

# DISTRIBUTED DEEP LEARNING STRATEGIES FOR AUTOMATIC SPEECH RECOGNITION

Wei Zhang, Xiaodong Cui, Ulrich Finkler, Brian Kingsbury, George Saon, David Kung, Michael Picheny

IBM Research AI

## ABSTRACT


In this paper, we propose and investigate a variety of distributed deep learning strategies for automatic speech recognition (ASR) and evaluate them with a state-of-the-art Long short-term memory (LSTM) acoustic model on the 2000-hour Switchboard (SWB2000), which is one of the most widely used datasets for ASR performance benchmark. We first investigate what are the proper hyper-parameters (e.g., learning rate) to enable the training with sufficiently large batch size without impairing the model accuracy. We then implement various distributed strategies, including Synchronous (SYNC), Asynchronous Decentralized Parallel SGD (ADPSGD) and the hybrid of the two HYBRID, to study their runtime/accuracy trade-off. We show that we can train the LSTM model using ADPSGD in 14 hours with 16 NVIDIA P100 GPUs to reach a 7.6% WER on the Hub5-2000 Switchboard (SWB) test set and a 13.1% WER on the Call-Home (CH) test set. Furthermore, we can train the model using HYBRID in 11.5 hours with 32 NVIDIA V100 GPUs without loss in accuracy.

**Index Terms**— automatic speech recognition, LSTM, deep learning, parallel computing, switchboard.

## 1. INTRODUCTION

Neural networks with deep architectures have been the dominant acoustic modeling approach for automatic speech recognition (ASR) in recent years. They have yielded state-of-the-art performance as compared to previous technologies based on hidden Markov models (HMMs) and Gaussian mixtures [1]. In some tasks, Deep Learning (DL) has achieved near human-level ASR performance [2, 3]. It is commonly agreed that the success of DL for ASR relies on the availability of large amount of training data and high-performance computing. Therefore, distributed DL is not only preferred but also necessary in DL ASR to guarantee fast turnaround time for model training.

Unlike the widely-studied computer vision tasks, such as ImageNet [4, 5, 6, 7, 8], few studies have been published regarding how to accelerate distributed training for ASR tasks on large public dataset (e.g., SWB2000) with the exception of [9, 10]. Compared to computer vision tasks, such as ImageNet, ASR tasks have very distinct behaviors in terms of distributed computing: (1) The state-of-the-art acoustic models conventionally can only be trained with relatively small batch size (e.g. 256) [2], unlike ImageNet where a ResNet model can be trained with a batch size of 8192 or larger [4, 6, 7]. (2) Computation/communication ratio is low in ASR tasks. In Section 3.1, we demonstrate that SWB2000 with a state-of-the-art LSTM is five times more challenging to scale out than a ResNet for ImageNet. Therefore, we need to revisit distributed training strategies other than the standard synchronous SGD training for acceleration. In this paper, we attempt to address the abovementioned



**Fig. 1:** A centralized distributed learning architecture and a decentralized distributed learning architecture

two issues by (1) increasing the batch size for a high-performance LSTM model without impairing model accuracy (2) using the Asynchronous Decentralized Parallel SGD (ADPSGD) [11] approach to reduce the communication cost and remove runtime bottlenecks.

## 2. BACKGROUND

A neural network training algorithm seeks to find a set of parameters  $\theta^*$  that minimizes the discrepancy between the network output  $\hat{Y}$  and the ground truth  $Y$ . This is usually accomplished by defining a differentiable cost function  $C(\hat{Y}, Y)$  and iteratively updating each of the model parameters using some variant of the gradient descent algorithm:

$$E_m = \frac{1}{m} \sum_{s=1}^m C(\hat{Y}_s, Y_s), \quad (1a)$$

$$\nabla \theta(t) = (\partial E_m / \partial \theta)(t), \quad (1b)$$

$$\theta(t+1) = \theta(t) - \alpha(t) \nabla \theta(t) \quad (1c)$$

where  $\theta(t)$  represents the model parameter at iteration  $t$ ,  $\alpha$  is the step size (also known as the learning rate), and  $m$  is the batch size.

