

Trade-offs of Local SGD at Scale: An Empirical Study

Jose Javier Gonzalez Ortiz

JOSEJG@MIT.EDU

MIT CSAIL

Jonathan Frankle

JFRANKLE@MIT.EDU

MIT CSAIL

Mike Rabbat

MIKERABBAT@FB.COM

Facebook AI Research

Ari Morcos

ARIMORCOS@FB.COM

Facebook AI Research

Nicolas Ballas

BALLASN@FB.COM

Facebook AI Research

Abstract

As datasets and models become increasingly large, distributed training has become a necessary component to allow deep neural networks to train in reasonable amounts of time. However, distributed training can have substantial communication overhead that hinders its scalability. One strategy for reducing this overhead is to perform multiple unsynchronized SGD steps independently on each worker between synchronization steps, a technique known as *local SGD*. We conduct a comprehensive empirical study of local SGD and related methods on a large scale image classification task. We find that performing local SGD comes at a price: lower communication costs (and thereby faster training) are accompanied by lower accuracy. This finding is in contrast from the smaller-scale experiments in prior work, suggesting that local SGD encounters challenges at scale. We further show that incorporating the slow momentum framework of Wang et al. (2020) consistently improves accuracy without requiring additional communication, hinting at future directions for potentially escaping this trade-off.

Keywords: Deep Learning, Distributed Optimization, Local SGD, Convolutional Neural Networks

1. Introduction

As datasets and models continue to grow in size, it has become a common practice to train deep neural networks in a distributed manner across multiple hardware *workers* Goyal et al. (2017); Shallue et al. (2018). Most deep learning models are currently optimized using some variant of stochastic gradient descent Robbins and Monro (1951), often using a mini-batch approach Bottou (2010); Dekel et al. (2012). In distributed scenarios, the communication overhead necessary to synchronize gradients between workers can quickly dominate the time necessary to compute the model updates, hindering the scalability of this approach. Moreover, because of the serial nature of neural network training, all the nodes must wait until the synchronization completes, and performance is therefore dependent on the slowest node Dutta et al. (2018); Ferdinand et al. (2020).

These issues have motivated the development of optimization algorithms that reduce the amount of communication between workers. A simple yet practical example is *local SGD* (closely related to *federated averaging* (McMahan et al., 2017)), where, instead of synchronizing the

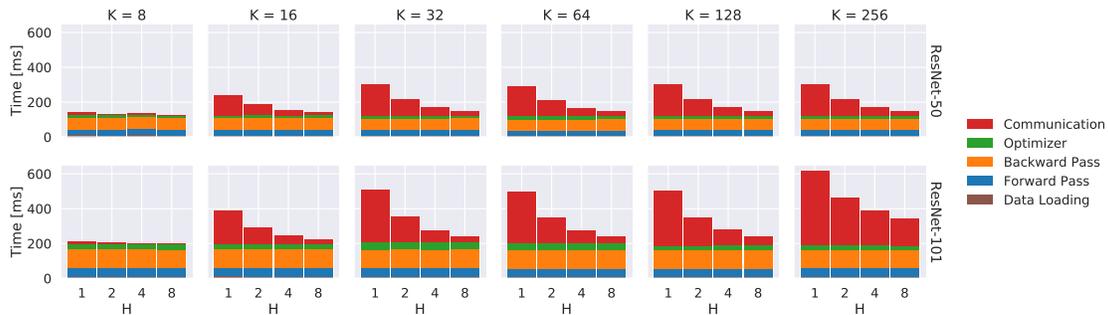


Figure 1: Breakdown of the average wall-clock time per iteration during the training process. Results are reported for various numbers of workers (K) and numbers of local steps (H). Every node has 8 workers. The main difference is the communication time, which decreases as we reduce the frequency of model averaging. For minibatch SGD (which we label as $H = 1$) with multiple nodes ($K > 8$), the communication time dominates other parts of training.

gradients at every iteration, each worker performs multiple SGD steps locally and then averages the model weights across all workers Zhang et al. (2016). Local SGD has been shown to have good optimization properties from a theoretical standpoint Stich (2018); Zhang et al. (2016); Woodworth et al. (2020). However, while local SGD does speed up training, the resulting models are often less accurate compared to a synchronous minibatch SGD baseline (Lin et al., 2018).

Post-local SGD is a variant of local SGD introduced by Lin et al. (2018) with the goal of remedying these problems. Post-local SGD divides training into two phases. In the first phase, workers perform synchronous minibatch SGD; in the second, they switch to local SGD. Lin et al. claim that this approach generalizes better on large batch training than both local SGD and minibatch SGD while reducing communication for the second phase of training. The majority of the analysis of local SGD and post-local SGD reported in Lin et al. (2018) is on CIFAR-10 Krizhevsky et al. (2009).

Motivated by these results, we perform a thorough analysis of local SGD and post-local SGD on the ImageNet-1k Russakovsky et al. (2015) classification task, a *de facto* benchmark for large-scale vision classification problems. As Figure 1 shows, inter-node communication can dominate training time, becoming a bottleneck in the training process. We complement the analysis of Lin et al. (2018), studying how the choice of learning rate schedule and the point at which to switch phases affect the generalization accuracy of models trained with post-local SGD. We find that post-local SGD at ImageNet-scale is a double-edged sword: decreases in communication costs (by increasing the number of local steps) are accompanied by decreases in accuracy. As a result, practitioners interested in post-local SGD must weigh the trade-offs between training speedup and reduction in the accuracy of the final model. Looking ahead, our analysis of the interaction between post-local SGD, learning rates, and momentum points toward potential opportunities to escape these trade-offs.

Contributions. Our main contributions are as follows:

1. We perform a comprehensive empirical study on ImageNet that identifies previously unreported scalability limitations of local and post-local SGD. Our analysis highlights how, when compared to the fully synchronous baseline, local and post-local SGD suffer from non-trivial accuracy drops as workers or local steps increase.

2. Our analysis is the first to identify that post-local SGD performance heavily relies on the choice of hyperparameters, including learning rate schedule and switching point.
3. We show that using slow momentum Wang et al. (2020) together with post-local SGD achieves a better quality-performance trade-off.
4. We show that switching to local SGD has a regularization effect on optimization that is only beneficial in the short term, suggesting it is always better to make the switch later in training.

