

Foundations and Trends® in Databases

Distributed Learning Systems with First-Order Methods

An Introduction

Suggested Citation: Ji Liu and Ce Zhang (2020), "Distributed Learning Systems with First-Order Methods", Foundations and Trends® in Databases: Vol. 9, No. 1, pp 1–100. DOI: 10.1561/1900000062.

Ji Liu
University of Rochester and
Kuaishou Inc., USA
ji.liu.uwisc@gmail.com

Ce Zhang
ETH Zurich, Switzerland
ce.zhang@inf.ethz.ch

This article may be used only for the purpose of research, teaching,
and/or private study. Commercial use or systematic downloading
(by robots or other automatic processes) is prohibited without ex-
plicit Publisher approval.

now
the essence of knowledge
Boston — Delft

Contents

1	Introduction	4
1.1	Gradient Descent	5
1.2	Stochastic Gradient Descent	11
1.3	A Simplified Distributed Communication Model	18
2	Distributed Stochastic Gradient Descent	31
2.1	A Simplified Performance Model for Distributed Synchronous Data-Parallel SGD	31
2.2	Theoretical Analysis	36
2.3	Caveats	36
3	System Relaxation 1: Lossy Communication Compression	39
3.1	System Implementation	40
3.2	Theoretical Analysis for CSGD	45
3.3	Error Compensated Stochastic Gradient Descent (EC-SGD) for Arbitrary Compression Strategies	48
3.4	Theoretical Analysis for EC-SGD	50
4	System Relaxation 2: Asynchronous Training	56
4.1	System Implementation	57
4.2	Theoretical Analysis	60

5 System Relaxation 3: Decentralized Communication	67
5.1 System Implementation	68
5.2 Theoretical Analysis	70
6 Further Reading	83
6.1 Communication and Data Compression	84
6.2 Decentralization and Approximate Averaging	84
6.3 Asynchronous Communication	85
6.4 Optimizing for Communication Rounds	85
6.5 System Optimization and Automatic Tradeoff Management	85
6.6 Other Topics	86
References	88

Distributed Learning Systems with First-Order Methods

Ji Liu¹ and Ce Zhang²

¹*University of Rochester and Kuaishou Inc., USA;*
ji.liu.uwisc@gmail.com

²*ETH Zurich, Switzerland; ce.zhang@inf.ethz.ch*

ABSTRACT

Scalable and efficient distributed learning is one of the main driving forces behind the recent rapid advancement of machine learning and artificial intelligence. One prominent feature of this topic is that recent progress has been made by researchers in *two* communities: (1) *the system community* such as database, data management, and distributed systems, and (2) *the machine learning and mathematical optimization community*. The interaction and knowledge sharing between these two communities has led to the rapid development of new distributed learning systems and theory.

In this monograph, we hope to provide a brief introduction of some distributed learning techniques that have recently been developed, namely *lossy communication compression* (e.g., quantization and sparsification), *asynchronous communication*, and *decentralized communication*. One special focus in this monograph is on making sure that it can be easily understood by researchers in *both* communities — on the system side, we rely on a simplified system model hiding many system details that are not necessary for the intuition behind the system speedups; while, on the theory side, we

rely on minimal assumptions and significantly simplify the proof of some recent work to achieve comparable results.

Notations and Definitions

Throughout this monograph, we make the following definitions.

- All vectors are assumed to be column vectors by default.
- α , β , and γ usually denote constants.
- Bold lowercase letters usually denote vectors, such as \mathbf{x} , \mathbf{y} , and \mathbf{v} .
- $\langle \mathbf{x}, \mathbf{y} \rangle$ denotes the dot product between two vectors \mathbf{x} and \mathbf{y} .
- Capital letters usually denote matrices, such as W .
- \lesssim means “small and equal to up to a constant factor”, for example, $a_t \lesssim b_t$ means that there exists a constant $\alpha > 0$ independent of t such that $a_t \leq \alpha b_t$.
- $\mathbf{1}$ denotes a vector with 1 at everywhere and its dimension depends on the context.
- $f'(\cdot)$ denotes the gradient or differential of the function $f(\cdot)$.
- $[M] := \{1, 2, \dots, M\}$ denotes a set containing integers from 1 to M .

1

Introduction

Real-world distributed learning systems, especially those relying on first-order methods, are often constructed in two “phases” — first comes the textbook (stochastic) gradient descent (**SGD**) algorithm, and then certain aspects of the system design are “relaxed” to remove the system bottleneck, be that communication bandwidth, latency, synchronization cost, etc. Throughout this work, we will describe multiple popular ways of system relaxation developed in recent years and analyze their system behaviors and theoretical guarantees.

In this section, we provide the background for both the theory and the system. On the theory side, we describe the intuition and theoretical properties of standard gradient descent (**GD**) and stochastic gradient descent (**SGD**) algorithms (we refer the reader to Bottou *et al.*, 2016 for more details). On the system side, we introduce a simplified performance model that hides many details but is just sophisticated enough for us to reason about the performance impact of the different system relaxation techniques that we will introduce in the later sections.

Summary of Results. In this monograph, we focus on three different system relaxation techniques, namely *lossy communication compression*,

Table 1.1: Summary of results covered in this monograph. For distributed settings, we assume that there are N workers and the latency and the bandwidth of the network are α and β , respectively. The lossy compression scheme has a compression ratio of $\eta (< 1)$ (which introduces additional variance σ' to the gradient estimator) and the decentralized communication scheme uses a communication graph g of degree $\deg(G)$. ς measures the data variation among workers in the decentralized scenario — $\varsigma = 0$ if all workers have the same dataset. We assume the simplified communication model and communication pattern as described in Subsection 1.3

Algorithm	System Optimization	# Iterations to ϵ	Communication Cost
GD	/	$O\left(\frac{1}{\epsilon}\right)$	N/A
SGD	/	$O\left(\frac{1}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right)$	N/A
mb-SGD	Distributed Baseline	$O\left(\frac{1}{\epsilon} + \frac{\sigma^2}{N\epsilon^2}\right)$	$O(N\alpha + \beta)$
CSGD	Compression	$O\left(\frac{1}{\epsilon} + \frac{\sigma^2}{N\epsilon^2} + \frac{\sigma'^2}{\epsilon^2}\right)$	$O(N\alpha + \beta\eta)$
EC-SGD	Compression	$O\left(\frac{1}{\epsilon} + \frac{\sigma^2}{N\epsilon^2} + \frac{\sigma'}{\epsilon^{2/3}}\right)$	$O(N\alpha + \beta\eta)$
ASGD	Asynchronization	$O\left(\frac{N}{\epsilon} + \frac{\sigma^2}{N\epsilon^2}\right)$	$O(N\alpha + \beta)$
DSGD	Decentralization	$O\left(\frac{1}{\epsilon} + \frac{\sigma^2}{N\epsilon^2} + \frac{\varsigma}{(1-\rho)\epsilon^{2/3}}\right)$	$O(\deg(G)(\alpha + \beta))$

asynchronous communication, and *decentralized communication*. For each system relaxation technique, we study their convergence behavior (i.e., # iterations we need to achieve ϵ precision) and the communication cost per iteration. Table 1.1 summarizes the results we will cover in this monograph.

1.1 Gradient Descent

Let us consider the generic machine learning objective that can be summarized by the following form

$$\min_{\mathbf{x} \in \mathbb{R}^d} \quad \left\{ f(\mathbf{x}) := \frac{1}{M} \sum_{m=1}^M F_m(\mathbf{x}) \right\}. \quad (1.1)$$

Let $f^* := \min_{\mathbf{x}} f(\mathbf{x})$ and assume that it exists by default. Each F_m corresponds to a *data sample* in the context of machine learning.

The gradient descent (GD) can be described as

$$(\text{GD}) \quad \mathbf{x}_{t+1} = \mathbf{x}_t - \gamma f'(\mathbf{x}_t) \quad (1.2)$$

where t is the iteration index and $f'(\mathbf{x}_t)$ is the gradient of f at \mathbf{x}_t .

1.1.1 Intuitions

We provide two intuitions about the gradient descent GD algorithm to indicate why it will work.

Steepest Descent Direction. The gradient (or a differential) of a function is the steepest direction to increase the function value given an infinitely small step, which can be seen from the property of the function gradient $\forall \|\mathbf{v}\| = 1$

$$\langle f'(\mathbf{x}), \mathbf{v} \rangle = f'_\mathbf{v}(\mathbf{x}) := \lim_{\delta \rightarrow 0} \frac{f(\mathbf{x} + \mathbf{v}\delta) - f(\mathbf{x})}{\delta}.$$

$f'_\mathbf{v}(\mathbf{x})$ is the directional gradient, which indicates how much increment there is on function value along the direction \mathbf{v} by a tiny unit step. To find the steepest unit descent direction is to maximize

$$\max_{\|\mathbf{v}\|=1} f'_\mathbf{v}(\mathbf{x}).$$

Since $f'_\mathbf{v}(\mathbf{x}) = \langle f'(\mathbf{x}), \mathbf{v} \rangle$, it is easy to verify that the steepest direction is $\mathbf{v}^* = \frac{f'(\mathbf{x})}{\|f'(\mathbf{x})\|}$. Note that our goal is to minimize the function value. Therefore, GD is a natural idea via moving the model \mathbf{x}_t along the steepest “descent” direction $-f'(\mathbf{x}_t)$.

Minimizing a Model Function. Another perspective from which to view gradient descent is based on the model function. Since the original objective function $f(\mathbf{x})$ is usually very complicated, it is very hard to minimize the objective function directly. A straightforward idea is to construct a model function to locally approximate (at \mathbf{x}_t) the original objective in each iteration. The model function needs to be simple and to approximate the original function well enough. Therefore, the most natural idea is to choose a quadratic function (that is usually simple

to solve)

$$M_{\mathbf{x}_t, \gamma}(\mathbf{x}) := f(\mathbf{x}_t) + \langle f'(\mathbf{x}_t), \mathbf{x} - \mathbf{x}_t \rangle + \frac{1}{2\gamma} \|\mathbf{x} - \mathbf{x}_t\|^2.$$

This model function is a good approximation in the sense that

- $f(\mathbf{x}_t) = M_{\mathbf{x}_t, \gamma}(\mathbf{x}_t)$
- $f'(\mathbf{x}_t) = M'_{\mathbf{x}_t, \gamma}(\mathbf{x}_t)$
- $f(\cdot) \leq M_{\mathbf{x}_t, \gamma}(\cdot)$ if the learning rate γ is sufficiently small.

For the first two, it is easy to understand why they are important. The last one is important to the convergence, which will be seen soon. Figure 1.1 illustrates the geometry of the model function. One can verify that the GD algorithm is nothing but iteratively update the optimization variable \mathbf{x} via minimizing the model function at the current point \mathbf{x}_t :

$$\begin{aligned} \mathbf{x}_{t+1} &= \operatorname{argmin}_{\mathbf{x}} M_{\mathbf{x}_t, \gamma}(\mathbf{x}) \\ &= \operatorname{argmin}_{\mathbf{x}} \frac{1}{2\gamma} \|\mathbf{x} - (\mathbf{x}_t - \gamma f'(\mathbf{x}_t))\|^2 + \text{constant} \\ &= \mathbf{x}_t - \gamma f'(\mathbf{x}_t). \end{aligned}$$

The convergence of GD can also be revealed by this intuition — \mathbf{x}_{t+1} always improves \mathbf{x}_t unless the gradient is zero

$$f(\mathbf{x}_{t+1}) \leq M_{\mathbf{x}_t, \gamma}(\mathbf{x}_{t+1}) \leq M_{\mathbf{x}_t, \gamma}(\mathbf{x}_t) = f(\mathbf{x}_t),$$

where $f(\mathbf{x}_{t=1}) = f(\mathbf{x}_t)$ holds if and only if $f'(\mathbf{x}_t) = 0$.

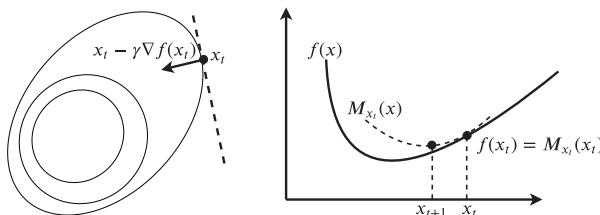


Figure 1.1: (Left) Illustration of gradient and steepest descent direction. (Right) Illustration of model function.

1.1.2 Convergence Rate

From the intuition of GD, the convergence of GD is automatically implied. This subsection provides the convergence *rate* via rigorous analysis. To show the convergence rate, let us first make some commonly used assumptions in the following.

Assumption 1. We assume:

- **(Smoothness)** All functions $F_m(\cdot)$'s are differentiable.
- **(L -Lipschitz gradient)** The objective function is assumed to have a Lipschitz gradient, that is, there exists a constant L satisfying $\forall \mathbf{x}, \forall \mathbf{y}$

$$\|f'(\mathbf{x}) - f'(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\| \quad (1.3)$$

$$f(\mathbf{y}) - f(\mathbf{x}) \leq \langle f'(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|^2. \quad (1.4)$$

The smoothness assumption on $F_m(\cdot)$'s implies that the overall objective function $f(\cdot)$ is differentiable or smooth too. The assumption (1.4) can be deduced from (1.3), and we refer readers to the textbook by Boyd and Vandenberghe (2004) or their course link.¹ The Lipschitz gradient assumption essentially assumes that the curvature of the objective function is bounded by L . We make the assumption of (1.4) just for convenience of use later.

We apply the Lipschitz gradient assumption and immediately obtain the following golden inequality:

$$\begin{aligned} f(\mathbf{x}_{t+1}) - f(\mathbf{x}_t) &\leq \langle f'(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \frac{L}{2}\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &= -\gamma\|f'(\mathbf{x}_t)\|^2 + \frac{\gamma^2 L}{2}\|f'(\mathbf{x}_t)\|^2 \\ &= -\gamma \left(1 - \frac{\gamma L}{2}\right) \|f'(\mathbf{x}_t)\|^2. \end{aligned} \quad (1.5)$$

We can see that as long as the learning rate γ is small enough such that $1 - \gamma L/2 > 0$, $f(\mathbf{x}_{t+1})$ can improve $f(\mathbf{x}_t)$. Therefore, the learning rate

¹<http://www.seas.ucla.edu/~vandenbe/236C/lectures/gradient.pdf>.

cannot be too large to guarantee the progress in each step. However, it is also a bad idea if the learning rate is too small, since the progress is proportional to $\gamma(1 - \gamma L/2)$. The optimal learning rate can be obtained by simply maximizing

$$\gamma(1 - \gamma L/2)$$

over γ , which gives the optimal learning rate for the gradient descent method as $\gamma^* = 1/L$. Substituting $\gamma = \gamma^*$ into (1.5) yields

$$f(\mathbf{x}_{t+1}) - f(\mathbf{x}_t) \leq -\frac{1}{2L} \|f'(\mathbf{x}_t)\|^2$$

or equivalently

$$f(\mathbf{x}_t) - f(\mathbf{x}_{t+1}) \geq \frac{1}{2L} \|f'(\mathbf{x}_t)\|^2. \quad (1.6)$$

Summarizing Eq. (1.6) over t from $t = 1$ to $t = T$ yields

$$\begin{aligned} \frac{1}{2L} \sum_{t=1}^T \|f'(\mathbf{x}_t)\|^2 &\leq \sum_{t=1}^T (f(\mathbf{x}_t) - f(\mathbf{x}_{t+1})) \\ &= f(\mathbf{x}_1) - f(\mathbf{x}_{t+1}) \\ &\leq f(\mathbf{x}_1) - f^*. \end{aligned}$$

Rearranging the inequality yields the following convergence rate for gradient descent.

Theorem 1.1. Under Assumption 1, the gradient descent method admits the following convergence rate

$$\frac{1}{T} \sum_{t=1}^T \|f'(\mathbf{x}_t)\|^2 \lesssim \frac{L}{T} \quad (1.7)$$

by choosing the learning rate $\gamma = \frac{1}{L}$. Here, we treat $f(\mathbf{x}_1) - f^*$ as a constant.

This result indicates that the averaged gradient norm converges in the rate of $1/T$. It is worth noting that, unlike the convex case, we are unable to use the commonly used criterion $f(\mathbf{x}_t) - f^*$ to evaluate the convergence (efficiency). That is to say, the algorithm guarantees

the convergence only to a stationary point ($\|f'(\mathbf{x}_t)\|^2 \rightarrow 0$) because of the nonconvexity. The connection between two criteria $f(\mathbf{x}_t) - f^*$ and $\|f'(\mathbf{x}_t)\|^2$ can be seen from

$$\frac{1}{L} \|f'(\mathbf{x}_t)\|^2 \leq f(\mathbf{x}_t) - f^*.$$

The proof can be found in the standard textbook or the course link.²

There are two major disadvantages for the GD method:

- The computational complexity and system overhead can be too high in each iteration to compute a single gradient.
- For nonconvex objectives, the gradient descent often sticks on a bad (shallow) local optimum.

1.1.3 Iteration/Query/Computation Complexity

The convergence rate is the key to analyzing the overall complexity. People usually consider three types of overall complexity: (1) iteration complexity, (2) query complexity, and (3) computation complexity. To evaluate the overall complexity to solve the optimization problem in (1.1), we need first to specify a precision of our solution, since in practice it is difficult (also not really necessary) to exactly solve the optimization problem. In particular, in our case the overall complexity must take into account how many iterations/queries/computations are required to ensure the average gradient norm $\frac{1}{T} \sum_{t=1}^T \|f'(\mathbf{x}_t)\|^2 \leq \epsilon$.

Iteration Complexity. From Theorem 1.1, it is straightforward to verify that the iteration complexity is

$$O\left(\frac{L}{\epsilon}\right). \quad (1.8)$$

Query Complexity. Here “query” refers to the number of queries of the data samples. GD needs to query all M samples in each iteration. Therefore, the query complexity can be computed from the iteration

²<https://www.cs.rochester.edu/~jliu/CSC-576/class-note-6.pdf>.

complexity by multiplying the number of queries in each iteration

$$O\left(\frac{LM}{\epsilon}\right). \quad (1.9)$$

Computation Complexity. Similarly, the computation complexity can be computed from the query complexity by multiplying the complexity of computing one sample gradient $F'_m(\mathbf{x})$. The typical complexity of computing one sample gradient is proportional to the dimension of the variable, which is d in our notation. To see the reason, let us imagine a naive linear regression with $F_m := \frac{1}{2}(\mathbf{a}_m^\top \mathbf{x} - b)^2$ and a sample gradient of $f'_m(\mathbf{x}) := \mathbf{a}_m(\mathbf{a}_m^\top \mathbf{x} - b)$. Therefore, the computation complexity of **GD** is

$$O\left(\frac{LMD}{\epsilon}\right).$$

It is worth pointing out that the computation complexity is usually proportional to the query complexity (no matter for what kinds of objective) if we consider and compare only first-order (or sample-gradient-based) methods. Therefore, in the remainder of this monograph, we compare only the query complexity and the iteration complexity.

1.2 Stochastic Gradient Descent

One disadvantage of **GD** is that it requires one to query all samples in an iteration, which could be overly expensive. To overcome this shortcoming, the stochastic gradient method **SGD** is widely used in machine learning training. Instead of computing a full gradient in each iteration, it is usual to compute only the gradient on a batch (or minibatch) of sampled data. In particular, people randomly sample an $m_t \in [M]$ independently each time and update the model by

$$(\text{SGD}) \quad \mathbf{x}_{t+1} = \mathbf{x}_t - \gamma F'_{m_t}(\mathbf{x}_t), \quad (1.10)$$

where $m_t \in [M]$ denotes the index randomly selected at the t th iteration. $F'_m(\mathbf{x})$ (or $F'_{m_t}(\mathbf{x}_t)$) is called the stochastic gradient (at the t th iteration). We use $\mathbf{g}(\cdot) := F'_m(\cdot)$ (or $\mathbf{g}_t(\cdot) := F'_{m_t}(\cdot)$) to denote the stochastic gradient (or at the t th iteration) for short. An important property for the stochastic gradient is that its expectation is equal to the true

gradient, that is,

$$\mathbb{E}[\mathbf{g}(\mathbf{x})] = \mathbb{E}_m[F'_m(\mathbf{x})] = f'(\mathbf{x}) \quad \forall \mathbf{x}.$$

An immediate advantage of SGD is that the computational complexity reduces to $O(d)$ per iteration. It is worth pointing out that the SGD algorithm is NOT a descent algorithm³ due to the randomness.

1.2.1 Convergence Rate

The next questions are whether it converges and, if it does, how quickly. We first make a typical assumption.

Assumption 2. We make the following assumption:

- (**Unbiased gradient**) The stochastic gradient is unbiased, that is,

$$\mathbb{E}_m[F'_m(\mathbf{x})] = f'(\mathbf{x}) \quad \forall \mathbf{x}.$$

- (**Bounded stochastic variance**) The stochastic gradient is with bounded variance, that is, there exists a constant σ satisfying

$$\mathbb{E}_m[\|F'_m(\mathbf{x}) - f'(\mathbf{x})\|^2] \leq \sigma^2 \quad \forall \mathbf{x}.$$

We first apply the Lipschitzian gradient property in Assumption 1:

$$\begin{aligned} f(\mathbf{x}_{t+1}) - f(\mathbf{x}_t) &\leq \langle f'(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \frac{L}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &= -\gamma \langle f'(\mathbf{x}_t), \mathbf{g}_t(\mathbf{x}_t) \rangle + \frac{L\gamma^2}{2} \|\mathbf{g}_t(\mathbf{x}_t)\|^2. \end{aligned} \quad (1.11)$$

Note two important properties:

- $\mathbb{E}[\langle f'(\mathbf{x}_t), \mathbf{g}_t(\mathbf{x}_t) \rangle] = \langle f'(\mathbf{x}_t), \mathbb{E}[\mathbf{g}_t(\mathbf{x}_t)] \rangle = \|f'(\mathbf{x}_t)\|^2$
- $\mathbb{E}[\|\mathbf{g}_t(\mathbf{x}_t)\|^2] = \|f'(\mathbf{x}_t)\|^2 + \mathbb{E}[\|\mathbf{g}_t(\mathbf{x}_t) - f'(\mathbf{x}_t)\|^2] \leq \|f'(\mathbf{x}_t)\|^2 + \sigma^2$,

³A descent algorithm means $f(\mathbf{x}_{t+1}) \leq f(\mathbf{x}_t)$, that is, \mathbf{x}_{t+1} is not always worse than \mathbf{x}_t for any iterate t .

where the second property uses the property of variance, that is, any random variable vector ξ satisfies

$$\mathbb{E}[\|\xi\|^2] = \|\mathbb{E}[\xi]\|^2 + \mathbb{E}[\|\xi - \mathbb{E}[\xi]\|^2]. \quad (1.12)$$

Apply these two properties to (1.11) and take expectation on both sides:

$$\begin{aligned} & \mathbb{E}[f(\mathbf{x}_{t+1})] - \mathbb{E}[f(\mathbf{x}_t)] \\ & \leq -\gamma \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] + \frac{L\gamma^2}{2} (\mathbb{E}[\|f'(\mathbf{x}_t)\|^2] + \sigma^2) \end{aligned} \quad (1.13)$$

$$\leq -\gamma \left(1 - \frac{\gamma L}{2}\right) \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] + \frac{\gamma^2}{2} L \sigma^2. \quad (1.14)$$

From (1.13), we can see that SGD does not guarantee “descent” in each iteration, unlike GD, but it does guarantee “descent” in the expectation sense in each iteration as long as γ is small enough and $\|f'(\mathbf{x}_t)\|^2 > 0$. This is because the first term in (1.13) is in the order of $O(\gamma)$ while the second term is in the order of $O(\gamma^2)$.