Distributed DL systems, the de facto approach to training large DL tasks, typically adopt a Parameter Server (PS) [12] architecture. Figure 1(a) depicts a PS architecture in which each learner calculates gradients and transfers them to the PS. The PS then updates the weights and sends them back to the learners. The timestamp, a scalar counter, of the PS's weights is increased by 1 (i.e. from  $t$  to  $t+1$ ) for each update. *Staleness* [13] is defined as the discrepancy between the timestamp of learners' weights which are used to calculate gradients, and the timestamp of the PS's weights. To achieve convergence that is matching to the single learner system, Synchronous SGD (SSGD) is often used. In SSGD, the PS weight update rule is given in Equation (2)<sup>1</sup>: each learner calculates gradients and receives updated weights in lockstep with the others. The weights used to calculate the gradients are always the same as the one on the PS,

<sup>1</sup>Throughout this paper, we use  $\lambda$  to denote the number of learners.

thus staleness is 0.

$$\begin{aligned}\nabla\theta(t) &= \frac{1}{\lambda} \sum_{i=1}^{\lambda} \nabla\theta_{L_i}, L_i \in L_1, \dots, L_{\lambda} \\ \theta(t+1) &= \theta(t) - \alpha \nabla\theta(t)\end{aligned}\quad (2)$$

The summation operation in Equation (2) is communicative and associative; this is known as the “Reduce” operation in the High Performance Computing field. All-Reduction is the operation that reduces (e.g., sums) all the elements and broadcasts the reduction results to each participant. When the message to be “AllReduced” is large, as in the DL case, the optimal algorithm maximizes the communication bandwidth utilization by breaking a large message to chunks and pipe-lining the reduction operation with message transferring in a ring topology[14]. Such an algorithm can finish the AllReduce operation in  $2 \times M/\text{Bandwidth}$  time, where  $M$  is the size of the message, regardless the number of participants. Many such implementations exist, most notably[15, 16, 5].

One key drawback of SSGD is that one slow learner can slow-down the entire training which is known as the *straggler problem*[17] in distributed computing. To avoid this problem, practitioners proposed Asynchronous SGD (ASGD) which allows each learner to calculate gradients and asynchronously push/pull the gradients/weights to/from PS. The weight update rule in ASGD is given in Equation (3):

$$\begin{aligned}\nabla\theta(t) &= \nabla\theta_{L_i}, L_i \in L_1, \dots, L_{\lambda} \\ \theta(t+1) &= \theta(t) - \alpha \nabla\theta(t)\end{aligned}\quad (3)$$

Staleness in ASGD is proportional to the number of learners in the system[13, 12] and can severely harm convergence[18, 19]. To achieve the best model accuracy, most distributed deep learning tasks adopt SSGD only[10, 4, 20, 21].

To avoid the straggler problem **and** maintain competitive model accuracy, decentralized distributed computing algorithm[22] is proposed, in both synchronous form Decentralized Parallel SGD (DPSGD)[23] and in asynchronous form Asynchronous Decentralized Parallel SGD (ADPSGD)[11]. The architecture of a decentralized SGD system is depicted in Figure 1(b), where each learner  $i$  calculates the gradients, updates its weights, and averages its weights with its neighbor  $j$  in a ring topology. DPSGD/ADPSGD weights update rule is defined in Equation (4).

$$\begin{aligned}\theta(t)'_{L_i} &= \theta(t)_{L_i} - \alpha \nabla\theta(t)_{L_i} \\ \theta(t+1)_{L_i} &= (\theta(t)'_{L_i} + \theta(t)'_{L_j})/2, L_i, L_j \in L_1, \dots, L_{\lambda}\end{aligned}\quad (4)$$


In DPSGD, assuming the pair-wise weight averaging can be executed multiple times after each gradient calculation, all the learners will reach the same weights, thus the staleness can be zero. In ADPSGD, since computation is overlapped with communication, the staleness is 1 at best. ADPSGD has shown excellent runtime and convergence performance on computer vision tasks with CNN type of models (e.g. ResNet and VGG)[11]. In this paper, we demonstrate that ADPSGD also achieves excellent runtime and convergence performance on the SWB2000 speech recognition task with an LSTM model.

### 3. DESIGN AND IMPLEMENTATION

We describe how to increase the batch size to enable efficient distributed computing for SWB2000-LSTM in Section 3.1. We then describe the design and implementation for different distributed learning strategies in Section 3.2.

