

Staleness-Aware Async-SGD for Distributed Deep Learning

Wei Zhang, Suyog Gupta

IBM T. J. Watson Research Center
Yorktown Heights, NY 10598, USA
{weiz,suyog}@us.ibm.com

Xiangru Lian, Ji Liu

Department of Computer Science
University of Rochester, NY 14627, USA
{liangru,ji.liu.uwisc}@gmail.com

Abstract

Deep neural networks have been shown to achieve state-of-the-art performance in several machine learning tasks. Stochastic Gradient Descent (SGD) is the preferred optimization algorithm for training these networks and asynchronous SGD (ASGD) has been widely adopted for accelerating the training of large-scale deep networks in a distributed computing environment. However, in practice it is quite challenging to tune the training hyperparameters (such as learning rate) when using ASGD so as to achieve convergence and linear speedup, since the stability of the optimization algorithm is strongly influenced by the asynchronous nature of parameter updates. In this paper, we propose a variant of the ASGD algorithm in which the learning rate is modulated according to the gradient staleness and provide theoretical guarantees for convergence of this algorithm. Experimental verification is performed on commonly-used image classification benchmarks: CIFAR10 and Imagenet to demonstrate the superior effectiveness of the proposed approach, compared to SSGD (Synchronous SGD) and the conventional ASGD algorithm.

1 Introduction

Large-scale deep neural networks training is often constrained by the available computational resources, motivating the development of computing infrastructure designed specifically for accelerating this workload. This includes distributing the training across several commodity CPUs ([Dean *et al.*, 2012],[Chilimbi *et al.*, 2014]), or using heterogeneous computing platforms containing multiple GPUs per computing node ([Seide *et al.*, 2014],[Wu *et al.*, 2015],[Strom, 2015]), or using a CPU-based HPC cluster ([Gupta *et al.*, 2015]).

Synchronous SGD (SSGD) is the most straightforward distributed implementation of SGD in which the master simply splits the workload amongst the workers at every iteration. Through the use of barrier synchronization, the master ensures that the workers perform gradient computation using the identical set of model parameters. The workers are forced to wait for the slowest one at the end of every iteration. This synchronization cost deteriorates the scalability and runtime

performance of the SSGD algorithm. Asynchronous SGD (ASGD) overcomes this drawback by removing any explicit synchronization amongst the workers. However, permitting this asynchronous behavior inevitably adds “staleness” to the system wherein some of the workers compute gradients using model parameters that may be several gradient steps behind the most updated set of model parameters. Thus when fixing the number of iterations, ASGD-trained model tends to be much worse than SSGD-trained model. Further, there is no known principled approach for tuning learning rate in ASGD to effectively counter the effect of stale gradient updates.

Prior theoretical work by [Tsitsiklis *et al.*, 1986] and [Agarwal and Duchi, 2011] [Liu *et al.*, 2013] and recent work by [Liu and Wright, 2015], [Lian *et al.*, 2015], [Zhang *et al.*, 2015] provide theoretical guarantees for convergence of stochastic optimization algorithms in the presence of stale gradient updates for *convex* optimization and *nonconvex* optimization, respectively. We find that adopting the approach of scale-out deep learning using ASGD gives rise to complex interdependencies between the training algorithm’s hyperparameters (such as learning rate, mini-batch size) and the distributed implementation’s design choices (such as synchronization protocol, number of learners), ultimately impacting the neural network’s accuracy and the runtime performance. In practice, achieving good model accuracy through distributed training requires a careful selection of the training hyperparameters and much of the prior work cited above lacks enough useful insight to help guide this selection process.

The work presented in this paper intends to fill this void by undertaking a study of the interplay between the different design parameters encountered during distributed training of deep neural networks. In particular, we focus our attention on understanding the effect of stale gradient updates during distributed training and developing principled approaches for mitigating these effects. To this end, we introduce a variant of the ASGD algorithm in which we keep track of the staleness associated with each gradient computation and adjust the learning rate on a per-gradient basis by simply dividing the learning rate by the staleness value. The implementation of this algorithm on a CPU-based HPC cluster with fast interconnect is shown to achieve a tight bound on the gradient staleness. We experimentally demonstrate the effectiveness of the proposed staleness-dependent learning rate scheme using

commonly-used image classification benchmarks: CIFAR10 and Imagenet and show that this simple, yet effective technique is necessary for achieving good model accuracy during distributed training. Further, we build on the theoretical framework of [Lian *et al.*, 2015] and prove that the convergence rate of the staleness-aware ASGD algorithm is consistent with SGD: $\mathcal{O}\left(1/\sqrt{T}\right)$ where T is the number of gradient update steps.

