

FireCaffe: near-linear acceleration of deep neural network training on compute clusters

Forrest N. Iandola, Khalid Ashraf, Matthew W. Moskewicz, Kurt Keutzer

DeepScale* and UC Berkeley

{forresti, ashrafkhalid, moskewicz, keutzer}@eecs.berkeley.edu

Abstract

Long training times for high-accuracy deep neural networks (DNNs) impede research into new DNN architectures and slow the development of high-accuracy DNNs. In this paper we present FireCaffe, which successfully scales deep neural network training across a cluster of GPUs. We also present a number of best practices to aid in comparing advancements in methods for scaling and accelerating the training of deep neural networks. The speed and scalability of distributed algorithms is almost always limited by the overhead of communicating between servers; DNN training is not an exception to this rule. Therefore, the key consideration here is to reduce communication overhead wherever possible, while not degrading the accuracy of the DNN models that we train. Our approach has three key pillars. First, we select network hardware that achieves high bandwidth between GPU servers – Infiniband or Cray interconnects are ideal for this. Second, we consider a number of communication algorithms, and we find that reduction trees are more efficient and scalable than the traditional parameter server approach. Third, we optionally increase the batch size to reduce the total quantity of communication during DNN training, and we identify hyperparameters that allow us to reproduce the small-batch accuracy while training with large batch sizes. When training GoogLeNet and Network-in-Network on ImageNet, we achieve a 47x and 39x speedup, respectively, when training on a cluster of 128 GPUs.

1. Introduction and Motivation

Since the publication of AlexNet [29], a variety of new deep neural network (DNN) architectures such as GoogleNet [41], Network-in-Network [32], and VGG [39] have been developed at a rapid pace. This is natural, because with the training and testing dataset fixed (e.g. ImageNet-1k [13]), it is the DNN architecture that is primarily responsible for improvements in accuracy. In other words, the race for improvements in accuracy in image classification and other contemporary problems of computer science, has become a race in the development of new DNN architectures. So, what is the bottleneck in the development of new architectures?

In the development of new DNN architectures, as in any human research endeavor, creativity is a key element. However, the impact of architectural variations in DNNs – such as number of layers, filter dimensions, and so forth – can be hard to predict, and experimentation is required to assess their impact. A high-accuracy deep neural network (DNN) model such as GoogLeNet [41] can take weeks to train on a modern GPU. This is true even when leveraging deep neural network primitives like cuDNN [7], maxDNN [30], or fbfft [43] – all of which operate near the theoretical peak computation per second achievable on GPUs. Thus, training time is a key challenge at the root of the development of new DNN architectures. This sentiment was voiced by Jeffrey Dean of Google in his recent keynote address [10].

The four key points that Dean makes are:

- We [i.e. DNN researchers and users] want results of experiments quickly
- There is a “patience threshold:” No one wants to wait more than a few days or a week for a result
- This significantly affects scale of problems that can be tackled
- We sometimes optimize for experiment turnaround time, rather than absolute minimal system resources for performing the experiment

*<http://deepscale.ai>

Given the considerable resources available to Google researchers, Dean’s comments indicate that simply throwing more computational resources at the problem is not sufficient to solve the DNN training problem. In the following, we will spend a little more time dimensionalizing the current problems with DNN training and the upside potential if these problems can be solved.

1.1. Accelerating DNN Research and Development

As a particular example of where long training times are limiting the pace of DNN research and productization, consider the following. ImageNet-1k has 1.2 million training images, distributed across 1000 different category labels. From first-hand conversations with engineers and executives, we know that several internet companies have internal databases containing billions of images with hundreds of thousands of different category labels. Due to long training times, these companies are facing serious delays in bringing DNN-based solutions to market. Accelerated DNN training solutions would solve a major pain point for these companies.

1.2. Real-Time DNN Training

So far, we have argued how accelerating DNN training would benefit applications where DNNs are in use today. Now, we consider ways in which accelerating DNN training would allow DNN-based techniques to be applied in entirely new ways. There are a number of situations where it is crucial to incorporate new data into a DNN model in real time. For example, reinforcement learning (RL) enables robots to learn things themselves with minimal supervision. A recent study by Levine *et al.* applied state-of-the-art DNN-based RL techniques to enable a robot to teach itself how to build lego structures and screw on bottle caps [31]. This technique is effective, and the robot does indeed learn to screw on bottle caps. However, it takes 3-4 hours for the robot to learn to screw on bottle caps, and the majority of this time is spent on DNN training. Faster DNN training would enable this and other reinforcement learning applications to move toward real-time.

Deep Neural Networks are used for an ever-broadening variety of problems, including classifying [41, 24] and detecting [22, 17] objects in images, writing sentences about images [15], identifying actions in videos [4], performing speech recognition [21], and gaining semantic understanding of text [33]. We anticipate that sophisticated reinforcement learning (RL) systems in robotics will eventually leverage all of these modalities, ideally in real-time.

1.3. Accelerating DNN Training with FireCaffe

In our work, we focus directly on the problem of DNN training. Since single-GPU efficiency has reached the hard limits of the hardware, the next frontier for accelerating DNN training is to scale it across a compute cluster. In this paper, we present FireCaffe, which scales DNN training across a cluster of 128 GPUs with speedups of more than 40x compared to a single GPU. Our strategy for scaling up DNN training is to focus on reducing communication overhead, and we make a number of design choices toward this goal. For example, we use fast interconnects such as Infiniband or Cray Gemini to accelerate communication among the GPUs. We also show that reduction trees are a faster method for communication than using parameter servers. We map our parallelization strategy to high-accuracy DNN architectures that require less communication.

The rest of the paper is organized as follows. In Section 2, we describe our choice of hardware for evaluating scalable DNN training, and Section 3 introduces key factors that we will use for analyzing communication among GPU workers. We describe tradeoffs between DNN parallelism strategies in Section 4, and Section 5 explains why certain high-accuracy DNN architectures are particularly amenable to parallelism. In Section 6, we describe our approach to efficiently implementing distributed DNN training. In Section 7, we describe good practices that facilitate the comparison of scalable DNN training techniques and we present our speedups for training the NiN and GoogLeNet architectures on ImageNet. Section 8 describes approaches that are complimentary to FireCaffe for further accelerating DNN training. Finally, we conclude in Section 9.

2. Hardware for scalable DNN training

It is both useful and possible to experiment with the scalability of DNN computations using theoretical or scale models. However, demonstration and verification of the correctness and real-world scalability of the proposed FireCaffe system requires using concrete hardware platforms. The speed at which data can be sent between nodes is a key consideration in selecting a hardware platform for scalable DNN training. This is because, the faster the interconnect between nodes is, the more scale we can achieve without being dominated by communication overhead. Hardware manufacturers such as Cray and Mellanox address this by developing high-bandwidth, low-latency interconnects that are substantially faster than typical Ethernet connections.