#### 3.1. Increase Batch Size

A sufficiently large batch size is necessary for enabling efficient distributed DL for two reasons: (1) The larger the batch, the more learners that can be used. (2) Gradient computation is more efficient with



(a) Computation time (hr) per epoch on one P100 GPU and minimum communication bandwidth requirement (GB/s) under different batch size per learner, SWB2000-LSTM

(b) Held-out loss w.r.t epoch for batch size 256 and batch size 2560, SWB2000-LSTM

**Fig. 2:** Computation time / Communication bandwidth requirement and Held-out loss under different batch size, SWB2000-LSTM

a larger batch size. In Figure 2a, blue bars show the computation time per epoch under different batch size for the SWB2000-LSTM<sup>2</sup> task measured on a P100 GPU. It takes 8.58 hrs to finish one epoch under batch size 256 as compared to 18.33 hrs under batch size 32. Furthermore, with a smaller batch size, more frequent communication is required. In Figure 2a, orange bars show the minimum bandwidth requirement to transfer the gradients so that communication time and computation time break even. Batch size 32 per learner requires almost 4X bandwidth (4.98GB/s) as compared with batch size 256 (1.33GB/s). Conventional wisdom on SWB2000-LSTM task is batch size larger than 256 significantly lowers the model accuracy[2]. The hyperparameter setup for the batch size 256 configuration is the learning rate is set to be 0.1, momentum is set as 0.9, and learning rate anneals by  $\frac{1}{\sqrt{2}}$  every epoch from the 11<sup>th</sup> epoch. The training finishes in 16 epochs. Inspired by the work proposed in[4], we are able to increase the batch size from 256 to 2560 without decreasing model accuracy by (1) linearly warming up the base learning rate from 0.1 to 1 in the first 10 epochs and (2) annealing the learning rate by  $\frac{1}{\sqrt{2}}$  from the 11<sup>th</sup> Epoch. Figure 2b plots the held-loss w.r.t epochs for batch size 256 and batch size 2560; they are indistinguishable by epoch 16.

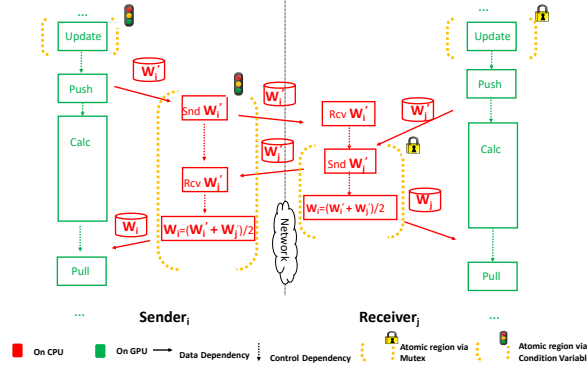

In the ImageNet-ResNet task, a batch size of 32 takes about 0.18 sec to compute on one P100 GPU, whereas the same batch size for SWB2000-LSTM task takes only 0.067sec to compute. Moreover, the ImageNet-ResNet model size is about 100MB, whereas the SWB2000-LSTM model size is 165MB. *The combination of shorter computation time and larger model size makes SWB2000-LSTM 5x more challenging to parallelize than the ImageNet-ResNet task.*

#### 3.2. System Design

Three distributed SGD algorithms are considered and implemented as follows: Synchronous (SYNC), Asynchronous Decentralized Parallel SGD (ADPSGD), and the hybrid of these two algorithms HYBRID. In Figure 3, we sketch our system APIs integration (in bold-font texts) with the underlying DL framework. We assume the underlying DL framework can provide the following functionalities:  $g = \text{calGrad}(W)$ , which calculates the gradients  $g$  based on weights  $W$  and  $W' = \text{apply}(W, g)$ , which applies gradients  $g$  to weights  $W$  and returns updated weights  $W'$ .

<sup>2</sup>We describe the details of the LSTM model used in this paper in Section 4.2

**Fig. 3:** Integration of communication protocol with solver.



**Fig. 4:** System architecture of ADPSGD.

**SYNC:** Summing of the weights and taking their average after every iteration, as shown in Figure 3a, is equivalent to applying weights update by using the averaged gradients. We use the fastest Allreduce implementation available to us (DDL-Allreduce[5]) to implement the SYNC strategy. As we will show in Section 5.2, DDL-Allreduce is 1.2X-3X faster than the open-source MPI-Allreduce implementation in OpenMPI[24].