2. Background and Related Work

Minibatch SGD. Neural networks are typically trained with *minibatch SGD*. In minibatch SGD, the dataset $\mathcal{D} = \{(x_i, y_i)\}_{i \in [N]}$ is divided into non-overlapping subsets of size B known as *minibatches*. Gradient descent is performed sequentially on these minibatches, passing through the entire dataset over the course of an *epoch*. The dataset is typically randomly shuffled before each epoch, meaning the minibatch composition and order are different on each pass through the dataset. See Algorithm 1 below for full details. In practice, networks are often trained in a *distributed* data-parallel fashion across K workers Li et al. (2020). Each worker has a separate copy of the weights and computes gradients using a disjoint subset of the data. The entire dataset is reshuffled and split among workers at the beginning of every epoch. After the backward pass, gradients are averaged across workers before updating the model weights.

Local SGD. In local SGD, described in Algorithm 2, the workers update the local copies of their weights, and every $H > 1$ iterations they *synchronize* the weights across workers by averaging the weights stored on each worker. Papers on local SGD typically credit McDonald et al. (2009), Zinkevich et al. (2010) and McDonald et al. (2010) with pioneering local SGD in pre-deep learning settings; these works train models to completion and average the final parameters. Zhang et al. (2016) explore local SGD with periodic averaging of models throughout training; they prove that it converges in convex settings (in fact, faster than minibatch SGD) in the face of gradient variance and show that it can optimize a LeNet-5 network on MNIST. Povey et al. (2014) and Su and Chen (2015) use local natural gradient descent with periodic synchronization to optimize deep models. Kamp et al. (2018) synchronize updates among neighboring workers using a gossiping protocol. Zhou and Cong (2017), Stich (2018), Wang and Joshi (2018b), Yu et al. (2019), and Woodworth et al. (2020) prove convergence properties of local SGD under various technical assumptions. Zhou and Cong (2017) experiment on CIFAR-10 and show that accuracy drops off if synchronization occurs too infrequently. Wang and Joshi (2018a) start with infrequent averaging and then increase the communication over the course of training. Periodic model averaging also takes place in *federated learning*, which focuses on performing SGD in settings where data is distributed across many devices and a model can only be updated on the device where the data resides. McMahan et al. (2017), who initiated this research literature, perform parallel SGD between devices with synchronization on every step, although they explore synchronizing after many iterations. Local SGD is common in federated learning systems for performance reasons (Smith et al., 2017).

Post-local SGD. Post-local SGD (Lin et al., 2018) is a hybrid in which training occurs synchronously for the first part of training and switches to local SGD later in training. Post-local SGD involves performing synchronous SGD for the first T steps of training and local

Algorithm 1 Minibatch SGD with learning rate schedule $\gamma(t)$, batch size B , initial weights w_0 , and loss ℓ .

```

1:  $S \leftarrow \lceil \frac{N}{B} \rceil$  (iterations per epoch)
2: for each epoch  $e$  do
3:   Shuffle dataset  $\mathcal{D} = \{x_i, y_i\}_{i \in N}$ 
4:   for each iteration  $m$  in 1 to  $S$  do
5:      $t \leftarrow eS + m$  (the current step of training)
6:      $o \leftarrow (m-1)B$ 
7:      $w_t \leftarrow w_{t-1} - \gamma(t) \frac{1}{B} \sum_{i=1}^B \nabla \ell(f(x_{o+i}; w_t), y_{o+i})$ 

```

Algorithm 2 Local SGD with learning rate schedule $\gamma(t)$, batch size B , initial weights w_0 , loss ℓ , and R workers.

```

1:  $S \leftarrow \lceil \frac{N}{BK} \rceil$  (iterations per epoch)
2: for each epoch  $e$  do
3:   Shuffle dataset  $\mathcal{D} = \{x_i, y_i\}_{i \in N}$ 
4:   for each iteration  $m$  in 1 to  $S$ , at worker  $k$  (in parallel) do
5:      $t \leftarrow eS + m$  (the current step of training)
6:      $o \leftarrow (m-1)B + (k-1)SB$ 
7:      $w_t^{(k)} \leftarrow w_{t-1}^{(k)} - \gamma(t) \frac{1}{B} \sum_{i=1}^B \nabla \ell(f(x_{o+i}; w_t), y_{o+i})$ 
8:     if  $t \bmod H = 0$  then
9:        $w_t^{(k)} \leftarrow \frac{1}{K} \sum_{k=1}^K w_t^{(k)}$  (synchronize)

```

SGD thereafter. The hyperparameters may differ between the synchronous and local phases. For example, the local phase may use a different learning rate or batch size. Lin et al. suggest post-local SGD as a replacement for standard, synchronous, large-batch training schemes that “significantly improves the generalization performance.”

Other approaches for reducing communication costs include using gradient compression to reduce the size of updates (Alistarh et al., 2017; Wen et al., 2017; Bernstein et al., 2018; Karimireddy et al., 2019; Vogels et al., 2019) and modifying the communication patterns between nodes Lian et al. (2017, 2018); Assran et al. (2019), often by introducing approximation or asynchrony.

3. Experimental Setup

Our aim in this work is to experimentally investigate the performance of post-local SGD and the related methods on a large-scale distributed training workload. Here, we describe the experimental setup for the experiments. Further implementation details are listed in Appendix A.4

Dataset. We focus on the computer vision task of image classification on the ImageNet dataset (Deng et al., 2009; Russakovsky et al., 2015). This is a 1000-way classification task with ~ 1.28 million training images. Performance is evaluated by top-1 error on 50,000 held out validation images. For the training examples, we follow the standard channel normalization and data augmentation scheme as in He et al. (2016b) and use the default training and validation splits from Russakovsky et al. (2015).

Models. We use the ResNet-50 and ResNet-101 architectures (He et al., 2016a). We initialize the weights as in He et al. (2015) and incorporate the modifications described in

Goyal et al. (2017) for initializing the fully connected layer and the last Batch Normalization layer of the residual blocks. We use L2 weight decay with $\lambda=10^{-4}$.

Training. Following recommendations from Goyal et al. (2017), we scale the learning rate linearly with the global batch size according to $\eta=KB0.1/256$, where K is the number of workers and B is the per-worker batch size, which is set to 32. In addition, we apply a learning rate warm-up strategy that starts from 0.1 and linearly increases the learning rate every iteration during the first five epochs of training so that the target learning rate η is reached at the end of the fifth epoch. All experiments are trained with Nesterov momentum (Nesterov, 2013) with $\beta=0.9$. We also apply momentum correction as described in Goyal et al. (2017) to stabilize training whenever the learning rate value is modified. Unless mentioned otherwise, we use a step-wise learning rate schedule that decays the learning rate by a factor of 10 at the end of the 30th, 60th and 80th epochs, as in He et al. (2016a).