For example, the Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF) has a high bandwidth, low latency Cray Gemini interconnect for communication among servers. The Titan supercomputer has a total of 18000 servers,

with one NVIDIA Kepler-based K20x GPU per server [2, 35]. With this in mind, we choose the OLCF Titan supercomputer for tuning and evaluating FireCaffe.

In this research, we use relatively small slices of the overall capacity of Titan for each training run. The additional computational capacity (~27 PetaFLOPS/s in total) enables us to conduct multiple training runs concurrently, where each training run utilizes 32 to 128 GPUs. When considering 32-node slices of Titan, we found that the interconnect speed (at least for the applications of this work) is similar to that provided by having all nodes in the slice connected to a single Infiniband-class switch.

3. Preliminaries and terminology

Deep neural network training is comprised of iterating between two phases: forward and backward propagation. In the forward phase, a batch of data items (e.g. images) is taken from the training set, and the DNN attempts to classify them. Then comes the backward phase, which consists of computing gradients with respect to the weights (∇W) and gradients with respect to the data (∇D). The weight gradients are used to update the model’s weights. Then, an other forward phase is performed, and so on. We train models using batched stochastic gradient descent (SGD), which is the standard choice for popular DNN models such as GoogLeNet [41].

We now present a few preliminaries, which we will use later in the paper for reasoning about data volume to communicate in distributed DNN training. In Equation 1, we show how to calculate the total size (in bytes) of the weights in all convolutional and fully-connected layers, combined.

$$|W| = \sum_{L=1}^{\#layers} ch_L * numFilt_L * filterW_L * filterH_L * 4 \tag{1}$$

where ch is the number of channels, $numFilt$ is the number of filters, $filterH$ is the filter height, and $filterW$ is the filter width. Next, Equation 2 expresses the size of activations produced by all layers, combined.

$$|D| = \sum_{L=1}^{\#layers} ch_L * numFilt_L * activationW_L * activationH_L * batch * 4 \tag{2}$$

where $activationH$ is the activation map height, $activationW$ is the activation width, and $batch$ is the batch size. Note the *4 in Equations 1 and 2 – this is because a floating-point number is 4 bytes.

To minimize confusion, we now define some terminology. In our terminology, the each following sets of words are synonyms: (weights = parameters = filters = W); (nodes = workers = GPU servers). We also sometimes use the terms “activations” and “data” (D) interchangeably. *Fully-connected layers* are a special case of convolutional layers where $filterH = activationH$ and $filterW = activationW$. We define an “epoch” as one pass through the training data. Finally, the word “performance” can be ambiguous, so we write in terms of specific metrics such as “accuracy” and “training time.”

4. Parallelism strategies

There are two commonly-used methods for parallelizing neural network training across multiple servers: model parallelism (e.g. [45]) and data parallelism (e.g. [51]).

For batched SGD training of DNNs, we define *data parallelism* as the case where each worker (e.g. GPU) gets a subset of the batch, and then the workers communicate by exchanging weight gradient updates ∇W . We define *model parallelism* as the case where each worker gets a subset of the model parameters, and the workers communicate by exchanging data gradients ∇D and exchanging activations D . Note that $|W| = |\nabla W|$ and $|D| = |\nabla D|$; in other words, the weights and weight gradients are the same size; and the data and data gradients are the same size.

Now, to maximize DNN training scalability, our goal is to select a parallelism strategy that requires the lowest possible quantity of communication between servers. The choice of whether it is ideal to use data parallelism, model parallelism, or both depends strongly on the DNN’s architectural characteristics. Commonly-used DNN architectures for speech recognition (e.g. [38]) consist primarily of fully-connected layers, where the activations and parameters have the same spatial resolution (typically 1x1). For typical batch sizes, these fully-connected models often have a similar quantity of weights W and activations D . For example, we observe in Table 1 that this property holds true for the MSFT-Speech DNN architecture [38].

In computer vision, some of the most popular and accurate DNN models (e.g. GoogLeNet [41]) consist primarily of convolution layers, where the spatial resolution of the filters is smaller than the resolution of the activations. For these convolutional models, data parallelism is typically preferable because it requires less communication – that is, $|\nabla W|$ is much smaller than $|\nabla D|$ at typical batch sizes. Notice that the computer vision DNNs in Table 1 all have this property. In FireCaffe,

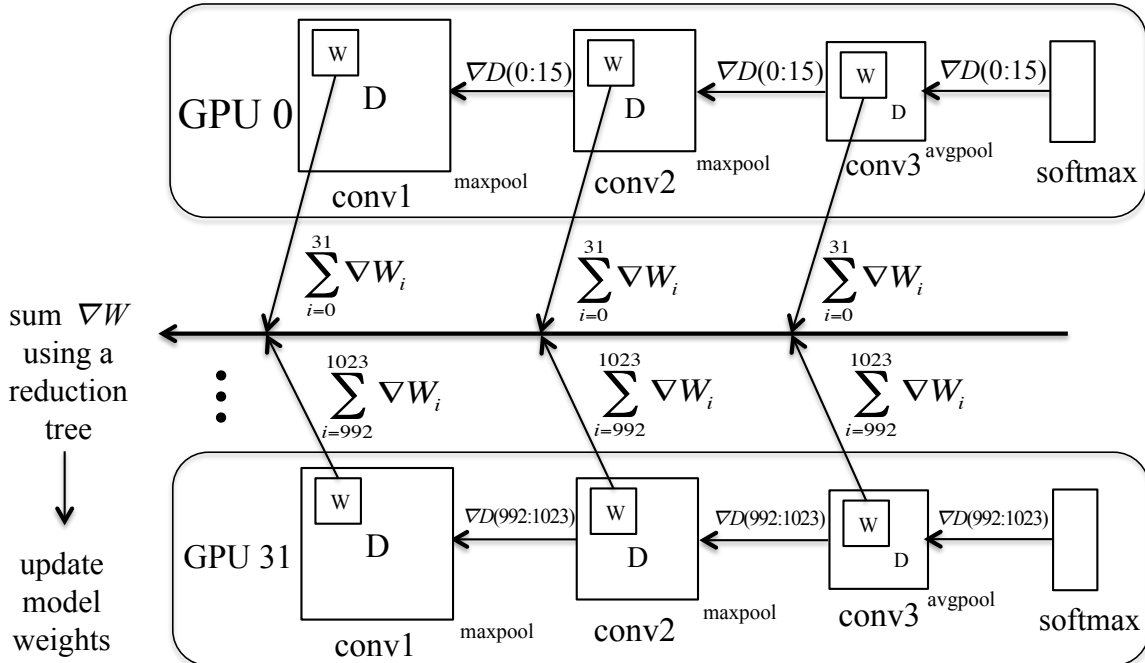


Figure 1. **Data parallel DNN training in FireCaffe:** Each worker (GPU) gets a subset of each batch.

we enable data parallelism across a cluster of GPUs, and we find that it produces ample speedups for training popular deep convolutional neural network architectures. We illustrate our data parallel approach in Figure 1. In this configuration, all the GPUs contain the full DNN model parameters. Each worker (GPU) gets a subset of each batch. The GPUs compute their share of the weight gradients. Once the gradients are calculated locally, they are added together using either a parameter server or a reduction tree communication (described in Section 6.2).