**ADPSGD:** Figure 4 shows the system architecture of ADPSGD. Assuming  $N$  learners in a system ( $N$  is an even number), we designate learners of odd id  $i$  ( $i \in 1, 3, \dots, N-1$ ) as senders and learners of even id  $j$  ( $j \in 2, 4, \dots, N$ ) as receivers. Bipartition of the communication graph guarantees acyclic communication, thus it is deadlock-free. Each sender communicates with its left and right neighbor in alternate iterations. Each sender process runs two threads: a main thread and a communication thread. The main thread calculates the gradients and applies the weight updates. When a new set of weights  $W'_i$  are generated, the sender signals its communication thread to send  $W'_i$  to its neighbor learner  $j$ , receives learner  $j$ 's weights  $W'_j$ , and then updates its weight  $W_i$  to be the average of  $W'_i$  and  $W'_j$ . Additionally, as required in the proof of [11], the weight matrix of learners need to be doubly stochastic and symmetric which implies the weight update on GPU and the weight averaging on the CPU must not interleave so that the receiver has a consistent view of weights on the sender. We enforce this atomicity via a condition variable[25]. Similarly, each receiver  $j$  runs a main thread and a communication thread. In each iteration, a receiver's main thread calculates gradients and then updates its weights in an atomic region. Meanwhile, the receiver's communication thread waits until the weights are received from its neighbor  $i$ . Then, the communication thread does the following in an atomic region: (i) sends its weights to its engaging neighbor and updates its weights by averaging its weights with the received weights. The atomicity is enforced via a mutex lock[26]. As compared to the implementation in [11], this implementation updates weights on GPU, which runs faster and no gradient information needs be extracted from the underlying solver; the disadvantage of this implementation is that each

	SYNC	HYBRID	ADPSGD
Comm/Compute Overlap	×	✓	✓
Straggler avoidance	×	×	✓
Staleness	0	1	at best 1

**Table 1:** Comparison of the runtime and staleness between different algorithms

sending operation is only triggered when new gradients are calculated and there is no way to exchange weights more frequently even when the network is free. Also, GPU weight updates need to wait if the weights are being changed in the communication thread.

**HYBRID:** Note that the communication threads in ADPSGD essentially runs an Allreduce over learner  $i$  and  $j$ . By replacing the point-to-point message passing with an Allreduce over all learners, we can leverage the optimized fast Allreduce implementation and also minimize the weights discrepancy among different learners. In essence, **push\_weights( $W'$ )** signals the communication thread to conduct an Allreduce which runs concurrently with gradients calculation, and **pull\_weights** simply retrieves the Allreduce-d results from the last push and takes an average. The handshaking between communication thread and main thread is a fast lock-free implementation[27]. When gradient calculation takes longer time than Allreduce, this scheme should completely overlap communication with computation.

Table 1 summarizes the runtime and staleness comparison between different algorithms.

## 4. METHODOLOGY

### 4.1. Software and Hardware


PyTorch 0.5.0 is the underlying DL framework. We use the CUDA 9.2 compiler, the CUDA-aware OpenMPI 3.1.1, and g++ 4.8.5 compiler to build our communication library, which connects with PyTorch via a Python-C interface.

We develop and experiment our systems on a production-run cluster ClusterA, which has 4 servers in total. Each server is equipped with 14-core Intel Xeon E5-2680 v4 2.40GHz processor, 1TB main memory, and 4 P100 GPUs. We also run a 32-GPU experiment on a 4-server high-GPU-density experimental cluster ClusterB. Each ClusterB server is equipped with 18-core Intel Xeon E5-2697 2.3GHz processor, 1TB main memory, and 8 V100 GPUs. Both servers are connected by 100Gbit/s ethernet. On both servers, GPUs and CPUs are connected via PCIe Gen3 bus, which has a 16GB/s peak bandwidth in each direction.

### 4.2. DL Models and Dataset

The acoustic model is an LSTM with 6 bi-directional layers. Each layer contains 1,024 cells (512 cells in each direction). On top of the LSTM layers, there is a linear projection layer with 256 hidden units followed by a softmax output layer with 32,000 units corresponding to context-dependent HMM states. The LSTM is unrolled with 21 frames and trained with non-overlapping feature subsequences of that length. The feature input is a fusion of FMLLR (40-dim), i-Vector (100-dim), and logmel with its delta and double delta (40-dim  $\times 3$ ).