Next we summarize (1.13) from $t = 1$ to $t = T$ and obtain

$$\mathbb{E}[f(\mathbf{x}_{T+1})] - f(\mathbf{x}_1) \leq -\gamma \left(1 - \frac{\gamma L}{2}\right) \sum_{t=1}^T \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] + \frac{\gamma^2}{2} TL \sigma^2. \quad (1.15)$$

We choose the learning rate $\gamma = \frac{1}{L + \sigma\sqrt{TL}}$ which implies that $(1 - \gamma L/2) > 1/2$. It follows

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] \\ & \lesssim \frac{f(\mathbf{x}_1) - \mathbb{E}[f(\mathbf{x}_{T+1})]}{T\gamma} + \gamma L \sigma^2 \\ & \lesssim \frac{f(\mathbf{x}_1) - f^\star}{T\gamma} + \gamma L \sigma^2 \\ & \lesssim \frac{(f(\mathbf{x}_1) - f^\star)L}{T} + \frac{(f(\mathbf{x}_1) - f^\star)\sqrt{L}\sigma}{\sqrt{T}}. \end{aligned}$$

Therefore the convergence rate of SGD can be summarized into the following theorem.

Theorem 1.2. Under Assumptions 1 and 2, the SGD method admits the following convergence rate

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] \lesssim \frac{L}{T} + \frac{\sqrt{L}\sigma}{\sqrt{T}}$$

by choosing the learning rate $\gamma = \frac{1}{L+\sigma\sqrt{TL}}$. Here we treat $f(\mathbf{x}_1) - f^*$ as a constant.

We highlight the following observations from Theorem 1.2.

- (**Consistent with GD**) If $\sigma = 0$, the SGD algorithm reduces to GD and the convergence rate becomes $L(f(\mathbf{x}_1) - f^*)/L$, which is consistent with the convergence rate for GD proven in Theorem 1.1.
- (**Asymptotic convergence rate**) The convergence rate of SGD achieves $O(1/\sqrt{T})$.

1.2.2 Iteration/Query Complexity

Using a similar analysis as Subsection 1.1.3, we can obtain the iteration complexity of SGD, which is also the query complexity (since there is only one query per one sample gradient)

$$O\left(\frac{L}{\epsilon} + \frac{L\sigma^2}{\epsilon^2}\right).$$

It is worse than GD in terms of the iteration complexity in (1.8), which is not a surprising result. The comparison of query complexity makes more sense since it is more related to the physical running time or the computation complexity. From the detailed comparison in Table 1.2, we can see that

- SGD is superior to GD, if $\frac{\sigma^2}{M} \ll \epsilon$;
- SGD is inferior to GD, if $\frac{\sigma^2}{M} \gg \epsilon$.

Table 1.2: Complexity comparison among GD, SGD, and mb-SGD

Algorithms	Iteration Complexity	Query Complexity
GD	$O\left(\frac{L}{\epsilon}\right)$	$O\left(\frac{ML}{\epsilon}\right)$
SGD	$O\left(\frac{L}{\epsilon} + \frac{L\sigma^2}{\epsilon^2}\right)$	$O\left(\frac{L}{\epsilon} + \frac{L\sigma^2}{\epsilon^2}\right)$
mb-SGD	$O\left(\frac{L}{\epsilon} + \frac{L\sigma^2}{B\epsilon^2}\right)$	$O\left(\frac{LB}{\epsilon} + \frac{L\sigma^2}{\epsilon^2}\right)$

It is worth pointing out that when the number of samples M is huge and a low precision solution is satisfactory,⁴ $\frac{\sigma}{M\sqrt{L}} \ll \epsilon$ usually holds. As a result, SGD is favored for solving big data problems.

1.2.3 Minibatch Stochastic Gradient Descent (mb-SGD)

A straightforward variant of the GD algorithm is to compute the gradient of a minibatch of samples (instead of a single sample) in each iteration, that is,

$$\mathbf{g}^{\mathcal{B}}(\mathbf{x}) = \frac{1}{B} \sum_{m \in \mathcal{B}} F'_m(\mathbf{x}), \quad (1.16)$$

where $B := |\mathcal{B}|$. The minibatch \mathcal{B} is obtained by using i.i.d. samples with (or without) replacement. One can easily verify that

$$\mathbb{E}[\mathbf{g}^{\mathcal{B}}(\mathbf{x})] = f'(\mathbf{x}).$$

Sample “With” Replacement. The stochastic variance (for the “with” replacement case) can be bounded by

$$\begin{aligned} & \mathbb{E}[\|\mathbf{g}^{\mathcal{B}}(\mathbf{x}) - f'(\mathbf{x})\|^2] \\ &= \mathbb{E} \left[\left\| \frac{1}{B} \sum_{m \in \mathcal{B}} (F'_m(\mathbf{x}) - f'(\mathbf{x})) \right\|^2 \right] \end{aligned} \quad (1.17)$$

⁴A low precision solution is satisfactory in many application scenarios, since a high precision solution may cause an unwanted overfitting issue.

$$\begin{aligned}
&= \frac{1}{B} \sum_{m \in \mathcal{B}} \mathbb{E} \left[\|F'_m(\mathbf{x}) - f'(\mathbf{x})\|^2 \right] \\
&\leq \frac{\sigma^2}{B} \quad (\text{from Assumption 2}).
\end{aligned}$$

Sample “Without” Replacement. The stochastic variance for the “without” replacement is even smaller, but it involves a bit more complicated derivation. We essentially need the following key lemma.

Lemma 1.3. Give a set including $M \geq 2$ real numbers $\{a_1, a_2, \dots, a_M\}$. Define a random variable

$$\bar{\xi}_{[B]} := \frac{1}{B} \sum_{m=1}^B \xi_m,$$

where ξ_1, \dots, ξ_B are uniformly randomly sampled from the set “without” replacement, and $B(1 \leq B \leq M)$ is the batch size. Then the following equality holds

$$\mathbf{Var}[\bar{\xi}] = \left(\frac{M-B}{M-1} \right) \frac{\mathbf{Var}[\xi_1]}{B}.$$

Proof. First, it is not hard to see that the marginal distributions of ξ_m 's are identical. For simplicity of notation, we assume that $\mathbb{E}[\xi_m] = 0$ without the loss of generality. Therefore, we have $\mathbf{Var}[\xi_m] = \mathbb{E}[\xi_m^2]$ for all k .

Next we have the following derivation:

$$\begin{aligned}
\mathbf{Var}[B\bar{\xi}_{[B]}] &= \mathbb{E}[(B\bar{\xi}_{[B]})^2] \quad (\text{due to } \mathbb{E}[\bar{\xi}_{[B]}] = 0) \\
&= \sum_{m=1}^B \mathbb{E}[\xi_m^2] + \sum_{k \neq l} \mathbb{E}[\xi_m \xi_l] \\
&= B\mathbf{Var}[\xi_1] + B(B-1)\mathbb{E}[\xi_1 \xi_2], \tag{1.18}
\end{aligned}$$

where the last equality uses the fact $\mathbb{E}[\xi_m^2] = \mathbf{Var}[\xi_k] = \mathbf{Var}[\xi_1]$ for any k and $E[\xi_k \xi_l] = E[\xi_1 \xi_2]$ for any $k \neq l$. Note that $\mathbf{Var}[M\bar{\xi}_{[M]}] = 0$, since

it has only one possible combination for $\{\xi_1, \xi_2, \dots, \xi_M\}$. Then letting $B = M$ obtains the following dependence from (1.18)

$$\mathbb{E}[\xi_1 \xi_2] = \frac{-1}{M-1} \text{Var}[\xi_1].$$

Plug this result into (1.18)

$$\text{Var}[B\bar{\xi}_{[B]}] = B \left(\frac{M-B}{M-1} \right) \text{Var}[\xi_1],$$

which implies the claimed result. \square

If a_m 's are vectors and satisfy $\frac{1}{M} \sum_{m=1}^M \xi_m = 0$, from Lemma 1.3 one can easily verify

$$\mathbb{E}[\|\bar{\xi}\|^2] = B \left(\frac{M-B}{M-1} \right) \mathbb{E}[\|\xi_1\|^2]. \quad (1.19)$$

Now we are ready to compute the stochastic variance for the “without” replacement sampling strategy. Let \mathcal{B} be a batch of samples “without” replacement. Then we let $\xi_m := F'_m(\mathbf{x}) - f'(\mathbf{x})$ and from (1.19) obtain

$$\begin{aligned} \mathbb{E}[\|\mathbf{g}^{\mathcal{B}}(\mathbf{x}) - f'(\mathbf{x})\|^2] &= \mathbb{E}[\|\bar{\xi}\|^2] \\ &= \left(\frac{M-B}{M-1} \right) \frac{\mathbb{E}[\|\xi_1\|^2]}{B} \\ &\leq \left(\frac{M-B}{M-1} \right) \frac{\sigma^2}{B} \\ &\leq \frac{\sigma^2}{B}. \end{aligned}$$

To sum up, we have the stochastic variance bounded by $\frac{\sigma^2}{B}$ no matter “with” or “without” replacement sampling.

We can observe that the effect of using a minibatch stochastic gradient is nothing but reduced variance. All remaining analysis for the convergence rate remains the same. Therefore, it is quite easy to obtain the convergence rate of `mb-SGD`

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] \lesssim \frac{L}{T} + \frac{\sqrt{LB}}{\sqrt{TB}}. \quad (1.20)$$

The iteration complexity and the query complexity are reported in Table 1.2.

1.3 A Simplified Distributed Communication Model

When scaling up the stochastic gradient descent (SGD) algorithm to a distributed setting, one often needs to develop *system relaxations* techniques to achieve better performance and scalability. In this monograph, we describe multiple popular system relaxation techniques that have been developed in recent years. In this subsection, we introduce a simple performance model of a distributed system, which will be used in later sections to reason about the performance impact of different relaxation techniques.

From a mathematical optimization perspective, all of the system relaxations that we will describe *do not make the convergence (loss vs. # iterations/epochs) faster*.⁵ Then *why do we even want to introduce these relaxations into our system in the first place?*

One common theme of the techniques we cover in this monograph is that their goal is not to improve the convergence rate in terms of # iterations/epochs; rather, their goal is to make each iteration finish faster in terms of wall-clock time. As a result, to reason about each system relaxation technique in this monograph, we need to first agree on a *performance model* of the underlying distributed system. In this subsection, we introduce a very simple performance model — it ignores many (if not most) important system characteristics, but it is just informative enough for readers to understand why each system relaxation technique in this monograph actually makes a system faster.

1.3.1 Assumptions

In practice, it is often the case that the bandwidth or latency of each worker’s network connection is the dominating bottleneck in the communication cost. As a result, in this monograph we focus on the following simplified communication model.

Figure 1.2 illustrates our communication model. Each worker (blue rectangle) corresponds to *one* computation device (worker), and all

⁵The reason that we emphasize the “mathematical optimization” perspective is that some researchers find that certain system relaxations can actually lead to better generalization performance. We do not consider generalization in this monograph.

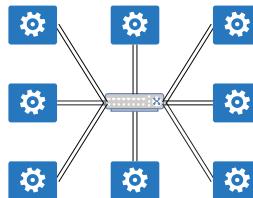


Figure 1.2: An illustration of the distributed communication model we use in this monograph. We assume that all devices (worker, machine) are connected via a “logical switch” whose property is defined in Subsection 1.3.

workers are connected via a “logical switch” that has the following property.

1. The switch has infinitely large bandwidth. We make this simplifying assumption to reflect the observation that, in practice, the bottleneck is often the bandwidth or latency of each worker’s network connection.
2. For each message that “passes through” the switch (sent by worker w_i and received by worker w_j), the switch adds a constant delay t_{latency} independently of the number of concurrent messages that this switch is serving. This delay is the timestamp difference between the sender sending out the first bit and the receiver receiving the first bit.

For each worker, we also assume the following properties.

1. Each worker can only send one message at the same time.
2. Each worker can only receive one message at the same time.
3. Each worker can concurrently receive one message and send one message at the same time.
4. Each worker has a fixed bandwidth, i.e., to send/receive one unit (e.g., MB) amount of data, it requires $t_{\text{transfer1MB}}$ seconds.

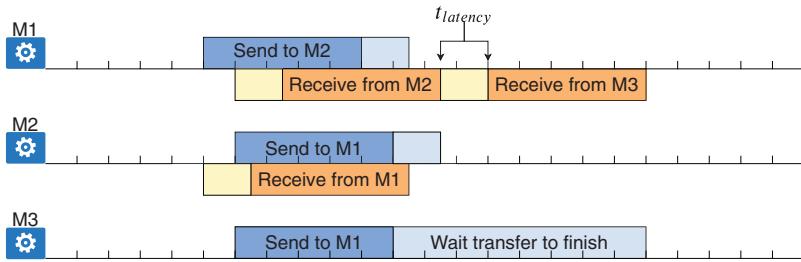


Figure 1.3: Illustration of the communication pattern of Example 1.2.

Example 1.1. Under the above communication model, consider the following three events:

time	event
0:05	M1 send 1MB to M2
0:06	M2 send 1MB to M1
0:06	M3 send 1MB to M2

We assume that the latency added by the switch $t_{latency}$ is 1.5 units of time and it took 5 units of time to transfer 1 MB of data. Figure 1.3 illustrates the timeline on three machines under our communication model. The yellow block corresponds to the *latency* added by the “logical switch”. We also see that the machine M1 can concurrently send (blue block) and receive (orange block) data at the same time; however, when the machine M3 tries to send data to M1, because the machine M2 is already sending data to M1, M3 needs to wait (the shallow blue block of M3).

Example 1.2. Figure 1.4 illustrates a hypothetical scenario in which all data sent in Example 1.2 are “magically” compressed by $2\times$ at the sender. As we will see in later sections, this is similar to what would happen if one were to compress the gradient by $2\times$ during training.

We make multiple observations from Figure 1.4.

1. First, compressing data does make the “system” faster. Without compression, all three events finish in 14 units of time (Figure 1.3) whereas it finishes in 9 units of time after compression. This is

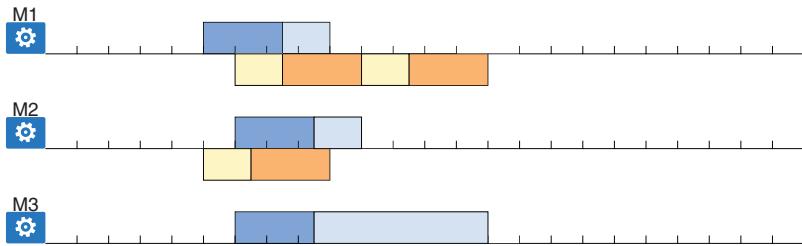


Figure 1.4: Illustration of the communication pattern of Example 1.2, with $2\times$ data compression.

because the time used to *transfer* the data is decreased by half in our communication model.

2. Second, even if the data are compressed by $2\times$, the speedup of the system is smaller than that; in fact, it is only $14/9 = 1.55\times$. This is because, even though the transfer time is cut by half, the communication latency does not decrease as a result of the data compression.

We now use the above communication model to describe the *communication patterns* of three popular ways to implement distributed stochastic gradient descent. These implementations will often serve as the baseline from which we apply different system relaxations to remove certain system bottlenecks that arise in different configurations of $(t_{\text{latency}}, t_{\text{transfer}})$ together with the relative computational cost on each machine.

Workloads. We focus on one of the core building blocks to implement a distributed SGD system — each worker M_i holds a parameter vector w_i , and they communicate to compute the sum of all parameter vectors: $S = \sum_i w_i$. At the end of communication, each worker holds one copy of S .

1.3.2 Synchronous Parameter Server

The parameter server is not only one of the most popular system architectures for distributed stochastic gradient descent; it is also one

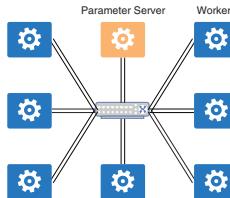


Figure 1.5: Illustration of the parameter server architecture with a single dedicated parameter server.

of the most popular communication models that researchers have in mind when they conduct theoretical analysis. In a parameter server architecture, one or more machines serve as the *parameter server(s)* and other machines serve as the workers processing data. Periodically, workers send updates to the parameters to the parameter servers and the parameter servers send back the updated parameters. Figure 1.5 illustrates this architecture: the orange machine is the parameter server and the blue machines are the workers.

A real-world implementation of a parameter server architecture usually involves many system optimizations to speed up the communication. In this subsection, we build our abstraction using the simplest implementation with only a single machine serving as the parameter server. We also scope ourselves and only focus on the synchronous communication case.

When using this simplified parameter server architecture to calculate the sum S , each worker M_i sends their local parameter vector w_i to the parameter server, and the parameter server collects all these local copies, sums them up, and sends back to each worker. In a simple example with three workers and one parameter server, the series of communication events looks like this:

Time=0	Worker1 send w_1 to PS
Time=0	Worker2 send w_2 to PS
Time=0	Worker3 send w_3 to PS
Time=T	PS send S to Worker1
Time=T	PS send S to Worker2
Time=T	PS send S to Worker3

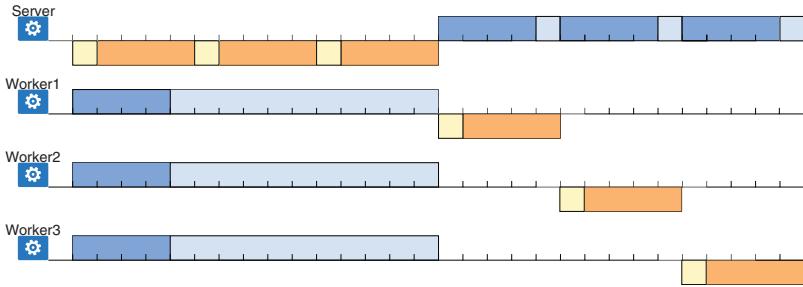


Figure 1.6: Illustration of the communication pattern of the parameter server architecture with a single dedicated parameter server.

Figure 1.6 illustrates the communication timeline of these events. We see that, in the first phase, all workers send their local parameter vectors to the parameter server at the same time. Because, in our communication model, the parameter server can receive data from only one worker at a time, it took $3(t_{\text{latency}} + t_{\text{transfer}})$ for the aggregation phase to finish. In the broadcast phase, because, in our communication model, the parameter server can send data to only one worker at a time, it took $3(t_{\text{latency}} + t_{\text{transfer}})$ for the broadcast phase to finish.

In general, when there are N workers and 1 parameter server, *as under our communication model*, a parameter server architecture in which all workers are *perfectly synchronized* takes

$$2n(t_{\text{latency}} + t_{\text{transfer}})$$

to compute and broadcast the sum S over all local copies $\{w_i\}$.

Discussions. As we see, the communication cost of the parameter server architecture grows linearly with respect to the total number of workers we have in the system. As a result, this architecture could be sensitive to both latency and transfer time. This motivates some system relaxations that could alleviate potential system bottlenecks.

1. When the network has small latency t_{latency} compared with the transfer time t_{transfer} , one could conduct lossy compression (e.g., via quantization, sparsification, or both) to decrease the transfer time. Usually, this approach can lead to a linear speedup with

respect to the compression rate, up to a point that t_{latency} starts to dominate.

- When the network has large latency t_{latency} , compression on its own won't be the solution. In this case, one could adopt a decentralized communication pattern, as we will discuss later in this monograph.

1.3.3 AllReduce

Calculating the sum over distributed workers is a very common operator used in distributed computing and high performance computing systems. In many communication frameworks, it can be achieved using the `AllReduce` operator. Optimizing and implementing the `AllReduce` operator has been studied by the HPC community for decades, and the implementation is usually different for different numbers of machines, different sizes of messages, and different physical communication topologies.

In this monograph, we focus on the simplest case, in which all workers form a *logical* ring and communicate only to their neighbors (all communications still go through the single switch all workers are connected to). We also assume that the local parameter vector is large enough.

Under these assumptions, we can implement an `AllReduce` operator in the following way. Each worker w_n partitions their local parameter vectors into N partitions (N is the number of workers): w_n^k is the k th partition of the local model w_n . The communication happens in two phases.

- Phase 1.** At the first iteration of Phase 1, each machine n sends w_n^n to its “next” worker in the logical ring, i.e., w_j where $j = n + 1 \bmod N$. Once machine j receives a partition k , it sums up the received partition with its local partition, and sends the aggregated partition to the next worker in the next iteration. After $N - 1$ communication iterations, different workers now have the sum of different partitions.
- Phase 2.** Phase 2 is similar to Phase 1, with the difference that when machine n receives a partition k , it replaces its local copy

with the received partition, and passes it onto the next machine in the next iteration.

At the end of communication, all workers have the sum S of all partitions.

Example 1.3. We walk through an example with four workers M_1, \dots, M_4 . The communication pattern of the above implementation is as follows. We use $w\{M_i, j\}$ to denote the j th partition on machine M_i .

```
# For the first partition w{M1 (worker id), 1 (partition id)}
Time= 0   M1 sends w{M1,1}                               to M2
Time= t    M2 sends w{M1,1} + w{M2,1}                  to M3
Time=2t   M3 sends w{M1,1} + w{M2,1} + w{M3,1}        to M4
Time=3t   M4 sends w{M1,1} + w{M2,1} + w{M3,1} + w{M4,1} to M1
Time=4t   M1 sends w{M1,1} + w{M2,1} + w{M3,1} + w{M4,1} to M2
Time=5t   M2 sends w{M1,1} + w{M2,1} + w{M3,1} + w{M4,1} to M3

# For the second partition w{M2 (worker id), 2 (partition id)}
Time= 0   M2 sends w{M2,2}                               to M3
Time= t    M3 sends w{M2,2} + w{M3,2}                  to M4
Time=2t   M4 sends w{M2,2} + w{M3,2} + w{M4,2}        to M1
Time=3t   M1 sends w{M2,2} + w{M3,2} + w{M4,2} + w{M1,2} to M2
Time=4t   M2 sends w{M2,2} + w{M3,2} + w{M4,2} + w{M1,2} to M3
Time=5t   M3 sends w{M2,2} + w{M3,2} + w{M4,2} + w{M1,2} to M4

# For the third partition w{M3 (worker id), 3 (partition id)}
Time= 0   M3 sends w{M3,3}                               to M4
Time= t    M4 sends w{M3,3} + w{M4,3}                  to M1
Time=2t   M1 sends w{M3,3} + w{M4,3} + w{M1,3}        to M2
Time=3t   M2 sends w{M3,3} + w{M4,3} + w{M1,3} + w{M2,3} to M3
Time=4t   M3 sends w{M3,3} + w{M4,3} + w{M1,3} + w{M2,3} to M4
Time=5t   M4 sends w{M3,3} + w{M4,3} + w{M1,3} + w{M2,3} to M1

# For the fourth partition w{M4 (worker id), 4 (partition id)}
Time= 0   M4 sends w{M4,4}                               to M1
Time= t    M1 sends w{M4,4} + w{M1,4}                  to M2
Time=2t   M2 sends w{M4,4} + w{M1,4} + w{M2,4}        to M3
Time=3t   M3 sends w{M4,4} + w{M1,4} + w{M2,4} + w{M3,4} to M4
Time=4t   M4 sends w{M4,4} + w{M1,4} + w{M2,4} + w{M3,4} to M1
Time=5t   M1 sends w{M4,4} + w{M1,4} + w{M2,4} + w{M3,4} to M2
```

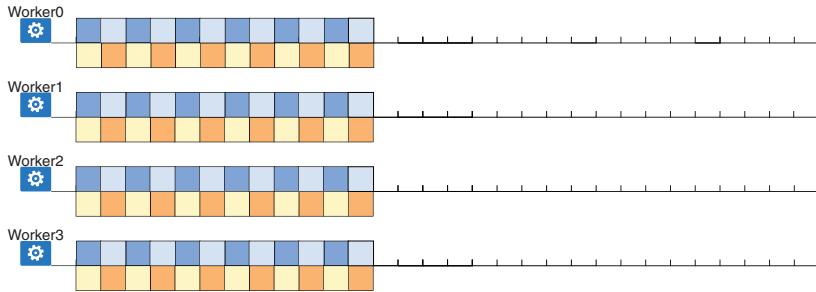


Figure 1.7: Illustration of the communication pattern of the AllReduce architecture with ring topology.