Table 1. Volumes of data and computation for four widely-used DNN architectures. The batch size impacts all numbers in this table except for $|W|$, and we use a batch size of 1024 in this table. Here, TFLOPS is the quantity of computation to perform.

DNN architecture	typical use-case	data_size $ D $	weight_size $ W $	data/weight ratio	Forward+Backward TFLOPS/batch
NiN [32]	computer vision	5800MB	30MB	195	6.7TF
AlexNet [28]	computer vision	1680MB	249MB	10.2	7.0TF
GoogLeNet [41]	computer vision	19100MB	54MB	358	9.7TF
VGG-19 [39]	computer vision	42700MB	575MB	71.7	120TF
MSFT-Speech [38]	speech recognition	74MB	151MB	0.49	0.00015TF

5. Choosing DNN architectures to accelerate

Of the popular deep convolutional neural network architectures for computer vision, some are more amenable to data parallel training than others. One might naïvely assume that DNN models with more parameters would produce higher classification accuracy. To evaluate this assumption, consider Figure 2, where we plot the total size of all parameters in bytes versus top-5 ImageNet accuracy for several popular DNN architectures. Observe that Network-in-Network (NiN) [32] and AlexNet [29] have similar accuracy, while NiN has 8x fewer parameters than AlexNet. Likewise, GoogLeNet [41] and VGG [39] have similar accuracy, yet GoogLeNet has 10x fewer parameters. In data parallel training, $|\nabla W|$ is the quantity of data sent by each GPU worker, so DNN architectures with fewer parameters require less communication and are more amenable to training at large scale.

You may wonder, what are the architectural choices that led to NiN and GoogLeNet having 8-10x fewer parameters than AlexNet and VGG? The answer is twofold. First, GoogLeNet and NiN are more judicious in their use of filters with spatial resolution: many of the filters in GoogLeNet and NiN have a resolution of 1x1 instead of 3x3 or larger. Second, while VGG and AlexNet each have more than 150MB of fully-connected layer parameters, GoogLeNet has smaller fully-connected layers, and NiN does not have fully-connected layers.

In summary, models with fewer parameters are more amenable to scalability in data parallel training, while still delivering

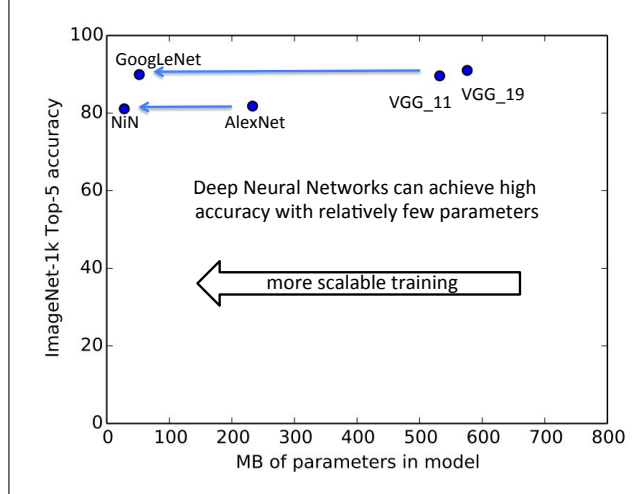


Figure 2. Deep neural network architectures with more parameters do not necessarily deliver higher accuracy.

high accuracy. Therefore, for the rest of the paper, we focus our efforts on accelerating the training of models with fewer parameters (e.g. NiN and GoogLeNet) while maintaining high accuracy.

6. Implementing efficient data parallel training

Our data-parallel distributed training strategy requires no communication among GPU workers in the forward pass. In the backward pass, a traditional single-GPU implementation (e.g. single-GPU Caffe [27]) sums the weight gradients over all images in the batch and then uses the weight gradient sum to update the model.¹ When we distribute the backward pass over a compute cluster, each GPU worker computes a sum of the weight gradients ($\sum \nabla W$) for its subset of the batch. Then, we sum the weight gradients across GPUs. This gradient aggregation scheme produces identical numerical results as you would find on a single GPU.

Now, our task is to find an efficient way to sum up ∇W among GPUs in a compute cluster. We consider two strategies for implementing this gradient aggregation: parameter servers, and reduction trees.

6.1. Parameter server

One strategy for communicating gradients is to appoint one node as a *parameter server*. The remaining *worker* nodes are each assigned a subset of the batch on which to perform forward and backward-propagation. After each backward pass, all the workers send their gradient updates to the parameter server. Then, the parameter server computes the sum of the gradients. Finally, the parameter server sends the summed gradients to the workers, and the workers apply these gradient updates to their local copies of the model. We illustrate the parameter server communication pattern in Figure 3(a).

The logical question here is, *what is the communication overhead of a parameter server, and how does that overhead scale as we increase the number of GPU workers?* Recall from Section 4 that each GPU worker provides $|W| = |\nabla W|$ bytes of weight gradients (Equation 1), which need to be summed with gradients from all other GPU workers. Now, the bottleneck is in sending and receiving all the gradients on one parameter server. If there are p GPU workers, the parameter server is responsible for sending and receiving $|\nabla W| * p$ bytes of data. If each node (GPU worker or parameter server) can send and receive data at a rate of BW bytes/s, then we can calculate the minimum communication time as follows:

$$param_server_communication_time = \frac{|\nabla W| * p}{BW} (sec) \quad (3)$$

In other words, the parameter server’s communication time scales linearly as we increase the number of GPU workers; doubling the number of workers leads to at least 2x more communication time per gradient update. We confirm this experimentally in Figure 4.

For the parameter server experiments in Figure 4, we have implemented a fully synchronous parameter server with the following characteristics. The parameter server is one arbitrarily-selected server in the cluster, while the other servers are workers; the parameter server and worker servers have identical hardware. After each batch, the workers send their weight

¹However, the data gradients (∇D) are not summed up.