The language model (LM) is rebuilt using publicly available training data, including Switchboard, Fisher, Gigaword, and Broadcast News, and Conversations. Its vocabulary has 85K words and it has 36M 4-grams.



(a) Held-out loss comparison between Baseline, SYNC, ADPSGD, and HYBRID (b) Speed-up comparison between different strategies

Fig. 5: Convergence/Speed-up comparison, 16 GPUs on ClusterA.

## 5. EXPERIMENTAL RESULTS

### 5.1. Convergence Results

	Baseline	SYNC	HYBRID	ADPSGD
SWB	7.5%	7.6%	7.6%	7.6%
CH	13.0%	13.1%	13.1%	13.2%

Table 2: WER comparison between baseline, SYNC, HYBRID, and ADPSGD after training for 16 epochs.

Our LSTM baseline trained on single GPU (batchsize 256) gives a WER of 7.5%/13.0% on the Switchboard/CallHome (SWB/CH) of the NIST Hub5 2000 evaluation test sets after the Cross-Entropy training, which is a competitive baseline. We compare this baseline with SYNC, HYBRID, and ADPSGD in Figure 5a for heldout loss and in Table 2 for WER.

### 5.2. Runtime Results

Figure 5b plots the speed-up for different algorithms up to 16 GPUs on ClusterA. SYNC-OpenMPI is found to be the slowest one. It is also found that ADPSGD achieves the best speed-up (about 11X over 16 GPUs) and finishes the training in 13.98 hours. ADPSGD did not achieve linear speed-up because it requires CPU-based weights averaging and GPU weights update to occur atomically which could be remedied by offloading GPU weights update to CPU. HYBRID does not outperform SYNC-DDL significantly even though computation is long enough to hide the communication. This is because HYBRID asynchronously calls DDL which relies on NVIDIA NCCL library[16] for the intra-server reduction. NCCL heavily competes with training for GPU resources (e.g., stream processor and memory) when used asynchronously.

Table 3 shows the speed-up measured for each algorithm when one of the 16 GPUs slows down. ADPSGD is immune to the straggler problem, whereas the speedup of other algorithms quickly diminishes. Figure 6 shows a snapshot of the number of minibatches processed by each GPU in one epoch when half of ClusterA are shared by other users. ADPSGD automatically balanced the workload per GPU. SYNC and HYBRID would enforce each GPU to process the same number of minibatches in this scenario.

### 5.3. Experiments on 32 GPUs

We ran an experiment on ClusterB with 32 GPUs and batch size 80 per GPU<sup>3</sup>. SYNC-DDL, HYBRID and ADPSGD complete one

<sup>3</sup>It takes 195 hours to finish training SWB2000 task on a V100 GPU, with a batchsize 80.

Slowdown of one GPU	ADPSGD		SYNC-DDL/HYBRID	
	Time/epoch (hr)	Speed-up	Time/epoch (hr)	Speed-up
no slowdown	0.87	10.88	1.09/1.03	8.70/9.26
2X	0.89	10.63	1.67/1.63	5.71/5.83
10X	0.91	10.42	6.24/6.46	1.52/1.47
100X	0.92	10.38	57.73/60.80	0.16/0.16

Table 3: Runtime and speedup when one GPU slows down by 2X-100X, 16 GPUs on ClusterA. ADPSGD is immune to the straggler problem.




Fig. 6: When half of ClusterA is shared to run other tasks, ADPSGD balances the workload for different GPUs for the SWB2000-LSTM task.

epoch in 0.75 hrs (16.25x speed-up), 0.71hrs (17.17x speed-up), and 0.83hrs (14.69x speed-up) respectively. *HYBRID* trains SWB2000 to reach WER 7.6% on SWB and WER 13.1% on CH in 11.5 hrs.

ADPSGD requires more CPU resources than SYNC and HYBRID to conduct weights averaging and passing weights. ClusterB has a less favorable CPU/GPU ratio. Furthermore, 8 GPUs share one PCIe bus on ClusterB and ADPSGD saturates the 16GB/s bandwidth quickly. ADPSGD would run significantly faster if deployed on clusters with higher CPU/GPU ratio, higher main memory bandwidth, and/or advanced GPU-CPU interconnect (e.g., NVLink[28]).

