





CRAY®

# Scaling Deep Learning

Mike Ringenburg ([mikeri@cray.com](mailto:mikeri@cray.com))

Slides courtesy of Peter Mendygral ([pjm@cray.com](mailto:pjm@cray.com))



# Legal Disclaimer

*Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.*

*Cray Inc. may make changes to specifications and product descriptions at any time, without notice.*

*All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.*

*Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.*

*Cray uses codenames internally to identify products that are in development and not yet publically announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.*

*Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.*

*The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, and URIKA. The following are trademarks of Cray Inc.: APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYPAT, CRAYPORT, ECOPHLEX, LIBSCI, NODEKARE, THREADSTORM, REVEAL. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.*



# Motivation

- **A trained neural network can be a powerful tool for**
  - Pattern recognition
  - Classification
  - Clustering
  - Others...
- **Scaling Deep Learning (DL) training is also a tool for**
  - Models that take a very long time to train (and have a very large training dataset)
  - Increasing the frequency at which models can be retrained with new or improved data
- **This talk reviews scaling DL training and topics that can be important to successfully applying it**


# Agenda

- **De-mystifying Deep Learning**
- **~~TensorFlow on Theta~~ (← covered earlier today)**
- **Parallelization Methods for TensorFlow**
- **Convergence Considerations at Scale**
- **CPE ML Plugin Example**

# De-mystifying Deep Learning

# What is Deep Learning: A Specific Example

- An organic gardener is building a robot in his garage to recognize the 10 insects found in his garden, and decide which ones to kill with a laser
- The robot will have a camera, and will capture JPEG files of the insects
- The robot needs a ‘program’ to classify each JPEG according to which of the 10 kinds of insect was photographed



# Inputs & Outputs

- **Our input is a JPEG**
  - 224x224 pixels, 3 colors → a 224x224x3 element vector of the pixel values
- **Our output is a classification**
  - One of 10 categories → a 10 element vector with a “1” in the position representing the category to which the image belongs

How many “IF” statements will we need to figure out that a bunch of pixel values is a Japanese beetle?



Bruce Marlin  
(CC BY 3.0: Attribution 3.0 Unported)

# This is an Artificial Intelligence Problem

- If you can't get the output from the input with a bunch of loops and conditionals, it's AI
- But, if that won't work, how can we do it?
  
- Hint #1: Any mapping of inputs to outputs is a function
- Hint #2: A function can be approximated using a (good) approximating function



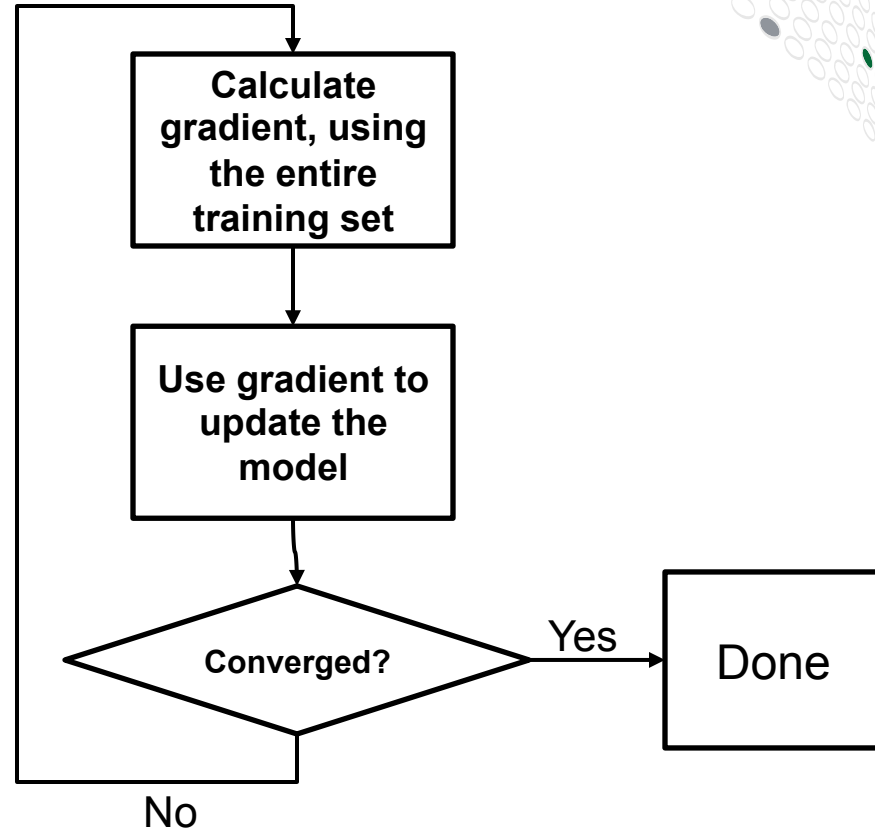
# An Approximating Function

- **How can we determine a good approximating function?**
  - Choose its form (linear, polynomial, ...)
  - Minimize the overall error at a finite number of inputs with known outputs - - **fit the curve**
    - We have to find the values of the free parameters of the function that minimize the error – it doesn't matter how we do it