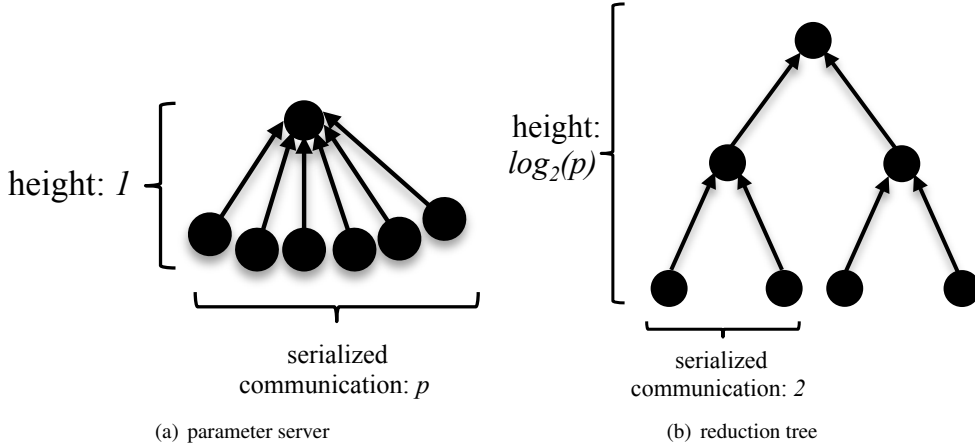


Figure 3. Illustrating how parameter servers and reduction trees communicate weight gradients. In this figure, we only show the summing-up of weight gradients. We distribute the weight gradient sums by going back down the tree.

gradients to the parameter server, the parameter server computes the sum, and then the parameter server sends the sum back to the workers.

There are a number of ways to augment the parameter server for greater scalability. For example, when having a single parameter server became a bottleneck, Microsoft Adam [8] and Google DistBelief [12] each defined a pool of nodes that collectively behave as a parameter server. The bigger the parameter server hierarchy gets, the more it looks like a reduction tree. This made us wonder: could we achieve greater scalability if we implement gradient aggregation as a reduction tree?

6.2. Reduction tree

There are various common patterns of communication in parallel programs; among such common patterns, a frequently occurring one is *allreduce*. This pattern occurs when each worker produces one or more data values that must be globally reduced (generally with a commutative binary element-wise operator) to produce a single result value, and then this single value must be broadcast to all workers before they can continue. In this work, each worker produces a single vector of length $|\nabla W|$ (the gradient updates for that worker), which must be reduced using element-wise vector addition (to sum the per-worker gradient updates for each parameter). Since this computation exactly fits the *allreduce* communication pattern it is convenient to use existing library support for such operations. While there are many possible implementations of *allreduce*, most share the key property that the time taken to perform the operation scales as the log of the number of workers (at least for large numbers of workers). Intuitively, this is because *allreduce* algorithms use binomial reduction tree and/or butterfly communication patterns internally [42]. Out of the possible allreduce implementation strategies, we find that the binomial reduction tree is particularly easy to reason about on a theoretical level. So, for the rest of this section, we focus on allreduce communication implemented with a reduction tree.

In Figures 3(a) and 3(b), we present the intuition on how parameter servers and reduction trees differ. We might think of a parameter server as a reduction tree with a height of 1 and a branching factor of p . However, many cluster computers and supercomputers have several dimensions of network fabric among nodes (e.g. an N - D Torus), which enable nodes to talk to each other via many different paths. With this in mind, we can sum gradients using a taller reduction tree, where nodes collaboratively sum the gradients. For example, consider a binary communication tree with a branching factor of 2 and a depth of $\log_2(p)$. In this case, the serialized communication is $2\log_2(p)$; the outer 2 represents the fact that each node receives data from 2 children, and the $\log_2(p)$ is the height of the tree. Therefore, unlike the parameter server model, the reduction tree’s communication time is:

$$\text{reduction_tree_communication_time} = \frac{|\nabla W| * 2\log_2(p)}{BW} (\text{sec}) \quad (4)$$

In practice, the base of $\log(p)$ depends on the branching factor in the reduction tree, but the basic idea here is straightforward: While the parameter server communication overhead scales *linearly* with p , reduction tree communication is much more efficient because it scales *logarithmically* as $O(\log(p))$. We confirm that reduction trees scale more efficiently than parameter servers in Figure 4.

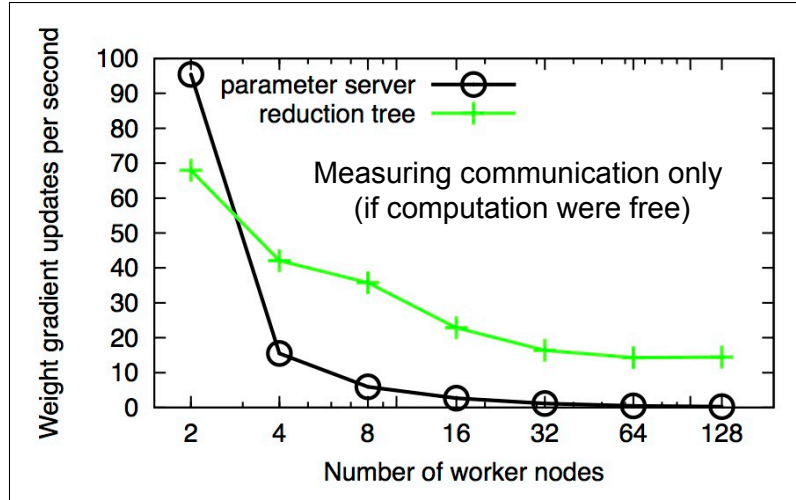


Figure 4. Comparing communication overhead with a parameter server vs. a reduction tree. This is for the Network-in-Network DNN architecture, so each GPU worker contributes 30MB of gradient updates.

7. Evaluation of FireCaffe-accelerated training on ImageNet

In this section, we evaluate how FireCaffe can accelerate DNN training on a cluster of GPUs. We train GoogLeNet [41] and Network-in-Network [32] on up to 128 GPU servers in the Titan supercomputer (described in Section 2), leveraging FireCaffe’s reduction tree data parallelism (Section 6.2). We begin by describing our evaluation methodology, and then we analyze the results.

7.1. Evaluation Methodology

We now describe a few practices that aid in comparing advancements in accelerating the training of deep neural networks.

1. Evaluate the speed and accuracy of DNN training on a publicly-available dataset.

In a recent study, Azizpour *et al.* applied DNNs to more than 10 different visual recognition challenge datasets, including human attribute prediction, fine-grained flower classification, and indoor scene recognition [5]. The accuracy obtained by Azizpour *et al.* ranged from 56% on scene recognition to 91% on human attribute prediction. As you can see, the accuracy of DNNs and other machine learning algorithms depends highly on the specifics of the application and dataset to which they are applied. Thus, when researchers report improvements in training speed or accuracy on proprietary datasets, there is no clear way to compare the improvements with the related literature. For example, Baidu [46] and Amazon [40] recently presented results on accelerating DNN training. Amazon and Baidu² each reported their training time numbers on a proprietary dataset, so it’s not clear how to compare these results with the related literature. In contrast, we conduct our evaluation on a publicly-available dataset, ImageNet-1k [13], which contains more than 1 million training images, and each image is labeled as containing 1 of 1000 object categories. ImageNet-1k is a widely-studied dataset, so we can easily compare our accuracy, training speed, and scalability results with other studies that use this data.