Previously, [Ho *et al.*, 2013] presented a parameter server based distributed learning system where the staleness in parameter updates is bounded by forcing faster workers to wait for their slower counterparts. Perhaps the most closely related prior work is that of [Chan and Lane, 2014] which presented a multi-GPU system for distributed training of speech CNNs and acknowledge the need to modulate the learning rate in the presence of stale gradients. The authors proposed an exponential penalty for stale gradients and show results for up to 5 learners, without providing any theoretical guarantee of the convergence rate. However, in larger-scale distributed systems, the gradient staleness can assume values up to a few hundreds ([Dean *et al.*, 2012]) and the exponential penalty may reduce the learning rate to an arbitrarily small value, potentially slowing down the convergence. In contrast, in this paper, we formally prove our proposed ASGD algorithm to converge as fast as SSGD. Further, our implementation achieves near-linear speedup while maintaining the optimal model accuracy. We demonstrate this on widely used image classification benchmarks.

2 System architecture

In this section we present an overview of our distributed deep learning system and describe the synchronization protocol design. In particular, we introduce the *n*-softsync protocol which enables a fine-grained control over the upper bound on the gradient staleness in the system. For a complete comparison, we also implemented the *Hardsync protocol* (aka SSGD) for model accuracy baseline since it generates the most accurate model (when fixing the number of training epochs), albeit at the cost of poor runtime performance.

2.1 Architecture Overview

We implement a parameter server based distributed learning system, which is a generalization¹ of Downpour SGD in [Dean *et al.*, 2012], to evaluate the effectiveness of our proposed staleness-dependent learning rate modulation technique. Throughout the paper, we use the following definitions:

- λ : number of learners (workers).
- μ : mini-batch size used by each learner to produce stochastic gradients.
- α : learning rate.
- Epoch: a pass through the entire training dataset.
- Timestamp: we use a scalar clock to represent weights timestamp i , starting from $i = 0$. Each weight update

increments the timestamp by 1. The timestamp of a gradient is the same as the timestamp of the weight used to compute the gradient.

- $\tau_{i,l}$: staleness of the gradient from learner l . A learner l pushes gradient with timestamp j to the parameter server of timestamp i , where $i \geq j$. We calculate the staleness $\tau_{i,l}$ of this gradient as $i - j$. $\tau_{i,l} \geq 0$ for any i and l .

Each learner performs the following sequence of steps. `getMinibatch`: Randomly select a mini-batch of examples from the training data; `pullWeights`: A learner pulls the current set of weights from the parameter server; `calcGradient`: Compute stochastic gradients for the current mini-batch. We divide the gradients by the mini-batch size; `pushGradient`: Send the computed gradients to the parameter server;

The parameter server group maintains a global view of the neural network weights and performs the following functions. `sumGradients`: Receive and accumulate the gradients from the learners; `applyUpdate`: Multiply the average of accumulated gradient by the learning rate (step length) and update the weights.

2.2 Synchronization protocols

We implemented two synchronization protocols: *hardsync* protocol (aka, SSGD) and *n*-softsync protocol (aka, ASGD). *Hardsync* protocol (SSGD) yields a model with the best accuracy number, however synchronization overheads deteriorate the overall runtime performance. *n*-softsync protocol is our proposed ASGD algorithm that automatically tunes learning rate based on gradient staleness and achieves model accuracy comparable with SSGD while providing a near-linear speedup in runtime.

Hardsync protocol: To advance the weights' timestamp θ from i to $i + 1$, each learner l compute a gradient $\Delta\theta_l$ using a mini-batch size of μ and sends it to the parameter server. The parameter server averages the gradients over λ learners and updates the weights according to equation 1, then broadcasts the new weights to all learners. The learners are forced to wait for the updated weights until the parameter server has received the gradient contribution from *all* the learners and finished updating the weights. This protocol guarantees that each learner computes gradients on the exactly the same set of weights and ensures that the gradient staleness is 0. The *hardsync* protocol serves as the baseline, since from the perspective of SGD optimization it is equivalent to SGD using batch size $\mu\lambda$.

$$g_i = \frac{1}{\lambda} \sum_{l=1}^{\lambda} \Delta\theta_l \quad (1)$$

$$\theta_{i+1} = \theta_i - \alpha g_i.$$

n-softsync protocol: Each learner l pulls the weights from the parameter server, calculates the gradients and pushes the gradients to the parameter server. The parameter server updates the weights after collecting at least $c = \lfloor(\lambda/n)\rfloor$ gradients from *any* of the λ learners. Unlike *hardsync*, there are no explicit synchronization barriers imposed by the parameter server and the learners work asynchronously and independently. The splitting parameter n can vary from 1 to λ . The