From the above pattern, it is not hard to see why, at the end, each worker has a copy of $S = \sum_{n=1}^N w_n$.

One interesting property of the above way of implementing the **AllReduce** operator is that, at any timestep, each machine concurrently sends and receives one partition of the data, which is possible in our communication model. Figure 1.7 illustrates the communication timeline.

We make multiple observations.

1. Compared with a parameter server architecture with a single parameter server (Figure 1.6), the total amount of data that each worker sends and receives is the same in both cases — in both cases, the amount of data sent and received by each machine is equal to the size of the parameter vector.
2. At any given time, each worker sends and receives data concurrently. At any given time, only the left neighbor w_n sends data to $w_{n+1 \bmod N}$ and $w_{n+1 \bmod N}$ only sends data to $w_{n+2 \bmod N}$. This allows the system to take advantage of the *aggregated* bandwidth of N machines (which grows linearly with respect to N) instead of being bounded by the bandwidth of a single central parameter server.

In general, when there are $N + 1$ workers, *as under our communication model* and assuming that the computation cost to sum up

parameter vectors is negligible, an `AllReduce` operator in which all workers are *perfectly synchronized* took

$$2Nt_{\text{latency}} + 2t_{\text{transfer}}$$

to compute and broadcast the sum S over all local copies $\{w_n\}$.

Discussions. As we see, the latency of an `AllReduce` operator grows linearly with respect to the total number of workers we have in the system. As a result, this architecture could be sensitive to network latency. This motivates some system relaxations that could alleviate potential system bottlenecks.

1. When the network has large latency t_{latency} , compression on its own won't be the solution. In this case, one could adopt a decentralized communication pattern, as we will discuss later in this monograph.
2. When the network has small latency t_{latency} and the parameter vector is very large, the transfer time t_{transfer} can still become the bottleneck. In this case, one could conduct lossy compression (e.g., via quantization, sparsification, or both) to decrease the transfer time. Usually, this approach can lead to a linear speedup with respect to the compression rate, up to a point that t_{latency} starts to dominate.

Caveats. We will discuss the case of asynchronous communication later in this monograph. Although it is quite natural to come up with an asynchronous parameter server architecture, making the `AllReduce` operator run in an asynchronous fashion is less natural. As a result, when there are stragglers in the system (e.g., one worker is significantly slower than all other workers), `AllReduce` can make it more difficult to implement a straggler avoidance strategy if one simply uses the off-the-shelf implementation.

Why Do We Partition the Parameter Vector? One interesting design choice in implementing the `AllReduce` operator is setting each local

parameter vector to be partitioned into N partitions. This decision is important if you want to fully take advantage of the aggregated bandwidth of *all* workers. Take the same four-worker example and assume that we do not partition the model. In this case, the series of communication events will look like this:

Time= 0	M1 sends $w\{M1\}$	to M2
Time= t	M2 sends $w\{M1\} + w\{M2\}$	to M3
Time=2t	M3 sends $w\{M1\} + w\{M2\} + w\{M3\}$	to M4
Time=3t	M4 sends $w\{M1\} + w\{M2\} + w\{M3\} + w\{M4\}$	to M1
Time=4t	M1 sends $w\{M1\} + w\{M2\} + w\{M3\} + w\{M4\}$	to M2
Time=5t	M2 sends $w\{M1\} + w\{M2\} + w\{M3\} + w\{M4\}$	to M3

In general, with $N + 1$ workers, the communication cost without partitioning becomes

$$2N(t_{\text{latency}} + t_{\text{transfer}}).$$

Comparing this with the $2Nt_{\text{latency}} + 2t_{\text{transfer}}$ cost of `AllReduce` with model partition, we see that model partition is the key reason for taking advantage of the full aggregated bandwidth provided by all machines.

1.3.4 Multi-Machine Parameter Server

One can extend the single-server parameter server architecture and use multiple machines serving as parameter servers instead. In this monograph, we focus on the scenario in which each worker also serves as a parameter server.

Under this assumption, we can implement a multi-server parameter server architecture in the following way. Each worker w_n partitions their local parameter vectors into N partitions (N is the number of workers): w_n^k . The communication happens in two phases.

1. Phase 1. All workers send their n th partition to worker w_n . Worker w_n aggregates all messages and calculates the n th partition of the sum S .
2. Phase 2. Worker w_n sends the n th partition of the sum S to all other workers.

With careful arrangement of communication events, we can also take advantage of the full aggregated bandwidth in this architecture, as illustrated in the following example.

Example 1.4. We walk through an example with four workers M_1, \dots, M_4 . The communication pattern of the above implementation is as follows:

```
# First partition: w{M1 (worker id), 1 (partition id)}
# First partition of the result: S{1}
Time= 0    M2 sends w{M2,1} to M1
Time= t    M3 sends w{M3,1} to M1
Time=2t    M4 sends w{M4,1} to M1
Time=3t    M1 sends S{1}      to M2
Time=4t    M1 sends S{1}      to M3
Time=5t    M1 sends S{1}      to M4

# Second partition: w{M2 (worker id), 2 (partition id)}
# Second partition of the result: S{2}
Time= 0    M1 sends w{M1,2} to M2
Time= t    M4 sends w{M4,2} to M2
Time=2t    M3 sends w{M3,2} to M2
Time=3t    M2 sends S{2}      to M3
Time=4t    M2 sends S{2}      to M4
Time=5t    M2 sends S{2}      to M1

# Third partition w{M3 (worker id), 3 (partition id)}
# Third partition of the result: S{3}
Time= 0    M4 sends w{M4,3} to M3
Time= t    M1 sends w{M1,3} to M3
Time=2t    M2 sends w{M2,3} to M3
Time=3t    M3 sends S{3}      to M4
Time=4t    M3 sends S{3}      to M1
Time=5t    M3 sends S{3}      to M2

# Fourth partition w{M4 (worker id), 4 (partition id)}
# Fourth partition of the result: S{4}
```

Time= 0	M3 sends $w\{M3,4\}$ to M4
Time= t	M2 sends $w\{M2,4\}$ to M4
Time=2t	M1 sends $w\{M1,4\}$ to M4
Time=3t	M4 sends $S\{4\}$ to M1
Time=4t	M4 sends $S\{4\}$ to M2
Time=5t	M4 sends $S\{4\}$ to M3

For the above communication events, it is not hard to see that, at the end, each machine has access to the sum $S = \sum_{n=1}^N w_n$. In terms of the communication pattern, under our communication model, the multi-server parameter server architecture has the same pattern as **AllReduce**, illustrated in Figure 1.7.

In general, when there are $N + 1$ workers, *as under our communication model* and assuming that the computation cost to sum up parameter vectors is negligible, a multi-server parameter server architecture in which all workers are *perfectly synchronized* took

$$2Nt_{\text{latency}} + 2t_{\text{transfer}}$$

to compute and broadcast the sum S over all local copies $\{w_n\}$.

2

Distributed Stochastic Gradient Descent

The previous section provides us with the background of stochastic gradient descent and a simple communication model. This allows us to start analyzing the performance of a simple, distributed stochastic gradient descent system, which will serve as the baseline for the remaining part of this monograph.

2.1 A Simplified Performance Model for Distributed Synchronous Data-Parallel SGD

Recall the optimization problem that we hope to solve:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \quad \left\{ f(\mathbf{x}) := \frac{1}{M} \sum_{m=1}^M F_m(\mathbf{x}) \right\}. \quad (2.1)$$

The stochastic gradient descent algorithm works by sampling, uniformly randomly with replacement, a term $m_t \in [M]$, and updating the current model \mathbf{x}_t with

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \gamma F'_{m_t}(\mathbf{x}_t)$$

until convergence.



Figure 2.1: Illustration of SGD on a single machine.

2.1.1 SGD on a Single Machine

Implementing SGD on a single machine, with a single thread, is easy. The system stores the current model \mathbf{x}_t in memory, and repeats two stages, as illustrated in Figure 2.1.

1. Computation. The system: (1) fetches \mathbf{x}_t from the main memory; (2) fetches all information necessary to compute $F'_{m_t}(\mathbf{x}_t)$; and (3) computes $F'_{m_t}(\mathbf{x}_t)$.
2. Update. The system updates \mathbf{x}_t with $F'_{m_t}(\mathbf{x}_t)$.

2.1.2 Data-Parallel SGD on Multiple Devices

When distributing the above algorithm on multiple devices, there are multiple ways of distributing the workload for different system bottlenecks. For example, when the model \mathbf{x}_t is too large and does not fit into the fast memory of a single device, it can be partitioned onto different devices for the computation phase. This strategy is called *model parallelism*. In this monograph, we focus on what is called *data parallelism* — each device has access to a partition of the data set (i.e., $[M]$), and repeats a three-stage process as illustrated in Figure 2.2.

1. Computation. Each worker (say worker n) samples, from its local partition, an $F'_{m_t^{(n)}}(\mathbf{x}_t)$ to compute. $m_t^{(n)}$ denotes the index sampled by worker n at iteration t .

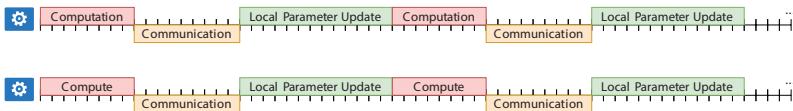


Figure 2.2: Illustration of the distributed synchronous data-parallel SGD.

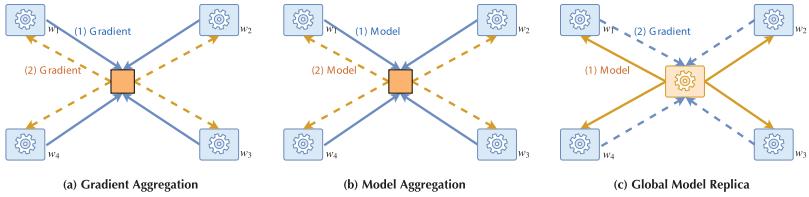


Figure 2.3: Illustration of three possible implementations of data-parallel SGD on multiple devices.

2. Communication. Workers communicate to compute the sum of all local gradients $\sum_n F'_{m_t^{(n)}}(\mathbf{x}_t)$.
3. Update. The system updates \mathbf{x}_t with $\sum_n F'_{m_t^{(n)}}(\mathbf{x}_t)$.

The above algorithm can be implemented in three ways, depending on whether the system aggregates the gradient or the model, and the locality of model. Figure 2.3 illustrates these three implementations.

Gradient Aggregation. One simple strategy to implement the above algorithm is for each worker w_n to maintain a local model replica $\mathbf{x}_t^{(n)}$. At the very beginning, the replica on all workers is the same:

$$\mathbf{x}_0^{(1)} = \dots = \mathbf{x}_0^{(N)}.$$

In the communication phase, the system aggregates the gradient using any one of the three communication primitives we introduced before (e.g., `AllReduce`). At the end of communication, each worker sees the same aggregated gradient

$$\sum_n F'_{m_t^{(n)}}(\mathbf{x}_t).$$

In the update phase, each worker applies the update locally, independently. Because all workers see the same aggregated gradient, their model replica stays equal after the updates.

Model Aggregation. It is possible to implement the system in a different way. Each worker w_n still maintains the local replica $\mathbf{x}_t^{(n)}$ and makes sure they are the same at the very beginning.

However, the system does the local update first by applying the local gradient to each local model:

$$\mathbf{x}_{t+1/2}^{(n)} = \mathbf{x}_t^{(n)} - \gamma F'_{m_t^{(n)}}(\mathbf{x}_t^{(n)}).$$

In the communication phase, all workers communicate the *model*, $x_{t+1/2}^{(n)}$ and calculate the sum

$$\mathbf{x}_{t+1} = \frac{1}{N} \sum_n \mathbf{x}_{t+1/2}^{(n)} = \mathbf{x}_t - \gamma \frac{1}{N} \sum_n F'_{m_t^{(n)}}(\mathbf{x}_t^{(n)}).$$

The system then uses this sum \mathbf{x}_{t+1} to update its local model replica.

Global Model Replica. The third way of implementing the system is to maintain a *global*, instead of *local*, model replica. This global replica is stored on one, or multiple, *parameter servers*. At the very beginning, each worker w_j fetches the global model replica \mathbf{x}_t from the parameter servers, calculates the local gradient, and sends the local gradient to the parameter servers. The parameter server then calculates the sum of all local gradients, and updates its global replica.

2.1.3 Discussion

It is easy to see that all three implementations logically implement the same algorithm. However, the system tradeoff among these three different implementations can be quite delicate.

1. **Size of Model = Size of Gradient.** When the size of the model is the same as the size of the gradient, these three implementations can have similar communication cost. When we implement the gradient aggregation and model aggregation approach using `AllReduce` and the global model replica using a multi-server parameter server, we see that all three implementations have the communication cost

$$2Nt_{\text{latency}} + 2t_{\text{transfer}}$$

where the transfer time t_{transfer} depends only on the size of the model and the gradient. In this case, we would expect the end-to-end performance of these three implementations to be similar.

2. **Size of Model \neq Size of Gradient.** The tradeoff between these three implementations becomes more complex when the size of the model does not equal the size of the gradient. For example, there could be a very dense model but a very sparse gradient, or vice versa. When this happens, these three implementations can have very different performance. Specifically, the communication cost can be summarized as follows:¹

(a) **Gradient Aggregation (`AllReduce`):**

$$2Nt_{\text{latency}} + 2t_{\text{transfer}}^{(\text{grad})}.$$

(b) **Model Aggregation (`AllReduce`):**

$$2Nt_{\text{latency}} + 2t_{\text{transfer}}^{(\text{model})}.$$

(c) **Global Model Replica (Multi-server PS):**

$$2Nt_{\text{latency}} + t_{\text{transfer}}^{(\text{model})} + t_{\text{transfer}}^{(\text{grad})}.$$

Obviously, different applications can have different performance under these three implementations.

In this monograph, *we assume that the size of the model and the size of the gradient are always the same*. Readers might wonder why we bother to introduce all three strategies, given that they all have similar performance under this assumption. As we will see later, different types of system relaxations might be more suitable to be applied to different implementations. For example, asynchronous relaxation might be easier to implement when there is a global model replica. Moreover, some system relaxations might require different theoretical analysis under different implementations. For example, the lossy quantization technique applied to gradient aggregation and model aggregation lead

¹This performance model assumes that the communication cost does not change during the communication process. One example that does not satisfy this assumption is when one wants to use `AllReduce` to aggregate a set of *sparse* gradients (Alistarh *et al.*, 2018b). In this case, the communication will become denser and denser during the communication process (i.e., the sum of two sparse gradients can only become denser).

to different convergence behavior; decentralized relaxation might not work at all if one uses gradient aggregation.

In this monograph, when introducing different system relaxations, we assume that we have the luxury of choosing one from these three implementations without worrying about the delicate tradeoff between them. In practice, given a specific type of ML models, choosing the right implementation and system relaxation is often an engaged, task-specific problem.

2.2 Theoretical Analysis

The theoretical analysis is very similar to the `mb-SGD` we introduced in Subsection 1.2.3. One can simply imagine that the minibatch stochastic gradient in (1.16) is computed by N workers and that the minibatch size is $N = B$. Then, the convergence rate can be easily obtained from (1.20)

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] \lesssim \frac{L}{T} + \frac{\sqrt{L}\sigma}{\sqrt{TN}}. \quad (2.2)$$

2.3 Caveats

These are multiple aspects of data-parallel SGD implementations that we will not consider in this monograph. These aspects are often critical in practice to achieving good performance; however, including them into theoretical analysis is either impossible (because of the lack of an analytical understanding) or more engaged (because of the sophisticated performance model introduced by these factors). The goal of this monograph is *not* to cover all aspects of distributed learning; instead, providing a minimalist tutorial of the basics. As a result, we briefly discuss these aspects. We refer readers to Section 6 for further reading about these topics.

- 1. Batch Size.** One aspect that we do not consider in our analysis is the impact of batch size on model accuracy. The underlying assumption of this monograph is that using a larger batch size, once converged, will lead to the same model accuracy as a smaller batch size. This assumption allows us to treat the impact of

larger batch size as a way to reduce the variance of the stochastic gradient. However, in practice, especially for models such as deep neural networks, the impact of batch size can be quite delicate and could also have an impact on accuracy and generalization.

2. **Hardware Parallelism.** Another assumption that we are making on the system side is that the time that one device needs to compute stochastic gradient over *one* data point is K times cheaper than computing stochastic gradient over a batch of K data points. This is true in terms of the number of floating point operations that one needs to conduct; however, it might not necessarily be true with real-world hardware in terms of wall-clock time. For hardware, especially those optimizing throughput such as GPUs, batching might have a significant impact on performance. In some cases, calculating a single data point might be as expensive as processing a batch of data points, in terms of wall-clock time. This aspect, together with the impact of batch size on model accuracy, could make it quite complicated to design distributed learning systems.
3. **Model Structure.** In this monograph, we assume that the model (\mathbf{x}_t) is a dense vector that does not have much structure. In practice, the model might have more structures. For example, if \mathbf{x}_t corresponds to a deep neural network, it can be decomposed into multiple layers. If one uses the standard back-propagation algorithm, the communication of layer can be overlapped with the computation (e.g., one can calculate the gradient of layer l when communicating for layer $l + 1$ in the backward pass). Recent research has tried to take advantage of this structure (See Section 6). As another example, when the data or model has some sparsity structure, one could also take advantage of it during (distributed) training.
4. **Communication Model.** Last but not least, we assume that communication via the “logical switch” is the only way that different devices can communicate. However, in practice, modern hardware is more complicated than this — different CPU cores

might “communicate” via the shared L3 cache; different CPU sockets might “communicate” via QPI; a GPU/CPU hybrid system has even more delicate ways of communicating. Designing a distributed learning system when each of the workers is a multi-socket CPU system together with multiple GPU devices is more complicated than the simple performance model that we cover in this monograph.

3

System Relaxation 1: Lossy Communication Compression

In distributed systems, data movement can often be significantly slower than computation. This is also often true for distributed learning — a forward and backward pass to calculate the gradient on a 152-layer ResNet requires roughly 3×11 GFLOPS and has a 230 MB model. In *theory*, with 16 GPUs, each of which provides 10 TFLOPS, and a 40 Gbit network, a 256-image batch would require 0.05 seconds for gradient calculation, but would require 0.09 seconds for network transfer time. In practice, both processes are often slower than the theoretical peak; however, their relative order remains similar.

In this section, we focus on one system relaxation technique to speed up the expensive gradient exchange step in distributed SGD — instead of exchanging the gradient as 32 bit floating point numbers, the system first quantizes it into a lower precision representation before communication. The impact of this relaxation is to decrease the time needed to transfer a model — if one quantizes the gradient into 8 bit fixed point representation, the transfer time would *theoretically* decrease by $4\times$.

3.1 System Implementation

To implement this system relaxation, two aspects of the process need to be addressed.

1. How can a floating point vector \mathbf{x} , representing either the model or the gradient, be compressed, using a function $Q(\cdot)$ such as the result $Q(\mathbf{x})$, which can be stored more efficiently (either more sparsely than \mathbf{x} via sparsification, or using fewer bits via quantization)?
2. How can the lossy compression function $Q(\cdot)$ be used to implement distributed SGD? As we will see later, compressing different communication channels for different implementations of distributed SGD (i.e., model aggregation, gradient aggregation, and a global model replica) actually leads to different algorithms and needs different analysis of convergence.

3.1.1 Lossy Communication Compression

There are multiple ways of conducting lossy compression that are popularly used for distributed SGD. In this monograph, we focus on *unbiased lossy compression*, which ensures that

$$\mathbb{E}_\xi [Q(\mathbf{x}; \xi)] = \mathbf{x}.$$

The term ξ denotes the randomness of compression, and the above equation ensures that the compressed vector $Q(\mathbf{x}; \xi)$ is *unbiased* with respect to the original vector x .

Unbiased Compression via Quantization. There are multiple ways of conducting unbiased compression, and quantization is arguably the simplest approach. Figure 3.1 illustrates this process.

Given an original input vector \mathbf{x} , stored as floating points (e.g., 32 bits per element), let $\min(\mathbf{x})$ and $\max(\mathbf{x})$ be the smallest and the largest element of \mathbf{x} . A simple way of quantizing the input vector using b -bits per element is simply to quantize each element of \mathbf{x} , \mathbf{x}_i independently. With b -bits, one can represent 2^b “knobs” partitioning

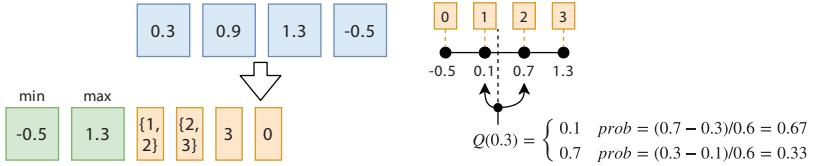


Figure 3.1: Illustration of one simple quantization method.

the range $[\min(\mathbf{x}), \max(\mathbf{x})]$. Assume these knobs are uniformly located as $\{c_0, \dots, c_{2^b-1}\}$, i.e.,

$$c_i = i \times \frac{\max(\mathbf{x}) - \min(\mathbf{x})}{2^b - 1} + \min(\mathbf{x}).$$

Assume that the i th element of \mathbf{x} , x_i falls between $[c_i, c_{i+1}]$; we can use the following randomized quantization (RQ) function:

$$Q(\mathbf{x}_i; \xi_i) = \begin{cases} c_i & \xi > \frac{x_i - c_i}{c_{i+1} - c_i} \\ c_{i+1} & \text{otherwise} \end{cases} \quad (3.1)$$

where ξ_i is a uniform random variable on $[0, 1]$.

Figure 3.1 illustrates this process for $b = 2$ (i.e., 2-bit per element) and $\min(\mathbf{x}) = -0.5$, $\max(\mathbf{x}) = 1.3$. Take the first element 0.3, for example: it falls in the internal $[0.1, 0.7]$. As a result, with probability $\frac{2}{3}$ the system quantizes the element 0.3 to 0.1 and with probability $\frac{1}{3}$, the system quantizes the element 0.3 to 0.7. If one calculates the expectation, we have

$$\frac{2}{3} \times 0.1 + \frac{1}{3} \times 0.7 = 0.3$$

Because there are only 2^b knobs that the quantization function $Q(-)$ can take values in, the output can be encoded using b bits per element.

