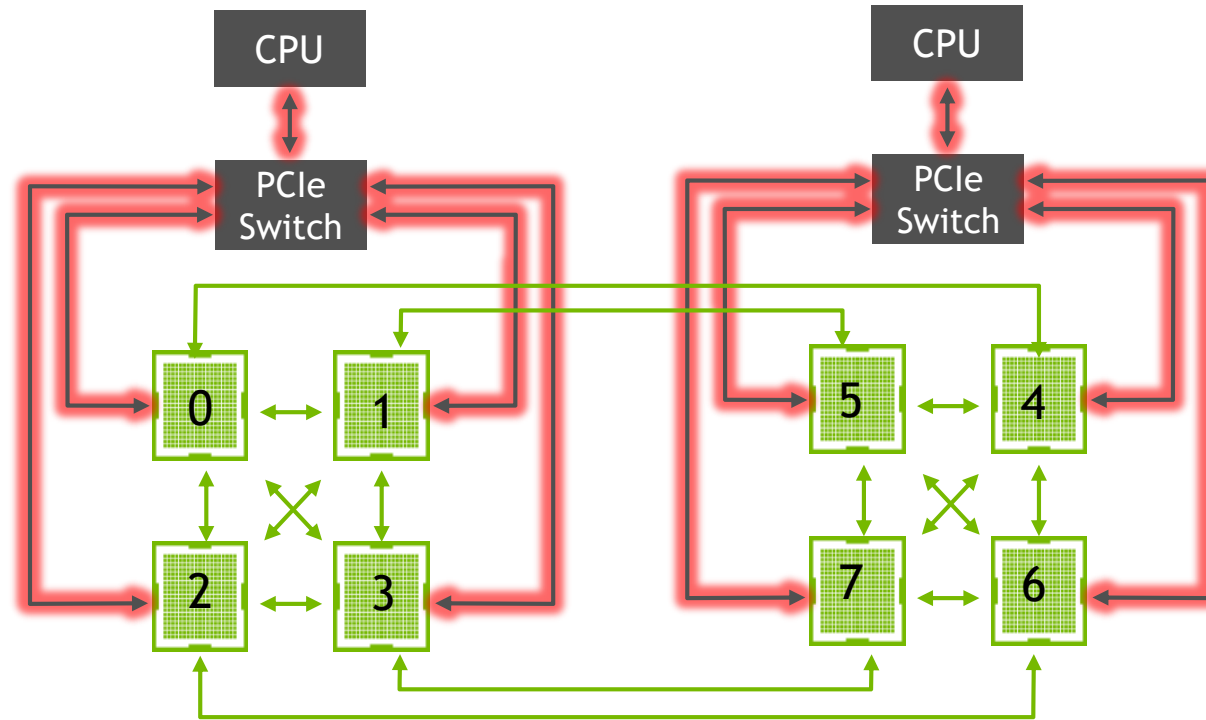
# DATA PARALLELISM - NVLINK

## A KEY TO SCALE



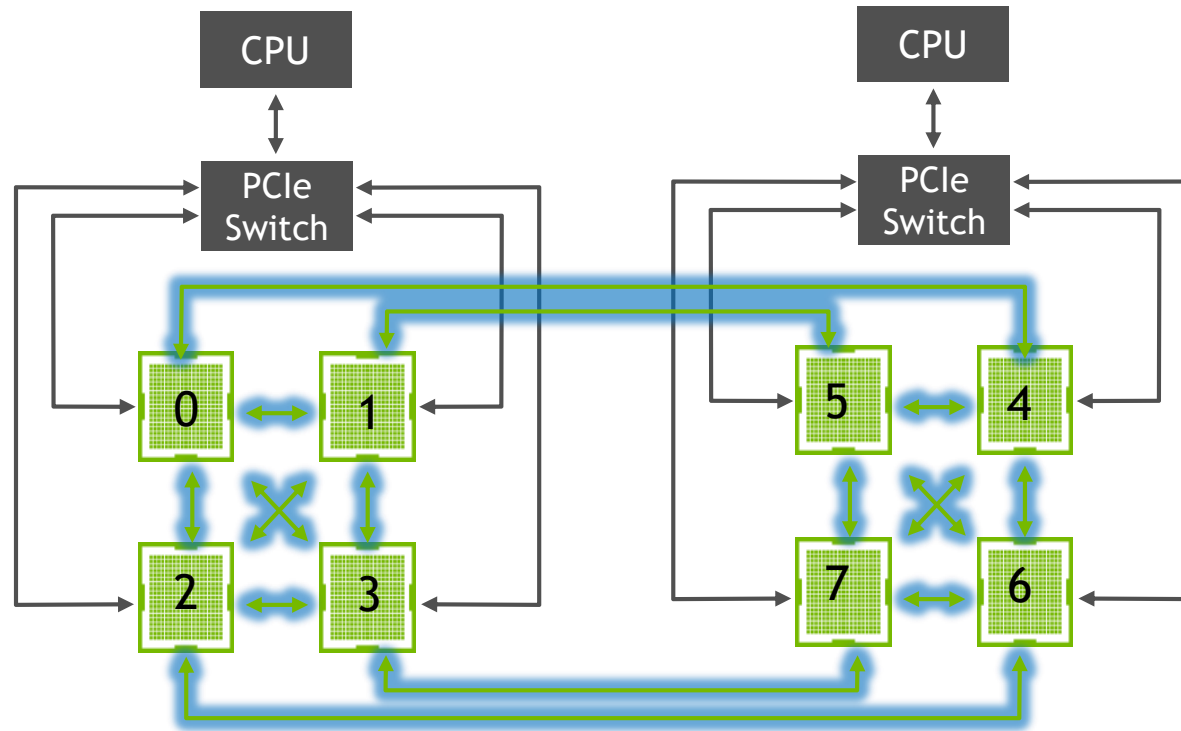


# DATA PARALLELISM - NVLINK



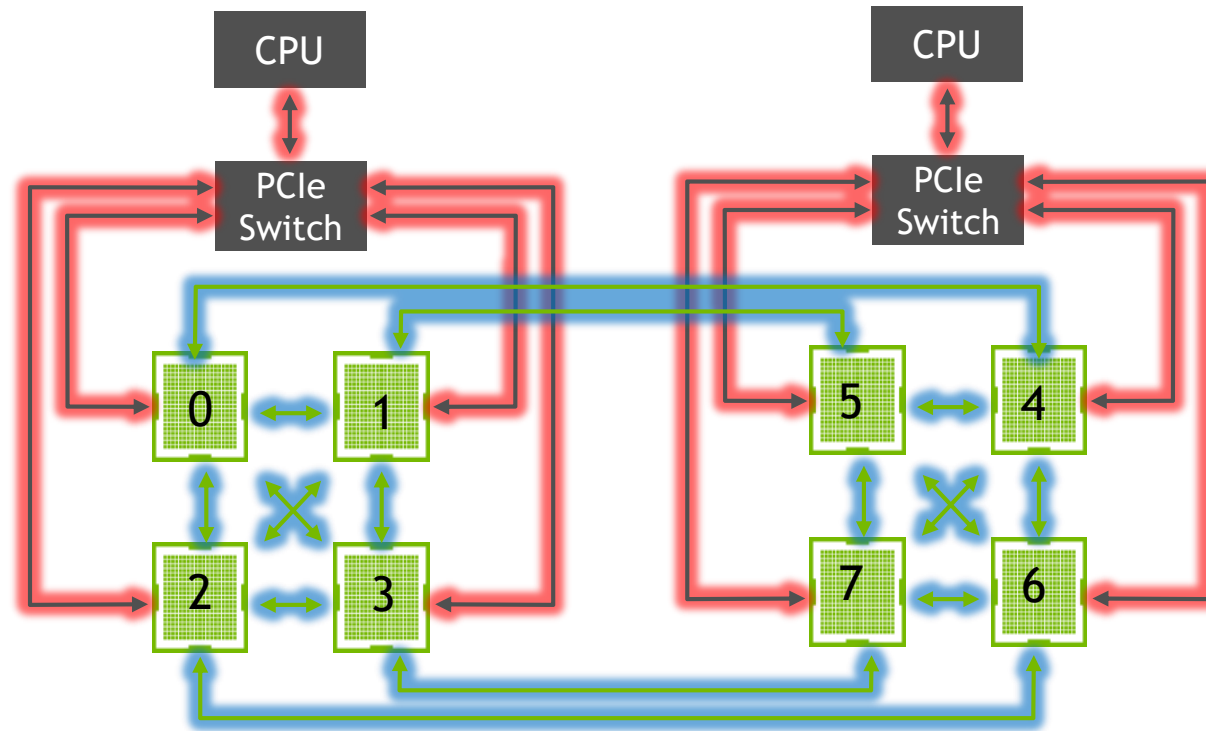
Data loading over PCIe

# DATA PARALLELISM - NVLINK



Gradient averaging over NVLink

# SPEEDUP TX TO NCCL AND NVLINK



Effective use of communication resources: No congestion

# MULTIGPU TRAINING /W TF/KERAS

## Recipe 1 - „Press the Sport Button“

1. Use the NVIDIA Container (Horovod, NCCL and MPI is installed)
2. Keep your model
3. Adapt the training procedure to make use of Horovod

1. import horovod.tensorflow as hvd
2. hvd.init()
3. Per tower/global optimizer
4. get\_or\_create\_global\_step()
5. global\_optimizer.minimize

```
7
8 import DenseNet as dn
9 slim = tf.contrib.slim
10
11 tf.logging.set_verbosity(tf.logging.INFO)
12
13 growth_k = 12
14 nb_block = 2 # how many (dense block + Transition Layer) ?
15 class_num = 2
16
17 IMAGE_HEIGHT, IMAGE_WIDTH = 223,223
18 class_directories= ["mouse-click-events/pos", "mouse-click-events/neg"]
19
20 def grab_dir(path):=
21
22 def Xy_shuffle_loader(pos_dir, neg_dir):=
23
24 def main():
25
26     hvd.init()
27
28     X_train, y_train, X_validation, y_validation = Xy_shuffle_loader(class_directories[0], class_directories[1])
29
30     graph = tf.Graph()
31
32     with graph.as_default():