Local SGD. For local SGD and post-local SGD, we average the model weights of all the workers before computing the validation performance at the end of every epoch. This model average is performed outside of the training process. We consistently found that averaging the models before computing validation was beneficial to generalization performance. For post-local SGD, before the switch to local SGD is performed, we average the gradients and not the model weights. During the post-local SGD regime the per-worker momentum buffers are kept local and not synchronized, this is a common approach Lian et al. (2017); Assran et al. (2019); Wang et al. (2020).

4. The Benefits of (Post-)Local SGD

While local SGD undoubtedly reduces communication and thus overall runtime, the specific speedup varies from task to task. In this section we illustrate the scalability challenges of minibatch SGD in distributed settings and show that local SGD and post-local SGD achieve substantially better performance by reducing communication. To reduce the effect of run-to-run variation, we execute ten independent runs (seeds) and report the median performance values. We consider local SGD where the model is averaged every H steps from the beginning of training and post-local SGD with a switching point at the first learning rate decay (epoch 30) as suggested in Lin et al. (2018).

Local SGD reduces communication overhead. The main performance cost of scaling minibatch SGD to a distributed setting is the communication overhead to synchronize gradients before updating model weights at every iteration. If this overhead is comparable to the other parts of training, the training process will be significantly slowed down, leading to reduced performance. We show that, for the classification task considered on a 10 Gb/s Ethernet interconnect connection, the performance benefits of local SGD are substantial.

Figure 1 presents a breakdown of the running times of different parts of the training loop for a variety of settings. Communication time refers to performing an all-reduce to average the weights across all workers. As expected, all except for the communication time stay approximately constant for different numbers of workers. For a single node, i.e., 8 GPUs, the communication time is quite low, indicating that intra-node communication is negligible with respect to other parts of training. However, for multi-node minibatch SGD ($H=1$ and $K>8$ in the figure), the communication time dominates the other steps as the number of workers K increases. On the other hand, for local SGD ($H>1$ in the figure) as we increase the number

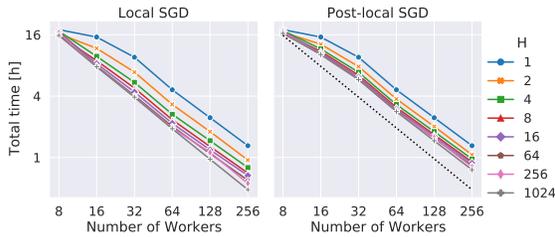


Figure 2: Total time in hours as a function of the number of workers for local SGD and post-local SGD. We report median time over 10 separate runs.

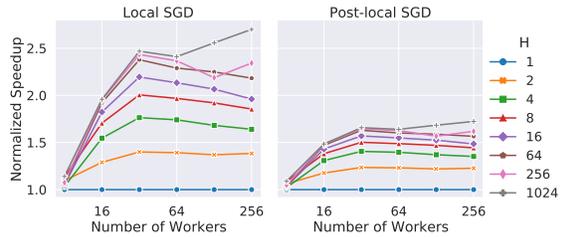


Figure 3: Normalized Speedup of local and post-local SGD for a given world size. I.e., the ratio between the time per epoch for local SGD and the corresponding minibatch SGD.

of local steps H the communication overhead is amortized over several iterations, leading to a reduced iteration time.

Local SGD scales to large distributed settings. While we see that communication overhead is a large fraction of the time per iteration in Figure 1, we also observe that the communication time is not proportional to the number of workers. Consequently, despite the introduced communication overhead when going from single-node to multi-node training, we should observe good scalability trends as we increase the overall number of workers. This is depicted in Figure 2, where the total time, in hours, to complete 90 epochs of training a ResNet-50 is presented for local and post-local SGD for different numbers of workers and local steps. The approximately linear trends indicate that we can successfully make use of additional nodes to achieve reduced training times without running into other performance bottlenecks. In general, we observe that in both minibatch SGD (indicated as $H = 1$ in Figure 2) and local SGD, doubling the number of workers halves the time taken to run an epoch. Whereas for minibatch SGD going from a single node to multiple nodes (i.e., 8 to 16 workers) has an effect on running time, local SGD shows better scalability asymptotically approaching linear speedup.

Reducing communication is effective but has diminishing returns. While we observed overall training time reductions, it is still unclear what specific speedup reducing communication is accounting for. To measure the contribution of local SGD, we normalize times when $H > 1$ by the time for $H = 1$ with the same number of workers, resulting in the normalized speedup shown in Figure 3. This normalized speedup corresponds to the speedup obtained by using local SGD with $H > 1$ relative to using $H = 1$ (i.e., synchronizing every iteration, equivalent to minibatch SGD) with the same number of workers. With more than one node ($K > 8$), local SGD provides a substantial speedup. For instance, 32 or 64 workers communicating every 4 or 8 local steps leads to relative speedups close to 1.75 and 2 respectively. However, as we keep increasing the number of steps between synchronizations, the speedup increases with diminishing returns since the amount of synchronizations decays exponentially.

5. Trade-offs of (Post-)Local SGD

5.1 Trade-offs of local model updates

The measured speedup benefits of reducing communication during the training process suggest that (post-)local SGD presents an opportunity to substantially reduce training times. As such,

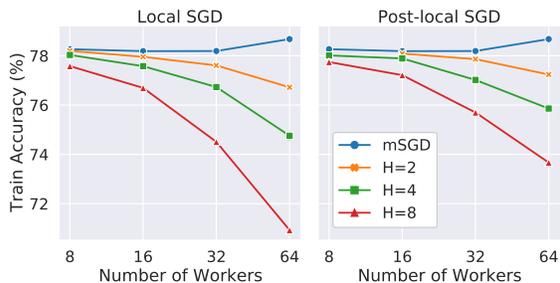


Figure 4: Training Accuracy as a function of the number of workers for local SGD (left) and post-local SGD (right). Separate curves refer to experiments with the same number of local steps H between model averages. mSGD corresponds to the minibatch SGD baseline.