Caveats. The above approach is one of the simplest ways to construct an unbiased compression function. There are multiple ways that it can be improved. We briefly summarize these caveats below and refer readers to Section 6 for further readings.

1. In practice, more sophisticated quantization functions can be constructed. For example, instead of normalizing the vector using l_∞ norm (i.e., $\min(\mathbf{x}), \max(\mathbf{x})$), one can normalize it using

other norms, say l_2 . One can also partition the vectors into different “buckets” and quantize each bucket independently. Moreover, the “knobs” do not need to be uniformly distributed over $[\min(\mathbf{x}), \max(\mathbf{x})]$, and one can even learn those knobs as part of the learning process. All these approaches have been used for communication quantization and can sometimes outperform the baseline quantization strategy that we just introduced.

2. If the original input vector is stored as 32-bit floating point numbers, quantization can only provide at most $32\times$ compression, because one cannot compress each element below a single bit. There are other approaches that do not have this limitation. One such approach is called *sparsification*, which maps the input vector x to another vector x' that is more sparse. This mapping can also be made to be unbiased, a strategy that has been popular in practice.
3. In this monograph, we focus on the scenario in which the lossy compression scheme is *unbiased*. However, a *biased* compression scheme can still be used while still achieving convergence of the training process. For example, a vector \mathbf{x} can be compressed by taking the top- k element. This operation is clearly biased, but there is recent work proving convergence using this type of biased compression scheme.

3.1.2 SGD with Lossy Communication Compression (CSGD)

It is possible to use a quantization function $Q(-)$ to compress communications for distributed training. In the previous sections, we described three different ways of implementing distributed SGD and different communication primitives; for each of these implementations and primitives, compressing the communication might actually lead to different algorithms and convergence behaviors.

Multi-Server Parameter Server for Gradient Aggregation. The easiest implementation for demonstrating the impact of lossy compression is probably gradient aggregation using a multi-server parameter server.

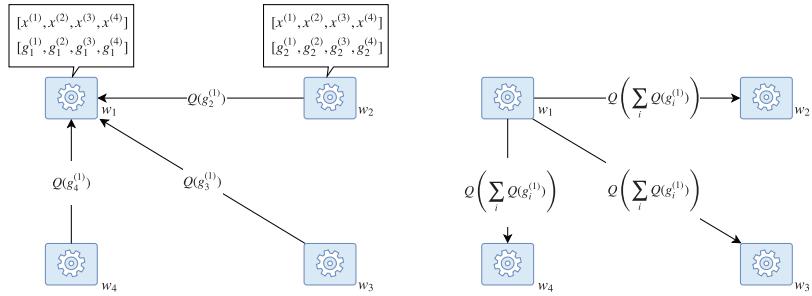


Figure 3.2: Illustration of multi-server parameter servers with lossy communication compression. In this example, there are four workers. Each worker is the parameter server of one partition of the model. This figure illustrates the communication for the first partition of the model $x^{(1)}$ hosted by the first worker w_1 .

In this strategy, each machine holds a replica of the full model. Given N machines, the gradient vector is partitioned into N chunks — each machine w_n is responsible for aggregating the n th partition, as illustrated in Figure 3.2.

Specifically, all machines send their n th partition of the local gradient to the worker w_n , which aggregates all incoming vectors and broadcasts back the sum. Both the incoming and outgoing messages can be compressed by using the quantization function. Let the local gradient on worker n be g_n ; after aggregation, each worker receives

$$Q\left(\frac{1}{N} \sum_{n=1}^N Q(g_n)\right). \quad (3.2)$$

AllReduce for Gradient Aggregation. We can also apply lossy compression to AllReduce, which makes the system behavior more complicated to analyze. Figure 3.3 illustrates the communication pattern.

Specifically, given n workers, the system divides the local gradient into n partitions. Each partition “flows through” a ring formed by all the machines, and keeps getting aggregated with the local gradient (Figure 3.3 (left)). At the end, one worker has the result of the aggregation. In the second step, the aggregated result “flows through” a ring of all the machines again — whenever a machine receives the aggregated

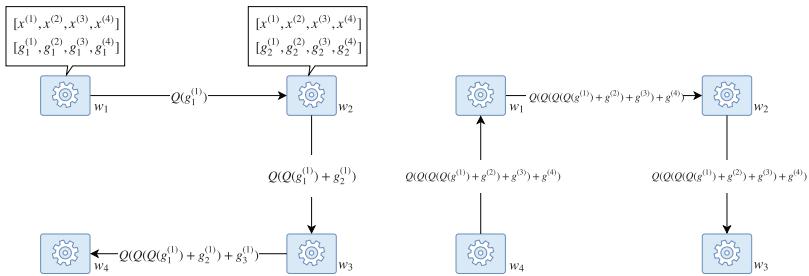


Figure 3.3: Illustration of AllReduce with lossy communication compression. In this example, there are four workers. Each worker initiates the communication of one partition of the model. This figure illustrates the communication for the first partition of the model $x^{(1)}$ initiated by the first worker w_1 .

result, it makes a local copy, and sends the aggregated result to the next machine in the ring (Figure 3.3 (left)).

One caveat of this strategy is that, in order to make sure all incoming and outgoing messages are compressed, the sum must be compressed N times if there are N machines in the ring. As a result, as illustrated in Figure 3.3, after aggregation, each worker receives

$$Q(\cdots Q(Q(Q(g_1) + g_2) + g_3) + \cdots + g_N). \quad (3.3)$$

Compared with the aggregated result in the parameter server case ($Q(\frac{1}{N} \sum_{n=1}^N Q(g_n))$), it is clear that the AllReduce case requires more engaged analysis.

Impact of Lossy Compression. Both implementations have a similar impact on the end-to-end performance. Figure 3.4 and Figure 3.5 illustrate the impact of a $2\times$ compression.

We see that lossy compression would decrease the transfer time, as the communicated messages are now smaller. It does not have any impact on latency, as the number of communications the system needs to conduct stays the same. In this example, we ignore the computation cost of conducting the compression, which is often small compared with the computation time of the gradient.

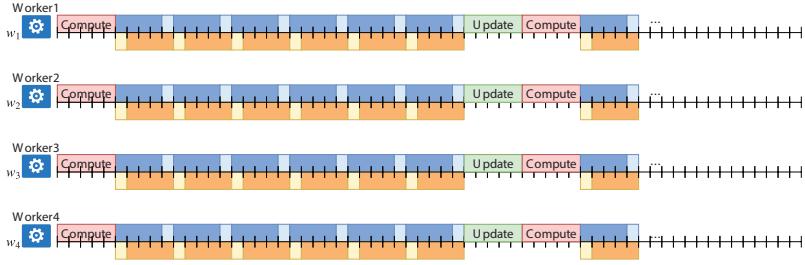


Figure 3.4: Illustration of the impact of lossy compression (without compression).

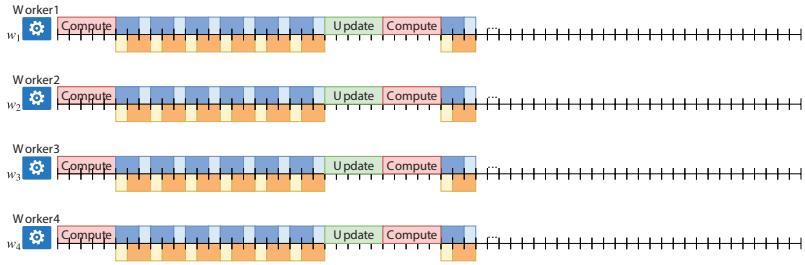


Figure 3.5: Illustration of the impact of lossy compression (with 2× compression).

Specifically, for AllReduce and a multi-server parameter server whose communication cost is

$$2nt_{\text{latency}} + 2t_{\text{transfer}},$$

compressing the communication by K times leads to a communication cost of

$$2nt_{\text{latency}} + \frac{2}{K}t_{\text{transfer}}.$$

When system performance is bounded by the communication cost, and the communication cost is dominated by the transfer time, we observe linear speedup with respect to the compression ratio, in terms of the time needed to finish a single iteration.

3.2 Theoretical Analysis for CSGD

In this subsection, we analyze CSGD, an SGD variant in which the stochastic gradient is compressed by some lossy compression scheme.

The analysis for the CSGD algorithm is very similar to the analysis for stochastic gradient based methods we have seen in previous sections. First, we can see that the basic updating rule for CSGD is nothing but

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \gamma \mathbf{g}_t, \quad (3.4)$$

where \mathbf{g}_t could take the form of either (3.2) or (3.3). But it needs to admit the following two key assumptions, which essentially serve the same purpose as Assumption 2.

- **(Unbiased gradient)** The stochastic gradient is unbiased, that is,

$$\mathbb{E}[\mathbf{g}_t] = f'(\mathbf{x}_t).$$

- **(Bounded stochastic variance)** The stochastic gradient has bounded variance, that is, there exists a constant σ_c satisfying

$$\mathbb{E}[\|\mathbf{g}_t - f'(\mathbf{x}_t)\|^2] \leq \sigma_c^2.$$

Following the same analysis procedure as for SGD in Subsection 1.2.1, we can obtain the following convergence rate for CSGD:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] \lesssim \frac{L}{T} + \frac{\sqrt{L}\sigma_c}{\sqrt{T}}, \quad (3.5)$$

by choosing the learning rate $\gamma = \frac{1}{L+\sigma_c\sqrt{TL}}$. The key difference is that the variance σ_c is different from the σ in Assumption 2. To take a closer look at σ_c , let us consider the form for \mathbf{g}_t to be (3.2) and assume that each stochastic gradient g_n is unbiased with bounded variance $\mathbb{E}[\|g_n - f'(\mathbf{x}_t)\|^2] \leq \sigma^2 \forall n$ and the compression is unbiased with bounded variance σ' , that is,

Assumption 3. (Unbiased compression) The (probably randomized) compression operator $Q(\cdot)$ is assumed to be unbiased:

$$\mathbb{E}[Q(\mathbf{y})] = \mathbf{y}, \quad \forall \mathbf{y}.$$

Assumption 4. The (probably randomized) compression operator $Q(\cdot)$ is assumed to be bounded: **(Bounded compression)**

$$\mathbb{E} [\|Q(\mathbf{y}) - \mathbf{y}\|^2] \leq \sigma'^2, \quad \forall \mathbf{y}.$$

Then we can bound the variance for \mathbf{g}_t by

$$\begin{aligned} & \mathbb{E} [\|\mathbf{g}_t - f'(\mathbf{x}_t)\|^2] \\ & \leq \mathbb{E} \left[\left\| \frac{1}{N} \sum_{n=1}^N Q(g_n) - f'(\mathbf{x}_t) \right\|^2 \right] + \sigma'^2 \\ & = \frac{1}{N^2} \sum_{n=1}^N \mathbb{E} [\|Q(g_n) - f'(\mathbf{x}_t)\|^2] + \sigma'^2 \\ & \leq \frac{1}{N^2} \sum_{n=1}^N (\mathbb{E} [\|Q(g_n) - g_n\|^2] + \mathbb{E} [\|g_n - f'(\mathbf{x}_t)\|^2]) + \sigma'^2 \\ & \leq \underbrace{\frac{\sigma^2}{N} + \left(1 + \frac{1}{N}\right) \sigma'^2}_{=: \sigma_c^2}. \end{aligned}$$

Therefore, from (3.5) the convergence rate of CSGD can be summarized into

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} [\|f'(\mathbf{x}_t)\|^2] \lesssim \frac{L}{T} + \frac{\sqrt{L}\sigma}{\sqrt{NT}} + \underbrace{\frac{\sqrt{L}\sigma'}{\sqrt{T}}}_{\text{caused by compression}}. \quad (3.6)$$

Comparing this approach to the parallel SGD's convergence rate in (2.2), we can see that the last term is the addition caused by using compression. When the stochastic variance σ^2/N dominates the compression variance σ'^2 , then the compression does not affect the convergence rate significantly.

The key to ensuring convergence is the unbiased assumption. One can verify that the randomized quantization compression strategy in (3.1)

satisfies this assumption. Randomized sparsification compression is another method satisfying the unbiased assumption.

- **Randomized Sparsification** (Wangni *et al.*, 2018). For any real number z , with probability p , set z to 0 and $\frac{z}{p}$ with probability p . This is also an unbiased compression operator.
- **Randomized Quantization** (Alistarh *et al.*, 2018b; Wang *et al.*, 2018b; Zhang *et al.*, 2017). For any real number $z \in [a, b]$ (a, b are pre-designed low-bit numbers), with probability $\frac{b-z}{b-a}$, compress z into a , and with probability $\frac{z-a}{b-a}$ compress z into b . This compression operator is unbiased.

However, some popular compression methods do not really satisfy the requisite unbiased assumption, for example:

- **1-Bit Quantization** (Bernstein *et al.*, 2018; Wen *et al.*, 2017). Compress a vector \mathbf{x} into $\|\mathbf{x}\|\text{sign}(\mathbf{x})$ or x , where $\text{sign}(\mathbf{x})$ is a vector whose element takes the sign of the corresponding element in \mathbf{x} . This compression operator is biased.
- **Clipping**. For any real number z , directly set its lower k bits to zero. For example, deterministically compress 1.23456 into 1.2 with its lower 4 bits set to zero. This compression operator is biased.

It is worth pointing that it is usually hard to ensure convergence using biased compression methods in theory, but they can still be used in practice and may yield positive results.

3.3 Error Compensated Stochastic Gradient Descent (EC-SGD) for Arbitrary Compression Strategies

The unbiased compression assumption in Assumption 3 is relatively restrictive. To overcome this limitation, we introduce a very recent algorithm, called error-compensated stochastic gradient descent EC-SGD or DoubleSqueeze (Tang *et al.*, 2019b), which is compatible to any reasonable (potentially biased) compression methods.

To illustrate this algorithm, let us first formally define the objective:

$$\min_{\mathbf{x}}: \quad \left\{ f(\mathbf{x}) := \frac{1}{N} \sum_{n=1}^N f_n(\mathbf{x}) \right\}, \quad (3.7)$$

where $f_n(\mathbf{x}) := \mathbb{E}_{\xi \sim \mathcal{D}_n} F_n(\mathbf{x}; \xi)$ and \mathcal{D}_n denotes the distribution of local data at node n . We do not assume that all nodes can access the whole dataset. Apparently, this is a more general loss function than (2.1).

Algorithm Description. We now consider the parameter server architecture for simplicity. The key idea of EC-SGD is to record the error (or bias) caused by compression and compensate the error in the next round iteratively. More specifically, at the t th iteration, **all workers** (indexed by n) compute their local gradients by

$$\mathbf{v}_t^{(n)} = F'_n(\mathbf{x}_t; \xi_t^{(n)}) + \boldsymbol{\delta}_{t-1}^{(n)} \quad (3.8)$$

$$\boldsymbol{\delta}_t^{(n)} = \mathbf{v}_t^{(n)} - Q(\mathbf{v}_t^{(n)}), \quad (3.9)$$

where $F'_n(\mathbf{x}_t; \xi_t^{(n)})$ is the local stochastic gradient, $\boldsymbol{\delta}_{t-1}^{(n)}$ is the compression error left by the previous iteration, $\mathbf{v}_t^{(n)}$ is the error-compensated gradient, and $\boldsymbol{\delta}_t^{(n)}$ is the new compression error. Note that \mathbf{x}_t is the global model at the t th iteration, which is retrieved by each individual worker. To reduce the communication cost, all workers send the compressed error-compensated gradient $Q(\mathbf{v}_t^{(n)})$ to the parameter server. At the t th iteration, the **parameter server** aggregates all received gradients plus the error $\boldsymbol{\delta}_t$ left in last iteration on the parameter server

$$\mathbf{v}_t = \frac{1}{N} \sum_{n=1}^N Q(\mathbf{v}_t^{(n)}) + \boldsymbol{\delta}_{t-1} \quad (3.10)$$

and then sends the compressed \mathbf{v}_t (that is, $Q(\mathbf{v}_t)$) to all workers recording the new error on the parameter server

$$\boldsymbol{\delta}_t = \mathbf{v}_t - Q(\mathbf{v}_t). \quad (3.11)$$

Note that the parameter server does not need to maintain a global model \mathbf{x}_t . Instead, the (virtual) global model \mathbf{x}_t is retrieved by each worker through the received $Q(\mathbf{v}_t)$, that is,

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \gamma Q(\mathbf{v}_t). \quad (3.12)$$

One can see that all communicated information is compressed.

3.4 Theoretical Analysis for EC-SGD

A big advantage of the EC-SGD algorithm is that it can be compatible with any reasonable compression method and its convergence efficiency is quite robust to the compression. To understand the magic of EC-SGD, we provide the essential mathematical iteration in the following lemma.

Lemma 3.1. The EC-SGD algorithm with N works by following the iteration rule shown below

$$\tilde{\mathbf{x}}_{t+1} = \tilde{\mathbf{x}}_t - \gamma \frac{1}{N} \sum_{n=1}^N F'_n \left(\mathbf{x}_t; \xi_t^{(n)} \right), \quad (3.13)$$

where

$$\begin{aligned} \tilde{\mathbf{x}}_t &:= \mathbf{x}_t - \gamma \Omega_{t-1} \\ \Omega_t &:= \boldsymbol{\delta}_t + \frac{1}{N} \sum_{n=1}^N \boldsymbol{\delta}_t^{(n)}. \end{aligned}$$

From (3.13) it is apparent that the EC-SGD algorithm essentially follows the principle of SGD or GD, except that \mathbf{x}_t gets a little bit of perturbation by Ω_t , which aggregates all compression errors at round t .

Proof. This can be proved by straightforward linear algebra computation, by plugging (3.10) into (3.12):

$$\begin{aligned} &\frac{1}{\gamma} (\mathbf{x}_t - \mathbf{x}_{t+1}) \\ &= Q \left(\boldsymbol{\delta}_{t-1} + \frac{1}{N} \sum_{n=1}^N Q \left(\mathbf{v}_t^{(n)} \right) \right) \\ &= \boldsymbol{\delta}_{t-1} + \frac{1}{N} \sum_{n=1}^N Q \left(\mathbf{v}_t^{(n)} \right) - \boldsymbol{\delta}_t \quad (\text{from (3.11)}) \\ &= \frac{1}{N} \sum_{n=1}^N \left(\mathbf{v}_t^{(n)} - \boldsymbol{\delta}_t^{(n)} \right) + \boldsymbol{\delta}_{t-1} - \boldsymbol{\delta}_t \quad (\text{from (3.9)}) \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{N} \sum_{n=1}^N \left(F'(\mathbf{x}_t; \xi_t^{(n)}) + \delta_{t-1}^{(n)} - \delta_t^{(n)} \right) + \Omega_{t-1} - \Omega_t \quad (\text{from (3.10)}) \\
&= \frac{1}{N} \sum_{n=1}^N F'(\mathbf{x}_t; \xi_t^{(n)}) + \Omega_{t-1} - \Omega_t.
\end{aligned}$$

This completes the proof. \square

Due to a similar updating form to CSGD, we can apply a similar proof strategy with particular consideration to the perturbation of \mathbf{x} .

Theorem 3.2. Under Assumptions 1 (L -Lipschitz gradient assumption for $f(\cdot)$) and 4, and the commonly used bounded stochastic gradient assumption,

$$\mathbb{E} \left[\left\| f'(\mathbf{x}) - \frac{1}{N} \sum_{n=1}^N F'_n(\mathbf{x}_t; \xi_t^{(n)}) \right\|^2 \right] \leq \sigma^2,$$

choose the learning rate to be

$$\gamma = \left(2L + \sqrt{\frac{T}{N}}\sigma + T^{1/3}\sigma'^{2/3} \right)^{-1}.$$

If T is sufficiently large such that the learning rate satisfies the following condition

$$\gamma \leq \min \left\{ \frac{1}{4L}, \sqrt{\frac{N}{T}} \frac{1}{\sigma}, \frac{1}{T^{1/3}\sigma'^{2/3}} \right\},$$

then the EC-SGD algorithm admits the following convergence rate:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|f'(\mathbf{x}_t)\|^2 \lesssim \frac{1}{T} + \frac{\sigma}{\sqrt{TN}} + \frac{\sigma'^{2/3}}{T^{2/3}},$$

where $f(\mathbf{x}_0)$ and L are treated to be constants.