33
34         tf.train.create_global_step()
35
36         images = tf.placeholder(tf.float32, shape=[None, IMAGE_HEIGHT, IMAGE_WIDTH, 3], name='input_image')
37         labels = tf.placeholder(tf.float32, shape=[None, class_num], name='training_labels')
38         training_flag = tf.placeholder(tf.bool, name='training_flag')
39
40         logits = dn.DenseNet(x=images, nb_blocks=nb_block, filters=growth_k, training=training_flag).model
41         logits = tf.identity(logits, name='output_logits')
42
43         loss = tf.reduce_mean(tf.nn.softmax_cross_entropy_with_logits_v2(logits=logits, labels=labels))
44
45         tower_optimizer = tf.train.RMSPropOptimizer(0.001 * hvd.size())
46         global_optimizer = hvd.DistributedOptimizer(tower_optimizer)
47
48         global_step = tf.train.get_or_create_global_step()
49
50         train_op = global_optimizer.minimize(loss=loss, global_step=global_step)
```

# MULTIGPU TRAINING /W TF/KERAS

## Recipe 1 - „Press the Sport Button“

1. sync\_hook to create stopping criteria and some logging
2. Some config (tf.ConfigProto())
3. Training session
  1. Use the Horovod tailored session
  2. Check if there is still data in the pipeline
4. Train

→ Your new bottleneck is likely the data pipeline now!

```
83 train_op = global_optimizer.minimize(loss=loss, global_step=global_step)