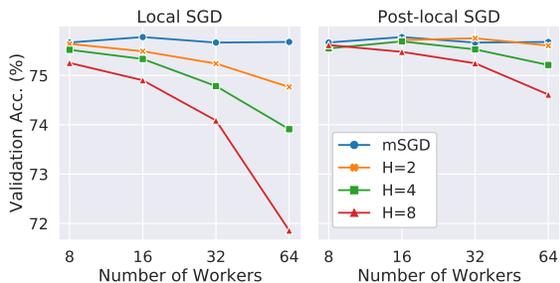


Figure 5: Validation Accuracy as a function of the number of workers for local SGD (left) and post-local SGD (right). Separate curves refer to experiments with the same number of local steps H between model averages. mSGD corresponds to the minibatch SGD baseline.

we now analyze the effect on generalization of these techniques to understand whether this performance comes at a cost. We explore the scalability of both local SGD and post-local SGD (switching at epoch 30 unless otherwise noted) as we increase the total number of workers (i.e., GPUs) and as we increase the number of local steps between model averaging. We further study varying the amount of communication of post-local SGD by changing the switching point between the phases. We show that on ImageNet-1k, local SGD and post-local SGD struggle to scale in any of these axes without causing a drop in validation performance.

Increasing the number of workers. Our first analysis focuses on the trade-off between increasing the number of workers and the effects on generalization as measured by the decrease in validation accuracy. We find that increasing workers has a negative effect on accuracy for both local and post-local SGD, consistently leading to worse training and validation results. Figures 4 and 5 show training and validation accuracy as the number of workers increases. We include results for various numbers of local steps. Neither local SGD nor post-local SGD maintain accuracy as the number of workers increases. Although post-local SGD experiences smaller absolute drops in accuracy compared to local SGD, this comes at the expense of more communication and consequently more runtime. The loss in accuracy seems to be a direct result of the optimization mechanics of local SGD, since the minibatch baseline does not experience this trade-off as the number of workers increases.

Reducing synchronization frequency. As we saw from the earlier timing results, reducing the frequency of synchronizations is the main mechanism that local SGD has for reducing the communication overhead. Here, we investigate the effect of increasing the number of local steps between model averaging synchronizations, identifying a very similar trade-off to increasing the number of workers. Figures 4 and 5 show training and validation accuracy for varying numbers of local steps H for local and post-local SGD. For both methods, as the number of local steps grows larger, training and validation accuracy monotonically decrease.

Towards a unified trade-off. We showed that increasing the number of workers K or local steps H led to lower accuracy. Interestingly, neither of these factors was clearly dominant and the absolute drop in training and validation accuracy is similar as we increase either of

them. We show that this phenomenon can be better viewed in terms of a different statistic: total number of local model updates between synchronizations, i.e., the product $K \times H$.

Both of the factors K and H affect the optimization mechanics. As we increase the number of workers, K , the global batch size linearly increases because every worker has a fixed local batch size. Consequently, since the number of epochs is fixed, the total number of model updates decreases by the same factor. In local SGD models are averaged after H local steps. While model averaging usually leads to better performance, it depends on how much the local models have diverged since the last synchronization.

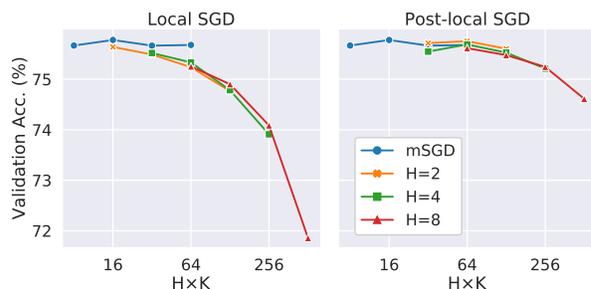


Figure 6: Validation Accuracy as a function of the number of local model updates ($H \times K$) for local SGD (left) and post-local SGD (right). mSGD values correspond to the minibatch SGD baseline.

Figure 6 shows validation accuracy as a function of the product $H \times K$. We found the same pattern holds for training accuracy (see Figure 12 in the Appendix). When viewed this way, we can observe that the trade-off curves for different numbers of workers and frequencies line up for both training and validation. This suggests that the main factor influencing final model quality is the total number of local model updates between synchronizations. Consequently, to maintain model accuracy, an increase in the number of local steps should be matched with a proportional decrease in the number of workers and *vice versa*.

5.2 The effect of hyperparameters

The post-local switching point poses a trade-off. Post-local SGD divides training into two phases. In the first phase, workers perform synchronous minibatch SGD; in the second, they switch to local SGD. However, the iteration at which to switch is another hyperparameter that affects the amount of communication; here, we study whether it also affects final accuracy. We find that the switching point presents a trade-off between training time and final accuracy, with later switching points improving accuracy at the cost of additional communication.

Up to this point, all the reported experiments switched to local SGD at epoch 30 out of 90. This is when the first learning rate decay happens with a step-wise schedule with decays at epochs 30, 60 and 80 (as recommended by Lin et al. (2018)). However, switching earlier can reduce the amount of communication. Conversely, previous results showed consistent improvements from local SGD to post-local SGD and, since local SGD can be thought of as a post-local SGD with a switching point at epoch 0, we also explore whether we can recover the lost model accuracy by switching at a later point.

To study this phenomena, we select switching points at epochs where the learning rate decreases (30, 60 and 80) and intermediate epochs (15, 45, 75, and 85). Figure 7 depicts these results for both ResNet-50 and ResNet-101 for various numbers of local steps. Both networks present similar patterns for the chosen values of switching points. Notably, the learning rate decay points seem to have a significant positive effect into the final performance for this task.



Figure 7: Validation accuracy for post-local SGD with step-wise learning rate schedule as a function of the epoch when the switch to local SGD is done. Results are for $K = 64$ workers and various numbers of local steps H . mSGD is the baseline minibatch SGD method, synchronizing at every update.

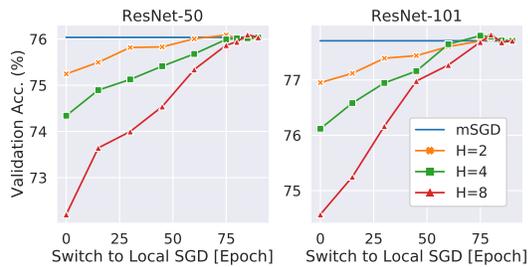


Figure 8: Validation accuracy for post-local SGD with half cosine learning rate schedule as a function of the epoch when the switch to local SGD is done. Results are for $K = 64$ workers and various numbers of local steps H . mSGD is the baseline minibatch SGD method, synchronizing at every update.