¹We restricted the scope of this paper to deal with data-level parallelism only

n -softsync weight update rule is given by:

$$c = \lfloor (\lambda/n) \rfloor$$

$$g_i = \frac{1}{c} \sum_{l=1}^c \alpha(\tau_{i,l}) \Delta\theta_l, \quad l \in \{1, 2, \dots, \lambda\} \quad (2)$$

$$\theta_{i+1} = \theta_i - g_i,$$

where $\alpha(\tau_{i,l})$ is the gradient staleness-dependent learning rate. Note that our proposed n -softsync protocol is a generalization of Downpour-SGD of [Dean et al., 2012] (a commonly-used ASGD implementation), in that when n is set to be λ , our implementation is equivalent to Downpour-SGD. By setting different n , ASGD can have different degrees of staleness, as demonstrated in Section 2.4.

2.3 Implementation Details

We use MPI as the communication mechanism between learners and parameter servers. Parameter servers are sharded. Each learner and parameter server are 4-way threaded. During the training process, a learner pulls weights from the parameter server, starts training when the weights arrive, and then calculates gradients. Finally it pushes the gradients back to the parameter server before it can pull the weights again. We do not “accrue” gradients at the learner so that each gradient pushed to the parameter server is always calculated out of one mini-batch size as accruing gradients generally lead to a worse model [Lian et al., 2015; Gupta et al., 2015]. In addition, the parameter server communicates with learners via MPI blocking-send calls (i.e., pullWeights and pushGradient), that is the computation on the learner is stalled until the corresponding blocking send call is finished. The design choice is due to the fact that it is difficult to guarantee making progress for MPI non-blocking calls and multi-thread level support to MPI communication is known not to scale [MPI-Forum, 2012]. Further, by using MPI blocking calls, the gradients’ staleness can be effectively bounded, as we demonstrate in Section 2.4. Note that the computation in parameter servers and learners are however concurrent (except for the learner that is communicating with the server, if any). No synchronization is required between learners and no synchronization is required between parameter server shards.

Since memory is abundant on each computing node, our implementation does not split the neural network model across multiple nodes (model parallelism). Instead, depending on the problem size, we pack either 4 or 6 learners on each computing node. Learners operate on homogeneous processors and run at similar speed. In addition, fast interconnect expedites pushing gradients and pulling weights. Both of these hardware aspects help bound (but does not guarantee) gradients’ staleness.

2.4 Staleness analysis

In the hardsync protocol, the update of weights from θ_i to θ_{i+1} is computed by aggregating the gradients calculated using weights θ_i . As a result, each of the gradients g_i in the i^{th} step carries with it a staleness $\tau_{i,l}$ equal to 0.

Figure 1 shows the measured distribution of gradient staleness for different n -softsync (ASGD) protocols when using

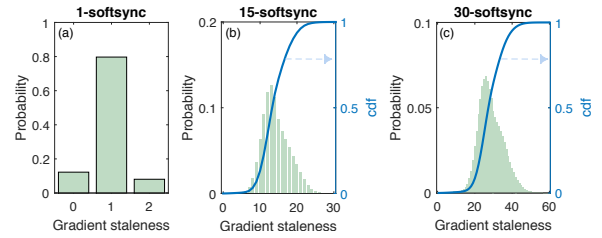


Figure 1: Distribution of gradient staleness for 1, 15, and 30-softsync protocols. $\lambda = 30$

$\lambda = 30$ learners. For the 1-softsync, the parameter server updates the current set of weights when it has received a total of 30 gradients from (any of) the learners. In this case, the staleness $\tau_{i,l}$ for the gradients computed by the learner l takes values 0, 1, or 2. Similarly, the 15-softsync protocol forces the parameter server to accumulate $\lambda/15 = 2$ gradient contributions from the learners before updating the weights. On the other hand, the parameter server updates the weights after receiving a gradient from any of the learners when the 30-softsync protocol is enforced. The average staleness $\langle \tau_i \rangle$ for the 15-softsync and 30-softsync protocols remains close to 15 and 30, respectively. Empirically, we have found that a large fraction of the gradients have staleness close to n , and only with a very low probability (< 0.0001) does τ exceed $2n$. These measurements show that, in general, $\tau_{i,l} \in \{0, 1, \dots, 2n\}$ and $\langle \tau_i \rangle \approx n$ for the n -softsync protocol². Clearly, the n -softsync protocol provides an effective mechanism for controlling the gradient staleness.