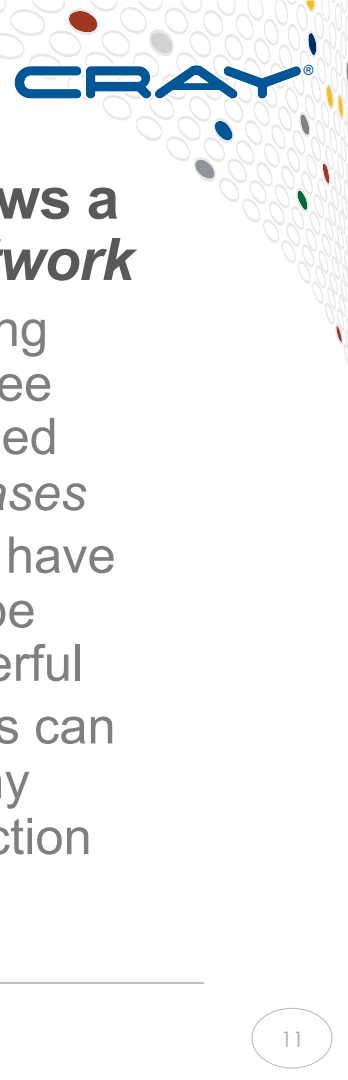
Fitting the curve is a lot like *training* the function to know the answer for arbitrary inputs

# Training via Gradient Descent


- **We want to approximate  $y=f(x)$** 
  - Find a function that maps a set of inputs to a set of outputs, to some level of accuracy
- **We know  $y_i=f(x_i)$ , for  $i=1,N$**
- **Iterate:**
  - First iteration only: initialize the free parameters of  $f$
  - Calculate error (over  $N$  known points)
  - Calculate gradient of error, as a function of the free parameters of  $f$
  - Adjust the free parameters of  $f$  a 'small' distance in the direction of negative of error gradient
  - Assess convergence & stop when 'good enough'



# A Really Useful Kind of Function



Deep neural network



- This image shows a ***deep neural network***

- An approximating function, with free parameters called *weights* and *biases*
- Deep networks have been found to be especially powerful
- Neural networks can approximate any continuous function arbitrarily well

- **DL training is a classic high-performance computing problem which demands:**
  - Large compute capacity in terms of FLOPs, memory capacity and bandwidth
  - A performant interconnect for fast communication of gradients and model parameters
  - Parallel I/O and storage with sufficient bandwidth to keep the compute fed at scale

# Data Parallelism - Collective-based Synchronous SGD

- Data parallel training divides a global mini-batch of examples across processes
- Each process computes gradients from their local mini-batch
- Average gradients across processes
- All processes update their local model with averaged gradients (all processes have the same model)

---

## Algorithm 1 Sync-SGD algorithm

---

for  $0 \leq step < max\_steps$  do

$G_{local} \leftarrow COMPUTE\_GRADIENTS(\text{mini batch})$

$G_{global} \leftarrow 1/N_{ranks} \times ALLREDUCE(G_{local})$

$APPLY\_GRADIENTS(G_{global})$

end for

---

Compute intensive

Communication intensive

Typically not much compute

- Not showing the activity of reading training samples (and possible augmentation)

# Why do we want to scale?

- **Deep Network Training**

- We can strong scale training time-to-accuracy provided
  - Number of workers (e.g., # nodes)  $\ll$  number of training examples
  - Learning rate for particular batch size / scale is known

- **Hyper-Parameter Optimization**

- For problems and datasets where baseline accuracy is not known
  - learning rate schedule
  - momentum
  - batch size
- Evolve topologies if good architecture is unknown (common with novel datasets / mappings)
  - Layer types, width, number filters
  - Activation functions, drop-out rates