2. Report hyperparameter settings such as weight initialization, momentum, batch size, and learning rate.

Glorot *et al.* [18], Breuel [6], and Xu *et al.* [48] have each shown that seemingly-subtle hyperparameter settings such as weight initialization can have a big impact on the speed and accuracy produced in DNN training. When training Network-in-Network (NiN) [32], we initialize the weights with a gaussian distribution centered at 0, and we set the standard deviation (std) to 0.01 for 1x1 convolution layers, and we use std=0.05 for other layers. For NiN, we initialize the bias terms to a constant value of 0, we set the weight decay to 0.0005, and we set momentum to 0.9. These settings are consistent with the Caffe configuration files released by the NiN authors [32].

Frustratingly, in Google’s technical reports on GoogLeNet [41, 25], training details such as batch size, momentum, and learning rate are not disclosed. Fortunately, Wu *et al.* [47] and Guadarrama [19] each reproduced GoogLeNet and released all the details of their training protocols. As in [19], we train GoogLeNet with momentum=0.9 and weight decay=0.0002, we

²Baidu evaluated their training times using proprietary dataset [46]. Baidu also did some ImageNet experiments, but Baidu did not report the training time on ImageNet.

use xavier [18] weight initialization, and we initialize the bias terms to a constant value of 0.2. We will address learning rate and batch size settings in the following sections.

Given a DNN architecture, there are a number of strategies that can further increase accuracy, albeit at a substantial computational cost. One such strategy is to train multiple independent copies of a DNN architecture (e.g. GoogLeNet), each with a different random number generator seed for initializing the parameters. At test time, these DNNs can be used as an *ensemble* – that is, all DNNs are run on the test data, and for each test data item, the DNN’s classification activations are averaged. For example, using an ensemble of 7 GoogLeNet DNNs, Szegedy *et al.* achieved a 2 percentage-point accuracy improvement on ImageNet, compared to a single GoogLeNet baseline [41]. An other such technique is to augment the data by adding deformations or color variations during training and/or testing [46]. Our focus in this paper is to show speedup on training single models and compare with reported baselines. Hence we avoid using exotic data augmentation or ensembles of multiple DNNs. In our experiments, we resize images to 256x256; at training time we use a 224x224 crop with a randomized offset, and at test time we classify the 224x224 crop in the center of the image; these settings are also commonly used in the AlexNet [29] and Network-in-Network [32] DNN architectures.

3. Measure *speedups* with respect to a single-server baseline.

In order to meaningfully measure how much we have accelerated DNN training by adding more GPUs, we must have a representative baseline, e.g. with a single GPU. When reporting results, we begin by considering time required to train a DNN on single GPU, and we report our multi-GPU speedups with respect to this single-GPU baseline. A recent study by Microsoft [8] reported training a custom DNN architecture (e.g. not GoogLeNet or NiN) on a cluster of CPU servers. This may sound impressive, but Microsoft did not report the time that the model would take to train on a single server. It could be that Microsoft achieved a 10x speedup by going from 1 server to 10 servers, or the speedup could be 2x – this isn’t clear from the information provided in Microsoft’s paper. This illustrates the importance of measuring the speed of scalable DNN training systems with respect to a single-server baseline.

4. Measure *accuracy* with respect to a single-server baseline.

In our experience, if hyperparameters such as learning rate and batch size are selected too aggressively, a DNN model may converge quickly, but fall short of the state-of-art accuracy. Therefore, in our experiments, we train multi-GPU models until they reach to the single-GPU accuracy baseline; this validates that we can accelerate DNN training without degrading accuracy. However, in cluster-scale multi-GPU training experiments by Baidu [46] and Flickr [34], the training is stopped prematurely before the DNNs converge. This leaves us wondering whether the Baidu and Flickr multi-GPU training experiments would have reproduced the accuracy produced on a single GPU. To avoid this type of confusion, we evaluate both the speed and accuracy of FireCaffe DNN training with respect to a single-server/single-GPU baseline.

7.2. Results: Midsized deep models

Using the settings described by Krizhevsky [29], we find that AlexNet achieves 58.9% top-1 ImageNet-1k accuracy after 100 epochs of training. After just 47 epochs of training, we find that NiN also converges to 58.9% top-1 accuracy. Each training iteration of NiN is more time-consuming than AlexNet, and AlexNet and NiN both take approximately 6 days to converge to this level of accuracy.

At Google, Krizhevsky developed a scheme for accelerating AlexNet training using multiple GPUs within a single server [28]. Krizhevsky’s strategy uses data parallelism in convolutional layers and model parallelism in fully-connected layers. As we show in Table 2, Krizhevsky achieves near-linear acceleration up to 8 GPUs, but it has not been shown to scale beyond a single server. For reasons that we don’t entirely understand, Krizhevsky’s accuracy drops by 1.8 percentage points when doing multi-GPU training [28].

In FireCaffe, we scale NiN training to 32 GPUs, which is the scale at which we find communication time and computation are approximately equal³. We begin by using the learning rate and batch size settings that were reported in the Caffe configuration file released by the NiN authors [32]: For a batch size of 256, we use an initial learning rate of 0.01, and we reduce this by a factor of 10x twice during our training. Using this configuration, we reproduce the single-GPU NiN accuracy in 11 hours (13x speedup) when training on 32 GPUs.

For a fixed number of epochs, increasing the batch size reduces the number of times we need to communicate weight gradients, thus reducing the overall training time. With this in mind, we now train NiN with a batch size of 1024.⁴ As in [28] when we increase the batch size, we increase the learning rate by an equal proportion. For example, when we use a batch

³at a batch size of 1024

⁴While keeping a fixed number of epochs. In other words, with a batch size of 1024, we perform 4x fewer training iterations than with a batch size of 256.

size of 1024, we initialize the learning rate to 0.04. In this configuration, we train NiN in just 6 hours (23x speedup) on 32 GPUs. By increasing the batch size to 1024, we achieved a substantial speedup, but this came at the price of reducing the final accuracy by $\frac{3}{10}$ of a percentage point. We expect that this $\frac{3}{10}\%$ of accuracy could be regained at a batch size of 1024 – while retaining a substantial speed advantage – by training for a few more epochs. Finally, on 128 GPUs, we achieve a 39x speedup over single-GPU training.