In our implementation, the parameter server uses the staleness information to modulate the learning rate on a per-gradient basis. For an incoming gradient with staleness $\tau_{i,l}$, the learning rate is set as:

$$\alpha_{i,l} = \frac{\alpha_0}{\tau_{i,l}} \quad \text{if } \tau_{i,l} > 0 \quad (3)$$

where α_0 is typically set as the ‘best-known’ learning rate when using SSGD. In section 4 we show experimental results comparing this staleness-dependent learning rate scheme with the case where the learning rate is kept constant at α_0 .

3 Theoretical Analysis

This section provides theoretical analysis of the ASGD algorithm proposed in section 2. More specifically, we will show the convergence rate and how the gradient staleness affects the convergence. In essence, we are solving the following generic optimization problem:

$$\min_{\theta} F(\theta) := \frac{1}{N} \sum_{i=1}^N f_i(\theta),$$

where θ is the parameter vector we are pursuing, N is the number of samples, and $f_i(\theta)$ is the loss function for the i^{th}

²We have found this empirical observation to hold true regardless of the mini-batch size per learner and the size of the model. The plots in Figure 1 were generated using the CIFAR10 dataset/model (see section 4) and mini-batch size per learner $\mu = 4$.

sample. If every learner computes μ gradients at once and the parameter server updates the parameter when it receives c mini-batches from learners, from the perspective of the parameter server, the update procedure of parameter θ mentioned can be written as

$$g_i = \underbrace{\frac{1}{c} \sum_{l=1}^c \frac{\alpha_0}{\tau_{i,l}} \left(\frac{1}{\mu} \sum_{s=1}^{\mu} \nabla f_{\xi_{i,s,l}}(\theta_{i-\tau_{i,l}}) \right)}_{\text{calculated in a learner}}, \quad (4)$$

$$\theta_{i+1} = \underbrace{\theta_i - g_i}_{\text{calculated in the parameter server}},$$

where $\tau_{i,l}$ is the staleness of the parameter used to calculate the gradients in the l^{th} learner in step i and α_0 is a constant. $\xi_{i,s,l}$ denotes the subscript of f used to calculate the s^{th} gradient in the l^{th} learner in the i^{th} step.

To simplify the analysis, we decompose every single step in (4) to c steps by updating only one batch of gradients in each step. Then the sequence $\{\theta_i\}_{i \in \mathbb{N}}$ becomes $\{\tilde{\theta}_t\}_{t \in \mathbb{N}}$, and $\tilde{\theta}_{ci} = \theta_i$. Formally, it will be

$$\tilde{g}_t = \frac{1}{c\mu} \frac{\alpha_0}{\tilde{\tau}_t - r_t} \left(\sum_{s=1}^{\mu} \nabla f_{\xi_{t,s}}(\tilde{\theta}_{t-\tau_t}) \right), \quad (5)$$

$$\tilde{\theta}_{t+1} = \underbrace{\tilde{\theta}_t - \tilde{g}_t}_{\text{calculated in the parameter server}},$$

where r_t is the remainder of t/c . $\xi_{t,s}$ denotes the subscript of f used to calculate the s^{th} gradient in the t^{th} step in our new formulation. One can verify that here t increases c times faster than the i in (4).

Note that the $\tilde{\tau}_t - r_t$ in (5) is always positive³. We use $\{p_t\}_{t \in \mathbb{N}}$ to denote the difference $p_t = \tilde{\tau}_t - r_t$. It immediately follows that $p_t = \tau_{\lfloor t/c \rfloor, r_t}$. From the Theorem 1 in [Lian et al., 2015] with some modification, we have the following theorem, which indicates the convergence rate and the linear speedup property of our algorithm.

Theorem 1. *Let C_1, C_2, C_3, C_4 be certain positive constants depending on the objective function $F(\theta)$. Under certain commonly used assumptions (please find in Theorem 1 in [Lian et al., 2015]), we can achieve an convergence rate of*

$$\frac{1}{\sum_{t=1}^T 1/p_t} \sum_{t=1}^T \frac{1}{p_t} \mathbb{E}(\|\nabla F(\tilde{\theta}_t)\|^2) \leq 2 \frac{\sqrt{\frac{2C_1C_2}{\mu} \sum_{t=1}^T \left(\frac{1}{p_t}\right)}}{\sum_{t=1}^T \frac{1}{p_t}} \quad (6)$$

³In (5) when a mini-batch is updated into the parameter, the counter (t) will increase by 1, while in (4) the counter (i) increases by 1 every c mini-batches updated into the parameter. For example if we have 3 mini-batches with staleness 1 pushed, in (4) all $\tau_{i,l}$ will be 1. However, in (5), if one mini-batch is updated in iteration t , the staleness $\tilde{\tau}_{t+1}$ of the other two mini-batches becomes 2, so we need to subtract the redundant part of the staleness caused by the difference in counting. Because the staleness after subtraction is exactly the original staleness in (4), it is always positive.