## 6. CONCLUSION AND FUTURE WORK

In this paper, we made the following contributions: (1) we first described the hyper-parameter setup for SWB2000-LSTM speech recognition task using batch size of 2560, which is a sufficiently large batch size that enables efficient distributed training. (2) We implemented and compared different distributed learning algorithms for this task. Our system trains a model to WER 7.6% on SWB and WER 13.1% on CH in less than 12 hours. To the best of our knowledge, this is the fastest system that trains these tasks to this level of accuracy. Our future work includes: (1) Implement wait-free ADPSGD as proposed in [11] to further improve convergence and runtime. (2) Experiment with hardware with better memory bandwidth, CPU/GPU ratio, and CPU-GPU inter-connect. (3) Experiment with different types of speech recognition workloads. (4) Explore other methods to increase batch size and/or use mixed-precision training as in [6, 7].

## 7. RELATED WORK

Distributed DL have been applied to speech recognition[10, 29], computer vision[4], language modeling[20], and machine translation[21] tasks. To reduce the cost of communication, researchers have proposed gradient quantization[9, 30] and gradient compression[31, 8]. All these works adopt a synchronous training method which would become unacceptably slow in a resource-sharing or Cloud environment[32]. Asynchronous SGD, based on the parameter-server architecture, is known to have inferior performance and should be avoided when possible[18, 4, 10, 19]. This work is the first that applies Asynchronous Decentralized Parallel SGD (ADPSGD), which has the theoretical guarantee to converge at the same rate as SGD[11], to the challenging SWB2000-LSTM task to achieve state-of-the-art model accuracy in a record time.