So far, we have compared FireCaffe to the cuda-convnet2 framework from Google [28], which runs on a single-server/multi-GPU platform but not in a multi-server distributed platform. In addition to cuda-convnet2, Google has developed the TensorFlow framework [1], which also supports single-server/multi-GPU training but not distributed multi-server training. Thus far, Google has not released training speed results for multi-GPU TensorFlow. Twitter [16] has also experimented with scaling DNN training to 8 GPUs, but speed and accuracy results have not been released. Tencent [51], Theano [14], and Facebook [49] have published AlexNet single-server/multi-GPU training times that are slower than Google [28]. Other than FireCaffe, we have not seen literature on training AlexNet/NiN-scale models in a multi-server/multi-GPU setting. On 32 GPUs, FireCaffe is at least 3x faster to train AlexNet/NiN-scale models than all of the aforementioned results.

Table 2. Accelerating the training of midsized deep models on ImageNet-1k.

	Hardware	Net	Epochs	Batch size	Initial Learning Rate	Train time	Speedup	Top-1 Accuracy
Caffe [27]	1 NVIDIA K20	AlexNet [29]	100	256	0.01	6.0 days	1x	58.9%
Caffe	1 NVIDIA K20	NiN [32]	47	256	0.01	5.8 days	1x	58.9%
Google cuda-convnet2 [28]	8 NVIDIA K20s (1 node)	AlexNet	100	varies	0.02	16 hours	7.7x	57.1%
FireCaffe (ours)	32 NVIDIA K20s (Titan supercomputer)	NiN	47	256	0.01	11 hours	13x	58.9%
FireCaffe-batch1024 (ours)	32 NVIDIA K20s (Titan supercomputer)	NiN	47	1024	0.04	6 hours	23x	58.6%
FireCaffe-batch1024 (ours)	128 NVIDIA K20s (Titan supercomputer)	NiN	47	1024	0.04	3.6 hours	39x	58.6%

7.3. Results: Ultra-deep models

Ultra-deep models such as GoogLeNet can produce higher accuracy, but they present an even bigger challenge in terms of training time. Internally, Google has trained GoogLeNet on a cluster of CPU servers, but they have not reported the time required to complete this training [41, 25]. Fortunately, Guadarrama reproduced GoogLeNet in Caffe, and he released his GoogLeNet Caffe configuration files [19]. Guadarrama trained for 64 epochs using a batch size of 32 and an initial learning rate of 0.01, and we use these settings in our single-GPU GoogLeNet training experiments. Instead of occasionally reducing the learning rate by 10x, Guadarrama used a polynomial learning rate – that is, the learning rate is gradually reduced after every iteration of training. More specifically, at a given iteration of training, the learning rate is calculated as $initialLearningRate(1 - \frac{iter}{max.iter})^{power}$, and we set *power* to 0.5 in all of our GoogLeNet training runs. Running this in Caffe on a single-GPU, GoogLeNet takes 21 days to train on ImageNet-1k, producing 68.3% top-1 accuracy and 88.7% top-5 accuracy. This is slightly lower than the 89.9% top-5 single-model accuracy reported by Google [41], and it will be interesting to see whether the open-source Caffe community will eventually be able reproduce or surpass Google’s GoogLeNet accuracy. Here, we use the single-GPU Caffe GoogLeNet accuracy (88.7% top-5 accuracy) as a baseline, and we aim to reproduce this rapidly on a cluster of GPUs.

Now, we consider how to accelerate GoogLeNet training using FireCaffe. We initially tried to run GoogLeNet with a batch size of 32 on a GPU cluster, but there just wasn’t enough work per batch to keep a GPU cluster saturated. As we learned earlier in the paper, larger batch sizes lead to less frequent communication and therefore enable more scalability in a distributed setting. When modifying the batch size, Breuel [6] and Krizhevsky [28] found that the choice of learning rate is crucial in order to preserve high accuracy. We trained five separate versions of GoogLeNet, each with a different initial learning rate (LR): {0.02, 0.04, 0.08, 0.16, and 0.32}, and all with a batch size of 1024. With LR=0.16 and LR=0.32, GoogLeNet didn’t ever learn anything beyond random-chance accuracy on the test set. Using LR=0.02 produced 66.1% top-1 ImageNet-1k accuracy, and LR=0.04 produced 67.2%. Finally, we declare victory with LR=0.08, where we achieved 68.3% accuracy (again, with a batch size of 1024), which matches the accuracy of the baseline that used a batch size of 32. With a batch size of 1024 and a fixed number of epochs, we find that FireCaffe on 32 GPUs can train GoogLeNet 23x faster than a single GPU. When we move from a batch size of 32 with LR=0.01 to a batch size of 1024 with LR=0.08, we find that GoogLeNet takes a few more epochs to converge (72 epochs instead of 64 epochs), so the absolute training speedup is 20x; we show these results in Table 3. In other words, FireCaffe can train GoogLeNet in 23.4 hours on 32 GPUs, compared

to 21 days on a single GPU. Finally, on 128 GPUs, we achieve a 47x speedup over single-GPU GoogLeNet training, while matching the single-GPU accuracy.

Table 3. Accelerating the training of ultra-deep, computationally intensive models on ImageNet-1k.

	Hardware	Net	Epochs	Batch size	Initial Learning Rate	Train time	Speedup	Top-1 Accuracy	Top-5 Accuracy
Caffe	1 NVIDIA K20	GoogLeNet [41]	64	32	0.01	21 days	1x	68.3%	88.7%
FireCaffe (ours)	32 NVIDIA K20s (Titan supercomputer)	GoogLeNet	72	1024	0.08	23.4 hours	20x	68.3%	88.7%
FireCaffe (ours)	128 NVIDIA K20s (Titan supercomputer)	GoogLeNet	72	1024	0.08	10.5 hours	47x	68.3%	88.7%

8. Complementary approaches to accelerate DNN training

We have discussed related work throughout the paper, but we now provide a brief survey of additional techniques to accelerate deep neural network training. Several of the following techniques could be used in concert with FireCaffe to further accelerate DNN training.

8.1. Accelerating convolution on GPUs

In the DNN architectures discussed in this paper, more than 90% of the floating-point operations in forward and backward propagation reside in convolution layers, so accelerating convolution is key to getting the most out of each GPU. Recently, a number of techniques have been developed to accelerate convolution on GPUs. Unlike CPUs, NVIDIA GPUs have an inverted memory hierarchy, where the register file is larger than the L1 cache. Volkov and Demmel [44] pioneered a *communication-avoiding* strategy to accelerate matrix multiplication on GPUs by staging as much data as possible in registers while maximizing data reuse. Iandola *et al.* [23] extended the communication-avoiding techniques to accelerate 2D convolution; and cuDNN [7] and maxDNN [30] extended the techniques to accelerate 3D convolution. FireCaffe can be coupled with current and future GPU hardware and convolution libraries for further speedups.