84 train_prediction = tf.nn.softmax(logits=logits)
85
86 sync_hook = [
87     hvd.BroadcastGlobalVariablesHook(0),
88     tf.train.StopAtStepHook(last_step=20000 // hvd.size()),
89     tf.train.LoggingTensorHook(tensors={'step': global_step, 'loss': loss}, every_n_iter=10),
90 ]
91
92 config = tf.ConfigProto()
93 config.gpu_options.allow_growth = True
94 config.gpu_options.visible_device_list = str(hvd.local_rank())
95
96 checkpoint_dir = './checkpoints' if hvd.rank() == 0 else None
97
98 batch_size = 128
99 total_batch = int(len(X_train) / batch_size)
100 step = 0
101
102 # do the following on the same graph as on the one above
103 with tf.train.MonitoredTrainingSession(checkpoint_dir=checkpoint_dir, hooks=sync_hook, config=config) as mtsess:
104
105     while not mtsess.should_stop():
106         if step > total_batch - 1:
107             step = 0
108
109         batch_x = X_train[step * batch_size: (step+1) * batch_size]
110         batch_y = np.eye(class_num)[y_train[step * batch_size: (step+1) * batch_size]]
111         train_feed_dict = {'images': batch_x, 'labels': batch_y, 'training_flag': True}
112
113         _, l, t_predict = mtsess.run([train_op, loss, train_prediction], feed_dict=train_feed_dict)
114         #print("loss:", l)
115
116     mtsess.close()
117
118 if __name__ == "__main__":
119     tf.app.run()
120
```