Proof. Denote by for short

$$\Delta_t := \frac{1}{N} \sum_{n=1}^N F'_n(\mathbf{x}_t; \xi_t^{(n)}) - f'(\mathbf{x}_t).$$

To apply the analysis strategy of GD, we first prove some preliminary results to estimate the perturbation. For Δ_t — the difference between the stochastic gradient and the true gradient — we have the following two properties:

$$\mathbb{E}[\Delta_t] = \frac{1}{N} \sum_{n=1}^N \left(f'(\mathbf{x}_t) - \mathbb{E} \left[F'_n(\mathbf{x}_t; \xi_t^{(n)}) \right] \right) = \mathbf{0}, \quad (3.14)$$

$$\begin{aligned} \mathbb{E}[\|\Delta_t\|^2] &= \frac{1}{N^2} \mathbb{E} \left\| \sum_{n=1}^N \left(f'(\mathbf{x}_t) - F'_n(\mathbf{x}_t; \xi_t^{(n)}) \right) \right\|^2 \\ &= \frac{1}{N^2} \sum_{n=1}^N \mathbb{E} \left\| f'(\mathbf{x}_t) - F'_n(\mathbf{x}_t; \xi_t^{(n)}) \right\|^2 \\ &\leq \frac{\sigma^2}{N} \quad (\text{from the bounded SG Assumption}). \end{aligned} \quad (3.15)$$

Then we estimate the upper bound of Ω_t , which measures the distance between \mathbf{x}_t and $\tilde{\mathbf{x}}_t$:

$$\begin{aligned} \mathbb{E}[\|\Omega_t\|^2] &= \mathbb{E} \left\| \boldsymbol{\delta}_t + \frac{1}{N} \sum_{n=1}^N \boldsymbol{\delta}_t^{(n)} \right\|^2 \\ &\leq 2\mathbb{E} \|\boldsymbol{\delta}_t\|^2 + 2\mathbb{E} \left\| \frac{1}{N} \sum_{n=1}^N \boldsymbol{\delta}_t^{(n)} \right\|^2 \\ &\leq 2\sigma'^2 + \frac{2}{N} \sum_{n=1}^N \mathbb{E} \left\| \boldsymbol{\delta}_t^{(n)} \right\|^2 \\ &\leq 4\sigma'^2. \end{aligned}$$

From the assumption that $f(\mathbf{x})$ has the L -Lipschitz gradient, we obtain

$$\begin{aligned} \mathbb{E} \|f'(\tilde{\mathbf{x}}_t) - f'(\mathbf{x}_t)\| &\leq L^2 \mathbb{E} \|\tilde{\mathbf{x}}_t - \mathbf{x}_t\|^2 \\ &= L^2 \gamma^2 \mathbb{E} \|\Omega_{t-1}\|^2 \\ &\leq 2L^2 \gamma^2 \sigma'^2. \end{aligned} \quad (3.16)$$

From Lemma (3.1), we have the updating rule

$$\tilde{\mathbf{x}}_{t+1} = \tilde{\mathbf{x}}_t - \gamma(f'(\mathbf{x}_t) + \Delta_t).$$

Next we can follow the standard analysis pipeline by considering $\mathbb{E}f(\tilde{\mathbf{x}}_{t+1}) - \mathbb{E}f(\tilde{\mathbf{x}}_t)$:

$$\begin{aligned} & \mathbb{E}f(\tilde{\mathbf{x}}_{t+1}) - \mathbb{E}f(\tilde{\mathbf{x}}_t) \\ & \leq \mathbb{E} \langle \tilde{\mathbf{x}}_{t+1} - \tilde{\mathbf{x}}_t, f'(\tilde{\mathbf{x}}_t) \rangle + \frac{L}{2} \mathbb{E} \|\tilde{\mathbf{x}}_{t+1} - \tilde{\mathbf{x}}_t\|^2 \\ & = -\gamma \mathbb{E} \langle f'(\mathbf{x}_t), f'(\tilde{\mathbf{x}}_t) \rangle + \gamma \mathbb{E} \langle \Delta_t, f'(\tilde{\mathbf{x}}_t) \rangle + \frac{L\gamma^2}{2} \mathbb{E} \|f'(\mathbf{x}_t) - \Delta_t\|^2 \\ & = -\gamma \mathbb{E} \langle f'(\mathbf{x}_t), f'(\tilde{\mathbf{x}}_t) \rangle + \frac{L\gamma^2}{2} \mathbb{E} \|f'(\mathbf{x}_t)\|^2 + \frac{L\gamma^2}{2} \mathbb{E} \|\Delta_t\|^2 \\ & \quad (\text{due to } \mathbb{E}\Delta_t = \mathbf{0} \text{ in (3.14)}) \\ & \leq -\gamma \mathbb{E} \langle f'(\mathbf{x}_t), f'(\tilde{\mathbf{x}}_t) \rangle + \frac{L\gamma^2}{2} \mathbb{E} \|f'(\mathbf{x}_t)\|^2 + \frac{L\gamma^2\sigma^2}{2N} \\ & \quad (\text{due to (3.15)}) \\ & = -\left(\gamma - \frac{L\gamma^2}{2}\right) \mathbb{E} \|f'(\mathbf{x}_t)\|^2 - \gamma \mathbb{E} \langle f'(\mathbf{x}_t), f'(\tilde{\mathbf{x}}_t) - f'(\mathbf{x}_t) \rangle \\ & \quad + \frac{L\gamma^2\sigma^2}{2N}. \end{aligned}$$

Next, using the following relaxation,

$$\begin{aligned} & |\mathbb{E} \langle f'(\mathbf{x}_t), f'(\tilde{\mathbf{x}}_t) - f'(\mathbf{x}_t) \rangle| \\ & \leq 2\mathbb{E} \|f'(\tilde{\mathbf{x}}_t) - f'(\mathbf{x}_t)\|^2 + \frac{1}{2} \mathbb{E} \|f'(\mathbf{x}_t)\|^2 \\ & \leq 2L^2 \mathbb{E} \|\tilde{\mathbf{x}}_t - \mathbf{x}_t\|^2 + \frac{1}{2} \mathbb{E} \|f'(\mathbf{x}_t)\|^2 \\ & \leq 2L^2\gamma^2 \mathbb{E} \|\Omega_t\|^2 + \frac{1}{2} \mathbb{E} \|f'(\mathbf{x}_t)\|^2 \end{aligned}$$

and (3.16) yields

$$\begin{aligned} & \mathbb{E}f(\tilde{\mathbf{x}}_{t+1}) - \mathbb{E}f(\tilde{\mathbf{x}}_t) \\ & \leq \left(-\frac{\gamma}{2} + \frac{L\gamma^2}{2}\right) \mathbb{E} \|\nabla f(\mathbf{x}_t)\|^2 + \frac{L\gamma^2\sigma^2}{2N} + 4L^2\sigma'^2\gamma^3. \end{aligned}$$

Summing up the inequality above from $t = 0$ to $t = T - 1$, we get

$$\begin{aligned} & \mathbb{E}f(\tilde{\mathbf{x}}_T) - \mathbb{E}f(\tilde{\mathbf{x}}_0) \\ & \leq -\left(\frac{\gamma}{2} - \frac{L\gamma^2}{2}\right) \sum_{t=0}^{T-1} \mathbb{E} \|f'(\mathbf{x}_t)\|^2 + \frac{L\gamma^2\sigma^2T}{2N} + 4L^2\sigma'^2\gamma^3T, \end{aligned}$$

which can be also written as

$$\begin{aligned} & \left(\frac{\gamma}{2} - \frac{L\gamma^2}{2}\right) \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|^2 \\ & \leq \mathbb{E}f(\tilde{\mathbf{x}}_0) - \mathbb{E}f(\tilde{\mathbf{x}}_T) + \frac{L\gamma^2\sigma^2T}{2N} + 4L^2\sigma'^2\gamma^3T \\ & \leq \mathbb{E}f(\mathbf{x}_0) - \mathbb{E}f(\mathbf{x}^*) + \frac{L\gamma^2\sigma^2T}{2N} + 4L^2\sigma'^2\gamma^3T. \end{aligned}$$

Since $\gamma \leq \frac{1}{4L}$, we can rewrite the above inequality as

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|^2 \lesssim \frac{1}{\gamma T} + \frac{\gamma\sigma^2}{N} + \sigma'^2\gamma^2.$$

Based on the choice of γ , one can verify that

$$\begin{aligned} \frac{1}{\gamma T} & \lesssim \frac{1}{T} + \frac{\sigma}{\sqrt{NT}} + \frac{\sigma'^{2/3}}{T^{2/3}} \\ \frac{\gamma\sigma^2}{N} & \lesssim \frac{\sigma}{\sqrt{NT}} \\ \sigma'^2\gamma^2 & \lesssim \frac{\sigma'^{2/3}}{T^{2/3}}, \end{aligned}$$

which implies the claim. \square

Comparing this approach to (3.6), we can now see one more advantage of EC-SGD over CSGD with respect to the convergence efficiency, in addition to the fact that EC-SGD does not require unbiased compression, as does CSGD.

Remark 3.1. A compression assumption that is probably more realistic than Assumption 4 could be

$$\mathbb{E} [\|Q(\mathbf{y}) - \mathbf{y}\|^2] \leq \alpha \|\mathbf{y}\|^2 \quad \forall \mathbf{y}.$$

However, this will involve more complicated analysis. Readers can refer to Liu *et al.* (2019). We can still see similar advantages over CSGD, even under this new assumption.

4

System Relaxation 2: Asynchronous Training

In implementations that we described in previous sections, the communications among all workers are perfectly synchronized — all workers conduct computation at the same time, and are all blocked until the communication among all machines is finished. However, when the number of machines in a distributed system grows, such a *synchronous* strategy might have several limitations. First, as the communication costs increase, the amount of time that each machine can spend on computation decreases, as they cannot conduct any computation during the communication phase. Second, if some machines are slower than other machines (i.e., there are stragglers), all machines need to wait for the slowest machine to finish computation.

In this section, we describe one system relaxation to accommodate these problems. In this relaxation, we remove the synchronization barriers among all the machines. This decreases the synchronization overhead, with the consequence that each machine now has access to a *staled* model.

4.1 System Implementation

There can be multiple ways of implementing an asynchronous communication strategy, each of which might have slight differences in their convergence behavior. In this section, we describe one of the simplest implementations, which is easiest in terms of both theoretical analysis and system implementation.

We focus on a scenario in which the system is implemented using a *single-server* parameter server.

1. The parameter server holds the global replica of the model.
2. Each worker w_i works in four phases. (1) At the beginning of each iteration, w_i asks the parameter server for the global replica of the model. (2) Upon receiving this model, the worker w_i uses it to compute the local gradient. (3) The worker then sends the local gradient to the parameter server, which then applies it to update the global model replica. In this step, we assume that the update of the global model is *atomic* — that is, the parameter won’t “mix” the updates from different workers and these updates won’t overwrite each other. (4) Worker w_i waits until the transfer is finished, and repeats from step (1) immediately.

In practice, the above process can be implemented in slightly different ways.

1. A worker w_i can start processing the next data point without waiting for the transfer to finish.
2. The parameter server does not need to conduct atomic updates on the global model replica.

Some of these implementations can make theoretical analysis more engaged. We choose to focus on the simple asynchronous communication scheme described above, as it is one of the simplest implementations that can demonstrate the advantage of asynchronous communication under our performance model.

Impact of Asynchronous Communication. Figure 4.1 and Figure 4.2 illustrate the impact of asynchronous communication, *under our simplified performance model for communication*. We see that, under our current performance model, asynchronous communication can make each iteration faster — instead of waiting for all workers to finish processing, a worker can start the next iteration of computation immediately after its communication with the parameter server is finished. On the other hand, we also observe that the model used by each machine for computation might be *staled*. Take, for example, the last model of worker 3 w_3 in Figure 4.2 — this model misses one update from w_1 because w_3 does not wait for w_1 before it starts the next iteration of computation.

As we can see from Figure 4.2, one potential system bottleneck is the parameter server: when the server is saturated in terms of network

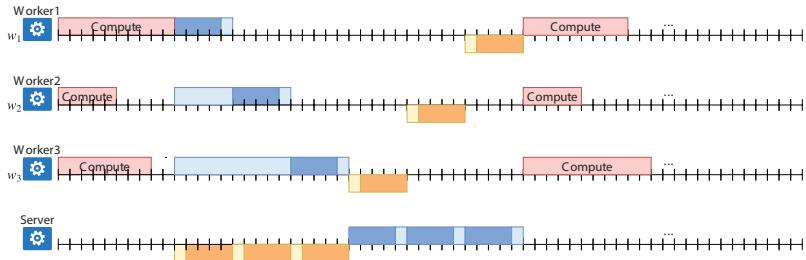


Figure 4.1: Illustration of the impact of asynchronous communication (without asynchrony).

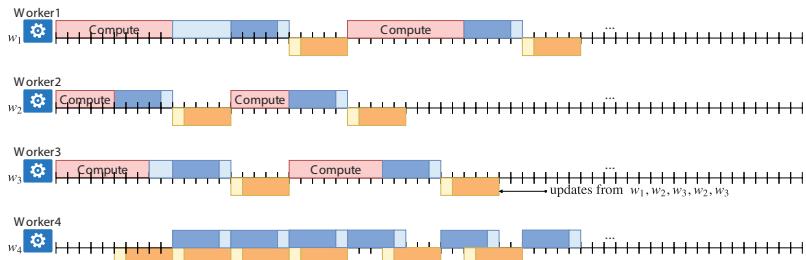


Figure 4.2: Illustration of the impact of asynchronous communication (with asynchrony).

bandwidth, adding more workers cannot make processing more data (i.e., calculating more gradients) faster without making the staleness larger. One way to accommodate this is to use multi-server parameter server architecture, in which each parameter server takes charge of one partition of the model. In this case, different partition of the model might have different *staleness*, which will, not surprisingly, make the theoretical analysis more engaged.

Caveats. In practice, there are other considerations that require careful attention when using asynchronous communication for distributed training. We briefly summarize these caveats and refer readers to Section 6 for further readings.

1. **Bounded Staleness.** If not careful, the above implementation does not necessarily provide bounded staleness — it is possible that one machine is always waiting for its communication to finish. In this case, some updates on the global replica might be using a gradient that is calculated based on a very old model, which could slow down the convergence. Some research have tried to accommodate this by enforcing bounded staleness in their implementation.
2. **Straggler.** When there are stragglers in the system (i.e., some machines are significantly slower than other machines), the staleness is systematic — one partition of the data is always more staled than another. In this case, staleness-aware/heterogeneous-aware algorithms could be designed to stabilize the convergence behavior.

One motivation for introducing asynchronous communication is to accommodate stragglers. There are other approaches to achieve the same objective without introducing asynchrony, however. For example, there has been research that attempts to use error correction code to make the synchronous approach more robust to stragglers. In this strategy, each worker conducts a small amount of redundant computation, such that the exact (or approximate) gradient can be recovered without waiting for the slowest k machines.

4.2 Theoretical Analysis

We now analyze **ASGD**, a SGD variant with asynchronous communications. The updating rule of **ASGD** can be cast into the following form:

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \gamma g_t(\mathbf{x}_{D(t)}),$$

where $D(t) \in \{1, 2, \dots, t\}$ and $\mathbf{x}_{D(t)}$ denote some early iterations of the model \mathbf{x}_t .

One important property for $g_t(\mathbf{x}_{D(t)})$ is

$$\mathbb{E}[g_t(\mathbf{x}_{D(t)})] = f'(\mathbf{x}_{D(t)}). \quad (4.1)$$

To ensure the convergence rate, let us make an additional assumption about staleness:

Assumption 5. We make the following assumption:

- (**Bounded staleness**) Assume that the staleness is bounded, that is,

$$0 \leq t - D(t) \leq \tau \quad \forall t,$$

where τ is a global staleness upper bound.

Intuitively, to ensure convergence, the staleness cannot be overlarge. In an extreme case, if $D(t) = 1 \forall t$, then the model is always updated from the true stochastic gradient in the first step, which has no hope of converging. Therefore, it is necessary to restrict the bound of $t - D(t)$. In practice, the staleness parameter $t - D(t)$ is proportional to the number of workers.

Next we start to show the convergence rate proof of **ASGD**, step by step.

$$f(\mathbf{x}_{t+1}) - f(\mathbf{x}_t) \leq -\gamma \langle f'(\mathbf{x}_t), g_t(\mathbf{x}_{D(t)}) \rangle + \frac{L\gamma^2}{2} \|g_t(\mathbf{x}_{D(t)})\|^2.$$

Take expectation on both sides:

$$\begin{aligned}\mathbb{E}[f(\mathbf{x}_{t+1})] - \mathbb{E}[f(\mathbf{x}_t)] &\leq -\gamma \mathbb{E}[\langle f'(\mathbf{x}_t), g_t(\mathbf{x}_{D(t)}) \rangle] \\ &\quad + \frac{L\gamma^2}{2} \mathbb{E}[\|g_t(\mathbf{x}_{D(t)})\|^2].\end{aligned}\quad (4.2)$$

We consider the items on the right-hand side respectively. For $\mathbb{E}[\|g_t(\mathbf{x}_{D(t)})\|^2]$, we have

$$\begin{aligned}\mathbb{E}_{D(t)}[\|g_t(\mathbf{x}_{D(t)})\|^2] &= \mathbb{E}_{D(t)}[\|\mathbf{g}_t(\mathbf{x}_{D(t)}) - f'(\mathbf{x}_{D(t)})\|^2] + \|\mathbf{f}_t(\mathbf{x}_{D(t)})\|^2 \\ &\leq \sigma^2 + \|\mathbf{f}_t(\mathbf{x}_{D(t)})\|^2,\end{aligned}\quad (4.3)$$

where the first equality is due to the fact (1.12), and the second is due to the bounded variance assumption, that is, Assumption 2.

For $\mathbb{E}[\langle f'(\mathbf{x}_t), g(\mathbf{x}_{D(t)}) \rangle]$, we have

$$\begin{aligned}\mathbb{E}[\langle f'(\mathbf{x}_t), g(\mathbf{x}_{D(t)}) \rangle] &= \mathbb{E}[\langle f'(\mathbf{x}_t), \mathbb{E}_{D(t)}[g(\mathbf{x}_{D(t)})] \rangle] \\ &= \mathbb{E}[\langle f'(\mathbf{x}_t), f'(\mathbf{x}_{D(t)}) \rangle] \\ &= \mathbb{E}\left[\frac{1}{2}\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2 - \frac{1}{2}\|f'(\mathbf{x}_t)\|^2 - \frac{1}{2}\|f'(\mathbf{x}_{D(t)})\|^2\right].\end{aligned}\quad (4.4)$$

The last equality uses an important property,

$$\langle \mathbf{a}, \mathbf{b} \rangle = \frac{1}{2}\|\mathbf{a} - \mathbf{b}\|^2 - \frac{1}{2}\|\mathbf{a}\|^2 - \frac{1}{2}\|\mathbf{b}\|^2,$$

which can be verified by a straightforward linear algebra computation. Next we take expectation on both sides of (4.2) and plug (4.3) and (4.4) into that:

$$\begin{aligned}\mathbb{E}[f(\mathbf{x}_{t+1})] - \mathbb{E}[f(\mathbf{x}_t)] &\leq -\frac{\gamma}{2} \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] - \frac{\gamma}{2} \mathbb{E}[\|f'(\mathbf{x}_{D(t)})\|^2] + \frac{\gamma}{2} \mathbb{E}[\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2] \\ &\quad + \frac{\gamma^2 L \sigma^2}{2} + \frac{\gamma^2 L}{2} \mathbb{E}[\|f'(\mathbf{x}_{D(t)})\|^2] \\ &\leq -\frac{\gamma}{2} \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] + \frac{\gamma^2 L \sigma^2}{2} - \frac{\gamma}{2}(1 - \gamma L) \mathbb{E}[\|f'(\mathbf{x}_{D(t)})\|^2] \\ &\quad + \frac{\gamma}{2} \mathbb{E}[\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2] \\ &\leq -\frac{\gamma}{2} \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] + \frac{\gamma^2 L \sigma^2}{2} + \frac{\gamma}{2} \mathbb{E}[\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2],\end{aligned}\quad (4.5)$$

where the last inequality is obtained by choosing a sufficiently small learning rate γ satisfying

$$1 \geq \gamma L. \quad (4.6)$$

We have seen the first two terms on the right-hand side of (4.5) (if ignoring the constant parts) in the proof to SGD. We can roughly treat this term as $\mathbb{E}\|f'(\mathbf{x}_t)\|^2$ plus some perturbation, depending on the staleness τ . The major term we need to treat very seriously is the last term $\mathbb{E}[\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2]$, whose upper bound is given by the following lemma. It basically shows that this item is bounded by a higher order of γ : $O(\gamma^3)$.

Lemma 4.1. Under Assumptions 1, 2, and 5, we have

$$\begin{aligned} \mathbb{E}[\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2] &\leq 2\gamma^2 L^2 \tau \sigma^2 \\ &\quad + 2\gamma^2 L^2 \tau \sum_{s=D(t)}^{t-1} \mathbb{E} \left[\|f'_s(\mathbf{x}_{D(t)})\|^2 \right]. \end{aligned}$$

Proof. We start by bounding the difference between $f'(\mathbf{x}_t)$ and $f'(\mathbf{x}_{D(t)})$

$$\begin{aligned} &\mathbb{E}[\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2] \\ &\leq L^2 \mathbb{E}[\|\mathbf{x}_t - \mathbf{x}_{t_{\tau_t}}\|^2] \quad (\text{Due to Assumption 1}) \\ &= \gamma^2 L^2 \mathbb{E} \left[\left\| \sum_{s=D(t)}^{t-1} g_s(\mathbf{x}_{D(s)}) \right\|^2 \right]. \\ &= \gamma^2 L^2 \mathbb{E} \left[\left\| \sum_{s=D(t)}^{t-1} g_s(\mathbf{x}_{D(s)}) - \sum_{s=D(t)}^{t-1} f'_s(\mathbf{x}_{D(s)}) + \sum_{s=D(t)}^{t-1} f'_s(\mathbf{x}_{D(s)}) \right\|^2 \right] \\ &\leq 2\gamma^2 L^2 \mathbb{E} \left[\left\| \sum_{s=D(t)}^{t-1} \underbrace{(g_s(\mathbf{x}_{D(s)}) - f'_s(\mathbf{x}_{D(s)}))}_{=: \xi_{s-D(t)}} \right\|^2 \right] \\ &\quad + 2\gamma^2 L^2 \mathbb{E} \left[\left\| \sum_{s=D(t)}^{t-1} f'_s(\mathbf{x}_{D(s)}) \right\|^2 \right], \end{aligned} \quad (4.7)$$

where the last inequality uses a variant of triangle inequality $\|\mathbf{a} + \mathbf{b}\|^2 \leq 2\|\mathbf{a}\|^2 + 2\|\mathbf{b}\|^2$. Note that the sequence of $\{\xi_0, \xi_1, \dots, \xi_{\tau_t - D(t) - 1}\}$ is a martingale sequence with

$$\xi_l := g_{s+l}(\mathbf{x}_{D(s+l)}) - f'_{s+l}(\mathbf{x}_{D(s+l)}).$$

A martingale sequence satisfies that

$$\mathbb{E}(\xi_{l+1} \mid \xi_0, \xi_1, \dots, \xi_l) = 0.$$

Therefore, it is easy to verify that

$$\begin{aligned} \mathbb{E} \left[\left\| \sum_{l=0}^{t-1-D(t)} \xi_l \right\|^2 \right] &= \mathbb{E} \left[\left\| \sum_{l=0}^{t-2-D(t)} \xi_l \right\|^2 \right] + \mathbb{E} [\|\xi_{t-1}\|^2] \\ &= \sum_{l=0}^{t-1-D(t)} \mathbb{E} [\|\xi_l\|^2]. \end{aligned}$$

In other words, we have

$$\mathbb{E} \left[\left\| \sum_{s=D(t)}^{t-1} (g_s(\mathbf{x}_{D(s)}) - f'_s(\mathbf{x}_{D(s)})) \right\|^2 \right] = \sum_{l=0}^{t-1-D(t)} \mathbb{E} [\|\xi_l\|^2] \leq \tau \sigma^2.$$

Applying this property to (4.7) yields

$$\begin{aligned} &\mathbb{E}[\|f'(\mathbf{x}_t) - f'(\mathbf{x}_{D(t)})\|^2] \\ &\leq 2\gamma^2 L^2 \tau \sigma^2 + 2\gamma^2 L^2 \mathbb{E} \left[\left\| \sum_{s=D(t)}^{t-1} f'_s(\mathbf{x}_{D(s)}) \right\|^2 \right] \\ &\leq 2\gamma^2 L^2 \tau \sigma^2 + 2\gamma^2 L^2 \tau \sum_{s=D(t)}^{t-1} \mathbb{E} \left[\|f'_s(\mathbf{x}_{D(s)})\|^2 \right], \end{aligned}$$

where the last inequality uses Assumption 5 the following property: for any vectors $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$, we have

$$\left\| \sum_{i=1}^n \mathbf{a}_i \right\|^2 \leq n \sum_{i=1}^n \|\mathbf{a}_i\|^2.$$

This completes the proof. \square

We choose the learning rate to be sufficiently small. In particular, let

$$\gamma L\tau \leq 1/2. \quad (4.8)$$

Using the following short notations:

$$\begin{aligned} a_t &:= \mathbb{E}[\|f'(\mathbf{x}_t)\|^2] \\ b_t &:= \mathbb{E}[f(\mathbf{x}_t)], \end{aligned}$$

and apply Lemma 5.4 to (4.5):

$$\begin{aligned} b_{t+1} - b_t &\leq -\frac{\gamma}{2}a_t + \frac{\gamma^2 L}{2}(1 + 2\gamma L\tau)\sigma^2 + \gamma^3 L^2 \tau \sum_{s=D(t)}^{t-1} a_s \\ &\leq -\frac{\gamma}{2}a_t + \frac{\gamma^2 L}{2}(1 + 2\gamma L\tau)\sigma^2 + \gamma^3 L^2 \tau \sum_{s=D(t)}^{t-1} a_s \\ &\leq -\frac{\gamma}{2}a_t + \gamma^2 L\sigma^2 + \frac{\gamma}{4\tau} \sum_{s=D(t)}^{t-1} a_s. \end{aligned} \quad (4.9)$$

Next we take sum of (4.9) over t from $t = 1$ to $t = T$

$$\begin{aligned} b_{T+1} - b_1 &\leq -\frac{\gamma}{2} \sum_{t=1}^T a_t + \gamma^2 L T \sigma^2 + \frac{\gamma}{4\tau} \sum_{t=1}^T \sum_{s=D(t)}^{t-1} a_s \\ &\leq -\frac{\gamma}{2} \sum_{t=1}^T a_t + \gamma^2 L T \sigma^2 + \frac{\gamma}{4} \sum_{t=1}^T a_t \\ &= -\frac{\gamma}{4} \sum_{t=1}^T a_t + \gamma^2 L T \sigma^2. \end{aligned}$$

Therefore, we have the following convergence rate:

$$\frac{1}{T} \sum_{t=1}^T a_t \leq \frac{4(b_1 - b_{T+1})}{\gamma T} + 4\gamma L\sigma^2.$$

If we treat $f(\mathbf{x}_1) - f^*$ to be constant, then we have

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|f'(\mathbf{x}_t)\|^2 \lesssim \frac{f(\mathbf{x}_1) - \mathbb{E}[f(\mathbf{x}_{T+1})]}{\gamma T} + \gamma L\sigma^2$$

$$\begin{aligned} &\lesssim \frac{f(\mathbf{x}_1) - f^*}{\gamma T} + \gamma L\sigma^2 \\ &\lesssim \frac{1}{\gamma T} + \gamma L\sigma^2. \end{aligned}$$

We choose the learning rate γ to be

$$\gamma = \frac{1}{L(\tau + 1) + \sqrt{TL}\sigma} \quad (4.10)$$

to satisfy the requirements of γ in (4.6) and (4.8), which leads to

$$\begin{aligned} \frac{1}{\gamma T} &\leq \frac{L(\tau + 1)}{T} + \frac{\sqrt{L}\sigma}{\sqrt{T}} \\ \gamma L\sigma^2 &\leq \frac{\sqrt{L}\sigma}{\sqrt{T}}. \end{aligned}$$

Therefore, we can summarize the convergence rate of ASGD in the following theorem.