where T is the total iteration number, if

$$\alpha_0 = \sqrt{C_1 c^2 \mu / \sum_{t=1}^T \left(\frac{2}{p_t^2} C_2 \right)}, \quad (7)$$

under the prerequisite that

$$\alpha_0 \leq \frac{cC_2}{C_3 p_t \sum_{j=t-2n}^{t-1} \frac{1}{p_j^2}}, \quad \forall t, \quad (8)$$

and

$$C_3 \frac{\alpha_0}{c p_t} + C_4 n \frac{\alpha_0^2}{c^2 p_t} \sum_{\kappa=1}^{2n} \frac{1}{p_{t+\kappa}} \leq 1, \quad \forall t. \quad (9)$$

First note that (8) and (9) can always be satisfied by selecting small enough α_0 (or equivalently, large enough T). Thus, if the learning rate is appropriately chosen in our algorithm, the weighted average of the gradients (which is the LHS of (6)) is guaranteed to converge. Also note that that to achieve this convergence rate, the batch size μ cannot be too large, since from (7) we can see that a larger μ leads to a larger α_0 , which may not satisfy the prerequisites (8) and (9).

A clearer dependence between the staleness and the convergence rate can be found by taking a closer look at the RHS (6):

Remark 1. *Note that the RHS of (6) is of the form $h(z_1, \dots, z_T) = O\left(\frac{\sqrt{z_1^2 + z_2^2 + \dots + z_T^2}}{z_1 + z_2 + \dots + z_T}\right)$ by letting $z_t = \frac{1}{p_t}$.*

If the summation $z_1 + \dots + z_T$ is fixed, one can verify that h is minimized when $z_1 = z_2 = \dots = z_T$. Therefore our theorem suggests that a centralized distribution of staleness p (or τ in (4)) leads to a better convergence rate.

Further, we have the following result by considering the ideal scenario of the staleness.

Remark 2. *Note that if we take p_t as a constant p , we have*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}(\|\nabla F(\tilde{\theta}_t)\|^2) \leq 2 \frac{\sqrt{2C_1C_2}}{\sqrt{T}\mu}.$$

Thus this convergence rate is roughly in the order of $O(1/\sqrt{\mu T})$, where T is the total iteration number and μ is the mini-batch size. Equivalently, a goal of $\frac{1}{T} \sum_{t=1}^T \mathbb{E}(\|\nabla F(\tilde{\theta}_t)\|^2) \leq \epsilon$ can be achieved by having $\mu T = O(1/\epsilon^2)$. This is consistent with the convergence rate of SGD in [Lian et al., 2015], which suggests a linear speedup can be achieved in our algorithm.

Proof. From (9) we have

$$\begin{aligned} & C_3 \frac{\alpha_0}{c p_t} + C_4 n \frac{\alpha_0^2}{c^2 p_t} \sum_{\kappa=1}^{2n} \frac{1}{p_{t+\kappa}} \\ &= C_3 \mu \frac{\alpha_0}{\mu c p_t} + C_4 \mu^2 n \frac{\alpha_0}{\mu c p_t} \sum_{\kappa=1}^{2n} \frac{\alpha_0}{\mu c p_{t+\kappa}} \\ &\leq 1, \quad \forall t. \end{aligned}$$

With (8) we have

$$\frac{\alpha_0^2}{\mu c^2 p_t^2} C_2 \geq \frac{\alpha_0^3}{c^3 p_t \mu} C_3 \sum_{j=t-2n}^{t-1} \frac{1}{p_j^2}, \forall t. \quad (10)$$

Note that the upperbound of the staleness is $2n$ in our setting. Then it follows from Theorem 1 in [Lian *et al.*, 2015] that

$$\begin{aligned} & \frac{1}{\sum_{t=1}^T 1/p_t} \sum_{t=1}^T \frac{1}{p_t} \mathbb{E}(\|\nabla F(\tilde{\theta}_t)\|^2) \\ & \leq \frac{C_1 + \sum_{t=1}^T \left(\frac{\alpha_0^2}{\mu c^2 p_t^2} C_2 + \frac{\alpha_0^3}{c^3 p_t \mu} C_3 \sum_{j=t-2n}^{t-1} \frac{1}{p_j^2} \right)}{\sum_{t=1}^T \frac{\alpha_0}{c p_t}} \\ & \stackrel{(10)}{\leq} \frac{C_1 + \alpha_0^2 \sum_{t=1}^T \left(\frac{2}{\mu c^2 p_t^2} C_2 \right)}{\sum_{t=1}^T \frac{\alpha_0}{c p_t}} \\ & \stackrel{(7)}{=} 2 \sqrt{C_1 c \sum_{t=1}^T \left(\frac{2}{\mu c p_t^2} C_2 \right)} / \sum_{t=1}^T \frac{1}{p_t} \\ & = 2 \sqrt{\frac{2C_1 C_2}{\mu} \sum_{t=1}^T \left(\frac{1}{p_t^2} \right)} / \sum_{t=1}^T \frac{1}{p_t}, \end{aligned}$$