# MULTIGPU TRAINING /W TF/KERAS

## Recipe 2 - Optimize fuel flow

1. Avoid using `sess.run([ops], feed_dict={...})`
  2. Load data using `tfrecords`
    1. Skeleton for efficient multithreaded input pipeline
    2. Move data augmentation in here
  3. Train
- GPU utilization should rise even more. Check if there is still a data pipeline problem



# A TFRECORD READER

## Recipe 2 - Optimize fuel flow

TFRecord is a binary format

1. Protocol Buffers
2. Deserialization
3. Augmentation
4. Returns graph nodes

```
14 def read_and_decode(filename_queue):
15
16     reader = tf.TFRecordReader()
17
18     _, serialized_example = reader.read(filename_queue)
19
20     features = tf.parse_single_example(serialized_example, features={
21         'height': tf.FixedLenFeature([], tf.int64),
22         'width': tf.FixedLenFeature([], tf.int64),
23         'img_raw': tf.FixedLenFeature([], tf.string),
24         'label': tf.FixedLenFeature([], tf.float32)
25     })
26
27     image = tf.decode_raw(features['img_raw'], tf.uint8)
28     image = tf.cast(image, tf.float32)
29     image = image/256.
30
31     label = tf.cast(features['label'], tf.int32)
32     label = tf.one_hot(label, depth=2)
33
34     height = tf.cast(features['height'], tf.int32)
35     width = tf.cast(features['width'], tf.int32)
36
37     width, height = IMAGE_WIDTH, IMAGE_HEIGHT
38
39     image_shape = tf.stack([height, width, tf.constant(3)], axis = 0)
40     image = tf.reshape(image, [height, width, 3])
41     image = tf.image.random_flip_left_right(image, seed=42)
42     image = tf.image.random_flip_up_down(image, seed=42)
43
44     vrot = 20
45     degrees = np.random.random() * vrot - vrot/2
46
47     image = tf.contrib.image.rotate(image, degrees * np.pi / 180., interpolation='BILINEAR')
48     #image_size_const = tf.constant([IMAGE_HEIGHT, IMAGE_WIDTH, 3], dtype=tf.int32)
49     #transform_image = tf.image.resize_image_with_crop_or_pad(image=image, target_height=IMAGE_HEIGHT, target_width=IMAGE_WIDTH)
50
51     images, labels = tf.train.shuffle_batch([image, label], batch_size=20, capacity=30, num_threads=16, min_after_dequeue=100)
52
53     return images, labels
54
```

# MULTIGPU TRAINING /W TF/KERAS

## Recipe 2 - Optimize fuel flow

1. We use the same model as before
2. Replace images and labels placeholders
3. TF Record nodes appear directly on the graph

```
1 import numpy as np
2 import skimage.io as io
3 import tensorflow as tf
4 import horovod.tensorflow as hvd
5
6 import DenseNet as dn
7 slim = tf.contrib.slim
8
9 tf.logging.set_verbosity(tf.logging.INFO)
10
11 IMAGE_HEIGHT, IMAGE_WIDTH = 223,223
12 tfrec_fname = "risse.tfrecords"
13
14 def read_and_decode(filename_queue):
15     growth_k = 12
16     nb_block = 2 # how many (dense block + Transition Layer) ?
17     class_num = 2
18
19     def main():
20         hvd.init()
21
22         graph = tf.Graph()
23
24         with graph.as_default():
25             filename_queue = tf.train.string_input_producer([tfrec_fname], num_epochs=10)
26             images, labels = read_and_decode(filename_queue)
27             training_flag = tf.placeholder(tf.bool, name="training_flag")
28             tf.train.create_global_step()
29
30             images = tf.identity(images, name='input_image')
31
32             logits = dn.DenseNet(x=images, nb_blocks=nb_block, filters=growth_k, training=training_flag).model
33
34             assert(logits.shape == labels.shape)
35
36             loss = tf.reduce_mean(tf.nn.softmax_cross_entropy_with_logits_v2(logits=logits, labels=labels))
37
38             tower_optimizer = tf.train.RMSPropOptimizer(0.001 * hvd.size())
39             global_optimizer = hvd.DistributedOptimizer(tower_optimizer)
40
41             global_step = tf.train.get_or_create_global_step()
42
43             train_op = global_optimizer.minimize(loss=loss, global_step=global_step)
44             train_prediction = tf.nn.softmax(logits=logits)
```