Theorem 4.2. Choose the learning rate in (4.10) for ASGD. Under Assumptions 1, 2, and 5, ASGD admits the following convergence rate:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|f'(\mathbf{x}_t)\|^2 \lesssim \frac{L}{T} + \frac{\sqrt{L}\sigma}{\sqrt{T}} + \frac{L\tau}{T},$$

where we treat $f(\mathbf{x}_1) - f^*$ to be constant.

We highlight the following observations from this convergence analysis:

- (**Consistency to SGD**) If there is only one worker, then $\tau = 0$ and the ASGD algorithm reduces the SGD algorithm. We can see that the convergence rate in Theorem 4.2 is consistent with SGD.
- (**Linear speedup**) The additional term in the rate is the last term, as compared to SGD. It is caused by the asynchronous parallelism. Recall that $\tau()$ is proportional to the total number of workers. If $\tau()$ satisfies $\tau \leq \frac{\sqrt{T}\sigma}{\sqrt{L}}$, then the convergence efficiency

would not be affected by using asynchronous updating. As a result, linear speedup can be achieved. Here, the linear speedup is in the sense that the average computational complexity per worker is asymptotically the computational complexity using one single worker divided by N — the total number of workers.

5

System Relaxation 3: Decentralized Communication

Lossy communication compression is designed to alleviate system bottlenecks caused by network bandwidth. Another type of network bottleneck is caused by *latency*. In that case, when there are N workers, both the `AllReduce` and the multi-server parameter server have an $O(N)$ dependency on the network latency. The fundamental reason for this is that all these approaches insist that the information on each worker be propagated to all other workers in a single round of communication.

There are multiple standard ways to get rid of the latency bottleneck, e.g., using a textbook reduction tree. In this section, we describe an alternative approach to improving the latency overhead. We call this method decentralized communication. Specifically, we form a logical ring among N workers (all workers are still connected via the same “logical switch”). At every single iteration, a worker sends one message to the neighbor on its immediate left and one message to the neighbor on its immediate right. With this method, the latency overhead becomes $O(1)$. On the downside, however, the information on a single worker only reaches its two adjacent neighbors in one round of communication.

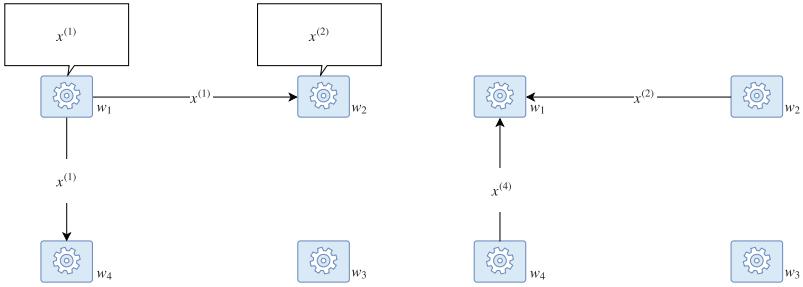


Figure 5.1: Illustration of decentralized communication pattern on the first worker \$w_1\$ (other workers are similar). Each worker holds a local model replica and only sends this model to the neighbors on its immediate right and left.

5.1 System Implementation

To implement this system relaxation, we rely on a communication pattern like that of the *model aggregation* approach for implementing distributed SGD instead of exchanging *gradients* as we did in the previous section on system exchange *models*, as illustrated in Figure 5.1. Specifically:

- At step \$t\$, each worker \$w_n\$ holds a model replica \$\mathbf{x}_t^{(n)}\$. The system first calculates the stochastic gradient using its local replica \$f'(\mathbf{x}^{(n)})\$ and then updates its local model using this standard SGD rule:

$$\mathbf{x}_{t+1/2}^{(n)} = \mathbf{x}_t^{(n)} - \gamma f'(\mathbf{x}^{(n)}).$$

- Each worker \$w_n\$ sends its locally updated model \$x_{t+1/2}^{(n)}\$ to the neighbor on its immediate right, \$w_{(n+1) \bmod N}\$, and to the neighbor on its immediate left, \$w_{(n-1) \bmod N}\$. Symmetrically, the worker will also receive models from the neighbor on its immediate right, \$w_{(n+1) \bmod N}\$, and from the neighbor on its immediate left, \$w_{(n-1) \bmod N}\$. Upon receiving the neighbors' models, the worker updates its local model as the average between its local model and its neighbors' models:

$$\mathbf{x}_{t+1}^{(n)} = \frac{1}{3} \left(\mathbf{x}_{t+1/2}^{(n)} + \mathbf{x}_{t+1/2}^{((n+1) \bmod N)} + \mathbf{x}_{t+1/2}^{((n-1) \bmod N)} \right).$$

Impact of Decentralized Communication. The communication cost for one round of communication in the decentralized setting is

$$2t_{\text{latency}} + 2t_{\text{transfer}}.$$

Recall that the communication cost of the multi-server parameter server or the `AllReduce` is

$$2(N - 1)t_{\text{latency}} + 2\frac{N - 1}{N}t_{\text{transfer}}.$$

With our method, we see a clear improvement with the decentralized communication in terms of communication latency. When the underlying network has a high latency, the decentralized strategy can be significantly faster.

Figure 5.2 and Figure 5.3 illustrate the communication patterns. Interestingly, in this specific example, the decentralized approach communication might seem to take longer. This happens because, when there are n workers, each worker only needs to send $\frac{N-1}{N}$ of its model in the centralized case instead of the full model. In the decentralized case, workers need to exchange their full models. As a result, the transfer time of the decentralized approach might be slightly higher. However, as illustrated in Figure 5.3, the decentralized approach has an advantage in terms of latency, and this improvement will be more significant when there are more workers and the underlying network has a higher latency.

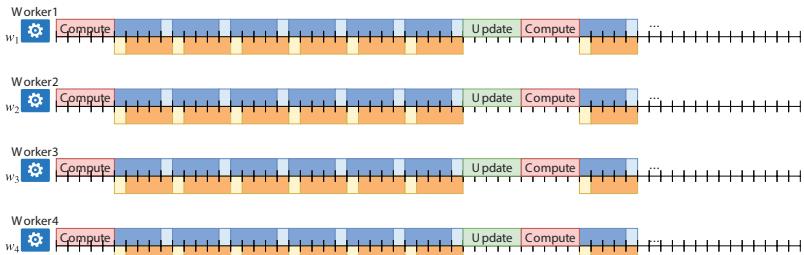


Figure 5.2: Illustration of the impact of decentralized communication (without decentralization).

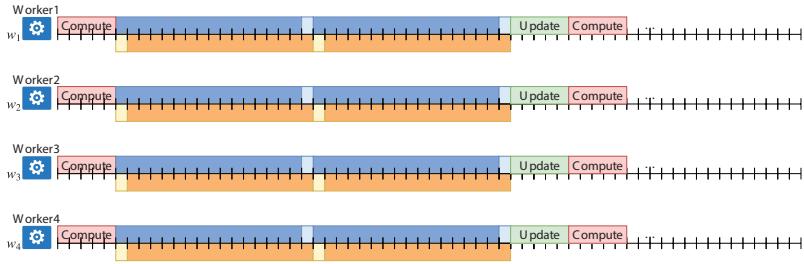


Figure 5.3: Illustration of the impact of decentralized communication (with decentralization).

More generally, one can extend the above example beyond a simple ring topology. We consider this below in the theoretical analysis.

5.2 Theoretical Analysis

We analyze DSGD, a SGD variant with decentralized communication. Let us consider the same objective as (3.7). For convenience of reference, we repeat it again here

$$\min_{\mathbf{x}} : \quad \left\{ f(\mathbf{x}) := \frac{1}{N} \sum_{n=1}^N f_n(\mathbf{x}) \right\} \quad (5.1)$$

where $f_n(\mathbf{x}) := \mathbb{E}_{\xi \sim \mathcal{D}_n} F_n(\mathbf{x}; \xi)$ and \mathcal{D}_n denotes the distribution of local data at node n . We do not assume that all nodes can access the whole dataset.

Based on the algorithm description in the previous subsection, the DSGD algorithm's updating rule can be cast in the following form:

$$\mathbf{X}_{t+1} = \mathbf{X}_t W - \gamma \mathbf{G} \left(\mathbf{X}_t; \{\xi_t^{(n)}\}_{n=1}^N \right). \quad (5.2)$$

Here we use the following notations:

$$\begin{aligned} \mathbf{X}_t &:= [\mathbf{x}_t^{(1)}, \mathbf{x}_t^{(2)}, \dots, \mathbf{x}_t^{(N)}] \\ \mathbf{G} \left(\mathbf{X}_t; \{\xi_t^{(n)}\}_{n=1}^N \right) &:= [F'_1 \left(\mathbf{x}_t^{(1)}; \xi_t^{(1)} \right), \dots, F'_N \left(\mathbf{x}_t^{(N)}; \xi_t^{(N)} \right)] \end{aligned}$$

where $\mathbf{x}_t^{(n)}$ denotes the local model on node n at time t . We also use \mathbf{G}_t to denote $\mathbf{G} \left(\mathbf{X}_t; \{\xi_t^{(n)}\}_{n=1}^N \right)$ for short. $W \in \mathbf{R}^{N \times N}$ is called the confusion matrix.

5.2.1 Assumptions

To show the convergence rate of DSGD, we need to make a few important assumptions about the objective function.

Assumption 6. We make the following assumptions.

- (**Smoothness and Lipschitzian gradient**). All functions $F_n(\cdot; \xi)$'s are smooth and all $f_n(\mathbf{x})$'s have an L -Lipschitzian gradient, that is, $\forall \mathbf{x}, \forall \mathbf{y} \forall n \in [N]$

$$\|f'_n(\mathbf{x}) - f'_n(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|.$$

- (**Unbiased sampling**). For all workers $n \in [N]$, the stochastic gradient is unbiased, that is,

$$\mathbb{E}_{\xi \sim \mathcal{D}_n}[F'_n(\mathbf{x}; \xi)] = f'_n(\mathbf{x}) \quad \forall \mathbf{x}.$$

- (**Bounded inner variance**). All local stochastic gradients have a bounded variance, that is,

$$\mathbb{E}_{\xi \sim \mathcal{D}_n} \|f'_n(\mathbf{x}) - F'_n(\mathbf{x}; \xi)\|^2 \leq \sigma^2 \quad \forall \mathbf{x}.$$

- (**Bounded outer variance**). The global gradient variance is bounded, that is,

$$\frac{1}{N} \sum_{n=1}^N \|f'_n(\mathbf{x}) - f'(\mathbf{x})\|^2 \leq \varsigma^2 \quad \forall \mathbf{x}.$$

The first assumption on the smoothness and the Lipschitz gradient is essentially the same as Assumption 1. The second assumption

is essentially that all local stochastic gradients are locally unbiased. The bounded inner variance assumption assumes that all local stochastic gradients have a bounded local variance that is similar to Assumption 2. The last assumption is that the total gradient difference among workers is bounded. If all workers access the same dataset, then $\varsigma = 0$.

We next make the necessary assumptions for the confusion matrix W .

Assumption 7. We also make the following assumptions for the confusion matrix W .

- (**Symmetric and doubly stochastic matrix**) W is a symmetric and doubly stochastic matrix, that is,

$$W^\top \mathbf{1} = \mathbf{1} \quad \text{and} \quad W = W^\top.$$

- (**Spectral gap**) denoted by ρ the second largest eigenvalue of W in term of the absolute value, that is,

$$\rho := \max_{n=2,3,\dots,N} |\lambda_n(W)|$$

where $\lambda_n(W)$ denotes the n largest eigenvalue of W . We assume that the spectral gap $1 - \rho$ is greater than 0.

$W^\top \mathbf{1} = \mathbf{1}$ is to ensure that the weighted sum is 1 when take the average of models, and the symmetry ensures that all eigenvalues of W are real numbers. The spectral gap $1 - \rho$ roughly measures how fast the information can be spread over the network. The smaller ρ is, the faster the communication is. Note that since W is doubly stochastic, the largest eigenvalue is always 1. Let us look at a few examples to get some sense of the value of ρ :

$$W_1 = \frac{\mathbf{1}\mathbf{1}^\top}{N}, \quad \rho = 0$$

$$W_2 = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \cdots & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & \cdots & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \cdots & 0 & \frac{1}{3} & \frac{1}{3} \end{bmatrix} \quad \rho \approx 1 - \frac{16\pi^2}{3N^2}$$

$$W_3 = \begin{bmatrix} \text{any doubly stochastic matrix} & \mathbf{0} \\ \mathbf{0}^\top & 1 \end{bmatrix} \quad \rho = 1.$$

W_1 (with $\rho = 1$) corresponds to the fully connected network, which is the fastest network for spreading information; W_2 (with a ρ value close to but smaller than 1) corresponds to the ring network; W_3 (with $\rho = 1$) corresponds to a disconnected network, which does work for DSGD.

The last assumption we make is about the initial local models, which we assume to be identical. This assumption is not necessary (we can obtain a similar result without it) but can simplify notations and derivations in the proof.

Assumption 8. We make the following assumption for the algorithm initialization.

- The initial values for all workers are identical, that is,

$$\mathbf{x}_1^{(1)} = \mathbf{x}_1^{(2)} = \cdots = \mathbf{x}_1^{(N)}.$$

5.2.2 Convergence Rate

The theoretical analysis for DSGD is relatively complicated. Readers can jump directly to Theorem 5.6 if they are not interested in the convergence proof.

To simplify the notations used in our proof, let us define some new notations for convenience. Denote by

$$\begin{aligned} f'(\mathbf{X}_t) &:= \left[f'_1(\mathbf{x}_t^{(1)}), f'_2(\mathbf{x}_t^{(2)}), \dots, f'_N(\mathbf{x}_t^{(N)}) \right] \\ \bar{f}'(\mathbf{X}_t) &:= f'(\mathbf{X}_t) \frac{\mathbf{1}}{N}, \\ \bar{\mathbf{x}}_t &:= \mathbf{X}_t \frac{\mathbf{1}}{N}, \\ \bar{\mathbf{X}}_t &:= \mathbf{X}_t \frac{\mathbf{1}\mathbf{1}^\top}{N}, \\ \bar{\mathbf{g}}_t &:= \mathbf{G}_t \frac{\mathbf{1}}{N}. \end{aligned}$$

Then it is not hard to see that

$$\mathbb{E} \left[\mathbf{G} \left(\mathbf{X}_t; \{\xi_t^{(n)}\}_{n=1}^N \right) \right] = f'(\mathbf{X}_t). \quad (5.3)$$

We next show some preliminary results that illustrate our final proof:

Lemma 5.1. Under Assumption 6, we have

$$\left\| \bar{f}'(\mathbf{X}_t) - f'(\bar{\mathbf{x}}_t) \right\|^2 \leq \frac{L^2}{N} \|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_F^2 \quad (5.4)$$

$$\left\| f'(\mathbf{X}_t) - f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 \leq L^2 \|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_F^2 + N\varsigma^2 \quad (5.5)$$

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{G}_t\|^2 \right] &\leq 2N \|f'(\bar{\mathbf{x}}_t)\|^2 \\ &\quad + 4L^2 \|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_F^2 + 4N\varsigma^2 + N\sigma^2. \end{aligned} \quad (5.6)$$

Proof. We first bound the difference between $\bar{f}'(\mathbf{X}_t)$ and $f'(\bar{\mathbf{x}}_t)$:

$$\begin{aligned} \left\| \bar{f}'(\mathbf{X}_t) - f'(\bar{\mathbf{x}}_t) \right\|^2 &= \left\| \frac{1}{N} \sum_{n=1}^N f'_n(\mathbf{x}_t^{(n)}) - \frac{1}{N} \sum_{n=1}^N f'_n(\bar{\mathbf{x}}_t) \right\|^2 \\ &= \frac{1}{N^2} \left\| \sum_{n=1}^N f'_n(\mathbf{x}_t^{(n)}) - \sum_{n=1}^N f'_n(\bar{\mathbf{x}}_t) \right\|^2 \end{aligned}$$

$$\begin{aligned}
& \leq \frac{1}{N} \sum_{n=1}^N \left\| f'_n(\mathbf{x}_t^{(n)}) - f'_n(\bar{\mathbf{x}}_t) \right\|^2 \\
& \leq \frac{L^2}{N} \sum_{n=1}^N \left\| \mathbf{x}_t^{(n)} - \bar{\mathbf{x}}_t \right\|^2 \\
& = \frac{L^2}{N} \left\| \mathbf{X}_t - \bar{\mathbf{X}}_t \right\|_F^2
\end{aligned}$$

where the first inequality uses the property

$$\left\| \sum_{n=1}^N \mathbf{z}^{(n)} \right\|^2 \leq N \sum_{n=1}^N \left\| \mathbf{z}^{(n)} \right\|^2.$$

Next, we prove the second inequality by using Assumption 6 (smoothness):

$$\begin{aligned}
& \frac{1}{2} \left\| f'(\mathbf{X}_t) - f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 \\
& = \frac{1}{2} \left\| f'(\mathbf{X}_t) - f'(\bar{\mathbf{X}}_t) + f'(\bar{\mathbf{X}}_t) - f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 \\
& \leq \left\| f'(\mathbf{X}_t) - f'(\bar{\mathbf{X}}_t) \right\|_F^2 + \left\| f'(\bar{\mathbf{X}}_t) - f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 \\
& \leq L^2 \left\| \mathbf{X}_t - \bar{\mathbf{X}}_t \right\|_F^2 + \left\| f'(\bar{\mathbf{X}}_t) - f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 \\
& = L^2 \left\| \mathbf{X}_t - \bar{\mathbf{X}}_t \right\|_F^2 + \sum_{n=1}^N \left\| f'_n(\bar{\mathbf{x}}_t) - f'(\bar{\mathbf{x}}_t) \right\|^2 \\
& \leq L^2 \left\| \mathbf{X}_t - \bar{\mathbf{X}}_t \right\|_F^2 + N\varsigma^2.
\end{aligned}$$

Next, we prove the last inequality:

$$\begin{aligned}
\mathbb{E} \left[\|\mathbf{G}_t\|^2 \right] & = \|\mathbb{E}[\mathbf{G}_t]\|_F^2 + \mathbb{E} \|\mathbf{G}_t - \mathbb{E}[\mathbf{G}_t]\|_F^2 \\
& \leq \|f'(\mathbf{X}_t)\|_F^2 + N\sigma^2 \\
& \leq 2 \left\| f'(\mathbf{X}_t) - f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 + 2 \left\| f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 + N\sigma^2 \\
& = 2 \left\| f'(\mathbf{X}_t) - f'(\bar{\mathbf{x}}_t) \mathbf{1}^\top \right\|_F^2 + 2N \|f'(\bar{\mathbf{x}}_t)\|^2 + N\sigma^2.
\end{aligned}$$

□

Next, we prove a useful inequality.

Lemma 5.2. Given two non-negative sequences $\{a_t\}_{t=1}^{\infty}$ and $\{b_t\}_{t=1}^{\infty}$ that satisfying

$$a_t = \sum_{s=1}^t \rho^{t-s} b_s, \quad (5.7)$$

with $\rho \in [0, 1]$, we have

$$\sum_{t=1}^k a_t^2 \leq \frac{1}{(1-\rho)^2} \sum_{s=1}^k b_s^2.$$

Proof. Consider the left-hand side:

$$\begin{aligned} \sum_{t=1}^k a_t^2 &= \sum_{t=1}^k \sum_{s=1}^t \rho^{t-s} b_s \sum_{r=1}^t \rho^{t-r} b_r \\ &= \sum_{t=1}^k \sum_{s=1}^t \sum_{r=1}^t \rho^{2t-s-r} b_s b_r \\ &\leq \sum_{t=1}^k \sum_{s=1}^t \sum_{r=1}^t \rho^{2t-s-r} \frac{b_s^2 + b_r^2}{2} \\ &= \sum_{t=1}^k \sum_{s=1}^t \sum_{r=1}^t \rho^{2t-s-r} b_s^2 \\ &\leq \frac{1}{1-\rho} \sum_{t=1}^k \sum_{s=1}^t \rho^{t-s} b_s^2 \\ &\leq \frac{1}{(1-\rho)^2} \sum_{s=1}^k b_s^2. \end{aligned}$$

This completes the proof. \square

Given these preliminary results, we can show the first inequality for the upper bound of the total model variation, that is, the distance between the local models \mathbf{X}_t and the average model $\bar{\mathbf{X}}_t$. In attempting to understand the decentralized algorithm, the first question people may ask is whether all local models will achieve consensus

eventually as the centralized counterpart always satisfies the consensus over all iterations. Lemma 5.4 essentially shows that DSGD can ensure that all $\mathbf{x}_t^{(n)}$'s converge to $\bar{\mathbf{x}}_t$. This happens because the total variation $\sum_{t=1}^T \mathbb{E} [\|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2]$ is roughly bounded by $O(\gamma^2 T)$. Based on our experiences of analyzing stochastic algorithms, we know that the learning rate is usually proportional to $1/\sqrt{T}$. Therefore, the total variation is bounded by a constant, which indicates that all $\mathbf{x}_t^{(n)}$ converges to $\bar{\mathbf{x}}_t$.