Post-local SGD performance heavily depends on learning rate schedule. One unexplored aspect of post-local SGD is the interaction between the switching time heuristic and the choice of learning rate schedule. Here, we investigate how the choice of learning rate schedule affects the generalization results of post-local SGD. Recently, learning rate schedules other than step-wise, such as cosine annealing, have gained popularity due to the competitive results they provide without having to carefully specify the learning rate decay points (He et al., 2019; Radosavovic et al., 2019). As an alternative to the step-wise schedule, we consider a half-period cosine schedule which sets the learning rate to $\eta_t = \eta_0(1 + \cos(\pi \cdot t / (T_{\max}))) / 2$, where η_0 is the initial learning rate, t is the current epoch, and T_{\max} is the total number of epochs. As with the step-wise schedule, we modify the cosine schedule to have a linear warm-up phase updated every iteration over the first five epochs of training. The learning rate is updated at the end of every epoch except for the warm-up phase, when it is updated at every iteration.

Figure 8 presents a trade-off analysis of how the switching point affects the final training accuracy of the model with a half cosine learning rate schedule. Unlike before, we see a more monotonic trade-off between the final validation accuracy and when the switch to local SGD is performed. When compared to Figure 7, the results in Figure 8 indicate that the scale of the learning rate when the switch to local SGD is performed is crucial to ensure that the drop in validation accuracy is minimal.

5.3 Post-local SGD behavior depends on the task scale

Post-local SGD was proposed as a way to remedy the generalization gap present in large batch training Chen and Huo (2016); Keskar et al. (2016); Hoffer et al. (2017); Shalloe et al. (2018). However, our observations do not align with the generalization improvements identified by Lin et al. (2018) for ResNet-20 on CIFAR-10 for increasing values of K and H even when exploring different switching points. We found similar patterns for a downscaled version of the previously presented experiments. We performed analogous experiments on TinyImageNet, a dataset which comprises a subset of ImageNet at a lower resolution, using a ResNet-50 (with

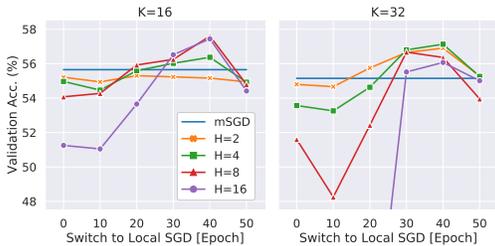


Figure 9: Validation accuracy for post-local SGD on TinyImageNet as a function of the epoch when switching to local SGD. Learning rate is decayed at epochs 30 and 50.

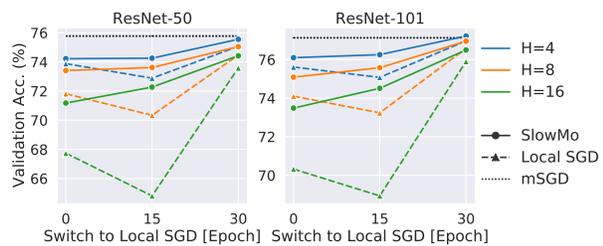


Figure 10: Validation accuracy for various local SGD settings with and without the slow momentum update. SlowMo yields better performance than local or post-local SGD.

slight modifications to work on the smaller dataset). Figure 9 shows that post-local SGD (when performed after the first learning rate decay point) does produce accuracy gains with respect to the synchronous minibatch baseline. This contrasts with our ImageNet experiments, where little to no accuracy gains are achieved when switching to local SGD during training. We have thus identified an example where distributed optimization behaves differently for tasks of different scales.

We note that our ImageNet results are not directly comparable to the results of Lin et al. (2018) on ImageNet since we do not use the LARS optimizer (You et al., 2017) for the main body of experiments; we use SGD with Nesterov momentum instead. To ensure that choice of optimizer is not a confounding factor in this analysis, we performed analogous experiments using LARS and the results are included in Figure 13 in the Appendix. Our analysis indicates that post-local SGD + LARS suffers from a similar trade-off. When switching to local SGD in the early phase of training, final model accuracy drops with respect to the minibatch baseline. In contrast, we do see that LARS trade-off is more linear and provides some accuracy improvements when the switch occurs later in training.

6. Shifting the Trade-off with Slow Momentum

We have identified several trade-offs of local and post-local SGD in terms of training time and model accuracy. We now explore modifying the optimization process as a way to mitigate the drop in performance that local-SGD and post-local SGD experience, arriving at better trade-off. To that end we make use of the slow momentum (SlowMo) method Wang et al. (2020), which builds on top of the BMUF method (Chen and Huo, 2016) by generalizing local SGD in a way similar to how SGD with momentum generalizes SGD. In particular, the slow momentum algorithm works by adding an additional momentum update after model averaging steps. SlowMo consistently provided improvements when incorporated into algorithms with reduced communication, including local SGD and other decentralized algorithms. SlowMo does not require additional communication since the model averages provide enough information to compute the slow gradients and perform the update.

We evaluate SlowMo on top of local SGD and post-local SGD with switching points at 15 or 30 epochs. Figure 10 shows that adding the slow momentum updates leads to improved validation accuracy in all cases. Incorporating SlowMo at the switching point at epoch 15 leads