# MULTIGPU TRAINING /W TF/KERAS

## Recipe 2 - Optimize fuel flow

1. Skipped manual data loading
2. Start reader threads to put data into our input nodes
  1. `tf.train.Coordinator()`
  2. `tf.train.start_queue_runners()`
  3. Reference input nodes as ops
3. Graceful shutdown
4. Train

→ Check GPU/CPU utilization. Further bottleneck analysis required?

```
88     sync_hook = [  
89         hvd.BroadcastGlobalVariablesHook(0),  
90         tf.train.StopAtStepHook(last_step=20000 // hvd.size()),  
91         tf.train.LoggingTensorHook(tensors={'step': global_step, 'loss': loss}, every_n_iter=10),  
92     ]  
93  
94     config = tf.ConfigProto()  
95     config.gpu_options.allow_growth = True  
96     config.gpu_options.visible_device_list = str(hvd.local_rank())  
97  
98     checkpoint_dir = './checkpoints' if hvd.rank() == 0 else None  
99  
100    # do the following on the same graph as on the one above  
101    with tf.train.MonitoredTrainingSession(checkpoint_dir=checkpoint_dir, hooks=sync_hook, config=config) as mtsess:  
102        coord = tf.train.Coordinator()  
103        threads = tf.train.start_queue_runners(sess=mtsess, coord=coord)  
104  
105        while not mtsess.should_stop():  
106            img, lbls, _, l, t_predict = mtsess.run([images, labels, train_op, loss, train_prediction], feed_dict={training_images: images, training_labels: lbls})  
107            #print("loss:", l)  
108  
109            coord.request_stop()  
110            coord.join(threads)  
111  
112            mtsess.close()  
113  
114    if __name__ == "__main__":  
115        tf.app.run()  
116
```

```
root@8eb16c5b074d:/workspace/home/Risse# mpirun -np 8 --allow-run-as-root -H localhost:8 -bind-to none -map-by slot -x NCCL_DEBUG=INFO -x LD_LIBRARY_PATH -x PATH -mca pm1 obi -mca btl ^openib python trainontfrecord.py  
u* 16.04 0:--* 1:-- 2:-- 40d1h 5.34 80x2.1GHz 503.8G1% 2018-06-13 16:13:11
```

# CALL TO ACTION

Roll up the sleeves!

NVIDIA GPU CLOUD: Register free @ [NGC.NVIDIA.COM](https://ngc.nvidia.com)

NVIDIA Docker Runtime: [github nvidia-docker](https://github.com/nvidia/nvidia-docker)

Turn your training code into MultiGPU capable training code!

Register for GTC - join more sessions on MultiGPU and Tensorcores!

<https://www.nvidia.com/en-eu/gtc/>

**Enjoy Early Bird Pricing until 10th July and use following code for 25% discount:**

**NVAHOHL**


# NVIDIA DEEP LEARNING INSTITUTE

Hands-on, self-paced and instructor-led training in deep learning and accelerated computing for developers and data scientists

Request onsite instructor-led workshops at your organization: [www.nvidia.com/requestdli](http://www.nvidia.com/requestdli)

Take self-paced courses online:  
[www.nvidia.co.uk/dlilabs](http://www.nvidia.co.uk/dlilabs)


Download the course catalog and view upcoming public workshops:  
[www.nvidia.co.uk/dli](http://www.nvidia.co.uk/dli)




Deep Learning Fundamentals




Autonomous Vehicles




Medical Image Analysis




Genomics



Finance



Game Development & Digital Content



Accelerated Computing Fundamentals

More industry-specific training coming soon...