## 8. REFERENCES

- [1] G Hinton, L Deng, D Yu, G Dahl, A Mohamed, N Jaitly, A Senior, V Vanhoucke, P Nguyen, T. N Sainath, and B Kingsbury, “Deep neural networks for acoustic modeling in speech recognition,” *IEEE Signal Processing Magazine*, pp. 82–97, November 2012.
- [2] G Saon, G Kurata, T Sercu, K Audhkhasi, S Thomas, D Dimitriadis, X Cui, B Ramabhadran, M Picheny, L.-L. Lim, B Roomi, and P Hall, “English conversational telephone speech recognition by humans and machines,” in *Interspeech*, 2017.
- [3] W Xiong, J Droppo, X Huang, F Seide, M. L. Seltzer, A Stolcke, D Yu, and G Zweig, “Toward human parity in conversational speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2410–2423, Dec 2017.
- [4] P Goyal, P Dollár, R. B Girshick, P Noordhuis, L Wesolowski, A Kyrola, A Tulloch, Y Jia, and K He, “Accurate, large minibatch SGD: training imagenet in 1 hour,” *CoRR*, vol. abs/1706.02677, 2017.
- [5] M Cho, U Finkler, S Kumar, D. S Kung, V Saxena, and D Sreedhar, “Powerai DDL,” *CoRR*, vol. abs/1708.02188, 2017.
- [6] Y You, I Gitman, and B Ginsburg, “Scaling SGD batch size to 32k for imagenet training,” *CoRR*, vol. abs/1708.03888, 2017.
- [7] X Jia, S Song, W He, Y Wang, H Rong, F Zhou, L Xie, Z Guo, Y Yang, L Yu, T Chen, G Hu, S Shi, and X Chu, “Highly scalable deep learning training system with mixed-precision: Training imagenet in four minutes,” *CoRR*, vol. abs/1807.11205, 2018.
- [8] W Wen, C Xu, F Yan, C Wu, Y Wang, Y Chen, and H Li, “Terngrad: Ternary gradients to reduce communication in distributed deep learning,” in *NIPS’2017*, pp. 1509–1519. Curran Associates, Inc., 2017.
- [9] F Seide, H Fu, J Droppo, G Li, and D Yu, “1-bit stochastic gradient descent and application to data-parallel distributed training of speech dnns,” in *Interspeech 2014*, September 2014.
- [10] D Amodei(et.al.), “Deep speech 2 : End-to-end speech recognition in english and mandarin,” in *ICML’16*. 2016, pp. 173–182, PMLR.
- [11] X Lian, W Zhang, C Zhang, and J Liu, “Asynchronous decentralized parallel stochastic gradient descent,” in *ICML*, 2018.
- [12] J Dean, G. S Corrado, R Monga, K Chen, M Devin, Q. V Le, M. Z Mao, M Ranzato, A Senior, P Tucker, K Yang, and A. Y Ng, “Large scale distributed deep networks,” in *NIPS*, 2012.
- [13] W Zhang, S Gupta, X Lian, and J Liu, “Staleness-aware async-sgd for distributed deep learning,” in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 2016, pp. 2350–2356.
- [14] P Patarasuk and X Yuan, “Bandwidth optimal all-reduce algorithms for clusters of workstations,” *J. Parallel Distrib. Comput.*, vol. 69, pp. 117–124, 2009.
- [15] Baidu, *Effectively Scaling Deep Learning Frameworks*, Available at <https://github.com/baidu-research/baidu-allreduce>.
- [16] Nvidia, *NCCL: Optimized primitives for collective multi-GPU communication*, Available at <https://github.com/NVIDIA/nccl>.
- [17] J Dean and S Ghemawat, “Mapreduce: Simplified data processing on large clusters,” *Commun. ACM*, vol. 51, no. 1, pp. 107–113, Jan. 2008.
- [18] J Chen, R Monga, S Bengio, and R Jozefowicz, “Revisiting distributed synchronous sgd,” in *International Conference on Learning Representations Workshop Track*, 2016.
- [19] W Zhang, S Gupta, and F Wang, “Model accuracy and runtime tradeoff in distributed deep learning: A systematic study,” in *IEEE International Conference on Data Mining*, 2016.
- [20] R Puri, R Kirby, N Yakovenko, and B Catanzaro, “Large scale language modeling: Converging on 40gb of text in four hours,” *CoRR*, vol. abs/1808.01371, 2018.
- [21] M Ott, S Edunov, D Grangier, and M Auli, “Scaling neural machine translation,” *EMNLP 2018 THIRD CONFERENCE ON MACHINE TRANSLATION*, vol. abs/1806.00187, 2018.
- [22] R Nair and S Gupta, “Wildfire: Approximate synchronization of parameters in distributed deep learning,” *IBM Journal of Research and Development*, vol. 61, no. 4/5, pp. 7:1–7:9, July 2017.
- [23] X Lian, C Zhang, H Zhang, C.-J Hsieh, W Zhang, and J Liu, “Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent,” in *NIPS*, 2017.
- [24] E Gabriel(et.al.), “Open MPI: Goals, concept, and design of a next generation MPI implementation,” in *Proceedings, 11th European PVM/MPI Users’ Group Meeting*, 2004.
- [25] B. W Lampson and D. D Redell, “Experience with processes and monitors in mesa,” *Commun. ACM*, vol. 23, no. 2, pp. 105–117, Feb. 1980.
- [26] R. H Arpaci-Dusseau and A. C Arpaci-Dusseau, *Operating Systems: Three Easy Pieces*, Arpaci-Dusseau Books, 0.91 edition, May 2015.
- [27] U Finkler, H Franke, and D. S Kung, “DYCE: A resilient shared memory paradigm for heterogeneous distributed systems without memory coherence,” in *Proceedings of the Computing Frontiers Conference, CF’17, Siena, Italy, May 15-17, 2017*, 2017, pp. 17–26.
- [28] Nvidia, *NVLINK FABRIC*, Available at <https://www.nvidia.com/en-us/data-center/nvlink/>.
- [29] K Chen and Q Huo, “Scalable training of deep learning machines by incremental block training with intra-block parallel optimization and blockwise model-update filtering,” in *ICASSP’2016*, March 2016.
- [30] N Wang, J Choi, D Brand, C.-Y Chen, and K Gopalakrishnan, “Training deep neural networks with 8-bit floating point numbers,” in *NIPS*, 2018.
- [31] C Chen, J Choi, D Brand, A Agrawal, W Zhang, and K Gopalakrishnan, “Adacomp : Adaptive residual gradient compression for data-parallel distributed training,” *AAAI*, 2017.
- [32] W Zhang, M Feng, Y Zheng, Y Ren, Y Wang, J Liu, P Liu, B Xiang, L Zhang, B Zhou, and F Wang, “Gadei: On scale-up training as a service for deep learning,” *The IEEE International Conference on Data Mining series(ICDM’2017)*, 2017.