# Parallelization Methods for DL

# Parallelization Techniques

- **Data Parallelism**

- As described earlier, divides global mini-batch among processes
- Two methods for this:
  - Synchronous: single model (possibly replicated across all processes) updated with globally averaged gradients every iteration
  - Asynchronous: processes provide gradients every iteration but are allowed to fall out of sync from one another. Processes each have their own model that may or may not be the same as any other process

- **Model Parallelism**

- Single model with layers decomposed across processes
- Activations communicated between processes

- **This talk will focus on synchronous data parallel approach**



# Distributed TensorFlow

- **TensorFlow has a native method for parallelism across nodes**
  - ClusterSpec API
  - Uses gRPC layer in TensorFlow based on sockets
- **Can be difficult to use and optimize**
- **User must specify**
  - hostnames and ports for all worker processes
  - hostnames and ports for all parameter server processes (see next slide)
  - # of workers
  - # of parameter server processes
  - Chief process of workers

# Distributed TensorFlow



- **Number of parameter servers (PS) processes to use is not clear**
  - Too few results in many-to-few comm pattern (very bad) and stalls delivering updated parameters
  - Too many results in many-to-many comm patten (also bad)
- **Users typically have to pick a scale and experiment for best performance**



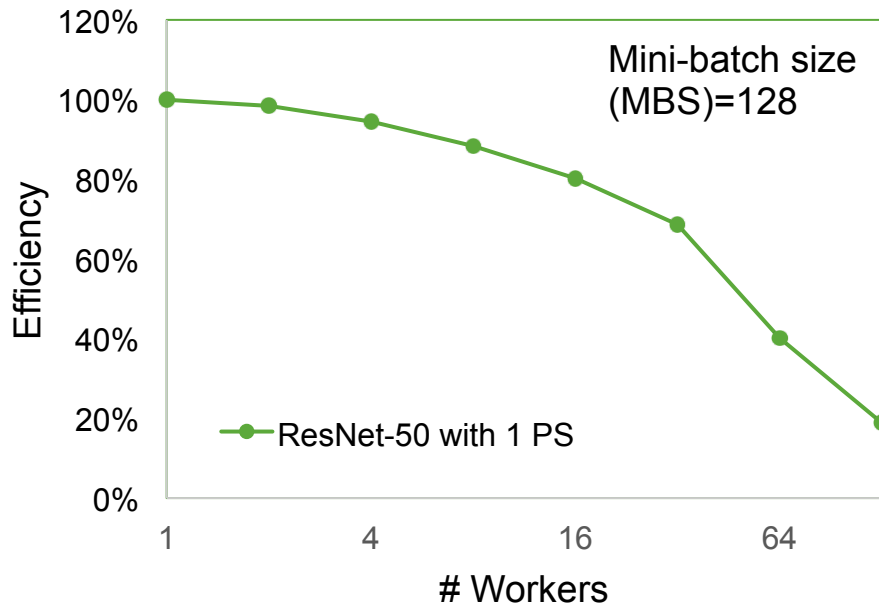


Figure 7: Synchronous and asynchronous data parallel training

# Distributed TensorFlow Scaling on Cray XC40 - KNL



From Mathuriya et al. @ NIPS 2017

COMPUTE

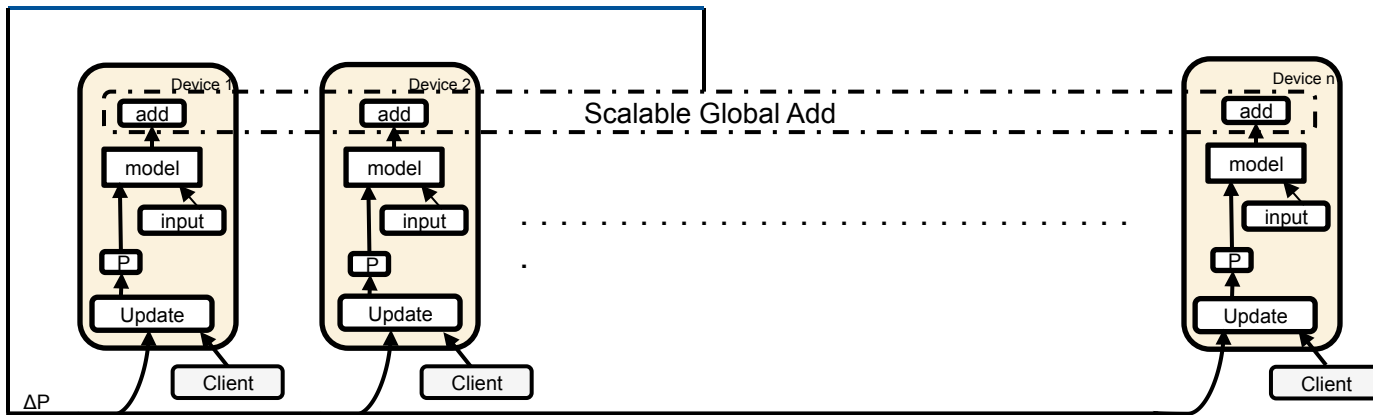
STORE

ANALYZE

# MPI-based Data Parallel TensorFlow

- **The performance and usability issues with distributed TensorFlow can be addressed by adopting an MPI communication model**