completing the proof. \square

4 Experimental Results

4.1 Hardware and Benchmark Datasets

We deploy our implementation on a P775 supercomputer. Each node of this system contains four eight-core 3.84 GHz IBM POWER7 processors, one optical connect controller chip and 128 GB of memory. A single node has a theoretical floating point peak performance of 982 Gflop/s, memory bandwidth of 512 GB/s and bi-directional interconnect bandwidth of 192 GB/s.

We present results on two datasets: CIFAR10 and ImageNet. The CIFAR10 [Krizhevsky and Hinton, 2009] dataset comprises of a total of 60,000 RGB images of size 32×32 pixels partitioned into the training set (50,000 images) and the test set (10,000 images). Each image belongs to one of the 10 classes, with 6000 images per class. For this dataset, we construct a deep convolutional neural network (CNN) with 3 convolutional layers each followed by a pooling layer. The output of the 3rd pooling layer connects, via a fully-connected layer, to a 10-way softmax output layer that generates a probability distribution over the 10 output classes. This neural network architecture closely mimics the CIFAR10 model available as a part of the open-source Caffe deep learning package ([Jia *et al.*, 2014]). The total number of trainable parameters in this network are ~ 90 K (model size of ~ 350 kB). The neural network is trained using momentum-accelerated mini-batch SGD with a batch size of 128 and momentum set to 0.9. As a data preprocessing step, the per-pixel mean is computed over the entire training dataset and subtracted from the input to the neural network.

For ImageNet [Russakovsky *et al.*, 2015], we consider the image dataset used as a part of the 2012 ImageNet

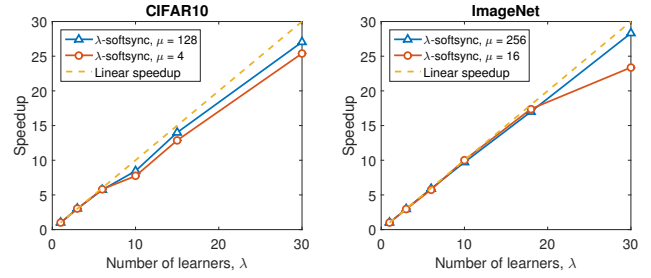


Figure 2: Measured speedup in training time per epoch for (a) CIFAR10 (model size ~ 350 kB) and (b) ImageNet (model size ~ 300 MB)

Large Scale Visual Recognition Challenge (ILSVRC 2012). The training set is a subset of the ImageNet database and contains 1.2 million 256×256 pixel images. The validation dataset has 50,000 images. Each image maps to one of the 1000 non-overlapping object categories. For this dataset, we consider the neural network architecture introduced in [Krizhevsky *et al.*, 2012] consisting of 5 convolutional layers and 3 fully-connected layers. The last layer outputs the probability distribution over the 1000 object categories. In all, the neural network has ~ 72 million trainable parameters and the total model size is 289 MB. Similar to the CIFAR10 benchmark, per-pixel mean computed over the entire training dataset is subtracted from the input image feeding into the neural network.

4.2 Runtime Evaluation

Figure 2 shows the speedup measured on CIFAR10 and ImageNet, for up to 30 learners. Our implementation achieves 22x-28x speedup over benchmarks and different batch sizes. On average, we find that the ASGD runs 50% faster than its SSGD counterpart.

4.3 Model Accuracy Evaluation

For each of the benchmarks, we perform two sets of experiments: (a) setting learning rate fixed to the best-known learning rate for SSGD, $\alpha = \alpha_0$, and (b) tuning the learning rate on a per-gradient basis depending on the gradient staleness τ , $\alpha = \alpha_0/\tau$. It is important to note that when $\alpha = \alpha_0$ and $n = \lambda$ (number of learners) in n -softrsync protocol, our implementation is equivalent to the Downpour-SGD of [Dean *et al.*, 2012]. Albeit at the cost of poor runtime performance, we also train using the hardsync protocol since it guarantees zero gradient staleness and achieves the best model accuracy. Model trained by Hardsync protocol provides the target model accuracy baseline for ASGD algorithm. Further, we perform distributed training of the neural networks for each of these tasks using the n -softrsync protocol for *different* values of n . This allows us to systematically observe the effect of stale gradients on the convergence properties.