8.2. Reducing communication among servers

Reducing the quantity of data communicated per batch is a useful way to increase the speed and scalability of DNN training. There is an inherent tradeoff here: as gradients are more aggressively quantized, training speed goes up, but the model’s accuracy may go down compared to a non-quantized baseline. While FireCaffe uses 32-bit floating-point values for weight gradients, Jeffrey Dean stated in a recent keynote speech that Google often uses 16-bit floating-point values for communication between servers in DNN training [11]. Along the same lines, Wawrzyniec *et al.* used 16-bit weights and 8-bit activations in distributed neural network training [45]. Going one step further, Seide *et al.* used 1-bit gradients for backpropagation, albeit with a drop in the accuracy of the trained model [37]. Finally, a related strategy to reduce communication between servers is to discard (and not communicate) gradients whose numerical values fall below a certain threshold. Amazon presented such a thresholding strategy in a recent paper on scaling up DNN training for speech recognition [40]. However, Amazon’s evaluation uses a proprietary dataset, so it is not clear how this type of thresholding impacts the accuracy compared to a well-understood baseline.

So far in this section, we have discussed strategies for compressing or quantizing data to communicate in distributed DNN training. There has also been a series of studies on applying dimensionality reduction to DNNs once they have been trained. Jaderberg *et al.* [26] and Zhang *et al.* [50] both use PCA to compress the weights of DNN models by up to 5x, albeit with a substantial reduction in the model’s classification accuracy. Han *et al.* [20] use a combination of pruning, quantization, and Huffman encoding to compress the weights of pretrained models by 35x with no reduction in accuracy. Thus far, these algorithms have only been able to accelerate DNNs at test time.

9. Conclusions

Long training times impose a severe limitation on progress in deep neural network research and productization. Accelerating DNN training has several benefits. First, faster DNN training enables models to be trained on ever-increasing dataset sizes in a tractable amount of time. Accelerating DNN training also enables product teams to bring DNN-based products to market more rapidly. Finally, there are a number of compelling use-cases for real-time DNN training, such as robot self-learning. These and other compelling applications led us to focus on the problem of accelerating DNN training, and our work has culminated in the FireCaffe distributed DNN training system.

Our approach to accelerating DNN training at scale has three key pillars. First, we select network hardware that achieves high bandwidth between GPU servers – Infiniband or Cray interconnects are ideal for this. Second, when selecting a communication algorithm, we find that reduction trees are more efficient and scalable than the traditional parameter server approach. Third, we optionally increase the batch size to reduce the total quantity of communication during DNN training, and we identify hyperparameters that allow us to reproduce the small-batch accuracy while training with large batch sizes. These three pillars helped us to achieve a near-linear speedup for a number of leading deep neural network architectures. In particular, we have achieved 39x speedup on NiN training, and a 47x speedup on GoogLeNet training on a 128 GPU cluster.

Acknowledgements

This research used resources of the Oak Ridge Leadership Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Thanks to Bryan Catanzaro, Evan Shelhamer, and Trevor Darrell for helpful discussions on scalable deep neural network training. Thanks to Sergio Guadarrama and Christian Szegedy for helpful discussions about GoogLeNet. Thanks to Thomas Breuel for sharing his intuition with us about the design space of deep neural network architectures, and in particular for discussions regarding the impact of learning rates. Thanks to Kostadin Ilov for IT support and for answering our numerous questions about hardware infrastructure. Thanks to Aditya Devarakonda for helpful discussions about supercomputer network interconnect hardware.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems. *Google Technical Report*, 2015. 9
- [2] V. Anantharaj, F. Foerster, W. Joubert, and J. Wells. Approaching exascale: application requirements for olcf leadership computing. https://www.olcf.ornl.gov/wp-content/uploads/2013/01/OLCF_Requirements_TM_2013_Final1.pdf, 2013. p43–45. 3
- [3] K. Asanovic and D. Patterson. FireBox: a hardware building block for 2020 warehouse-scale computers. In *USENIX FAST*, 2014. 13
- [4] K. Ashraf, B. Elizalde, F. Iandola, M. Moskewicz, J. Bernd, G. Friedland, and K. Keutzer. Audio-based multimedia event detection with dnns and sparse sampling. In *ICMR*, 2015. 2
- [5] H. Azizpour, A. S. Razavian, J. Sullivan, A. Maki, and S. Carlsson. From generic to specific deep representations for visual recognition. In *CVPR Deep Vision Workshop*, 2015. 7
- [6] T. M. Breuel. The effects of hyperparameters on SGD training of neural networks. *arXiv:1508.02788*, 2015. 7, 9
- [7] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer. cuDNN: efficient primitives for deep learning. *arXiv:1410.0759*, 2014. 1, 10
- [8] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman. Project Adam: building an efficient and scalable deep learning training system. In *OSDI*, 2014. 6, 8
- [9] H. Chu. MDB: A memory-mapped database and backend for openldap. In *LDAPCon*, 2011. 13
- [10] J. Dean. Keynote: Large scale deep learning. In *ACM Conference on Information and Knowledge Management (CIKM)*, 2014. 1
- [11] J. Dean. Keynote: Large scale deep learning. In *GPU Technology Conference*, 2015. 10
- [12] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, Q. V. Le, and A. Y. Ng. Large scale distributed deep networks. In *NIPS*, 2012. 6
- [13] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 7
- [14] W. Ding, R. Wang, F. Mao, and G. W. Taylor. Theano-based large-scale visual recognition with multiple gpus. *arXiv:1412.2302*, 2014. 9
- [15] H. Fang, S. Gupta, F. Iandola, R. Srivastava, L. Deng, P. Dollar, J. Gao, X. He, M. Mitchell, J. C. Platt, C. L. Zitnick, and G. Zweig. From captions to visual concepts and back. In *CVPR*, 2015. 2
- [16] C. Farabet. Real-time, content-driven representations at twitter. In *GPU Technology Conference*, 2015. 9
- [17] R. B. Girshick, F. N. Iandola, T. Darrell, and J. Malik. Deformable part models are convolutional neural networks. In *CVPR*, 2015. 2
- [18] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. *JMLR*, 2010. 7, 8
- [19] S. Guadarrama. BVLC googlenet. https://github.com/BVLC/caffe/tree/master/models/bvlc_googlenet, 2015. 7, 9