- **TensorFlow does have an MPI option, but it only replaces point to point operations in gRPC with MPI**
  - Collective algorithm optimization in MPI not used
- **Other frameworks, such as Caffe and CNTK, include MPI collectives**
- **An MPI collective based approach would eliminate the need for PS processes and likely be optimized without intervention from the user**

# Scalable Synchronous Data Parallelism



- Note there are no PS processes in this model
- Resources dedicated to gradient calculation

# Uber Horovod


- **Uber open source addon for TensorFlow, PyTorch, and Keras-on-TF that replaces native optimizer class with a new class**
  - Horovod adds an allreduce between gradient computation and model update in this class
- **New Python class includes NCCL and MPI collective reductions for gradient aggregation**
- **<https://github.com/uber/horovod>**
- **No modifications to TensorFlow source required**
  - User modifies Python training script instead



# Cray Programming Environment Machine Learning Plugin (CPE ML Plugin)



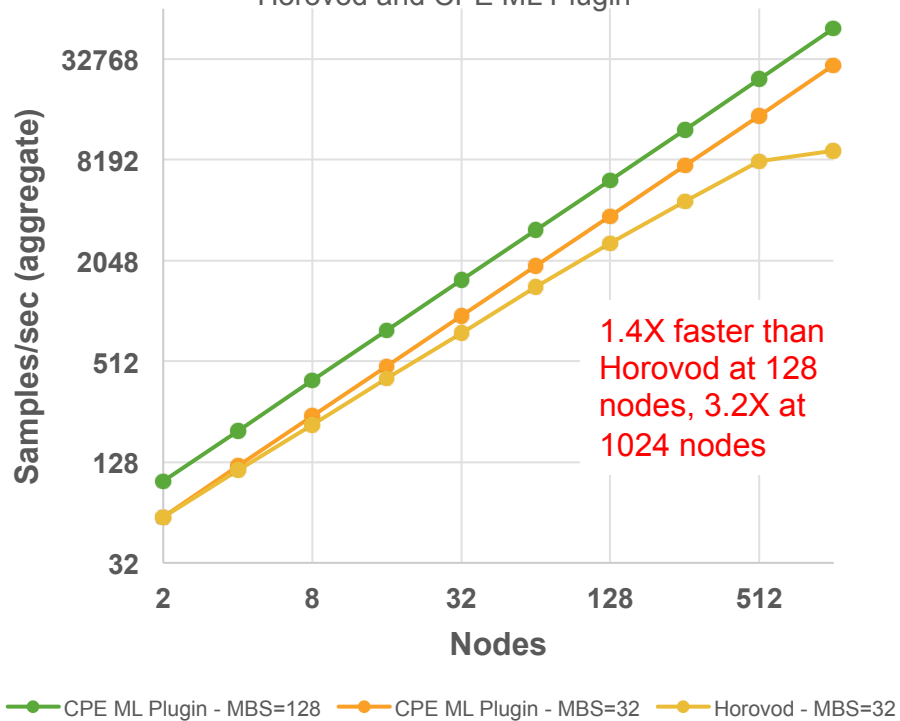
- DL communication plugin with Python and C APIs
- Originally developed for TensorFlow but also portable to other frameworks (PyTorch and Keras tested)
  - Callable from C/C++ source
  - Called from Python if data stored in NumPy arrays or Tensors
- Like Horovod does not require modification to TensorFlow source
  - User modifies training script
- Uses custom allreduce specifically optimized for DL workloads
  - Optimized for Cray Aries interconnect and IB for Cray clusters
- Tunable through API and environment variables
- Supports multiple gradient aggregations at once with thread teams
  - Useful for Generative Adversarial Networks (GAN), for example