Lemma 5.3. Under Assumptions 7 and 8, we have

$$\sum_{t=1}^T \mathbb{E} [\|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2] \leq \sum_{t=1}^T \frac{\gamma^2}{(1-\rho)^2} \mathbb{E} [\|\mathbf{G}_t\|^2].$$

Proof. From the updating rule, we obtain the following closed form for \mathbf{X}_t and $\bar{\mathbf{X}}_t$:

$$\begin{aligned}\mathbf{X}_t &= \gamma \sum_{s=1}^{t-1} \mathbf{G}_s W^{t-s-1} \\ \bar{\mathbf{X}}_t &= \mathbf{X}_t \frac{\mathbf{1}\mathbf{1}^\top}{N}.\end{aligned}$$

Next, we bound the difference between $\bar{\mathbf{X}}_t$ and \mathbf{X}_t as follows:

$$\begin{aligned}&\frac{1}{\gamma} \|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_F \\&= \frac{1}{\gamma} \left\| \mathbf{X}_t \left(I - \frac{\mathbf{1}\mathbf{1}^\top}{N} \right) \right\|_F \\&= \left\| \sum_{s=1}^{t-1} \mathbf{G}_s W^{t-s-1} \left(I - \frac{\mathbf{1}\mathbf{1}^\top}{N} \right) \right\|_F \\&= \left\| \sum_{s=1}^{t-1} \mathbf{G}_s \left(W^{t-s-1} - \frac{\mathbf{1}\mathbf{1}^\top}{N} \right) \right\|_F \quad (\text{due to } W\mathbf{1} = \mathbf{1}) \\&\leq \sum_{s=1}^{t-1} \left\| \mathbf{G}_s \left(W^{t-s-1} - \frac{\mathbf{1}\mathbf{1}^\top}{N} \right) \right\|_F \quad (\text{due to the triangle inequality})\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{s=1}^{t-1} \|\mathbf{G}_s\|_F \left\| \left(W^{t-s-1} - \frac{\mathbf{1}\mathbf{1}^\top}{N} \right) \right\| \quad (\text{due to } \|AB\|_F \leq \|A\|_F \|B\|) \\
&\leq \sum_{s=1}^{t-1} \rho^{t-s-1} \|\mathbf{G}_s\|_F.
\end{aligned}$$

Taking the square and the sum over t , the equation above becomes

$$\begin{aligned}
&\sum_{t=1}^T \frac{1}{\gamma^2} \|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_F^2 \\
&\leq \sum_{t=1}^T \left(\sum_{s=1}^{t-1} \rho^{t-s-1} \|\mathbf{G}_s\|_F \right)^2 \\
&\leq \frac{1}{(1-\rho)^2} \sum_{t=1}^T \|\mathbf{G}_t\|_F^2 \quad (\text{due to Lemma 5.2}).
\end{aligned}$$

This completes the proof. \square

Lemma 5.4. Under Assumptions 1, 2, and 5, if the learning rate is chosen to satisfy

$$\gamma \leq \frac{1-\rho}{4L}$$

we have

$$\begin{aligned}
\sum_{t=1}^T \mathbb{E} [\|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2] &\leq \sum_{t=1}^T \frac{8N(\sigma^2 + \varsigma^2)\gamma^2}{(1-\rho)^2} \\
&\quad + \sum_{t=1}^T \frac{2N\gamma^2}{(1-\rho)^2} \mathbb{E} [\|f'(\bar{\mathbf{x}}_t)\|^2].
\end{aligned}$$

Proof. Applying Lemma 5.1 to Lemma 5.3 gives

$$\begin{aligned}
\frac{(1-\rho)^2}{\gamma^2} \mathbb{E} [\|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2] &\leq 2N\mathbb{E} [\|f'(\bar{\mathbf{x}}_t)\|^2] \\
&\quad + 4L^2\mathbb{E} [\|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2] + 4N\varsigma^2 + N\sigma^2.
\end{aligned}$$

Given the restriction on the learning rate γ , we have $4L^2 \leq (1-\rho)^2/(2\gamma^2)$ and thus obtain

$$\begin{aligned} \frac{(1-\rho)^2}{\gamma^2} \mathbb{E} [\|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2] &\leq 2N \mathbb{E} [\|f'(\bar{\mathbf{x}}_t)\|^2] \\ &\quad + 8N(\varsigma^2 + \sigma^2), \end{aligned}$$

which completes the proof. \square

Lemma 5.5. Under Assumption 6, if $\gamma \leq \frac{1}{L}$, we have

$$\begin{aligned} \mathbb{E} [\|f'(\bar{\mathbf{x}}_t)\|^2] &\leq \frac{2}{\gamma} \mathbb{E} [f(\bar{\mathbf{x}}_t)] - \mathbb{E} [f(\bar{\mathbf{x}}_{t+1})] \\ &\quad + \frac{L^2}{N} \mathbb{E} [\|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2] + \frac{L\gamma\sigma^2}{N}. \end{aligned}$$

Proof. Let us start with a basic property:

$$\mathbb{E}[\bar{\mathbf{g}}_t] = \frac{1}{N} \sum_{n=1}^N f'_n(\mathbf{x}_t^{(n)}) =: \bar{f}'(\mathbf{X}_t). \quad (5.8)$$

We consider the improvement of $\mathbb{E} [f(\bar{\mathbf{x}}_{t+1})]$ over $\mathbb{E} [f(\bar{\mathbf{x}}_t)]$ in the following:

$$\begin{aligned} &\mathbb{E} [f(\bar{\mathbf{x}}_{t+1})] - \mathbb{E} [f(\bar{\mathbf{x}}_t)] \\ &\leq \mathbb{E} \langle f'(\bar{\mathbf{x}}_t), -\gamma \bar{\mathbf{g}}_t \rangle + \frac{L\gamma^2}{2} \mathbb{E} [\|\bar{\mathbf{g}}_t\|^2] \\ &= -\gamma \mathbb{E} \langle f'(\bar{\mathbf{x}}_t), \bar{f}'(\mathbf{X}_t) \rangle + \frac{L\gamma^2}{2} \mathbb{E} [\|\bar{\mathbf{g}}_t\|^2] \\ &= -\frac{\gamma}{2} \mathbb{E} \|f'(\bar{\mathbf{x}}_t)\|^2 - \frac{\gamma}{2} \mathbb{E} \|\bar{f}'(\mathbf{X}_t)\|^2 + \frac{\gamma}{2} \mathbb{E} \|\bar{f}'(\mathbf{X}_t) - f'(\bar{\mathbf{x}}_t)\|^2 \\ &\quad + \frac{L\gamma^2}{2} \mathbb{E} [\|\bar{\mathbf{g}}_t\|^2]. \end{aligned} \quad (5.9)$$

We look at the last term first in the above inequality:

$$\begin{aligned} \mathbb{E} [\|\bar{\mathbf{g}}_t\|^2] &= \|\mathbb{E}[\bar{\mathbf{g}}_t]\|^2 + \mathbb{E} [\|\bar{\mathbf{g}}_t - \mathbb{E}[\bar{\mathbf{g}}_t]\|^2] \\ &\leq \|\bar{f}'(\mathbf{X}_t)\|^2 + \frac{\sigma^2}{N} \end{aligned} \quad (5.10)$$

where the inequality uses the property of $\bar{\mathbf{g}}_t$ in (5.8) and the property for any i.i.d. random variables $z^{(1)}, \dots, z^{(n)}$

$$\mathbf{Var} \left(\frac{1}{N} \sum_{n=1}^N z^{(n)} \right) = \frac{1}{N} \mathbf{Var}(z^{(n)}),$$

together with the boundedness in Assumption 6.

Plugging (5.10) and (5.4) into (5.9) yields

$$\begin{aligned} & \mathbb{E}[f(\bar{\mathbf{x}}_{t+1})] - \mathbb{E}[f(\bar{\mathbf{x}}_t)] \\ & \leq -\frac{\gamma}{2} \mathbb{E}\|f'(\bar{\mathbf{x}}_t)\|^2 - \left(\frac{\gamma}{2} - \frac{L\gamma^2}{2} \right) \mathbb{E}\|\bar{f}'(\mathbf{X}_t)\|^2 \\ & \quad + \frac{L^2\gamma}{2N} \mathbb{E}\|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_F^2 + \frac{L\gamma^2\sigma^2}{2N} \\ & \leq -\frac{\gamma}{2} \mathbb{E}\|f'(\bar{\mathbf{x}}_t)\|^2 + \frac{L^2\gamma}{2N} \mathbb{E}\|\mathbf{X}_t - \bar{\mathbf{X}}_t\|_F^2 + \frac{L\gamma^2\sigma^2}{2N}. \end{aligned}$$

This completes the proof. \square

Now we are ready to show the main result.

Theorem 5.6. Under Assumptions 6 and 7, we choose the learning rate to be

$$\gamma = \left(1 + \sqrt{TN}\sigma + T^{1/3}\varsigma^{2/3}(1-\rho)^{-2/3} \right)^{-1}.$$

If T is sufficiently large such that the learning rate satisfies the following condition

$$\gamma \leq \min \left\{ \frac{1-\rho}{4L}, \frac{(1-\rho)^2 N}{L} \right\}, \quad (5.11)$$

then the DSGD algorithm admits the following convergence rate:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}\|f'(\bar{\mathbf{x}}_t)\|^2 \lesssim \frac{1}{T} + \frac{\sigma}{\sqrt{NT}} + \left(\frac{\varsigma}{T(1-\rho)} \right)^{2/3}.$$

Proof. Taking the summarization over t from $t = 1$ to $t = T$ for Lemma 5.5 yields

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|f'(\bar{\mathbf{x}}_t)\|^2 \\ & \leq \frac{2}{T\gamma} (f(\bar{\mathbf{x}}_0) - f^\star) + \frac{L\sigma^2\gamma}{N} + \frac{L^2}{NT} \sum_{t=1}^T \|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2 \\ & \leq \frac{2}{T\gamma} (f(\bar{\mathbf{x}}_0) - f^\star) + \frac{L\sigma^2\gamma}{N} \\ & \quad + \frac{8L^2\varsigma^2\gamma^2}{(1-\rho)^2} + \frac{8L^2\sigma^2\gamma^2}{(1-\rho)^2} + \frac{2L^2\gamma^2}{(1-\rho)^2 T} \sum_{t=1}^T \mathbb{E} \|f'(\bar{\mathbf{x}}_t)\|^2. \end{aligned}$$

From the restriction on the learning rate, we have

$$\frac{8L^2\sigma^2\gamma^2}{(1-\rho)^2} \leq \frac{L\sigma^2\gamma}{N}$$

and

$$\frac{2L^2\gamma^2}{(1-\rho)^2} \leq \frac{1}{2}.$$

It follows that

$$\begin{aligned} & \frac{1}{2T} \sum_{t=1}^T \mathbb{E} \|f'(\bar{\mathbf{x}}_t)\|^2 \\ & \leq \frac{2}{T\gamma} (f(\bar{\mathbf{x}}_0) - f^\star) + \frac{L\sigma^2\gamma}{N} + \frac{L^2}{NT} \sum_{t=1}^T \|\bar{\mathbf{X}}_t - \mathbf{X}_t\|_F^2 \\ & \leq \frac{2}{T\gamma} (f(\bar{\mathbf{x}}_0) - f^\star) + \frac{2L\sigma^2\gamma}{N} + \frac{8L^2\varsigma^2\gamma^2}{(1-\rho)^2}. \end{aligned}$$

For simplicity, we treat $f(\bar{\mathbf{x}}_0) - f^\star$ and L as constants and obtain the following simplified inequality:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|f'(\bar{\mathbf{x}}_t)\|^2 \lesssim \frac{1}{T\gamma} + \frac{\sigma^2\gamma}{N} + \frac{\varsigma^2\gamma^2}{(1-\rho)^2}.$$

Due to the form of the learning rate in (5.11), we have

$$\frac{1}{T\gamma} \leq \frac{1}{T} + \frac{\sigma}{\sqrt{NT}} + \frac{\varsigma^{2/3}}{T^{2/3}(1-\rho)^{2/3}}$$

$$\frac{\sigma^2 \gamma}{N} \leq \frac{\sigma}{\sqrt{NT}}$$

$$\frac{\varsigma^2 \gamma^2}{1 - \rho} \leq \frac{\varsigma^{2/3}}{T^{2/3}(1 - \rho)^{2/3}},$$

which completes the proof. \square

We highlight the following observations from this convergence analysis.

- **(Consistency to mb-SGD)** If we use the fully connected network, then $\rho = 1$ and the DSGD algorithm reduces the (centralized) mb-SGD algorithm. We can see that the convergence rate in Theorem (5.6) is consistent with mb-SGD, since the last term becomes zero.
- **(Linear speedup)** The additional term in the rate is the last term, comparing to mb-SGD. It is caused by using decentralized updating. As long as T is sufficiently large, then the last term will be dominated by the second term. As a result, the linear speedup can be achieved.

6

Further Reading

Developing efficient distributed learning systems is an emerging topic that has received intensive interests in recent years. The goal of this monograph is by no means to provide a complete summary of the recent development in this area — instead, our goal is merely to provide an “overly simplified” overview.

In this section, we assemble a *best-effort* reading list to provide readers pointers for further reading. This list is incomplete — it is more like a set of “*pointers to pointers*” whose goal is to provide a bird-eye view on the trend of research and to provide a starting point for readers to start their own navigation. To assemble this list, we went through *most* papers published in ICML, N(eur)IPS, VLDB, SIGMOD and SysML since 2015 (up to 2019), and *tried our best* to summarize relevant papers into multiple categories. This selection method means that this list inevitably misses many early seminal work on this topic. However, we believe that the union of all papers cite by papers in this list should provide a reasonable coverage.

6.1 Communication and Data Compression

Data movement during training can be one of the largest system bottleneck, especially when there are many workers in the system or the computation device is significantly more powerful than the peak throughput of data movements. One collection of work focus on optimizing data movements via compression, and popular compression strategy includes quantization, sparsification, sketching, and other noise-corrupted transformation. Some examples of recent work in this direction include Acharya *et al.* (2019), Zhu and Lafferty (2018), Wu *et al.* (2018), Cohen *et al.* (2018), Bernstein *et al.* (2018), Zhang *et al.* (2017), Chen *et al.* (2017), Zhang *et al.* (2016), Gupta *et al.* (2015), Zhu and Gu (2015), Jiang *et al.* (2018), Elgohary *et al.* (2016), Kara *et al.* (2018), Wang *et al.* (2019), Lim *et al.* (2019), Wang *et al.* (2018a), Wang *et al.* (2018c), Agarwal *et al.* (2018), Alistarh *et al.* (2018b), Banner *et al.* (2018), Yu *et al.* (2018), Stich *et al.* (2018), Jiang and Agrawal (2018), Alistarh *et al.* (2017), Wen *et al.* (2017), and De Sa *et al.* (2015).

6.2 Decentralization and Approximate Averaging

Another line of work focuses on the scenario in which calculating the exact average among all workers is difficult in a single round of communication. One example is in peer-to-peer networks in which each worker can only communicate with its neighbor. One collection of work focus on analyzing the system behavior and designing novel algorithms in such a scenario. Some examples of recent work in this direction include Yu *et al.* (2019), Assran *et al.* (2019), Li *et al.* (2019), Li and Yan (2019), Tang *et al.* (2018b), Li *et al.* (2019), He *et al.* (2018), Lian *et al.* (2017), Li and Yan (2017), Simonetto *et al.* (2017), Colin *et al.* (2016), Yuan *et al.* (2016), Shi *et al.* (2015,b), Ling and Ribeiro (2013), and Nedić and Ozdaglar (2009). It is also worth pointing out that the researchers in federated learning also borrow the idea of decentralization for the data privacy purpose (He *et al.*, 2019).

There are also works which try to combine both communication compression and decentralization, for example Beznosikov *et al.* (2020), Tang *et al.* (2019a), Koloskova *et al.* (2019), and Tang *et al.* (2018a).

6.3 Asynchronous Communication

The synchronization barrier among all workers is often a system bottleneck, especially when there are many workers or there are stragglers (i.e., some workers are slower than other workers). One collection of work focus on removing the synchronization barrier by allowing the workers to proceed in an *asynchronous* fashion. Some examples of recent work in this direction include Zhou *et al.* (2018d), Simsekli *et al.* (2018), Nguyen *et al.* (2018), Gu *et al.* (2018), Zheng *et al.* (2017), Peng *et al.* (2017), Aybat *et al.* (2015), Hsieh *et al.* (2015), Jiang *et al.* (2017), Wangni *et al.* (2018), Sun *et al.* (2017), You *et al.* (2016), Chen *et al.* (2016), Pan *et al.* (2016), Lian *et al.* (2015), Chaturapruek *et al.* (2015), Liu and Wright (2015), Liu *et al.* (2014), Sridhar *et al.* (2013), Agarwal and Duchi (2012), and Lian *et al.* (2016).

There are also works which try to combine both decentralization and asynchronous communication, for example Lian *et al.* (2019).

6.4 Optimizing for Communication Rounds

There is a collection of work that tries to optimize for the number of communication rounds during training — instead of communicating every iteration, these methods allow each worker to run longer *locally* for multiple iterations. Some examples of recent work in this direction include Garber *et al.* (2017), Ma *et al.* (2015), Wang *et al.* (2017), Kamp *et al.* (2017), and Zhang *et al.* (2015).

6.5 System Optimization and Automatic Tradeoff Management

On the system side, one line of work is to further optimize the system communication primitives and communication strategies to take advantage of the property of the underlying ML workload. Some examples of recent work in this direction include Cho *et al.* (2019), Hashemi *et al.* (2019), Jayarajan *et al.* (2019), Jia *et al.* (2019), and Wang *et al.* (2018d). Another line of work tries to automatically optimize in the tradeoff introduced by these system relaxation techniques (e.g., the communication frequency which is often a hyperparameter). Some examples

of recent work in this direction include Dünner *et al.* (2018), Kaoudi *et al.* (2017), Li *et al.* (2018), Mahajan *et al.* (2018), Xin *et al.* (2018), and Wang and Joshi (2019).

6.6 Other Topics

One line of work focuses on designing variance reduction techniques for stochastic first-order methods, e.g., Ji *et al.* (2019), Horváth and Richtarik (2019), Zou *et al.* (2018), Zhou *et al.* (2018,c), Hazan and Luo (2016), Li *et al.* (2016), Allen-Zhu and Yuan (2016), Allen-Zhu and Hazan (2016), Reddi *et al.* (2016), Jothimurugesan *et al.* (2018), Zhou *et al.* (2018a), Liu *et al.* (2018), Hanzely *et al.* (2018), Arjevani (2017), Bietti and Mairal (2017), Palaniappan and Bach (2016), and Hofmann *et al.* (2015). One interesting way of achieving this is to change the sample distribution during training (i.e., sample more informative data points more frequently than non-informative data points), e.g., Katharopoulos and Fleuret (2018), Namkoong *et al.* (2017), Gopal (2016), Zhao and Zhang (2015), Johnson and Guestrin (2018), Cutkosky and Busa-Fekete (2018), Zhu (2016), and Qu *et al.* (2015).

Another line of work focuses on designing distributed learning algorithms that is robust to failures, e.g., by making the system Byzantine-resilient or tolerant to stragglers via Gradient coding. Examples of recent work in this direction include Xie *et al.* (2019), Yin *et al.* (2018), Damaskinos *et al.* (2018), Chen *et al.* (2018), Ye and Abbe (2018), Raviv *et al.* (2018), Tandon *et al.* (2017), Alistarh *et al.* (2018a), Karakus *et al.* (2017), and Blanchard *et al.* (2017).

Many, if not most, theoretical analysis of system relaxations of distributed learning systems assumes that the system samples data points with replacement. However, this is different from how most learning systems are implemented in practice. There have been efforts in trying to close this gap by analyzing different strategies of scan order. Examples of recent work in this direction include Haochen and Sra (2019), Nagaraj *et al.* (2019), and Shamir (2016).

In this monograph, we mainly focus on techniques for distributed learning that optimize for the communication among workers. However, there are other, orthogonal directions, in optimizing for the performance

and scalability of distributed learning systems. For example, when one focuses on distributed learning on the edge (Zhang *et al.*, 2018) or in geo-distributed setting (Hsieh *et al.*, 2017), additional considerations are often necessary. Another line of work tries to further take advantage of specific structures of the given task, e.g., deep neural networks. One example is a very interesting line of work that uses large batch size for training deep learning models (Ghadimi *et al.*, 2013; Goyal *et al.*, 2017; You *et al.*, 2020). For more techniques in this direction, we refer the reader to a comprehensive survey paper Ben-Nun and Hoefer (2018) on this topic.

References

- Acharya, J., C. De Sa, D. Foster, and K. Sridharan (2019). “Distributed learning with sublinear communication”. *International Conference on Machine Learning*.
- Agarwal, A. and J. C. Duchi (2012). “Distributed delayed stochastic optimization”. *IEEE Conference on Decision and Control*.
- Agarwal, N., A. T. Suresh, F. X. X. Yu, S. Kumar, and B. McMahan (2018). “cpSGD: Communication-efficient and differentially-private distributed SGD”. *Annual Conference on Neural Information Processing Systems*. 7575–7586.
- Alistarh, D., Z. Allen-Zhu, and J. Li (2018a). “Byzantine stochastic gradient descent”. *Annual Conference on Neural Information Processing Systems*. 4618–4628.
- Alistarh, D., D. Grubic, J. Li, R. Tomioka, and M. Vojnovic (2017). “QSGD: Communication-efficient SGD via gradient quantization and encoding”. *Annual Conference on Neural Information Processing Systems*. 1707–1718.
- Alistarh, D., T. Hoefer, M. Johansson, N. Konstantinov, S. Khirirat, and C. Renggli (2018b). “The convergence of sparsified gradient methods”. *Annual Conference on Neural Information Processing Systems*. 5977–5987.

- Allen-Zhu, Z. and E. Hazan (2016). “Variance reduction for faster non-convex optimization”. *International Conference on Machine Learning*.
- Allen-Zhu, Z. and Y. Yuan (2016). “Improved SVRG for non-strongly-convex or sum-of-non-convex objectives”. *International Conference on Machine Learning*.
- Arjevani, Y. (2017). “Limitations on variance-reduction and acceleration schemes for finite sums optimization”. *Annual Conference on Neural Information Processing Systems*.
- Assran, M., N. Loizou, N. Ballas, and M. Rabbat (2019). “Stochastic gradient push for distributed deep learning”. *International Conference on Machine Learning*.
- Aybat, N., Z. Wang, and G. Iyengar (2015). “An asynchronous distributed proximal gradient method for composite convex optimization”. *International Conference on Machine Learning*.
- Banner, R., I. Hubara, E. Hoffer, and D. Soudry (2018). “Scalable methods for 8-bit training of neural networks”. *Annual Conference on Neural Information Processing Systems*. 5151–5159.
- Ben-Nun, T. and T. Hoefer (2018). “Demystifying parallel and distributed deep learning: An in-depth concurrency analysis”. *arXiv:1802.09941*.
- Bernstein, J., Y.-X. Wang, K. Azizzadenesheli, and A. Anandkumar (2018). “signSGD: Compressed optimisation for non-convex problems”. *International Conference on Machine Learning*.
- Beznosikov, A., S. Horváth, P. Richtárik, and M. Safaryan (2020). “On biased compression for distributed learning”. *arXiv:2002.12410*.
- Biotti, A. and J. Mairal (2017). “Stochastic optimization with variance reduction for infinite datasets with finite sum structure”. *Annual Conference on Neural Information Processing Systems*. 1622–1632.
- Blanchard, P., E. M. El Mhamdi, R. Guerraoui, and J. Stainer (2017). “Machine learning with adversaries: Byzantine tolerant gradient descent”. *Annual Conference on Neural Information Processing Systems*. 118–128.
- Bottou, L., F. E. Curtis, and J. Nocedal (2016). “Optimization methods for large-scale machine learning”. *arXiv:1606.04838*.