CIFAR10

When using a single learner, the mini-batch size is set to 128 and training for 140 epochs using momentum accelerated SGD (momentum = 0.9) results in a model that achieves $\sim 18\%$ misclassification error rate on the test dataset. The

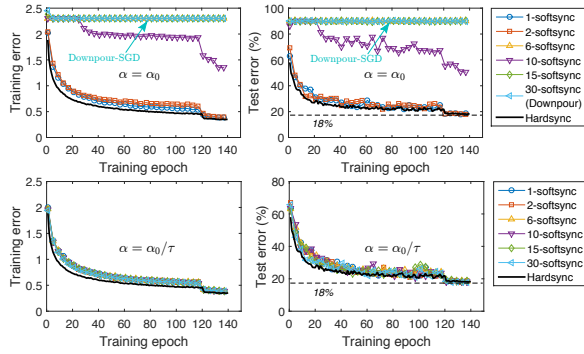


Figure 3: CIFAR10 *results*: (a) Top: training error, test error for different n -softsync protocols, learning rate set as α_0 (b) Bottom: staleness-dependent learning rate of Equation 3. Hardsync (SSGD, black line), Downpour-SGD shown as baseline for comparison. $\lambda = 30$, $\mu = 4$.

base learning rate α_0 is set to 0.001 and reduced by a factor of 10 after the 120th and 130th epoch. In order to achieve comparable model accuracy as the single-learner, we follow the prescription of [Gupta *et al.*, 2015] and reduce the mini-batch size per learner as more learners are added to the system in order to keep the product of mini-batch size and number of learners approximately invariant.

Figure 3 shows the training and test error obtained for different synchronization protocols: hardsync and n -softsync, $n \in (1, \lambda)$ when using $\lambda = 30$ learners. The mini-batch size per learner is set to 4 and all the other hyperparameters are kept unchanged from the single-learner case. Figure 3 top half shows that as the gradient staleness is increased (achieved by increasing the splitting parameter n in n -softsync protocol), there is a gradual degradation in SGD convergence and the resulting model quality. In the presence of large gradient staleness (such as in 15, and 30-softsync protocols), training fails to converge and the test error stays at 90%. In contrast, Figure 3 bottom half shows that when these experiments are repeated using our proposed staleness-dependent learning rate scheme of Equation 3, the corresponding curves for training and test error for different n -softsync protocols are virtually indistinguishable (see Figure 3 bottom half). Irrespective of the gradient staleness, the trained model achieves a test error of $\sim 18\%$, showing that *proposed learning rate modulation scheme is effective in bestowing upon the training algorithm a high degree of immunity to the effect of stale gradients*.

ImageNet

With a single learner, training with mini-batch size of 256, momentum 0.9 results in top-1 error of 42.56% and top-5 error of 19.18% on the validation set at the end of 35 epochs. The initial learning rate α_0 is set to 0.01 and reduced by a factor of 5 after the 20th and again after the 30th epoch. Next, we train the neural network using 18 learners, different n -softsync protocols and reduce the mini-batch size per learner to 16.⁴

⁴For this benchmark, we use a variant of SGD – adaptive learning rate method (AdaGrad [Duchi *et al.*, 2011; Dean *et al.*, 2012]) – to

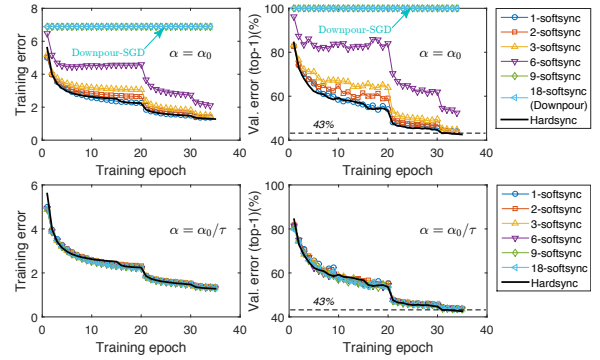


Figure 4: ImageNet *results*: (a) Top: training error, top-1 validation error for different n -softsync protocols, learning rate set as α_0 (b) Bottom: staleness-dependent learning rate of Equation 3. Hardsync(SSGD, black line), Downpour-SGD shown as baseline for comparison. $\lambda = 18$, $\mu = 16$.