# Horovod / CPE ML Plugin – Throughput Scaling



ResNet50 Performance on XC40 (Cori KNL at NERSC)  
Horovod and CPE ML Plugin



Inception v3 Performance on XC50 (Piz Daint at CSCS)  
– CPE ML Plugin ONLY



COMPUTE

STORE

ANALYZE



# Convergence Considerations at Scale



# Problems in Scaling DL Training

- **Increasing workers increases the global batch size**
  - This reduces the number of updates to the model (iterations) per epoch (full pass through dataset)
  - Can require more iterations to converge to same validation accuracy for models trained at smaller batch sizes
- **Large-batch (LB) training can have different convergence properties than Small-batch (SB) training**
  - LB training can lead to models which fail to generalize to validation datasets
  - LB training error can look similar to SB training error, but validation error fails to improve

# Problems in Scaling DL Training

- **Possible reasons for the observed failure to generalize using large batch methods [2]:**
  - The model overfits
  - Optimization is attracted to saddle-points
  - Loss of the explorative properties gained with small batches

# Observations on Scaled Learning Rates

- **Step 1) Start with common initial learning rate for selected optimizer (from Keras documentation)**
  - Adam -> 0.001
  - RMSProp -> 0.001
  - SGD -> 0.01
  - Adagrad -> 0.01
  - Adadelta -> 1.0
  
- **Step 2) Multiply learning rate by N or Sqrt(N)**
  - N is the number of parallel processes
  - Discussed in further detail on next slide
  
- **Step 3) Decay learning rate during training (e.g., exponential decay)**
  - Setup a learning rate schedule using your initial learning rate as the starting state
  - Learning rate typically lowered periodically or continuously
  - Helps improve final accuracy
  - Likely very important to reduce learning rate over time when initial learning rate scaled large
  
- **Step 4) Run it**
  - Train and observe loss or training accuracy, check validation accuracy
  - Adjust initial learning rate up if learning too slowly or down if model is not learning
  - Repeat steps as needed to improve convergence and accuracy

# Learning Rate Scaling Rules

- **Sqrt Scaling Rule:**

- When the local minibatch size is multiplied by  $N_{\downarrow workers}$ , multiply the learning rate by  $\sqrt{N_{\downarrow workers}}$ .

$$\eta_{\downarrow init} = \eta_{init} * \sqrt{N_{\downarrow workers}}$$

- Error on the mean only improves as  $\text{sqrt}(N_{\text{workers}})$

- **Linear Scaling Rule:**

- When the minibatch size is multiplied by  $N$ , multiply the learning rate by  $N$ .

$$\eta_{\downarrow init} = \eta_{init} * N_{\downarrow workers}$$

- Naïve rule for scaling learning rate in distributed training but it works for some problems
- More attractive (when it works) because it shouldn't require many additional iterations to reach same accuracy



# Other Considerations

- **Linearly scaled learning rate causes most problems early in training [3]**
  - Design a warm-up set of iterations to reduce these errors
  - Once training settled on good path, transition to larger learning rate
- **SGD momentum can be useful at scale [1]**
- **Layer-wise Adaptive Rate Scaling (LARS) allows different learning rates at each layer [1]**



# Useful References

[1] LARGE BATCH TRAINING OF CONVOLUTIONAL NETWORKS --  
<https://arxiv.org/pdf/1708.03888.pdf>

[2] ON LARGE-BATCH TRAINING FOR DEEP LEARNING: GENERALIZATION GAP AND SHARP MINIMA -- <https://openreview.net/pdf?id=H1oyRIYgg>

[3] Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour --  
<https://research.fb.com/wp-content/uploads/2017/06/imagenet1kin1h5.pdf>

[4] Train longer, generalize better: closing the generalization gap in large batch training of neural networks -- <https://arxiv.org/pdf/1705.08741.pdf>

[5] Extremely Large Minibatch SGD: Training ResNet-50 on ImageNet in 15 Minutes -- <https://arxiv.org/pdf/1711.04325.pdf>

# CPE ML Plugin Example




# Training Script Modifications

- **Both Horovod and CPE ML Plugin require some modifications to a serial training script**
- **For the CPE ML Plugin the changes are**
  - Importing the Python module
  - Initialize the module
    - Possibly configure the thread team(s) for specific uses
  - Broadcast initial model parameters
  - Incorporate gradient aggregation between gradient computation and model update
  - Finalize the Python module



# MNIST Example

- Dataset of handwritten digits from 0-9
- Simple CNN can be used to identify handwritten digits