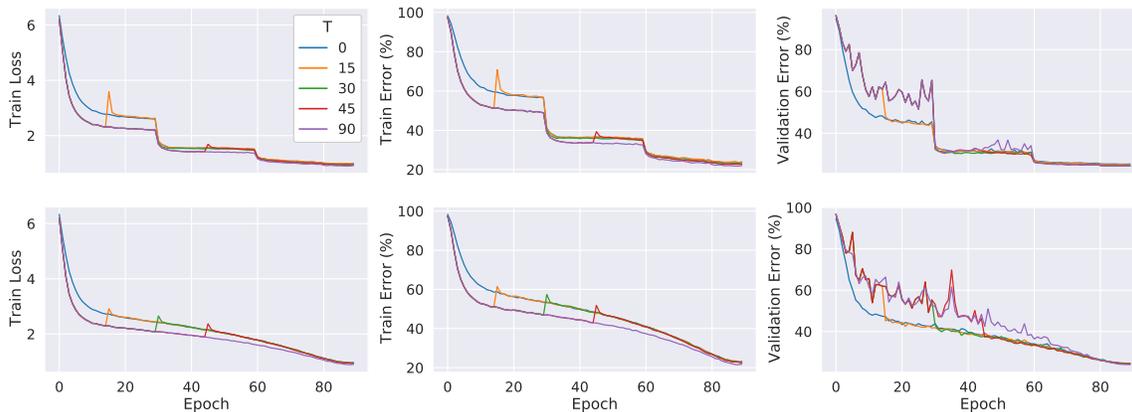


Figure 11: Training Loss (first column), Training Error (second column) and Validation Error (third column) for post-local SGD experiments with different switching points T and number of local steps $H = 4$. Top row corresponds to step-wise learning rate schedule with decays at 30, 60 and 80 epochs. Bottom row corresponds to a half cosine schedule. For a given number of local steps, regardless of when the switch to local SGD is performed, the loss and error curves follow the same general trajectory. We found this pattern consistently across different models, numbers of workers and numbers of local steps.

to noticeably better accuracy, outperforming the local SGD case and inverting the behavior previously seen. Although SlowMo is not able to fully bridge the accuracy gap compared to synchronous minibatch SGD, it does mitigate the drop by reliably boosting performance for different settings.

7. The Effect of Local SGD on Generalization

Up to this point, our analysis has focused on the accuracy achieved by post-local SGD. However, the best validation accuracy of the model might not be representative of how local SGD and post-local SGD affect the optimization of the model over the course of training. Here, we study how the loss and the errors evolve over time for different switching points revealing that switching to local-SGD has a regularization effect regardless of when it is performed. We show that switching to local SGD has a regularization effect on optimization that is only beneficial in the short term, suggesting it is always better to make the switch later in training.

In Figure 11 we plot the training loss, top-1 training error and top-1 validation error for post-local SGD for various switching points. $T = 0$ corresponds to local-SGD and $T = 90$ corresponds to the minibatch baseline. We observe a regularization effect when the switch to local SGD is performed for the post-local SGD method: the training error is higher than the minibatch baseline while the validation error is lower. Results are for 64 workers and 4 local steps between model averages but we found these patterns consistently for different numbers of workers and number of local steps. More interestingly, after the switching point, the curves' various post-local SGD models closely follow the same trajectory for both training and validation metrics. This pattern holds for both learning rate schedules although it is more noticeable for the half cosine learning rate schedule because of the higher absolute change.

If we now consider these results along with the earlier analysis that looked at the final validation accuracy, we can conclude that, despite the switch to local-SGD improving generalization in the short term, the long term effects are detrimental to final model accuracy. We believe that this interplay has not been studied thoroughly before and that it appears to be a critical component when making use of strategies that switch from a fully synchronous regime to local training, like post-local SGD.

8. Limitations

A limitation of this work is that it only studies the ImageNet-1k classification task. Nevertheless, with over 1.28 million images and 1000 classes, ImageNet-1k is considered a *de facto* benchmark for natural image vision classification tasks. Our experiments study a wide variety of settings and account for more than 47,000 GPU hours of compute. Moreover, existing analysis on CIFAR-10 along with our experiments using TinyImageNet seem to indicate that the behavior of post-local SGD depends on the task scale. Since methods such as local SGD are used in a distributed setting mostly required by large networks and datasets, we deem important that a benchmark of this size is used. Lastly, while we do not propose new algorithmic approaches, our analysis identifies trade-offs of local and post-local SGD not reported before which impact its applicability by researchers and practitioners alike. For instance, the dependence between final model accuracy and factors like learning rate or switching point is critical to the viability of post-local SGD.

9. Conclusion

In this paper we present a thorough analysis of local and post-local SGD at scale by evaluating them on the ImageNet-1k classification task. We identify several scalability limitations that local and post-local SGD experience and analyze the trade-offs they present in terms of final model accuracy. We find that the number of workers and the number of local steps for local SGD present accuracy trade-offs, and we unify these hyperparameters into a trade-off expressed in terms of number of local model updates between synchronizations. We characterize the behavior of post-local SGD for different switching points along the training process, revealing a similar trade-off between communication and model accuracy. Furthermore, we identify that the choice of learning rate schedule has a large impact for the resulting accuracy of post-local SGD. We further show that incorporating the slow momentum framework of Wang et al. (2020) consistently improves accuracy without requiring additional communication. Lastly, our analysis reveals that switching to local SGD has a regularization effect on optimization that is beneficial in the short term but detrimental to final model accuracy.

References

- Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*, pages 1709–1720, 2017.
- Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Mike Rabbat. Stochastic gradient push for distributed deep learning. In *International Conference on Machine Learning*, pages

- 344–353. PMLR, 2019.
- Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, and Anima Anandkumar. signsgd with majority vote is communication efficient and fault tolerant. *arXiv preprint arXiv:1810.05291*, 2018.
- Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010*, pages 177–186. Springer, 2010.
- Kai Chen and Qiang Huo. Scalable training of deep learning machines by incremental block training with intra-block parallel optimization and blockwise model-update filtering. In *2016 IEEE international conference on acoustics, speech and signal processing (icassp)*, pages 5880–5884. IEEE, 2016.
- Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal distributed online prediction using mini-batches. *The Journal of Machine Learning Research*, 2012.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- Sanghamitra Dutta, Gauri Joshi, Soumyadip Ghosh, Parijat Dube, and Priya Nagpurkar. Slow and stale gradients can win the race: Error-runtime trade-offs in distributed sgd. *arXiv preprint arXiv:1803.01113*, 2018.
- Nuwan Ferdinand, Haider Al-Lawati, Stark C Draper, and Matthew Nokleby. Anytime minibatch: Exploiting stragglers in online distributed optimization. *arXiv preprint arXiv:2006.05752*, 2020.
- Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016a.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016b.
- Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2019.
- Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *Advances in Neural Information Processing Systems*, pages 1731–1741, 2017.