Figure 4 top half shows the training and top-1 validation error when using the learning rate that is the same as the single learner case α_0 . The convergence properties progressively deteriorate as the gradient staleness increases, failing to converge for 9 and 18-softsync protocols. On the other hand, as shown in Figure 4 bottom half, automatically tuning the learning rate based on the staleness results in nearly identical behavior for all the different synchronization protocols. These results echo the earlier observation that the proposed learning rate strategy is effective in combating the adverse effects of stale gradient updates. Furthermore, adopting the staleness-dependent learning rate helps avoid the laborious manual effort of tuning the learning rate when performing distributed training using ASGD.

Summary With the knowledge of the initial learning rate for SSGD (α_0), our proposed scheme can automatically tune the learning rate so that distributed training using ASGD can achieve accuracy comparable to SSGD while benefiting from near linear-speedup.

5 Conclusion

In this paper, we study how to effectively counter gradient staleness in a distributed implementation of the ASGD algorithm. We prove that by using our proposed staleness-dependent learning rate scheme, ASGD can converge at the same rate as SSGD. We quantify the distribution of gradient staleness in our framework and demonstrate the effectiveness of the learning rate strategy on standard benchmarks (CIFAR10 and ImageNet). The experimental results show that our implementation achieves close to linear speedup for up to 30 learners while maintaining the same convergence rate in spite of the varying degree of staleness in the system and across vastly different data and model sizes.

Acknowledgement

Xiangru Lian and Ji Liu are supported by the NSF grant CNS-1548078 and the startup fund at University of Rochester.

improve the stability during distributed training.

References

- [Agarwal and Duchi, 2011] Alekh Agarwal and John C Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems*, pages 873–881, 2011.
- [Chan and Lane, 2014] William Chan and Ian Lane. Distributed asynchronous optimization of convolutional neural networks. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [Chilimbi *et al.*, 2014] Trishul Chilimbi, Yutaka Suzue, Johnson Apacible, and Karthik Kalyanaraman. Project adam: Building an efficient and scalable deep learning training system. OSDI’14, pages 571–582. USENIX Association, 2014.
- [Dean *et al.*, 2012] Jeffrey Dean, Greg S. Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Quoc V. Le, Mark Z. Mao, MarcAurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, and Andrew Y. Ng. Large scale distributed deep networks. In *NIPS*, 2012.
- [Duchi *et al.*, 2011] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [Gupta *et al.*, 2015] S. Gupta, W. Zhang, and J. Milthorpe. Model accuracy and runtime tradeoff in distributed deep learning. *ArXiv e-prints*, September 2015.
- [Ho *et al.*, 2013] Qirong Ho, James Cipar, Henggang Cui, Seunghak Lee, Jin Kyu Kim, Phillip B Gibbons, Garth A Gibson, Greg Ganger, and Eric P Xing. More effective distributed ml via a stale synchronous parallel parameter server. In *Advances in Neural Information Processing Systems*, pages 1223–1231, 2013.
- [Jia *et al.*, 2014] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [Krizhevsky and Hinton, 2009] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Computer Science Department, University of Toronto, Tech. Rep*, 1(4):7, 2009.
- [Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [Lian *et al.*, 2015] Xiangru Lian, Yijun Huang, Yuncheng Li, and Ji Liu. Asynchronous parallel stochastic gradient for nonconvex optimization. In *NIPS*, 2015.
- [Liu and Wright, 2015] Ji Liu and Stephen J Wright. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.
- [Liu *et al.*, 2013] Ji Liu, Stephen J. Wright, Christopher Ré, Victor Bittorf, and Srikrishna Sridhar. An asynchronous parallel stochastic coordinate descent algorithm. *ArXiv*, 2013.
- [MPI-Forum, 2012] MPI-Forum. MPI: A Message-Passing Interface Standard Version 3.0, 2012.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, pages 1–42, April 2015.
- [Seide *et al.*, 2014] Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [Strom, 2015] Nikko Strom. Scalable distributed dnn training using commodity gpu cloud computing. In *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [Tsitsiklis *et al.*, 1986] John N Tsitsiklis, Dimitri P Bertsekas, Michael Athans, et al. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE transactions on automatic control*, 31(9):803–812, 1986.
- [Wu *et al.*, 2015] Ren Wu, Shengen Yan, Yi Shan, Qingqing Dang, and Gang Sun. Deep image: Scaling up image recognition. *CoRR*, abs/1501.02876, 2015.
- [Zhang *et al.*, 2015] Ruiliang Zhang, Shuai Zheng, and James T. Kwok. Fast distributed asynchronous SGD with variance reduction. *CoRR*, abs/1508.01633, 2015.