- This example is adapted from the TensorFlow official MNIST example
- <https://github.com/tensorflow/models/tree/master/official/mnist>
- Modified script included with CPE ML Plugin
  - `module load /projects/datascience/kristyn/modulefiles/craype-ml-plugin-py3/1.1.2`
  - `module load tensorflow`
  - `less $CRAYPE_ML_PLUGIN_BASEDIR/examples/tf_mnist/mnist.py`



# CPE ML Plugin – Execution Example

- Once the script is modified job launch is just like a typical MPI job

```
module load /projects/datascience/kristyn/modulefiles/craype-ml-plugin-py3/1.1.2
module load tensorflow
```

```
aprun -n4 -N1 -cc none -b python \  
$CRAYPE_ML_PLUGIN_BASEDIR/examples/tf_mnist/mnist.py \  
--enable_ml_comm \  
--data_dir=/lus/theta-fs0/projects/SDL_Workshop/mendygra/mnist_data \  
--model_dir=[train dir]
```

The Cray logo is rendered in a bold, blue, sans-serif font. The letters are closely spaced, with the 'Y' having a distinctive shape where the two bottom strokes are separated.

**CRAY**

COMPUTE



STORE



ANALYZE

The background is a complex, abstract digital landscape. It features a grid of glowing white dots that recede into the distance, creating a sense of depth. Overlaid on this are various blue and white digital elements, including binary code (0s and 1s), glowing lines, and a central bright light source that creates a lens flare effect. The overall color palette is dominated by blues and whites, giving it a high-tech, futuristic feel.

**Questions?**

**Thank You!**

# Backup

# CPE ML Plugin - Import

- Access the Python API by importing the module

```
import tensorflow as tf
# CRAY ADDED
import ml_comm as mc
import math
#
```

# CPE ML Plugin - Initialization

- **Compute the number of trainable variables in the model**
  - Required for the CPE ML Plugin to pre-allocate needed communication buffers
  - Example sets up a single thread team with one thread

```
# CRAY ADDED
if FLAGS.enable_ml_comm:

    # initialize the Cray PE ML Plugin
    tosize = sum([reduce(lambda x, y: x*y, v.get_shape().as_list()) for v in tf.trainable_variables()])
    mc.init(1, 1, tosize, "tensorflow")
```

# CPE ML Plugin – Team Configuration

- Set the maximum number of steps (mini batches) to train for
  - Verbose output every 200 steps
- Also set output path to rank-specific location

```
# config the thread team (correcting the number of epochs for the effective batch size)
FLAGS.train_epochs = int(FLAGS.train_epochs / mc.get_n_ranks())
max_steps = int(math.ceil(FLAGS.train_epochs *
    ( NUM_IMAGES['train'] + NUM_IMAGES['validation'] ) / FLAGS.batch_size))
mc.config_team(0, 0, 100, max_steps, 2, 200)
```

```
# give each rank its own directory to save in
FLAGS.model_dir = FLAGS.model_dir + '/rank' + str(mc.get_rank())
```



# CPE ML Plugin – Broadcast Initial Model

- Broadcast initial model parameter values from rank 0 to all other ranks
- Then assign broadcasted values locally

```
# CRAY ADDED
# since this script uses a monitored session, we need to create a hook to initialize
# variables after the session is generated
class BcastTensors(tf.train.SessionRunHook):

    def __init__(self):
        self.bcast = None

    def begin(self):
        if not self.bcast:
            new_vars = mc.broadcast(tf.trainable_variables(),0)
            self.bcast = tf.group(*[tf.assign(v,new_vars[k]) for k,v in enumerate(tf.trainable_variables())])
```

# CPE ML Plugin – Gradient Aggregation

- Perform gradient averaging across all ranks between local gradient calculation and model update

```
# CRAY ADDED
if FLAGS.enable_ml_comm:

    # we need to split out the minimize call below so we can modify gradients
    grads_and_vars = optimizer.compute_gradients(loss)

    grads      = mc.gradients([qv[0] for qv in grads_and_vars], 0)
    gs_and_vs  = [(g,v) for (_,v), g in zip(grads_and_vars, grads)]

    train_op = optimizer.apply_gradients(gs_and_vs,
                                         global_step=tf.train.get_or_create_global_step())
# END CRAY ADDED
```

# CPE ML Plugin – Finalize

- After all training steps are complete clean up data structures and MPI

```
# CRAY ADDED  
if FLAGS.enable_ml_comm:  
    mc.finalize()  
# END CRAY ADDED
```