- Michael Kamp, Linara Adilova, Joachim Sicking, Fabian Hüger, Peter Schlicht, Tim Wirtz, and Stefan Wrobel. Efficient decentralized deep learning by dynamic model averaging. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 393–409. Springer, 2018.
- Sai Praneeth Karimireddy, Quentin Rebjock, Sebastian U Stich, and Martin Jaggi. Error feedback fixes signsgd and other gradient compression schemes. *arXiv preprint arXiv:1901.09847*, 2019.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Matthew L Leavitt and Ari Morcos. Selectivity considered harmful: evaluating the causal impact of class selectivity in dnns. *arXiv preprint arXiv:2003.01262*, 2020.
- Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, et al. Pytorch distributed: experiences on accelerating data parallel training. *arXiv preprint arXiv:2006.15704*, 2020.
- Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 5330–5340, 2017.
- Xiangru Lian, Wei Zhang, Ce Zhang, and Ji Liu. Asynchronous decentralized parallel stochastic gradient descent. In *International Conference on Machine Learning*, pages 3043–3052. PMLR, 2018.
- Tao Lin, Sebastian U Stich, Kumar Kshitij Patel, and Martin Jaggi. Don’t use large mini-batches, use local sgd. *arXiv preprint arXiv:1808.07217*, 2018.
- Ryan McDonald, Mehryar Mohri, Nathan Silberman, Dan Walker, and Gideon S Mann. Efficient large-scale distributed training of conditional maximum entropy models. In *Advances in neural information processing systems*, pages 1231–1239, 2009.
- Ryan McDonald, Keith Hall, and Gideon Mann. Distributed training strategies for the structured perceptron. In *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*, pages 456–464, 2010.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Daniel Povey, Xiaohui Zhang, and Sanjeev Khudanpur. Parallel training of dnns with natural gradient and parameter averaging. *arXiv preprint arXiv:1410.7455*, 2014.
- Ilija Radosavovic, Justin Johnson, Saining Xie, Wan-Yen Lo, and Piotr Dollár. On network design spaces for visual recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1882–1890, 2019.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- Christopher J Shallue, Jaehoon Lee, Joseph Antognini, Jascha Sohl-Dickstein, Roy Frostig, and George E Dahl. Measuring the effects of data parallelism on neural network training. *arXiv preprint arXiv:1811.03600*, 2018.
- Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. In *Advances in Neural Information Processing Systems*, pages 4424–4434, 2017.
- Sebastian U Stich. Local sgd converges fast and communicates little. *arXiv preprint arXiv:1805.09767*, 2018.
- Hang Su and Haoyu Chen. Experiments on parallel training of deep neural network using model averaging. *arXiv preprint arXiv:1507.01239*, 2015.
- Thijs Vogels, Sai Praneeth Karimireddy, and Martin Jaggi. Powersgd: Practical low-rank gradient compression for distributed optimization. In *Advances in Neural Information Processing Systems*, pages 14259–14268, 2019.
- Jianyu Wang and Gauri Joshi. Adaptive communication strategies to achieve the best error-runtime trade-off in local-update sgd. *arXiv preprint arXiv:1810.08313*, 2018a.
- Jianyu Wang and Gauri Joshi. Cooperative sgd: A unified framework for the design and analysis of communication-efficient sgd algorithms. *arXiv preprint arXiv:1808.07576*, 2018b.
- Jianyu Wang, Vinayak Tantia, Nicolas Ballas, and Michael Rabbat. Slowmo: Improving communication-efficient distributed sgd with slow momentum. In *International Conference on Learning Representations*, 2020.

- Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in neural information processing systems*, pages 1509–1519, 2017.
- Blake Woodworth, Kumar Kshitij Patel, Sebastian U. Stich, Zhen Dai, Brian Bullins, H. Brendan McMahan, Ohad Shamir, and Nathan Srebro. Is local SGD better than minibatch SGD? *arXiv preprint arXiv:2002.07839*, 2020.
- Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.
- Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5693–5700, 2019.
- Jian Zhang, Christopher De Sa, Ioannis Mitliagkas, and Christopher Ré. Parallel sgd: When does averaging help? *arXiv preprint arXiv:1606.07365*, 2016.
- Fan Zhou and Guojing Cong. On the convergence properties of a k -step averaging stochastic gradient descent algorithm for nonconvex optimization. *arXiv preprint arXiv:1708.01012*, 2017.
- Martin Zinkevich, Markus Weimer, Lihong Li, and Alex J Smola. Parallelized stochastic gradient descent. In *Advances in neural information processing systems*, pages 2595–2603, 2010.

Appendix A. Experimental Setup - Additional Details

A.1 LARS

We make use of the LARS optimizer as implemented in NVIDIA APEX library with clipping enabled and a trust coefficient of 0.02 and $\epsilon=10^{-8}$ for all the reported experiments. Despite LARS being proposed with a polynomial learning rate schedule, we use the same step-wise learning rate schedule as with the other experiments to reduce the chance of confounding factors. We found the polynomial schedule to perform slightly better ($< 0.5\%$) than the step-wise but both are able to achieve competitive top 1 validation accuracy results. All other hyperparameters are identical to the setup described for SGD with momentum.

A.2 TinyImageNet

As discussed in the main body of the paper, we perform experiments on TinyImageNet, a spatially downsampled subset of ImageNet. Here we include details of our experimental setup for results involving TinyImageNet. Unlike with other datasets, to our knowledge there is no rigorous analysis like Goyal et al. (2017) for TinyImageNet that contains validated recommendations of hyperparameters. Therefore, we performed a series of hyperparameter searches and arrived at the following choices after cross-validating several settings.

We use a slightly modified version of ResNet-50 and ResNet-101 where the first maxpool layer has been discarded (as in Leavitt and Morcos (2020)) to account for the difference in spatial size of the images (224x224 in ImageNet-1k vs 64x64 in TinyImageNet). Similarly, we halve the size of the learnable filters to account to the reduction in dataset complexity. Models are trained for 60 epochs with learning rate decays at epochs 30 and 50. We use L2 weight decay with $\lambda=10^{-3}$. We found weight decay to be one of the most influential hyperparameters for this task. Apart from these changes, the rest of the training setup uses the same choices as with the ImageNet experiments: learning rate warmup, momentum correction, learning rate scaling and model initialization.

A.3 SlowMo

SlowMo introduces two additional hyperparameters: α , the “slow” learning rate, and β , the “slow” momentum coefficient. For computational reasons, we do not sweep the α and β hyperparameters of SlowMo and set them to recommended defaults of $\alpha=1$ and $\beta=0.5$.

A.4 Platform Details

All methods are implemented using PyTorch 1.6 (Paszke et al., 2019) and we use the ResNet implementations from torchvision 0.7 with CUDA 10.1 and the NCCL communication library. Our experiments run on nodes with eight NVIDIA V100 GPUs each. The nodes communicate over 10 Gb/s Ethernet links. Throughout, a worker refers to a process in a node that makes use of one GPU exclusively. Thus, the number of workers is equivalent to the number of GPUs and, in our setup, is equivalent to eight times the number of nodes.