- Boyd, S. and L. Vandenberghe (2004). *Convex Optimization*. Cambridge University Press.
- Chaturapruk, S., J. C. Duchi, and C. Ré (2015). “Asynchronous stochastic convex optimization: The noise is in the noise and SGD don’t care”. *Annual Conference on Neural Information Processing Systems*. 1531–1539.
- Chen, C., N. Ding, C. Li, Y. Zhang, and L. Carin (2016). “Stochastic gradient MCMC with stale gradients”. *Annual Conference on Neural Information Processing Systems*. 2945–2953.
- Chen, L., H. Wang, Z. Charles, and D. Papailiopoulos (2018). “DRACO: Byzantine-resilient distributed training via redundant gradients”. *International Conference on Machine Learning*.
- Chen, X., M. R. Lyu, and I. King (2017). “Toward efficient and accurate covariance matrix estimation on compressed data”. *International Conference on Machine Learning*.
- Cho, M., U. Finkler, D. Kung, and H. Hunter (2019). “BlueConnect: Decomposing AllReduce for deep learning on heterogeneous network hierarchy”. *Conference on Systems and Machine Learning*.
- Cohen, M., J. Diakonikolas, and L. Orecchia (2018). “On acceleration with noise-corrupted gradients”. *International Conference on Machine Learning*.
- Colin, I., A. Bellet, J. Salmon, and S. Clémenccon (2016). “Gossip dual averaging for decentralized optimization of pairwise functions”. *International Conference on Machine Learning*.
- Cutkosky, A. and R. Busa-Fekete (2018). “Distributed stochastic optimization via adaptive SGD”. *Annual Conference on Neural Information Processing Systems*. 1914–1923.
- Damaskinos, G., E. M. El Mhamdi, R. Guerraoui, R. Patra, and M. Taziki (2018). “Asynchronous Byzantine machine learning (the case of SGD)”. *International Conference on Machine Learning*.
- De Sa, C. M., C. Zhang, K. Olukotun, and C. Ré (2015). “Taming the wild: A unified analysis of Hogwild-style algorithms”. *Annual Conference on Neural Information Processing Systems*.

- Dünner, C., T. Parnell, D. Sarigiannis, N. Ioannou, A. Anghel, G. Ravi, M. Kandasamy, and H. Pozidis (2018). “Snap ML: A hierarchical framework for machine learning”. *Annual Conference on Neural Information Processing Systems*. 250–260.
- Elgohary, A., M. Boehm, P. J. Haas, F. R. Reiss, and B. Reinwald (2016). “Compressed linear algebra for large-scale machine learning”. *Proceedings of the VLDB Endowment*. 9(12): 960–971.
- Garber, D., O. Shamir, and N. Srebro (2017). “Communication-efficient algorithms for distributed stochastic principal component analysis”. *International Conference on Machine Learning*.
- Ghadimi, S., G. Lan, and H. Zhang (2013). “Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization”. *arXiv:1308.6594*.
- Gopal, S. (2016). “Adaptive sampling for SGD by exploiting side information”. *International Conference on Machine Learning*.
- Goyal, P., P. Dollár, R. B. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He (2017). “Accurate, large minibatch SGD: Training ImageNet in 1 hour”. *arXiv:1706.02677*.
- Gu, B., Z. Huo, C. Deng, and H. Huang (2018). “Faster derivative-free stochastic algorithm for shared memory machines”. *International Conference on Machine Learning*.
- Gupta, S., A. Agrawal, K. Gopalakrishnan, and P. Narayanan (2015). “Deep learning with limited numerical precision”. *International Conference on Machine Learning*: 1737–1746.
- Hanzely, F., K. Mishchenko, and P. Richtarik (2018). “SEGA: Variance reduction via gradient sketching”. *Annual Conference on Neural Information Processing Systems*. 2086–2097.
- Haochen, J. and S. Sra (2019). “Random shuffling beats SGD after finite epochs”. *International Conference on Machine Learning*. 2624–2633.
- Hashemi, S. H., S. A. Jyothi, and R. H. Campbell (2019). “TicTac: Accelerating distributed deep learning with communication scheduling”. *Conference on Systems and Machine Learning*.
- Hazan, E. and H. Luo (2016). “Variance-reduced and projection-free stochastic optimization”. *International Conference on Machine Learning*.

- He, C., C. Tan, H. Tang, S. Qiu, and J. Liu (2019). “Central server free federated learning over single-sided trust social networks”. *arXiv:1910.04956*.
- He, L., A. Bian, and M. Jaggi (2018). “COLA: Decentralized linear learning”. *Annual Conference on Neural Information Processing Systems*. 4541–4551.
- Hofmann, T., A. Lucchi, S. Lacoste-Julien, and B. McWilliams (2015). “Variance reduced stochastic gradient descent with neighbors”. *Annual Conference on Neural Information Processing Systems*. 2305–2313.
- Horváth, S. and P. Richtarik (2019). “Nonconvex variance reduced optimization with arbitrary sampling”. *International Conference on Machine Learning*.
- Hsieh, C.-J., H.-F. Yu, and I. Dhillon (2015). “PASSCoDe: Parallel ASynchronous stochastic dual co-ordinate descent”. *International Conference on Machine Learning*.
- Hsieh, K., A. Harlap, N. Vijaykumar, D. Konomis, G. R. Ganger, P. B. Gibbons, and O. Mutlu (2017). “Gaia: Geo-distributed machine learning approaching LAN speeds”. *Proceedings of the 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI ’17)*.
- Jayarajan, A., J. Wei, G. Gibson, A. Fedorova, and G. Pekhimenko (2019). “Priority-based parameter propagation for distributed DNN training”. *Conference on Systems and Machine Learning*.
- Ji, K., Z. Wang, Y. Zhou, and Y. Liang (2019). “Improved zeroth-order variance reduced algorithms and analysis for nonconvex optimization”. *International Conference on Machine Learning*.
- Jia, Z., M. Zaharia, and A. Aiken (2019). “Beyond data and model parallelism for deep neural networks”. *Conference on Systems and Machine Learning*.
- Jiang, J., B. Cui, C. Zhang, and L. Yu (2017). “Heterogeneity-aware distributed parameter servers”. *International Conference on Management of Data*.

- Jiang, J., F. Fu, T. Yang, and B. Cui (2018). “SketchML: Accelerating distributed machine learning with data sketches”. *Proceedings of the 2018 International Conference on Management of Data. SIGMOD’18.* 1269–1284.
- Jiang, P. and G. Agrawal (2018). “A linear speedup analysis of distributed deep learning with sparse and quantized communication”. *Annual Conference on Neural Information Processing Systems.* 2530–2541.
- Johnson, T. B. and C. Guestrin (2018). “Training deep models faster with robust, approximate importance sampling”. *Annual Conference on Neural Information Processing Systems.* 7276–7286.
- Jothimurugesan, E., A. Tahmasbi, P. Gibbons, and S. Tirthapura (2018). “Variance-reduced stochastic gradient descent on streaming data”. *Annual Conference on Neural Information Processing Systems.* 9928–9937.
- Kamp, M., M. Boley, O. Missura, and T. Gärtner (2017). “Effective parallelisation for machine learning”. *Annual Conference on Neural Information Processing Systems.* 6480–6491.
- Kaoudi, Z., J.-A. Quiane-Ruiz, S. Thirumuruganathan, S. Chawla, and D. Agrawal (2017). “A cost-based optimizer for gradient descent optimization”. *International Conference on Management of Data.*
- Kara, K., K. Eguro, C. Zhang, and G. Alonso (2018). “ColumnML: Column-store machine learning with on-the-fly data transformation”. *Proceedings of the VLDB Endowment.* 12(4): 348–361.
- Karakus, C., Y. Sun, S. Diggavi, and W. Yin (2017). “Straggler mitigation in distributed optimization through data encoding”. *Annual Conference on Neural Information Processing Systems.* 5440–5448.
- Katharopoulos, A. and F. Fleuret (2018). “Not all samples are created equal: Deep learning with importance sampling”. *International Conference on Machine Learning.*
- Koloskova, A., S. Stich, and M. Jaggi (2019). “Decentralized stochastic optimization and Gossip algorithms with compressed communication”. *International Conference on Machine Learning.*
- Li, X., T. Zhao, R. Arora, H. Liu, and J. Haupt (2016). “Stochastic variance reduced optimization for nonconvex sparse learning”. *International Conference on Machine Learning.*

- Li, Y. and M. Yan (2019). “On linear convergence of two decentralized algorithms”. *arXiv:1906.07225*.
- Li, Y., M. Yu, S. Li, S. Avestimehr, N. S. Kim, and A. Schwing (2018). “Pipe-SGD: A decentralized pipelined SGD framework for distributed deep net training”. *Annual Conference on Neural Information Processing Systems*. 8056–8067.
- Li, Z., W. Shi, and M. Yan (2019). “A decentralized proximal-gradient method with network independent step-sizes and separated convergence rates”. *IEEE Transactions on Signal Processing*. 67(17): 4494–4506.
- Li, Z. and M. Yan (2017). “A primal-dual algorithm with optimal stepsizes and its application in decentralized consensus optimization”. *arXiv:1711.06785*.
- Lian, X., Y. Huang, Y. Li, and J. Liu (2015). “Asynchronous parallel stochastic gradient for nonconvex optimization”. *Annual Conference on Neural Information Processing Systems*. 2737–2745.
- Lian, X., C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu (2017). “Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent”. *Annual Conference on Neural Information Processing Systems*. 5336–5346.
- Lian, X., H. Zhang, C.-J. Hsieh, Y. Huang, and J. Liu (2016). “A comprehensive linear speedup analysis for asynchronous stochastic parallel optimization from zeroth-order to first-order”. *Annual Conference on Neural Information Processing Systems*. 3062–3070.
- Lian, X., W. Zhang, C. Zhang, and J. Liu (2019). “Asynchronous decentralized parallel stochastic gradient descent”. *International Conference on Machine Learning*.
- Lim, H., D. G. Andersen, and M. Kaminsky (2019). “3LC: Lightweight and effective traffic compression for distributed machine learning”. *Conference on Systems and Machine Learning*.
- Ling, Q. and A. Ribeiro (2013). “Decentralized dynamic optimization through the alternating direction method of multipliers”. *IEEE Transactions on Signal Processing*. 62(5): 1185–1197.

- Liu, J. and S. J. Wright (2015). “Asynchronous stochastic coordinate descent: Parallelism and convergence properties”. *SIAM on Optimization*.
- Liu, J., S. J. Wright, C. Ré, V. Bittorf, and S. Sridhar (2014). “An asynchronous parallel stochastic coordinate descent algorithm”. *International Conference on Machine Learning*.
- Liu, S., B. Kailkhura, P.-Y. Chen, P. Ting, S. Chang, and L. Amini (2018). “Zeroth-order stochastic variance reduction for nonconvex optimization”. *Annual Conference on Neural Information Processing Systems*. 3731–3741.
- Liu, X., Y. Li, J. Tang, and M. Yan (2019). “A double residual compression algorithm for efficient distributed learning”. *arXiv:1910.07561*.
- Ma, C., V. Smith, M. Jaggi, M. Jordan, P. Richtarik, and M. Takac (2015). “Adding vs. averaging in distributed primal-dual optimization”. *International Conference on Machine Learning*.
- Mahajan, D., J. K. Kim, J. Sacks, A. Ardalan, A. Kumar, and H. Esmaeilzadeh (2018). “In-RDBMS hardware acceleration of advanced analytics”. *Proceedings of the VLDB Endowment*. 11(11): 1317–1331.
- Nagaraj, D., P. Jain, and P. Netrapalli (2019). “SGD without replacement: Sharper rates for general smooth convex functions”. *International Conference on Machine Learning*.
- Namkoong, H., A. Sinha, S. Yadlowsky, and J. C. Duchi (2017). “Adaptive sampling probabilities for non-smooth optimization”. *International Conference on Machine Learning*.
- Nedić, A. and A. Ozdaglar (2009). “Distributed subgradient methods for multi-agent optimization”. *IEEE Transactions on Automatic Control*. 54(1): 48–61.
- Nguyen, L., P. H. A. Nguyen, M. van Dijk, P. Richtarik, K. Scheinberg, and M. Takac (2018). “SGD and Hogwild! Convergence without the bounded gradients assumption”. *International Conference on Machine Learning*.
- Palaniappan, B. and F. Bach (2016). “Stochastic variance reduction methods for saddle-point problems”. *Annual Conference on Neural Information Processing Systems*. 1416–1424.

- Pan, X., M. Lam, S. Tu, D. Papailiopoulos, C. Zhang, M. I. Jordan, K. Ramchandran, and C. Ré (2016). “Cyclades: Conflict-free asynchronous machine learning”. *Annual Conference on Neural Information Processing Systems*. 2576–2584.
- Peng, H., S. Zhe, X. Zhang, and Y. Qi (2017). “Asynchronous distributed variational Gaussian process for regression”. *International Conference on Machine Learning*.
- Qu, Z., P. Richtarik, and T. Zhang (2015). “Quartz: Randomized dual coordinate ascent with arbitrary sampling”. *Annual Conference on Neural Information Processing Systems*. 865–873.
- Raviv, N., R. Tandon, A. Dimakis, and I. Tamo (2018). “Gradient coding from cyclic MDS codes and expander graphs”. *International Conference on Machine Learning*.
- Reddi, S. J., A. Hefny, S. Sra, B. Poczos, and A. Smola (2016). “Stochastic variance reduction for nonconvex optimization”. *International Conference on Machine Learning*.
- Shamir, O. (2016). “Without-replacement sampling for stochastic gradient methods”. *Annual Conference on Neural Information Processing Systems*. 46–54.
- Shi, W., Q. Ling, G. Wu, and W. Yin (2015a). “A proximal gradient algorithm for decentralized composite optimization”. *IEEE Transactions on Signal Processing*. 63(22): 6013–6023.
- Shi, W., Q. Ling, G. Wu, and W. Yin (2015b). “Extra: An exact first-order algorithm for decentralized consensus optimization”. *SIAM Journal on Optimization*. 25(2): 944–966.
- Simonetto, A., A. Koppel, A. Mokhtari, G. Leus, and A. Ribeiro (2017). “Decentralized prediction-correction methods for networked time-varying convex optimization”. *IEEE Transactions on Automatic Control*. 62(11): 5724–5738.
- Simsekli, U., C. Yildiz, T. H. Nguyen, T. Cemgil, and G. Richard (2018). “Asynchronous stochastic quasi-Newton MCMC for non-convex optimization”. *International Conference on Machine Learning*.
- Sridhar, S., V. Bittorf, J. Liu, C. Zhang, C. Ré, and S. J. Wright (2013). “An approximate efficient solver for LP rounding”. *Annual Conference on Neural Information Processing Systems*. 2895–2903.

- Stich, S. U., J.-B. Cordonnier, and M. Jaggi (2018). “Sparsified SGD with memory”. *Annual Conference on Neural Information Processing Systems*. 4452–4463.
- Sun, T., R. Hannah, and W. Yin (2017). “Asynchronous coordinate descent under more realistic assumptions”. *Annual Conference on Neural Information Processing Systems*.
- Tandon, R., Q. Lei, A. G. Dimakis, and N. Karampatziakis (2017). “Gradient coding: Avoiding stragglers in distributed learning”. *International Conference on Machine Learning*.
- Tang, H., S. Gan, C. Zhang, T. Zhang, and J. Liu (2018a). “Communication compression for decentralized training”. *Annual Conference on Neural Information Processing Systems*. 7663–7673.
- Tang, H., X. Lian, S. Qiu, L. Yuan, C. Zhang, T. Zhang, and J. Liu (2019a). “DeepSqueeze: Parallel stochastic gradient descent with double-pass error-compensated compression”. *arXiv:1907.07346*.
- Tang, H., X. Lian, M. Yan, C. Zhang, and J. Liu (2018b). “ D^2 : Decentralized training over decentralized data”. *International Conference on Machine Learning*.
- Tang, H., X. Lian, T. Zhang, and J. Liu (2019b). “Doublesqueeze: Parallel stochastic gradient descent with double-pass error-compensated compression”. *International Conference on Machine Learning*.
- Wang, H., S. Sievert, S. Liu, Z. Charles, D. Papailiopoulos, and S. Wright (2018a). “ATOMO: Communication-efficient learning via atomic sparsification”. *Annual Conference on Neural Information Processing Systems*. 9872–9883.
- Wang, H., S. Sievert, S. Liu, Z. Charles, D. Papailiopoulos, and S. Wright (2018b). “Atom: Communication-efficient learning via atomic sparsification”. *Annual Conference on Neural Information Processing Systems*. 9850–9861.
- Wang, J., M. Kolar, N. Srebro, and T. Zhang (2017). “Efficient distributed learning with sparsity”. *International Conference on Machine Learning*.
- Wang, J. and G. Joshi (2019). “Adaptive communication strategies to achieve the best error-runtime trade-off in local-update SGD”. *Conference on Systems and Machine Learning*.

- Wang, N., J. Choi, D. Brand, C.-Y. Chen, and K. Gopalakrishnan (2018c). “Training deep neural networks with 8-bit floating point numbers”. *Annual Conference on Neural Information Processing Systems*. 7686–7695.
- Wang, S., D. Li, Y. Cheng, J. Geng, Y. Wang, S. Wang, S.-T. Xia, and J. Wu (2018d). “BML: A high-performance, low-cost gradient synchronization algorithm for DML training”. *Annual Conference on Neural Information Processing Systems*. 4243–4253.
- Wang, Z., K. Kara, H. Zhang, G. Alonso, O. Mutlu, and C. Zhang (2019). “Accelerating generalized linear models with MLWeaving: A one-size-fits-all system for any-precision learning”. *Proceedings of the VLDB Endowment*. 12(7): 807–821.
- Wangni, J., J. Wang, J. Liu, and T. Zhang (2018). “Gradient sparsification for communication-efficient distributed optimization”. *Annual Conference on Neural Information Processing Systems*. 1306–1316.
- Wen, W., C. Xu, F. Yan, C. Wu, Y. Wang, Y. Chen, and H. Li (2017). “TernGrad: Ternary gradients to reduce communication in distributed deep learning”. *Annual Conference on Neural Information Processing Systems*. 1508–1518.
- Wu, J., W. Huang, J. Huang, and T. Zhang (2018). “Error compensated quantized SGD and its applications to large-scale distributed optimization”. *International Conference on Machine Learning*.
- Xie, C., S. Koyejo, and I. Gupta (2019). “Zeno: Distributed stochastic gradient descent with suspicion-based fault-tolerance”. *International Conference on Machine Learning*.
- Xin, D., S. Macke, L. Ma, J. Liu, S. Song, and A. Parameswaran (2018). “HELIX: Holistic optimization for accelerating iterative machine learning”. *Proceedings of the VLDB Endowment*. 12(4): 446–460.
- Ye, M. and E. Abbe (2018). “Communication-computation efficient gradient coding”. *International Conference on Machine Learning*.
- Yin, D., Y. Chen, R. Kannan, and P. Bartlett (2018). “Byzantine-robust distributed learning: Towards optimal statistical rates”. *International Conference on Machine Learning*.
- You, Y., J. Li, J. Hseu, X. Song, J. Demmel, and C. Hsieh (2020). “Reducing BERT pre-training time from 3 days to 76 minutes”. *arXiv:1904.00962*.

- You, Y., X. Lian, J. Liu, H.-F. Yu, I. S. Dhillon, J. Demmel, and C.-J. Hsieh (2016). “Asynchronous parallel greedy coordinate descent”. *Annual Conference on Neural Information Processing Systems*. 4689–4697.
- Yu, C., H. Tang, C. Renggli, S. Kassing, A. Singla, D. Alistarh, C. Zhang, and J. Liu (2019). “Distributed learning over unreliable networks”. *International Conference on Machine Learning*.
- Yu, M., Z. Lin, K. Narra, S. Li, Y. Li, N. S. Kim, A. Schwing, M. Annavaram, and S. Avestimehr (2018). “GradiVeQ: Vector quantization for bandwidth-efficient gradient aggregation in distributed CNN training”. *Annual Conference on Neural Information Processing Systems*.
- Yuan, K., Q. Ling, and W. Yin (2016). “On the convergence of decentralized gradient descent”. *SIAM Journal on Optimization*. 26(3): 1835–1854.
- Zhang, C., P. Patras, and H. Haddadi (2018). “Deep learning in mobile and wireless networking: A survey”. *arXiv:1803.04311*.
- Zhang, H., J. Li, K. Kara, D. Alistarh, J. Liu, and C. Zhang (2017). “ZipML: Training linear models with end-to-end low precision, and a little bit of deep learning”. *International Conference on Machine Learning*.
- Zhang, L., T. Yang, R. Jin, Y. Xiao, and Z.-H. Zhou (2016). “Online stochastic linear optimization under one-bit feedback”. *International Conference on Machine Learning*.
- Zhang, S., A. E. Choromanska, and Y. LeCun (2015). “Deep learning with elastic averaging SGD”. *Annual Conference on Neural Information Processing Systems*. 685–693.
- Zhao, P. and T. Zhang (2015). “Stochastic optimization with importance sampling for regularized loss minimization”. *International Conference on Machine Learning*.
- Zheng, S., Q. Meng, T. Wang, W. Chen, N. Yu, Z.-M. Ma, and T.-Y. Liu (2017). “Asynchronous stochastic gradient descent with delay compensation”. *International Conference on Machine Learning*.
- Zhou, D., P. Xu, and Q. Gu (2018a). “Stochastic nested variance reduced gradient descent for nonconvex optimization”. *Annual Conference on Neural Information Processing Systems*.

- Zhou, D., P. Xu, and Q. Gu (2018b). “Stochastic variance-reduced cubic regularized Newton methods”. *International Conference on Machine Learning*.
- Zhou, K., F. Shang, and J. Cheng (2018c). “A simple stochastic variance reduced algorithm with fast convergence rates”. *International Conference on Machine Learning*.
- Zhou, Z., P. Mertikopoulos, N. Bambos, P. Glynn, Y. Ye, L.-J. Li, and L. Fei-Fei (2018d). “Distributed asynchronous optimization with unbounded delays: How slow can you go?” *International Conference on Machine Learning*.
- Zhu, R. (2016). “Gradient-based sampling: An adaptive importance sampling for least-squares”. *Annual Conference on Neural Information Processing Systems*. 406–414.
- Zhu, R. and Q. Gu (2015). “Towards a lower sample complexity for robust one-bit compressed sensing”. *International Conference on Machine Learning*.
- Zhu, Y. and J. Lafferty (2018). “Distributed nonparametric regression under communication constraints”. *International Conference on Machine Learning*.
- Zou, D., P. Xu, and Q. Gu (2018). “Stochastic variance-reduced Hamilton Monte Carlo methods”. *International Conference on Machine Learning*.