- [20] S. Han, H. Mao, and W. J. Dally. A deep neural network compression pipeline: Pruning, quantization, huffman encoding. *arXiv:1510.00149*, 2015. 10
- [21] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng. Deep speech: Scaling up end-to-end speech recognition. *arXiv:1412.5567*, 2014. 2
- [22] F. N. Iandola, M. W. Moskewicz, S. Karayev, R. B. Girshick, T. Darrell, and K. Keutzer. Densenet: Implementing efficient convnet descriptor pyramids. *arXiv:1404.1869*, 2014. 2
- [23] F. N. Iandola, D. Sheffield, M. Anderson, P. M. Phothilimthana, and K. Keutzer. Communication-minimizing 2d convolution in gpu registers. In *ICIP*, 2013. 10
- [24] F. N. Iandola, A. Shen, P. Gao, and K. Keutzer. DeepLogo: Hitting logo recognition with the deep neural network hammer. *arXiv:1510.02131*, 2015. 2
- [25] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *JMLR*, 2015. 7, 9
- [26] M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. *arXiv:1405.3866*, 2014. 10
- [27] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv:1408.5093*, 2014. 5, 9
- [28] A. Krizhevsky. One weird trick for parallelizing convolutional neural networks. *arXiv:1404.5997*, 2014. 4, 8, 9
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*, 2012. 1, 4, 8, 9, 13
- [30] A. Lavin. maxDNN: an efficient convolution kernel for deep learning with maxwell gpus. *arXiv:1501.06633*, 2015. 1, 10
- [31] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *arXiv:1504.00702*, 2015. 2
- [32] M. Lin, Q. Chen, and S. Yan. Network in network. *arXiv:1312.4400*, 2013. 1, 4, 7, 8, 9, 13
- [33] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013. 2
- [34] C. Noel, J. Shi, and A. Feng. Large scale distributed deep learning on hadoop clusters. 2015. <http://yahoohadoop.tumblr.com/post/129872361846>. 8
- [35] Oak Ridge Leadership Computing Facility (OLCF). Introducing titan: Advancing the era of accelerated computing. <http://www.olcf.ornl.gov/titan>, 2015. 3
- [36] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *NIPS*, 2011. 13
- [37] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *INTERSPEECH*, 2014. 10
- [38] F. Seide, G. Li, and D. Yu. Conversational speech transcription using context-dependent deep neural networks. In *INTERSPEECH*, 2011. 3, 4
- [39] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. 1, 4
- [40] N. Strom. Scalable distributed dnn training using commodity gpu cloud computing. In *INTERSPEECH*, 2015. 7, 10
- [41] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv:1409.4842*, 2014. 1, 2, 3, 4, 7, 8, 9, 10
- [42] R. Thakur, R. Rabenseifner, and W. Gropp. Optimization of collective communication operations in mpich. *International Journal of High Performance Computing Applications*, 2005. 6
- [43] N. Vasilache, J. Johnson, M. Mathieu, S. Chintala, S. Piantino, and Y. LeCun. Fast convolutional nets with fbfft: A gpu performance evaluation. *arXiv:1412.7580*, 2014. 1
- [44] V. Volkov and J. W. Demmel. Benchmarking GPUs to tune dense linear algebra. In *Supercomputing*, 2008. 10
- [45] J. Wawrzynek, K. Asanovic, B. Kingsbury, D. Johnson, J. Beck, and N. Morgan. Spert-ii: A vector microprocessor system. *Computer*, 1996. 3, 10
- [46] R. Wu, S. Yan, Y. Shan, Q. Dang, and G. Sun. Deep image: Scaling up image recognition. *arXiv:1501.02876*, 2015. 7, 8
- [47] Z. Wu, Y. Zhang, F. Yu, and J. Xiao. A gpu implementation of googlenet. <http://vision.princeton.edu/pvt/GoogLeNet>, 2014. 7
- [48] B. Xu, N. Wang, T. Chen, and M. Li. Empirical evaluation of rectified activations in convolutional network. *arXiv:1505.00853*, 2015. 7
- [49] O. Yadan, K. Adams, Y. Taigman, and M. Ranzato. Multi-gpu training of convnets. *arXiv:1312.5853*, 2013. 9
- [50] X. Zhang, J. Zou, X. Ming, K. He, and J. Sun. Efficient and accurate approximations of nonlinear convolutional networks. *arXiv:1411.4229*, 2014. 10
- [51] Y. Zou, X. Jin, Y. Li, Z. Guo, E. Wang, and B. Xiao. Mariana: Tencent deep learning platform and its applications. In *VLDB*, 2014. 3, 9

Appendix: Frequently Asked Questions about FireCaffe

1. *I use Caffe. Will my DNN models and network definitions work with FireCaffe?*

Yes.

2. *What are some good DNN architectures to use with FireCaffe?*

We offer intuition and analysis of this in Sections 4 and 6, but here we offer a short summary of what you need to know. As a general rule, DNN architectures with fewer parameters train faster in FireCaffe. We have found that NiN [32] and AlexNet [29] provide similar accuracy on ImageNet and other popular data benchmarks. However, NiN has 8x fewer parameters and therefore incurs 8x less communication cost than AlexNet. If you were planning to use AlexNet, we recommend trying NiN for faster training.

3. *I want to design my own DNN architectures that will scale well nicely in FireCaffe while producing good accuracy on my problem. What design tradeoffs should I consider?*

If you're designing your own DNN architectures, do your best to economize on parameters. Also, reducing the convolution and pooling strides may improve accuracy; this doesn't require additional parameters and doesn't hurt training scalability.

4. *Is it better to use FireCaffe with one server that contains many GPUs, or is it better to distribute the GPUs distributed over many servers?*

FireCaffe is compatible with both of these scenarios. FireCaffe can even run across many servers that each contain many GPUs. To avoid stragglers during backpropagation, we prefer running FireCaffe on a collection of identical GPUs (e.g. all Titan X or all K80; but preferably not a mixture of K80 and Titan X). For optimal speed and utilization, FireCaffe prefers a low-latency interconnect such as Infiniband between servers.

5. *How does FireCaffe handle the storage and loading of training data?*

As in Caffe, FireCaffe can ingest LMDB databases [9] of training data. This data format is agnostic to the type of data (e.g. images, audio, text), so long as each training data item consists of a vector (e.g. pixels, audio waveform, text trigram) and a label (e.g. dog or cat). We store the LMDB database of training data on a distributed filesystem that is accessible to all workers. During DNN training, each worker is responsible for loading its own training data from the distributed filesystem.

6. *Does FireCaffe use nondeterministic techniques such as Hogwild [36]?*

No. Given a specific seed for the random number generator, FireCaffe produces repeatable numerical results, just like ordinary Caffe. In fact, for a model with Dropout disabled, FireCaffe will give you the exact same numerics as ordinary Caffe. With Dropout enabled, FireCaffe handles randomization in a slightly different way than Caffe, but the numerics are still deterministic for FireCaffe with a constant number of GPUs.

7. *How did FireCaffe get its name?*

Asanovic and Patterson presented a roadmap for warehouse-scale computing in the year 2020, which they call *FireBox* [3]. The grand vision is that a typical large-scale commercial datacenter will have extremely low latency connections within and across racks, likely using photonic networking hardware. Low-latency network hardware is crucial not only for today's mainstream applications like distributed databases and search engines, but also for emerging applications such as deep neural network training at large scale. We have designed FireCaffe with FireBox-style warehouse-scale computing in mind.