Appendix B. Additional Experiments

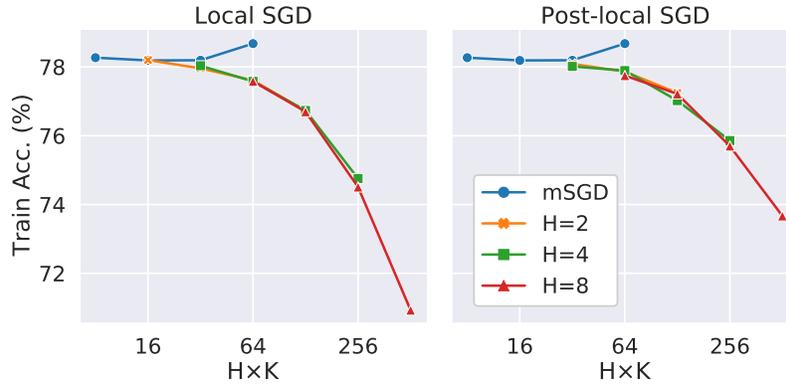


Figure 12: Training Accuracy as a function of the number of local model updates ($H \cdot K$) for local SGD (left) and post-local SGD (right). mSGD values correspond to the minibatch SGD baseline.

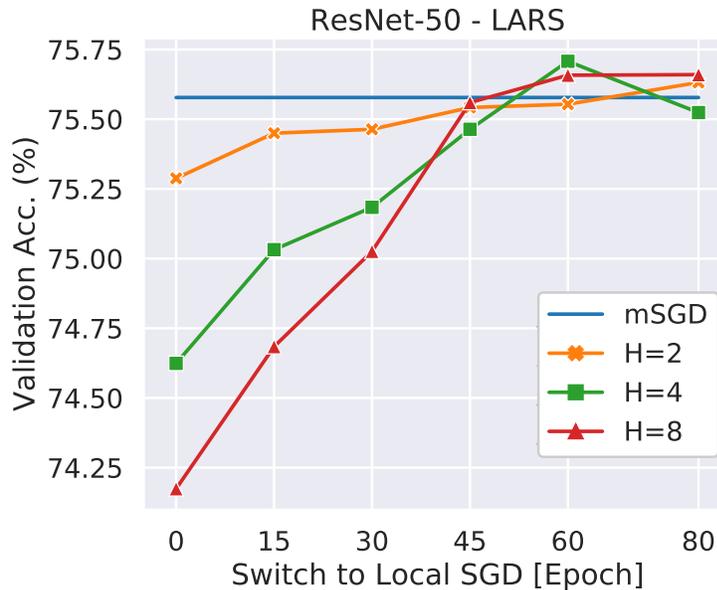


Figure 13: Validation accuracy for post-local SGD with the LARS optimizer as a function of the epoch when the switch to local SGD is done. Results are for $K = 32$ workers and step-wise learning rate schedule. LARS presents a similar trade-off to SGD, with switching points later into training providing better accuracy. LARS seems to have a more linear trade-off without a direct impact.

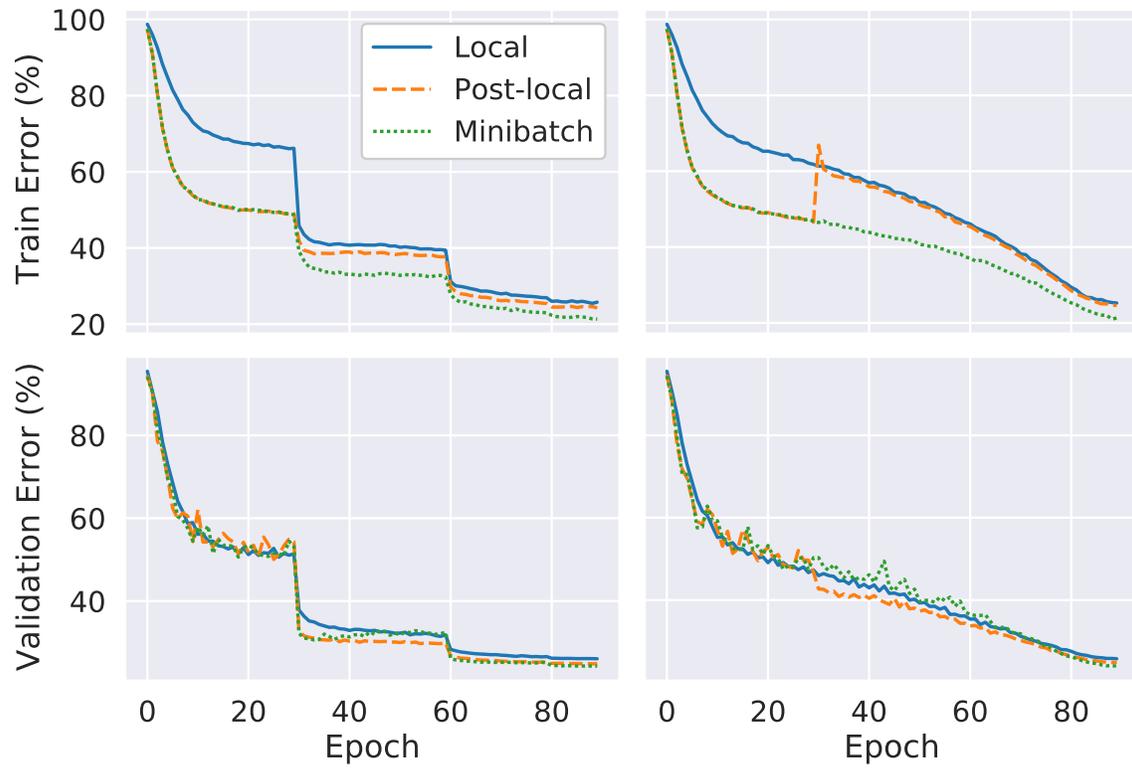


Figure 14: Training and validation error minibatch, local SGD, and post-local SGD for step-wise and half cosine learning rate schedules. The left column is for step-wise and the right column is for half cosine. Results are for 64 workers and 4 local steps for local and post-local SGD. Switching to local SGD has a detrimental effect on the training error but a beneficial effect on the validation error.